Auction Timing and Market Thickness*

Isaías N. Chaves† Shota Ichihashi‡

August 18, 2020

Abstract

A seller faces a pool of potential bidders that changes over time. She can delay the auction at a cost, in the hopes of having a thicker market later on. The seller imposes both static distortions (through her choice of reserve prices) and dynamic distortions (through her choice of market thickness). We isolate the effects of partial regulations: regulations that restrict the seller's reserve price, but leave market thickness unconstrained, and regulations that restrict market thickness, but leave the choice of reserve price unconstrained. We show that regulating only the static distortion can harm efficiency. We also show that dynamically responding to changes in the bidder pool is essential: committing to delay until an optimal deadline can waste most of the achievable revenue.

JEL Classification Numbers: D42, D44, D47, D82, L12.

Keywords: Auctions, Market Design, Optimal Stopping, Monopoly, Order Statistics.

*Many thanks to Kyle Bagwell, Gabriel Carroll, Matt Jackson, Fuhito Kojima, Maciej Kotowski, Paul Milgrom, Ilya Segal, Andy Skrzypacz, and Takuo Sugaya for insightful comments. We also thank audiences at Stanford, IIOC, and Michigan. Chaves thanks the Forman Family Fellowship in Economics. Ichihashi is grateful to Yoshida Scholarship Foundation and the Stanford Graduate Fellowship. An earlier version of this paper circulated under the title "When to Hold an Auction?" The opinions expressed in this article are the authors' own and do not reflect the views of the Bank of Canada.

†Kellogg School of Management, Northwestern University.

Email: isaias.chaves@kellogg.northwestern.edu

‡Bank of Canada.

Email: shotaichihashi@gmail.com

1 Introduction

The study of auctions usually begins with a known set of bidders. This leaves out a crucial part of the design problem: the auctioneer is often uncertain about who will show up for the auction (and by when) and must trade off the delay costs from searching for bidders against the benefits of market thickness.

Consider, for instance, the market for distressed corporate assets, in which a substantial portion of sales happen via auction (Boone and Mulherin, 2009). After the firm decides to restructure and sell a division, it cannot take for granted that there will be a given set of bidders. Interested bidders for these assets are hard to find, and even when found, it takes considerable time and effort to motivate them to bid. The process usually involves both hiring an investment bank at significant cost and several months' delay. Boone and Mulherin (2009) note that the process takes more than six months on average, a long and costly wait for a company in distress. These search frictions are especially salient in corporate bankruptcy sales. LoPucki and Doherty (2007) remark that, in their sample,

> On average, eighty prospects are contacted for each sale, and thirty sign confidentiality agreements. But the average number of bidders is only 1.6 per sale. In fifteen of the twenty-six cases for which we have data (58%), there was only one bidder (p.35).

An earlier study had found that only eighteen of fifty-five sales in the sample had multiple bidders (Hotchkiss and Mooradian, 1998).

Delaying these auctions may be necessary to attract bidders, but the delay imposes substantial discounting costs. By delaying, the company risks missing out on a favorable market moment, and it exposes itself to adverse events—e.g., an employee exodus or a loss of key suppliers—that destroy value and make a successful reorganization impossible. During Polaroid's 2002 bankruptcy court case, its management argued for a speedy sale rather than a reorganization because its "revenues [were] falling off" and "[the company] was like a melting ice cube" (LoPucki and Doherty, 2007, p. 54).

Incorporating these delay costs leads to new auction design questions. The firm in our previous vignette must choose not only how to run the auction, but *when* to hold it: how long should the firm search for bidders?[1] Likewise, the relevant market design concerns also change. In a static setting, the main efficiency concern is whether the good goes to the highest value bidder. In contrast, the firm in this example can cause distortions both through its choice of auction format (specifically, its use of reserve prices) and through its choice of timing. Moreover, these two choices interact: the (static) distortions caused by a reserve price affect the firm's incentives to delay the auction and can induce further (dynamic) distortions.

Motivated by these examples, we study the problem of auction timing. The pool of potential bidders changes stochastically over time. Bidders have symmetric, independent private values for an indivisible good. An auctioneer chooses a timing policy for an auction, and a reserve price to use once the auction is run. We assume the auctioneer uses a second price auction, and we look at timing policies that rule out indicative bidding, i.e., bids are only taken after the auctioneer "stops" and runs the auction format.[2] However, we allow the pool of potential bidders to evolve according to either (i) an arbitrary arrival process, or (ii) an arrival-departure process in a broad class. Considering this class of policies lets us focus squarely on the optimal choice of *market thickness*, and it leads to a clean distinction between static (auction format/reserve price) and dynamic (auction timing) distortions.

We ask two main questions. First, do revenue-maximizing sellers over- or under-value market thickness, i.e., do they wait too long or too little to hold auctions, relative to the social planner? A revenue-maximizing seller will typically use a different reserve price than a social planner, which in turn feeds into the choice of market thickness, so a full answer requires comparing the different auctioneers' choices across reserve prices.

[1] Real estate auctions have some similar features—the seller has a joint choice of auction format and auction timing—but we emphasize auctions for corporate auctions because discounting costs, which are crucial to our model, seem more significant in that market: sellers of single family homes may or may not care about delaying an additional week or three to obtain one more bidder.

[2] Therefore, we do not solve the full dynamic mechanism design problem (see, e.g., Gershkov, Moldovanu, and Strack, 2018 or Board and Skrzypacz, 2016), but we can consider a very general class of bidder dynamics. We discuss the advantages of our modeling approach in more detail in Section 2. Note that if the auctioneer can commit in advance to a mechanism that will be run upon stopping, then a revenue maximizing auctioneer will use (WLOG) a second price auction with an optimal reserve price, whereas a surplus-maximizing auctioneer will use (WLOG) a second price auction with an efficient reserve price.

Second, we investigate the equilibrium effects of *partial regulation*: what is the impact of regulatory interventions that only target the static distortion (choice of reserve price) but leave the dynamic distortion (choice of timing and market thickness) unregulated? Conversely, what about policy interventions that directly target the dynamic distortion (i.e., enforce a timing rule or level of market thickness), but leave the static distortion unregulated?

For example, in internet search advertising markets, one could think of regulators intervening in an advertising platform's ranking algorithm without affecting the "latency" (i.e., how long the webpage waits for bids before displaying search results to the user). Recently, large advertising platforms have been accused of downgrading results from competing services on their search advertising platforms; EU regulators, in particular, have considered directly influencing Google's ranking policies to prevent it from downranking competing services. This intervention targets the auction for search results without affecting market thickness.[3] In contrast, regulators require stock exchanges to impose temporary trading halts following market upheaval; trade resumes with a "halt auction" (often uniform price double auctions).[4] Regulators could impose requirements on the length of the trading halt, without tinkering with the mechanics of the halt auction. The longer the trading halt, the more orders (bidders) would accumulate for the eventual halt auction.

We identify a property of the distribution of bidder values that governs both the dynamic distortions caused by revenue maximization and the interactions between static and dynamic distortions. Let F be the distribution of valuations with density f, and define $\xi_F(v) := \frac{f(v)}{1-F(v)}v$. For an auctioneer $a \in \{\text{seller}, \text{planner}\}$ and a reserve price $r \in \{\text{optimal}, \text{efficient}\}$,

[3] When a large advertising platform like Google Adwords sells an impression, it queries different ad exchanges that aggregate bids from advertisers. The longer platform waits for responses, the more bids it will have for a particular impression, but the slower the page will load. This downgrades the internet user's experience and reduces future traffic for the platform. Since advertisers bid for slots, we can think of ranking regulations as an intervention that targets static distortions (i.e., lowering the reserve price for Yelp on a restaurant query or for Expedia on a flight query), but leaves the choice of market thickness (how long to wait for exchange responses) constant.

[4] For instance, see the description of the halt auction for the Chicago Board Options Exchange at https://cdn.cboe.com/resources/membership/Cboe_US_Equities_Auction_Process.pdf

say $(a, r) \succeq_{stop} (a', r')$ if, for any arrival process for bidders in a large class, a always holds an auction with reserve r later (with more bidders) than a' holds auction with reserve r'.[5] Our main result (Theorem 1), informally stated, says that

$$\begin{pmatrix} \text{seller,} \\ \text{efficient } r \end{pmatrix} \succeq_{stop} \begin{pmatrix} \text{seller,} \\ \text{optimal } r \end{pmatrix} \succeq_{stop} \begin{pmatrix} \text{planner,} \\ \text{optimal } r \end{pmatrix} \succeq_{stop} \begin{pmatrix} \text{planner,} \\ \text{efficient } r \end{pmatrix}. \tag{1}$$

whereas if ξ_F is decreasing, the string of "inequalities" is exactly reversed.

In particular, for ξ_F increasing, a seller inefficiently over-invests in market thickness, whereas for ξ_F decreasing, she inefficiently under-invests.[6] Regular distributions in the sense of Myerson (1981) can have increasing or decreasing ξ_F's, and therefore result in sellers that over- or under-invest in market thickness. On balance, however, we find that a revenue-maximizing seller inefficiently over-invests in market thickness for many common distributions. In particular, ξ_F is trivially increasing for all increasing hazard rate distributions.

Fully ordering delay or "thickness" decisions across auctioneers and reserve prices allows us to evaluate partial regulation that targetes only the static distortion or only the dynamic distortion. We focus on the case of increasing ξ_F, which is arguably the natural one. Our results highlight a harmful unintended consequence from regulating *only* the reserve price if the planner cannot control the choice of market thickness. When forced to use an efficient reserve price and sell more often than she would like, a seller with monopoly power could respond by (i) delaying more, in the hopes that additional bidders will make up for the low-revenue auction, or (ii), delay *less*, in the hopes of collecting the receipts from the low-revenue auction sooner. We show that the seller always responds by doing (i). Therefore, depending on the discount factor and valuation distribution, partial regulation can *lower* social surplus by generating additional delay costs.

[5]Here, we use the minimum of the support of F, which can be positive, as an efficient reserve price for the second price auction. This uniquely pins down the revenue from the efficient auction.

[6] Recall Bulow and Roberts (1989)'s monopoly pricing interpretation of auction theory, where an auction problem facing a distribution of values F is equivalent to a third-degree price discrimination problem facing a demand curve $1 - F$. Then ξ_F corresponds to the price elasticity of that demand curve. Prior literature refers to ξ_F as a generalized failure rate, and it explores the implications of an increasing generalized failure rate. We postpone that discussion to the related literature section.

Theorem 1 also gives some guidance on how a regulator should intervene to affect market thickness if she cannot control the auction mechanism. Facing the loss of surplus from being forced to use an inefficient reserve price, the planner could respond by (i) delaying more, in the hope that drawing from a larger pool of bidders makes up for the low-surplus auction, or (ii), delaying *less*, to realize the surplus from that low-surplus auction faster. We show that for "standard" valuation distributions, the planner should always choose (i).

Finally, suppose that a regulator who could intervene on market thickness suddenly gained the additional ability to regulate the auction mechanism (i.e., the reserve price). Should she delay more than when she was constrained, to use the high-surplus auction on a larger pool, or less, to realize the surplus from the efficient auction sooner? We show that she should always do (ii). In other words, the regulator who can intervene on both margins will delay less than the fully unregulated seller. In that sense, *market thickness is inefficiently high.*[7]

The following example summarizes the above discussion:

Example 1. *Suppose that the auctioneer chooses the auction date over a time horizon* $t \in \{1, 2, 3, 4\}$*. Denoting the number of bidders in period* t *by* n_t*, let* $(n_1, n_2, n_3, n_4) = (1, 2, 4, 9)$*. The discount factor is* $\delta = 0.6$*, so at time* t*, both revenue and surplus are discounted by* δ^{t-1}*. Bidders' values are drawn iid from an exponential distribution with mean* 1*. For each* $n \in \{1, 2, 3, 4\}$*, let* $R(n, p)$ *and* $W(n, p)$ *denote the expected revenue and surplus, respectively, from a second price auction with a reserve price of* p *and* n *bidders. Let* p^* *denote the optimal reserve price. Table 1 then presents the discounted expected revenue and surplus from each possible timing choice and reserve price* $p \in \{0, p^*\}$ *(rounding to two decimal places). The optimal auction time for each row is highlighted in gray.*

In this example, $\xi_F(v) = v$ *is increasing. The optimal timing choices in gray satisfy* (1).

[7] If the auctioneers are choosing a stopping time adapted to bidder arrivals that happen according to a homogenous Poisson process, one can show that marginal decreases in thickness are socially beneficial. It is easy to show that the optimal stopping time consists of a time-independent cutoff in the number of bidders. The proof techniques used in Appendix C show that the planner's objective is single peaked in the choice of cutoff. Therefore, unless the planner and seller have the same cutoff, reducing that cutoff on the margin increases efficiency.

Table 1: Expected Discounted Payoffs

t	1	2	3	4
$\delta^t W(t,0)$	1.00	0.90	0.75	0.61
$\delta^t W(t,p^*)$	0.74	0.76	0.70	0.61
$\delta^t R(t,p^*)$	0.37	0.40	0.41	0.40
$\delta^t R(t,0)$	0	0.30	0.39	0.40

Forcing the seller to use an efficient reserve price reduces discounted total surplus. With an optimal reserve price ($p^* = 1$), the seller holds the auction at $t = 3$ and generates discounted total surplus of $\delta^3 W(3, p^*) = 0.70$, whereas with an efficient reserve price ($p = 0$), the seller holds the auction at $t = 4$ and generates surplus of $\delta^4 W(4, 0) = 0.61$. In contrast, holding the reserve price fixed, (i) forcing the seller to hold an auction at an efficient time increases discounted total surplus (from 0.61 *to* 1.00 *under zero reserve, and from* 0.70 *to* 0.76 *under the optimal reserve price), and (ii) this surplus is always maximized at an earlier time than was chosen by the seller.*

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes the environment and formally sets up the auction timing problem. Section 4 describes a notion of right-skewness—the star-shaped order—for a value distribution. Focusing on the case with no departures, this section establishes several comparative statics using that notion, and it studies the inefficiencies that arise from revenue-maximization. Section 5 considers several extensions. There we consider comparative statics in the distribution of values, and we extend our auction timing results to the case when bidders depart stochastically. Finally, we show that optimal nonadaptive timing rules (i.e., holding the auction at a fixed deadline) can perform arbitrarily worse than the fully optimal adaptive rule—the gains from having access to the history of bidder arrivals can therefore be arbitrarily large. Section 7 concludes. All omitted proofs are relegated to the Appendices.

2 Related Literature

Wang (1993) appears to be the first paper to study the timing of an auction as an endogenous choice by a seller. Wang analyzes a model in which the seller has one unit of an object to sell to buyers who arrive according to a homogenous Poisson process. The seller who bears a constant flow cost from holding inventory can decide to sell the object either by posting a price or by holding an auction at a fixed deadline chosen in advance. Wang shows that the seller's choice of timing typically differs from the socially optimal one. In contrast, we consider general arrival and departure dynamics, and we allow the seller much more leeway in choosing when to hold the auction. In our model, the auctioneer can use any stopping time that is adapted to the history of arrivals and departures. Cong (2019) studies an auction timing problem in which a seller sells real options to a fixed set of long-lived bidders. Among other things, he finds that the seller inefficiently delays the auction. Arefeva and Meng (2020) study a two-period auction timing problem in which delaying an auction enables bidders to learn about their outside options and update their values.

Our results relate to prior static models on bidder solicitation and auctions (Szech, 2011; Fang and Li, 2015). These papers study models where the auctioneer can pay a linear cost to increase the number of bidders in the auction. In contrast to these papers, the solicitation cost in our paper comes from time discounting.

We give a formal comparison between our results and this prior work in Section 4. At a high level, there are four main differences. First, these prior results only rank choices over both market thickness and reserve prices for the less natural case of a decreasing hazard rate. Second, the results in Szech (2011) do not fully rank market thickness decisions across auctioneers and reserve prices. By giving such a complete ranking, we can cleanly separate the role of static and dynamic distortions arising from revenue maximization; as a consequence, we can also study the effects of partial interventions that target only one of those two margins. Third, in terms of modeling choices, whether linear or discounting costs are more appropriate depends on the application, but in many leading examples, the key delay cost is the vanishing of a trading opportunity. For example, in the case of bankruptcy auctions, a major source of delay costs is that before the assets are sold, adverse contingencies may arise that destroy the value of the assets. Delay costs will therefore enter multiplicatively.[8] Finally, we argue that excessive delay costs are a more fundamental inefficiency than excessive (linear) solicitation costs. Excessive expenses on advertising and solicitation are a pure transfer to an advertising agency or an investment bank, whereas excessive delay, since it causes the vanishing of efficient trading opportunities, directly destroys surplus that is not captured by any party.

This paper fits within a broader literature on how to sell (or allocate) to stochastically arriving agents with private information (Gallien, 2006; Board and Skrzypacz, 2016; Loertscher, Muir, and Taylor, 2016; Gershkov et al., 2018). These papers derive a fully optimal dynamic mechanism for a more restricted class of agents' arrival processes, and they focus especially on when these mechanisms can be implemented in posted prices. In contrast, we focus on the question of how to time the transaction. At the cost of looking at a more restricted class of mechanisms, this modeling approach has several advantages. Aside from the substantive focus it allows—emphasizing issues of market thickness/timing, as well as cleanly distinguishing between static and dynamic distortions—the approach enables us to consider very general arrival dynamics and some departure dynamics,[9] and it makes possible monotone comparative statics results. In contrast with this literature, we can also identify conditions under which revenue maximization leads to inefficiently *low* market thickness.[10]

Other papers in economics and operations research have used the star-shaped order in auction models. Moldovanu, Sela, and Shi (2008) study sellers who decide on their discrete supply of a homogeneous good in multi-unit auctions. They use the star-shaped order to provide comparative statics with respect the distribution of values in the monopoly problem.

[8] Also note that in our setting, not only do the seller and the planner face different marginal benefits from "soliciting" bidders, but they also incur different marginal costs, which complicates the comparison between the two. This follows from discounting: the planner's cost from delaying an efficient auction by dt is $r \times dt \times$ (expected surplus), whereas the seller's cost is $r \times dt \times$ (expected revenue).

[9] The literature has mostly focused on the case without departures of agents. See Pai and Vohra (2013) and Mierendorff (2016) for exceptions.

[10] In a model of a dynamic two-sided market with random arrivals of finite-type traders, Loertscher et al. (2016) find that revenue maximization leads to inefficiently high market-thickness. The models are substantially different, since we consider a one-sided market but have a richer type space and more general bidder dynamics. We see the approaches as complementary.

Paul and Gutierrez (2004) and Li (2005) use the star-shaped order to study how the winner's rent in an auction depends on the number of bidders. Using the star-shaped order, we provide new comparative statics on how the growth rate of the revenue (as a function of the number of bidders) depends on the value distribution and the auction format.

The condition that $\xi_F(\cdot)$ is increasing was introduced by Lariviere and Porteus (2001) in the revenue management literature under the name of "increasing generalized failure rate" (IGFR). For example, the condition ensures that the monopoly price is unique. Ziya, Ayhan, and Foley (2004) provide an insightful review of this literature and the relationship between IGFR and other distributional assumptions.[11] In the auction theory literature, Kleinberg and Yuan (2013) and Schweizer and Szech (2019) use the condition to obtain quantitative bounds on equilibrium outcomes. For example, Kleinberg and Yuan (2013) call the condition hyper-regularity and obtain a worst-case bound bound on the ratio of expected revenue to expected welfare. Our results highlight a new economic implication of (IGFR): it governs the *relative growth rate* of revenue and total welfare.

3 Setting

An auctioneer chooses between allocating one indivisible good to any of n bidders in $t = 0$, or allocating it to any of $n + d$ bidders in $t = 1$, where $n, d \in \mathbb{N}$. The auctioneer has one-period discount factor $\delta \in (0, 1)$. We focus on comparative statics on timing that are uniform in (n, d, δ). Therefore, delay decisions obtained in this simple setting extend to the richer problem where (i) bidders arrive over time according to any continuous-time birth process $\{N_t, t \geq 0\}$, and (ii) the auctioneer discounts time at a constant rate rate ρ, and (iii), she chooses an optimal stopping time, adapted to $\{N_t, t \geq 0\}$, at which to hold the auction (see Appendix B for details).

All bidders are symmetric, and have an independent private value v_i drawn from a cumulative distribution function F. Assume that (i) F has a finite expectation, (ii) F has a positive and continuous density f over its support $[\underline{v}, \overline{v})$ for some $\underline{v} \geq 0$ and $\overline{v} \in (\underline{v}, \infty]$, and (iii) F

[11] Other contributions include Lariviere (2006) and Banciu and Mirchandani (2013).

is regular, i.e., $v - \frac{1-F(v)}{f(v)}$ is increasing.[12] We use "increasing" to mean "non-decreasing," and "decreasing" to mean "non-increasing." Following Bulow and Roberts (1989), we define type v's marginal revenue as $MR(v) := v - \frac{1-F(v)}{f(v)}$.

The auctioneer runs a second price auction. The auction can be either *optimal*, i.e., with a reserve price of $p^* := \inf\{p \in [\underline{v}, \overline{v}) : MR(p) \geq 0\}$, or *efficient*, i.e., with a reserve price of $\underline{v}$.[13] Denote by $r = \text{Opt}$ that the auction is optimal, and by $r = \text{Eff}$ that the auction is efficient. The auctioneer is constrained to using the same kind of auction in either period. The auctioneer $a \in \{S, P\}$ can be a *seller* who maximizes discounted expected revenue ($a = S$), or a *(social) planner* who maximizes discounted expected total surplus ($a = P$).

Let $R(n, p)$ and $W(n, p)$ denote, respectively, the (undiscounted) expected revenue and expected surplus from a second price auction with reserve price p. Note that $R(n, p^*) = \mathbb{E}[\max\{MR(v_1), \ldots, MR(v_n), 0\}]$, $R(n, \underline{v}) = \mathbb{E}[\max\{MR(v_1), \ldots, MR(v_n)\}]$, and $W(n, \underline{v}) = \mathbb{E}[\max\{v_1, \ldots, v_n\}]$.

For any $(a, r) \in \{S, P\} \times \{\text{Opt}, \text{Eff}\}$, we can consider the incentive of the auctioneer to delay the auction. For example, a seller using an optimal auction (i.e., (S, Opt)) delays if $R(n, p^*) \leq \delta R(n + d, p^*)$, whereas a planner with an efficient auction (i.e., (P, Eff)) delays if $W(n, \underline{v}) \leq \delta W(n + d, \underline{v})$.[14]

We compare the delay decisions of different auctioneers for different auctions. The following notation captures the fact some auctioneer with one format delays more, in a uniform sense, than another auctioneer with another format.

Definition 1. Take two auctioneers $a, a' \in \{S, P\}$ and two auction formats $r, r' \in \{\text{Opt}, \text{Eff}\}$.

[12] Assumption (i) implies $\mathbb{E}[\max\{v_1, \ldots, v_n\}] \leq n\mathbb{E}[v_i] < +\infty$. Since the i-th order statistic is always lower than the highest order statistic, this inequality implies that all the order statistics have finite moments.

[13] Any reserve price $r \leq \underline{v}$ leads to an efficient auction. We focus on $r = \underline{v}$ for simplicity, because this allows us to write the expected revenue from an efficient auction as the expected marginal revenue of the winner (since the expected payoff of the lowest type ($\underline{v}$) bidder equals zero).

[14] Note that, in our specification, the planner discounts the value of allocating to bidder i from the perspective of $t = 0$, not from the perspective of the time at which the bidder arrived at the market. If the planner discounted a bidder from his time of arrival, both revenue calculations and i's incentives for reporting v_i truthfully would be unaffected. However, this utility function has the unintuitive feature that the planner would privilege "young" bidders at the expense of "old" ones. We thank Gabriel Carroll for pointing out this interpretation.

Then, write

$$(a, r) \succeq (a', r')$$

if, for any $(n, d, \delta) \in \mathbb{N}^2 \times (0, 1)$, whenever a weakly prefers to hold an auction r in $t = 1$, a' weakly prefers to hold an auction r' in $t = 1$. If $(a, r) \succeq (a', r')$ and $(a', r') \succeq (a, r)$, then write $(a, r) = (a', r')$.

4 Right-Skewness and Comparative Statics

To study the incentives to delay an auction, we use the following stochastic order that captures right-skewness of distributions:

Definition 2 (Star-Shaped Order)**.** Take two random variables X and Y. We say Y is greater in the **star-shaped order** if there exist non-negative, increasing functions g and h and a third random variable Z such that $X \sim g(Z)$, $Y \sim h(Z)$, and for any $c > 0$, $h - cg$ crosses 0 at most once and from below.

To understand Definition 2 and why it captures right-skewness, suppose h and g are strictly positive.[15] In that case, $h - cg$ crosses 0 once and from below for all $c > 0$ iff h/g is increasing. Whenever h/g is increasing, $h(Z)$ pushes large draws of Z further out into its right tail (relative to small draws) than $g(Z)$ does, and $h(Z)$ also shrinks small draws of Z towards the left tail more than $g(Z)$ does. This comparison is scale-invariant.

The following crucial lemma connects the star-shaped order to the choice of auction timing.

Lemma 1. *Consider two planners P and P' facing bidders with valuations X_i and X_i', respectively. If X_i is greater in the star-shaped order than X_i' (i.e., is "more right-skewed" then X_i'), then*

$$(P, \text{Eff}) \succeq (P', \text{Eff}).$$

[15] See Barlow and Proschan (1966) for an insightful reference on the star-shaped order and connections to order statistics. The star-shaped order generalizes the convex transformation order, which was first introduced by van Zwet (1964) as a formalization of right-skewness.

The intuition for Lemma 1 is easiest to see when $h(Z) = Z$ in Definition 2, i.e., $X_i = g(X_i')$ with $x \mapsto g(x)/x$ increasing. In this case, adding more bidders (i.e., more draws from the distribution) will raise total surplus (i.e., the sample maximum) proportionally more quickly for planner P, giving her a stronger incentive to delay an auction. For instance, if a new draw exceeds the current sample maximum by 10% under X_i', the sample maximum under $X_i = g(X_i')$ grows by more than 10%. This intuition is phrased realization by realization, but Lemma 1 ensures that the comparison carries over to expected highest draws.

The proof of Lemma 1 requires the following technical result on order statistics:

Lemma 2. *Let $W_A(n-k;n)$ and $W_B(n-k;n)$, $k \in \{0, \dots, n-1\}$, denote the expected $k+1$-th highest draws from n iid samples drawn according to the law of continuously distributed random variables X_i^A and X_i^B, respectively. Suppose $W_A(n-k;n)$ and $W_B(n-k;n)$ are positive for each (n,k). If there exists an iid sequence of random variables $Z_i, i = 1, \dots, n$ and two increasing functions h and g such that*

1. *for any $c \in \mathbb{R}_{++}$, $h - cg$ changes sign at most once and from negative to positive;*
2. *$g \geq 0$; and*
3. *for all $i = 1, \dots, n$, $X_i^A = h(Z_i)$ and $X_i^B = g(Z_i)$,*

then $W_A(n-k;n)/W_B(n-k;n)$ is increasing in n.

Appendix A proves Lemma 2. We use this result to prove Lemma 1:

Proof of Lemma 1. Consider two planners A and B facing bidders with valuations X_i^A and X_i^B, respectively. Suppose that X_i^A is greater than X_i^B in the star-shaped order and that planner B prefers to delay an efficient auction. Borrowing the notation in Lemma 2, we can write this condition as $\delta W_B(n+d;n+d) \geq W_B(n;n)$. Lemma 2 implies $\frac{W_A(n+d;n+d)}{W_B(n+d;n+d)} \geq \frac{W_A(n;n)}{W_B(n;n)}$. It follows that $\delta \frac{W_A(n+d;n+d)}{W_A(n;n)} \geq \delta \frac{W_B(n+d;n+d)}{W_B(n;n)} \geq 1$, and thus $\delta W_A(n+d;n+d) \geq W_A(n;n)$. Therefore, planner A also delays the auction. $\square$

Remark 1. The ranking of delay decisions in Lemma 1 applies uniformly in (n, d, δ). Thus, it carries over, path by path, to the case where planners with a constant discount rate ρ face a bidder pool that varies according to any continuous-time birth process $\{N_t, t \geq 0\}$ and choose an optimal stopping time, adapted to $\{N_t, t \geq 0\}$, at which to hold the auction. Formally, if X_i's are greater in the star-shaped order (more right-skewed) than the X_i''s and bidder values are drawn independently from the arrival process, then for any $\{N_t, t \geq 0\}$ planner P stops later than planner P' (chooses an almost surely larger stopping time). Appendix B provides formal details. In all the results in the paper, it is therefore correct to read our partial order $\succeq$ as an order on stopping times.

We now compare the timing decisions of a seller using an optimal auction to that of a planner using an efficient auction. This is the relevant comparison whenever the auctioneers can choose both the timing and the reserve price so as to maximize their own objectives.[16] Given the number n of bidders, the seller's optimal revenue, $\mathbb{E}[\max_i \max(MR(v_i), 0)]$, is the total surplus from a hypothetical efficient auction in which bidders' values are given by $\max(MR(v), 0)$. Thus, by Lemma 1, whether the seller holds an auction later or earlier than the planner depends on whether the virtual valuation $MR(v)$ is more or less right skewed than the valuation v itself:

Proposition 1. *If $\xi_F(v) := \frac{f(v)}{1-F(v)}v$ is increasing, then $(S, \text{Opt}) \succeq (P, \text{Eff})$. If $\xi_F(v)$ is decreasing, then $(P, \text{Eff}) \succeq (S, \text{Opt})$.*

Proof. Suppose $\xi_F(v)$ is increasing. Then,

$$\frac{\max(MR(v), 0)}{v} = \mathbb{1}_{\{MR(v)) \geq 0\}} \cdot \left(1 - \frac{1 - F(v)}{f(v)v}\right) = \mathbb{1}_{\{\xi_F(v) \geq 1\}} \cdot \left(1 - \frac{1}{\xi_F(v)}\right)$$

is increasing in v, which implies that for any $c > 0$, $\max(MR(v), 0) - cv$ crosses 0 at most once and from below. Lemma 1 then implies $(S, \text{Opt}) \succeq (P, \text{Eff})$. A symmetric argument implies that if $\xi_F(v)$ is decreasing, $(P, \text{Eff}) \succeq (S, \text{Opt})$. $\square$

Whether $\xi_F(v)$ is increasing or decreasing determines whether the revenue from the optimal auction grows proportionally faster or slower than the welfare from an efficient auction.[17]

[16] Given that n and d are exogenous and independent of the reserve price, it is optimal for a seller to use an optimal reserve price and a planner to use the highest efficient reserve price, regardless of timing.

[17] A regular distribution can have an increasing or decreasing $\xi_F(\cdot)$. Section 5.2 of Schweizer and Szech (2019) provide a simple example: Consider the parametric class of distributions $F(x) = 1 - (x + a)^{-b}$ with support $(1 - a, \infty)$ where $a < 1$ and $b > 0$. For any $b \geq 1$, F is regular. Also, $\xi_F(x) = \frac{bx}{x+a}$ is increasing if $a \geq 0$ and decreasing if $a < 0$.

This pins down the relative preferences of the seller and the planner over whether to delay and have more bidders. In fact, our main result provides a complete ranking of delay decisions across auctioneers and reserve prices:

Theorem 1.

1. *If* $\xi_F(v) = \frac{f(v)}{1-F(v)}v$ *is increasing, then*

$$(S, \text{Eff}) \succeq (S, \text{Opt}) \succeq (P, \text{Opt}) \succeq (P, \text{Eff}). \tag{2}$$

2. *If* $\xi_F(v)$ *is decreasing, then*

$$(S, \text{Eff}) = (S, \text{Opt}) \preceq (P, \text{Eff}) = (P, \text{Opt}).$$

Suppose $\xi_F(\cdot)$ is increasing, which holds when F has an increasing hazard rate. Point 1 of Theorem 1 says that for a fixed reserve price $p \in \{\underline{v}, p^*\}$, the seller wants to delay more than the planner, while raising the reserve price from $\underline{v}$ to p^* has opposite effects on planner's and seller's delay decision. Below we highlight the key implications of (2).

1. **Partial regulation can lower social surplus.** When forced to use an efficient auction, the seller waits *more* than when she is free to use an optimal inefficient auction. As shown in Example 1, even for standard distributions, the additional discounting costs can overwhelm the surplus gain from a large bidder pool. Thus, forcing the seller to use an efficient reserve price can lower the discounted total surplus.

2. **Planner should rush inefficient auctions.** When forced to use an optimal auction, the planner should wait less than the seller with that auction.

3. **Thickness is inefficiently high.** An unconstrained planner should wait less to hold her planner-preferred auction (i.e., an efficient one) than the unconstrained seller would wait to hold her seller-preferred auction.

Comparison to Bidder Solicitation Literature Our setup is closely related to static models of bidder solicitation. To highlight our contribution, it is useful to formally compare our results to this prior work (Szech, 2011; Fang and Li, 2015). For brevity, we focus our discussion on the results in Szech (2011). That paper studies an auctioneer who can pay an additive cost c_n to solicit n bidders, where $c_n - c_{n-1}$ is weakly increasing and c_n is strictly so. To compare Szech (2011)'s results to ours, introduce the following order: Let $(a, r) \succeq_{solicit} (a', r')$ if, for any $\{c_n\}_{n\geq}$ satisfying the above property, auctioneer a with reserve price r solicits more bidders than auctioneer a' with reserve price r'. Also, let $\lambda(v) = \frac{f(v)}{1-F(v)}$ denote the hazard rate. Szech (2011) shows that, for $\lambda(v)$ increasing[18]

$$(S, \text{Eff}) \succeq_{solicit} \{(S, \text{Opt}), (P, \text{Eff})\}$$

while for $\lambda(v)$ decreasing,[19]

$$\{(S, \text{Eff}), (P, \text{Eff})\} \succeq_{solicit} (S, \text{Opt}) \text{ and } (P, \text{Eff}) \succeq_{solicit} (S, \text{Eff}).$$

5 Extensions

5.1 Comparative Statics on Value Distribution

Next, we describe how the seller's timing decision responds to changes in the valuation distribution:

Proposition 2. *Let S_F and S_H denote the seller who faces value distribution F and H, respectively. Let p_F^* and p_H^* denote the corresponding optimal reserve prices.*

1. *If F is greater in the star-shaped order (more right-skewed) than H, then $(S_F, \text{Eff}) \succeq (S_H, \text{Eff})$.*

2. *If $p_F^* \geq p_H^*$ and $MR_F(x)/MR_H(x)$ is increasing whenever $x \geq p_F^*$, then $(S_F, \text{Opt}) \succeq (S_H, \text{Opt})$.*

[18] Corollary 2 for (S, Opt) vs (S, Eff) ranking, Proposition 1(ii) for (S, Eff) and (P, Eff) ranking.

[19] (S, Opt) vs (S, Eff) by Corollary 2, (S, Eff) and (P, Eff) by Proposition 1(i), and (S, Opt) vs (P, Eff) by Corollary 3.

The proof of the first point is almost identical to that of Lemma 1: expected revenue is the expected second highest value, so one applies Lemma 2 with $k = 1$ as opposed to $k = 0$. The second point follows from directly from Lemma 1. Thinking about right-skewness also illuminates why the condition $p_F^* \geq p_H^*$ is necessary. Without this additional assumption, if $MR_H(v_i^H) \vee 0$'s atom at p_H^* is sufficiently far to the right of $MR_F(v_i^F) \vee 0$'s atom at p_F^*, $MR_F(v_i^F) \vee 0$ could fail to be more right-skewed than $MR_H(v_i^H) \vee 0$,even though $MR_F(\cdot)$ pushes large draws further into its right tail than $MR_H(\cdot)$.

5.2 Bidder Departures

Lemma 1 and its consequences apply only for $d \geq 0$. Therefore, these "pointwise" theorems imply optimal timing comparative statics results only for stochastic arrival processes that are pathwise increasing. In fact, the results might not extend when bidders can depart stochastically: when $\frac{f(v)}{1-F(v)}v$ is increasing, the seller gains proportionally more than the planner from adding an additional bidder, but this necessarily means that she *loses* proportionally more from losing a marginal bidder. Therefore, when bidders can depart, it is possible that the seller may want to hold the auction sooner than the planner. Nevertheless, Appendix C shows that Theorem 1 extends to a class of non-monotone birth-death processes: bidder arrivals given by certain renewal processes, and bidder departures given by exponentially distributed lifetimes. As a special case, Theorem 1 extends to the situation with Poisson arrivals and departures.

5.3 Impossibility of α-approximation

In the spirit of a large literature on approximation in mechanism design, one might hope that the seller could do reasonably well by restricting herself to simpler mechanisms.[20] One plausible conjecture is that, with full knowledge of the environment, the gain from timing policies that react to the history of arrivals is limited in the following sense: simply holding the auction at an optimal fixed deadline captures a large share of the payoffs from a fully

[20] See, for instance, Hartline and Roughgarden (2009) and Hartline (2012) for approximation results in static auctions.

adaptive policy. In the jargon of approximate mechanism design (Hartline, 2012), do fixed-in-advance deadlines can provide an "α-approximation" to the auction timing problem, i.e., do fixed deadlines always achieve at least a share α of the fully optimal expected discounted revenue, regardless of the distribution of values or arrivals? Alas, the following result shows that the relative loss to the seller from restricting herself to holding the auction at a fixed date (even if the date is set optimally in advance, with full knowledge of the environment) is unbounded. In the worst case over arrival processes, if the seller cannot react to the size of the bidder pool, she will obtain only a vanishing fraction of the revenue generated by an optimal adaptive policy.

We focus revenue-maximization problems where the bidders arrive according to a renewal process $(N_t)_{t\geq 0}$ with inter-arrival distribution G but do not depart. Thus, an auction timing problem is described by a triplet (ρ, F, G) of discount rate ρ, value distribution F, the inter-arrival distribution G.

Now, given (ρ, F, G), let $OPT(\rho, F, G)$ be the ex-ante expected revenue from the optimal timing and auction format policy. Also, let $DET(\rho, F, G)$ be the maximum expected revenue when the seller is restricted to using a deterministic timing policy, i.e., a *deadline*.[21] That is, $DET(\rho, F, G)$ is the maximized value of the problem $\max_{t\in\mathbb{R}_+} e^{-\rho t}\mathbb{E}[R(N_t, p^*)]$. The following result states that, without further restriction on the class of auction timing problems, even the best deterministic stopping rule can waste almost all of the revenue attained by an optimal history-contingent policy.

Theorem 2. *Fix any $\rho > 0$. Then,*

$$\inf_{F\in\mathscr{F}, G\in\mathscr{G}} \frac{DET(\rho, F, G)}{OPT(\rho, F, G)} = 0,$$

where $\mathscr{F}$ is the set of regular distributions, and $\mathscr{G}$ is the set of probability distributions whose supports are subsets of $\mathbb{R}_{++}$.

The proof, in Appendix D, proceeds by explicitly constructing a sequence of examples that attains the infimum. G is supported on a finite grid of points, and along the sequence, all mass is moving to the right end of the grid. If we take a nearly degenerate value distribution (so that revenue from an auction does not really increase by adding additional

[21] This is the class of policies studied by Wang (1993) in a model with exponential G.

bidders beyond the first one), we can then choose grid points so that along the sequence, the optimal deterministic deadline is always to plan on holding the auction at the first time grid point, whereas the fully optimal policy is to wait for the first arrival. Therefore, as we move along the sequence, the best deterministic policy leaves the good unsold with probability approaching one, and achieves a payoff approaching zero. Meanwhile, the optimal history-contingent policy makes sure the good is sold by always waiting for the first bidder, guaranteeing a payoff that is always positive, but far in the future and shrinking along the sequence.

As shown in the proof, along the sequence $\{(F_n, G_n)\}_{n=1}^{\infty}$, both $DET(\rho, F_n, G_n)$ and $OPT(\rho, F_n, G_n)$ converge to zero: that is, by using a fixed (deterministic) deadline, the seller ends up capturing only a vanishing fraction of a vanishing quantity. However, the seller's loss may not be small in absolute terms because for each fixed n, scaling up v would make $OPT - DET$ arbitrarily large. Given that we can vary the loss in absolute terms by manipulating F, Theorem 2 should not be read as a quantitative statement about the size of losses incurred by deterministic deadlines. Rather, we interpret the result in two ways: unless we restrict the primitives of the timing problem, (i) deterministic deadlines are not an appropriate subclass of stopping times to focus on; and (ii), the gain to the seller from being able to access and react to the history of arrivals can be arbitrarily large.

6 Conclusion

We have characterized the inefficiency consequences of revenue maximization in an auction timing model, identifying a condition on the distribution of values that determines wedge between a seller and a social planner's relative incentives to trade off delay costs and market thickness. Our model of the auction market is highly stylized. However, it helps separate static and dynamic distortions from revenue maximization, and it identifies unintended consequences of partial interventions that focus on only one of the distortions. As an example, going back to the case of bankruptcy auctions, the legal literature is mostly concerned with the phenomenon of too few bidders (LoPucki and Doherty, 2007). However, to the extent that senior creditors' private benefits from bankruptcy auctions diverge from social benefits (they care about total proceeds more than efficient reorganization), our results suggest that there is also a risk of *excessive* delay and “too many” bidders.

References

Arefeva, Alina and Delong Meng (2020), “How to set a deadline for auctioning a house.” *Available at SSRN.*

Banciu, Mihai and Prakash Mirchandani (2013), “Technical Note: New Results Concerning Probability Distributions with Increasing Generalized Failure Rates.” *Operations Research*, 61, 925–931.

Barlow, Richard E. and Frank Proschan (1966), “Inequalities for Linear Combinations of Order Statistics from Restricted Families.” *The Annals of Mathematical Statistics*, 37, 1574–1592.

Board, Simon and Andrzej Skrzypacz (2016), “Revenue Management with Forward-Looking Buyers.” *Journal of Political Economy*, 124, 1046–1087.

Boone, Audra L. and J. Harold Mulherin (2009), “Is There One Best Way to Sell a Company? Auctions Versus Negotiations and Controlled Sales.” *Journal of Applied Corporate Finance*, 21, 28–37.

Bulow, Jeremy and Paul Klemperer (2009), “Why Do Sellers (Usually) Prefer Auctions?” *The American Economic Review*, 99, 1544–1575.

Bulow, Jeremy and John Roberts (1989), “The Simple Economics of Optimal Auctions.” *The Journal of Political Economy*, 1060–1090.

Che, Yeon-Koo and Ian Gale (1998), “Standard Auctions with Financially Constrained Bidders.” *The Review of Economic Studies*, 65, 1–21.

Cong, Lin William (2019), “Timing of auctions of real options.” *Management Science*.

Dughmi, Shaddin, Tim Roughgarden, and Mukund Sundararajan (2009), "Revenue Submodularity." In *Proceedings of the 10th ACM conference on Electronic commerce*, 243–252, ACM.

Fang, Rui and Xiaohu Li (2015), "Advertising a Second-Price Auction." *Journal of Mathematical Economics*, 61, 246–252.

Gallien, Jérémie (2006), "Dynamic Mechanism Design for Online Commerce." *Operations Research*, 54, 291–310.

Gershkov, Alex, Benny Moldovanu, and Philipp Strack (2015), "Efficient Dynamic Allocation with Strategic Arrivals." Working Paper.

Gershkov, Alex, Benny Moldovanu, and Philipp Strack (2018), "Revenue- Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand." *Management Science*.

Hartline, Jason D. (2012), "Approximation in Mechanism Design." *The American Economic Review*, 102, 330–336.

Hartline, Jason D. and Tim Roughgarden (2009), "Simple Versus Optimal Mechanisms." In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234, ACM.

Hotchkiss, Edith S and Robert M Mooradian (1998), "Acquisitions as a Means of Restructuring Firms in chapter 11." *Journal of Financial Intermediation*, 7, 240–262.

Jewitt, Ian (1987), "Risk Aversion and the Choice Between Risky Prospects: The Preservation of Comparative Statics Results." *The Review of Economic Studies*, 54, 73–85.

Kleinberg, Robert and Yang Yuan (2013), "On the ratio of revenue to welfare in single-parameter mechanism design." In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 589–602.

Lariviere, Martin A. (2006), "A Note on Probability Distributions with Increasing Generalized Failure Rates." *Operations Research*, 54, 602–604.

Lariviere, Martin A. and Evan L. Porteus (2001), “Selling to the Newsvendor: An Analysis of Price-Only Contracts.” *Manufacturing & Service Operations Management*, 3, 293–305.

Li, Xiaohu (2005), “A Note on Expected Rent in Auction Theory.” *Operations Research Letters*, 33, 531–534.

Loertscher, Simon, Ellen V. Muir, and Peter G. Taylor (2016), “Optimal Market Thickness and Clearing.” Technical report, Working Paper, University of Melbourne, Deparment of Mathematics and Statistics.

LoPucki, Lynn M and Joseph W Doherty (2007), “Bankruptcy Fire Sales.” *Michigan Law Review*, 106, 1.

Mierendorff, Konrad (2016), “Optimal Dynamic Mechanism Design With Deadlines.” *Journal of Economic Theory*, 161, 190–222.

Moldovanu, Benny, Aner Sela, and Xianwen Shi (2008), “Competing Auctions with Endogenous Quantities.” *Journal of Economic Theory*, 141, 1–27.

Myerson, Roger B. (1981), “Optimal Auction Design.” *Mathematics of Operations Research*, 6, 58–73.

Pai, Mallesh M. and Rakesh Vohra (2013), “Optimal Dynamic Auctions and Simple Index Rules.” *Mathematics of Operations Research*, 38, 682–697.

Paul, Anand and Genaro Gutierrez (2004), “Mean sample spacings, sample size and variability in an auction-theoretic framework.” *Operations Research Letters*, 32, 103–108.

Quah, John K. H. and Bruno Strulovici (2013), “Discounting, Values, and Decisions.” *Journal of Political Economy*, 121, 896–939.

Schweizer, Nikolaus and Nora Szech (2019), “Performance bounds for optimal sales mechanisms beyond the monotone hazard rate condition.” *Journal of Mathematical Economics*, 82, 202–213.

Szech, Nora (2011), “Optimal Advertising of Auctions.” *Journal of Economic Theory*, 146, 2596–2607.

van Zwet, W. R. (1964), “Convex Transformations: A New Approach to Skewness and Kurtosis.” *Statistica Neerlandica*, 18, 433–441.

Wang, Ruqu (1993), “Auctions Versus Posted-Price Selling.” *The American Economic Review*, 838–851.

Ziya, Serhan, Hayriye Ayhan, and Robert D Foley (2004), “Relationships among three assumptions in revenue management.” *Operations Research*, 52, 804–809.

Appendix

A Comparative Statics on Auction Timing: Proofs for Section 4

To prove Lemma 2, we use the following result by Barlow and Proschan (1966).[22]

Lemma 3 (From Lemma 3.5 in (Barlow and Proschan, 1966)). *Let $X^{(1:n)} \leq X^{(2:n)} \leq \cdots \leq X^{(n:n)}$ be an ordered iid sample from a continuous distribution. Suppose a function $\gamma : \mathbb{R} \to \mathbb{R}$ changes sign k times, and that the expectation of $\gamma(X^{(i:n)})$ exists. Then $\pi(n-i,n) = \mathbb{E}[\gamma(X^{(n-i:n)})]$ changes sign at most k times as a function of n; if $\pi(n-i,n)$ changes sign in n exactly k times, then the changes occur in the same order as those of $\gamma(x)$.*

We now prove Lemma 2.

Proof of Lemma 2. The proof is a slight generalization of Theorem 3.6 in Barlow and Proschan (1966). From Lemma 3 with $\gamma(z) = h(z) - cg(z)$, we have that

$$\begin{aligned}
\mathbb{E}[h(Z^{(n-k:n)})] - c\mathbb{E}[g(Z^{(n-k:n)})] &= W_A(n-k;n) - cW_B(n-k;n) \\
&= W_B(n-k;n)\left(\frac{W_A(n-k;n)}{W_B(n-k;n)} - c\right)
\end{aligned}$$

changes sign at most once and from negative to positive, as a function of n (the first equality uses the fact that h and g are non-decreasing). Because $W_B(n-k;n) > 0$ and this holds for any $c > 0$, we conclude that $W_A(n-k;n)/W_B(n-k;n)$ is increasing in n. $\square$

Proof of Theorem 1. First, suppose $\xi_F(v)$ is increasing.
$(S, \text{Eff}) \succeq (S, \text{Opt})$: Take n iid draws of the marginal revenue, $X_1, \ldots, X_n$, where $X_i = MR(v_i)$. For each $x \in \overline{\mathbb{R}}$ and $p \in \overline{\mathbb{R}}$, define $g_p(x) = \mathbb{E}[\max\{X_1, x, p\}]$, where the expectation is with respect to X_1. It holds that $g_{-\infty}(-\infty) = R(1, \underline{v})$ and $g_0(-\infty) = R(1, p^*)$. Also, we have $\mathbb{E}[g_{-\infty}(X^{(n-1:n-1)})] = R(n, \underline{v})$ and $\mathbb{E}[g_0(X^{(n-1:n-1)})] = R(n, p^*)$.

Note that $g_{-\infty}(x)/g_0(x)$ is increasing in x.[23] Thus, we have $\frac{g_{-\infty}(-\infty)}{g_0(-\infty)} \leq \frac{g_{-\infty}(X_2)}{g_0(X_2)}$ for any realization of X_2, which implies $\frac{g_{-\infty}(-\infty)}{g_0(-\infty)} \leq \frac{\mathbb{E}[g_{-\infty}(X_2)]}{\mathbb{E}[g_0(X_2)]}$. As a result, we obtain $\frac{R(2,p^*)}{R(1,p^*)} \leq \frac{R(2,\underline{v})}{R(1,\underline{v})}$.

[22] In the economics literature, the variation diminishing property of totally positive kernels, on which this Lemma depends, is used, e.g., in Jewitt (1987) and Bulow and Klemperer (2009).

[23] Indeed, $\frac{g_{-\infty}(x)}{g_0(x)}$ is strictly increasing for $x \leq 1$ and equal to 1 for $x \geq 1$.

Next, take any $n \geq 3$. Because $g_{-\infty}(x)/g_0(x)$ is increasing, for each $c > 0$, $g_{-\infty}(x) - cg_0(x)$ crosses 0 at most once and from below. Also, because $g_p(x)$ is increasing, we have $g_p(X^{(k:k)}) = \max_{k=1,\dots,n} g_p(X_k)$ for any p and k. Then Lemma 2 implies $\frac{\mathbb{E}[g_{-\infty}(X^{(n-2:n-2)})]}{\mathbb{E}[g_0(X^{(n-2:n-2)})]} \leq \frac{\mathbb{E}[g_{-\infty}(X^{(n-1:n-1)})]}{\mathbb{E}[g_0(X^{(n-1:n-1)})]}$, which is equivalent to $\frac{R(n,p^*)}{R(n-1,p^*)} \leq \frac{R(n,\underline{v})}{R(n-1,\underline{v})}$. To sum up, we have $\frac{R(n+1,p^*)}{R(n,p^*)} \leq \frac{R(n+1,\underline{v})}{R(n,\underline{v})}$ for any $n \in \mathbb{N}$. By the same argument used in Lemma 1, we have $(S, \text{Eff}) \succeq (S, \text{Opt})$.

$(S, \text{Opt}) \succeq (P, \text{Opt})$: Note that

$$R(n, p^*) = \mathbb{E}[\max_i MR(v_i)\mathbb{1}_{\{MR(v_i)\geq 0\}}] \quad \text{and} \quad W(n, p^*) = \mathbb{E}[\max_i v_i \mathbb{1}_{\{MR(v_i)\geq 0\}}].$$

Here, $W(n, p^*)$ is the expected total surplus in the second price auction with a revenue-optimal reserve price p^*. We show that for any $c > 0$, $MR(v)\mathbb{1}_{\{MR(v)\geq 0\}} - v\mathbb{1}_{\{MR(v)\geq 0\}}$ changes sign at most once from negative to positive. Take any v such that $MR(v)\mathbb{1}_{\{MR(v)\geq 0\}} - cv\mathbb{1}_{\{MR(v)\geq 0\}} > 0$. This implies $\mathbb{1}_{\{MR(v)\geq 0\}} = 1$, and thus $\mathbb{1}_{\{MR(v')\geq 0\}} = 1$ for all $v' > v$. We show $MR(v') - cv' > 0$. Suppose to the contrary that $MR(v') - cv' \leq 0$. Then, $\frac{MR(v')}{v'} \leq c$. As $\lambda(x)x$ is increasing, $\left(x - \frac{1}{\lambda(x)}\right) x^{-1} = 1 - (\lambda(x)x)^{-1}$ is increasing. Thus, $\frac{MR(v)}{v} \leq c$, which is a contradiction. It follows that $MR(v)\mathbb{1}_{\{MR(v)\geq 0\}} - v\mathbb{1}_{\{MR(v)\geq 0\}}$ changes sign at most once from negative to positive. Lemmata 1 and 2 imply $(S, \text{Opt}) \succeq (P, \text{Opt})$.

$(P, \text{Opt}) \succeq (P, \text{Eff})$: We have $W(n, p^*) = \mathbb{E}[\max_i v_i \mathbf{1}_{\{MR(v_i)\geq 0\}}]$ and $W(n, \underline{v}) = \mathbb{E}[\max_i v_i]$. If $v\mathbb{1}_{\{MR(v)\geq 0\}} - cv > 0$, then $\mathbb{1}_{\{MR(v_i)\geq 0\}} = 1$ and $1 > c$, and thus $v'\mathbb{1}_{\{MR(v')\geq 0\}} - cv' > 0$ for any $v' > v$. By the same logic as before, $(P, \text{Opt}) \succeq (P, \text{Eff})$.

Next, suppose $\xi_F(v)$ is decreasing. We show that the optimal auction is efficient, i.e., $p^* = \underline{v}$. It suffices to show that $v - \frac{1-F(v)}{f(v)} \geq 0$ for all $v \geq \underline{v}$. Suppose to the contrary that $v - \frac{1-F(v)}{f(v)} = v(1 - \frac{1}{\xi_F(v)}) < 0$ for some v. Then, since $v \geq 0$, $1 - \frac{1}{\xi_F(v)} < 0$ holds. Because $\xi_F(v)$ is decreasing, we obtain

$$v'\left(1 - \frac{1}{\xi_F(v')}\right) < v\left(1 - \frac{1}{\xi_F(v)}\right).$$

if $v' > v$. This contradicts the regularity of F.

The above result implies that $(S, \text{Opt}) = (S, \text{Eff})$ and $(P, \text{Opt}) = (P, \text{Eff})$. Finally, $(P, \text{Eff}) \succeq (S, \text{Eff})$ follows from by arguments symmetric to the increasing $\xi_F(v)$ case. $\square$

B Comparative Statics Under General Arrival Processes

In the main text, we mentioned that Lemma 1 and Theorem 1 could be reinterpreted as monotone comparative statics for the following optimal stopping problem: an auctioneer a with fixed format r observes a stochastic process $\{N_t, t \geq 0\}$ for bidder arrivals, and she must choose a stopping time adapted to $\{N_t, t \geq 0\}$ at which to hold the auction. We now formalize this idea.

Recall that $R(n, p)$ and $W(n, p)$ respectively denote the expected revenue and total surplus from the second price auction with reserve price p. Also, p^* is the revenue-optimal reserve price. By iterated expectations, the auction timing problem for $(a, r) = (S, \text{Opt})$ is

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-\rho\tau} R(N_\tau, p^*)\right],$$

and the one for (P, Eff) is

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-\rho\tau} W(N_\tau, \underline{v})\right],$$

where $\mathcal{T}$ is the set of stopping times adapted to $\mathcal{F}^N$, the filtration generated by $\{N_t, t \geq 0\}$. The stopping problems for (S, Eff) and (P, Opt) are defined similarly.

Our comparative statics results can be stated in terms of strong set order comparisons on sets of stopping times with a natural lattice structure. However, to simplify the exposition, we follow Quah and Strulovici (2013) in assuming that an auctioneer stops the first time she is indifferent between holding the auction and continuing to search for bidders. That is, we focus on the *essential infimum* of the solutions to the relevant optimal stopping problem.[24]

Definition 3. Take two auctioneers $a, a' \in \{S, P\}$ and two auction formats $r, r' \in \{\text{Opt}, \text{Eff}\}$. Let $\mathcal{T}^*_{a,r}$ and $\mathcal{T}^*_{a',r'}$ be the solution sets over $\mathcal{T}$ to their respective optimal stopping problems. Then, write

$$(a, r) \succeq_{stop} (a', r')$$

if, for any pathwise increasing $\{N_t, t \geq 0\}$, $\operatorname{ess\,inf} \mathcal{T}^*_{a,r} \geq \operatorname{ess\,inf} \mathcal{T}^*_{a',r'}$. If $(a, r) \succeq_{stop} (a', r')$ and $(a', r') \succeq_{stop} (a, r)$, then write $(a, r) =_{stop} (a', r')$.

[24] As in Quah and Strulovici (2013), the assumptions on $\{N_t, t \geq 0\}$ and $\mathcal{F}^N$ ensure that the essential infimum of solutions to the stopping problem is itself a solution to the stopping problem.

It suffices to extend Lemma 1 to this setting:

Lemma 4. *Let $\{N_t, t \geq 0\}$ be any birth process for bidder arrivals. Consider two planners A and B facing bidders with valuations X_i and X_i', respectively. If X_i is greater in the star-shaped order (more right-skewed) than X_i', then*

$$(A, \text{Eff}) \succeq_{stop} (B, \text{Eff})$$

Proof. To simplify exposition, let $S_A(n)$ and $S_B(n)$ respectively denote the expected surplus for planner A and B in the efficient auction with n bidders. Let τ_A and τ_B be the essential infima over optimal $\mathcal{F}^N$-stopping times for A and B's stopping problems. To derive a contradiction, suppose that $\tau_A < \tau_B$ with a positive probability. The inequalities below are evaluated on this event. Since τ_A is optimal for A,

$$S_A(N_{\tau_A}) \geq \mathbb{E}[e^{-\rho\tau} S_A(N_{\tau_A+\tau}) | \mathcal{F}_{\tau_A}]$$

for any feasible stopping time τ such that $\{\tau = 0\}$ is $\mathcal{F}^N_{\tau_A}$ measurable. In particular, this inequality must hold for $\tau_{B|A}$, B's optimal continuation stopping time starting at $\mathcal{F}^N_{\tau_A}$ (since $\tau_B > \tau_A$, $\{\tau_{B|A} = 0\}$ is $\mathcal{F}_{\tau_A}$-measurable). Therefore,

$$1 \geq \mathbb{E}\Big[e^{-\rho\tau_{B|A}} \frac{S_A(N_{\tau_A+\tau_{B|A}})}{S_A(N_{\tau_A})} \Big| \mathcal{F}^N_{\tau_A}\Big] \geq \mathbb{E}\Big[e^{-\rho\tau_{B|A}} \frac{S_B(N_{\tau_A+\tau_{B|A}})}{S_B(N_{\tau_A})} \Big| \mathcal{F}^N_{\tau_A}\Big],$$

where the second inequality follows from $N_{\tau_A+\tau_{B|A}} \geq N_{\tau_A}$ (since there are no departures) and Lemma 1 (applied pathwise inside the expectations). Altogether,

$$S_B(\tau_A) \geq \mathbb{E}[e^{-\rho\tau_{B|A}} S_B(N_{\tau_A+\tau_{B|A}}) | \mathcal{F}^N_{\tau_A}],$$

so $\tau_{B|A} = 0$ is an optimal continuation policy for B at τ_A, which contradicts $\tau_B > \tau_A$. □

By Lemma 4, Theorem 1 continues to hold with $\succeq$ replaced by $\succeq_{stop}$.

C Auction Timing Under Stochastic Departures

This section consists of two parts. First, we extend the model to incorporate stochastic departure of bidders, and we state a version of Theorem 1 for this richer setting. The proof of this new result requires a partial characterization of the optimal timing policy in a class of arrival-departure processes. We prove that characterization in a second section.

Setting With Bidder Departures To circumvent the issues described in section 5.2, we impose some structure on the arrival-departure process $\mathcal{N} = \{N_t, t \geq 0\}$. First, bidders arrive according to a *renewal process*: the time between the arrival of the $k-1$-th and k-th bidders, $k = 1, 2, \ldots$, is given by a non-negative random variable $W_k \overset{iid}{\sim} G$. We assume that $\mathbb{E}[e^{-\rho W_k} | W_k \geq t]$ is well-defined for any $t \geq 0$. Second, each bidder leaves after a random time in the pool: each bidder i has a sojourn time D_i in the pool, so that if i arrives at time a_i, he departs at calendar time $a_i + D_i$. Sojourn times D_i are drawn iid from an exponential distribution, independently of all values and arrival times. Bidder i and the auctioneer observe a_i, and at time $t \geq a_i$, they both know whether $t \leq a_i + D_i$, but neither can observe D_i. In other words, both i and the auctioneer know when i arrives and whether he is still available, but neither knows how much longer i might be available for. Remark 3 discusses how our results extend to the case in which i can misreport his arrival and availability. For example, D_i could represent the sudden arrival of a better outside option for a bidder, or, in the corporate takeover example, it could model the occurrence of an unexpected negative shock to liquidity that prevents the bidder from taking part in the auction.

Let $\mathbb{L}_t$ denote the time elapsed since the most recent arrival at calendar time t, so that the last arrival occurred at $t - \mathbb{L}_t$. By the renewals assumption, the auctioneer only needs to keep track of $\mathbb{L}_t$ in order to predict future arrival times. Moreover, because the D_i's are exponential, by the memoryless property, current bidders' time in the market is irrelevant for predicting their departure times.[25] Therefore, to know the future law of motion for $\mathcal{N}$ at t, the auctioneer only needs to know N_t and $\mathbb{L}_t$. Since bidder values are iid and independent of the arrivals process, the auctioneer thus faces a Markov optimal stopping problem with state variable $(N_t, \mathbb{L}_t)$ (likewise for the planner).

To characterize the optimal stopping policy, we will require the following condition on the inter-arrival-time distribution:

[25] If we consider a general distribution for D_i, then the auctioneer needs to keep track of the time that each current bidder has spent in the market: that time affects her beliefs about how likely the bidder is to leave in the next dt units of time, which in turn would affect the payoffs from stopping in the next dt. In this case, the auctioneer has a large and growing state space in her optimal stopping problem, which makes further analysis intractable.

Definition 4. G is "New Better than Used" (NBU) if, for any $a \sim G$, $t > 0$, we have $a \geq_{FOSD} (a-t)|\{a \geq t\}$.

NBU is a generalization of the increasing hazard rate (IHR) property: under NBU, an auctioneer is more optimistic about arrivals in between arrivals than right after an arrival, while under IHR, the auctioneer gets increasingly optimistic about arrivals the longer it has been since the last arrival.

For this richer setting, we state a "uniform" order on delay decisions analogous to $\succeq$ and $\succeq_{stop}$ in Definitions and 1 and 3:

Definition 5. Let $\mathcal{T}^0_{P,\text{Eff}}$ and $\mathcal{T}^0_{P,\text{Opt}}$ denote the planner's optimal stopping times in

$$\sup_{\tau \in \mathcal{T}} \quad \mathbb{E}\left[e^{-\rho\tau} W(N_\tau, \underline{v})\right],$$
$$N_0 = 0,$$

and

$$\sup_{\tau \in \mathcal{T}} \quad \mathbb{E}\left[e^{-\rho\tau} W(N_\tau, p^*)\right],$$
$$N_0 = 0,$$

respectively, where $\mathcal{T}$ is the set of stopping times adapted to $\mathcal{F}^N$. Define $\mathcal{T}^0_{S,\text{Eff}}$, and $\mathcal{T}^0_{S,\text{Opt}}$ for the seller similarly, with W replaced by R. Then, for two auctioneers $a, a' \in \{S, P\}$ and auction formats $r, r' \in \{\text{Opt}, \text{Eff}\}$, write

$$(a, r) \succeq_d (a', r')$$

if, for any $\{N_t, t \geq 0\}$ generated by G satisfying NBU and D_i that is iid exponential,

$$\operatorname{ess\,inf} \mathcal{T}^0_{a,r} \geq \operatorname{ess\,inf} \mathcal{T}^0_{a',r'}.$$

If $(a, r) \succeq_d (a', r')$ and $(a', r') \succeq_d (a, r)$, then write $(a, r) =_d (a', r')$.

That is $(a, r) \succeq_d (a', r')$ if a always holds auction r later than a' holds r', *when both start from an empty bidder pool.*

Our main result for this setting is as follows:

Theorem 3. *All the comparative statics results in Theorem 1 hold with $\succeq_d$ in place of $\succeq$.*

Theorem 3 relies on an explicit characterization of the optimal stopping region. To state that auxiliary result, let $\tau(n, n+1) := \inf\{t \geq 0 : N_t = n+1\} - \inf\{t \geq 0 : N_t = n\}$ denote the random time between when the pool first reaches size n and when it first reaches $n+1$. (The possibility of departures makes the distribution of $\tau(n, n+1)$ different from G.) We then show

Theorem 4. *Let $Q(\cdot)$ be any one of $W(\cdot, \underline{v})$, $W(\cdot, p^*)$, $R(\cdot, \underline{v})$ and $R(\cdot, p^*)$. Suppose that G is NBU and D_i is drawn iid from an exponential distribution, and let $\mathcal{T}_Q^0$ denote the solution to*

$$\sup_{\tau \in \mathcal{T}} \quad \mathbb{E}\left[e^{-\rho\tau} Q(N_\tau)\right],$$

$$N_0 = 0.$$

Let $\beta(n+1) := \mathbb{E}[e^{-\rho\tau(n,n+1)}]$ denote the (unconditional) expected discount factor between the first time the bidder pool reaches size n and the first time it reaches size $n+1$, and let

$$n^* := \inf\left\{n \in \mathbb{N} : 1 \geq \beta(n+1)\frac{Q(n+1)}{Q(n)}\right\}. \tag{3}$$

Then

$$\operatorname{ess\,inf} \mathcal{T}_Q^0 = \inf\{t : N_t \geq n^*\}.$$

The one-step characterization of n^* in (3) has an especially tractable structure. The critical threshold function $\beta(n+1)\frac{Q(n+1)}{Q(n)}$ decouples into a term that depends only on the physical properties of the arrival-departure process ($\beta(n+1)$), and a term that depends only on the proportional gain from having an additional bidder in the pool ($\frac{Q(n+1)}{Q(n)}$). Without additional structure on G, we have not been able to solve for the optimal timing policy for all initial states (n, w). However, we can fully characterize the optimal policy for all sufficiently small initial bidder pools, and as we show below, all our comparative statics results from Section 4 extend to this case.[26]

[26] By the usual dynamic programming arguments, the policy in Theorem 4 must be time-consistent. Hence, starting from any (n, w) that is reachable from $(0, 0)$ before stopping, stopping the first time n^* is reached must also be optimal.

Using the characterization in Theorem 4, we can one again apply the machinery of Lemma 1 to obtain comparative statics results:

Proof of Theorem 3. Let $S_A(n)$ and $S_B(n)$ denote the expected surplus for planner A and B in the efficient auction with n bidders, respectively. Let n_A^* and n_B^* denote the cutoff in (3) for the corresponding planner and auction. From the expression in (3), we can see that, starting from an empty pool, planner A waits for a larger number of bidders than planner B if $\frac{S_A(n+1)}{S_A(n)} \geq \frac{S_B(n+1)}{S_B(n)}$ for all n, because any n that satisfies the inequality in the infimum defining n_A^* will satisfy the inequality defining n_B^*. Hence, if the value distribution faced by planner A is more right-skewed (greater in the star-shaped order) than that faced by B, she will stop later. The same argument as in Theorem 1 then establishes the result. □

Remark 2 (Role of Distributional Assumptions)**.** Here, we explain the role of the distributional assumptions used for Theorems 3 and 4. First, when G is NBU, continued search is least attractive right after an arrival, since at that point, the seller expects to wait the longest for an additional arrival. This leads the seller to stop only at arrivals, at least for small enough initial pools. Second, the two assumptions of renewal-process arrivals and exponential departures imply that, for each n, the evolution of N_t after first reaching a new record n is independent of all prior history (other than the fact that n was reached). This memorylessness allows us to decompose the waiting time for n bidders into independent one-step increments between successive records, which leads to the one-step characterization of the optimal stopping threshold.

Remark 3 (Strategic Arrivals and Departures)**.** Throughout the paper, we have assumed that the seller can observe the arrivals and availability of bidders. However, this assumption might not always be innocuous. For instance, a sophisticated bidder in an auction for corporate assets may pretend that getting financial backing for a bid takes longer than is actually necessary, if it thinks that by delaying it can place the seller in a disadvantageous position. Whether or not the optimal timing policy is implementable when arrivals and availability are privately known will depend on the form of that policy, and therefore indirectly on the law of motion for the bidder pool.[27]

[27] For example, Gershkov, Moldovanu, and Strack (2015) show that an efficient dynamic mechanism with

To study this possibility, assume i privately observes his arrival time a_i, and only he knows at t whether $a_i + D_i \leq t$. (So even though D_i is unobserved, i observes whether or not "his time is up.") Therefore, i can delay his arrival until any $a_i' \geq a_i$ such that his time is not up at a_i', and conditional on reporting an arrival at a_i', he can pretend to have left at any $t \geq a_i'$ if his time is not up by t.[28]

For the class of arrival-departure processes we have considered in this section, bidders will in fact want to report their arrivals and availability truthfully. The proof adapts an argument by Gershkov et al. (2015) to the case of auction timing. Let N_t' denote the number of bidders observed by the auctioneer at t when bidder i misreports his arrival or departure time. Note that, however i chooses to strategically time his arrival and/or departure, $N_t' \leq N_t$ always holds. In particular, N_t' never hits n^* strictly earlier than N_t. Thus, bidder i either (i) misses the opportunity to participate in the auction and obtains a payoff of zero, or (ii) joins an n^*-bidder auction that will be held later than if i had truthfully reported his arrival and departure times. Bidder i does not benefit from the deviation in (i), because the payoff from the auction is non-negative. Nor does he benefit in the case (ii), because he must incur the cost of additional discounting without changing the number of opposing bidders he faces in the auction. Therefore, the cut-off policy in Theorem 4 is robust to bidders' strategic incentive to "misreport" the timing of their arrival and departure.[29]

C.1 Characterizing the Optimal Timing Policy

To prove Theorem 4, we require four technical lemmas. The first one is a preliminary fact regarding payoffs upon stopping.

Lemma 5. *Let $Q(\cdot)$ be any one of $W(\cdot), R(\cdot, p^*)$ and $R(\cdot, 0)$. Then, $Q(n)$ is*

observable arrivals may not be implementable if later arrivals make the seller pessimistic about the time of future arrivals.

[28] Similar to the literature on auctions with private budget constraints (Che and Gale, 1998), we assume that agents can only engage in one-sided deviations in reporting their arrival and departure status.

[29] Here, we show implementability by using specific properties of the stopping region derived from the NBU assumption. However, an almost identical argument shows that bidders will report arrivals truthfully whenever arrivals are given by an arbitrary renewal process, in which case the shape of the stopping region is *a priori* unclear.

1. *non-negative and increasing in n;*

2. *concave in n, i.e., $Q(n+1) - Q(n) \leq Q(n) - Q(n-1), \forall n \in \mathbb{N}$;*

3. *their one-step ratios are decreasing, i.e., $\frac{Q(n)}{Q(n-1)}$ is decreasing in $n \in \mathbb{N}$.*[30]

Proof. Point 1 is obvious for $W(n, \underline{v})$, since it equals $\mathbb{E}[v^{(n)}]$, where $v^{(n)}$ denotes the highest draw out of n samples from F. Regularity of F (i.e., increasing $MR(\cdot)$) ensures that $\mathbb{E}(\max\{MR(v^{(n)}), x\})$ is increasing in n, and thus $R(n, p)$ is increasing for any p. To see that $\mathbb{E}(\max\{MR(v^{(n)}), x\})$ is non-negative, note that $\max\{MR(v^{(n)}), x\} \geq MR(v_1)$ and $E[MR(v_1)] = 0$. To show Point 2, take subsets S, T of the set of bidders B with $S \supsetneq T$, and a bidder $i \in B \setminus S$. Then, we claim

$$\max_{j \in S \cup \{i\}} v_j - \max_{j \in S} v_j \leq \max_{j \in T \cup \{i\}} v_j - \max_{j \in T} v_j \tag{4}$$

where $\max_{j \in T} v_j = 0$ if $T = \emptyset$. Indeed, when $v_i = \max_{j \in S \cup \{i\}} v_j$, the comparison becomes $-\max_{j \in S} v_j \leq -\max_{j \in T} v_j$, which holds by the assumption that $T \subset S$; when $v_i < \max_{j \in S \cup \{i\}}$, the left hand side is 0, and the right hand side is non-negative. Taking expectations on both sides of (4) delivers the concavity of $W(n, \underline{v})$.[31] The result extends to $R(n, \underline{v})$ and $R(n, p^*)$ by substituting marginal revenues $MR(v_j)$ and $MR(v_j) \vee 0$, respectively, for v_j.[32] To show Point 3, note that diminishing returns of $n \mapsto Q(n)$ imply

$$Q(n+1)\left(1 - \frac{Q(n)}{Q(n+1)}\right) \leq Q(n)\left(1 - \frac{Q(n-1)}{Q(n)}\right) \tag{5}$$

By revenue monotonicity, the terms in the parentheses are non-negative, so for (5) to hold even though $Q(n+1) \geq Q(n)$, $\frac{Q(n)}{Q(n-1)}$ must be decreasing, as required. $\square$

The following lemma presents a useful property of the auctioneer's value function. Let $V(n, w)$ denote the value function at state $(N_t, \mathbb{L}_t) = (n, w)$.

Lemma 6. *For any n, $a > 0$, $V(n, a) \geq V(n, 0)$.*

30 $\frac{x}{0} = \infty$.

31 The proof strategy follows Corollary 2.7. in Dughmi, Roughgarden, and Sundararajan (2009).

32 We sometimes use the shorthand $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$, for notational simplicity.

Proof. For any $x \geq 0$, define $W^x := W - x | \{W \geq x\}$. Fix some $a > 0$. First, note that by the NBU assumption, $W^0 \geq_{FOSD} W^a$, so we can always find a probability space such that W^a and W^0 are both defined on that space and $W^0 \geq W^a$ almost surely.

With that in mind, consider two stochastic processes defined on a common probability space on which $W^0 \geq W^a$ almost surely. First, define N_a as follows.

- At $t = 0$, initialize $N_a(0) = n$. Also, start n independent exponential departure clocks simultaneously.
- If arrival clock W^a "ticks," increase N_a by 1 and draw a new arrival clock $W_1 \sim G$ and an exponential departure clock D_1.
- For each $k \geq 1$, if arrival clock W_k "ticks," increase N_a by 1 and draw a new arrival clock $W_{k+1} \sim G$ and an exponential departure clock D_{k+1}.
- If any departure clock ticks, decrease N_a by 1.

Second, define N_0 in the same way as N_a except that we use W^0 instead of W^a for the first arrival clock. Note that the values from following the policies that are optimal for N_a and N_0 are $V(n, a)$ and $V(n, 0)$, respectively, but given the construction of N_a and N_0 on a common probability space, we can compare $V(n, a)$ and $V(n, 0)$ realization by realization.

We prove the result by induction on n. First, as $W^0 \geq W^a$ for sure, we get

$$V(0, a) = \mathbb{E}[e^{-\rho W^a}]V(1, 0) \geq \mathbb{E}[e^{-\rho W^0}]V(1, 0) = V(0, 0).$$

Second, apply the induction hypothesis that $V(k, a) \geq V(k, 0)$, $k = 0, \ldots, n-1$. We refer to the auctioneers who confront N_a and N_0 as Auctioneers a and 0, respectively. Likewise, we refer the optimal policies of Auctioneers a and 0 as Policies a and 0, respectively. Suppose that Auctioneer 0 takes Policy 0, which yields $V(n, 0)$. Meanwhile, Auctioneer a uses the following "hybrid" timing policy.

1. As long as no bidders arrive or depart, adopt Policy 0.
2. If some bidder arrives or departs while imitating Auctioneer 0, from that point on adopt Policy a.

We show that this policy gives the Auctioneer a a greater continuation value than $V(n,0)$ for any realizations of the stochastic processes. Under the first scenario in this hybrid policy, if Auctioneer a holds an auction (say, at time t) while imitating Auctioneer 0, then both of them obtain a discounted payoff of $e^{-\rho t}Q(n)$ because $W^0 \geq W^a$. Under the second scenario, there are two cases. If Auctioneer a switches to Policy a at t because a bidder has departed, then the time-t continuation payoffs of Auctioneer's 0 and a are $V(n-1,t)$ and $V(n-1,a+t)$, respectively. By the inductive hypothesis, $V(n-1,a+t) \geq V(n-1,t)$. Finally, if Auctioneer a switches to policy a because a new bidder has arrived at time W^a, then the time-W^a continuation payoff of Auctioneer a is $V(n+1,0)$, while that of Auctioneer 0 is $V(n,W^a) \leq V(n+1,0)$.[33] □

Lemma 6 implies that, when starting from state $(0,0)$, the optimal policy stops the first time the bidder pool reaches some n^*. To complete the proof of Theorem 4, it remains to prove the auxiliary results used in the one-step characterization of n^* in (3). Recall that $\beta(j) := \mathbb{E}[e^{-\rho\tau(j-1,j)}]$, where $\tau(j-1,j)$ is the random time elapsed between when the pool size first reaches $j-1$ to when it first reaches j. Then we have the following payoff decomposition:

Lemma 7. *The auctioneer's expected discounted payoff at time* 0 *from holding the auction upon the n-th bidder's arrival is given by $\prod_{j=1}^{n}\beta(j)Q(n)$.*

Proof. Let $\tau_n = \inf\{t : N_t = n\}$. Then the expected discounted payoff from stopping at the n-th bidder's arrival becomes

$$\mathbb{E}[e^{-\rho\tau_n}]Q(n) = \mathbb{E}[e^{-\rho\sum_{j=1}^{n}\tau(j-1,j)}]Q(n). \tag{6}$$

To calculate (6), construct the following "fictitious" process $\mathbf{N}' = \{N'_t, t \geq 0\}$:

- Initialize $N'_0 = 0$.

[33] It must be that $V(n+1,0) \geq V(n,t)$ for any t. Indeed, starting from (n,t), let q be the discounted probability that, under the optimal policy, the state reaches $(n+1,0)$ before the auctioneer holds an auction, and let q' be the discounted probability, under the optimal policy, that the seller holds the auction before the state reaches $(n+1,0)$, i.e., the auction takes place with at most n bidders. Clearly, $q + q' \leq 1$, so $V(n,t) = qV(n+1,0) + q'Q(n) \leq V(n+1,0)$.

- Draw arrival clocks sequentially: first draw $W_1 \sim G$ at $t = 0$, then when the first clock ticks ($t = W_1$) draw $W_2 \sim G$ independently at $t = W_1$, and so on.
- If an arrival clock W_j "ticks":
  1. Increase $\mathbf{N}'$ by 1.
  2. Remove (ignore thereafter) all remaining departure clocks, and replace them with the same number of new, independent departure clocks.
  3. Add an additional exponential departure clock, independent of all arrival clocks.
- If any existing departure clock ticks:
  1. Decrease $\mathbf{N}'$ by 1.
  2. Remove (ignore thereafter) all remaining departure clocks, and replace them with the same number of new, independent exponential departure clocks.

By the memoryless property of the Poisson clocks, $\mathbf{N}$ and $\mathbf{N}'$ will have the same marginal distributions, so letting $\tau'(j-1,j)$ denote the successive record times for N', $\tau'(j-1,j) \sim \tau(j-1,j)$ and $e^{-r\sum_{j=1}^{n}\tau(j-1,j)} \sim e^{-r\sum_{j=1}^{n}\tau'(j-1,j)}$. In addition, N''s first-increment times $\{\tau(k-1,k)\}_{k\in\mathbb{N}}$ will be mutually independent, since the increasing part of $\mathbf{N}$ is a renewal, and by the way we construct $\mathbf{N}'$ (replacing all "old" clocks with fresh independent ones at each point of change), all the dependence in between successive records of $\mathbf{N}'$ has been removed. Therefore

$$\mathbb{E}[e^{-r\sum_{j=1}^{n}\tau(j-1,j)}] = \mathbb{E}[e^{-r\sum_{j=1}^{n}\tau'(j-1,j)}] = \Pi_{j=1}^{n}\mathbb{E}[e^{-\rho\tau'(j-1,j)}] = \Pi_{j=1}^{n}\beta(j),$$

as required. □

Finally, we show that $\beta(\cdot)$ in 4 is decreasing.

Lemma 8. *Let $\tau(n-1,n)$ be the time between when the bidder pool first reaches $n-1$ bidders and when it first reaches n bidders. For any $n \in \mathbb{N}$, $\tau(n,n+1) \geq_{FOSD} \tau(n-1,n)$. Thus, $\beta(j) = \mathbb{E}[e^{-\rho\tau(j-1,j)}]$ is decreasing in j.*

Proof. The proof is by coupling. At $t = 0$, start n independent exponential departure clocks simultaneously. Label one of these clocks "first." Independently of these clocks, draw arrival clocks sequentially: first draw $W_1 \sim G$ at $t = 0$, then when the first clock ticks ($t = W_1$) draw $W_2 \sim G$ independently at $t = W_1$, and so on. Add an additional independent departure clock every time a new arrival clock $W_j, j \geq 2$ is drawn. Then define two stochastic processes $\mathbf{M} = \{M_t, t \geq 0\}$ and $\mathbf{M}' = \{M'_t, t \geq 0\}$ on this space such that

- Initialize M_0 at n and M'_0 at $n - 1$
- If an arrival clock W_j "ticks," increase $\mathbf{M}$ and $\mathbf{M}'$ by 1.
- If any departure clock ticks, decrease $\mathbf{M}$ by 1.
- If any departure clock other than the first one in the original set ticks, decrease $\mathbf{M}'$ by 1.

Note that $M_t = M'_t + 1$ before the first departure clock ticks, and $M_t = M'_t$ thereafter, i.e., $\mathbf{M}$ and $\mathbf{M}'$ eventually "couple."

By the renewals and exponential departures assumptions, M_t is distributed as N_t started from state $(n, 0)$, while M'_t is distributed as N_t started from state $(n - 1, 0)$. Therefore, the time $\mathbf{M}$ first crosses $n+1$, denoted σ_{n+1}, has the same distribution as $\tau(n, n+1)$. Similarly, the time $\mathbf{M}'$ first crosses n, denoted σ'_n, has the same distribution as $\tau(n - 1, n)$.

We claim $\sigma_{n+1} \geq \sigma'_n$ almost surely. First, if $\mathbf{M}$ and $\mathbf{M}'$ have not coupled by $\sigma_{n+1} \wedge \sigma'_n$, M will reach $n + 1$ at σ'_n (since $M_t = M'_t + 1$ before coupling). Hence, $\sigma_{n+1} = \sigma'_n$. Second, suppose that $\mathbf{M}$ and $\mathbf{M}'$ couple before time $\sigma_{n+1} \wedge \sigma'_n$ (so that they meet at some state $k \leq n - 1$). Then $\sigma_{n+1} > \sigma'_n$, since $M_t = M'_t$ after coupling, and $\mathbf{M}$ must reach n before reaching $n + 1$. Therefore, we conclude that $\tau(n, n + 1) \geq_{FOSD} \tau(n - 1, n)$ by the usual argument, and the result follows. □

We are now ready to prove Theorem 4.

Proof of Theorem 4. Let $\tau^*_{0,0}$ be the essential infimum over optimal stopping times starting from state $(0, 0)$. Suppose, contrary to the theorem, that $\tau^*_{0,0}$ dictates stopping in between arrivals. That is, there exists a state (n, w) with $w > 0$, such that $Q(n) \geq V(n, w)$. Then, using Lemma 6, $Q(n) \geq V(n,w) \geq V(n,0)$, so if the auctioneer weakly prefers stopping at (n,w), she would also prefer it at $(n,0)$. In any history starting from state $(0,0)$, $(n,0)$ is always reached before (n,w), so $\tau^*_{0,0}$ would have dictated stopping at $(n,0)$ without ever reaching (n,w), a contradiction. Therefore, if starting from $(0,0)$ the auctioneer stops at all, she stops at an arrival. Accordingly, we need to check that the auctioneer would indeed stop in finite time almost surely. But this is immediate because $Q(1) > 0$ and the payoff from never stopping is 0. Altogether, we have that the auction will happen the first time N_t reaches an acceptable threshold, i.e., some n such that $Q(n) \geq V(n,0)$. Let n^* denote the smallest such n.

It remains to prove that in fact n^* is given by (3). The optimal stopping policy must be time-consistent, so if the auctioneer stops at n^*, then n^* must be the smallest maximizer of the ex ante expected discounted payoff from waiting for exactly n bidders. That is, n^* must be the smallest maximizer of $\delta(n)Q(n)$, where $\delta(n)$ is the expected discount factor from waiting for n bidders. By Lemma 7, our assumptions on bidder dynamics imply that the n-step discount (from waiting for n bidders) decomposes into $\delta(n) = \prod_{i=1}^{n} \beta(i)$. The inequality $Q(n) \geq \beta(n+1)Q(n+1)$ is therefore a necessary local condition for optimization of $\delta(n)Q(n)$, i.e., it implies that $\delta(n)Q(n) \geq \delta(n+1)Q(n+1)$. Moreover, a coupling argument (Lemma 8) implies that the successive record times $\tau(n-1,n)$ increase in the first-order stochastic dominance sense as n grows, so $\beta(n) = \mathbb{E}[e^{-\rho\tau(n-1,n)}]$ is decreasing. Together with Point 3 of Lemma 5, since $\beta(\cdot)$ is decreasing, so is $n \mapsto \beta(n+1)\frac{Q(n+1)}{Q(n)}$. It follows that $\prod_{i=1}^{n} \beta(i)Q(n)$ is single-peaked. Therefore the condition in (3) is in fact a sufficient condition for maximizing $\delta(n)Q(n)$, and the optimal policy starting from $(0,0)$ stops the first time n^* is reached. $\square$

D Worst-Case Performance of Non-Adaptive Policies: Proof of Theorem 2

Proof. We construct a sequence $\{F_n, G_n\}_{n=1}^{\infty}$ such that $\frac{DET(r,F_n,G_n)}{OPT(r,F_n,G_n)} \to 0$. Suppose first that $F_n = F$, where F is a degenerate distribution at 1. The degeneracy of F will be without loss of generality because we can approximate any such F by a sequence of non-degenerate, continuous, and regular distributions, and the same logic used here will apply along that sequence.

Now, consider the following inter-arrival distribution G_n. Fix any positive integer M. Let the support of G_n be $\{t_1, \dots, t_M\}$ for any n. Let p_k denote the probability that G_n puts on t_k for $k = 1, \dots, M$. (We suppress the dependence of p_j on n for notational simplicity.) Then define each p_i inductively as follows.

$$\begin{aligned} p_1 &= \frac{1}{n} \\ p_k &= \sqrt{p_{k-1}} - p_1 - p_2 - \cdots - p_{k-1}, \\ p_M &= 1 - p_1 - \cdots - p_{M-1}. \end{aligned}$$

There are two remarks. First, for large n, $p_k \geq 0$ for each k, because the first term $\sqrt{p_{k-1}}$ converges to zero strictly slower than the other terms $-p_1 - p_2 - \cdots - p_{k-1}$. Second, note that $\sum_{j=1}^{k} p_j = \sqrt{p_{k-1}}$ holds for each $k < M$.

Now, suppose that the seller, who has to choose a fixed deadline, holds the auction at t_1 (later we show that we can construct $t_1, \dots, t_M$ so that choosing t_1 is indeed optimal). Let $\alpha_k := e^{-rt_k}$ for each $k = 1, \dots, M$. Then it holds that

$$\begin{aligned} \frac{OPT}{DET} &= \frac{\sum_{i=1}^{M} p_i \alpha_i}{\max_{1 \leq k \leq M} \sum_{i=1}^{k} p_i \alpha_k} \\ &= \frac{\sum_{i=1}^{M} p_i \alpha_i}{\alpha_1 p_1} \\ &= 1 + \sum_{j=2}^{M} \frac{p_j}{p_1} \cdot \frac{\alpha_j}{\alpha_1}. \end{aligned} \tag{7}$$

For $t = t_1$ to be optimal, it is enough that for $j = 1, \dots, M$,

$$p_1 \alpha_1 \geq \alpha_j \sum_{i=1}^{j} p_i \iff \frac{p_1}{\sum_{i=1}^{j} p_i} \geq \frac{\alpha_j}{\alpha_1}.$$

Now, since $\frac{p_1}{\sum_{i=1}^{j} p_i}$ is less than 1 and is strictly decreasing in j, we can choose $0 < t_1 < t_2 < \cdots < t_M$ so that $\alpha_k = e^{-rt_k}$ for $k = 1, \dots, M$ makes the above inequalities binding.

Combining this (in)equality and equation (7), we obtain

$$\frac{OPT}{DET} = 1 + \sum_{k=2}^{M} \frac{p_k}{\sum_{i=1}^{k} p_i}. \tag{8}$$

Now, we show that as $n \to +\infty$, the RHS of (8) converges to M. Indeed, each summand in the RHS can be written as

$$\frac{p_k}{\sum_{i=1}^{k} p_i} = \frac{p_k}{\sqrt{p_{k-1}}} = 1 - \sum_{j=1}^{k-1} \frac{p_j}{\sqrt{p_{k-1}}} \to 1 \quad \text{as} \quad n \to +\infty. \tag{9}$$

The last convergence follows from the fact that p_j is of order $n^{-\frac{1}{2j}}$ by construction. Now, (8) and (9) imply that

$$\frac{DET}{OPT} \to \frac{1}{M} \quad \text{as} \quad n \to +\infty.$$

That is, for each F, we can choose a sequence of G_n so that $\frac{DET}{OPT}$ converges to $1/M$. Because M is arbitrary, there exists $\{F_n, G_n\}_{n=1}^{\infty}$ such that $\frac{DET(r,F_n,G_n)}{OPT(r,F_n,G_n)} \to 0$. (We can construct such a sequence explicitly by a diagonalization argument.) $\square$