Demand Uncertainty and Dynamic Allocation with Strategic Agents

Alex Gershkov and Benny Moldovanu*

29.11.2010

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

In this paper we revisit a simple, canonical model of dynamic trading where an uninformed designer allocates an indivisible object to a stream of randomly arriving, privately informed agents. 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. We are grateful to the European Research Council and to the German Science Foundation for financial support. Gershkov: Department of Economics and Center for the Study of Rationality, The Hebrew University of Jerusalem, alexg@huji.ac.il; Moldovanu: Department of Economics, University of Bonn, mold@uni-bonn.de.

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 profitability of future transactions.

The main question we ask is whether the designer is able to implement, under the above conditions, the first-best dynamic policy, i.e., the efficient policy he would choose in the dynamic setting where he is perfectly informed about both valuations and arrival times.

Our particular setting builds upon models appearing in the dynamic pricing literature, but the implications of strategic timing decisions on the ability of a planner to learn about the environment and to implement a particular optimal policy are important for a variety of dynamic settings, as detailed below.

Dynamic pricing has been pioneered on an industrial scale by airline companies in the mid 70's. The so called revenue (or yield) management practices have rapidly spread to hotel booking, freight transportation, car rentals, retailing (e.g., apparel and groceries), electricity generation, e-business, advertising and broadcasting, and event management (sports, concerts, etc...). In their excellent overview, Talluri and van Ryzin [2004] repeatedly point to the fact that the various workhorse models in the accompanying literature have myopic customers who buy a product as soon as the price drops below their willingness to pay. Alternatively, it is assumed that agents are short lived and only have one immediate opportunity to buy the product. The main reason behind these assumptions is technical: allowing for strategic consumers and opportunistic planners usually yields a complex, difficult to analyze game whose predictions are not easily amenable to practical, industrial implementations. But, it is of course important to know how the models' predictions and recommendations change once long lived, strategic agents are taken into account: the more expensive and durable the purchase, the more significant this aspect seems to be.

Other important settings where strategic timing decisions affects learning and implementability of optimal policies include:

a In financial markets specialist market makers face sequences of traders, some of them better informed than the market. Each transaction affects the bid-ask spread since the market maker has to protect himself from adverse selection. The possibility of strategic timing of trading decisions differentiates two classical models of the market micro-structure literature. In Glosten and Milgrom [1985] the market maker faces a known, exogenously given arrival process of informed and uninformed traders. Traders are short lived, and they do not incorporate the effect of their own and others' trades in the decision when to trade. In contrast, Kyle [1985] considers an exogenous arrival process of short lived uninformed traders, but the informed trader - there is only one of these, which greatly simplifies the analysis - can be thought as optimizing his trading times[1].

b Monetary policy is often gradually implemented, with new information arriving from the market being incorporated along the way. Investors who anticipate how their behavior changes future monetary policy will strategically choose investment times, sometimes completely defeating the purpose of the planned policy. Caplin and Leahy[1996] model the search for an optimal monetary stimulus: a policy maker who is uninformed about the state of the economy repeatedly sets interest rates and gathers information by observing whether agents invest or not. A difficulty arises when strategic investors realize that their current action influences future investment opportunities: less investment today

[1] Back and Baruch [2004] provide a comparison between the equilibrium outcomes of the two models.

makes the policy maker more pessimistic, which leads to a lower future interest rate, and thus to more profitable future investment opportunities. Such behavior hinders the policy maker's ability to implement the optimal policy, e.g., a gradual reduction in interest rates may be ineffectual. As we shall see below, this kind of difficulty - strategies that makes the decision maker more pessimistic about the future - plays an important role in our paper.

c In another type of models, the agents themselves learn over time by observing others[2]. For example potential investors in an infant industry,whose profitability is very uncertain, may want to delay their investment decisions until the market conditions (which may include the consequences of a targeted industrial policy) become clearer. This kind of behavior may hinder the implementation of a targeted industrial policy. Rob (1991) rationalizes subsidies to early entrants in infant industries by showing that they can restore efficiency by correctly adjusting the informational externality. Such subsidy schemes play an important role also in our current setting.

In the present paper we expand on a classical continuous-time search model: 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 for correlations in arrival times. The precise probabilistic nature of the arrival process may not be known to the planner. Both agents and planner discount the future, and the planner wishes to maximize the expected value of the allocation.

We ask whether the complete information efficient policy (first-best) is implementable via individually rational monetary transfers that do not depend on events posterior to the physical allocation. We also identify scenarios where the implementing scheme involves only transfers that accompany physical trades. Such requirements are very natural in the applications mentioned above, e.g., for the trading of securities among a market maker and other agents.

[2] See Chamley [2004] for a good survey.

The main insights derived in the paper can be summarized as follows:

1. If agents' arrivals are observable, the dynamically efficient policy can be implemented by requiring payments equal to the imposed **allocative** externality (Proposition 5). This insight holds although the independence assumptions usually invoked in the dynamic mechanism design literature (see review below) are not satisfied in our model. In order to calculate the externality price, the planner needs to consider alternative scenarios where allocations to other agents occur at possibly later points in time.

2. 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 **??**. 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 times[3]. Together with the allocative externality payment discussed above, this subsidy implements the complete information efficient dynamic policy (Proposition 6). The subsidy has the flavor of an "advanced booking" discount, but here the discount is independent of a physical transaction.

3. In contrast to the case where there is learning via past observed values (see Gershkov and Moldovanu [2009b]), a positive implementation

[3]Non-negativity is required by the individual rationality constraint.

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 before. But practical applications of this positive result hinge on a verifiability 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, e.g., when a license to operate is awarded or when an infant industry is supported by the state. When such identification is impossible it makes sense to restrict attention to mechanisms that couple monetary transfers to physical transactions (*winner-pays mechanism*). Proposition 13 shows that no individual rational winner-pay scheme can implement the efficient allocation if, roughly speaking, a later arrival of some agent makes the designer more pessimistic about future arrivals. Finally, Proposition 9 and Proposition 11 identify large and important classes of arrival processes - renewals with a known distribution of inter–arrival times, and a non-homogenous Poisson process combined with a pure birth process, respectively - where the complete information efficient policy can be implemented via winner-pays mechanism even if arrivals are not observable. In all cases where winner-pay mechanisms are able to implement the efficient dynamic allocation, our methods immediately yield also corresponding revenue maximizing scheme.

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 and show that an expected externality scheme implements the efficient allocation for the case where values are private, but arrivals are observable. In Section 5 we show that the expected externality 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. Several proofs are relegated to an Appendix.

1.1 Related Literature

As mentioned above, most of the classical yield management literature (see Talluri and van Ryzin [2004] for a comprehensive survey) abstracts from strategic effects on the buyer side. 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 [2010b] allow the planner to learn about the distribution of values from past observations. 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.

Mason and Välimäki [2009] focus on revenue maximizing, posted-price mechanisms in a model with unknown demand and with stochastic and unobservable arrivals of short-lived buyers. 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.

A small literature considers the effect of 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, but restrict the monopolist seller to choosing two prices, independently of the past history of sales.

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. But, 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. Thus the strategic element in the time dimension and the ensuing learning which are the focus of our paper does not appear there. Gallien's solution necessarily coincides with the one where arrivals are observable, and where agents are short lived (see Albright [1974] and Gershkov and Moldovanu [2009a]).

Board and Skrypacz [2010] characterize revenue maximization in a finite horizon, discrete time model where agents are patient. The main difference to our framework is that arrivals are described by a known process where current arrivals are independent of past ones. The optimal mechanism awards the good to the agent with the highest valuation exceeding a cutoff. In their model the cutoff is constant in all periods prior to the last.

Pai and Vohra [2008] and Mierendorff [2010a] 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.

A recent literature looks at efficient dynamic mechanism design. Bergemann and Valimäki [2010], Cavallo, Parkes and Singh [2010], Parkes and Singh [2003], Said [2010] and Athey and Segal [2007] 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[4]. Roughly speaking, these authors use various independence assumptions in order to stay within a private values framework, and 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").[5] Mierendorff [2010b] an-

[4]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.

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

alyzes an independent, private value model where an 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.

Finally, we note that 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. Monetary transfers are therefore not necessary in order to execute the optimal policy.

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[6]. The time horizon is potentially infinite, but the framework is rich enough to embedd the finite horizon case by considering arrival process where after some time T no more arrivals occur, i.e., where $\mathcal{N}(t)$ is constant for any $t \geq T$. 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's private information is two-dimensional: the arrival time $t \geq 0$ and the value $v \geq 0$ he gets if allocated the object. In other words, we assume that 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 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.

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

The agents' values are assumed to be represented by I.I.D. random variables $\widetilde{v}_i$ on support $[0,a]$ where $a \leq \infty$, with common c.d.f. F, with continuos density f. We assume that each $\widetilde{v}_i$ has a finite mean, and a finite 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.

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[7].

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 $\mathbf{t}_{\mathcal{N}(T)} = \left(t_1, ..., t_{\mathcal{N}(T)}\right)$ 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 $X'(T) \geq X(T)$.

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

Proposition 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 arrival[8]. We illustrate below the main difficulty the planner faces when trying to implement the first-best allocation rule.

Example 2 *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.[9] The Poisson process is a special case where G is exponential (see Mortensen [1986]). 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 policy[10]. Nevertheless, Zuckerman

[8]Even if the designer is not completely informed, the arrival process is still a counting process.

[9]Note that, although inter-arrival times are independently distributed random variables, arrival times are correlated in a renewal process.

[10]For example, consider a process where times between consecutive arrivals can be either ε or Δ where $\varepsilon << \Delta$. Assume that a buyer with a moderately high value arrives at t.

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 reserve[11].

Definition 3 *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 $(W - y/W \geq y)$.*[12]

Theorem 4 *(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 4 implies that the optimal complete information

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 $t + \Delta$).

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

[12] Note that any random variables with an increasing hazard (or failure) rate is NBU.

policy for the case where the designer knows the relevant distribution of interarrival 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 that

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

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 [13]:

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 5 below) generate 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 recall that the object is allocated upon arrival in the complete information optimal policy (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

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

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 $e^{-\delta t'}\left(v - x_{[2,3]}(\delta)\right) > 0$ [14]. Hence overall 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 subsidies 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) and the subsidy scheme in 2 - which deals with the informational externality- does implements 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.

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

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.

4 Direct Revelation Mechanisms

We now come back to the general, incomplete information model. Without loss of generality (see Myerson [1986]) 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 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

[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.

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., $h = \lim_{T\to\infty} \eta(T)$.

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 is 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]

Finally, 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'

[16]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.

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

types form history η_{-i}, and when the designer uses the optimal stopping rule.

4.1 Externality Payments for Observable Arrivals

Let us first consider the case where values are private information, but arrival times are observable. We show that 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. Note that the independence assumptions imposed by Bergemann and Valimaki [2010] are not satisfied here because of the possible correlations in arrival times (which implicitly determine whether the value for the object at a certain period is positive or not). Nevertheless, the described mechanism has a similar flavor to their dynamic pivot mechanism.

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 $\eta_{-i}(T)$. Observe (also for later uses) that

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

Proposition 5 *The payment scheme*

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

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

Recall here we assumed (without loss of generality) 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 in Proposition 5 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.

5 A Subsidy Scheme for Early Arrivals

In this section we 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

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

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).

Under a very mild condition about the variation of $U\left(t,t',v\right)$, 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, if the optimal stopping time is almost surely finite, efficiency can be approximated by taking T arbitrarily large.

Proposition 6 *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*

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

Then the subsidy $S(t')=e^{\delta t'}M(T-t')$ together with the payment scheme given in Proposition (5) implements the dynamically efficient allocation for any history h where the optimal stopping time is less than T.

Remark 7 *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 9 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. The result of Proposition 6 can be extended to the time interval $[0, \infty)$ (hence for all possible histories) under some other conditions: Assume for example that $\frac{\partial U(t,t',v)}{\partial t'}$ exists for any $v, t \leq t'$, and define $L(t') = \sup_{v,t \leq t'} \frac{\partial U(t,t',v)}{\partial t'}$. Assume further that there exists a function $c(t') \geq L(t')$ such that $C = \int_0^\infty c(z)dz < \infty$. We need to ensure that the function $U(t, t', v) + e^{-\delta t'} S(t')$ decreases in t' for any v and $t \leq t'$. Hence, it is sufficient to find a function S such that

$$e^{-\delta t'} \left(S'(t') - \delta S(t') \right) + L(t') \leq 0.$$

In particular, we can choose an S that satisfies the differential equation:

$$e^{-\delta t'} \left(S'(t') - \delta S(t') \right) + L(t') = 0.$$

The non-negative subsidy function

$$S(t') = e^{t'\delta} \left[C - \int_0^{t'} L(z)dz \right]$$

solves the differential equation, and therefore, together with the payment defined in Proposition 5, implements then the efficient allocation for any history of arrivals.

3. 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 2, where we illustrated such a case.

The above rather permissive result is in stark contrast with the restricted possibility result obtained by Gershkov and Moldovanu [2010b] 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 8 *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.

We first look at renewals. 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 9 *Assume that arrivals are unobservable, and that the arrival process is a renewal with inter-arrival distribution G.*

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.

Remark 10 *Let us briefly look at the infinite-horizon variant of a problem first posed in a discrete time, finite horizon and undiscounted model by Pai and Vohra [2008]: each buyer i's private information is characterized by a triple (v_i, t_i, d_i) where v_i is the value for the object, t_i is the arrival time, and $d_i \geq t_i$ is the exit time*[18]*. The designer's goal is to maximize revenue. We are able to completely solve this problem for our setting if arrivals are governed by an NBU renewal process. Assume that the virtual valuation $x - \frac{1-F(x)}{f(x)}$ is strictly increasing, and denote by H the implied distribution of virtual values. Then, for any functions d_i, $d_i \geq t_i$, $i = 1, 2, ..,$ the revenue maximizing policy is to charge a constant price P where P is the unique solution to the equation*
$$p = \frac{\phi(\delta)}{1-\phi(\delta)} \int_{p}^{\infty} (z - p) dH(z)$$
The proof follows immediately from the above Theorem, and from the standard mechanism design exercise of replacing revenue maximization with welfare

[18]This last component can be thought of as a deadline after which his value for the object drops to zero. This deadline may be either a deterministic, or a stochastic function of t_i. The case $d_i = t_i$ for all i corresponds to the case witout recall, while the case $d_i = \infty$ for all i corresponds to the case with perfect recall analyzed above

maximization with respect to virtual values. Since in this case no recall occurs in the revenue optimizing policy, the solution coincides with the one found by Gallien [2006] for the case $d_i = \infty$ (perfect recall) and by Gershkov and Moldovanu [2009a] for the case $d_i = t_i$ (no recall)[19]*. Gallien's analysis is involved since he derived the above result without using Zuckerman's [1988] insight on the continuos-time search problem.*

More generally, as long as the implementation of the efficient allocation in our basic model with types (v_i, t_i) does not require a subsidy - as is the case in Proposition 9 above, or in Propositions 11 below- our results immediately yield the revenue maximizing scheme[20]*. In particular, the revenue maximizing mechanism is given then by a sequence of posted prices that depends on the arrival history.*

We now turn our attention to non-homogenous Poisson/pure birth processes where the arrival rate is non-decreasing in both the elapsed time t and the number of arrivals i up to t. Here an optimal policy that stops almost surely in finite time may not exist because the designer gets more and more optimistic over time. In order to prove existence we need to bound the expected discounted number of arrivals. The proof of the next result is somewhat involved: it uses a combination of concepts and results from majorization theory, and from the theory of order statistics.

Proposition 11 *Assume that arrivals are governed by a non-homogenous Poisson/pure-birth process with arrival rate $\lambda_i(t)$ such that λ is non-decreasing in both elapsed time t and the number of arrivals i up to t. Suppose that there exists a non-decreasing function $\beta(t)$ such that $\lambda_i(t) \leq \beta(t)$ for all i, t, and such that $\int_0^\infty \beta(t)e^{-\delta t}dt < \infty$. Then:*

1. *An optimal stopping rule exists for the setting with observable arrivals.*
2. *Charging for the object a payment of $v_T^*(\mathbf{t}_{\mathcal{N}(T)})$, the optimal cutoff under observable arrivals, implements the efficient allocation also if arrivals*

[19] Gallien offers a solution also for the case where the seller has several identical units: then the revenue maximizing price jumps up after each sale. Gershkov and Moldovanu allow for multiple heterogenous units. In their no-recall case, the solution has the same form for any renewal process (not necessarily NBU).

[20] Technically, this require the virtual value to be increasing in value, as above.

are non-observable. In other words, a subsidy is not needed for efficient implementation.

Example 12 *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 [2010a], 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* $lim_{t\to\infty}(y(t)e^{-\delta t}) = 0$ *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.*

Our last result generalizes the insight at the end of Section 3, and has a negative flavor: roughly speaking, 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.

Proposition 13 *Assume that there exist an agent* i *, a type* (t, v) *of that agent, and* $t' > t$ *such that:*

Condition 1 *For any* $h_{-i} \in H_{-i}(t, v)$ *it holds that* $\tau_v(t, h_{-i}) \geq \tau_v(t', h_{-i})$ *with strict inequality for a non-zero measure of histories* h_{-i}

Condition 2 *For any* $t'' \geq t'$ *and any* $\eta_{-i}(t'')$ *it holds that* $V(t, 0, \eta_{-i}(t''), t'') \geq V(t', 0, \eta_{-i}(t''), t'')$

Then there is no winner-pay mechanism that implements the dynamically efficient allocation.

Remark 14 *Instead of Condition 2 in the statement of Proposition 13, it would be enough to require the much weaker average condition*

$$\begin{aligned} & \int_{H_{-i}(t,v)} e^{-\delta\tau_v(t',h_{-i})} V(t,0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))d\mu(h_{-i}/t) \\ \geq & \int_{H_{-i}(t,v)} e^{-\delta\tau_v(t',h_{-i})} V(t',0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))d\mu(h_{-i}/t) \end{aligned}$$

which is implied by Condition 2. Since Condition 2 implies that, for any possible history, a later arrival of i *will make the planner more pessimistic about the future arrivals, one may conjecture that this condition already implies Condition 1. However, this is not the case. For example, consider a homogeneous Poisson arrival process with an unknown parameter. In this case, the state of the process at time* T *only depends on* T*,* $X(T)$*, and on the total number of arrivals (but not on their precise timing) - see for example Pratt, Raiffa and Schlaifer [1995], Chapter 15. Thus, for a given history of arrivals of agents other than* i*, at any time* $t'' > t' \geq t$ *, the planner's belief is the same, no matter whether agent* i *arrived at* t *or at* t'*. In particular, Condition 2 always holds with equality for any agent* i *and for any* $t' \geq t$*. But, there is no agent* i*, and value* v *of that agent such that* $\tau_v(t,h_{-i}) > \tau_v(t',h_{-i})$*.*

7 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 where late arrivals induce pessimism about future arrivals. For practical applications, it is of interest to further study "second-best" policies in such environments.

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 on pricing have been mostly neglected by the existing literature since the standard assumption was that buyers are short-lived.

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.

8 Appendix

Proof for Example 2. 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

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

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 Proposition 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 $X'(T)\alpha\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right) + \beta\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right)$ where $\alpha\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right)$ is the discounted probability that the object will be allocated to the agent with value $X'(T)$ and where $\beta\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right)$ 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 then one (otherwise it would be optimal to stop at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$) and therefore $\alpha\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right) < 1$. Since Υ prescribes to continue search at state $\{X'(T), T, \mathbf{t}_{\mathcal{N}(T)}\}$, it must be the case that

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

which implies

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

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 $X(T)\alpha\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right) + \beta\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right)$. Since we know that $\beta\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right) > \left(1 - \alpha\left(T, \mathbf{t}_{\mathcal{N}(T)}, \Upsilon\right)\right) X'(T)$ we obtain

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

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 5. Observe first that the optimal stopping time $\tau_v(t, h_{-i})$ is decreasing in v for any $t,\ h_{-i}$. To see this, assume that $\tau_v(t, h_{-i}) = T$. Then $V(t, v, \eta^h_{-i}(T), T) \geq v^*_T(\mathbf{t}_{\mathcal{N}(T)})$. Hence, if agent i arrives at t with value $v' \geq v$, we obtain

$$V(t, v', \eta^h_{-i}(T), T) \geq V(t, v, \eta^h_{-i}(T), T) \geq v^*_T(\mathbf{t}_{\mathcal{N}(T)})$$

and therefore $\tau_{v'}(t, h_{-i}) \leq T$, as desired.

If agent i arrives at time t and values the object at $v = 0$, the principal's expected utility under the complete information optimal stopping rule is given by

$$U_0 = \int e^{-\delta\tau_0(t,h_{-i})} V(t, 0, \eta^h_{-i}\left(\tau_0(t, h_{-i})\right), \tau_0(t, h_{-i})) d\mu(h_{-i}/t)$$

Assume now that agent i arrives at time t with value v. The planner's utility under the complete information optimal stopping rule is given by:

$$\begin{aligned}
&\int_{H^c_{-i}(v)} e^{-\delta\tau_0(t,h_{-i})} V(t, 0, \eta^h_{-i}\left(\tau_0(t, h_{-i})\right), \tau_0(t, h_{-i})) d\mu(h_{-i}/t) + \\
&\int_{H_{-i}(v)} e^{-\delta\tau_v(t,h_{-i})} v d\mu(h_{-i}/t) dh_{-i} = U_0 + \\
&\int_{H_{-i}(v)} e^{-\delta\tau_v(t,h_{-i})} \left[v - e^{-\delta\tau_0(t,h_{-i}+\delta\tau_v(t,h_{-i})} V(t, 0, \eta^h_{-i}\left(\tau_0(t, h_{-i})\right), \tau_0(t, h_{-i}))\right] d\mu(h_{-i}/t) = \\
&U_0 + \int_{H_{-i}(v)} e^{-\delta\tau_v(t,h_{-i})} \left[v - V(t, 0, \eta^h_{-i}\left(\tau_v(t, h_{-i})\right), \tau_v(t, h_{-i}))\right] d\mu(h_{-i}/t)
\end{aligned}$$

The last equality follows by the definition of $V(t, 0, \eta^h_{-i}\left(\tau_{v'}(t, h_{-i})\right), \tau_{v'}(t, h_{-i})$, by the law of iterated expectation, and because $\tau_v(t, x) \leq \tau_0(t, x)$.

If agent i arrives at time t with value v, but the designer uses the policy as if the value would be v', her utility is analogously:

$$U_0 + \int_{H_{-i}(v')} e^{-\delta\tau_{v'}(t,h_{-i})} \left[v - V(t, 0, \eta^h_{-i}\left(\tau_{v'}(t, h_{-i})\right), \tau_{v'}(t, h_{-i}))\right] d\mu(h_{-i}/t)$$

By the optimality of the stopping time $\tau_v(t, x)$ for an agent arriving at time t with value v, we obtain that

$$
\begin{aligned}
& U_0 + \int_{H_{-i}(v)} e^{-\delta \tau_v(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_v(t,h_{-i})\right), \tau_v(t,h_{-i})) \right] d\mu(h_{-i}/t) \\
\geq \;& U_0 + \int_{H_{-i}(v')} e^{-\delta \tau_{v'}(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_{v'}(t,h_{-i})\right), \tau_{v'}(t,h_{-i})) \right] d\mu(h_{-i}/t)
\end{aligned}
$$

$$
\begin{aligned}
\Leftrightarrow & \int_{H_{-i}(v)} e^{-\delta \tau_v(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_v(t,h_{-i})\right), \tau_v(t,h_{-i})) \right] d\mu(h_{-i}/t) \\
\geq & \int_{H_{-i}(v')} e^{-\delta \tau_{v'}(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_{v'}(t,h_{-i})\right), \tau_{v'}(t,h_{-i})) \right] d\mu(h_{-i}/t).
\end{aligned} \tag{5}
$$

Incentive compatibility for agent i requires that

$$
\begin{aligned}
& \int_{H_{-i}(v)} e^{-\delta \tau_v(t,h_{-i})} [v - P(t,v,\eta^h_{-i}\left(\tau_v(t,h_{-i})\right), \tau_v(t,h_{-i}))] d\mu(h_{-i}/t) \\
\geq \;& \int_{H_{-i}(v')} e^{-\delta \tau_{v'}(t,h_{-i})} [v - P(t,v',\eta^h_{-i}\left(\tau_{v'}(t,h_{-i})\right), \tau_{v'}(t,h_{-i}))] d\mu(h_{-i}/t).
\end{aligned}
$$

Plugging the payment scheme from (3) gives

$$
\begin{aligned}
& \int_{H_{-i}(v)} e^{-\delta \tau_v(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_v(t,h_{-i})\right), \tau_v(t,h_{-i})) \right] d\mu(h_{-i}/t) \\
& \int_{H_{-i}(v')} e^{-\delta \tau_{v'}(t,h_{-i})} \left[v - V(t,0,\eta^h_{-i}\left(\tau_{v'}(t,h_{-i})\right), \tau_{v'}(t,h_{-i})) \right] d\mu(h_{-i}/t)
\end{aligned}
$$

which is exactly inequality (5).

Note that $V(t,v,\eta^h_{-i}(T),T) \geq v$ for any T, with equality at $T = \tau_v(t,h_{-i})$ by the definition of the optimal stopping rule. Hence individual rationality follows because no money flows between agents that do not get the object and the planner, and because the payment of a winner, which is given by $V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))$, satisfies

$$
V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i})) \leq V(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i})) = v.
$$

■

Proof of Proposition 6. 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

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

Given truthful reporting in the arrival time dimension, the payment scheme of Proposition 5 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 5, and because $S(t') \geq 0$. ■

Proof of Proposition 9. **1.** We show that $U(t,t',v)$ decreases in t' for any v and $t \leq t'$ (see Proposition 6 and Remark 7-1). 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 4. ■

For the proof of Proposition 11 we first need the following concepts and results from majorization theory:

Definition 15 *1. For any $n-$tuple $\gamma = (\gamma_1, \gamma_2, .., \gamma_n)$ let $\gamma_{(j)}$ denote the jth largest coordinate (so that $\gamma_{(n)} \leq \gamma_{(n-1)} \leq ... \leq \gamma_{(1)}$). 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:*

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

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

Theorem 16 (Marschall and Proschan [1965]) *Let $X = (X_1, ..., X_n)$ be an n-dimensional random vector with a permutation invariant joint distribution. Let ϕ: $R^n \to R$ be a continuous, convex function that is permutation invariant in its arguments. Then the function $E\phi(\alpha_1 X_1, ..., \alpha_n X_n)$ is Schur convex, i.e.,*

$$E\phi(a_1 X_1, ..., a_n X_n) \leq E\phi(\beta_1 X_1, ..., \beta_n X_n).$$

whenever $\alpha \prec \beta$.

Proof of Proposition 11. **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 then 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.[21]

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

[21] The 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.

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 $X_1, ..., X_k$ with mean μ, representing the agents' values. Then the designer's expected utility in Problem A is given by $\sum_{k=0}^{\infty} P(k)E(X_{(k)})$.[22] Since $E\left(\sum_{i=1}^{k} X_i\right) = k\mu$, we know that $E(X_{(k)}) \leq k\mu$. Hence

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

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 $(t_1, t_2, ...)$, then the planner's expected utility is lower than that in Problem A with a number of agents given by $\left\lceil \sum_{i=1}^{\infty} e^{-\delta t_i} \right\rceil$ 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 $E[\max\{\alpha_1 X_1, \alpha_2 X_2, ...\}]$ where $\alpha_i = e^{-\delta t_i}$ and where t_i is the arrival time of the $i-$th agent[23]. If $\left\lceil \sum_{i=1}^{\infty} e^{-\delta t_i} \right\rceil$ is finite then $E[\max\{\alpha_1 X_1, \alpha_2 X_2, ...\}]$ is finite as well. To see that observe that

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

[22] We set $E(X_{(0)}) = 0$.

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

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 $EX_{(l)}$ where $l = \left\lceil \sum_{i=1}^{K} e^{-\delta t_i} \right\rceil$.

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

$$E[\max\{\alpha_1 X_1, \alpha_2 X_2, ..., \alpha_K X_K\}] \leq E[\max\{\widetilde{\alpha}_1 X_1, \widetilde{\alpha}_2 X_2, ..., \widetilde{\alpha}_K X_K\}] \quad (6)$$

where $1 \geq \widetilde{\alpha}_i \geq \alpha_i$ for $i \in \{1, ..., K\}$ and $\sum_{i=1}^{K} \widetilde{\alpha}_i = \left\lceil \sum_{i=1}^{K} \alpha_i \right\rceil$.[24]

Consider now $\sum_{i=1}^{K} \widetilde{\alpha}_i = l$ as above. Since $0 \leq \widetilde{\alpha}_i \leq 1$ for any i, the vector $(\widetilde{\alpha}_1, ..., \widetilde{\alpha}_K)$ is majorized by the vector $(\underbrace{1, .., 1}_{l}, \underbrace{0, ..0}_{K-l})$. Since the X_i's are I.I.D. , and since the maximum is a continuous, permutation invariant convex function, we can apply Theorem 16 to obtain

$$E[\max\{\widetilde{\alpha}_1 X_1, \widetilde{\alpha}_2 X_2, ..., \widetilde{\alpha}_K X_K\}] \leq E[\max\{X_1, ..., X_l\}] = EX_{(l)}. \quad (7)$$

Inequalities (6) and (7) 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.

2. The above argument 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.

[24] Since $\alpha_i \leq 1$ and $\sum_{i=1}^{K} \alpha_i \leq K$ such $\widetilde{\alpha}_i$s always exist.

As a consequence, 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. ■

Proof of Proposition 13. Since agent i with value v gets the object at history h_{-i} if $v = V(t, v, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))$, incentive compatibility requires that:

$$\begin{aligned} &\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}[v - P(t, v, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))]d\mu(h_{-i}/t) \\ \geq &\int_{H_{-i}(t',v')} e^{-\delta\tau_{v'}(t',h_{-i})}[v - P(t', v', \eta^h_{-i}(\tau_{v'}(t', h_{-i})), \tau_{v'}(t', h_{-i}))]d\mu(h_{-i}/t) \end{aligned}$$

where $t' \geq t$. In particular, this requires that deviations only in the value coordinate are suboptimal:

$$\begin{aligned} &\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}[v - P(t, v, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))]d\mu(h_{-i}/t) \\ \geq &\int_{H_{-i}(t,v')} e^{-\delta\tau_{v'}(t,h_{-i})}[v - P(t, v', \eta^h_{-i}(\tau_{v'}(t, h_{-i})), \tau_{v'}(t, h_{-i}))]d\mu(h_{-i}/t). \end{aligned}$$

Since the above inequality corresponds to the incentive compatibility constraint in the case of observable arrivals, payoff equivalence implies that any price $P(t, v, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))$ that implements the efficient allocation must satisfy

$$\begin{aligned} &\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}P(t, v, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))]d\mu(h_{-i}/t) \\ = &\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}V(t, 0, \eta^h_{-i}(\tau_v(t, h_{-i})), \tau_v(t, h_{-i}))]d\mu(h_{-i}/t) + C \end{aligned} \tag{8}$$

where C is a constant (independent of v). For winner-pay, individually rational mechanisms we necessarily have $C = 0$.

Incentive compatibility also implies that misrepresenting the arrival time only should not be profitable:

$$\begin{aligned}&\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}[v-P(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))]d\mu(h_{-i}/t)\\ \geq &\int_{H_{-i}(t',v)} e^{-\delta\tau_v(t',h_{-i})}[v-P(t',v,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))]d\mu(h_{-i}/t)\end{aligned}$$

for $t' \geq t$. Plugging (8) into the last expression yields:

$$\begin{aligned}&\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}[v-V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))]d\mu(h_{-i}/t) \quad (9)\\ \geq &\int_{H_{-i}(t',v)} e^{-\delta\tau_v(t',h_{-i})}[v-V(t',0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))]d\mu(h_{-i}/t).\end{aligned}$$

We now show that the last inequality cannot hold if the conditions in the statement of the Proposition are satisfied, yielding a contradiction to the possibility of efficient implementation.

Condition 1 implies that there exists v and $t' \geq t$ such that for any h_{-i} we have

$$V(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i})) \geq V(t',v,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i})).$$

This follows from the fact that $V(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))$ - the expected utility of the planner evaluated at the optimal stopping time - is just the highest value of the agents that arrived before that time. Since by assumption $\tau_v(t,h_{-i}) \geq \tau_v(t',h_{-i})$ we obtain

$$\begin{aligned}&V(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))\\ = \;& X(\tau_v(t,h_{-i})) \geq X(\tau_v(t',h_{-i})) = V(t',v,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))\end{aligned}$$

as desired. Since $t \leq t' \leq \tau_v(t',h_{-i}) \leq \tau_v(t,h_{-i})$, this also yields:

$$\begin{aligned}H_{-i}(t,v) &\equiv \{h_{-i}/v = V(t,v,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))\}\\ &\subseteq \{h_{-i}/v = V(t',v,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))\} \equiv H_{-i}(t',v)\end{aligned}$$

Since the integrands in (9) are positive, we must have:

$$\begin{aligned}&\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}[v-V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))]d\mu(h_{-i}/t) (10)\\ \geq &\int_{H_{-i}(t,v)} e^{-\delta\tau_v(t',h_{-i})}[v-V(t',0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))]d\mu(h_{-i}/t).\end{aligned}$$

By Condition 1 we know that

$$e^{-\delta\tau_v(t,h_{-i})}v \leq e^{-\delta\tau_v(t',h_{-i})}v, \tag{11}$$

with strict inequality for a non-zero measure of histories. Condition 2 implies that

$$\begin{aligned}
& \int_{H_{-i}(t,v)} e^{-\delta\tau_v(t,h_{-i})}V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))d\mu(h_{-i}/t) \\
= & \int_{H_{-i}(t,v)} e^{-\delta\tau_v(t',h_{-i})}V(t,0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))d\mu(h_{-i}/t) \\
\geq & \int_{H_{-i}(t,v)} e^{-\delta\tau_v(t',h_{-i})}V(t',0,\eta^h_{-i}(\tau_v(t',h_{-i})),\tau_v(t',h_{-i}))d\mu(h_{-i}/t). \quad (12)
\end{aligned}$$

The first equality follows from the definition of $V(t,0,\eta^h_{-i}(\tau_v(t,h_{-i})),\tau_v(t,h_{-i}))$ and from the law of iterated expectation since by Condition 1 we have $\tau_v(t',h_{-i}) \leq \tau_v(t,h_{-i})$. But, taken together, inequalities (11) and (12) contradict inequality (10). ■

References

[1974] Albright S.C.: "Optimal Sequential Assignments with Random Arrival Times", *Management Science* **21** (1), 60-67.

[2007] Athey, S. and Segal, I.: "An Efficient Dynamic Mechanism", discussion paper, Stanford University.

[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.

[2004] Back, K. and Baruch, S.: "Information in Securities Markets: Kyle meets Glosten and Milgrom", *Econometrica* **72**, 433-465.

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

[2010] Board, S. and Skrzypacz, A.: "Optimal Dynamic Auctions for Durable Goods", 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: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.

[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.

[2010b] Gershkov, A. and Moldovanu, B.: "Optimal Search, Learning, and Implementation", discussion paper, University of Bonn.

[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.

[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.

[2009] Mason, R. and Välimäki, J.: "Learning and Smooth Stopping", discussion paper, University of Exeter

[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] 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

[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)

[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.

[2010] Said, M.: "Auctions with Dynamic Populations: Efficiency and Revenue Maximization", discussion paper, Microsoft Research.

[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,

[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.

[2004] Talluri, K.T., and Van Ryzin, G.: *The Theory and Practice of Revenue Management,* Springer: New York.