



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

12 July 2010

Online at <https://mpra.ub.uni-muenchen.de/23852/>
MPRA Paper No. 23852, posted 13 Jul 2010 09:07 UTC

AUCTIONS WITH DYNAMIC POPULATIONS: EFFICIENCY AND REVENUE MAXIMIZATION

MAHER SAID[†]

JULY 12, 2010

ABSTRACT: We study indirect mechanisms in a setting where both objects and privately-informed buyers arrive stochastically to a market. The seller in this setting faces a sequential allocation problem, and must elicit the private information of the dynamic population of buyers in order to achieve her desired outcome—either an efficient or a revenue-maximizing allocation. We propose a simple indirect mechanism, the sequential ascending auction, which yields outcomes identical to those of an efficient dynamic Vickrey-Clarke-Groves mechanism. We construct equilibria in memoryless strategies where strategic bidders are able to reach the efficient outcome by revealing all private information in every period, inducing behavior that is symmetric across both incumbent and newly entered buyers. In contrast to settings with a static population, a sequence of second-price auctions cannot yield this outcome, as these auctions do not reveal sufficient information to symmetrize different cohorts. We also extend our results to revenue-maximization, showing that the sequential ascending auction with a reserve price is optimal.

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

JEL CLASSIFICATION: C73, D44, D82, D83.

[†]MICROSOFT RESEARCH NEW ENGLAND AND OLIN BUSINESS SCHOOL, WASHINGTON UNIVERSITY IN ST. LOUIS, SAID@WUSTL.EDU

This paper is a substantial revision of the third chapter of my doctoral dissertation at Yale University, and also supersedes and extends results first reported in Said (2008). I am indebted to Dirk Bergemann, Johannes Hörner, Ben Polak, and Larry Samuelson for their generous advice, comments, and encouragement. I also thank Rossella Argenziano, Alessandro Bonatti, Rahul Deb, Eduardo Faingold, James Fenske, Philip Haile, Sergei Izmalkov, Katrina Jessoe, Nenad Kos, David Parkes, Mallesh Pai, Kareen Rozen, Ilya Segal, Itai Sher, Juuso Välimäki, and Rakesh Vohra, as well as seminar participants at the AMMA 2009 conference, an NBER Market Design Working Group Meeting, and several universities, for their many valuable comments and discussions. Financial support from the National Science Foundation, the Cowles Foundation, and Yale University is gratefully acknowledged.

To an economic theorist, the above auction design is not a difficult concept. However, experience has shown that even economics Ph.D. students have trouble understanding the above description. . . . The problem is that if people do not understand the payment rules of the auction then we do not have any confidence that the end result will be efficient.

Nalebuff and Bulow (1993, p. 29)

1. INTRODUCTION

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 may not know when the next auction for a similar item will take place, nor how many such future auctions will occur. Similarly, future competition—the number of potential competitors, as well as their strength—is unknown to our buyer.

At the same time, this hypothetical buyer on eBay has available to her a wealth of other 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 potential 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 a question of fundamental importance: how may we achieve efficient or revenue-maximizing outcomes in markets with dynamic populations of privately-informed buyers?

We answer this question by developing a reasonably general model of a dynamic environment that reflects some key features of markets where dynamic populations are important. This model abstracts away many of the details of specific dynamic markets, focusing instead on 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 elicit the buyers' private information in order to achieve a desirable allocation.

It is well-known that dynamic variants of the classic Vickrey-Clarke-Groves mechanism may be used to achieve efficient outcomes in dynamic environments, and that they may also be “tweaked” for the purpose of revenue maximization. Such mechanisms 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.” This echoes a criticism made by [Nalebuff and Bulow \(1993\)](#) in their proposal to the FCC for the sale of wireless spectrum. As demonstrated by the quote in the epigraph above, [Nalebuff and Bulow](#) expressed concerns that “real-world” bidders, even reasonably sophisticated ones, would experience difficulties in understanding the price-determination schemes that form the basis of VCG mechanisms, thereby undermining the mechanisms' desirability.

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 look to design simple, transparent, and decentralized indirect mechanisms that serve as viable alternatives to their direct counterparts. In particular, we consider the possibility of achieving efficient or revenue-maximizing outcomes via a sequence of auctions. Despite the similarity of dynamic VCG-based 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 sealed-bid auction. In an environment with a dynamic population of buyers, however, a sequence of second-price sealed-bid auctions cannot yield outcomes equivalent to those of the dynamic VCG mechanism. This is due to the nature of the “option value” associated with losing in one of a sequence of auctions, as a losing buyer may win in a future auction. The value of this option depends on expected future prices, which are in turn determined by the private information of other competitors. Thus, despite an underlying independent-private-values framework, the strategic environment faced by individual bidders is more complicated, as the auction market dynamics induce interdependence in (option) values.¹

This interdependence implies that a standard second-price sealed-bid auction does not reveal sufficient information for the determination of buyers’ option values, as different “generations” of buyers have observed different histories and hence have asymmetric beliefs. In contrast, the ascending auction is a simple *open* auction format that does allow for the gradual revelation of buyers’ private information in a way that symmetrizes beliefs across cohorts.² 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 feature a strong “no regret” property: they form a periodic ex post equilibrium. Given expectations about future arrivals and behavior, each buyer’s behavior in any period remains optimal even after having observing her current opponents’ actions.

¹Similar phenomena have been noted in auctions with downstream interaction among buyers or with resale opportunities. See, for instance, Jehiel and Moldovanu (2000) or Haile (2003).

²This is in line with findings from the sequential common agency literature (as in Calzolari and Pavan (2006, 2009)) which shows that mechanisms must control not only for the allocations induced, but also the information revealed.

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.³ Bergemann and Välimäki (2010) 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, 2010) 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.⁴ Moreover, our fundamental point of departure from this literature is our focus on the design of *indirect* mechanisms.

This paper also relates to recent work on dynamic auctions and revenue management. Li (2009) and Board and Skrzypacz (2010) consider the efficient and optimal mechanisms, respectively, in a finite-horizon version of our model with storable goods; Li also finds (as we do) that an open auction format is required for indirect implementation. Mierendorff (2010) solves for the revenue-maximizing auction when buyers have privately-known departure times from the market, while in a more general setting Pai and Vohra (2009) 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 (2009, 2010a) derive efficient and revenue-maximizing 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 assume 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 (2010) 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

³Bergemann and Said (2010) and Parkes (2007) survey much of this literature, with the former focusing primarily on work in the economics literature, and the latter on the contributions from computer science.

⁴There 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 (2010).

turned attention towards sequential 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. We introduce the general model and environment in Section 2. In Section 3, we describe the benchmark efficient direct mechanism and then propose an indirect mechanism, the sequential ascending auction. We demonstrate that it yields equilibrium outcomes equivalent to the benchmark direct mechanism. Then, in Section 4 we discuss the extension of our results to revenue-maximization (leaving the formal derivation to an appendix). Finally, we conclude with Section 5.

2. MODEL

We consider an infinite-horizon, discrete-time environment with a single seller; time periods are indexed by $t \in \mathbb{N}$. In each period t , the seller has K_t units of a homogenous and indivisible good available for sale. The number of objects available in each period is a random variable drawn independently from the distribution μ_t on the nonnegative integers. We assume that objects are non-storable: any objects that are not allocated or sold “expire” at the end of each period, and hence cannot be carried over to future periods.

The beginning of each period t is marked by the arrival to the market of N_t buyers from a countable set \mathcal{I} . As with the number of objects available in each period, the number of arriving buyers may also vary, with N_t an independent draw from the distribution λ_t on the nonnegative integers. We denote by \mathcal{I}_t the finite subset of buyers that arrive on the market in period t . Each buyer i present on the market wishes to obtain a single unit of the seller’s good, and is endowed with a privately-known value v_i for that single unit. We assume that v_i is an independent draw from the distribution F_i on $\mathbf{V} := [0, \bar{v}]$.

In addition, we assume that buyers may exogenously 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 on the market until they are able to obtain an object. Finally, we assume that buyers are risk neutral, and that their preferences are quasilinear and time separable. All buyers, as well as the seller, discount the future with the common discount factor $\delta \in (0, 1)$.

As described above, new buyers may arrive to the market (and “old” buyers may depart) at the beginning of each period. These arrivals and departures yield a stochastic process $\{\alpha_t\}_{t \in \mathbb{N}}$, 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 .⁵ Simultaneously,

⁵This assumption is merely for simplicity. As will be discussed later in Section 3.3, the equilibria we construct and describe remain equilibria in the “larger” game where buyers may conceal their presence. This need not be true, however, if buyer arrivals are correlated over time—see Gershkov and Moldovanu (2010b).

new objects arrive, replacing any unallocated objects left over from the previous period. 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 . It will be convenient to denote the “state” of the market at the beginning of each period t by $\omega_t := (\alpha_t, K_t)$. Note, however, that buyers arriving in period t do *not* observe the state of the market in previous periods—buyers do not know when other buyers arrived on the market or how many objects were allocated in previous periods. Once the current state is known, the seller then allocates objects to buyers and makes (monetary) transfers, and we move on to the following period.

3. EFFICIENT IMPLEMENTATION

3.1. A Benchmark Efficient (Direct) Mechanism

We begin by examining the case in which the seller is a benevolent social planner whose goal is to maximize allocative efficiency; that is, the planner wishes to choose a feasible allocation rule $\{x_{i,t}\}$ to maximize

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}} \delta^{t-1} x_{i,t} v_i \right],$$

where $x_{i,t}$ is the probability that the seller allocates an object to buyer i in period t , and the expectation is taken with respect to the arrival and departure processes, as well as agents’ valuations. An allocation is *feasible* if

- (1) for all $t \in \mathbb{N}$, $x_{i,t} = 0$ if $\alpha_t(i) = 0$;
- (2) $\sum_{i \in \mathcal{I}} x_{i,t} \leq K_t$ for all $t \in \mathbb{N}$;
- (3) $\sum_{t \in \mathbb{N}} x_{i,t} \leq 1$ for all $i \in \mathcal{I}$; and
- (4) $x_{i,t}$ is adapted to the seller’s information.

Condition (1) requires that no only buyers that are physically present on the market are allocated an object. Conditions (2) requires that no more objects be allocated at any time than are actually available, while condition (3) requires that no no buyer receives more than the single object that she desires. Finally, condition (4) reflects the fact that the seller is not prescient, and hence cannot make allocations that are contingent on future information—period- t allocations must be made based only on the reports of buyers arriving by period t .

Recall that in a 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 similarly straightforward. Notice that objects are perishable, implying that there is no potential benefit to be gained by “withholding” an object or not allocating as many objects as possible. Note also that buyers’ values are persistent over time and are independent of the common departure rate. Therefore, since all buyers discount the future with the same discount factor, the cost of delaying allocation to any given buyer is symmetric and strictly increasing in values.

Thus, the efficient allocation rule is akin to an assortative matching: in each period t , allocate all K_t objects to the K_t highest-valued buyers present on the market. In the event that multiple

buyers have precisely the same value, ties may be broken arbitrarily. Notice that this policy is history-independent—the period- t efficient allocation depends only on the number of objects available (K_t), the buyers present (indicated by α_t), and these agents' reported values (denoted by \mathbf{v}_t). Thus, we will denote this rule by $\widehat{\mathbf{x}} := \{\widehat{x}_{i,t}(\omega_t, \mathbf{v}_t)\}$.

Since buyers' private information is single-dimensional and values are independent, a dynamic version of the standard Vickrey-Clarke-Groves mechanism can implement this efficient policy $\widehat{\mathbf{x}}$.⁶ In particular, we will focus on the *dynamic pivot mechanism* of Bergemann and Välimäki (2010). They show that the dynamic pivot mechanism implements socially efficient policies in a general dynamic setting with independent private values. The mechanism sets transfers such that, in each period, agents receive their flow (per-period) marginal contribution to the social welfare. This payment rule is both *periodic ex post incentive compatible* and *individually rational*: in each period, given expectations about the future, every agent present finds it optimal to participate truthfully in the mechanism, even after observing all private information up to and including that period.

In order to fully describe the dynamic pivot mechanism in our setting, we require the following definitions. For any state $\omega_t = (\alpha_t, K_t)$ and reported values \mathbf{v}_t , let

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

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

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

for the social welfare (from period t on) when i is removed from the market and the efficient policy $\widehat{\mathbf{x}}$ is implemented. Agent i 's (total) marginal contribution to the social welfare is then

$$w_i(\omega_t, \mathbf{v}_t) := W(\omega_t, \mathbf{v}_t) - W_{-i}(\omega_t, \mathbf{v}_t).$$

Similarly, her flow contribution in period t is simply the difference between her marginal contribution in the present and her expected *future* marginal contribution, which is given by

$$\begin{aligned} \widehat{w}_i(\omega_t, \mathbf{v}_t) &:= w_i(\omega_t, \mathbf{v}_t) - \delta \mathbb{E} [w_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t] \\ &= \underbrace{W(\omega_t, \mathbf{v}_t) - W_{-i}(\omega_t, \mathbf{v}_t)}_{\text{total contribution}} - \underbrace{\delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) - W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t]}_{\text{expected future contribution}}. \end{aligned}$$

Notice that when $\widehat{x}_{i,t}(\omega_t, \mathbf{v}_t) = 1$, buyer i leaves the market, and therefore she has no future contribution to the social welfare. Similarly, when $\widehat{x}_{i,t}(\omega_t, \mathbf{v}_t) = 0$, buyer i 's presence on the market has no impact on the current-period allocation, and hence her flow contribution is zero. Thus, for all (ω_t, \mathbf{v}_t) we have

$$\widehat{w}_i(\omega_t, \mathbf{v}_t) = \widehat{x}_{i,t}(\omega_t, \mathbf{v}_t) w_i(\omega_t, \mathbf{v}_t).$$

⁶Dolan (1978) was the first to propose the use of VCG-like mechanisms in settings with a dynamic population of buyers.

The dynamic pivot mechanism is then simply the mechanism $\widehat{\mathcal{M}} := \{\widehat{\mathbf{x}}, \widehat{\mathbf{p}}\}$, where the payment rule $\widehat{\mathbf{p}}$ is defined by

$$\begin{aligned}\widehat{p}_{i,t}(\omega_t, \mathbf{v}_t) &:= \widehat{x}_{i,t}(\omega_t, \mathbf{v}_t)v_i - \widehat{w}_i(\omega_t, \mathbf{v}_t) \\ &= \widehat{x}_{i,t}(\omega_t, \mathbf{v}_t)(v_i - w_i(\omega_t, \mathbf{v}_t))\end{aligned}\tag{1}$$

for all (ω_t, \mathbf{v}_t) . Thus, this mechanism yields a flow payoff to each buyer equal to her flow marginal contribution (implying a net payoff equal to her total marginal contribution), thereby incentivizing truthful behavior and allowing the implementation of the efficient policy. Unlike the VCG mechanism in static settings, however, the dynamic pivot mechanism is *not* dominant-strategy incentive compatible—even though buyers make only a single report upon their arrival. This is because, in contrast to static VCG, the payments $\widehat{p}_{i,t}$ are *not* distribution-free, and so beliefs about other players and their strategies matter for the determination of optimal behavior. However, [Bergemann and Välimäki \(2010, Theorem 1\)](#) show that the dynamic pivot mechanism is periodic ex post incentive compatible and individually rational. We restate their result here for completeness.

LEMMA 1 (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.*⁷

3.2. Towards an Efficient Indirect Mechanism

The dynamic pivot mechanism described above 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.⁸

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 are present in the initial period, [Kittsteiner, Nikutta, and Winter \(2004\)](#) show that a sequence of such auctions is, in fact, equivalent to dynamic 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

⁷[Bergemann and Välimäki \(2010\)](#) 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.

⁸Of 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.

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? Since it is an expectation of future payoffs from participating in the sequence of auctions, it must incorporate the buyer's expectations of future prices. But these future 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.⁹ 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; loosely speaking, buyer i will bid

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

where the expectation is conditional on her bid being pivotal (conditional on being tied for the K_t -th and $(K_t + 1)$ -st highest bids). In sharp contrast, however, to the sequential sealed-bid 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.

⁹It is well-known (see, for instance, [Dasgupta and Maskin \(2000\)](#) or [Jehiel and Moldovanu \(2001\)](#)) 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—our underlying environment is a standard independent private values setting, and hence the efficient policy is implementable by the *direct* mechanism.

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.

These difficulties persist in the face of other attempts to symmetrize beliefs. Another natural avenue by which to attempt this—in the case when buyers are *ex ante* symmetric—is to publicly release information about winning bids to both incumbent buyers and new entrants, leading to conditional beliefs about competing buyers that are truncations of the initial distributions. However, since buyers do not observe the arrival times of buyers that precede them, members of different cohorts who have had different experiences will continue to be differentially informed. Therefore, when buyers condition their bids on being pivotal, they will truncate at different points and continue to bid asymmetrically.

These asymmetries suggest 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, symmetrizing buyers’ beliefs and leading to bids that correctly account for the interdependence generated by market dynamics: buyers can essentially submit bids

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

perfectly accounting for interdependence and thereby allowing for an efficient outcome.

3.3. An Efficient Sequential Auction

To be more formal, 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. 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 $K_t \geq 1$ objects for sale, the auction ends whenever at most K_t 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 K_t 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 $k < K_t$ bidders active. The auction ends at this point, and $K_t - k$ of the “tied” bidders

are selected uniformly at random to receive an object—along with the remaining active bidders—paying the price at which the auction closed.¹⁰ 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.¹¹ 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 a best response to their competitors’ behavior.

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 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 on this information as it is revealed (or re-revealed) *during* the current period.

¹⁰The 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.

¹¹This is a crucial difference with the scheduling example in Bergemann and Välimäki (2010). Their example has no population uncertainty, and hence no additional informational asymmetries; therefore, the efficient outcome is also achievable by a sequence of first- or second-price sealed-bid auctions as in Kittsteiner, Nikutta, and Winter (2004).

With this in mind, we now informally describe the memoryless strategies used by each player in the sequential ascending auction mechanism. In each period, buyers will bid sincerely, remaining active in the auction until the price reaches the point at which they are 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:

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

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, let

$$\mathbf{y}_t := (y_t^1, \dots, y_t^{n_t})$$

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

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

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 $\mathbf{y}_t^{>k}$, allowing their bids to be conditioned on this information.

Finally, we define, for each $k = 2, \dots, n_t$,

$$w^{t+1}(\omega_t, v, \mathbf{y}_t^{>k}) := \delta \mathbb{E} [w_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (v, \dots, v, \mathbf{y}_t^{>k})]$$

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

With these preliminaries in hand, we may now define the strategies used by each bidder. We define, for each $k = 2, \dots, n_t$,

$$\hat{\beta}_{k, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) := v_i - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}). \quad (2)$$

In each period $t \in \mathbb{N}$, each agent $i \in \mathcal{A}_t$ bids according to $\hat{\beta}_{k, n_t}^t$ when she is one of k active competitors in the auction. Thus, each buyer i initially bids up to the point at which she is indifferent between winning the object at the current price and receiving her discounted expected marginal contribution in the next period.¹² At the outset of the auction, when no other competitors have dropped out yet, a buyer is only able to win at the current price if the other buyers exit simultaneously. Thus, she is pivotal and will win only if she is exactly tied with all other buyers present on the market. 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.

¹²See Said (2008) for an explicit closed-form example of these strategies (derived from equilibrium considerations alone) in a special case of the present model.

Recall that an efficient allocation assigns objects to buyers in decreasing order of their values, implying that the presence of a buyer i with value v_i does not affect the timing of allocation to buyers j with $v_j > v_i$. Furthermore, since ties among buyers may be broken arbitrarily without affecting the social surplus, there is no trade-off between allocating to i or to any other buyer j with value $v_j = v_i$. Thus, the marginal contribution to the social welfare made by a buyer i with value v_i is independent of the values of any buyers j with $v_j \geq v_i$. As this is true for any state and any configuration of values (and since $v_i \leq \bar{v}$), this implies that

$$\delta \mathbb{E} \left[w_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (v_i, \dots, v_i, \mathbf{y}_t^{>k}) \right] = \delta \mathbb{E} \left[w_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{k-1}, v, \mathbf{y}_t^{>k}) \right];$$

that is, the expected future marginal contribution $w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k})$ of a buyer i with $n_t - k$ competitors whose values $\mathbf{y}_t^{>k}$ are less than v_i and $k - 1$ competitors whose values at least as large as v_i does *not* depend upon the values of those latter competitors.

This observation implies that, for any $v_i > v'_i$, the difference $w^{t+1}(\omega_t, v_i) - w^{t+1}(\omega_t, v'_i)$ equals

$$\begin{aligned} & \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i)] - \delta \mathbb{E} [W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i)] \\ & - \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v'_i)] + \delta \mathbb{E} [W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v'_i)] \\ & = \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i)] - \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v'_i)], \end{aligned}$$

since the social welfare when i is not present on the market does not depend upon i 's value (and hence the second and fourth terms are equal). Moreover, it is easy to find an upper bound for this difference. In particular, note that the second term is the social welfare when taking the efficient policy for the case where i 's value is v'_i . By definition, this is no smaller than the social welfare when i 's value is v'_i but she is treated as though it is v_i . Therefore, we have

$$\begin{aligned} & \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i)] - \delta \mathbb{E} [W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v'_i)] \\ & \leq \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] (v_i - v'_i). \end{aligned}$$

Thus, $v_i > v'_i$ implies that

$$\hat{\beta}_{n_t, n_t}^t(\omega_t, v_i) - \hat{\beta}_{n_t, n_t}^t(\omega_t, v'_i) \geq \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \hat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-1}, v_i) \right] \right) (v_i - v'_i).$$

This difference is, of course, strictly positive since the expected discounted probability of receiving an object in the future is bounded above by $\delta < 1$. Thus, if all buyers follow the bidding strategies described in [Equation \(2\)](#), the buyer 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 the lowest-ranked buyer present has value $y_t^{n_t} < v_i$. As before, the buyer is pivotal and can win only if she is exactly tied with the remaining buyers. Thus, buyer i will remain active until she is indifferent between winning and continuing on to the future. So 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

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

However, we may rewrite the second difference term above may be as

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

Thus, the difference in future expected contributions $w^{t+1}(\omega_t, v_i, v_j) - w^{t+1}(\omega_t, v_j)$ is the sum of three differences: the first is the change in social welfare when i 's value increases from v_i to \bar{v} ; the second is the change in social welfare—with i removed from the market—when j 's value increases from v_j to v_i ; and finally, the third is the change in social welfare—with j removed from the market—when i 's value decreases 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 exactly offsets the loss from the third, implying that

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

A bounding argument similar to the one applied above with v_i and v_i' may then be used to show that the difference in bids

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

is positive. Thus, there is “continuity” at the first drop out point, in the sense that the exit of the lowest-valued buyer j 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 $\hat{\beta}_{k, n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k})$ is strictly increasing in v_i for all k , and that the “continuity” property described above holds after every exit from the auction. We summarize these facts in the following lemma.

LEMMA 2 (Bids are fully separating).

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

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

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 in [Theorem 2](#)), a standard revenue-equivalence result 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: the prices paid by bidders who follow the strategies described in Equation (2) are exactly equal to the prices they would pay in the truthful equilibrium of the dynamic pivot mechanism. To understand why this is the case, it is helpful to consider the externality imposed by a buyer i who receives an object. If that buyer i is removed from the market, a lower-ranked buyer j will instead be allocated an object, yielding an immediate gain in the surplus equal to her value v_j . However, buyer j will no longer be available for allocations in future periods, implying that her *future* marginal contribution is lost. This difference, however, is exactly the bid being made by buyer j in the ascending auction. Therefore, the higher-ranked buyer i pays a price equal to the externality she imposes on the market, yielding a payoff equal to her marginal contribution to the social welfare.¹³

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

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

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

Therefore, following the memoryless bidding strategies prescribed in Equation (2) leads to an outcome that is identical to that of truth-telling in the dynamic pivot mechanism. Moreover, we know from Lemma 1 that truth-telling is an equilibrium of the dynamic pivot mechanism. It remains to be shown, however, that the bidding strategies described in Equation (2) form an equilibrium of the sequential ascending auction mechanism.

This follows from two main observations. First, the use of memoryless strategies by a buyer's opponents incentivizes sincere bidding, as the competitors' future behavior is not contingent on information revealed in the current period. Thus, there is no value to be found in trying to alter others' beliefs about one's own private information. Secondly, buyers' marginal contributions to the social welfare are sufficiently well-behaved to guarantee that the induced interdependence due to downstream interaction still satisfies standard single-crossing conditions, thereby admitting an efficient equilibrium in the ascending auction. Thus, sincere bidding in accordance with Equation (2) is an equilibrium of the sequential ascending auction mechanism.

THEOREM 2 (Equilibrium in the sequential ascending auction).

Suppose that in each period, buyers bid according to the memoryless strategies described in Equation (2). 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.

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

¹³Marginal contributions and externalities were first used to construct efficient bidding equilibria by Bergemann and Välimäki (2003, 2006) in complete-information settings with first-price auctions.

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.¹⁴ 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 $\hat{\beta}_{k,n_t}^t$ are not rational for any agent, even when conditioning on competitors' values in the current period. Thus, the strategy profile specified in Equation (2) forms a periodic ex post equilibrium—the strong “no regret” property of the ascending auction in static settings is inherited here. Since agents essentially “re-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 mechanism 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.¹⁵

4. REVENUE MAXIMIZATION

In the static setting, Myerson (1981) provided an optimal mechanism for selling a single indivisible unit. Essentially, Myerson's optimal auction simply uses the standard Vickrey-Clarke-Groves mechanism to maximize virtual surplus instead of social surplus. When agents report their values v_i , the mechanism computes their virtual values $\varphi_i(v_i)$ and applies the VCG mechanism to these virtual values. This yields an allocation (the buyer with the highest non-negative virtual value) and a “virtual price” such that the winning buyer's virtual value less the virtual price equals her marginal contribution to the virtual surplus. A large literature building upon the Myersonian tradition has shown more generally that, in the presence of single-dimensional private information, optimal mechanisms are often just efficient mechanisms acting on virtual values.

Notice that in our setting, while the objects are individual units of a homogeneous good, they are differentiated products from the perspective of an individual buyer. 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.¹⁶

¹⁴We should point out that this equilibrium is not unique among efficient equilibria. Analogous to the multiplicity of symmetric separating 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.”

¹⁵The independence of arrivals across time periods is crucial here, as there are no informational externalities from arrivals that can induce additional interdependence. For more on this, see Gershkov and Moldovanu (2010b).

¹⁶This is similar to the heterogeneity in advertising slots in the keyword auction of Edelman and Schwarz (2010). There, slots are differentiated by their click-through-rates, whereas here objects are distinguished by their arrival times.

Notice, however, that buyers' preferences are additively separable across time and risk-neutral, implying that buyers' preferences over the temporally differentiated objects may be summarized by their expected discounted probability of receiving an object and their single-dimensional value for that object. This implies that the problem faced by a revenue-maximizing seller (who fully commits to a mechanism at time 0) may be solved by application of an efficient mechanism to virtual values. In particular, the mechanism we propose is the *dynamic virtual pivot mechanism*, the dynamic pivot mechanism discussed above applied to virtual values. Under the standard assumption that virtual values are monotone, this mechanism maximizes revenues and inherits the incentive compatibility and individual rationality properties of the efficient dynamic pivot mechanism. This is summarized by the following lemma (which is developed in detail in [Appendix B](#)).

LEMMA 3 (Revenue maximization via the dynamic virtual pivot mechanism).

Suppose that virtual values are increasing. Then the dynamic virtual pivot mechanism is periodic ex post incentive compatible and individually rational. Moreover, this mechanism maximizes the seller's revenue.

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 *transparent* and *decentralized* auction mechanism? In light of the mechanism's relationship to 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 $(k + 1)$ -st-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} , the solution to

$$\varphi(v) := v - \frac{1 - F(v)}{f(v)} = 0. \quad (3)$$

Such a mechanism is outcome equivalent to a $(k + 1)$ -st-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 is 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.¹⁷ More formal arguments similar to those of [Theorem 1](#) and [Theorem 2](#) (which we discuss in detail in [Appendix B](#)) yield the following result.

¹⁷In 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 for asymmetric bidders—see [Caillaud and Robert \(2005, Proposition 1\)](#).

THEOREM 3 (Revenue maximization via sequential ascending auctions).

Suppose that $F_i = F$ for all $i \in \mathcal{I}$ and that φ is strictly increasing. Then the sequential ascending auction with reserve price $\tilde{r} := \varphi^{-1}(0)$ admits a (periodic ex post) perfect Bayesian equilibrium that is outcome equivalent to the dynamic virtual pivot mechanism.

Thus, analogous to the case of [Section 3.3](#), 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 with an appropriately chosen reserve price.

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 buyers demand multiple units of potentially complementary goods. This 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 approaches cannot resolve the fundamental asymmetry in preferences and objectives that arise when some buyers have already satisfied a portion of their demand. Therefore, different institutional forms will be required. 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 2. Fix an arbitrary period $t \in \mathbb{N}$, 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 $n_t - k$ buyers have dropped out of the period- t auction, revealing values $\mathbf{y}_t^{>k}$, where $k \in \{1, \dots, n_t\}$. We wish to show first that $v_i > v'_i > y_t^{k+1}$ (where we take $y_t^{n_t+1} := 0$ by convention) implies that

$$\widehat{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) := v_i - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) > v'_i - w^{t+1}(\omega_t, v'_i, \mathbf{y}_t^{>k}) =: \widehat{\beta}_{k,n_t}^t(\omega_t, v'_i, \mathbf{y}_t^{>k}).$$

Notice that

$$\begin{aligned} & w^{t+1}(\omega_t, v'_i, \mathbf{y}_t^{>k}) - w^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) \\ &= \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k})] - \delta \mathbb{E}[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k})] \\ & \quad - \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k})] + \delta \mathbb{E}[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k})] \\ &= \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k})] - \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k})], \end{aligned}$$

since the social welfare when i is not present on the market does not depend upon her value (and hence the second and fourth terms are equal). Moreover, by treating buyer i with value v'_i as though her true value were v_i , we can provide a bound on the difference above. In particular, we have

$$\begin{aligned} & \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k})] - \delta \mathbb{E}[W(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k})] \\ & \geq \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \widehat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right] (v'_i - v_i). \end{aligned}$$

Thus, if $v_i > v'_i$, then

$$\begin{aligned} & \widehat{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) - \widehat{\beta}_{k,n_t}^t(\omega_t, v'_i, \mathbf{y}_t^{>k}) \\ & \geq (v_i - v'_i) \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \widehat{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, \mathbf{y}_t^{>k}) \right] \right) > 0, \end{aligned}$$

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

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

However, the second difference above may be rewritten as

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

Thus, the difference in expected future contributions $w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>k+1}) - w^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>k+1})$ 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 $w^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>k+1}) - w^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>k+1}) =$

$$\delta(\mathbb{E}[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, v_j, \mathbf{y}_t^{>k+1})] - \mathbb{E}[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, v_i, \mathbf{y}_t^{>k+1})]).$$

Moreover, by treating buyer j as though her true value were v_i (as we did above with v'_i and v_i), we can provide a bound on the difference above, which may be used to show that

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

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

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

PROOF OF THEOREM 1. We established with [Lemma 2](#) that following the bidding strategies $\hat{\beta}_{k, n_t}^t$ prescribed in [Equation \(2\)](#) 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 $t \in \mathbb{N}$, and let K_t denote the number of objects present, and $n_t := |\mathcal{A}_t|$ denote the number of agents present. As shown by [Lemma 2](#), the auction allocates the K_t objects to the group of buyers with the K_t highest values. (Recall that the total and flow marginal contributions are equal for buyers who receive an object.) 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. 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 sequential auction 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_k the bidder with the k -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

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

Expanding the summation terms above yields

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

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 $(K_t + 1)$ -st highest value, conditional on agents with higher values (which includes i) receiving an object today.

Thus, bidding according to $\widehat{\beta}_{k,n_t}^t$ leads to the allocations and prices identical to those of the dynamic pivot mechanism. \square

PROOF OF THEOREM 2. As the sequential ascending auction is a dynamic game of incomplete information, we use perfect Bayesian equilibrium as our solution concept. This 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 $\widehat{\beta}_{k,n_t}^t$, 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. Such post-deviation histories are zero probability events, and so 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* behaving sincerely in accordance with $\widehat{\beta}_{k,n_t}^t$.

Note that these beliefs are consistent with Bayes' rule even after probability zero histories.¹⁸ 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 *preconsistency* of beliefs in an extensive form game of incomplete information, a condition put forth by Hendon, Jacobsen, and Sloth (1996) and shown by Perea (2002) to be both necessary and sufficient for the one-shot-deviation principle to hold.¹⁹ Since perfect Bayesian equilibrium need not satisfy the one-shot deviation principle, this observation allows us to use the principle in our proof.

¹⁸These 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.

¹⁹This 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.

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 $\{\hat{\beta}_{k,n_t}^t\}$. 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 $\hat{\beta}_{k,n_t}^t$ 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,

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

Rearranging this inequality yields

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

where the equality comes from the definition of $\hat{\beta}_{K_t+1,n_t}^t$ in [Equation \(2\)](#). Since $v_i > v_{i_{K_t+1}}$, this contradicts the efficiency property established by [Lemma 2](#). 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_{K_t} , whereas if she exits, she receives as her continuation payoff her marginal contribution to the social welfare. So, suppose that $i = i_k$ for some $k > K_t$. Then a deviation to remaining active in the auction is profitable if, and only if,

$$v_k - \hat{\beta}_{K_t+1,n_t}^t(\omega_t, v_{i_{K_t}}, \dots, v_{i_{k-1}}, v_{i_{k+1}}, \dots, v_{i_{n_t}}) \geq w^{t+1}(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\hat{\beta}_{K_t+1,n_t}^t(\omega_t, v_{i_{K_t}}, \dots, v_{i_{k-1}}, v_{i_{k+1}}, \dots, v_{i_{n_t}}) \leq v_k - w^{t+1}(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}) = \hat{\beta}_{k,n_t}^t(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\hat{\beta}_{k,n_t}^t$ in [Equation \(2\)](#). As above, the fact that $v_k < v_{i_{K_t}}$ contradicts the efficiency property established by [Lemma 2](#). 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 $\hat{\beta}_{k,n_t}^t$ 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 $\hat{\beta}_{k,n_i}^t$. 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 (2) 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

APPENDIX B. OPTIMAL MECHANISMS

In this appendix, we formally present the results relating to revenue-maximization. In particular, we derive a revenue-maximizing direct mechanism (the dynamic virtual pivot mechanism) and then show that an identical outcome may be achieved as an equilibrium of a sequence of ascending price auctions with an appropriately chosen reserve price.

B.1. Preliminaries

It is straightforward to see that, with full commitment power, the revelation principle applies directly in this setting. Therefore, we may restrict attention to direct mechanisms without loss of generality. Since buyers' private information is persistent, a dynamic direct mechanism asks each agent i to make a *single* report $v'_i \in \mathbf{V}$, upon arrival to the market, of her type v_i . A dynamic direct mechanism is then a collection

$$\mathcal{M} = \{x_{i,t}, p_{i,t}\}_{i \in \mathcal{I}, t \in \mathbb{N}},$$

where $x_{i,t}$ and $p_{i,t}$ are mappings from reported valuations and observed states to the probability of allocation to agent i in period t and her payment in that period, respectively. A mechanism is *feasible* if

- (1) for all $t \in \mathbb{N}$, $x_{i,t} = 0$ and $p_{i,t} = 0$ if $\alpha_t(i) = 0$;
- (2) $\sum_{i \in \mathcal{I}} x_{i,t} \leq K_t$ for all $t \in \mathbb{N}$;
- (3) $\sum_{t \in \mathbb{N}} x_{i,t} \leq 1$ for all $i \in \mathcal{I}$; and
- (4) $x_{i,t}$ and $p_{i,t}$ are adapted to the seller's information.

When a buyer $i \in \mathcal{I}_t$ arrives on the market in period t , we assume that she observes only the current state of the market ω_t ; that is, the set \mathcal{A}_t of agents present on the market and the number K_t of objects available. The agent does not observe the history of past 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, she chooses a report v'_i to maximize her expected payoff

$$U_i(v'_i, v_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} (x_{i,s}(v'_i, \mathbf{v}_{-i})v_i - p_{i,s}(v'_i, \mathbf{v}_{-i})) \middle| \omega_t \right], \quad (\text{B.1})$$

where the expectation is taken with respect to the arrival and departure processes of buyers and objects, as well as the reports of all other agents that may be present on the market. A mechanism $\mathcal{M} = \{x_{i,t}, p_{i,t}\}$ is incentive compatible if buyers prefer to report their type truthfully (that is, choose $v'_i = v_i$), and \mathcal{M} is individually rational if a buyer's expected utility is nonnegative.

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

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

where

$$q_i(v'_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} x_{i,s}(v'_i, \mathbf{v}_{-i}) \middle| \omega_t \right] \quad (\text{B.2})$$

is the expected discounted sum of object allocation probabilities, and

$$m_i(v'_i, \omega_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} p_{i,s}(v'_i, \mathbf{v}_{-i}) \middle| \omega_t \right] \quad (\text{B.3})$$

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—the incentive problem faced by a seller in this setting is simplified by reducing the problem to a single-dimensional allocation problem.

Standard arguments as found in Myerson (1981, Lemma 2) and Mas-Colell, Whinston, and Green (1995, Proposition 23.D.2) may be used to characterize incentive compatible allocation rules. Moreover, we arrive immediately at a standard integral utility representation and revenue-equivalence result. The following summarizes (without proof) these standard preliminaries:

LEMMA B.1 (Incentive Compatibility, Individual Rationality, and Revenue Equivalence).

Fix a feasible dynamic direct mechanism $\mathcal{M} = \{x_{i,t}, p_{i,t}\}$. \mathcal{M} is incentive compatible and individually rational if, and only if, for all buyers $i \in \mathcal{I}$ and all states ω_t ,

- (1) the induced expected discounted probability of allocation $q_i(v_i, \omega_t)$ is nondecreasing in v_i ; and
- (2) the induced expected discounted payment by the lowest-possible type $m_i(0, \omega_t)$ is nonpositive.

Moreover, the expected discounted payment of a buyer $i \in \mathcal{I}_t$ with value v_i in state ω_t is given by

$$m_i(v_i, \omega_t) = m_i(0, \omega_t) + q_i(v_i, \omega_t) v_i - \int_0^{v_i} q_i(v'_i, \omega_t) dv'_i. \quad (\text{B.4})$$

B.2. A Benchmark Optimal (Direct) Mechanism

We now make the additional assumptions, for ease of exposition, that the distribution of each buyer i 's values F_i has a strictly positive and continuous density f_i and that each buyer's virtual valuation function

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

is strictly increasing in v_i .²⁰

Consider a single monopolist seller who commits, at time zero, to a dynamic direct mechanism $\mathcal{M} = \{x_{i,t}, p_{i,t}\}$. The seller's expected revenue from this mechanism is the expected discounted sum of payments made by each buyer. Recall that the expected discounted payment of a buyer i , conditional on entry, is denoted by $m_i(v_i, \omega_t)$. We may use Lemma B.1 and integrate Equation (B.4) by parts to write the seller's expected profits as

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^{t-1} \alpha_t(i) m_i(0, \omega_t) \right] + \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}_t} \delta^{t-1} \alpha_t(i) q_i(v_i, \omega_t) \varphi_i(v_i) \right], \quad (\text{B.5})$$

where the expectation is taken with respect to the arrival and departure processes, as well as buyers' values.

²⁰If these assumptions are violated, one may use the procedure of Monteiro and Svaiter (2010) to generate an "ironed" generalized virtual value that may be used in place of φ_i whenever it is not-well defined (when the assumptions on the densities f_i are violated) or is decreasing.

The seller's objective function above is the sum of two terms: the first term is a discounted sum of expected payments by agents who report value zero, while the second is a discounted sum of weighted virtual values. Since individual rationality requires that $m_i(0, \omega_t) \leq 0$ for all i and all ω_t , the first term is maximized by setting $m_i(0, \omega_t) = 0$ for all $i \in \mathcal{I}$ and all ω_t .

Moreover, the second term is identical to the efficiency-oriented social planner's objective function in Equation (B.5), 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 equivalent to maximizing the *virtual surplus*. As was the case with the efficient allocation rule $\hat{\mathbf{x}}$, the perishability of objects combined with the persistence of values across time immediately yields the optimal allocation rule $\tilde{\mathbf{x}}$: in each period t , allocate all K_t objects to the buyers with the K_t highest (nonnegative) virtual values.²¹

It should be clear that the revenue-maximizing allocation rule $\tilde{\mathbf{x}}$ is incentive compatible. The optimal policy $\tilde{\mathbf{x}}$ allocates an object to i after a given history if, and only if, i 's virtual value is positive and among the highest virtual values of buyers present at that history. Since we have assumed the standard regularity condition that the virtual valuation function is strictly increasing, $\tilde{x}_{i,t}$ 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 induced expected discounted probability of receiving an object, denoted by \tilde{q}_i , 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.

In particular, the payment rule we propose is the *dynamic virtual pivot mechanism*, the dynamic pivot mechanism discussed above applied to virtual values. This parallels the relationship, in the static single-object setting, between the Myerson (1981) optimal auction and the Vickrey-Clarke-Groves mechanism. Just as the optimal (static) auction yields buyers their marginal contribution to the virtual surplus, the dynamic virtual pivot mechanism will provide each buyer with her *flow* marginal contribution to the virtual surplus.

So, for each $i \in \mathcal{I}$, we denote by

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

the lowest value that agent i can report and still potentially receive an object under the revenue-maximizing allocation policy. Furthermore, for all $i \in \mathcal{I}$, we define for any state $\omega_t = (\alpha_t, K_t)$ and reported values \mathbf{v}_t the function

$$\Pi^i(\omega_t, \mathbf{v}_t) := \mathbb{E} \left[\sum_{s=t}^{\infty} \sum_{j \in \mathcal{I}} \delta^{s-t} \tilde{x}_{j,s}(\omega_s, \mathbf{v}_s) \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \middle| \omega_t \right]. \quad (\text{B.6})$$

This expression is the same as the virtual surplus in the seller's objective function in Equation (B.5), except that it measures the virtual surplus in units determined by F_i , the distribution from which i 's value is drawn.²² Since we have assumed that φ_i is increasing for all buyers i , φ_i^{-1} is also increasing; therefore, transforming the virtual values of all agents by φ_i^{-1} preserves their ordering.

²¹Notice that, as in the classic intertemporal price discrimination problem of Stokey (1979), a revenue-maximizing monopolist will commit to *never* sell to buyers with negative virtual values.

²²Note that when all buyers are ex ante symmetric, $\tilde{r}_i = \tilde{r}_j$, $\varphi_i^{-1}(\varphi_j(v_j)) = v_j$, and $\Pi^i(\omega_t, \mathbf{v}_t) = \Pi^j(\omega_t, \mathbf{v}_t)$ for all $i, j \in \mathcal{I}$.

Moreover, $\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \geq 0$ if, and only if, $\varphi_j(v_j) \geq 0$. Therefore,

$$\tilde{\mathbf{x}} \in \arg \max_{\{x_{j,t}\}} \left\{ \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \sum_{j \in \mathcal{I}_t} \delta^{t-1} x_{j,t} \left(\varphi_i^{-1}(\varphi_j(v_j)) - \tilde{r}_i \right) \right] \right\};$$

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).

Recalling that ω_s^{-i} denotes that buyer i has been removed from the market in period s , we write

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

for the virtual surplus (in i 's “units”) when i is removed from the market. Thus, the presence of agent i on the market when the state is ω_t and reported values are given by \mathbf{v}_t yields an expected marginal contribution to the virtual surplus equal to

$$\pi_i(\omega_t, \mathbf{v}_t) := \Pi^i(\omega_t, \mathbf{v}_t) - \Pi_{-i}^i(\omega_t, \mathbf{v}_t).$$

Similarly, her *flow* marginal contribution to the virtual surplus—again, in i 's units—equal to

$$\tilde{\pi}_i(\omega_t, \mathbf{v}_t) := \pi_i(\omega_t, \mathbf{v}_t) - \delta \mathbb{E} [\pi_i(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t].$$

Notice that, as with the efficient policy, a buyer's flow contribution is zero if they are not allocated an object in the current period, and equals their entire marginal contribution if they are.

The dynamic virtual pivot mechanism is then simply defined as the mechanism $\tilde{\mathcal{M}} := \{\tilde{\mathbf{x}}, \tilde{\mathbf{p}}\}$, where the payment rule $\tilde{\mathbf{p}}$ is defined by

$$\begin{aligned} \tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) &:= \tilde{x}_{i,t}(\omega_t, \mathbf{v}_t) v_i - \tilde{\pi}_i(\omega_t, \mathbf{v}_t) \\ &= \tilde{x}_{i,t}(\omega_t, \mathbf{v}_t) (v_i - \pi_i(\omega_t, \mathbf{v}_t)) \end{aligned} \tag{B.7}$$

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

LEMMA B.2 (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. Since $\tilde{q}_i(v_i, \omega_t)$ is nondecreasing in v_i , the virtual dynamic pivot mechanism is incentive compatible. This is not sufficient, however, for the claim that $\tilde{\mathcal{M}}$ is *periodic ex post* incentive compatible. To show that this is indeed the case, fix an arbitrary agent $i \in \mathcal{I}_t$ for any $t \in \mathbb{N}$, and suppose that i knows the reported values \mathbf{v}_t^{-i} of all agents other than i who are also on the market at time t . Then, by reporting a value v'_i upon her arrival in state ω_t , agent i 's payoff is

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

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 reporting $v'_i = v_i$. 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.

Finally, recall that an incentive compatible mechanism is individually rational if $m_i(0, \omega_t) \leq 0$ for all i and all ω_t . However, the payments defined in Equation (B.7) are such that

$$\tilde{p}_{i,t}(\omega_t, (0, \mathbf{v}_t^{-i})) = -\tilde{x}_{i,t}(\omega_t, (0, \mathbf{v}_t^{-i}))\pi_i(\omega_t, (0, \mathbf{v}_t^{-i})) \leq 0,$$

since the optimal allocation rule only allocates to agents with nonnegative contributions to the virtual surplus, and hence $\pi_i(v_i, \mathbf{v}_t^{-i}) \geq 0$ for all v_i and \mathbf{v}_t^{-i} . Since m_i is the expected discounted sum of payments $p_{i,t}$, we have (as desired) $m_i(0, \omega_t) \leq 0$. \square

B.3. An Optimal Sequential Auction

We now make the additional assumptions that buyers are ex ante symmetric; that is, we assume that $F_i = F_j = F$ for all $i, j \in \mathcal{I}$. We will show that a sequence of ascending auctions with reserve price $\tilde{r} := \varphi^{-1}(0)$ is equivalent to the dynamic virtual pivot mechanism.²³

We again make use of the multi-unit, uniform-price variant of the Milgrom and Weber (1982) button auction, but now with a reserve price equal to \tilde{r} . For simplicity, we assume that the price clock starts at zero and rises continuously.²⁴ When there are $K_t \geq 1$ units available in a given period, the auction will end whenever there are at most K_t 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 $\mathbf{y}_t := (y_t^1, \dots, y_t^{n_t})$ denotes the ordered valuations of all buyers present in period t , where y_t^1 is the largest value and $y_t^{n_t}$ is the smallest. As before, we let

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

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

$$\tilde{\pi}_i^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) := \delta \mathbb{E} \left[\pi_i(\omega_{t+1}, \mathbf{v}_{t+1}) \mid \omega_t, \mathbf{v}_t = (v_i, \dots, v_i, \mathbf{y}_t^{>k}) \right].$$

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 $k - 1$ competitors with the same value v_i and $n_t - k$ buyers ranked below i with values $\mathbf{y}_t^{>k}$. Notice that this is exactly i 's expected contribution to the social welfare over a replacement agent with value \tilde{r} , conditional on being pivotal in the

²³As mentioned in Section 4, it is straightforward to show that a sequence of ascending auctions with asynchronous clocks is optimal when buyers are *not* ex ante symmetric.

²⁴One could equivalently model each auction as a two-stage game in which buyers make a participation decision for the current-period auction, followed by the price clock starting at \tilde{r} .

current period. Moreover, observe that $\tilde{\pi}_i^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) = 0$ for all buyers with values $v_i \leq \tilde{r}$, regardless of the realization of $\mathbf{y}_t^{>k}$.

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

$$\tilde{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) := v_i - \tilde{\pi}_i^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}). \quad (\text{B.8})$$

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 $\mathbf{y}_t^{>k}$ 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 is exactly $\tilde{\beta}_{k,n_t}^t$. We may then use arguments similar to those of Theorems 1 and 2 to prove the following result.

THEOREM B.1 (Revenue maximization via sequential ascending auctions).

Suppose that $F_i = F$ for all $i \in \mathcal{I}$ and that φ is strictly increasing. Bidding according to $\{\tilde{\beta}_{k,n_t}^t\}$ 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 of this theorem parallels the developments of Section 3.3. In particular, we will first show, as in Lemma 2, 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 $\tilde{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k})$ are increasing in v_i for all $k = 1, \dots, n_t$. Moreover, if $v_i > y_t^{k+1}$, then

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

PROOF OF CLAIM. Fix an arbitrary period $t \in \mathbb{N}$, 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 $n_t - k$ buyers have dropped out of the period- t auction, revealing values $\mathbf{y}_t^{>k}$, where $n_t := |\mathcal{A}_t|$ is the number of agents present, and $k \in \{1, \dots, n_t\}$. We wish to show first that $v_i > v'_i > y_t^{k+1}$ implies that

$$\tilde{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) := v_i - \tilde{\pi}_i^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) > v'_i - \tilde{\pi}_i^{t+1}(\omega_t, v'_i, \mathbf{y}_t^{>k}) =: \tilde{\beta}_{k,n_t}^t(\omega_t, v'_i, \mathbf{y}_t^{>k}).$$

Notice that

$$\begin{aligned} & \tilde{\pi}_i^{t+1}(\omega_t, v'_i, \mathbf{y}_t^{>k}) - \tilde{\pi}_i^{t+1}(\omega_t, v_i, \mathbf{y}_t^{>k}) \\ &= \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k}) \right] - \delta \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k}) \right] \\ & \quad - \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right] + \delta \mathbb{E} \left[\Pi_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right]. \end{aligned}$$

This, however, is equal to

$$\delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k}) \right] - \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right],$$

since the virtual surplus when i is removed from the market does not depend upon her value (and hence the second and fourth terms are equal). Moreover, by treating buyer i with value v'_i as though her true value were v_i , we can provide a bound on the difference above. In particular, we have

$$\begin{aligned} & \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v'_i, \mathbf{y}_t^{>k}) \right] - \delta \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right] \\ & \geq \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \tilde{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{n_t-k}, v_i, \mathbf{y}_t^{>k}) \right] (v'_i - v_i). \end{aligned}$$

Thus, if $v_i > v'_i$, then

$$\begin{aligned} & \tilde{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) - \tilde{\beta}_{k,n_t}^t(\omega_t, v_i, \mathbf{y}_t^{>k}) \\ & \geq (v_i - v'_i) \left(1 - \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \tilde{x}_{i,s}(\omega_s, \mathbf{v}_s) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, \mathbf{y}_t^{>k}) \right] \right) > 0 \end{aligned}$$

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

$$\begin{aligned} & \text{Additionally, note that if } v_i > v_j = y_t^{>k+1}, \text{ then } \tilde{\pi}_i^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>k+1}) - \tilde{\pi}_i^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>k+1}) = \\ & \delta \left(\mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{k+1}, v_j, \mathbf{y}_t^{>k+1}) \right] - \mathbb{E} \left[\Pi(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, v_j, \mathbf{y}_t^{>k+1}) \right] \right) \\ & - \delta \left(\mathbb{E} \left[\Pi_{-i}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^k, v_i, v_j, \mathbf{y}_t^{>k+1}) \right] - \mathbb{E} \left[\Pi_{-j}(\omega_{t+1}, \mathbf{v}_{t+1}) | \omega_t, \mathbf{v}_t = (\bar{\mathbf{v}}^{k+1}, v_j, \mathbf{y}_t^{>k+1}) \right] \right). \end{aligned}$$

However, the second difference above may be rewritten as

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

Thus, $\tilde{\pi}_i^{t+1}(\omega_t, v_j, \mathbf{y}_t^{>k+1}) - \tilde{\pi}_i^{t+1}(\omega_t, v_i, v_j, \mathbf{y}_t^{>k+1})$ 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 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

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

Moreover, by treating buyer j as though her true value were v_i instead of v_j (as we did above with v'_i and v_i), we can provide a bound on the difference above, which may be used to show that

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

Thus, the exit of the buyer with rank $(k+1)$ does not induce the immediate exit of any buyer with a higher value. Therefore, since $\tilde{\beta}_{k, n_t}^t(\omega_t, v_i, v_j, \mathbf{y}_t^{>k+1})$ is strictly increasing in v_i , the price at which this exit occurs fully reveals the value of the $(k+1)$ -st 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 (B.8) are fully revealing of the buyers' values. \square

CLAIM. *Following the bidding strategies $\tilde{\beta}_{k, n_t}^t$ 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 $t \in \mathbb{N}$, 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 $\tilde{\beta}_{k, n_t}^t$ are strictly increasing; therefore, the auction allocates the K_t objects to the group of buyers with the K_t highest values greater than the reserve. 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 direct mechanism, each buyer i with $v_i > \tilde{r}$ receives an object and pays

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

where π_i is the agent's marginal contribution to the virtual surplus. 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, $\pi_i(\omega_t, \mathbf{v}_t) = v_i - \tilde{r}$, implying that $\tilde{p}_{i,t}(\omega_t, \mathbf{v}_t) = \tilde{r}$. In this case, then, the direct and indirect mechanisms yield the same outcome.

Suppose instead that $K_t < n_t$; that is, there are more agents present than objects. Denote by i_k the bidder with the k -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+1} drops out of the auction, which is given by

$$\tilde{\beta}_{K_t+1, n_t}^t(\omega_t, v_{i_{K_t+1}}, \dots, v_{i_{n_t}}) = v_{i_{K_t+1}} - \tilde{\pi}_i^{t+1}(\omega_t, v_{i_{K_t+1}}, \dots, v_{i_{n_t}}).$$

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 direct mechanism, on the other hand, an agent i who receives an object pays

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

This may be rewritten as

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

where the final equality follows from the fact that $\tilde{\pi}_i^{t+1}(\omega_t, v_{i_{K_t+1}}, \dots, v_{i_{n_t}})$ is defined to be the expected future marginal contribution to the virtual surplus of the agent with the $(K_t + 1)$ -st highest value. Therefore, following the bidding strategies $\tilde{\beta}_{k,n_t}^t$ 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 holds after each history. Thus, the two mechanisms are outcome equivalent. \square

Finally, it remains to be seen that the bidding strategies in Equation (B.8) do, in fact, form an equilibrium. As in the case of the sequential ascending auction with no reserve, the bidding strategies $\tilde{\beta}_{k,n_t}^t$ 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 $\tilde{\beta}_{k,n_t}^t$.

CLAIM. *Suppose that in each period, buyers bid according to the cutoff strategies given in Equation (B.8). 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 $\{\tilde{\beta}_{k,n_t}^t\}$. 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 $\{\tilde{\beta}_{k,n_t}^t\}$ 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 above, 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 $\tilde{\pi}_i^{t+1}(\omega_t, v_i, v_{i_{K_t+2}}, \dots, v_{i_{n_t}})$, 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,

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

Rearranging this inequality yields

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

where the equality comes from the definition of $\tilde{\beta}_{K_t, n_t}^t$ in Equation (B.8). 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_{K_t} , whereas if she exits, she receives as her continuation payoff her marginal contribution to the virtual surplus. So, suppose that $i = i_k$ for some $k > K_t$. Then a deviation to remaining active in the auction is profitable if, and only if,

$$v_k - \tilde{\beta}_{K_t, n_t}^t(\omega_t, v_{i_{K_t}}, \dots, v_{i_{k-1}}, v_{i_{k+1}}, \dots, v_{i_{n_t}}) \geq \tilde{\pi}_i^{t+1}(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}).$$

Rearranging this inequality yields

$$\tilde{\beta}_{K_t, n_t}^t(\omega_t, v_{i_{K_t}}, \dots, v_{i_{k-1}}, v_{i_{k+1}}, \dots, v_{i_{n_t}}) \leq v_{i_k} - \tilde{\pi}_i^{t+1}(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}) = \tilde{\beta}_{k, n_t}^t(\omega_t, v_{i_k}, \dots, v_{i_{n_t}}),$$

where the equality comes from the definition of $\tilde{\beta}_{k, n_t}^t$ in Equation (B.8). As above, the fact that $v_{i_k} < v_{i_{K_t}}$ 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 $\tilde{\beta}_{k, n_t}^t$ 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. 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 (B.8) is optimal at all histories: 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

Thus, bidding in each period according to the strategy described in Equation (B.8) 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 M. SAID (2010): "Dynamic Auctions: A Survey," Cowles Foundation Discussion Paper 1757R, Yale University.
- BERGEMANN, D., AND J. VÄLIMÄKI (2003): "Dynamic Common Agency," *Journal of Economic Theory*, 111(1), 23–48.
- (2006): "Dynamic Price Competition," *Journal of Economic Theory*, 127(1), 232–263.
- (2010): "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771–789.
- BIKHCHANDANI, S., P. A. HAILE, AND J. G. RILEY (2002): "Symmetric Separating Equilibria in English Auctions," *Games and Economic Behavior*, 38(1), 19–27.
- BOARD, S., AND A. SKRZYPACZ (2010): "Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-Sales," Unpublished manuscript, University of California, Los Angeles.
- CAILLAUD, B., AND J. ROBERT (2005): "Implementation of the Revenue-Maximizing Auction by an Ignorant Seller," *Review of Economic Design*, 9(2), 127–143.
- CALZOLARI, G., AND A. PAVAN (2006): "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, 130(1), 168–204.
- (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 144(2), 503–531.
- 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.
- DOLAN, R. J. (1978): "Incentive Mechanisms for Priority Queuing Problems," *Bell Journal of Economics*, 9(2), 421–436.
- EDELMAN, B., AND M. SCHWARZ (2010): "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100(2), 597–602.
- 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 (2009): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, 1(2), 168–198.
- (2010a): "Efficient Sequential Assignment with Incomplete Information," *Games and Economic Behavior*, 68(1), 144–154.

- (2010b): “Incentive Compatible Optimal Stopping,” Unpublished manuscript, University of Bonn.
- HAILE, P. A. (2003): “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 108(1), 72–110.
- 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.
- JEHIEL, P., AND B. MOLDOVANU (2000): “Auctions with Downstream Interaction Among Buyers,” *RAND Journal of Economics*, 31(4), 768–791.
- (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69(5), 1237–1259.
- 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.
- 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.
- LI, A. (2009): “Selling Storable Goods to a Dynamic Population of Buyers: A Mechanism Design Approach,” Unpublished manuscript, Stanford University.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*. Oxford University Press, Oxford.
- 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. (2010): “Optimal Dynamic Mechanism Design with Deadlines,” 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.
- MONTEIRO, P. K., AND B. F. SVAITER (2010): “Optimal Auction with a General Distribution: Virtual Valuation without Densities,” *Journal of Mathematical Economics*, 46(1), 21–31.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6(1), 58–73.
- NALEBUFF, B. J., AND J. I. BULOW (1993): “Designing the PCS Auction,” Attached to *Comments of Bell Atlantic Personal Communications, Inc.*, PP Docket 93-253.
- PAI, M., AND R. VOHRA (2009): “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.
- (2010): "Infinite-Horizon Mechanism Design: The Independent-Shock Approach," 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.
- SAID, M. (2008): "Information Revelation and Random Entry in Sequential Ascending Auctions," in *Proceedings of the 9th ACM Conference on Electronic Commerce (EC '08)*, Chicago.
- (2010): "Sequential Auctions with Randomly Arriving Buyers," Unpublished manuscript, Microsoft Research New England.
- STOKEY, N. L. (1979): "Intertemporal Price Discrimination," *Quarterly Journal of Economics*, 93(3), 355–371.
- ÜNVER, M. U. (2010): "Dynamic Kidney Exchange," *Review of Economic Studies*, 77(1), 372–414.
- 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.