



Munich Personal RePEc Archive

Auctions with Dynamic Populations: Efficiency and Revenue Maximization

Said, Maher

Microsoft Research New England, Olin Business School, Washington
University in St. Louis

20 August 2009

Online at <https://mpra.ub.uni-muenchen.de/16895/>
MPRA Paper No. 16895, posted 22 Aug 2009 19:17 UTC

AUCTIONS WITH DYNAMIC POPULATIONS: EFFICIENCY AND REVENUE MAXIMIZATION

MAHER SAID[†]
AUGUST 20, 2009

ABSTRACT: We study a setting where objects and privately-informed buyers arrive stochastically to a market. A seller in this setting faces a sequential allocation problem with a dynamic population. We derive both efficient and revenue-maximizing incentive compatible direct mechanisms. Our main result shows that the sequential ascending auction is a simple indirect mechanism that achieves these desirable objectives. We construct equilibria in memoryless strategies where, in every period, bidders reveal all private information. These equilibria are outcome equivalent to the direct mechanisms. In contrast to static settings, sequential second-price auctions cannot yield these outcomes, as they do not reveal sufficient information.

KEYWORDS: Dynamic mechanism design, Sequential allocation, Random arrivals, Revenue equivalence, Indirect mechanisms, Sequential ascending auctions.

JEL CLASSIFICATION: C73, D44, D82, D83.

[†]MICROSOFT RESEARCH NEW ENGLAND AND OLIN BUSINESS SCHOOL, WASHINGTON UNIVERSITY IN ST. LOUIS
E-mail address: msaid@microsoft.com.

This paper is based on a chapter of my doctoral dissertation at Yale University. I am indebted to Dirk Bergemann, Johannes Hörner, Ben Polak, and Larry Samuelson for their generous advice, comments, and encouragement. I would also like to thank Rossella Argenziano, Alessandro Bonatti, Rahul Deb, Eduardo Faingold, James Fenske, Sergei Iz-malkov, Katrina Jessoe, Nenad Kos, Mallesh Pai, Kareen Rozen, Ilya Segal, Itai Sher, Juuso Välimäki, and Rakesh Vohra for their many helpful discussions and suggestions. Financial support from the National Science Foundation, the Cowles Foundation, and Yale University is gratefully acknowledged.

1. INTRODUCTION

In this paper, we study the mechanism design problem of a seller in a dynamic market. In each period, a random number of buyers and objects arrive to the market. Buyers are risk-neutral and patient, while objects are homogeneous and perishable. Each buyer desires a single unit of the good in question, and valuations for the good vary across buyers. A mechanism designer must elicit the private information of these buyers in order to achieve her desired outcome—either an efficient or a revenue-maximizing allocation.

We first discuss direct revelation mechanisms which may be used to achieve these outcomes. In particular, we show that a dynamic analogue of the Vickrey-Clarke-Groves (VCG) mechanism is efficient, and that the optimal (revenue-maximizing) mechanism in this dynamic setting is essentially a pivot mechanism with a reserve price. We then consider the possibility of achieving the mechanism designer's goals via a simple *indirect* mechanism. In sharp contrast to the static world, second-price sealed-bid auctions do not lead to efficient outcomes, as they do not correspond to the dynamic version of the VCG mechanism. We show that decentralized implementation via a simple auction format is still possible, however: a sequence of ascending (English) auctions is outcome equivalent to dynamic VCG. Adding an optimally-chosen reserve price to the sequential ascending auction then leads to an intuitive revenue-maximizing indirect mechanism.

The role of population dynamics in markets is an under-studied topic that is of great importance. This is especially true because the vast majority of “real-world” markets are asynchronous: not all buyers and sellers are available or present at the same time. Rather, agents arrive at the market at different times, interact with various segments of the population, and then transact at different times. This fact, in conjunction with the potential arrival or departure of agents from the market in the future, leads to a trade-off: competition in the future may be higher or lower than at the present time, and opportunities to trade may arise more or less frequently. Thus, agents must choose between transacting now or waiting until the (uncertain) future.

In addition to this dynamic trade-off, an additional strategic element arises due to competition between agents *across* time. Buyers and sellers may face the same competitors repeatedly, implying that individuals will want to learn the private information of others. Moreover, each agent may be concerned about how her competitors will make use of any information that she reveals about herself.

To make these trade-offs and considerations more concrete, consider for a moment the problem faced by a buyer searching for a product on an online auction market such as eBay. Upon her arrival to the market, this buyer will have available to her a variety of auctions to participate in. Moreover, she can choose to “wait and see,” postponing her participation until a future date. Supposing that our buyer does, in fact, choose to participate in an auction immediately, she must then decide how much to bid. However, her willingness to pay will depend on her expectations about the future. From her perspective, future supply is random—she does not know when the next auction for a similar item will take place, nor how many such future auctions may occur. Similarly, future competition—the number of potential competitors, as well as their strength—is unknown to our buyer.

In addition, this hypothetical buyer on eBay has available to her a wealth of information. She may observe the prices at which similar items have sold for in previous auctions, as well as the actual bids submitted by various competitors. While rational bidding behavior requires the incorporation of such information into a submitted bid, our buyer may also be concerned with how her bid, given its observability, will affect others' behavior. She could, for instance, try to strategically alter her competitors' expectations about the future—submitting a relatively high bid, for example, could serve as a “signal” of high future competition.

Taking these considerations into account, it is clear that population dynamics can have a significant impact on issues such as competition, price determination, efficiency, and revenue. And given this impact, it is natural to question how this impact varies across different institutions or market forms. Therefore, in the present work, we are concerned with two main questions. First, can we achieve efficient or revenue-maximizing outcomes in markets with dynamic populations of privately-informed buyers? And secondly, and equally importantly, can we achieve these outcomes using natural or simple “real world” institutions?

The approach we take to answering these two questions is to develop a reasonably general model of a dynamic environment that reflects some key features of markets where dynamic populations are important. Note that we do not model eBay or some other specific market “X.” Rather, we are interested in determining how far the intuitions provided by static models may be pushed, where those intuitions break down, and what new insights and approaches are necessitated by market dynamics.

Thus, the model we present abstracts away many of the details of such dynamic markets, focusing instead on what we view as their essential features. In particular, demand is not constant, as the set and number of buyers change over time, with patient buyers entering and exiting the market according to a stochastic process. Similarly, supply is random. In some periods there may be many units available, while in others none. Finally, each buyer's valuation—her willingness to pay—is her private information. Therefore, a welfare- or revenue-maximizing seller must provide appropriate incentives for information revelation to this dynamic population. The seller then makes use of this information to dynamically allocate goods to buyers.

A natural candidate for achieving a welfare-maximizing allocation is the Vickrey-Clarke-Groves mechanism. It is well-known that, in static environments, the VCG mechanism is efficient. By choosing a transfer payment for each buyer that equals the externality imposed by her report on other participants in the mechanism, the VCG mechanism aligns the incentives of the buyer with those of a welfare-maximizing social planner. This leads to efficiency and dominant-strategy incentive compatibility, as truthful reporting now maximizes both the planner's and the buyers' objective functions. In the dynamic environment we consider, the arrival of a new buyer imposes an externality on her competitors by reordering the (anticipated) schedule of allocations to those buyers currently present on the market, as well as to those buyers expected to arrive in future periods. By charging each agent, upon her arrival, a price equal to this *expected* externality, the buyer's incentives can be aligned with those of the forward-looking planner. Therefore, this dynamic version of the VCG mechanism implements the efficient policy.

We also construct a revenue-maximizing direct mechanism for this setting. Making use of the risk-neutrality of buyers, we show that the optimal policy for a revenue-maximizing seller is equivalent to that of a social planner who wishes to maximize allocative efficiency, except that buyers' values are replaced by their *virtual* values. Each buyers' incentives may then be aligned with those of the seller by adding (optimal) reserve prices to the dynamic VCG mechanism. These reserve prices are chosen so as to provide each buyer with an expected payoff equal to her expected marginal contribution to the *virtual* surplus. This allows the seller to discriminate between buyers in such a way as to maximize revenue.

Notice that both of the mechanisms discussed above are direct revelation mechanisms, requiring buyers to report their values to the mechanism upon their arrival to the market. In practice, however, direct revelation mechanisms may be difficult to implement. For instance, the multi-unit Vickrey auction—the (static) multi-unit generalization of the standard VCG mechanism—is a direct revelation mechanism in which truth-telling is not just equilibrium behavior, but is in fact a dominant strategy for all participants. Despite this, [Ausubel \(2004\)](#) points out that the Vickrey auction lacks simplicity and transparency, explaining that “many [economists] believe it is too complicated for practitioners to understand.” Moreover, [Rothkopf, Teisberg, and Kahn \(1990\)](#) explain that concerns about privacy or the potential for future misuse of information revealed in a direct mechanism may preclude the real-world use of direct mechanisms.

These criticisms are corroborated by experimental evidence. According to [Kagel, Harstad, and Levin \(1987\)](#), who examined single-unit auctions with affiliated private values, the theoretical predictions about bidding behavior are significantly more accurate in ascending price-auctions than in second-price auctions, despite the existence of a dominant-strategy equilibrium in the second-price (Vickrey) auction. [Kagel and Levin \(2009\)](#) find a similar result in multi-object auctions with independent private values: ascending-type clock auctions significantly outperform the dominant-strategy solvable Vickrey auction in terms of efficiency. In another study examining the efficiency properties of several mechanisms in a resource allocation problem similar to the one we consider here, [Banks, Ledyard, and Porter \(1989\)](#) find that “the transparency of a mechanism . . . is important in achieving more efficient allocations.” In their experiments, a simple ascending auction dominated both centralized administrative allocation processes as well as decentralized markets in terms of both efficiency and revenues.

With these criticisms and “real-world feasibility” constraints in mind, we turn to the design of simple, transparent, and decentralized indirect mechanisms. In particular, we consider the possibility of achieving efficient or revenue-maximizing outcomes via a sequence of auctions. Despite the similarity of our direct mechanisms to their single-unit static counterparts, we show that this resemblance does not hold for the corresponding auction formats. Recall that, in the canonical static allocation problem, the analogue of the VCG mechanism is the second-price auction. In an environment with a dynamic population of buyers, however, a sequence of second-price auctions cannot yield outcomes equivalent to those of the dynamic VCG mechanism. In a sequential auction, there is an “option value” associated with losing in a particular period, as buyers may win an auction in a future period. The value of this option depends on expected future prices, which is determined by the private information of other competitors. Thus, despite working in a framework

with independent private values, the strategic environment faced by individual bidders is more complicated, as the dynamics of the auction market induce interdependence in (option) values.

This interdependence implies that a standard second-price sealed-bid auction does not reveal sufficient information for the determination of buyers' option values. In contrast, the ascending auction is a simple *open* auction format that does allow for the gradual revelation of buyers' private information. We use this fact to construct intuitive equilibrium bidding strategies for buyers in a sequence of ascending auctions. In each period, buyers bid up to the price at which they are indifferent between winning an object and receiving their expected future contribution to the social welfare. As buyers drop out of the auction, they (indirectly) reveal their private information to their competitors, who are then able to condition their current-period bids on this information.

When this process of information revelation is repeated in *every* period, newly arrived buyers are able to learn about their competitors without being privy to the events of previous periods. This information renewal is crucial for providing the appropriate incentives for new entrants to also reveal their private information. This allows for “memoryless” behavior—incumbent buyers willingly ignore payoff-relevant information from previous periods, as they correctly anticipate that it will be revealed again in the course of the current auction. These memoryless strategies are *not* the result of an a priori restriction on the strategy space, but are instead the result of fully-rational optimization on the part of individual buyers, leading to prices and allocations identical to the truth-telling equilibrium of the efficient direct mechanism. Moreover, these strategies form a periodic ex post equilibrium: given her expectations about future competition, each buyer's behavior in any period remains optimal even after observing her current opponents' values.

Similar arguments apply when considering revenue-maximizing indirect mechanisms. When buyers' values are drawn from the same distribution, the sequential ascending auction with an optimally-chosen reserve price admits an equilibrium that is equivalent to truth-telling in the revenue-maximizing direct mechanism. Thus, the sequential ascending auction is a natural decentralized institution for achieving either efficient or optimal outcomes.

The present work contributes to a recent literature exploring dynamic allocation problems and dynamic mechanism design.¹ Bergemann and Välimäki (2008) develop the dynamic pivot mechanism, a dynamic generalization of the Vickrey-Clarke-Groves mechanism that yields efficient outcomes when agents' private information evolves stochastically over time. Athey and Segal (2007) characterize an efficient dynamic mechanism that is budget-balanced and incentive compatible, again in the presence of evolving private information. In a similar dynamic setting, Pavan, Segal, and Toikka (2009) and Rahman (2009) consider the more general question of characterizing incentive-compatible mechanisms. While these papers study dynamic mechanisms for a fixed set of buyers whose types may change over time, we examine a setting where the number and set of buyers may change over time but types are fixed.² Moreover, a fundamental departure from the previous literature is our focus on the design of *indirect* mechanisms.

¹Parkes (2007) surveys much of this literature, including the many contributions of the computer science community.

²There is also a recent literature on dynamic allocation problems and mechanism design with an evolving population, but *without* money. See, for instance, Kurino (2009) or Ünver (2009).

This paper also relates to recent work on dynamic auctions and revenue management. Mierendorff (2008) characterizes an auction mechanism that efficiently allocates a single storable object when buyers arrive over the course of the auction. Pai and Vohra (2008) derive the revenue-maximizing mechanism for allocating a finite number of storable objects to buyers whose arrival to and departure from the market is also private information. Vulcano, van Ryzin, and Maglaras (2002) also examine optimal mechanisms for selling identical objects to randomly arriving buyers. When the objects are heterogeneous but commonly-ranked, Gershkov and Moldovanu (2008) and (2009) derive revenue-maximizing and efficient mechanisms. In contrast to the present work, the buyers in these models are impatient, and there is a fixed number of storable objects to be allocated.

Finally, our analysis of indirect mechanisms is linked to the sequential auctions literature. The seminal work is Milgrom and Weber (2000), which examines the properties of a variety of auction formats for the (simultaneous or sequential) sale of a fixed set of objects to a fixed set of buyers. Kittsteiner, Nikutta, and Winter (2004) extend that model to one in which buyers discount the future. Unlike the present work, however, they require the presence of all buyers in the initial period, leading to dramatically different conclusions about the equivalence (or lack thereof) of second-price and ascending auctions. Said (2009) examines the role of random entry in a model of sequential second-price auctions when objects are stochastically equivalent; that is, when values are independently and identically distributed across both buyers and objects. The computer science literature, motivated in part by the emergence of online auction sites such as eBay, has also turned attention towards sequential ascending auctions. Lavi and Nisan (2005) and Lavi and Segev (2009) examine the “worst-case” performance of sequential ascending auctions with dynamic buyer populations. Their prior-free, non-equilibrium analysis provides a lower bound on the efficiency of the allocations achieved via sequential ascending auctions.

The remainder of this paper is structured as follows. Section 2 introduces the general model and environment we consider. In Section 3, we fully characterize the efficient allocation rule and discuss its implementation via a dynamic variant of the Vickrey-Clarke-Groves mechanism. We then construct an efficient equilibrium of the sequential ascending auction and show that it is outcome equivalent to that of the efficient direct mechanism. Section 4 parallels the development in Section 3, characterizing the revenue-maximizing allocation policy and constructing an optimal direct mechanism for its implementation. We then show that the indirect sequential ascending auction mechanism is revenue-maximizing when used with an optimally-chosen reserve price. Finally, Section 5 concludes.

2. MODEL

2.1. Buyers, Objects, and Random Arrivals

We consider an infinite-horizon discrete-time environment; time periods are indexed by t , where $t \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$. There is a countable set \mathcal{I} of buyers, where each agent $i \in \mathcal{I}$ desires a single unit of a homogeneous, indivisible good. Each buyer i 's valuation v_i for this good is her private information, and v_i is independently distributed according to the distribution F_i . We assume that F_i has a strictly positive and continuous density f_i and support $\mathbf{V} := [0, \bar{v}]$, and that each buyer's

virtual valuation

$$\varphi_i(v_i) := v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

is a strictly increasing function of v_i .³ Moreover, we assume that buyers are risk neutral, and that their preferences are quasilinear and time separable. The future is discounted exponentially with the (common) discount factor $\delta \in (0, 1)$.

Not all buyers are present in each period. Rather, buyers arrive stochastically to the market. In particular, the set \mathcal{I} of buyers is partitioned into disjoint subsets $\{\mathcal{I}_t\}_{t \in \mathbb{N}_0}$, where \mathcal{I}_t is the finite subset of agents who may arrive in period t . The arrival of agent $i \in \mathcal{I}_t$ in period t is governed by an independent draw from a Bernoulli distribution, where $\pi_i \in [0, 1]$ denotes the probability that i is present. In addition, buyers may depart from the market after each period, where the (common) probability of any buyer i “surviving” to the following period is denoted by $\gamma \in [0, 1]$. Otherwise, buyers remain present in the market until they receive an object. Note that, unlike the probability of arriving to the market, the survival rate is identical across agents.

Thus, buyer arrivals and departures yield a stochastic process $\{\alpha_t\}_{t \in \mathbb{N}_0}$, where $\alpha_t : \mathcal{I} \rightarrow \{0, 1\}$ is an indicator function that tracks the presence of each agent on the market at time t , and

$$\mathcal{A}_t := \{i \in \mathcal{I} : \alpha_t(i) = 1\}$$

is the subset of agents present in period t . We assume that buyers cannot conceal their presence, and so α_t (equivalently, \mathcal{A}_t) is commonly known to the agents present at time t .⁴

In addition to the random arrival of buyers, several units of a homogeneous, indivisible, and non-storable good may also arrive on the market. Let $k_t \in \mathcal{K} := \{0, 1, \dots, \bar{K}\}$ denote the number of objects that arrive in period t , where $\bar{K} \in \mathbb{N}$ is the maximal number of objects potentially available in any given period. As with buyers, the arrival of objects is governed by a stochastic process, where $\mu_t(k) \in [0, 1]$ denotes the probability that exactly $k \in \mathcal{K}$ objects are available in period t . Moreover, these objects are non-storable; any unallocated objects “expire” at the end of each period, and hence cannot be carried over to future periods. This assumption plays an important role in the determination of the efficient policy, providing a great deal of tractability. As with the buyer arrival process, we assume that the arrival of objects is publicly observed, and so k_t is commonly known to those agents present on the market at time t .

At the beginning of each period, new buyers arrive to the market (and old buyers may depart). Simultaneously, new objects arrive, replacing any unallocated objects left over from the previous period. It will be useful to denote the “state” of the market at the beginning of each period t by $\omega_t := (\alpha_t, k_t)$. The realizations of the arrival and departure processes are publicly observed by all agents present on the market, implying that ω_t becomes common knowledge to all agents present at time t . The mechanism designer may then allocate objects to agents, and we move on to the following period.

³The assumptions on F_i are merely for expositional convenience. All the efficiency-related results continue to be true with general distributions. Moreover, for revenue-maximization, we may use the procedure of Skreta (2007) to define “ironed” virtual values that may be used whenever φ_i is decreasing or not well-defined.

⁴This assumption is merely for simplicity. As will be discussed later in Section 3.2, the equilibria we construct and describe remain equilibria in the larger game where buyers may conceal their presence.

2.2. Direct Mechanisms

In this setting, a dynamic direct mechanism asks each agent i to make a *single* report, upon arrival to the market, of her type v_i .⁵ We denote by \emptyset the “report” of an agent who has not arrived to the market. Thus, the mechanism designer has available to her in each period a collection of reports $\mathbf{r}_t : \mathcal{I}_t \rightarrow \mathbf{V} \cup \{\emptyset\}$, where \mathcal{R} is the set of all such reports. Note that the report $r_i \in \mathbf{V}$ of an agent i who *has* arrived need not be truthful, as this will depend upon the incentives provided by the mechanism.

Let \mathcal{H}_t denote the set of period- t histories, where each history $h_t \in \mathcal{H}_t$ is a sequence of arrivals and departures (of buyers and objects), agent reports, and allocations up to, and including, period $t - 1$. Thus, we have

$$h_t = (\omega_0, \mathbf{r}_0, \mathbf{x}_0; \omega_1, \mathbf{r}_1, \mathbf{x}_1; \dots; \omega_{t-1}, \mathbf{r}_{t-1}, \mathbf{x}_{t-1}),$$

where $\mathbf{x}_s = \{x_{i,s}\}_{i \in \mathcal{I}} \in \mathbf{X} := \{0, 1\}^{\mathcal{I}}$ is the allocation in period s .

A *dynamic direct mechanism* is then a sequence of feasible allocations and feasible monetary transfers $\mathcal{M} = \{\mathbf{x}_t, \mathbf{p}_t\}_{t \in \mathbb{N}_0}$, where we abuse notation and denote by

$$\mathbf{x}_t : \mathcal{H}_t \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \times \mathcal{R} \rightarrow \Delta(\mathbf{X})$$

a collection of allocation probabilities for each agent, and denote by

$$\mathbf{p}_t : \mathcal{H}_t \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \times \mathcal{R} \rightarrow \mathbb{R}^{\mathcal{I}}$$

a collection of monetary transfers *from* each agent. The period- t allocation $\mathbf{x}_t = \{x_{i,t}\}_{i \in \mathcal{I}}$ is a *feasible allocation* if, and only if,

$$\sum_{i \in \mathcal{I}} x_{i,t} \leq k_t \text{ and } x_{i,t} = 0 \text{ for all } i \notin \mathcal{A}_t.$$

These two conditions require, respectively, that no more objects than are available in period t are allocated at that time, and that objects are only allocated to agents that are present on the market. Notice that we have implicitly ruled out the possibility of allocating multiple objects to any agent as a consequence of the single-unit demand assumption. Similarly, $\mathbf{p}_t = \{p_{i,t}\}_{i \in \mathcal{I}}$ is a *feasible monetary transfer* if, and only if,

$$p_{i,t} = 0 \text{ for all } i \notin \mathcal{A}_t;$$

that is, agents who are not present on the market cannot make or receive payments.

We assume that, upon her arrival to the market in period t , agent $i \in \mathcal{I}_t$ observes only the current state of the market ω_t ; that is, the set \mathcal{A}_t of agents present on the market (equivalently, the indicator α_t) and the number k_t of objects available at time t . Agent i does not observe the history of arrivals and departures in previous periods or the history of allocative decisions, nor does she observe the reports of agents who have arrived before her. Thus, a reporting strategy for agent i , conditional on having arrived to the market, is simply a mapping

$$r_i : \mathbf{V} \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \rightarrow \mathbf{V}.$$

⁵It is straightforward to see that the revelation principle applies in this setting, and so the restriction to direct mechanisms is without loss of generality.

Let \mathbf{r}_{-i} denote the reports of all agents other than agent $i \in \mathcal{I}_t$. The expected payoff to i when she reports $r_i \in \mathbf{V}$ to the mechanism \mathcal{M} and all other agents report according to \mathbf{r}_{-i} is then

$$\mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(x_{i,s} (h_s, \omega_s, (r_i, \mathbf{r}_{-i})) v_i - p_{i,s} (h_s, \omega_s, (r_i, \mathbf{r}_{-i})) \right) \right],$$

where the expectation is taken with respect to the arrival and departure processes of buyers and sellers, as well the history h_t and the reports of all other agents that may be present on the market. Note that we have dropped the dependence of reporting strategies on histories and market presence to simplify notation.

2.3. Incentive Compatibility, Individual Rationality, and Revenue Equivalence

Consider a direct mechanism $\mathcal{M} = \{\mathbf{x}_t, \mathbf{p}_t\}_{t \in \mathbb{N}_0}$ and fix any period t and any agent $i \in \mathcal{I}_t$. Since we will be examining Bayesian implementation as opposed to dominant strategy implementation, suppose that all other agents $j \neq i$ are reporting truthfully; that is, suppose that

$$r_j(v_j, \omega_s) = v_j$$

for all $s \in \mathbb{N}_0$, all $j \in \mathcal{I}_s \setminus \{i\}$, and every $(v_j, \omega_s) \in \mathbf{V} \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$. For notational convenience, we will denote this strategy by \mathbf{v}_{-i} . Recall that agent $i \in \mathcal{I}_t$, upon her arrival, observes only the set \mathcal{A}_t of agents present and the number k_t of objects available in period t . We may then define

$$U_i(v'_i, v_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(x_{i,s} (h_s, \omega_s, (v'_i, \mathbf{v}_{-i})) v_i - p_{i,s} (h_s, \omega_s, (v'_i, \mathbf{v}_{-i})) \right) \right]$$

to be the expected payoff of agent $i \in \mathcal{I}_t$ from reporting $v'_i \in \mathbf{V}$ when her true type is $v_i \in \mathbf{V}$ and the (observed) current market state is given by $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$.

The mechanism \mathcal{M} is *incentive compatible* if, for all $t \in \mathbb{N}_0$ and all $i \in \mathcal{I}_t$,

$$U_i(v_i, v_i, \omega_t) \geq U_i(v'_i, v_i, \omega_t) \text{ for all } v_i, v'_i \in \mathbf{V} \text{ and all } \omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}.$$

Thus, \mathcal{M} is incentive compatible if truthful reporting by all agents, regardless of their time of entry, the agents present upon their arrival, and the number of objects available, is an equilibrium. Notice that this condition is equivalent to requiring interim (Bayesian) incentive compatibility for each agent, for *every* realization of the arrival processes and every realization of the agent's values.

Similarly, the mechanism \mathcal{M} is *individually rational* if, for all $t \in \mathbb{N}_0$ and all $i \in \mathcal{I}_t$,

$$U_i(v_i, v_i, \omega_t) \geq 0 \text{ for all } v_i \in \mathbf{V} \text{ and all } \omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}.$$

Thus, \mathcal{M} is individually rational if all agents prefer to participate (truthfully) in the mechanism than not, where we have normalized the outside option of each player to zero. As with incentive compatibility, this must hold for *every* realization of the arrival processes and the agent's values.

Notice that, due to the agents' risk neutrality and the quasilinearity of payoffs, we may rewrite the payoff functions U_i as

$$U_i(v'_i, v_i, \omega_t) = q_i(v'_i, \omega_t) v_i - m_i(v'_i, \omega_t),$$

where

$$q_i(v'_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} x_{i,s} (h_s, \omega_s, (v'_i, \mathbf{v}_{-i})) \right] \quad (1)$$

is the expected discounted sum of object allocation probabilities, and

$$m_i(v'_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} p_{i,s} (h_s, \omega_s, (v'_i, \mathbf{v}_{-i})) \right] \quad (2)$$

is the expected discounted sum of payments. Since buyers ultimately care only about the expected discounted probability of receiving an object and their expected discounted payment, the seller can restrict attention to these two functions when designing incentive schemes—we are able to simplify the incentive problem faced by a seller in this setting by reducing the problem to a single-dimensional allocation problem.

Therefore, we may rewrite the incentive compatibility and individual rationality constraints as

$$q_i(v_i, \omega_t)v_i - m_i(v_i, \omega_t) \geq q_i(v'_i, \omega_t)v_i - m_i(v'_i, \omega_t) \text{ for all } v_i, v'_i \in \mathbf{V}$$

and

$$q_i(v_i, \omega_t)v_i - m_i(v_i, \omega_t) \geq 0 \text{ for all } v_i \in \mathbf{V},$$

for all $t \in \mathbb{N}_0$, all $i \in \mathcal{I}_t$, and all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$.

Define for all $t \in \mathbb{N}_0$ and $i \in \mathcal{I}_t$ the function $\widehat{U}_i : \mathbf{V} \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \rightarrow \mathbb{R}$ by

$$\widehat{U}_i(v_i, \omega_t) := q_i(v_i, \omega_t)v_i - m_i(v_i, \omega_t).$$

\widehat{U}_i is then the expected payoff of agent i from truthfully reporting her value v_i . Making use of this function, we are able to extend the classic Myerson (1981) characterization of incentive compatibility and expected payoffs to this dynamic setting. In particular, we obtain the following results.

LEMMA 1 (Characterization of implementable mechanisms).

A direct mechanism $\mathcal{M} = \{\mathbf{x}_t, \mathbf{p}_t\}_{t \in \mathbb{N}_0}$ is incentive compatible and individually rational if, and only if, the following conditions are satisfied for all $t \in \mathbb{N}_0$ all $i \in \mathcal{I}_t$, and all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$:

- (1) $q_i(v_i, \omega_t)$ is nondecreasing in v_i ;
- (2) $\widehat{U}_i(v_i, \omega_t) = \widehat{U}_i(0, \omega_t) + \int_0^{v_i} q_i(v'_i, \omega_t) dv'_i$ for all $v_i \in \mathbf{V}$; and
- (3) $\widehat{U}_i(0, \omega_t) \geq 0$.

PROOF. The proof may be found in [Appendix A](#). □

COROLLARY 1 (Revenue equivalence).

If the dynamic direct mechanism \mathcal{M} is incentive compatible, then for all $t \in \mathbb{N}_0$ all $i \in \mathcal{I}_t$, and all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$, the expected payment of type $v_i \in \mathbf{V}$ of buyer i , conditional on entry, is

$$m_i(v_i, \omega_t) = m_i(0, \omega_t) + q_i(v_i, \omega_t)v_i - \int_0^{v_i} q_i(v'_i, \omega_t) dv'_i.$$

If, in addition, \mathcal{M} is individually rational, then $m_i(0, \omega_t) \leq 0$.

These results are the dynamic population analogues of the standard Myerson (1981) results for static allocation problems. Recall that in static settings with single-unit demand, incentive compatibility requires that increasing a buyer’s type should increase (weakly) her probability of receiving an object. In our dynamic setting with single-unit demand, incentive compatibility instead requires that increasing a buyer’s type should, roughly speaking, increase (weakly) her probability of receiving an object *sooner*.⁶ More precisely, the expected discounted sum of each agent’s allocation probabilities must be nondecreasing in that agent’s value, regardless of the state of the market upon her arrival. Moreover, the expected payoffs of a buyer in any two mechanisms with the same allocation rule can differ only by a constant, and hence expected payoffs and revenue are pinned down by the expected discounted probabilities of receiving an object.

3. EFFICIENT IMPLEMENTATION

3.1. Preliminaries: An Efficient Direct Mechanism

Before discussing the properties of efficient mechanisms in this setting, we first describe the socially efficient policy. Recall that, in the static single-object allocation setting, allocative efficiency is equivalent to allocating the object to the highest-valued buyer. In our dynamic setting, the structure of the environment—the nature of the arrival processes and the non-storability of objects—implies that the socially efficient policy is essentially an assortative matching. In particular, objects are ordered by their arrival time and buyers are ordered by their values, and “earlier” objects are allocated to higher-valued buyers. Of course, the feasibility constraints imposed by the dynamic nature of the agent population have an impact on the nature of the efficient policy, as the ordering of buyers by valuation need not correspond to the sequential ordering of buyers by their periods of availability. Thus, the socially efficient allocation policy is, in any given period, to allocate all available objects to the set of buyers currently present that have the highest values.

Formally, we consider a social planner who, at time zero, chooses a feasible dynamic direct mechanism $\mathcal{M} = \{\mathbf{x}_t, \mathbf{p}_t\}_{t \in \mathbb{N}_0}$. The planner’s goal is to maximize allocative efficiency; that is, the planner wishes to choose allocations $x_{i,t}$ to maximize

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{i \in \mathcal{I}} \delta^t x_{i,t}(h_t, \omega_t, \mathbf{v}) v_i \right],$$

where the expectation is taken with respect to the arrival and departure processes, as well as agents’ valuations. Recalling that $q_i(v_i, \omega_t)$ from Equation (1) is agent i ’s expected discounted probability of receiving an object (conditional on entry), we may rewrite this objective function as

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^t \alpha_t(i) q_i(v_i, \omega_t) v_i \right],$$

where $\alpha_t(i) = 1$ if $i \in \mathcal{I}_t$ arrives to the market and zero otherwise. (Recall that this arrival occurs with probability $\pi_i \in [0, 1]$ for each agent i .)

⁶This is a consequence of the fact that there are now multiple opportunities to receive an object: if a buyer does not receive an object in the period of her arrival, she may still receive one in future periods.

Before a formal statement of our result, a few additional definitions are necessary. Fix any state $z_t = (h_t, \omega_t, \mathbf{v}) \in \mathcal{H}_t \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \times \mathcal{R}$, where \mathbf{v} denotes the truthful reporting strategy by all agents. We denote by

$$\mathcal{A}_+(z_t) := \{i \in \mathcal{A}_t : |\{j \in \mathcal{A}_t : v_j \geq v_i\}| \leq k_t\}$$

the set of agents who are among the k_t highest-ranked buyers at state z_t . Similarly, the set of agents who are ranked *outside* the top k_t agents is denoted by

$$\mathcal{A}_-(z_t) := \{i \in \mathcal{A}_t : |\{j \in \mathcal{A}_t : v_j > v_i\}| \geq k_t\}.$$

Finally, we let

$$\mathcal{A}_\sim(z_t) := \mathcal{A}_t \setminus (\mathcal{A}_+(z_t) \cup \mathcal{A}_-(z_t))$$

denote the set of agents who are “on the boundary”—the set of agents that are tied for the k_t -th highest rank.

LEMMA 2 (Efficient allocation rules).

Suppose all buyers, upon arrival, report their true values. A feasible allocation rule $\{\mathbf{x}_t\}_{t \in \mathbb{N}_0}$ is efficient if, and only if, for all states $z_t = (h_t, \omega_t, \mathbf{v}) \in \mathcal{H}_t \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} \times \mathcal{R}$,

$$x_{i,t}(z_t) = 1 \text{ for all } i \in \mathcal{A}_+(z_t)$$

and

$$\sum_{i \in \mathcal{A}_\sim(z_t)} x_{i,t}(z_t) = k_t - |\mathcal{A}_+(z_t)| \text{ if } |\mathcal{A}_+(z_t)| < k_t.$$

PROOF. The proof may be found in [Appendix A](#). □

Note that the conditions in this lemma pin down the behavior of efficient allocation rules after almost all histories.⁷ The second condition applies only in the case of “ties” among the agents, which are probability zero events. Additionally, notice that the period- t efficient allocation does not depend on past allocations or history; only the set of objects available (indicated by k_t), the set of agents present at time t (indicated by α_t), and these agents’ reported values (denoted by \mathbf{v}_t) are relevant. Therefore, we will henceforth restrict attention to the efficient allocation rule which breaks ties with equal probability, defined by

$$\hat{x}_{i,t}(\omega_t, \mathbf{v}_t) := \begin{cases} 1 & \text{if } i \in \mathcal{A}_+ \\ 0 & \text{if } i \in \mathcal{A}_- \\ \frac{k_t - |\mathcal{A}_+|}{|\mathcal{A}_\sim|} & \text{if } i \in \mathcal{A}_\sim \end{cases}$$

for all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$ and $\mathbf{v}_t \in \mathbf{V}^{\mathcal{A}_t}$ (and where we have dropped the dependence on histories).

It is straightforward to show that, as in the standard static allocation problem, this first-best socially efficient policy is implementable. In particular, we may use the dynamic pivot mechanism of [Bergemann and Välimäki \(2008\)](#), a dynamic variant of the standard Vickrey-Clarke-Groves mechanism, in order to achieve the efficient outcome. By choosing payments which leave agents with

⁷While it is straightforward to do so, we do not formally account for the zero-probability event in which a buyer’s value equals zero. This simplifies both notation and exposition while leaving our results unchanged.

their *flow* contribution to the social welfare in each period, the dynamic pivot mechanism obtains truth-telling as an equilibrium which implements the efficient policy.

In order to fully describe the dynamic pivot mechanism, we require the following definitions. For any $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$ and truthful $\mathbf{v}_t \in \mathbf{V}^{\mathcal{A}_t}$, let

$$W(\omega_t, \mathbf{v}_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s, \mathbf{v}_s) v_j \right]$$

denote the social welfare (from period $t \in \mathbb{N}_0$ on) when the efficient policy $\hat{\mathbf{x}}$ is implemented. Denoting by ω_s^{-i} the state of the market in period $s \in \mathbb{N}_0$ when agent i has been removed from the market (that is, where we impose $\alpha_s(i) = 0$), we write

$$W_{-i}(\omega_t^{-i}, \mathbf{v}_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I} \setminus \{i\}} \delta^{s-t} \hat{x}_{j,s}(\omega_s^{-i}, \mathbf{v}_s) v_j \right].$$

for the social welfare (from period $t \in \mathbb{N}_0$ on) when i is removed from the market and the efficient policy $\hat{\mathbf{x}}$ is implemented. Agent i 's flow contribution in period $t \in \mathbb{N}_0$ is simply the difference between i 's total contribution to the social welfare and her expected *future* contribution:

$$\hat{w}_i(\omega_t, \mathbf{v}_t) := \underbrace{W(\omega_t, \mathbf{v}_t) - W_{-i}(\omega_t^{-i}, \mathbf{v}_t)}_{\text{total contribution}} - \underbrace{\delta \left(\mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) \right] - \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) \right] \right)}_{\text{expected future contribution}}.$$

Thus, we may define the dynamic pivot mechanism as the mechanism $\widehat{\mathcal{M}} := \{\hat{\mathbf{x}}_t, \hat{\mathbf{p}}\}_{t \in \mathbb{N}_0}$, where the payment rule $\hat{\mathbf{p}}$ is defined by

$$\hat{p}_{i,t}(\omega_t, \mathbf{v}_t) := \hat{x}_{i,t}(\omega_t, \mathbf{v}_t) v_i - \hat{w}_i(\omega_t, \mathbf{v}_t)$$

for all (ω_t, \mathbf{v}_t) . This mechanism yields to each agent flow payoffs equal to her flow contribution to the social welfare, and has been shown to be efficient. In addition, the dynamic pivot mechanism is periodic ex post incentive compatible. Introduced by Bergemann and Välimäki (2008), this notion requires that, given expectations about future behavior and arrivals, each buyer's current-period behavior remains optimal even after observing all current-period private information.

LEMMA 3 (Implementability and efficiency of the dynamic pivot mechanism).

Truth-telling in the dynamic pivot mechanism is periodic ex post incentive compatible and individually rational, thereby implementing the socially efficient policy.

PROOF. The result follows immediately from Bergemann and Välimäki (2008, Theorem 1).⁸ \square

Note that, unlike the VCG mechanism in static settings, the dynamic pivot mechanism $\widehat{\mathcal{M}}$ is *not* dominant-strategy incentive compatible. This is true even though buyers participating in our variant of this mechanism make only a single report upon their arrival (as compared to Bergemann and Välimäki (2008), for instance, where buyers make a report in every period). This is because,

⁸Bergemann and Välimäki (2008) show this result in the context of a fixed agent population. Parkes and Singh (2003) and Cavallo, Parkes, and Singh (2009) demonstrate that "Online VCG," another dynamic variant of the VCG mechanism, implements the efficient policy in the presence of an evolving agent population.

in contrast to the static VCG mechanism, payments are *not* distribution-free, and hence agents' beliefs about other players and their strategies matter for the determination of optimal behavior.⁹

3.2. An Efficient Sequential Auction

It is important to keep in mind that the dynamic pivot mechanism is a direct revelation mechanism, relying on a planner to aggregate the reported values of each buyer in order to determine allocations and payments. This raises an important question: does this efficient mechanism correspond to a familiar auction format? In the static single-object case, Vickrey (1961) provided a clear answer: the analogue of the Vickrey-Clarke-Groves mechanism for the allocation of a single indivisible good is the second-price auction. Both the sealed-bid second-price auction and the ascending (English) auction admit equilibria that are outcome equivalent to the VCG mechanism and are compelling prescriptions for “real-world” behavior.¹⁰

A reasonable conjecture is that a sequence of auctions would be useful in the context of a sequential allocation problem. But what auction format would be desirable? The “standard” analogue of the VCG mechanism in static settings is the second-price sealed-bid auction. In a dynamic setting where all buyers arrive before the initial period, Kittsteiner, Nikutta, and Winter (2004) show that a sequence of such auctions is equivalent to VCG. When buyers arrive stochastically as in our model, however, a sequence of second-price auctions does *not* correspond to the dynamic formulation of the Vickrey-Clarke-Groves mechanism discussed above.

This failure of equivalence is due to the fact that a buyer participating in a sequence of auctions has available to her an option: by losing in the current auction, she gains the ability to participate in future elements of the sequence. Let us denote the expected value of this future participation by δV . Rational bidding behavior in a second-price sealed-bid auction then requires shading one's bid downwards by the value of this option—our bidder chooses her bid b_i to maximize her expected payoff, solving

$$\max_{b_i} \left\{ \Pr \left(b_i > \max_{j \neq i} \{b_j\} \right) \mathbb{E} \left[v_i - \max_{j \neq i} \{b_j\} \right] + \Pr \left(b_i < \max_{j \neq i} \{b_j\} \right) \delta V \right\}.$$

Since the probability of winning and the probability of losing sum to one, we may rearrange the above expression into an equivalent optimization problem:

$$\max_{b_i} \left\{ \Pr \left(b_i > \max_{j \neq i} \{b_j\} \right) \mathbb{E} \left[(v_i - \delta V) - \max_{j \neq i} \{b_j\} \right] \right\} + \delta V.$$

This, however, is exactly the problem faced by a bidder in a static second-price sealed-bid auction when her true value is given by $v_i - \delta V$; standard dominance-type arguments show that it is then optimal to bid

$$b_i^* = v_i - \delta V.$$

What, then, is this option value? Clearly, it is an expectation of future payoffs from participating in the sequence of auctions, and hence will depend upon expected future prices. But these future

⁹Thanks are due to Ilya Segal for pointing out the difficulty in achieving dominant-strategy implementation in sequential-move mechanisms, thereby correcting an error in a previous version of this work.

¹⁰Of course, the revenue equivalence theorem applies, and several other standard auction mechanisms are able to yield efficient outcomes in the single-object static setting. However, they are not *outcome equivalent* to the VCG mechanism.

prices are determined by the valuations of one's competitors, and hence the continuation value V is itself a function of those valuations; that is,

$$V = V(v_i, v_{-i}).$$

Thus, despite the fact that we have started in an independent private-values framework, market dynamics (in particular, repeated competition across time) *generate* interdependence: buyers must learn their competitors' values in order to correctly "price" the option of future participation.¹¹ Moreover, this learning is not possible when using a second-price sealed-bid auction (or any other sealed-bid auction format, for that matter), as the auction format simply does not reveal sufficient information to market participants, and buyers will have to bid based on their expectations and beliefs about their competitors:

$$b_i^* = v_i - \mathbb{E} [\delta V(v_i, v_{-i})].$$

In sharp contrast, however, to the sequential auctions of Kittsteiner, Nikutta, and Winter (2004) (where all buyers are present throughout the sequence of auctions), bidders that arrive to the market at different times will have observed different histories. These bidders will therefore have asymmetric beliefs about their competitors, and hence asymmetric expectations. This leads to asymmetry in bidding behavior, which in turn generates inefficient outcomes.

A similar problem arises when considering the use of a second-price auction where bids are revealed each period after the allocation of objects. In particular, in any period in which there are new entrants, there will be buyers who are uninformed—and about whom incumbent buyers are uninformed. Again, these two groups will have differential information, and hence differential beliefs, thereby leading to inefficient outcomes. Note that this occurs despite the fact that bids are being revealed. Since information revelation is occurring *after* the auction is over, buyers are unable to condition their bids on that information. Instead, information revealed in the current period can be used only in subsequent periods; information revelation is occurring too "slowly" for information about others to be incorporated into current-period bidding.

This suggests the need for an *open* auction format, and in particular the ascending price auction. In such an auction, a price clock rises continuously and buyers drop out of the auction at various points. This allows buyers to observe the points at which their competitors exit the auction and make inferences about their valuations. These inferences can then be incorporated into *current-period* bidding, leading to bids that correctly account for the interdependence generated by market dynamics: buyers can arrive at the efficient outcome by submit bids

$$b_i^* = v_i - \delta V(v_i, v_{-i}),$$

To be specific, we make use of a simple generalization of the Milgrom and Weber (1982) "button" model of ascending auctions. In particular, we consider a multi-unit, uniform-price variant of their model. The auction begins, *in each period*, with the price at zero and with all agents present participating in the auction. Each bidder may choose any price at which to drop out of the auction.

¹¹It is well-known (see, for instance, Dasgupta and Maskin (2000) or Krishna (2003)) that efficient implementation is not guaranteed when values are interdependent. Note, however, that interdependence arises in our setting due to the dynamics of the indirect mechanism and the resulting option values—the underlying environment for direct mechanisms is a standard independent private values setting.

This exit decision is irreversible (in the current period), and is observable by all agents currently present. Thus, the current price and the set of active bidders is commonly known throughout the auction. When there are $m \geq 1$ objects for sale, the auction ends whenever at most m active bidders remain, with each remaining bidder receiving an object and paying the price at which the auction ended. Note that if there are fewer than m bidders initially, then the auction ends immediately at a price of zero. In addition, suppose that several bidders drop out of the auction simultaneously, leaving $m' < m$ bidders active. The auction ends at this point, and $m - m'$ of the “tied” bidders are selected with equal probability to receive an object—along with the remaining active bidders—paying the price at which the auction closed.¹² With this in mind, each bidder’s decision problem *within* a given period is not the choice of a single bid, but is instead the choice of a sequence of functions, each of which determines an exit price contingent on the (observed) exit prices of the bidders who have already exited the current auction. Therefore, over the course of the auction, buyers gradually reveal their private information to their competitors.

This process of gradual information revelation leads to an additional asymmetry, however. Consider the group of buyers who participated in an auction in period t but lost. At the end of the auction, they will have observed each others’ drop-out prices and inferred each others’ values, implying that at the beginning of period $t + 1$, they have essentially perfect and complete information about one another. But in period $t + 1$, a new group of buyers, about whom nothing is known, arrives on the market. We therefore have two differentially informed groups of buyers. Moreover, if we want to achieve an efficient outcome, these new entrants must be induced to reveal their private information despite being asymmetrically informed.

This asymmetry may be resolved via a process of *information renewal*: full revelation of *all* private information in *every* period. This is achieved by using “memoryless” strategies: incumbent buyers disregard their observations and information from previous periods and behave “as though” they are uninformed. By doing so, all buyers are able to *behave* symmetrically, thereby allowing newly arrived buyers to learn about their current competitors without knowledge of the events of previous periods. All buyers, incumbents and new entrants alike, are thus provided with the appropriate incentives to participate in the process of information revelation.

Note that this equilibrium in memoryless strategies is *not* the result of an a priori restriction on the set of strategies available to buyers. Rather, the use of memoryless strategies is the result of fully rational and unconstrained optimizing behavior. Buyers have perfect recall of the past, and also have available to them the option of conditioning on information revealed in previous periods. Ignoring that information, however, is an equilibrium best response to the behavior of their competitors.

Why would a fully-rational, utility-maximizing buyer “throw away” payoff-relevant information from previous periods and behave as though she were uninformed? Recall that bidders

¹²The sequential ascending auction mechanism we propose bears some resemblance to the multi-unit auction mechanism of Demange, Gale, and Sotomayor (1986). In fact, in an environment where all buyers are present in the initial period and without supply uncertainty, the two mechanisms arrive at equivalent outcomes. Recall, however, that their mechanism operates by making use of an auctioneer who raises the prices on “over-demanded” sets of objects. This is precluded in our setting with a dynamic population, as it is impossible to determine which future objects will be over-demanded before either the objects or the agents who desire them arrive to the market.

engage in the process of information revelation and renewal in *every* period. Buyers therefore anticipate that (in equilibrium) all private information they may have observed in the past will be revealed again. In particular, there is no need for buyers to condition their behavior at the outset of the current period on the past—buyers expect any payoff-relevant information to be revealed anew as the price clock rises over the course of the current-period auction, allowing them to condition their strategies on this information as it is revealed (or re-revealed) *during* the current period.

With this in mind, we now informally describe the memoryless strategies used by each player in the sequential ascending auction mechanism.¹³ In each period, buyers will remain active in the auction until the price reaches the point at which they are exactly indifferent between winning immediately and participating in future periods. Moreover, these buyers believe that any future prices they pay will equal the externality that they impose on the market; equivalently, they believe that the option value of future participation equals their expected future marginal contribution to the social welfare:

$$\delta V(v_i, v_{-i}) = \delta \mathbb{E} [\widehat{w}_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (v_i, v_{-i})].$$

As competitors drop out of the auction and reveal their private information, each remaining buyer recalculates this expected continuation value and redetermines her optimal drop-out point.

To formally describe the strategies, let $n_t := |\mathcal{A}_t| = \sum_{j \in \mathcal{I}} \alpha_t(j)$ denote the number of buyers present in period t . In addition, taking the perspective of an arbitrary bidder i , let

$$\mathbf{y}_t := (y_1^t, \dots, y_{n_t-1}^t)$$

denote the ordered valuations of all *other* buyers present in period t , where y_1^t is the largest value, and $y_{n_t-1}^t$ is the smallest. Finally, for each $m = 1, \dots, n_t - 1$, let

$$\bar{\mathbf{v}}^m := (\bar{v}, \dots, \bar{v}) \in \mathbf{V}^m \text{ and } \mathbf{y}_t^{>m} := (y_{m+1}^t, \dots, y_{n_t-1}^t).$$

If all buyers use symmetric strictly increasing bidding strategies within a period, the prices at which buyers exit the auction will reveal their values. Over the course of the current-period auction, buyers will then observe $\mathbf{y}_t^{>m}$, allowing their bids to be conditioned on this information.

Finally, we define, for each $m = 1, \dots, n_t - 1$,

$$w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) := \delta \mathbb{E} [\widehat{w}_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, \mathbf{y}_t^{>m})]$$

to be the (discounted) expected future marginal contribution of an agent $i \in \mathcal{A}_t$ with value v_i , where the expectation is conditional on the period- t presence of m competitors each with the highest possible value \bar{v} and $n_t - m - 1$ buyers ranked below i with values $\mathbf{y}_t^{>m}$.

With these preliminaries in hand, we may now define the strategies used by each bidder. We define, for each $m = 1, \dots, n_t - 1$,

$$\widehat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) := v_i - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}). \quad (3)$$

We will assume that, in each period $t \in \mathbb{N}_0$, each agent $i \in \mathcal{A}_t$ bids according to $\widehat{\beta}_{m, n_t}^t$ whenever she has m active competitors in the auction. Thus, each buyer i initially bids up to the point at which

¹³See Said (2008) for an explicit (closed-form) example of these strategies in a special case of the present model.

she is indifferent between winning the object at the current price and receiving her discounted expected marginal contribution in the next period, where the expectation is conditional on being the lowest-ranked of the n_t bidders currently present and all other bidders having the highest possible valuation.¹⁴ Note that these initial bids depend only upon the current state of the market and each buyer's *own* private information, and not on the past history of the market.

Observe that, if $v_i > v_j$, then

$$\begin{aligned} & w^{t+1}(\omega_t, v_j) - w^{t+1}(\omega_t, v_i) \\ &= \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] - \delta \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] \\ &\quad - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] + \delta \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] \\ &= \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right], \end{aligned}$$

since removing either i or j in the next period, conditional on her being the lowest-ranked agent, does not affect the order of anticipated future allocations to any other agents. Moreover, note that by treating buyer j as though her true value were v_i , we can provide a bound for the difference above. In particular, we have

$$\begin{aligned} & \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] \\ & \geq \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] (v_j - v_i). \end{aligned}$$

Thus, if $v_i > v_j$, then

$$\hat{\beta}_{n_t-1, n_t}^t(\omega_t, v_i) - \hat{\beta}_{n_t-1, n_t}^t(\omega_t, v_j) \geq \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] \right) (v_i - v_j) > 0$$

since the discounted expected probability of receiving an object in the future is bounded above by $\delta < 1$. Thus, the agent who is, in fact, the lowest-ranked buyer present in period t will be the first to drop out of the period- t auction, publicly revealing her value.

At this point, each remaining buyer i bids until she is indifferent between winning the object at the current price and receiving her discounted expected marginal contribution, conditional on the knowledge that she is the second-lowest ranked of the $n_t - 1$ bidders remaining active in the auction, that all remaining active bidders have the highest possible valuation and that the lowest-ranked buyer present has value $y_t^{n_t-1} < v_i$. In addition, suppose that buyer j with value v_j was the first to exit the auction. Then $y_t^{n_t-1} = v_j < v_i$ implies

$$\begin{aligned} & w^{t+1}(\omega_t, v_j) - w^{t+1}(\omega_t, v_i, v_j) \\ &= \delta \left(\mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] - \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_j) \right] \right) \\ &\quad - \delta \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_j) \right] - \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] \right). \end{aligned}$$

¹⁴This is not strictly necessary; any beliefs about the valuations of her opponents will suffice as long as the support of those beliefs is contained in the interval $(v_i, \bar{v}]$.

However, the second difference above may be rewritten as

$$\begin{aligned} & \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_j) \right] - \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_i) \right] \right) \\ & + \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_i) \right] - \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_j) \right] \right). \end{aligned}$$

Thus, the difference in future expected contributions

$$w^{t+1}(\omega_t, v_j) - w^{t+1}(\omega_t, v_i, v_j)$$

is the sum of three differences: the first is the gain in social welfare when increasing i 's value from v_i to \bar{v} ; the second is the gain in social welfare (when i is not on the market) from increasing j 's value from v_j to v_i ; and finally, the third is the loss in social welfare (when j is not present) from decreasing i 's value from \bar{v} to v_i . However, since $v_j < v_i$, the presence or absence of j from the market has no influence on when the efficient policy allocates to i , regardless of whether i 's value is v_i or \bar{v} . Therefore, the gain from the first difference equals the loss from the third, implying that

$$\begin{aligned} & w^{t+1}(\omega_t, v_j) - w^{t+1}(\omega_t, v_i, v_j) \\ & = \delta \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_j) \right] - \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-2}, v_i, v_i) \right] \right). \end{aligned}$$

A bounding argument similar to the one previously applied may then be used to show that the difference in bids

$$\hat{\beta}_{n_t-2, n_t}^t(\omega_t, v_i, v_j) - \hat{\beta}_{n_t-1, n_t}^t(\omega_t, v_j) > 0.$$

Thus, there is ‘‘continuity’’ at the first drop out point, in the sense that the exit of the lowest-valued buyer does not induce the immediate exit of any buyer with a (strictly) higher value. Therefore, if $\hat{\beta}_{n_t-2, n_t}^t(\omega_t, v_i, v_j)$ is increasing in v_i , the price at which the second exit occurs fully reveals the value of the second-lowest ranked buyer. Similar logic may be used to show that $\hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{> m})$ is strictly increasing in v_i for all m , and that the ‘‘continuity’’ property described above holds after every exit from the auction.

LEMMA 4 (Bids are fully separating).

For all $m = 1, \dots, n_t - 1$, $\hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{> m})$ is increasing in v_i . Moreover, if $v_i > \mathbf{y}_t^{m+1}$, then

$$\hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{> m}) > \hat{\beta}_{m+1, n_t}^t(\omega_t, \mathbf{y}_t^{> m}).$$

PROOF. The proof may be found in [Appendix A](#). □

Since the bids are fully separating, the efficient allocation is achieved when all buyers follow the prescribed strategies. If these strategies form an equilibrium of the sequential ascending auction (an assumption we will shortly verify), [Corollary 1](#) implies that expected payments by buyers in this mechanism must be the same, up to a constant, as those in the dynamic pivot mechanism. We may prove, however, a stronger equivalence result.

THEOREM 1 (Outcome equivalence of direct and indirect mechanisms).

Following the bidding strategies $\hat{\beta}_{m, n_t}^t$ in every period t in the sequential ascending auction mechanism is outcome equivalent to the dynamic pivot mechanism.

PROOF. In order to establish that the bidding strategies $\hat{\beta}_{m,n_t}^t$ lead to an efficient allocation, we require the following result demonstrating that the bids are, in fact, fully separating.

Thus, following the bidding strategies prescribed in Equation (3) leads to an efficient allocation. To see that these strategies also lead to payments equal to those of the dynamic pivot mechanism, fix an arbitrary period $t \in \mathbb{N}_0$, and let k_t denote the number of objects present, and $n_t := |\mathcal{A}_t|$ denote the number of agents present. As discussed above, the bidding strategies $\hat{\beta}_{m,n_t}^t$ are strictly increasing; therefore, the multi-unit uniform-price ascending auction ends allocates the k_t objects to the group of buyers with the k_t highest values. Recall that if $k_t \geq n_t$, the auction ends immediately, and all buyers present receive an object for free. Similarly, in the dynamic pivot mechanism, each buyer i receives an object, and makes a payment $\hat{p}_{i,t}$ given by

$$\hat{p}_{i,t}(\omega_t, \mathbf{v}_t) = v_i - w_i(\omega_t, \mathbf{v}_t),$$

where w_i is the agent's marginal contribution to the social welfare. Note that since i is receiving an object, her total and flow marginal contributions are equal. However, since there are sufficient objects present for each agent to receive one, i does not impose any externalities on the remaining agents; thus,

$$w_i(\omega_t, \mathbf{v}_t) = \hat{w}_i(\omega_t, \mathbf{v}_t) = v_i,$$

implying that $\hat{p}_{i,t}(\omega_t, \mathbf{v}_t) = 0$. In this case, then, the allocation and payments of the auction mechanism and the dynamic pivot mechanism are the same.

Suppose instead that $k_t < n_t$; that is, there are more agents present than objects. Denote by i_m the bidder with the m -th highest value. Then each agent who receives an object pays the price at which buyer i_{k_t+1} drops out of the auction, which is given by

$$\hat{\beta}_{k_t+1,n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}) = v_{i_{k_t+1}} - w^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}).$$

In the dynamic pivot mechanism, on the other hand, each agent i who receives an object pays

$$\begin{aligned} \hat{p}_{i,t}(\omega_t, \mathbf{v}_t) &= v_i - w_i(\omega_t, \mathbf{v}_t) \\ &= v_i - \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s, \mathbf{v}_s) v_j \right] + \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I} \setminus \{i\}} \delta^{s-t} \hat{x}_{j,s}(\omega_s^{-i}, \mathbf{v}_s) v_j \right] \\ &= v_i - \left(\sum_{m=1}^{k_t} v_m + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s, \mathbf{v}_s) v_j \right] \right) \\ &\quad + \left(\sum_{m=1}^{k_t} v_m + (v_{i_{k_t+1}} - v_i) + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s^{-i, -i_{k_t+1}}, \mathbf{v}_s) v_j \right] \right) \\ &= v_{i_{k_t+1}} - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s, \mathbf{v}_s) v_j \right] + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \hat{x}_{j,s}(\omega_s^{-i, -i_{k_t+1}}, \mathbf{v}_s) v_j \right] \\ &= v_{i_{k_t+1}} - w^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}), \end{aligned}$$

where the final equality follows from the fact that $w^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}})$ is defined to be the expected future marginal contribution of the agent with the $(k_t + 1)$ -th highest value, conditional on agents with higher values (which includes i) receiving an object today.

Thus, following the bidding strategies $\hat{\beta}_{m,n_t}^t$ leads to period- t prices and allocations identical to those of the dynamic pivot mechanism. Since the period t was arbitrary, as was the state ω_t , this equivalence holds after each history. Thus, the two mechanisms are outcome equivalent. \square

Therefore, following the memoryless bidding strategies prescribed in Equation (3) leads to an outcome that is identical to that of truth-telling in the dynamic pivot mechanism. Moreover, we know from Lemma 3 that truth-telling is an equilibrium of the dynamic pivot mechanism. It remains to be shown, however, that the bidding strategies described in Equation (3) form an equilibrium of the sequential ascending auction mechanism.

Since the sequential ascending auction mechanism is a dynamic game of incomplete information, the equilibrium concept we use is perfect Bayesian equilibrium. This solution concept requires that behavior be sequentially rational with respect to agents' beliefs, and that agents' beliefs be updated in accordance with Bayes' rule wherever possible. Since all buyers use the strictly increasing bidding strategies $\hat{\beta}_{m,n_t}^t$, behavior along the equilibrium path is perfectly separating, implying that Bayesian updating fully determines beliefs. To determine optimality off the equilibrium path, however, we need to consider the beliefs of bidders after a deviation. Since such post-deviation histories are zero probability events, we are free to choose arbitrary off-equilibrium beliefs. Therefore, we will suppose that, after a deviation, buyers disregard their previous observations, believing that the deviating agent is *currently* sincerely revealing her value in accordance with $\hat{\beta}_{m,n_t}^t$.

This particular specification of off-equilibrium beliefs is particularly useful. Note that these beliefs are consistent with Bayes' rule even after probability zero histories. This follows immediately from the fact that, generally, this system of beliefs consists of point-mass beliefs about the types of other agents. The only agents about whom beliefs do not take this form are those that have yet to arrive to the market and those who win an object in the period of their arrival—these agents reveal only a lower bound on their value.

Moreover, this property is equivalent to the condition of *preconsistency* of beliefs in an extensive form game of incomplete information put forth by Hendon, Jacobsen, and Sloth (1996), which Perea (2002) shows to be both necessary and sufficient for the one-shot-deviation principle to hold.¹⁵ This is an important observation, as perfect Bayesian equilibrium, in contrast to sequential equilibrium, need not satisfy the one-shot-deviation principle.¹⁶ We can therefore prove the following result.

THEOREM 2 (Equilibrium in the sequential ascending auction).

Suppose that in each period, buyers bid according to the memoryless strategies described in Equation (3). This strategy profile, combined with the system of beliefs described above, forms a (periodic ex post) perfect Bayesian equilibrium of the sequential ascending auction mechanism.

¹⁵This condition is called *updating consistency* by Perea (2002), and is also equivalent to part 3.1(1) of Fudenberg and Tirole (1991)'s definition of a *reasonable* assessment.

¹⁶These off-equilibrium beliefs also satisfy the “no-signaling-what-you-don't know condition” in Fudenberg and Tirole (1991). This suggests that (aside from measurability issues) one could construct a conditional probability system for this equilibrium that satisfies Fudenberg and Tirole's conditions for perfect extended Bayesian equilibrium. The set of all such equilibria coincides, in finite games, with the set of sequential equilibria.

PROOF OF THEOREM 2. We prove this proposition by making use of the one-shot deviation principle. Consider any period with $n_t := |\mathcal{A}_t|$ buyers on the market and k_t objects present. Suppose that all bidders other than player i are using the conjectured equilibrium strategies. We must show that bidder i has no profitable one-shot deviations from the collection of cutoff points $\{\hat{\beta}_{m,n_t}^t\}$. More specifically, we must show that i does not wish to exit the auction earlier than prescribed, nor does she wish to remain active later than specified.

Once again labeling agents such that buyer i_1 has the highest value and buyer i_{n_t} has the lowest, note that if $v_i < v_{i_{k_t}}$, bidding according to $\{\hat{\beta}_{m,n_t}^t\}$ implies that i does not win an object in the current period. Therefore, exiting earlier than specified does not affect i 's current-period returns. Moreover, since the bidding strategies are memoryless, neither future behavior by i 's competitors nor i 's future payoffs will be affected by an early exit. Suppose, on the other hand, that i has one of the k_t highest values; that is, that $v_i \geq v_{i_{k_t}}$. As established by [Theorem 1](#), i receives an object, paying a price such that her payoff is exactly equal to her marginal contribution to the social welfare. Deviating to an early exit, however, leads to agent i_{k_t+1} winning an object instead of buyer i . Moreover, i 's expected payoff is then $w^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}})$, which we defined as i 's future expected marginal contribution. This is a profitable one-shot deviation for i if, and only if,

$$w^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}) \geq v_i - \hat{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\hat{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}) \geq v_i - w^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}) = \hat{\beta}_{k_t, n_t}^t(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\hat{\beta}_{k_t, n_t}^t$ in [Equation \(3\)](#). Since $v_i > v_{i_{k_t+1}}$, this contradicts the efficiency property established by [Theorem 1](#). Thus, i does not wish to exit early.

Alternately, if $v_i \geq v_{i_{k_t}}$, then planning to remain active in the auction *longer* than specified does not change i 's payoffs, as i will win an object regardless. If, on the other hand, $v_i < v_{i_{k_t}}$, then delaying exit from the period- t auction can affect i 's payoffs. Since bids in future periods do not depend on information revealed in the current period, this only occurs if i remains in the auction long enough to win an object. If i wins, she pays a price equal to the exit point of $i_{v_{k_t}}$, whereas if she exits, she receives as her continuation payoff her marginal contribution to the social welfare. So, suppose that $i = i_m$ for some $m > k_t$. Then a deviation to remaining active in the auction is profitable if, and only if,

$$v_m - \hat{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t}}, \dots, v_{i_{m-1}}, v_{m+1}, \dots, v_{i_{n_t}}) \geq w^{t+1}(\omega_t, v_m, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\hat{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t}}, \dots, v_{i_{m-1}}, v_{m+1}, \dots, v_{i_{n_t}}) \leq v_m - w^{t+1}(\omega_t, v_m, \dots, v_{i_{n_t}}) = \hat{\beta}_{m-1, n_t}^t(\omega_t, v_m, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\hat{\beta}_{m-1, n_t}^t$ in [Equation \(3\)](#). As above, the fact that $v_m < v_{i_{k_t}}$ contradicts the efficiency property established by [Theorem 1](#). Therefore, i does not desire to remain active in the auction long enough to receive an object.

Thus, we have shown that no player has any incentive to deviate from the prescribed strategies when on the equilibrium path. In particular, using the bidding strategies $\hat{\beta}_{m,n_t}^t$ is sequentially rational given players' beliefs along the equilibrium path. Recall, however, that we have specified off-equilibrium beliefs such that buyers "ignore" their past observations when they observe a deviation from equilibrium play, updating their beliefs to place full probability on the valuation that rationalizes the deviation; they believe that the deviating agent is *currently* being truthful with regards to the strategies $\hat{\beta}_{m,n_t}^t$. The argument above then implies that continuing to bid according to the specified strategies remains sequentially rational with respect to these updated beliefs. Thus, bidding according to the cutoffs in Equation (3) is optimal along the entire game tree: this strategy profile forms a perfect Bayesian equilibrium of the sequential ascending auction mechanism. Moreover, observe that the arguments above consider the ex post profitability of deviations; the lack of profitable deviations even when there is no uncertainty about the valuations of current-period competitors implies that this equilibrium is, in fact, a periodic ex post equilibrium. \square

Theorems 1 and 2 jointly imply that the sequential ascending auction admits an efficient equilibrium that also yields prices identical to those of the dynamic pivot mechanism.¹⁷ The sequential ascending auction is therefore a natural, intuitive institution that yields efficient outcomes.

It is interesting to note several additional properties of this equilibrium. First, the proof shows that deviations from the bidding strategies $\hat{\beta}_{m,n_t}^t$ are not rational for any agent, even when conditioning on competitors' values in the current period. Thus, the strategy profile specified in Equation (3) forms a periodic ex post equilibrium. Since agents essentially "report" their values in each auction, the extensive-form structure of the indirect mechanism allows for a much larger number of potential deviations from truthful behavior as compared to the direct mechanisms discussed earlier. Despite this, however, there is no loss in the "strength" of implementation: equilibrium in both the direct and indirect mechanisms involves the same notion of periodic ex post equilibrium.

Furthermore, observe that in the sequential ascending auction, buyers have the ability to drop out immediately once an auction begins. Since the bidding strategies discussed above form a periodic ex post equilibrium, buyers do not wish to take advantage of this possibility, even if they know their opponents' values. Thus, although we have assumed that buyers cannot conceal their presence when arriving to the market, we may conclude that, in equilibrium, they would not take advantage of that opportunity were it afforded to them—the equilibrium we have constructed remains an equilibrium in the "larger" game where buyers may conceal their presence.

4. REVENUE MAXIMIZATION

While the previous section discusses efficient direct and indirect mechanisms, we have said little about revenue or optimality. In the static setting, Myerson (1981) showed that the optimal mechanism for selling a single indivisible unit is a Vickrey-Clarke-Groves mechanism, with the caveat that instead of allocating the good to the agent with the highest value, the seller allocates the object to the agent with the highest *virtual* value. Maskin and Riley (1989) extend Myerson's

¹⁷We should point out that this equilibrium is not unique. Analogous to the multiplicity of equilibria described by Bikhchandani, Haile, and Riley (2002), there exists a continuum of outcome-equivalent equilibria that differ only in the "speed" of information revelation within each period, among which the equilibrium we describe is the "slowest."

insights to the setting in which multiple identical units are offered for sale and show that, as in the single-unit case, the objects are allocated to the set of buyers with the highest virtual valuations.

In our setting, however, while the objects are individual units of a homogeneous good, from the perspective of an individual buyer, they are differentiated products. To make this clear, consider a buyer i with value v_i who is present at period t . If this buyer receives an object in period t , this yields her utility v_i . However, if she anticipates receiving an object in period $t + 1$, her valuation for that object is δv_i . Thus, she does not value the two objects identically. While there does exist a literature on auctions for multiple heterogeneous objects, much of the focus has been on efficiency and not revenue maximization.¹⁸ Thus, paralleling the previous section, we will describe an optimal dynamic direct mechanism for a revenue-maximizing seller. We show that revenue maximization in our setting with a dynamic population is achieved, as in Myerson (1981), by applying an efficient mechanism to virtual values. We will then discuss the indirect implementation of the revenue-maximizing policy, showing that this can be achieved via a sequence of ascending auctions with appropriately chosen reserve prices.

4.1. Preliminaries: An Optimal Direct Mechanism

We consider a single monopolist seller who commits, at time zero, to a dynamic direct mechanism $\mathcal{M} = \{\mathbf{x}_t, \mathbf{p}_t\}_{t \in \mathbb{N}_0}$.¹⁹ The seller's expected revenue from this mechanism is the expected discounted sum of payments made by each buyer. Recalling from Equation (2) that the expected payment of a buyer i , conditional on entry, is denoted by $m_i(v_i, \omega_t)$, the seller's payoff may be written as

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^t \alpha_t(i) m_i(v_i, \omega_t) \right],$$

where $\alpha_t(i) = 1$ if $i \in \mathcal{I}_t$ arrives to the market (which occurs with probability $\pi_i \in [0, 1]$) and $\alpha_t(i) = 0$ otherwise. Applying the revenue equivalence result from Corollary 1 and following Myerson (1981), standard techniques imply that the revenue-maximizing seller is faced with the problem of choosing a feasible mechanism to maximize

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^t \alpha_t(i) m_i(0, \omega_t) \right] + \mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^t \alpha_t(i) q_i(v_i, \omega_t) \varphi_i(v_i) \right], \quad (4)$$

subject to the incentive compatibility and individual rationality constraints discussed in Lemma 1.²⁰

Notice that the seller's objective function above is the sum of two terms: the first term is a discounted sum of expected payments, while the second is a discounted sum of weighted virtual values. Moreover, this second term is identical to the efficiency-oriented social planner's objective function in Equation (4), except that values have been replaced with virtual values. Therefore, the insights of Myerson (1981) carry over from the static world to this context: maximizing revenue is

¹⁸An exception is the literature on sponsored search keyword auctions. See, for instance, Edelman and Schwarz (2006).

¹⁹As is standard, we assume that the seller fully commits to the mechanism. This ensures that the revelation principle applies and that there is no loss of generality in considering only direct mechanisms.

²⁰Recall that $\varphi_i(v_i)$ is the virtual valuation of buyer i with value v_i , which we have assumed to be a strictly increasing function of v_i . This assumption is without loss of generality, as we may use the procedure of Skreta (2007) to define and use "ironed" virtual values if φ_i is decreasing or not well-defined.

equivalent to maximizing the virtual surplus. Objects are ordered by their arrival time and buyers are ordered by their *virtual* values, and “earlier” objects are allocated to agents with higher virtual values. Again, this matching must respect the feasibility constraints placed on the allocation rule. Thus, the revenue-maximizing allocation policy is, in each period, to allocate all available objects to the buyers currently present on the market that have the highest virtual values.²¹

Before proceeding to the formal statement of this result, some additional definitions are necessary. For any state $z_t = (h_t, \omega_t, \mathbf{v}) \in \mathcal{H}_t \times \{0, 1\}^I \times \mathcal{K} \times \mathcal{R}$, where \mathbf{v} denotes truthful reporting by all agents, we denote by

$$\mathcal{A}^\pi(z_t) := \{i \in \mathcal{A}_t : \varphi_i(v_i) \geq 0\}$$

the set of all agents present with non-negative virtual values. The set of agents $i \in \mathcal{A}^\pi(z_t)$ whose virtual value is among the k_t highest currently present is given by

$$\mathcal{A}_+^\pi(z_t) := \{i \in \mathcal{A}^\pi(z_t) : |\{j \in \mathcal{A}_t : \varphi_j(v_j) \geq \varphi_i(v_i)\}| \leq k_t\}.$$

Similarly, the set of agents $i \in \mathcal{A}^\pi(z_t)$ whose virtual value is ranked strictly *outside* the top k_t agents is

$$\mathcal{A}_-^\pi(z_t) := \{i \in \mathcal{A}^\pi(z_t) : |\{j \in \mathcal{A}_t : \varphi_j(v_j) > \varphi_i(v_i)\}| \geq k_t\}.$$

Finally,

$$\mathcal{A}_\sim^\pi(z_t) := \mathcal{A}^\pi(z_t) \setminus \left(\mathcal{A}_+^\pi(z_t) \cup \mathcal{A}_-^\pi(z_t) \right)$$

is the set of agents tied for the k_t -th highest ranking virtual value. With these definitions in hand, we may describe the set of optimal allocation rules.

LEMMA 5 (Revenue-maximizing allocation rules).

Suppose all buyers, upon arrival, report their true values. A feasible allocation rule $\{\mathbf{x}_t\}_{t \in \mathbb{N}_0}$ is optimal (revenue-maximizing) if, and only if, for all $z_t = (h_t, \omega_t, \mathbf{v}) \in \mathcal{H}_t \times \{0, 1\}^I \times \mathcal{K} \times \mathcal{R}$,

$$\begin{aligned} x_{i,t}(z_t) &= 1 \text{ for all } i \in \mathcal{A}_+^\pi(z_t), \text{ and} \\ \sum_{i \in \mathcal{A}_\sim^\pi(z_t)} x_{i,t}(z_t) &= k_t - |\mathcal{A}_+^\pi(z_t)| \text{ if } |\mathcal{A}_+^\pi(z_t)| < k_t. \end{aligned}$$

PROOF. The result follows directly from [Lemma 2](#) and the discussion above. \square

As with the efficient allocation rules described by [Lemma 2](#), all revenue-maximizing allocation rules agree after almost all histories. The only possible disagreements occur after the zero probability histories in which multiple agents have identical (positive) virtual values.²² Additionally, these allocation policies are independent of past history, as optimal allocations are functions only of the values of the agents currently present on the market and on the number of objects currently available. We will therefore refer to *the* revenue-maximizing allocation rule $\tilde{\mathbf{x}}$. By this, we mean the revenue-maximizing allocation rule which breaks ties with equal probability, which is given

²¹Note, however, that allocating objects to agents with negative virtual values decreases the seller’s payoff. Thus, the matching described above must restrict attention to buyers with non-negative virtual values.

²²As with the efficient policies discussed in [Section 3](#), we (without loss of generality) disregard the zero-probability events in which an agent’s virtual value is equal to zero.

by

$$\tilde{x}_{i,t}(\omega_t, \mathbf{v}_t) = \begin{cases} 1 & \text{if } i \in \mathcal{A}_+^\pi \\ 0 & \text{if } i \in \mathcal{A}_-^\pi \\ \frac{k_t - |\mathcal{A}_+^\pi|}{|\mathcal{A}_-^\pi|} & \text{if } i \in \mathcal{A}_\sim^\pi \end{cases}$$

for all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$ and $\mathbf{v}_t \in \mathbf{V}^{\mathcal{A}_t}$, where we drop (for convenience) the dependence on h_t .

It should be clear that the revenue-maximizing allocation rule \tilde{x} satisfies the requirements of incentive compatibility. From the perspective of a given buyer $i \in \mathcal{I}$, \tilde{x} allocates an object to i after a given history if, and only if, i is among the highest-ranking (by virtual value) buyers present at that history. Since we have assumed the standard regularity condition of increasing virtual valuations, $\tilde{x}_{i,t}$ is nondecreasing in v_i , given the values of the other agents present on the market. Since this property holds for any arbitrary history and realization of competitors' values, it is straightforward to show that the resulting expected discounted probability of receiving an object is also nondecreasing in v_i . Thus, by choosing an appropriate payment rule, it is possible to design an incentive compatible mechanism that implements the revenue-maximizing allocation policy.

Let us now examine the first term in the seller's objective function in [Equation \(4\)](#). Given the incentive compatibility of the revenue-maximizing allocation policy, the revenue equivalence result of [Corollary 1](#) implies that the individual rationality constraint faced by our seller is

$$m_i(0, \omega_t) \leq 0$$

for all $i \in \mathcal{I}$ and all $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$. Since m_i enters the seller's objective function additively, this constraint must be binding.

Notice that the problem of choosing a payment rule that satisfies this constraint in this dynamic setting is similar to the static optimal auction problem. In the static setting, the [Myerson \(1981\)](#) optimal mechanism can be reinterpreted as a Vickrey-Clarke-Groves mechanism: instead of maximizing surplus, the revenue-maximizing single-object allocation mechanism maximizes *virtual* surplus. When agents report their values v_i , the mechanism computes their virtual values $\varphi_i(v_i)$ and then applies the VCG mechanism to these virtual values. This yields an allocation and a "virtual price" such that the winning buyer's virtual value less the virtual price is equal to her marginal contribution to the virtual surplus. These insights can be applied in our setting; however, care must be taken to correctly account for the discounting-induced heterogeneity across goods available in different periods.

So, for each $i \in \mathcal{I}$, define

$$\tilde{r}_i := \varphi_i^{-1}(0).$$

This is the minimal value required for agent i to potentially receive an object under the revenue-maximizing allocation policy. Furthermore, for each $t \in \mathbb{N}_0$ and all $i \in \mathcal{I}_t$, we define for any $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$ and truthful $\mathbf{v}_t \in \mathbf{V}^{\mathcal{A}_t}$ the function

$$\Pi^i(\omega_t, \mathbf{v}_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, \mathbf{v}_s) \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \right]. \quad (5)$$

This expression is the same as the virtual surplus in the seller's objective function in Equation (4), except that instead of a weighted sum of virtual values, it is a weighted sum of the corresponding "real" values of agent i , less the reservation value \tilde{r}_i applied to agent i ; that is, Π^i measures the virtual surplus in the same units as i 's utility function and i 's distribution function.²³ Since we have assumed that virtual values are increasing for all agents, φ_i^{-1} is increasing. Therefore, transforming the virtual values of all agents by φ_i^{-1} preserves their ordering; moreover, $\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \geq 0$ if, and only if, $\varphi_j(v_j) \geq 0$. Therefore,

$$\tilde{\mathbf{x}} \in \arg \max_{\{\mathbf{x}_s\}_{s=t}^{\infty}} \left\{ \mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{j \in \mathcal{I}_t} \delta^t x_{j,t}(\omega_t, \mathbf{v}_t) \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \right] \right\};$$

that is, the optimal policy $\tilde{\mathbf{x}}$ is an "efficient" allocation rule for an environment in which a planner wishes to maximize the virtual surplus (as evaluated from the perspective of agent i).

Denoting by ω_s^{-i} the market state in period $s \in \mathbb{N}_0$ when buyer i is removed from the market (that is, where we impose $\alpha_s(i) = 0$), we write

$$\Pi_{-i}^i(\omega_t^{-i}, \mathbf{v}_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I} \setminus \{i\}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s^{-i}, \mathbf{v}_s) \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \right]$$

for the virtual surplus (in terms of i 's utility and distribution) when i is removed from the market. Thus, the presence of agent $i \in \mathcal{I}_t$ on the market yields an expected *flow* marginal contribution—again, in units of i 's utility and distribution functions—equal to

$$\tilde{w}_i(\omega_t, \mathbf{v}_t) := \underbrace{\Pi^i(\omega_t, \mathbf{v}_t) - \Pi_{-i}^i(\omega_t^{-i}, \mathbf{v}_t)}_{\text{total contribution}} - \delta \underbrace{\left(\mathbb{E} \left[\Pi^i(\omega_{t+1}, \mathbf{v}_{t+1}) \right] - \mathbb{E} \left[\Pi_{-i}^i(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) \right] \right)}_{\text{expected future contribution}}.$$

We now define the *dynamic virtual pivot mechanism* $\tilde{\mathcal{M}} := \{\tilde{\mathbf{x}}_t, \tilde{\mathbf{p}}_t\}_{t \in \mathbb{N}_0}$ to be the dynamic direct mechanism with the optimal allocation rule $\tilde{\mathbf{x}}$ and the payment rule $\tilde{\mathbf{p}}$ defined by

$$\tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) := \tilde{x}_{i,t}(\omega_t, \mathbf{v}_t)v_i - \tilde{w}_i(\omega_t, \mathbf{v}_t)$$

for all (ω_t, \mathbf{v}_t) . This mechanism will provide to each agent flow payoffs equal to her flow contribution to the virtual surplus. It is then straightforward to show that it implements the revenue-maximizing policy.

LEMMA 6 (Implementability and optimality of the dynamic virtual pivot mechanism).

Suppose that virtual values φ_i are increasing for all $i \in \mathcal{I}$. Then the dynamic virtual pivot mechanism $\tilde{\mathcal{M}}$ is periodic ex post incentive compatible and individually rational, thereby implementing the optimal policy.

PROOF. The proof may be found in Appendix A. □

4.2. An Optimal Sequential Auction

The dynamic virtual pivot mechanism is, of course, a direct mechanism. Again, the question of indirect implementation arises: is it possible to achieve the outcomes of the dynamic virtual pivot mechanism using a *decentralized* auction mechanism? In light of the mechanism's relationship to

²³Note that when all agents are ex ante symmetric, $\varphi_i^{-1}(\varphi_j(v_j)) = v_j$.

the (efficient) dynamic pivot mechanism, as well as the results of Section 3, a natural candidate for a revenue-maximizing indirect mechanism is the sequential ascending auction.

It is well-known that in a static setting with k units of a homogenous good to be allocated, efficiency is achievable by a Vickrey-Clarke-Groves mechanism. This mechanism is outcome equivalent to a k -th price sealed-bid or ascending auction. As established by Myerson (1981) in the case of a single object, and by Maskin and Riley (1989) with multiple units of a homogenous good, the revenue-maximizing mechanism when values are independently and identically drawn from the same distribution F is a pivot mechanism with a reserve price equal to $\tilde{r} := \varphi^{-1}(0)$. Such a mechanism is outcome equivalent to a k -th-price sealed-bid or ascending auction with a reserve price equal to \tilde{r} . In our dynamic setting with randomly arriving and departing buyers, the dynamic pivot mechanism are efficient. Moreover, the outcome of the dynamic pivot mechanism may be implemented via a sequence of ascending auctions. Reasoning by analogy, we may conclude that, since the dynamic virtual pivot mechanism is revenue-maximizing and corresponds (when buyers are ex ante symmetric) to the dynamic pivot mechanism with a reserve of \tilde{r} , a sequence of ascending auctions with reserve price \tilde{r} is the corresponding revenue-maximizing auction.²⁴

Let us formalize this analogy. We again make use of the multi-unit, uniform-price variant of the Milgrom and Weber (1982) button auction. However, we introduce a reserve price equal to \tilde{r} . For simplicity, we assume that the price clock starts at zero and rises continuously.²⁵ When there are $m \geq 1$ units available in a given period, the auction will end whenever there are at most m bidders still active *and* the price is at least \tilde{r} . At that time, each remaining bidder receives an object and pays the price at which the auction ended. As before, ties are broken fairly.

Recall that we denote by n_t the number of buyers present in period t , and that $\mathbf{y}_t := (y_t^1, \dots, y_t^{n_t-1})$ denotes the ordered valuations of all other buyers present in period t (taking the perspective of an arbitrary bidder i), where y_t^1 is the largest value and $y_t^{n_t-1}$ is the smallest. As before, we let

$$\bar{\mathbf{v}}^m := (\bar{v}, \dots, \bar{v}) \in \mathbf{V}^m \text{ and } \mathbf{y}_t^{>m} := (y_t^{m+1}, \dots, y_t^{n_t-1}) \text{ for each } m = 1, \dots, n_t - 1.$$

Finally, we define, for each $m = 1, \dots, n_t - 1$,

$$\tilde{w}^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) := \delta \mathbb{E} [\tilde{w}_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, \mathbf{y}_t^{>m})].$$

This is the (discounted) expected future marginal contribution of an agent $i \in \mathcal{A}_t$ to the virtual surplus, conditional on the period- t presence of m competitors with the highest possible value \bar{v} and $n_t - m - 1$ buyers ranked below i with values $\mathbf{y}_t^{>m}$. Notice that this is exactly i 's expected contribution to the social welfare over a replacement agent with value \tilde{r} . Moreover, observe that $\tilde{w}^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) = 0$ for all buyers with values $v_i \leq \tilde{r}$, regardless of the realization of $\mathbf{y}_t^{>m}$.

²⁴In the case where buyers are *not* ex ante symmetric, a sequence of ascending auctions will again be equivalent to the dynamic virtual pivot mechanism, with the proviso that buyers' price clocks run asynchronously at speeds corresponding to the rate of change in their virtual value functions. Such an auction corresponds to the Myerson (1981) optimal auction when buyers are not symmetric—see Proposition 1 of Caillaud and Robert (2005).

²⁵One could equivalently model each auction as a two-stage game in which buyers first make a participation decision for the current-period auction, and then the price clock starts at \tilde{r} .

In each period $t \in \mathbb{N}_0$, we assume that each agent $i \in \mathcal{A}_t$ bids up to the cutoffs $\tilde{\beta}_{m,n_t}^t$ whenever she has m active competitors in the auction, where

$$\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) := v_i - \tilde{w}^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}). \quad (6)$$

These (symmetric across agents) cutoffs are strictly increasing in v_i , implying that buyers can infer the values of those competitors that have already exited the auction. Note that when buyer i is active and knows the values $\mathbf{y}_t^{>m}$ of her inactive opponents, the price at which she is indifferent between winning an object and receiving her discounted marginal contribution in the next period (conditional on all remaining active buyers having values greater than hers) is exactly $\tilde{\beta}_{m,n_t}^t$. We may then use arguments similar to those of Theorems 1 and 2 to prove the following result.

THEOREM 3 (Revenue maximization via sequential ascending auctions).

Suppose that $F_i = F$ for all $i \in \mathcal{I}$. Following the bidding strategies $\tilde{\beta}_{m,n_t}^t$ in Equation (6) in every period of the sequential ascending auction with reserve price $\tilde{r} := \varphi^{-1}(0)$ is a (periodic ex post) perfect Bayesian equilibrium that is outcome equivalent to the dynamic virtual pivot mechanism.

PROOF. The proof mirrors Section 3.2. Full details may be found in Appendix B. \square

Theorem 3 therefore implies that the sequential ascending auction with a reserve price admits an equilibrium with prices and allocations identical to those of the dynamic virtual pivot mechanism. Therefore, analogous to the case of Section 3.2, we find that a seller who wishes to maximize revenues via a transparent, decentralized mechanism may do so by using a sequence of ascending auctions. Moreover, the method of proof of the proposition above shows that the strategy profile specified in Equation (6) forms a periodic ex post equilibrium. Given expectations about future arrivals and behavior, each buyer's current-period bid is a best response to the strategies of her opponents, regardless of the realization of their values or the history of the mechanism.

5. CONCLUSION

In this paper, we examine a private-values, single-unit-demand environment where buyers and objects arrive at random times. We discuss the implementation of the efficient allocation policy via the dynamic pivot mechanism, a dynamic variant of the classic Vickrey-Clarke-Groves mechanism. Moreover, by extending the static Myerson (1981) payoff- and revenue-equivalence results to our dynamic setting, we are able to derive an optimal direct mechanism. This mechanism succeeds in maximizing the seller's profits in this dynamic environment by applying the efficient mechanism to buyers' virtual values.

While direct mechanisms are useful theoretical devices, there is much evidence demonstrating that they may be of limited value in practice. We therefore consider indirect mechanisms in this setting and propose using a sequence of one-shot auctions instead. We show that a sequence of ascending auctions serves as a simple, natural, and intuitive institution that corresponds to the dynamic Vickrey-Clarke-Groves mechanism. Unlike the standard second-price auction, this open auction format allows each buyer to learn her competitors' values, and hence determine her own marginal contribution to the social welfare. When each buyer exits each auction at the price such that she is indifferent between winning the object and obtaining her future marginal contribution,

we obtain a decentralized price discovery mechanism that yields equilibrium outcomes identical to those of the centralized direct mechanism. Moreover, this equilibrium behavior is memoryless—buyers rationally ignore payoff-relevant information from previous periods, correctly anticipating that it will be revealed anew.

These results set the stage for several additional avenues of inquiry. For instance, suppose that objects need not be allocated in the period of their arrival, but can instead be placed in inventory and allocated in future periods. Such a situation provides a seller with an additional strategic tool: the ability to withhold an item in the current period in hopes of “better” demand in future periods. While some properties of our solution are maintained (for instance, the efficient policy continues to allocate objects to higher-ranked buyers before moving onto competitors with lower values, and a dynamic VCG mechanism will continue to be efficient), the indirect implementation results do not follow immediately. For example, in the case of storable objects, the sale of an object to a particular agent imposes an additional externality on her competitors, as the number of objects available in the future decreases. This reduces the incentives for buyers to truthfully reveal their private information, as this information may be of great strategic value to competitors. Moreover, reserve prices are necessary even for achieving efficient outcomes, and the manner in which these reserve prices fluctuate over time with changes in supply will be a crucial factor in the possibility of attaining efficient outcomes via an auction mechanism.

Another natural extension of our model is the generalization to the case in which agents may demand multiple units. This, however, introduces additional intertemporal tradeoffs in any auction mechanism, as expected future payoffs in individual valuations are no longer identical functions of individual values when buyers have differential demands or are faced with multi-unit “exposure” risks. While informational asymmetries may be resolved via information renewal and memoryless strategies, such strategies cannot resolve the fundamental asymmetry in objectives that arise when some buyers have already satisfied a portion of their demand. An alternative line of research relaxes the assumption that buyer entries and exits are exogenous, instead allowing buyers to condition their participation on market conditions. Such a model would provide an important building block to an understanding of competing marketplaces and platforms. We leave these questions, however, for future work.

APPENDIX A. OMITTED PROOFS

PROOF OF LEMMA 1. We first show the necessity of the three conditions for incentive compatibility and individual rationality. So, suppose that the mechanism \mathcal{M} is both incentive compatible and individually rational. Fix any $t \in \mathbb{N}_0$, any $i \in \mathcal{I}_t$, and arbitrary $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$. Then incentive compatibility implies that, for all $v_i \in \mathbf{V}$,

$$\widehat{U}_i(v_i, \omega_t) = \max_{v'_i \in \mathbf{V}} \{q_i(v'_i, \omega_t)v_i - m_i(v'_i, \omega_t)\}.$$

Thus, $\widehat{U}_i(v_i, \cdot)$ is an affine maximizer, and is hence a convex function of v_i . Moreover, for all $v_i, v'_i \in \mathbf{V}$, incentive compatibility is equivalent to

$$\begin{aligned} \widehat{U}_i(v'_i, \omega_t) &\geq q_i(v_i, \omega_t)v'_i - m_i(v_i, \omega_t) \\ &= q_i(v_i, \omega_t)v_i - m_i(v_i, \omega_t) + q_i(v_i, \omega_t)(v'_i - v_i) \\ &= \widehat{U}_i(v_i, \omega_t) + q_i(v_i, \omega_t)(v'_i - v_i). \end{aligned}$$

Thus, $q_i(v_i, \cdot)$ is a subderivative of $\widehat{U}_i(v_i, \cdot)$ at v_i . Since $\widehat{U}_i(v_i, \cdot)$ is convex in v_i , it is absolutely continuous and hence differentiable almost everywhere, implying that at every point of differentiability,

$$\frac{\partial}{\partial v_i} \widehat{U}_i(v_i, \omega_t) = q_i(v_i, \omega_t).$$

Since $\widehat{U}_i(v_i, \cdot)$ is convex, this implies that q_i must be nondecreasing in v_i .

Moreover, every absolutely continuous function is equal to the definite integral of its derivative, implying that

$$\widehat{U}_i(v_i, \omega_t) = \widehat{U}_i(0, \omega_t) + \int_0^{v_i} q_i(v'_i, \omega_t) dv'_i$$

for all $v_i \in \mathbf{V}$. Finally, since q_i is nondecreasing in v_i , the requirement of individual rationality is then satisfied for all v_i only if

$$\widehat{U}_i(0, \omega_t) \geq 0.$$

Hence, the three conditions are necessary conditions for \mathcal{M} to be incentive compatible and individually rational.

We now show the sufficiency of the three conditions for incentive compatibility and individual rationality. Suppose that \mathcal{M} satisfies the three conditions, and fix any $t \in \mathbb{N}_0$, any $i \in \mathcal{I}_t$, and arbitrary $\omega_t \in \{0, 1\}^{\mathcal{I}} \times \mathcal{K}$. Note first that q_i is nondecreasing in v_i and $\widehat{U}_i(0, \omega_t) \geq 0$ immediately imply that $\widehat{U}_i(v_i, \omega_t) \geq 0$ for all $v_i \in \mathbf{V}$, and so \mathcal{M} is individually rational.

Now, for any $v_i, v'_i \in \mathbf{V}$, the second condition implies that

$$\widehat{U}_i(v'_i, \omega_t) = \widehat{U}_i(v_i, \omega_t) + \int_{v_i}^{v'_i} q_i(v''_i, \omega_t) dv''_i.$$

If $v_i < v'_i$, then q_i nondecreasing implies that $q_i(v_i, \omega_t) \leq q_i(v'_i, \omega_t)$. Thus,

$$\widehat{U}_i(v'_i, \omega_t) \geq \widehat{U}_i(v_i, \omega_t) + q_i(v_i, \omega_t)(v'_i - v_i).$$

Similarly, if $v_i > v'_i$, then q_i nondecreasing implies that $q_i(v_i, \omega_t) \geq q_i(v'_i, \omega_t)$. Therefore,

$$\begin{aligned}\widehat{U}_i(v'_i, \omega_t) &= \widehat{U}_i(v_i, \omega_t) - \int_{v'_i}^{v_i} q_i(v''_i, \omega_t) dv''_i \\ &\geq \widehat{U}_i(v_i, \omega_t) - q_i(v_i, \omega_t)(v_i - v'_i) = \widehat{U}_i(v_i, \omega_t) + q_i(v_i, \omega_t)(v'_i - v_i).\end{aligned}$$

However, this inequality is, as shown above, equivalent to incentive compatibility. Since $v_i, v'_i \in \mathbf{V}$ were chosen arbitrarily, this implies that \mathcal{M} is incentive compatible. \square

PROOF OF LEMMA 2.²⁶ Note first that any two allocation rule that satisfy the conditions of the lemma yield the same expected payoff to the social planner. To see this, note that the only variation permitted is in the allocation of objects to agents with zero valuation and in the breaking of ties. Given the allocations to all other agents, choosing to allocate to agents with value zero yields neither an increase nor decrease in the realized surplus. Moreover, although the second condition regarding the allocation to agents in \mathcal{A}_\sim allows for various mixtures over this set of agents, the outcome of these mixtures is always the same: exactly $k_t - |\mathcal{A}_+|$ of these agents receive an object. Different choices among these outcomes does not affect future payoffs, as the arrival process of agents and objects is orthogonal to these allocative decisions, and all agents depart the system at the same exogenously given rate $1 - \gamma$.

With this in mind, let $\widehat{\mathbf{x}}$ denote a deterministic allocation rule that allocates an object to the highest-ranking agents (including those with value equal to zero), where ties are broken arbitrarily (but without randomization). Fix any policy \mathbf{x}_0 that yields the planner a strictly higher payoff than the policy $\widehat{\mathbf{x}}$, and define

$$\mathcal{Z}^0 := \left\{ (h_t, \omega_t) \in \mathcal{H} \times \{0, 1\}^{\mathcal{I}} \times \mathcal{K} : \mathbf{x}_0(z) \neq \widehat{\mathbf{x}}(z), \mathbf{x}_0(z') = \widehat{\mathbf{x}}(z') \text{ for all } z' \rightarrow z \right\}.$$

Thus, \mathcal{Z}^0 is the set of all histories and arrivals z such that \mathbf{x}_0 and $\widehat{\mathbf{x}}$ disagree at z , but agree on all of z 's prefixes; that is, \mathcal{Z}^0 is the set of “first” or “earliest” disagreements between \mathbf{x}_0 and $\widehat{\mathbf{x}}$. Since \mathbf{x}_0 does strictly better than $\widehat{\mathbf{x}}$, this set must have nonzero measure (with respect to the measure induced by the arrival processes), as otherwise the two policies would agree almost everywhere (and hence yield identical payoffs).

For each $z \in \mathcal{Z}^0$, note that the policy \mathbf{x}_0 induces a probability distribution over outcomes, where an outcome is an assignment of objects to agents. Denote by $\Sigma^0(z)$ the set of outcomes induced by \mathbf{x}_0 at history z . Thus, an outcome $\sigma \in \Sigma^0(z)$ is associated with a subset of buyers present that receive an object. Let $a_j(\sigma)$ denote the j -th highest-valued agent that receives an object under \mathbf{x}_0 in outcome σ . Similarly, let $b_j(z)$ denote the j -th highest-valued agent overall.

Define for each $z \in \mathcal{Z}^0$ and for each $\sigma \in \Sigma^0(z)$, we define the “continuation policy” $\mathbf{x}_1^\sigma(z)$ to be the allocation rule that allocates to the highest-ranking agents present at time z , and is equal to \mathbf{x}_0 at all successors of z except that it allocates to agent $a_j(\sigma)$ whenever \mathbf{x}_0 allocates to agent $b_j(z)$. Thus, $\mathbf{x}_1^\sigma(z)$ is the same as \mathbf{x}_0 except that it “swaps” the allocation decisions of $a_j(\sigma)$ and $b_j(z)$. Since $v_{a_j(\sigma)} \leq v_{b_j(z)}$ (with a strict inequality for at least one j), the expected payoff to the planner under \mathbf{x}_1^σ is greater than that of \mathbf{x}_0 along this branch of the mechanism tree. To see why this is true,

²⁶Thanks are due to Larry Samuelson for suggesting the method of proof used below.

consider any $v > v'$ and $t < t'$. Since $\delta < 1$, we have

$$(\delta^t v + \delta^{t'} v') - (\delta^t v' + \delta^{t'} v) = (\delta^t - \delta^{t'})(v - v') > 0.$$

Thus, even if agents do not depart the market, the planner's payoff along this path is increases.

Thus, define the allocation policy \mathbf{x}_1 to be equal to \mathbf{x}_0 at all histories that are not successors to histories in \mathcal{Z}^0 . Furthermore, for each $z \in \mathcal{Z}^0$, we define $\mathbf{x}_1(z)$ to be the stochastic policy that chooses $\mathbf{x}_1^\sigma(z)$ with the probability that \mathbf{x}_0 leads to outcome σ . Since this leads to an increase in the planner's payoff over \mathbf{x}_0 along every successor history to those in \mathcal{Z}^0 , and this set has positive measure, it must be the case that \mathbf{x}_1 yields a strictly greater payoff than \mathbf{x}_0 .

If \mathbf{x}_1 yields the planner a payoff less than or equal to that of $\hat{\mathbf{x}}$, transitivity of the planner's payoffs leads to a contradiction, implying that there does not exist a policy \mathbf{x}_0 such that \mathbf{x}_0 does strictly better than $\hat{\mathbf{x}}$, and hence that $\hat{\mathbf{x}}$ is optimal.

On the other hand, if \mathbf{x}_1 yields a payoff greater than that of $\hat{\mathbf{x}}$, we define the set

$$\mathcal{Z}^1 := \left\{ (h_t, \omega_t) \in \mathcal{H} \times \{0, 1\}^T \times \mathcal{K} : \mathbf{x}_1(z) \neq \hat{\mathbf{x}}(z), \mathbf{x}_1(z') = \hat{\mathbf{x}}(z') \text{ for all } z' \rightarrow z \right\}$$

to be the set of \mathbf{x}_1 's "first disagreements" with $\hat{\mathbf{x}}$. We may repeat the procedure above to then define a new policy \mathbf{x}_2 that agrees with $\hat{\mathbf{x}}$ at every $z \in \mathcal{Z}^1$, but does strictly better than either \mathbf{x}_1 . Notice that, if \mathbf{x}_1 does better than $\hat{\mathbf{x}}$, then we have arrived at a contradiction.

Proceeding in this manner, we construct a sequence of policies $\{\mathbf{x}_s\}_{s=0}^\infty$ with associated expected payoffs $\{W_s\}_{s=0}^\infty$ such that $W_s < W_{s+1}$ for all $s \in \mathbb{N}_0$. Note, however, that for all $s \in \mathbb{N}_0$, \mathbf{x}_s agrees with $\hat{\mathbf{x}}$ on *at least* all histories of length s . Since δ^s approaches zero as s becomes increasingly large, this implies that

$$\lim_{s \rightarrow \infty} W_s = \hat{W},$$

where \hat{W} is the planner's expected payoff from following policy $\hat{\mathbf{x}}$. Moreover, since $\{W_s\}$ is an increasing sequence, this implies that

$$\hat{W} \geq W_s \text{ for all } s \in \mathbb{N}_0,$$

a contradiction. It must therefore be the case that there does not exist a policy \mathbf{x}_0 that yields the planner a strictly higher payoff than $\hat{\mathbf{x}}$. Therefore, we may conclude that $\hat{\mathbf{x}}$ is, in fact, a socially optimal policy. \square

PROOF OF LEMMA 4. Fix an arbitrary period $t \in \mathbb{N}_0$, and let $\omega_t := (\alpha_t, k_t)$ denote the state of the market at time t . Consider an agent $i \in \mathcal{A}_t$ with value v_i , and suppose that $n_t - m - 1$ buyers have dropped out of the period- t auction, revealing values $\mathbf{y}_t^{>m}$, where $m \in \{1, \dots, n_t - 1\}$. We wish to show first that $v_i > v_j > y_t^{m+1}$ implies that

$$\hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) := v_i - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) > v_j - w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m}) =: \hat{\beta}_{m, n_t}^t(\omega_t, v_j, \mathbf{y}_t^{>m}).$$

$$\begin{aligned}
 \text{Notice that } w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m}) - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) &= \\
 \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] - \delta \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] \\
 - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] + \delta \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] \\
 = \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right],
 \end{aligned}$$

since removing either i or j in the following period, conditional on their being the m -th highest-ranked agent, does not differentially affect the order of anticipated future allocations to any other agents. In particular, since the efficient allocation rule $\hat{\mathbf{x}}$ makes assignments based solely on the ranking of valuations, it will choose the same assignments in future periods when i or j have been removed from the market.

Moreover, by naively treating buyer j as though her true value were v_i , we can provide a bound on the difference above. In particular, we have

$$\begin{aligned}
 \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] - \delta \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] \\
 \geq \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{\mathbf{x}}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] (v_j - v_i).
 \end{aligned}$$

Thus, if $v_i > v_j$, then

$$\begin{aligned}
 \hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) - \hat{\beta}_{m, n_t}^t(\omega_t, v_j, \mathbf{y}_t^{>m}) \\
 \geq (v_i - v_j) \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{\mathbf{x}}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, \mathbf{y}_t^{>m}) \right] \right) > 0
 \end{aligned}$$

since the discounted expected probability of receiving an object in the future is bounded above by $\delta < 1$. Thus, $\hat{\beta}_{m, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m})$ is strictly increasing in v_i .

$$\begin{aligned}
 \text{Also, note that if } v_i > v_j = y_t^{m+1}, \text{ then } w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m+1}) - w^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1}) = \\
 \delta \left(\mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] \right) \\
 - \delta \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right] \right).
 \end{aligned}$$

However, the second difference above may be rewritten as

$$\begin{aligned}
 \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{>m+1}) \right] \\
 + \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[W_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right].
 \end{aligned}$$

Thus,

$$w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m+1}) - w^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1})$$

is the sum of three differences. The first is the expected gain in social welfare when increasing i 's value from v_i to \bar{v} . The second is the expected gain in social welfare (when i is not on the market) from increasing j 's value from v_j to v_i . Finally, the third difference is the expected loss in social welfare (when j is not present) from decreasing i 's value from \bar{v} to v_i . However, since $v_j < v_i$, the

presence or absence of j from the market has no influence on when the efficient policy allocates to i , regardless of whether i 's value is v_i or \bar{v} . Therefore, the gain from the first difference equals the loss from the third difference, implying that

$$\begin{aligned} & w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{> m+1}) - w^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{> m+1}) \\ &= \delta \left(\mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{> m+1}) \right] - \mathbb{E} \left[W_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{> m+1}) \right] \right). \end{aligned}$$

Moreover, by (again) naively treating buyer j as though her true value were v_i , we can provide a bound on the difference above, which may be used to show that

$$\hat{\beta}_{m, n_t}^t(\omega_t, v_i, v_j, \mathbf{y}_t^{> m+1}) - \hat{\beta}_{m+1, n_t}^t(\omega_t, v_j, \mathbf{y}_t^{> m+1}) > 0.$$

Thus, the exit of the buyer with rank $(m+1)$ does not induce the immediate exit of any buyer with a higher value. Therefore, since $\hat{\beta}_{m, n_t}^t(\omega_t, v_i, v_j, \mathbf{y}_t^{> m+1})$ is strictly increasing in v_i , the price at which this exit occurs fully reveals the value of the $(m+1)$ -th highest-ranked buyer.

Since m was arbitrary, we may conclude that bids are fully separating. \square

PROOF OF LEMMA 6. As discussed within the text, the discounted expected probability of receiving an object under the revenue-maximizing allocation policy $\tilde{q}_i(v_i, \omega_t)$ is nondecreasing in v_i . This implies, applying [Lemma 1](#), that the virtual dynamic pivot mechanism is incentive compatible. To see that it this mechanism is also individually rational, observe that the law of iterated expectations implies that

$$\mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \tilde{w}_i(\omega_s, \mathbf{v}_s) | v_i = 0 \right] = \mathbb{E} [\tilde{w}_i(\omega_t, \mathbf{v}_t) | v_i = 0].$$

Simple arithmetic then implies that the expected utility from participating in the mechanism of an agent i with value $v_i = 0$ simplifies to

$$\hat{U}_i(0, \omega_t) = -\tilde{m}_i(0, \omega_t) = \mathbb{E} [\tilde{w}_i(\omega_t, \mathbf{v}_t) | v_i = 0].$$

Notice, however, the revenue-maximizing allocation rule *never* allocates an object to agent i since

$$\varphi_i(0) = -\frac{1}{f_i(0)} < 0.$$

Therefore, the optimal policy yields exactly the same outcome whether or not i is present, implying that $\tilde{w}_i = 0$ regardless of the realizations of other buyers' values. Thus, by [Lemma 1](#), the dynamic virtual pivot mechanism $\tilde{\mathcal{M}}$ is individually rational.

In order to show that truth-telling is a periodic ex post optimal strategy for all agents, fix an arbitrary agent $i \in \mathcal{I}_t$ for arbitrary $t \in \mathbb{N}_0$, and suppose that i knows the reported values \mathbf{v}_{-i}^t of all agents other than i who are also on the market at time t . Then, by reporting a value v_i^t upon her

arrival in state ω_t , agent i 's payoff under the dynamic virtual pivot mechanism is

$$\begin{aligned} & \left(\tilde{x}_{i,t}(\omega_t, (v'_i, \mathbf{v}_{-i}^t)) + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \tilde{x}_{i,s}(\omega_s, (v'_i, \mathbf{v}_{-i}^s)) \right] \right) (v_i - v'_i) + \tilde{w}_i(\omega_t, (v'_i, \mathbf{v}_{-i}^t)) \\ & = \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, (v'_i, \mathbf{v}_{-i}^s)) \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \right] - \Pi_{-i}^i(\omega_t^{-i}, \mathbf{v}_{-i}^t), \end{aligned}$$

where the expectation is taken with respect to the true distributions of values for agents arriving in periods $s > t$. Since $\tilde{\mathbf{x}}$ is an efficient policy for maximizing the above sum of “transformed” virtual values, the first term above is maximized by setting $v'_i = v_i$. Moreover, the second term does not depend on v'_i . Hence, i 's expected payoff is maximized by truthful reporting of her value, regardless of the reports of the other agents present or the state upon i 's arrival; that is, given the truth-telling behavior of agents arriving in every future period, truthful reporting is optimal regardless of the realizations of all other agents already present on the market. \square

APPENDIX B. PROOF OF THEOREM 3

The proof of this proposition parallels the developments of Section 3.2. In particular, we will first show, as in Lemma 4, that bids are fully separating. Then we will show that, analogous to Theorem 1, following the postulated bidding strategies leads to an identical outcome as the dynamic virtual pivot mechanism. Finally, we will show, as in Theorem 2, that these strategies form a (periodic ex post) perfect Bayesian equilibrium of the sequential auction mechanism.

CLAIM. *The bid functions $\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m})$ are increasing in v_i for all $m = 1, \dots, n_t - 1$. Moreover, if $v_i > y_t^{m+1}$, then*

$$\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) > \tilde{\beta}_{m+1,n_t}^t(\omega_t, \mathbf{y}_t^{>m}).$$

PROOF OF CLAIM. Fix an arbitrary period $t \in \mathbb{N}_0$, and let α_t and k_t indicate the set of agents and objects present on the market, respectively. Consider an agent $i \in \mathcal{A}_t$ with value v_i , and suppose that $n_t - m - 1$ buyers have dropped out of the period- t auction, revealing values $\mathbf{y}_t^{>m}$, where $n_t := |\mathcal{A}_t|$ is the number of agents present, and $m \in \{1, \dots, n_t - 1\}$. We wish to show first that $v_i > v_j > y_t^{m+1}$ implies that

$$\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) := v_i - \tilde{w}^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) > v_j - \tilde{w}^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m}) =: \tilde{\beta}_{m,n_t}^t(\omega_t, v_j, \mathbf{y}_t^{>m}).$$

Notice that

$$\begin{aligned} & \tilde{w}^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m}) - \tilde{w}^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>m}) \\ &= \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] - \delta \mathbb{E} \left[\Pi_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] \\ & \quad - \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] + \delta \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right]. \end{aligned}$$

This, however, is equal to

$$\delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] - \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right],$$

since removing either i or j in the following period, conditional on their being the m -th highest-ranked agent, does not differentially affect the order of anticipated future allocations to any other agents. In particular, since the the revenue-maximizing allocation rule $\tilde{\mathbf{x}}$ makes assignments based solely on the ranking of valuations, it will choose the same assignments in future periods when i or j have been removed from the market.

Moreover, by naïvely treating buyer j as though her true value were v_i , we can provide a bound on the difference above. In particular, we have

$$\begin{aligned} \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_j, \mathbf{y}_t^{>m}) \right] &\geq \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] \\ & \quad + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \tilde{\mathbf{x}}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-m}, v_i, \mathbf{y}_t^{>m}) \right] (v_j - v_i). \end{aligned}$$

Thus, if $v_i > v_j$, then

$$\begin{aligned} & \tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) - \tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m}) \\ & \geq (v_i - v_j) \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \tilde{x}_{i,s}(\omega_s, \mathbf{v}_s) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, \mathbf{y}_t^{>m}) \right] \right) > 0 \end{aligned}$$

since the discounted expected probability of receiving an object in the future is bounded above by $\delta < 1$. Thus, $\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>m})$ is strictly increasing in v_i .

Additionally, note that if $v_i > v_j = y_t^{m+1}$, then

$$\begin{aligned} & \tilde{w}^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m+1}) - \tilde{w}^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1}) \\ & = \delta \left(\mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] \right) \\ & \quad - \delta \left(\mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[\Pi_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right] \right). \end{aligned}$$

However, the second difference above may be rewritten as

$$\begin{aligned} & \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{>m+1}) \right] \\ & + \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{>m+1}) \right] - \mathbb{E} \left[\Pi_{-j}(\omega_{t+1}^{-j}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^{m+1}, v_j, \mathbf{y}_t^{>m+1}) \right]. \end{aligned}$$

Thus,

$$\tilde{w}^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m+1}) - \tilde{w}^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1})$$

is the sum of three differences. The first is the expected gain in virtual surplus when increasing i 's value from v_i to \bar{v} . The second is the expected gain in virtual surplus (when i is not on the market) from increasing j 's value from v_j to v_i . Finally, the third difference is the expected loss in virtual surplus (when j is not present) from decreasing i 's value from \bar{v} to v_i . However, since $v_j < v_i$, the presence or absence of j from the market has no influence on when the optimal (revenue-maximizing) policy allocates to i , regardless of whether i 's value is v_i or \bar{v} . Therefore, the gain from the first difference equals the loss from the third difference, implying that

$$\begin{aligned} \tilde{w}^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>m+1}) - \tilde{w}^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1}) & = \delta \left(\mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_j, \mathbf{y}_t^{>m+1}) \right] \right. \\ & \quad \left. - \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}^{-i}, \mathbf{v}_{t+1}) | \mathbf{v}_t = (\bar{\mathbf{v}}^m, v_i, v_i, \mathbf{y}_t^{>m+1}) \right] \right). \end{aligned}$$

Moreover, by (again) naively treating buyer j as though her true value were v_i , we can provide a bound on the difference above, which may be used to show that

$$\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1}) - \tilde{\beta}_{m+1,n_t}^t(\omega_t, v_j, \mathbf{y}_t^{>m+1}) > 0.$$

Thus, the exit of the buyer with rank $(m+1)$ does not induce the immediate exit of any buyer with a higher value. Therefore, since $\tilde{\beta}_{m,n_t}^t(\omega_t, v_i, v_j, \mathbf{y}_t^{>m+1})$ is strictly increasing in v_i , the price at which this exit occurs fully reveals the value of the $(m+1)$ -th highest-ranked buyer.

Since m was arbitrarily chosen, this implies that the drop-out points of buyers bidding according to the strategy described by Equation (6) are fully revealing of the buyers' values. \square

CLAIM. Following the bidding strategies $\tilde{\beta}_{m,n_t}^t$ in every period t in the sequential ascending auction mechanism is outcome equivalent to the dynamic virtual pivot mechanism.

PROOF OF CLAIM. Fix an arbitrary period $t \in \mathbb{N}_0$, and let k_t denote the number of objects present, and $n_t := |\mathcal{A}_t|$ denote the number of agents present. As shown above, the bidding strategies $\tilde{\beta}_{m,n_t}^t$ are strictly increasing; therefore, the multi-unit uniform-price ascending auction ends allocating the k_t objects to the group of buyers with the k_t highest values greater than the reserve.²⁷ Recall that if $k_t \geq n_t$, the auction ends immediately upon the price reaching the reserve value \tilde{r} , and all buyers present receive an object at that price. Similarly, in the dynamic virtual pivot mechanism, each buyer i with $v_i > \tilde{r}$ receives an object, and makes a payment $\tilde{p}_{i,t}$ given by

$$\tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) = v_i - \tilde{w}_i(\omega_t, \mathbf{v}_t),$$

where \tilde{w}_i is the agent's marginal contribution to the virtual surplus.²⁸ However, since there are sufficient objects present for each agent with a non-negative virtual value to receive one, i does not impose any externalities on the remaining agents; thus,

$$\tilde{w}_i(\omega_t, \mathbf{v}_t) = \tilde{w}_i(\omega_t, \mathbf{v}_t) = v_i - \tilde{r},$$

implying that $\tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) = \tilde{r}$. In this case, then, the allocation and payments of the auction mechanism and the dynamic pivot mechanism are the same.

Suppose instead that $k_t < n_t$; that is, there are more agents present than objects. Denote by i_m the bidder with the m -th highest value. Then each agent who receives an object pays the greater of the reserve price \tilde{r} and the price at which buyer i_{k_t+1} drops out of the auction, which is given by

$$\hat{\beta}_{k_t+1,n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}) = v_{i_{k_t+1}} - w^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}).$$

If $v_{i_{k_t+1}} < \tilde{r}$, then the situation is identical to the previous case. Therefore, assume that $v_{i_{k_t+1}} \geq \tilde{r}$.

In the dynamic virtual pivot mechanism, on the other hand, each agent i who receives an object pays a price

$$\begin{aligned} \tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) &= v_i - w_i(\omega_t, \mathbf{v}_t) \\ &= v_i - \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, \mathbf{v}_s) (v_j - \tilde{r}) \right] + \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I} \setminus \{i\}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s^{-i}, \mathbf{v}_s) (v_j - \tilde{r}) \right]. \end{aligned}$$

This may be rewritten as

$$\begin{aligned} v_i - \left(\sum_{m=1}^{k_t} (v_m - \tilde{r}) + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, \mathbf{v}_s) (v_j - \tilde{r}) \right] \right) \\ + \left(\sum_{m=1}^{k_t} (v_m - \tilde{r}) + (v_{i_{k_t+1}} - v_i) + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s^{-i, -i_{k_t+1}}, \mathbf{v}_s) (v_j - \tilde{r}) \right] \right). \end{aligned}$$

²⁷Recall that buyers with values less than \tilde{r} bid up to their true value, as they are never allocated an object, and so their future expected contribution to the virtual surplus is zero.

²⁸Note that since i is receiving an object, her total and flow marginal contributions are equal.

Rearranging the above expression allows us to rewrite it as

$$\begin{aligned} \tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) &= v_{i_{k_t+1}} - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, \mathbf{v}_s) (v_j - \tilde{r}) \right] \\ &\quad + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s^{-i, -i_{k_t+1}}, \mathbf{v}_s) (v_j - \tilde{r}) \right] \\ &= v_{i_{k_t+1}} - \tilde{w}^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}), \end{aligned}$$

where the second equality follows from the fact that $w^{t+1}(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}})$ is defined to be the expected future marginal contribution to the virtual surplus of the agent with the $(k_t + 1)$ -th highest value, conditional on agents with higher values (which includes i) receiving an object today. Therefore, following the bidding strategies $\tilde{\beta}_{m,n_t}^t$ leads to period- t prices and allocations identical to those of the dynamic pivot mechanism. Since the period t was arbitrary, as was the state ω_t , this equivalence holds after each history. Thus, the two mechanisms are outcome equivalent. \square

Finally, it remains to be seen that the bidding strategies in Equation (6) do, in fact, form an equilibrium. As in the case of the sequential ascending auction with no reserve, the bidding strategies $\tilde{\beta}_{m,n_t}^t$ are strictly increasing. Behavior along the equilibrium path is therefore perfectly separating, implying that Bayesian updating fully determines beliefs. In order to determine optimality off the equilibrium path, we again suppose that, after a deviation, buyers ignore their past observations and the history of the mechanism, and instead believe that the deviating agent is *currently* truthfully revealing her value in accordance with the bidding strategies $\tilde{\beta}_{m,n_t}^t$.

CLAIM. Suppose that in each period, buyers bid according to the cutoff strategies given in Equation (6). This strategy profile, combined with the system of beliefs described above, forms a perfect Bayesian equilibrium of the sequential ascending auction mechanism with reserve price \tilde{r} .

PROOF OF CLAIM. We prove this claim by making use of the one-shot deviation principle. Consider any period with $n_t := |\mathcal{A}_t|$ buyers on the market and k_t objects present. Suppose that all bidders other than player i are using the conjectured equilibrium strategies. We must show that bidder i has no profitable one-shot deviations from the collection of cutoff points $\{\tilde{\beta}_{m,n_t}^t\}$. More specifically, we must show that i does not wish to exit the auction earlier than prescribed, nor does she wish to remain active later than specified.

Once again labeling agents such that buyer i_1 has the highest value and buyer i_{n_t} has the lowest, note that if $v_i < \max\{v_{i_{k_t}}, \tilde{r}\}$, bidding according to $\{\tilde{\beta}_{m,n_t}^t\}$ implies that i does not win an object in the current period. Therefore, exiting earlier than specified does not affect i 's current-period returns. Moreover, since the bidding strategies are memoryless, neither future behavior by i 's competitors nor i 's future payoffs will be affected by an early exit.

Suppose, on the other hand, that $v_i > \tilde{r}$ and that i has one of the k_t highest values; that is, that $v_i \geq \max\{v_{i_{k_t}}, \tilde{r}\}$. As established by Theorem 1, i receives an object, paying a price such that her payoff is exactly equal to her marginal contribution to the virtual surplus. Deviating to an early exit, however, leads either to agent i_{k_t+1} winning an object (if $v_{i_{k_t+1}} \geq \tilde{r}$) instead of buyer i , or to an object being discarded. Moreover, i 's expected payoff is then $\tilde{w}^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}})$, which

we defined as i 's future expected marginal contribution to the virtual surplus. This is a profitable one-shot deviation for i if, and only if,

$$\tilde{w}^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}) \geq v_i - \tilde{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\tilde{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t+1}}, \dots, v_{i_{n_t}}) \geq v_i - \tilde{w}^{t+1}(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}) = \tilde{\beta}_{k_t, n_t}^t(\omega_t, v_i, v_{i_{k_t+2}}, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\tilde{\beta}_{k_t, n_t}^t$ in Equation (6). Since $v_i > v_{i_{k_t+1}}$, this contradicts the conclusion of the first claim above. Thus, i does not wish to exit the auction early.

Alternately, if $v_i \geq \max\{v_{i_{k_t}}, \tilde{r}\}$, then planning to remain active in the auction *longer* than specified does not change i 's payoffs, as i will win an object regardless. If, on the other hand, $v_i < \max\{v_{i_{k_t}}, \tilde{r}\}$, then delaying exit from the period- t auction can affect i 's payoffs. Since bids in future periods do not depend on information revealed in the current period, this only occurs if i remains in the auction long enough to win an object. If i wins, she pays a price equal to the larger of \tilde{r} and the exit point of $i_{v_{k_t}}$, whereas if she exits, she receives as her continuation payoff her marginal contribution to the virtual surplus. So, suppose that $i = i_m$ for some $m > k_t$. Then a deviation to remaining active in the auction is profitable if, and only if,

$$v_m - \tilde{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t}}, \dots, v_{i_{m-1}}, v_{m+1}, \dots, v_{i_{n_t}}) \geq \tilde{w}^{t+1}(\omega_t, v_m, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\tilde{\beta}_{k_t, n_t}^t(\omega_t, v_{i_{k_t}}, \dots, v_{i_{m-1}}, v_{m+1}, \dots, v_{i_{n_t}}) \leq v_m - \tilde{w}^{t+1}(\omega_t, v_m, \dots, v_{i_{n_t}}) = \tilde{\beta}_{m-1, n_t}^t(\omega_t, v_m, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\tilde{\beta}_{m-1, n_t}^t$ in Equation (6). As above, the fact that $v_m < v_{i_{k_t}}$ contradicts the conclusion of the claim above regarding the monotonicity of bids. Therefore, i does not desire to remain active in the auction long enough to receive an object.

Thus, we have shown that no player has any incentive to deviate from the prescribed strategies when on the equilibrium path. In particular, using the bidding strategies $\tilde{\beta}_{m, n_t}^t$ is sequentially rational given players' beliefs along the equilibrium path. Recall, however, that we have specified off-equilibrium beliefs such that buyers "ignore" their past observations when they observe a deviation from equilibrium play, updating their beliefs to place full probability on the valuation that rationalizes the deviation; they believe that the deviating agent is *currently* being truthful with regards to the strategies $\tilde{\beta}_{m, n_t}^t$. The argument above then implies that continuing to bid according to the specified strategies remains sequentially rational with respect to these updated beliefs. Thus, bidding according to the cutoffs in Equation (6) is optimal along the entire game tree: this strategy profile forms a perfect Bayesian equilibrium of the sequential ascending auction mechanism. \square

Thus, bidding in each period according to the strategy described in Equation (6) forms a perfect Bayesian equilibrium of the sequential ascending auction with reserve price \tilde{r} ; moreover, this equilibrium is outcome equivalent to the dynamic virtual pivot mechanism. \square

REFERENCES

- ATHEY, S., AND I. SEGAL (2007): "An Efficient Dynamic Mechanism," Unpublished manuscript, Harvard University.
- AUSUBEL, L. M. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94(5), 1452–1475.
- BANKS, J. S., J. O. LEDYARD, AND D. P. PORTER (1989): "Allocating Uncertain and Unresponsive Resources: An Experimental Approach," *The RAND Journal of Economics*, 20(1), 1–25.
- BERGEMANN, D., AND J. VÄLIMÄKI (2008): "The Dynamic Pivot Mechanism," Cowles Foundation Discussion Paper 1672, Yale University.
- BIKHCHANDANI, S., P. A. HAILE, AND J. G. RILEY (2002): "Symmetric Separating Equilibria in English Auctions," *Games and Economic Behavior*, 38(1), 19–27.
- CAILLAUD, B., AND J. ROBERT (2005): "Implementation of the Revenue-Maximizing Auction by an Ignorant Seller," *Review of Economic Design*, 9(2), 127–143.
- CAVALLO, R., D. C. PARKES, AND S. SINGH (2009): "Efficient Mechanisms with Dynamic Populations and Dynamic Types," Unpublished manuscript, Harvard University.
- DASGUPTA, P., AND E. MASKIN (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115(2), 341–388.
- DEMANGE, G., D. GALE, AND M. SOTOMAYOR (1986): "Multi-Item Auctions," *Journal of Political Economy*, 94(4), 863–872.
- EDELMAN, B., AND M. SCHWARZ (2006): "Optimal Auction Design in a Multi-Unit Environment: The Case of Sponsored Search Auctions," Unpublished manuscript, Harvard Business School.
- FUDENBERG, D., AND J. TIROLE (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53(2), 236–260.
- GERSHKOV, A., AND B. MOLDOVANU (2008): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, forthcoming.
- (2009): "Efficient Sequential Assignment with Incomplete Information," *Games and Economic Behavior*, forthcoming.
- HENDON, E., H. J. JACOBSEN, AND B. SLOTH (1996): "The One-Shot-Deviation Principle for Sequential Rationality," *Games and Economic Behavior*, 12(2), 274–282.
- KAGEL, J. H., R. M. HARSTAD, AND D. LEVIN (1987): "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55(6), 1275–1304.
- KAGEL, J. H., AND D. LEVIN (2009): "Implementing Efficient Multi-Object Auction Institutions: An Experimental Study of the Performance of Boundedly Rational Agents," *Games and Economic Behavior*, 66(1), 221–237.
- KITTSTEINER, T., J. NIKUTTA, AND E. WINTER (2004): "Declining Valuations in Sequential Auctions," *International Journal of Game Theory*, 33(1), 89–106.
- KRISHNA, V. (2003): "Asymmetric English Auctions," *Journal of Economic Theory*, 112(2), 261–288.
- KURINO, M. (2009): "House Allocation with Overlapping Agents: A Dynamic Mechanism Design Approach," Unpublished manuscript, University of Pittsburgh.

- LAVI, R., AND N. NISAN (2005): "Online Ascending Auctions for Gradually Expiring Items," in *Proceedings of the 16th Symposium on Discrete Algorithms (SODA 2005)*, Vancouver.
- LAVI, R., AND E. SEGEV (2009): "Efficiency Levels in Sequential Auctions with Dynamic Arrivals," Unpublished manuscript, Technion.
- MASKIN, E., AND J. G. RILEY (1989): "Optimal Multi-Unit Auctions," in *The Economics of Missing Markets, Information, and Games*, ed. by F. Hahn, pp. 312–335. Oxford University Press.
- MIERENDORFF, K. (2008): "An Efficient Intertemporal Auction," Unpublished manuscript, University of Bonn.
- MILGROM, P. R., AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5), 1089–1122.
- (2000): "A Theory of Auctions and Competitive Bidding, II," in *The Economic Theory of Auctions*, ed. by P. Klemperer, vol. 2, pp. 179–194. Edward Elgar Publishing, Cheltenham, UK.
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73.
- PAI, M., AND R. VOHRA (2008): "Optimal Dynamic Auctions," Unpublished manuscript, Northwestern University.
- PARKES, D. C. (2007): "Online Mechanisms," in *Algorithmic Game Theory*, ed. by N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, chap. 16, pp. 411–439. Cambridge University Press, Cambridge, UK.
- PARKES, D. C., AND S. SINGH (2003): "An MDP-Based Approach to Online Mechanism Design," in *Proceedings of the 17th Annual Conference on Neural Information Processing Systems (NIPS '03)*.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2009): "Dynamic Mechanism Design: Revenue Equivalence, Profit Maximization and Information Disclosure," Unpublished manuscript, Northwestern University.
- PEREA, A. (2002): "A Note on the One-Deviation Property in Extensive Form Games," *Games and Economic Behavior*, 40(2), 322–338.
- RAHMAN, D. (2009): "Dynamic Implementation," Unpublished manuscript, University of Minnesota.
- ROTHKOPF, M. H., T. J. TEISBERG, AND E. P. KAHN (1990): "Why Are Vickrey Auctions Rare?," *Journal of Political Economy*, 98(1), 94–109.
- SAID, M. (2008): "Information Revelation in Sequential Ascending Auctions," Unpublished manuscript, Yale University.
- (2009): "Sequential Auctions with Randomly Arriving Buyers," Unpublished manuscript, Yale University.
- SKRETA, V. (2007): "Optimal Auctions with General Distributions," Unpublished manuscript, New York University.
- ÜNVER, M. U. (2009): "Dynamic Kidney Exchange," *Review of Economic Studies*, forthcoming.
- VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16(1), 8–37.
- VULCANO, G., G. VAN RYZIN, AND C. MAGLARAS (2002): "Optimal Dynamic Auctions for Revenue Management," *Management Science*, 48(11), 1388–1407.