Available online at www.sciencedirect.com

SciVerse ScienceDirect

Journal of Economic Theory 147 (2012) 2419–2438

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Notes

Auctions with dynamic populations: Efficiency and revenue maximization ☆

Maher Said

Olin Business School, Washington University in St. Louis, United States

Received 12 January 2011; final version received 29 March 2012; accepted 1 April 2012

Available online 18 May 2012

Abstract

A seller has an uncertain number of perishable goods to sell in each period. Privately informed buyers arrive stochastically to the market. Buyers are risk neutral, patient, and have persistent private values for consuming a single unit. We show that the seller can implement the efficient allocation using a sequence of ascending auctions. The buyers use memoryless strategies to reveal all private information in every period, inducing symmetric behavior across different cohorts. We extend our results to revenue maximization, showing that a sequence of ascending auctions with asynchronous price clocks is an optimal mechanism.

JEL classification: C73; D44; D82; D83

Keywords: Dynamic mechanism design; Indirect mechanisms; Sequential ascending auctions; Sequential allocation; Random arrivals

☆ This paper is based on the third chapter of my dissertation at Yale University; it also supersedes and extends results first reported in Said [40]. I am especially indebted to Dirk Bergemann, Johannes Hörner, Ben Polak, and Larry Samuelson for their generous advice, comments, and encouragement. For many helpful discussions, I thank Rossella Argenziano, Alessandro Bonatti, Rahul Deb, Eduardo Faingold, James Fenske, Philip Haile, Sergei Izmalkov, Nenad Kos, David Parkes, Mallesh Pai, Kareen Rozen, Ilya Segal, Itai Sher, Juuso Välimäki, and Rakesh Vohra, as well as participants at various seminars and conferences. This work has also benefited greatly from the thoughtful comments of an associate editor and three anonymous referees. Financial support from the National Science Foundation, the Cowles Foundation, and Yale University is gratefully acknowledged, as is the generous hospitality of Microsoft Research New England.

E-mail address: said@wustl.edu.

http://dx.doi.org/10.1016/j.jet.2012.05.006

1. Introduction

In many "real-world" markets, not all buyers and sellers are available or present at the same time. Rather, agents arrive at different times, interact with various segments of the population, and then transact at different times. This fact, in conjunction with potential arrivals or departures in the future, leads to a trade-off: competition in the future may be more or less stiff than in the present, and opportunities to trade may be more or less frequent. Thus, agents must choose between transacting in the present or waiting until the (uncertain) future. In addition to this dynamic trade-off, competition between privately informed agents *across* time introduces additional dimensions for strategic behavior: when buyers and sellers face the same competitors repeatedly, the information transmitted by market institutions can have dramatic impacts on behavior and outcomes. Market participants may benefit from learning the private information of others or from manipulating others' beliefs about themselves. With these considerations in mind, it is clear that population dynamics can greatly influence competition, price determination, efficiency, and revenue in markets. It is therefore natural to question how this influence manifests, and what market structures are best able to achieve desirable outcomes in its presence.

We approach these questions in the present work with a model that features a single seller who maximizes either welfare or revenue in a dynamic market. At each point in time, a random number of buyers and objects arrive to the market. Buyers are risk neutral, patient, and have persistent private values for consuming a single unit of the seller's good. Objects, on the other hand, are perishable and can only be sold at the time of their arrival. While these modeling choices are not sufficiently general to encompass all dynamic environments of interest, there are still many settings for which our assumptions are satisfied.[1] For instance, many dynamic scheduling problems fit into our framework. Examples of such problems include the queuing of newly submitted tasks to a central server or supercomputing facility; the scheduling of departure times from a small airfield serving private jets and other charter aircraft (as opposed to large airlines with regularly scheduled flights); and the allocation of excess factory time or assembly lines in a manufacturing facility or foundry to new production requests.[2] "Last-minute" sales of time-sensitive goods (such as seats at professional sporting events, tickets to an opera or symphony, and cabins on a train or cruise ship) also fit into our framework and can be analyzed using our model.

It is well known that dynamic variants of the Vickrey–Clarke–Groves mechanism are efficient (and easily adapted for revenue maximization) in settings such as our own. These direct revelation mechanisms require buyers to report their values to the mechanism upon arrival to the market. In practice, however, direct revelation mechanisms may be impractical. 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 a dominant strategy. Despite this theoretical desirability, Ausubel [3] 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 critique by Nalebuff and Bulow [35]: when even reasonably sophisticated bidders "do not understand the payment rules of the [Vickrey] auction then we do not have any confidence that the end result will be efficient," thereby undermining the direct mechanism's desirability.

[1] Moreover, these choices may be a reasonable first-order approximation to reality. See Backus and Lewis [4] and Hendricks, Onur, and Wiseman [24] for recent uses of similar assumptions in empirical work.

[2] This latter setting arises naturally in the market for cloud computing services. Amazon uses a bidding system to allocate unused capacity, and SpotCloud is a new entrant that serves as a clearinghouse for excess computing capacity.

These criticisms are corroborated by experimental evidence. Kagel, Harstad, and Levin [26] examine single-unit auctions with affiliated private values, finding that theoretical predictions about bidding behavior are significantly more accurate in ascending auctions than in second-price auctions. Kagel and Levin [27] find a similar result in multi-object auctions with independent private values: ascending auctions are significantly more efficient than the dominance-solvable Vickrey auction. In another study examining the efficiency properties of several mechanisms in a resource allocation problem similar to our own, Banks, Ledyard, and Porter [5] conclude 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 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. In a setting *without* the entry of new buyers, Kittsteiner, Nikutta, and Winter [28] have shown that these goals are attainable using a sequence of second-price sealed-bid auctions. This is not the case, however, in settings such as our own where new buyers arrive over time: population dynamics lead to interdependence and asymmetry among buyers, which can jointly preclude the second-price sealed-bid auction from generating efficient (or revenue-maximizing) outcomes.

In contrast, the ascending auction is a natural auction format that, by allowing for the gradual revelation of buyers’ private information, avoids some of the problems of the second-price sealed-bid auction. We construct an equilibrium bidding strategy profile for buyers in a sequence of ascending auctions that generates an efficient outcome. 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. This yields prices and allocations identical to the truthful equilibrium of the efficient direct mechanism. Similar arguments show that a sequence of ascending auctions with asynchronous price clocks admits an equilibrium with the same prices and allocations as the optimal direct mechanism. (When buyers are ex ante symmetric, an optimally chosen reserve price may be used instead of asynchronous clocks.) Thus, the sequential ascending auction is a natural decentralized institution for achieving either efficient or optimal outcomes.

The present work contributes to the recent literature on dynamic mechanism design.[3] Athey and Segal [1,2] and Bergemann and Välimäki [9] develop efficient mechanisms for general dynamic environments, while Pavan, Segal, and Toikka [39] characterize necessary and sufficient conditions for dynamic incentive compatibility. While these papers study direct revelation mechanisms, we are primarily focused on the design of *indirect* mechanisms. This focus highlights the role of transparency and information revelation in dynamic markets with repeated competition. Our results are in line with the findings of Calzolari and Pavan [15,16], who show that mechanisms in sequential contracting environments must control for the information they reveal as well as the allocations they induce.[4]

This paper also relates to recent work on dynamic auctions and revenue management with a fixed supply of nonperishable goods. Board and Skrzypacz [11] and Li [30] design opti-

[3] See Bergemann and Said [6] for a survey of much of this literature.

[4] This is related to results for auctions with downstream interaction among buyers, as in Jehiel and Moldovanu [25], or those with resale opportunities, as in Haile [23]. See also Pancs [37,38], who studies transparency and information disclosure in settings where agents can also take nonenforceable actions.

mal and efficient mechanisms, respectively, in finite-horizon versions of our model; Li finds (as we do) that an open auction format is required for indirect implementation. Garrett [19], Mierendorff [32], and Pai and Vohra [36] derive optimal mechanisms in models with unobservable buyer arrivals and departures. Similar issues are explored by Gershkov and Moldovanu [20,22] and Vulcano, van Ryzin, and Maglaras [42] in markets with impatient buyers.

Finally, our focus on indirect mechanisms ties into the literature initiated by Milgrom and Weber [33], who examine the properties of several sequential auction formats. Kittsteiner, Nikutta, and Winter [28] introduce discounting to that framework; they assume, however, that buyers are ex ante homogeneous and are all present in the initial period. Budish and Zeithammer [12] examine auction sequencing and the timing of information revelation, while Said [41] and Zeithammer [43] study the impact of population dynamics in sequences of second-price auctions. Meanwhile, Lavi and Nisan [29] and Compte, Lavi, and Segev [17] provide "worst-case" lower bounds on the efficiency of sequential ascending auctions with dynamic populations.

2. Model

We consider an infinite-horizon, discrete-time environment with a single seller. In each period $t \in \mathbb{N}$, the seller has K_t units of a homogeneous 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. We assume that objects are nonstorable: any objects that are not allocated or sold "expire" at the end of each period, and hence cannot be carried over to future periods.

Each period t begins with the arrival 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 vary, with N_t an independent draw from the distribution λ_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 independently drawn from the distribution F_i on $[0, \bar{v}]$.

In addition, buyers may exogenously depart from the market (and never return) after each period, where the (common) probability of any buyer i "surviving" from period t to $t+1$ is denoted by $\gamma_t \in [0, 1]$. Otherwise, buyers remain present on the market until they obtain an object. Finally, we assume that buyers are risk neutral, with quasilinear and time-separable preferences. All buyers, as well as the seller, discount the future with the common discount factor $\delta \in (0, 1)$.

The buyers' arrivals and departures generate a stochastic process $\{\alpha_t\}_{t\in\mathbb{N}}$, where $\alpha_t : \mathcal{I} \to \{0, 1\}$ is an indicator function that tracks the presence of each buyer on the market at time t. We assume that buyers cannot conceal their presence, and so α_t is common knowledge among all agents present at time t. We also assume that the arrival of objects is publicly observed, and so K_t is also commonly known by all agents present at time t. It will be convenient to denote the "state" of the market at the beginning of each period t by $\omega_t := (\alpha_t, K_t)$. Once the current state is known, the seller allocates objects to buyers and makes transfers, and we move on to the following period.

3. Efficiency

We begin with the problem faced by a benevolent social planner whose goal is to maximize allocative efficiency; that is, the planner wishes to choose a feasible allocation rule $\{x_{i,t}\}_{i\in\mathcal{I},t\in\mathbb{N}}$ to

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

where $x_{i,t}$ is the probability of allocating to agent i in period t, and the expectation above is taken with respect to the arrival and departure processes, as well as agents' valuations.[5] Since objects are perishable, there is no potential benefit to be gained by "withholding" an object or not allocating as many objects as possible. In addition, buyers' values are persistent and are independent of the arrival and departure processes. Therefore, since all buyers share the same discount factor and survive to future periods with the same probability, the cost of delaying allocation to any given buyer is a symmetric and strictly increasing function of her value.

Thus, the efficient allocation rule simply allocates all K_t objects in period t to the K_t highest-valued buyers currently present, breaking ties 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 (α_t), and these agents' reported values (denoted by $\mathbf{v}_t$). Thus, we may write the efficient allocation rule as $\widehat{\mathbf{x}} := \{\widehat{x}_{i,t}(\omega_t, \mathbf{v}_t)\}$.

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] We will focus on the *dynamic pivot mechanism* of Bergemann and Välimäki [9]. In order to fully describe this mechanism in our setting (which will help facilitate the development and interpretation of our subsequent results), we let, for any state $\omega_t = (\alpha_t, K_t)$ and reported values $\mathbf{v}_t$,

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

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

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

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

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

The dynamic pivot mechanism is then simply the mechanism combining the efficient allocation rule $\widehat{\mathbf{x}}$ with the payment rule $\widehat{\mathbf{p}} := \{\widehat{p}_{i,t}\}_{i \in \mathcal{I}, t \in \mathbb{N}}$, where for all i, all t, and all $(\omega_t, \mathbf{v}_t)$,

$$\widehat{p}_{i,t}(\omega_t, \mathbf{v}_t) := \widehat{x}_{i,t}(\omega_t, \mathbf{v}_t) v_i - \Big(w_i(\omega_t, \mathbf{v}_t) - \delta \mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1}) \,\big|\, \omega_t, \mathbf{v}_t\big]\Big). \tag{1}$$

This mechanism yields each buyer a flow payoff equal to her flow marginal contribution to the social welfare, the difference between her marginal contribution to the social welfare and her expected *future* marginal contribution to the social welfare. As the payments $\widehat{p}_{i,t}$ are *not*

[5] An allocation rule $\{x_{i,t}\}$ is *feasible* if, for all $t \in \mathbb{N}$ and all $i \in \mathcal{I}$: (1) $x_{i,t} = 0$ if $\alpha_t(i) = 0$, so only buyers physically present on the market receive an object; (2) $\sum_{t \in \mathbb{N}} x_{i,t} \leqslant 1$, so no buyer receives more than the single unit she desires; (3) $\sum_{i \in \mathcal{I}} x_{i,t} \leqslant K_t$, so no more objects are allocated than are actually available; and (4) $\{x_{i,t}\}$ is adapted to the seller's information at time t, so the seller is not "prescient."

[6] Dolan [18] was the first to propose the use of VCG-like mechanisms in settings with a dynamic population of buyers.

distribution-free, the dynamic pivot mechanism is not dominant strategy incentive compatible. It *is*, however, periodic ex post incentive compatible and individually rational.[7]

Lemma 1. *(See Bergemann and Välimäki [9, Theorem 1].) The dynamic pivot mechanism is periodic ex post incentive compatible and individually rational, and hence implements the efficient policy.*

Of course, the dynamic pivot mechanism is a direct revelation mechanism, relying on the planner to aggregate buyers' reports to determine allocations and payments. In light of the concerns discussed in the introduction, a natural question arises: does this mechanism correspond to a familiar indirect mechanism? A reasonable conjecture is that a sequence of auctions may suffice. But what auction format would be desirable? In a static world, the second-price sealed-bid auction is efficient; Kittsteiner, Nikutta, and Winter [28] prove that a sequence of such auctions is also efficient in a dynamic setting without entry. With a dynamic population of buyers, however, this is no longer the case—we present an example in which the efficient allocation rule is not implementable by a sequence of second-price sealed-bid auctions where bids are never revealed. We then show that the efficient allocation *is* implementable by a sequence of ascending auctions.

3.1. An inefficient sequential auction

Suppose a single object is available in each of three periods, after which no further objects arrive. Two buyers are present in the initial period and two new entrants arrive in the second period. No further entry occurs in the third period, and buyers exit the market only upon winning an object. All buyers' values are independently drawn from the uniform distribution on [0, 1]. A second-price sealed-bid auction is conducted in each period, and no information about bids is revealed.

Since no objects arrive after the third period, all buyers, regardless of "vintage," have zero continuation value. Thus, standard dominance arguments imply that each buyer will bid her own value, yielding an efficient outcome in the final period regardless of buyers' beliefs.

Now consider the second period, and let $f^i_{1,2}(z_1, z_2)$ denote the joint density of z_1 and z_2, the highest- and second-highest values of bidder i's competitors. Suppose these competitors follow the strictly increasing strategy $\beta(\cdot)$. Then the payoff to buyer i (with value v_i) who bids $\beta(v')$ is

$$\int_0^{v'}\int_0^{z_1}\big(v_i-\beta(z_1)\big)f^i_{1,2}(z_1,z_2)\,\mathrm{d}z_2\,\mathrm{d}z_1+\delta\int_{v'}^{1}\int_0^{v_i}(v_i-z_2)f^i_{1,2}(z_1,z_2)\,\mathrm{d}z_2\,\mathrm{d}z_1.$$

In an efficient (and hence necessarily symmetric) equilibrium, i's payoff is maximized when she bids as though her value is $v'=v_i$. The resulting (necessary) first-order condition implies that

$$\beta(v_i)=(1-\delta)v_i+\delta\frac{\int_0^{v_i} z_2 f^i_{1,2}(v_i,z_2)\,\mathrm{d}z_2}{\int_0^{v_i} f^i_{1,2}(v_i,z_2)\,\mathrm{d}z_2}. \tag{2}$$

[7] A mechanism is *periodic ex post incentive compatible* if, for all $t\in\mathbb{N}$ and all buyers present in period t, truthful reporting is a best response when all other buyers report truthfully, regardless of the history and the private information of all other agents that are, or have been, present on the market. A mechanism is *periodic ex post individually rational* if, after every history, each buyer's equilibrium payoff is nonnegative.

From the perspective of the "incumbent" buyer (denoted by I), the two entrants' values are independent draws from the uniform distribution on [0, 1], and so

$$f^I_{1,2}(z_1, z_2) = \begin{cases} 2 & \text{if } z_1 \geqslant z_2, \\ 0 & \text{otherwise.} \end{cases}$$

Now consider either of the new entrants (denoted by E). She knows that the incumbent lost in the first period, and—in an efficient equilibrium—that this *must* be because the incumbent's value was the *lower* of two independent draws from a uniform distribution on [0, 1]. Meanwhile, the other entrant's value is uniformly distributed on [0, 1]. Thus, each entrant perceives her competitors' values as independent draws from different distributions, and her beliefs are given by

$$f^E_{1,2}(z_1, z_2) = \begin{cases} 2(2 - z_1 - z_2) & \text{if } z_1 \geqslant z_2, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the optimal bids from Eq. (2) for the incumbent and entrants, respectively, are

$$\beta^I(v_i) = (1-\delta)v_i + \delta\frac{v_i}{2} \quad \text{and} \quad \beta^E(v_i) = (1-\delta)v_i + \delta\frac{(6-5v_i)v_i}{12-9v_i}.$$

Clearly, the incumbent and entrant bid functions do not coincide. Thus, there is no equilibrium with symmetric and strictly increasing bids in the second period, implying that a sequence of second-price sealed-bid auctions with no bid disclosure does not admit an efficient equilibrium.[8]

The inefficiency in this example stems from the asymmetry in beliefs between the incumbent and entrants. Note that buyers shade their bids downward in a trade-off between winning immediately and potentially winning in the future. The value of future participation, however, depends on future prices, and hence competitors' values. Thus, even in an independent private values environment, market dynamics and repeated competition *generate* interdependence. Moreover, buyers arriving to the market at different times will have different beliefs about each other stemming from their different market experiences. The entry of new bidders thus implies that buyers participate in an asymmetric second-price auction with interdependent values. As first noted by Maskin [31], efficient equilibria need not exist in such a setting.

Efficiency can be restored, however, by replacing the sealed-bid auction with its ascending counterpart. 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 prices at which their competitors exit the auction and thereby make inferences about their valuations. Since all buyers observe the same exits, they make the same inferences. Thus, the ascending auction "symmetrizes" buyers' beliefs about their lower-valued competitors and permits an efficient equilibrium.

To see this concretely, consider the second period in the example above and suppose that, instead of a second-price sealed-bid auction, the good is sold using a standard "button" ascending auction. Suppose further that buyers (rationally) anticipate truthful bidding in the third period. Therefore, continuing to denote the second-highest value of an arbitrary buyer's competitors by z_2, buyer i's expected payoff in the next period if she loses in the current period is given by

$$V(v_i, z_2) = \begin{cases} 0 & \text{if } v_i \leqslant z_2, \\ v_i - z_2 & \text{if } v_i > z_2. \end{cases}$$

[8] The logic of this example extends to less opaque bid disclosure policies. The incumbent's beliefs about future competition (and hence her bid) are unaffected by the disclosure policy. The entrants' beliefs *are* responsive to additional information, but since they derive from a different prior distribution, asymmetry and inefficiency persist.

Suppose that, as long as all three buyers are active in the second-period auction, they each continue to be active until the price reaches

$$\beta_3(v_i) = v_i - \delta V(v_i, v_i) = v_i.$$

Note that β_3 is strictly increasing, implying that the first bidder to drop out of the auction has the lowest value for the object. Thus, each of the two remaining bidders learns the value z_2 of her weakest competitor and (since her own value is greater than z_2) that she would be certain to win in the next period. So, let the two remaining bidders stay active until the price reaches

$$\beta_2(v_i, z_2) = v_i - \delta V(v_i, z_2) = (1-\delta)v_i + \delta z_2.$$

Since β_2 is strictly increasing in v_i, the highest-valued bidder will in fact be the winner. It is easy to show that, given expectations about future play, bidding according to β_3 and β_2 is an ex post equilibrium in the second period that is independent of players' beliefs about each other (see Theorem 2 below for the more general statement of this result). Notice that the winner's payoff is

$$v_i - \beta_2(z_1, z_2) = (v_i - z_1) + \delta(z_1 - z_2),$$

which is exactly her marginal contribution to the social welfare. Finally, it is straightforward to construct strategies for the first period that, with the second- and third-period strategies above, yield an efficient equilibrium. Thus, by revealing buyers' private information to their competitors, the ascending auction symmetrizes beliefs across cohorts and generates an efficient outcome.

Note, however, that information revelation alone is not sufficient for efficiency: the *timing* of information revelation is crucial. Consider, for instance, a sequence of second-price auctions where, after each auction closes, all bids are revealed. Cai, Wurman, and Chao [13] have shown (in a setting with ex ante symmetric buyers and no entry) that this alternative mechanism is inefficient. When information revelation occurs only after each auction ends, buyers have an incentive to manipulate opponents' beliefs: by pretending that she has a low value in the first period, a buyer induces her opponents to decrease their bids in future periods. She can then increase her own bid in a future period and, since the deviation is unobserved until after bidding concludes, win an object at a relatively low price.[9] In an ascending auction, however, such an attempt to capitalize on weakened bidding is immediately detected by the "manipulated" buyers, who can immediately respond and adjust their current-period bids.

3.2. An efficient sequential auction

Returning to the more general environment of Section 2, we now show that a sequence of ascending auctions implements the efficient allocation rule in a periodic ex post incentive compatible and individually rational equilibrium. We employ a sequence of multi-unit, uniform-price "button" auctions. An auction begins, in each period, with the price at zero and all agents present participating in the auction. Each bidder may drop out of the auction at any price. Exits are irreversible (within the current-period auction) and observable by all other bidders. Thus, the current price and set of active bidders are commonly known throughout the auction.

When there are $K_t \geqslant 1$ objects for sale, the auction ends whenever at most K_t active bidders remain, each of whom receives an object and pays the auction closing price. If there are fewer

[9] A similar effect is noted by Zhang and Wang [44] in their analysis of sequential all-pay auctions (representing contests) with full revelation of bids between rounds, again leading to the nonexistence of an efficient equilibrium.

than K_t bidders, the auction ends immediately at a price of zero. The auction also ends if several bidders drop out of the auction simultaneously, leaving $k < K_t$ bidders active. These bidders—along with $K_t - k$ randomly chosen "tied" bidders—receive an object at the auction closing price.

The information revelation in an ascending auction may lead to an additional asymmetry not discussed in the example above. In an efficient equilibrium, buyers that lose in period t will have inferred each others' values from exit prices: in period $t+1$, they are perfectly informed about one another. But a new group of buyers, about whom nothing is known, arrives in period $t+1$. We therefore have differential information about two groups of buyers, and (in order to achieve an efficient outcome) the new entrants must be induced to also reveal their private information.

This asymmetry is resolved by the 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 act "as though" they are uninformed. By doing so, buyers are able to *behave* symmetrically. This provides all buyers, incumbents and new entrants alike, with the appropriate incentives to gradually reveal their private information.

Note that the equilibrium we construct in memoryless strategies is *not* the result of an a priori restriction on buyers' strategies. Buyers have perfect recall of the past and are free to condition on information revealed in previous periods. Ignoring that information, however, is a fully rational and unconstrained best response to their competitors' behavior.

Why then would a fully rational buyer ignore payoff-relevant information and behave as though she were uninformed? Recall that bidders engage in the process of information revelation in *every* period. Therefore, there is no need for buyers to condition their behavior on past observations—in equilibrium, any payoff-relevant information will be revealed anew over the course of the current auction, allowing buyers to condition on this information as it is revealed (or re-revealed) *during* the current period. Thus, as in the efficient ex post equilibria of a static ascending auction, beliefs about one's current competitors are irrelevant.

To formally describe the strategies, let $n_t := \sum_{j \in \mathcal{I}} \alpha_t(j)$ denote the number of buyers present at time t. For each $k = 1, \dots, n_t$, denote by y_t^k the k-th highest value among those n_t buyers, and by

$$\mathbf{y}_t^{>k} := \left(y_t^{k+1}, \dots, y_t^{n_t}\right)$$

the collection of the $(n_t - k)$ lowest values. Then in each period t, every buyer i bids up to

$$\widehat{\beta}_t^k\left(\omega_t, v_i, \mathbf{y}_t^{>k}\right) := v_i - \delta \mathbb{E}\left[w_i(\omega_{t+1}, \mathbf{v}_{t+1}) \middle| \omega_t, \mathbf{v}_t = \left(v_i, \dots, v_i, \mathbf{y}_t^{>k}\right)\right] \tag{3}$$

whenever she is one of k bidders still active. This cutoff is the price at which buyer i is indifferent between winning the object in the current period and receiving her expected future marginal contribution to the social welfare in the next period when the currently inactive buyers' values are given by $\mathbf{y}_t^{>k}$. In addition, this expectation is conditioned on the event that buyer i is pivotal; that is, on the event that she is exactly tied with the $k-1$ other active bidders. When all buyers use these cutoffs, allocations and prices are the same as those produced by truth-telling in the dynamic pivot mechanism.[10]

Theorem 1. *The allocations and transfers generated by the bidding strategies $\{\widehat{\beta}_t^k\}$ in the sequential ascending auction are identical to those of the truthful equilibrium of the dynamic pivot mechanism.*

[10] Marginal contributions and externalities were first used to construct efficient bidding equilibria by Bergemann and Välimäki [7,8] in complete-information first-price auctions.

The intuition behind this result is straightforward. Each of the bid functions $\widehat{\beta}_t^k(\omega_t, v_i, \mathbf{y}_t^{>k})$ is strictly increasing in v_i. Thus, when there are k active buyers in the period-t auction, each bidding up to $\widehat{\beta}_t^k(\omega_t, v_i, \mathbf{y}_t^{>k})$, the next buyer to exit the auction will be the one with the k-th lowest value. This exit will publicly reveal her value, making y_t^k common knowledge among all buyers. The $k-1$ remaining bidders then bid up to $\widehat{\beta}_t^k(\omega_t, v_i, \mathbf{y}_t^{>k-1})$, and so the next bidder to drop out will have the $(k-1)$-th lowest remaining value. This chain of reasoning leads to the conclusion that the bid functions in Eq. (3) yield an efficient outcome.

To see why the prices induced by these strategies equal those of the dynamic pivot mechanism, consider the externality imposed by a buyer i who receives an object. If i is removed from the market, her object will instead be allocated to j, the buyer with (K_t+1)-th highest value. This yields an immediate gain in social surplus equal to j's value v_j. However, j will no longer be available for allocations in future periods, implying that her *future* marginal contribution to the social welfare is lost. This difference, however, is exactly j's bid in the ascending auction. Therefore, 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.

It remains to be shown that the proposed strategy profile is, in fact, an equilibrium. We consider perfect Bayesian equilibrium, thereby requiring that behavior is sequentially rational with respect to agents' beliefs, and that these beliefs are updated in accordance with Bayes' rule wherever possible. Since all buyers use the strictly increasing bidding strategies $\{\widehat{\beta}_t^k\}$, behavior along the equilibrium path is perfectly separating and Bayesian updating fully determines beliefs. To determine optimality *off* the equilibrium path, however, we need to consider the beliefs of bidders after a deviation. As such histories occur with zero probability, we are free to choose arbitrary off-equilibrium beliefs. We will therefore suppose that, after a deviation, buyers disregard their previous observations and believe that the deviating agent is *currently* acting in "sincere" accordance with $\{\widehat{\beta}_t^k\}$.

That the proposed strategies and beliefs form a perfect Bayesian equilibrium follows from two key observations. First, the use of memoryless strategies by a buyer's opponents negates any incentive to manipulate competitors' beliefs, as future behavior is independent of information revealed in the present. Second, buyers' marginal contributions to the social welfare are sufficiently well-behaved to guarantee that the induced interdependence from downstream interaction still satisfies standard single-crossing conditions, thereby admitting an efficient equilibrium in each period's ascending auction. Thus, bidding in accordance with Eq. (3) is an equilibrium of the sequential ascending auction mechanism.

Theorem 2. *Suppose that in each period, buyers bid according to $\{\widehat{\beta}_t^k\}$ in Eq. (3). This strategy profile, combined with the system of beliefs described above, forms a (periodic ex post) perfect Bayesian equilibrium of the sequential ascending auction mechanism.*

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.[11] The sequential ascending auction is therefore a natural, intuitive institution that yields efficient outcomes.

[11] Analogous to the multiplicity of symmetric separating equilibria described by Bikhchandani, Haile, and Riley [10], there exists a continuum of efficient equilibria in our setting. The equilibrium we construct is that in which bids are maximal; final transaction *prices* for winning buyers, however, are unique across all efficient equilibria.

It is interesting to note some additional properties of this equilibrium. First, notice that the extensive-form structure of the indirect mechanism allows for a much larger number of potential deviations from truthful behavior relative to its direct counterpart, as agents essentially "re-report" their values in each auction. Despite this, there is no loss in the "strength" of implementation: the efficient equilibria in both the direct and indirect mechanisms are periodic ex post equilibria. Thus, the strong "no regret" property of the ascending auction in static settings is inherited in the sequential setting—no buyer has an incentive to deviate from $\{\widehat{\beta}_t^k\}$, even when conditioning on her current-period competitors' values.

Furthermore, observe that in the sequential ascending auction, buyers are free to drop out immediately once an auction begins. But in the equilibrium we construct, buyers do not wish to do so, even if they were to know their opponents' values. So although we have assumed that buyers' arrivals are publicly observable, the proposed strategy profile remains an equilibrium in the "larger" game with unobservable arrivals. The independence of buyer arrivals is crucial here, however: Gershkov and Moldovanu [21] have shown that if buyer arrivals are both unobservable and correlated, then expected-externality mechanisms (like the dynamic pivot mechanism) need not be efficient, and subsidies to "losing" buyers may be required to account for the informational externality arising from their arrival.

4. Revenue maximization

We now turn to the case of a revenue-maximizing monopolist with full commitment power. The seller now wants to choose a feasible allocation rule $\{x_{i,t}\}$ to maximize

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

where the expectation is taken with respect to buyer arrivals, departures, and values. We impose the additional assumption that, for all $i\in\mathcal{I}$, the distribution F_i has a strictly positive and continuous density f_i. In addition, we assume that, for all $i\in\mathcal{I}$, i's "virtual value," defined by

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

is strictly increasing in v_i. Standard techniques may be used to rewrite the seller's expected revenue in any incentive compatible direct mechanism as

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

that is, the seller maximizes revenue by maximizing the virtual surplus. Recall that objects are perishable, and that buyers' values are persistent and independent of the arrival and departure processes. Moreover, all buyers have the same discount and departure rates. These observations immediately imply that the optimal allocation rule $\widetilde{\mathbf{x}} := \{\widetilde{x}_{i,t}(\omega_t, \mathbf{v}_t)\}$ allocates all K_t objects in each period t to the buyers with the K_t highest (nonnegative) virtual values.

This optimal allocation rule can be implemented by the *dynamic virtual pivot mechanism*, the dynamic pivot mechanism described above applied to virtual values. Just as the optimal (static) auction yields buyers their marginal contribution to the virtual surplus, the dynamic virtual pivot mechanism provides each buyer with her flow marginal contribution to the virtual surplus.

So, for each $i \in \mathcal{I}$, we define for all states $\omega_t = (\alpha_t, K_t)$ and reported values $\mathbf{v}_t$ the function

$$\Pi^i(\omega_t, \mathbf{v}_t) := \mathbb{E}\left[\sum_{s=t}^{\infty} \sum_{j\in\mathcal{I}} \delta^{s-t} \widetilde{x}_{j,s}(\omega_s, \mathbf{v}_s)\big(\varphi_i^{-1}\big(\varphi_j(v_j)\big) - \varphi_i^{-1}(0)\big)\Big|\omega_t, \mathbf{v}_t\right].$$

This is the virtual surplus generated by the optimal allocation rule $\widetilde{\mathbf{x}}$, but denominated in units determined by F_i, the distribution of i's values. Similarly,

$$\Pi^i_{-i}(\omega_t, \mathbf{v}_t) := \mathbb{E}\left[\sum_{s=t}^{\infty} \sum_{j\in\mathcal{I}\setminus\{i\}} \delta^{s-t} \widetilde{x}_{j,s}\big(\omega_s^{-i}, \mathbf{v}_s\big)\big(\varphi_i^{-1}\big(\varphi_j(v_j)\big) - \varphi_i^{-1}(0)\big)\Big|\omega_t^{-i}, \mathbf{v}_t\right]$$

is the virtual surplus (in i's "units") when i is removed from the market. Then buyer i's marginal contribution to the virtual surplus when the state is ω_t and reported values are given by $\mathbf{v}_t$ is

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

The dynamic virtual pivot mechanism is then simply the mechanism combining the optimal allocation rule $\widetilde{\mathbf{x}}$ with the payment rule $\widetilde{\mathbf{p}} := \{\widehat{p}_{i,t}\}_{i\in\mathcal{I}, t\in\mathbb{N}}$, where for all i, all t, and all $(\omega_t, \mathbf{v}_t)$,

$$\widetilde{p}_{i,t}(\omega_t, \mathbf{v}_t) := \widetilde{x}_{i,t}(\omega_t, \mathbf{v}_t)v_i - \big(\pi_i(\omega_t, \mathbf{v}_t) - \delta\mathbb{E}\big[\pi_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t\big]\big)$$

for all i and all $(\omega_t, \mathbf{v}_t)$. By providing each buyer flow payoffs equal to her flow contribution to the virtual surplus, this mechanism inherits the desirable properties of the dynamic pivot mechanism.

Lemma 2. *Suppose that virtual values φ_i are increasing for all i. Then the dynamic virtual pivot mechanism is periodic ex post incentive compatible and individually rational, and hence maximizes revenue.*

Recall that in a static single-object setting, Myerson [34] demonstrated that revenues are maximized by using a virtual analogue of VCG. Moreover, just as the VCG mechanism may be replicated by using an ascending auction, Myerson's optimal mechanism can be replicated by adding a reserve price to that auction. In the case where bidders are ex ante heterogeneous, the optimal ascending auction uses asynchronous price clocks—see Caillaud and Robert [14, Proposition 1].

In our setting with a dynamic population of buyers, the dynamic pivot mechanism is efficient, and it corresponds to a sequence of ascending auctions. Meanwhile, Lemma 2 shows that the dynamic virtual pivot mechanism maximizes revenues. It stands to reason, then, that the corresponding indirect mechanism is an appropriately "tweaked" sequence of ascending auctions.

This analogy is, in fact, true when we use ascending auctions with asynchronous price clocks. A central clock, unobserved by the bidders, begins at zero in each period and increases continuously. Meanwhile, each bidder i faces a clock whose speed (relative to the central clock) corresponds to the rate of change of φ_i: when the central clock displays a price p_0, i's clock displays the price $p_i = \varphi_i^{-1}(p_0)$. Thus, each bidder i faces an individualized reserve price $\widetilde{r}_i := \varphi_i^{-1}(0)$.[12] Since $\widetilde{r}_i > 0$ for all i, each bidder must also choose whether or not to participate in the auction. In all other respects, the sequence of ascending auctions proceeds as in Section 3.2.

[12] If $F_i = F$ for all $i \in \mathcal{I}$, then an ascending auction with a "standard" price clock and optimal reserve price suffices.

Suppose that each buyer i present at time t participates in the period-t auction if, and only if, $v_i \geqslant \widetilde{r}_i$. Moreover, suppose that each buyer i remains active in that auction until the price reaches

$$\widetilde{\beta}_t^k\left(\omega_t, v_i, \mathbf{y}_t^{>k}\right) := v_i - \delta\mathbb{E}\left[\pi_i(\omega_{t+1}, \mathbf{v}_{t+1})\middle|\omega_t, \mathbf{v}_t = \left(v_i, \ldots, v_i, \mathbf{y}_t^{>k}\right)\right], \tag{4}$$

where k denotes the number of active bidders, $\omega_t = (\alpha_t, K_t)$ is the state of the market, and $\mathbf{y}_t^{>k}$ denotes the values already inferred from other buyers' exits in the current auction.[13] Thus, each buyer i remains active in the auction until she is just indifferent between winning at the current (individualized) price and receiving her discounted expected contribution to the virtual surplus in the following period. Analogous to Theorems 1 and 2, a seller who wishes to maximize revenues may do so by using a sequence of ascending auctions.

Theorem 3. *Suppose that virtual values φ_i are strictly increasing for all i. The strategy profile in which each buyer i participates if (and only if) $v_i \geqslant \widetilde{r}_i$ and bids according to $\{\widetilde{\beta}_t^k\}$ in Eq. (4) yields a periodic ex post perfect Bayesian equilibrium of the asynchronous sequential ascending auction. Moreover, this equilibrium yields the same allocations and prices as the dynamic virtual pivot mechanism.*

5. Concluding remarks

A critical assumption in our model is that the goods are perishable. When objects are storable, allocating an object in the current period implies forgoing the possibility of allocating it to a higher-valued or more profitable buyer in a future period. In such a setting, the efficient allocation rule is a cutoff policy: a buyer receives an object only when her value is greater than the option value of keeping the good for the future. Li [30] shows that efficient indirect implementation with storable goods continues to require an ascending-type auction format, albeit one more complicated than our own: since the efficient cutoffs differ from the externalities imposed by buyers, marginal-contribution bidding alone is no longer sufficient to guarantee efficiency. On the other hand, Board and Skrzypacz [11] show that, when there is only a single storable unit, a sequence of second-price sealed-bid auctions (with optimally chosen reserve prices) is revenue-maximizing. Unlike the present work, a closed auction format suffices with a single storable good since any bid above the reserve price ends the strategic interaction and eliminates any interdependence.

Another important assumption is that all buyers share a common discount rate and exit probabilities. These restrictions guarantee that the efficient policy allocates to higher-valued buyers before lower-valued ones. This implies that each buyer's marginal contribution to the social surplus depends only on the values of lower-valued buyers and the *number* of higher-valued buyers (but not on the latter's actual valuations). Relaxing these assumptions can preclude the efficiency of our indirect mechanism. When buyers have heterogeneous discount or survival rates, the efficient policy may choose between two agents in some period based on the value of a third agent who (in either case) does not receive an object in that period; this externality is reflected in the fact that the third agent's flow marginal contribution to the social welfare is nonzero despite not being allocated an object. Any auction format in which only winning bidders pay is unable to "price" this externality, ruling out efficient implementation. Clearly, this intuition extends immediately to revenue maximization.

[13] We treat the inferences about a nonparticipating buyer j as a point instead of a distribution. This is without loss of generality as the expectation in Eq. (4) is constant across all beliefs about v_j with support in $[0, \widetilde{r}_j]$.

Our 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 trade-offs in any auction mechanism, as expected future payoffs are no longer symmetric when buyers have differential demands or are faced with multi-unit "exposure" risks. Even if informational asymmetries can be resolved via information revelation and memoryless strategies, such approaches cannot resolve the asymmetries 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 work would provide an important building block to an understanding of competing marketplaces and platforms. We leave these questions, however, for future work.

Appendix A

Proof of Theorem 1. We will first show that the bidding strategies $\{\widehat{\beta}_t^k\}$ are strictly increasing. Fix an arbitrary $t \in \mathbb{N}$, and let $\omega_t := (\alpha_t, K_t)$ denote the state at time t. Consider an agent i 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 $k \in \{1, \dots, n_t\}$. Since $\widehat{\mathbf{x}}$ assigns objects to buyers in decreasing order of values, the presence of buyer i with value v_i does not affect the timing of allocation to buyers j with $v_j > v_i$. Furthermore, since ties may be broken arbitrarily without affecting the social surplus, there is no trade-off between allocating to i or to any other j with $v_j = v_i$. Thus, i's marginal contribution to the social welfare is independent of the values of any buyers j with $v_j \geqslant v_i$. As this is true for any state and any configuration of values (and since $v_i \leqslant \bar{v}$), this implies that

$$\begin{aligned}
&\mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(v_i, \dots, v_i, \mathbf{y}_t^{>k}\big)\big] \\
&\quad = \mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_i, \mathbf{y}_t^{>k}\big)\big],
\end{aligned}$$

where we define $\bar{\mathbf{v}}^m := (\bar{v}, \dots, \bar{v}) \in [0, \bar{v}]^m$. Therefore, for any $v_i' > v_i$, we may write

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

where the equality is because the surplus when i is not present is independent of i's value.

Moreover, by treating buyer i with value v_i' as though her true value were v_i, we can bound the difference above. In particular, we have

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

Thus, whenever $v_i' > v_i$, we may conclude that

$$\widehat{\beta}_t^k\big(\omega_t, v_i', \mathbf{y}_t^{>k}\big) - \widehat{\beta}_t^k\big(\omega_t, v_i, \mathbf{y}_t^{>k}\big) \geqslant \big(v_i' - v_i\big)\left(1 - \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t}\widehat{x}_{i,s}(\omega_s, \mathbf{v}_s)\Big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_i', \mathbf{y}_t^{>k}\big)\right]\right).$$

This lower bound is strictly positive, as the discounted expected probability of future allocation is bounded above by $\delta < 1$. Thus, $\widehat{\beta}_t^k(\omega_t, v_i, \mathbf{y}_t^{>k})$ is strictly increasing in v_i.

Also, note that if $v_i > v_j = y_t^k$, then we may write

$$\begin{aligned}
&\mathbb{E}\big[w_j(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(v_j, \dots, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(v_i, \dots, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&= \mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W_{-j}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad + \mathbb{E}\big[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&= \mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad + \mathbb{E}\big[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W_{-j}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big].
\end{aligned}$$

This expression may be then rewritten as

$$\begin{aligned}
&\mathbb{E}\big[w_j(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(v_j, \dots, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(v_i, \dots, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&= \big(\mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big]\big) \\
&\quad + \big(\mathbb{E}\big[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_j, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W_{-i}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_i, \mathbf{y}_t^{>k}\big)\big]\big) \\
&\quad + \big(\mathbb{E}\big[W_{-j}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-2}, v_i, v_i, \mathbf{y}_t^{>k}\big)\big] \\
&\quad - \mathbb{E}\big[W_{-j}(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t = \big(\bar{\mathbf{v}}^{k-1}, v_j, \mathbf{y}_t^{>k}\big)\big]\big),
\end{aligned}$$

which is simply a sum of three differences. The first is the expected gain in surplus when increasing i's value from v_i to $\bar{v}$. The second is the expected gain in 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 surplus (when j is absent) from decreasing i's value from $\bar{v}$ to v_i. However, since $v_j < v_i$, j's presence or absence has no influence on when the efficient policy allocates to i, regardless of whether i's value is v_i or $\bar{v}$. Therefore, the gain from the first difference equals the loss from the third, implying that

$$\begin{aligned}
&\mathbb{E}\big[w_j(\omega_{t+1},\mathbf{v}_{t+1})\big|\omega_t,\mathbf{v}_t=\big(v_j,\dots,v_j,\mathbf{y}_t^{>k}\big)\big]\\
&\quad-\mathbb{E}\big[w_i(\omega_{t+1},\mathbf{v}_{t+1})\big|\omega_t,\mathbf{v}_t=\big(v_i,\dots,v_i,v_j,\mathbf{y}_t^{>k}\big)\big]\\
&=\mathbb{E}\big[W_{-i}(\omega_{t+1},\mathbf{v}_{t+1})\big|\omega_t,\mathbf{v}_t=\big(\bar{\mathbf{v}}^{k-2},v_i,v_j,\mathbf{y}_t^{>k}\big)\big]\\
&\quad-\mathbb{E}\big[W_{-i}(\omega_{t+1},\mathbf{v}_{t+1})\big|\omega_t,\mathbf{v}_t=\big(\bar{\mathbf{v}}^{k-2},v_i,v_i,\mathbf{y}_t^{>k}\big)\big].
\end{aligned}$$

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

$$\widehat{\beta}_t^{k-1}\big(\omega_t,v_i,v_j,\mathbf{y}_t^{>k}\big)-\widehat{\beta}_t^{k}\big(\omega_t,v_j,\mathbf{y}_t^{>k}\big)>0.$$

Thus, the k-th ranked buyer's exit does not cause the immediate exit of any higher-valued buyers; therefore, the bid functions in Eq. (3) are fully separating and yield an efficient allocation.

To see that these strategies lead to prices equal to the dynamic pivot mechanism's transfers, again fix an arbitrary period t and state ω_t. Notice that when $\widehat{x}_{i,t}(\omega_t,\mathbf{v}_t)=1$, i's entire contribution to the social welfare is realized in period t. Likewise, when $\widehat{x}_{i,t}(\omega_t,\mathbf{v}_t)=0$ buyer i's contribution to the social welfare is realized entirely in the future. Thus, for all $(\omega_t,\mathbf{v}_t)$, we may write the payment from Eq. (1) as

$$\widehat{p}_{i,t}(\omega_t,\mathbf{v}_t)=\widehat{x}_{i,t}(\omega_t,\mathbf{v}_t)\big(v_i-w_i(\omega_t,\mathbf{v}_t)\big).$$

When $K_t\geqslant n_t$, the auction ends immediately, and all buyers present receive an object for free. Similarly, in the dynamic pivot mechanism, i's presence does not impose any externalities on the other agents as there are more objects than buyers; that is, $w_i(\omega_t,\mathbf{v}_t)=v_i$ and $\widehat{p}_{i,t}(\omega_t,\mathbf{v}_t)=0$.

Suppose instead that $K_t<n_t$, and denote by i_k the bidder with the k-th highest value. In the direct mechanism, a buyer i who receives an object pays $\widehat{p}_{i,t}(\omega_t,\mathbf{v}_t)=v_i-w_i(\omega_t,\mathbf{v}_t)$. However,

$$\begin{aligned}
w_i(\omega_t,\mathbf{v}_t)&=\mathbb{E}\Bigg[\sum_{s=t}^{\infty}\sum_{j\in\mathcal{I}}\delta^{s-t}\widehat{x}_{j,s}(\omega_s,\mathbf{v}_s)v_j\Bigg|\omega_t,\mathbf{v}_t\Bigg]\\
&\qquad-\mathbb{E}\Bigg[\sum_{s=t}^{\infty}\sum_{j\in\mathcal{I}\setminus\{i\}}\delta^{s-t}\widehat{x}_{j,s}\big(\omega_s^{-i},\mathbf{v}_s\big)v_j\Bigg|\omega_t^{-i},\mathbf{v}_t\Bigg]\\
&\qquad=\Bigg(\sum_{m=1}^{K_t}v_m+\mathbb{E}\Bigg[\sum_{s=t+1}^{\infty}\sum_{j\in\mathcal{I}}\delta^{s-t}\widehat{x}_{j,s}\big(\omega_s^{-i},\mathbf{v}_s\big)v_j\Bigg|\omega_t,\mathbf{v}_t\Bigg]\Bigg)\\
&\qquad-\Bigg(\sum_{m=1}^{K_t+1}v_m-v_i+\mathbb{E}\Bigg[\sum_{s=t+1}^{\infty}\sum_{j\in\mathcal{I}}\delta^{s-t}\widehat{x}_{j,s}\big(\omega_s^{-i,-i_{K_t+1}},\mathbf{v}_s\big)v_j\Bigg|\omega_t^{-i},\mathbf{v}_t\Bigg]\Bigg).
\end{aligned}$$

Therefore, we may write i's payment $\widehat{p}_{i,t}(\omega_t,\mathbf{v}_t)$ as

$$\begin{aligned}
&v_{i_{K_t+1}}-\mathbb{E}\Bigg[\sum_{s=t+1}^{\infty}\sum_{j\in\mathcal{I}}\delta^{s-t}\widehat{x}_{j,s}(\omega_s,\mathbf{v}_s)v_j\Bigg|\omega_t,\mathbf{v}_t\Bigg]\\
&\qquad+\mathbb{E}\Bigg[\sum_{s=t+1}^{\infty}\sum_{j\in\mathcal{I}}\delta^{s-t}\widehat{x}_{j,s}\big(\omega_s^{-i_{K_t+1}},\mathbf{v}_s\big)v_j\Bigg|\omega_t,\mathbf{v}_t\Bigg]\\
&\quad=v_{i_{K_t+1}}-\delta\mathbb{E}\big[w_{i_{K_t+1}}(\omega_{t+1},\mathbf{v}_{t+1})\big|\omega_t,\mathbf{v}_t\big]=\widehat{\beta}_t^{K_t+1}(\omega_t,v_{i_{K_t+1}},\dots,v_{i_{n_t}}),
\end{aligned}$$

which is precisely the bid made by the buyer with the $(K_t + 1)$-st highest value. Thus, whether K_t is more or less than n_t, the payments of the direct and indirect mechanisms are identical. □

Proof of Theorem 2. Consider any period $t \in \mathbb{N}$ with $n_t := \sum_{j\in\mathcal{I}} \alpha_t(j)$ 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 $\{\widehat{\beta}_t^k\}$. 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.

We label agents such that buyer i_k has the k-th highest value. Note that if $v_i < v_{i_{K_t}}$, bidding according to $\widehat{\beta}_t^k$ 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 future bidding does not depend on information revealed in the present, i's future payoffs are also unaffected by an early exit. Suppose, on the other hand, that $v_i \geqslant v_{i_{K_t}}$. As established in Theorem 1, i receives an object and pays 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 her future marginal contribution to the social welfare. This is a profitable one-shot deviation (ex post) for i if, and only if,

$$\delta\mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t\big] \geqslant v_i - \widehat{\beta}_t^{K_t+1}\big(\omega_t, v_{i_{K_t+1}}, \mathbf{y}_t^{>K_t+1}\big),$$

or (equivalently) if, and only if,

$$\widehat{\beta}_t^{K_t+1}\big(\omega_t, v_{i_{K_t+1}}, \mathbf{y}_t^{>K_t+1}\big) \geqslant v_i - \delta\mathbb{E}\big[w_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \mathbf{v}_t\big] = \widehat{\beta}_t^{K_t+1}\big(\omega_t, v_i, \mathbf{y}_t^{>K_t+1}\big),$$

where the equality follows from the fact that i's future marginal contribution to the social welfare does not depend on the values of buyers ranked above her. Since $v_i > v_{i_{K_t+1}}$, this contradicts the efficiency property established by Theorem 1. Thus, i never benefits from exiting early.

Alternately, if $v_i \geqslant v_{i_{K_t}}$, 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, $i = i_k$ for some $k > K_t$ (so $v_i < v_{i_{K_t}}$), 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 present, this only occurs if i delays exit by 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. Then a deviation to remaining active in the auction is profitable (ex post) if, and only if,

$$v_{i_k} - \widehat{\beta}_t^{K_t+1}\big(\omega_t, v_{i_{K_t}}, \ldots, v_{i_{k-1}}, \mathbf{y}_t^{>k}\big) \geqslant \delta\mathbb{E}\big[w_{i_k}(\omega_{t+1}, \mathbf{v}_t)\big|\omega_t, \mathbf{v}_t\big].$$

Rearranging this inequality yields

$$\widehat{\beta}_t^{K_t+1}\big(\omega_t, v_{i_{K_t}}, \ldots, v_{i_{k-1}}, \mathbf{y}_t^{>k}\big) \leqslant v_{i_k} - \delta\mathbb{E}\big[w_{i_k}(\omega_{t+1}, \mathbf{v}_t)\big|\omega_t, \mathbf{v}_t\big] = \widehat{\beta}_t^{k}(\omega_t, v_{i_k}, \ldots, v_{i_{n_t}}),$$

where we have again made use of the fact that i's future marginal contribution to the social welfare is independent of the values of those buyers ranked above her. As above, the fact that $v_{i_k} < v_{i_{K_t}}$ contradicts the efficiency property established by Theorem 1. Therefore, i does not prefer to remain active in the auction long enough to receive an object.

Thus, no player has an ex post incentive to deviate from the prescribed strategies when on the equilibrium path—bidding according to $\{\widehat{\beta}_t^k\}$ is ex post 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, and instead believe that the deviating agent is *currently* being "sincere" with regards to $\{\widehat{\beta}_t^k\}$. 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 Eq. (3) is a (periodic ex post) perfect Bayesian equilibrium of the sequential ascending auction mechanism. □

Proof of Lemma 2. To show that the dynamic virtual pivot mechanism is periodic ex post incentive compatible, fix an arbitrary agent i arriving in period $t \in \mathbb{N}$, and suppose that i knows the reported values $\mathbf{v}_t^{-i}$ of all other agents who are present at time t. By reporting a value v_i' upon her arrival in state ω_t, i's payoff is

$$\begin{aligned}
&\left(\widetilde{x}_{i,t}\big(\omega_t, \big(v_i', \mathbf{v}_t^{-i}\big)\big) + \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t}\widetilde{x}_{i,s}\big(\omega_s, \big(v_i', \mathbf{v}_s^{-i}\big)\big)\Big|\omega_t, \mathbf{v}_t\right]\right)\big(v_i - v_i'\big) \\
&\quad + \big(\pi_i\big(\omega_t, \big(v_i', \mathbf{v}_t^{-i}\big)\big) - \delta\mathbb{E}\big[\pi_i(\omega_{t+1}, \mathbf{v}_{t+1})\big|\omega_t, \big(v_i', \mathbf{v}_t^{-i}\big)\big]\big) \\
&= \mathbb{E}\left[\sum_{s=t}^{\infty}\sum_{j\in\mathcal{I}} \delta^{s-t}\widetilde{x}_{j,s}\big(\omega_s, \big(v_i', \mathbf{v}_s^{-i}\big)\big)\big(\varphi_i^{-1}\big(\varphi_j(v_j)\big) - \varphi_i^{-1}(0)\big)\Big|\omega_t, \mathbf{v}_t\right] \\
&\quad - \Pi_{-i}^i\big(\omega_t, \mathbf{v}_t^{-i}\big),
\end{aligned}$$

where the expectation is taken with respect to the true value distributions of agents arriving in periods $s > t$. Moreover, the monotonicity of φ_i implies that $\varphi_i^{-1}(\varphi_j(v_j)) - \varphi_i^{-1}(0) \geqslant 0$ if, and only if, $\varphi_j(v_j) \geqslant 0$. Therefore, $\widetilde{\mathbf{x}}$ is an "efficient" policy for a "planner" maximizing the above sum of transformed virtual values. Applying Lemma 1, the first term above is thus maximized by reporting $v_i' = v_i$. In addition, the second term does not depend on v_i'. Hence, given truth-telling by agents arriving in the future, truthful reporting is (ex post) optimal for buyer i.

In addition, recall that the dynamic virtual pivot mechanism yields buyer i a flow payoff in each period equal to her flow marginal contribution to the virtual surplus. By definition of the optimal allocation rule $\widetilde{\mathbf{x}}$, this payoff is always nonnegative. Hence, the dynamic virtual pivot mechanism is periodic ex post individually rational. □

Proof of Theorem 3. The proof is omitted, as it parallels those of Theorems 1 and 2. □

References

[1] S. Athey, I. Segal, Designing efficient mechanisms for dynamic bilateral trading games, Amer. Econ. Rev. 97 (2) (2007) 131–136.
[2] S. Athey, I. Segal, An efficient dynamic mechanism, unpublished manuscript, Harvard University, 2007.
[3] L.M. Ausubel, An efficient ascending-bid auction for multiple objects, Amer. Econ. Rev. 94 (5) (2004) 1452–1475.
[4] M. Backus, G. Lewis, An estimable demand system for a large auction platform market, unpublished manuscript, Harvard University, 2010.
[5] J.S. Banks, J.O. Ledyard, D.P. Porter, Allocating uncertain and unresponsive resources: An experimental approach, RAND J. Econ. 20 (1) (1989) 1–25.
[6] D. Bergemann, M. Said, Dynamic auctions, in: J.J. Cochran (Ed.), Wiley Encyclopedia of Operations Research and Management Science, Wiley, New York, 2011.
[7] D. Bergemann, J. Välimäki, Dynamic common agency, J. Econ. Theory 111 (1) (2003) 23–48.
[8] D. Bergemann, J. Välimäki, Dynamic price competition, J. Econ. Theory 127 (1) (2006) 232–263.
[9] D. Bergemann, J. Välimäki, The dynamic pivot mechanism, Econometrica 78 (2) (2010) 771–789.
[10] S. Bikhchandani, P.A. Haile, J.G. Riley, Symmetric separating equilibria in English auctions, Games Econ. Behav. 38 (1) (2002) 19–27.

[11] S. Board, A. Skrzypacz, Revenue management with forward-looking buyers, unpublished manuscript, University of California, Los Angeles, 2010.
[12] E. Budish, R. Zeithammer, An efficiency ranking of markets aggregated from single-object auctions, unpublished manuscript, UCLA Anderson School of Management, 2011.
[13] G. Cai, P.R. Wurman, X. Chao, The non-existence of equilibrium in sequential auctions when bids are revealed, J. Electron. Commer. Res. 8 (2) (2007) 141–156.
[14] B. Caillaud, J. Robert, Implementation of the revenue-maximizing auction by an ignorant seller, Rev. Econ. Design 9 (2) (2005) 127–143.
[15] G. Calzolari, A. Pavan, Monopoly with resale, RAND J. Econ. 37 (2) (2006) 362–375.
[16] G. Calzolari, A. Pavan, On the optimality of privacy in sequential contracting, J. Econ. Theory 130 (1) (2006) 168–204.
[17] O. Compte, R. Lavi, E. Segev, Efficiency of sequential English auctions with dynamic arrivals, in: Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12), Valencia, 2012.
[18] R.J. Dolan, Incentive mechanisms for priority queuing problems, Bell J. Econ. 9 (2) (1978) 421–436.
[19] D. Garrett, Durable goods sales with dynamic arrivals and changing values, unpublished manuscript, Northwestern University, 2011.
[20] A. Gershkov, B. Moldovanu, Dynamic revenue maximization with heterogeneous objects: A mechanism design approach, Am. Econ. J. Microecon. 1 (2) (2009) 168–198.
[21] A. Gershkov, B. Moldovanu, Demand uncertainty and dynamic allocation with strategic agents, unpublished manuscript, University of Bonn, 2010.
[22] A. Gershkov, B. Moldovanu, Efficient sequential assignment with incomplete information, Games Econ. Behav. 68 (1) (2010) 144–154.
[23] P.A. Haile, Auctions with private uncertainty and resale opportunities, J. Econ. Theory 108 (1) (2003) 72–110.
[24] K. Hendricks, I. Onur, T. Wiseman, Last-minute bidding in sequential auctions with unobserved, stochastic entry, Rev. Ind. Organ. 40 (1) (2012) 1–19.
[25] P. Jehiel, B. Moldovanu, Auctions with downstream interaction among buyers, RAND J. Econ. 31 (4) (2000) 768–791.
[26] J.H. Kagel, R.M. Harstad, D. Levin, Information impact and allocation rules in auctions with affiliated private values: A laboratory study, Econometrica 55 (6) (1987) 1275–1304.
[27] J.H. Kagel, D. Levin, Implementing efficient multi-object auction institutions: An experimental study of the performance of boundedly rational agents, Games Econ. Behav. 66 (1) (2009) 221–237.
[28] T. Kittsteiner, J. Nikutta, E. Winter, Declining valuations in sequential auctions, Int. J. Game Theory 33 (1) (2004) 89–106.
[29] R. Lavi, N. Nisan, Online ascending auctions for gradually expiring items, in: Proceedings of the 16th Symposium on Discrete Algorithms (SODA 2005), Vancouver, 2005.
[30] A. Li, Selling storable goods to a dynamic population of buyers: A mechanism design approach, unpublished manuscript, Stanford University, 2012.
[31] E. Maskin, Auctions and privatization, in: H. Siebert (Ed.), Privatization: Symposium in Honor of Herbert Giersch, J.C.B. Mohr (Paul Siebeck), Tübingen, 1992, pp. 115–136.
[32] K. Mierendorff, Optimal dynamic mechanism design with deadlines, unpublished manuscript, University of Zurich, 2011.
[33] P.R. Milgrom, R.J. Weber, A theory of auctions and competitive bidding, II, in: P. Klemperer (Ed.), The Economic Theory of Auctions, Edward Elgar Publishing, Cheltenham, UK, 2000.
[34] R.B. Myerson, Optimal auction design, Math. Oper. Res. 6 (1) (1981) 58–73.
[35] B.J. Nalebuff, J.I. Bulow, Designing the PCS auction, Attached to comments of Bell Atlantic Personal Communications, Inc., PP Docket 93-253, 1993.
[36] M. Pai, R. Vohra, Optimal dynamic auctions and simple index rules, unpublished manuscript, University of Pennsylvania, 2011.
[37] R. Pancs, Designing order-book transparency in electronic communication networks, unpublished manuscript, University of Rochester, 2011.
[38] R. Pancs, Sequential negotiations with costly information acquisition, unpublished manuscript, University of Rochester, 2011.
[39] A. Pavan, I. Segal, J. Toikka, Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure, unpublished manuscript, Northwestern University, 2011.
[40] M. Said, Information revelation and random entry in sequential ascending auctions, in: Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08), Chicago, 2008.

[41] M. Said, Sequential auctions with randomly arriving buyers, Games Econ. Behav. 73 (1) (2011) 236–243.
[42] G. Vulcano, G. van Ryzin, C. Maglaras, Optimal dynamic auctions for revenue management, Manage. Sci. 48 (11) (2002) 1388–1407.
[43] R. Zeithammer, Strategic bid-shading and sequential auctioning with learning from past prices, Manage. Sci. 53 (9) (2007) 1510–1519.
[44] J. Zhang, R. Wang, The role of information revelation in elimination contests, Econ. J. 119 (536) (2009) 613–641.