Optimal Search Auctions with a Deadline*

Joosung Lee† Daniel Z. Li‡

Abstract

This paper studies optimal auction design, in which a seller searches for potential bidders to allocate a product by a finite deadline. We fully characterise revenue-maximising mechanisms for both long- and short-lived bidders. For long-lived bidders, the cutoff revenue for stopping is time-invariant, while the sample size (the number of bidders to invite) increases over time and decreases with the current highest bid. We implement the optimal mechanism through a sequence of second-price auctions with *"clock-and-counter"* reserve prices, which depend jointly on a bidder's entry time and the number of subsequent competitors his bid induces. For short-lived bidders, the cutoff revenue for stopping decreases, and the optimal sample size increases over time. In this case, a sequence of second-price auctions with purely *"clock-based"* reserve prices suffices for implementation. Comparing the two cases, we quantify the seller's value of recalling past bids, which is single-peaked in the current highest bid. Finally, we identify a source of inefficiency in search auctions: a revenue-maximising seller over-searches relative to a welfare-maximising benchmark.

Keywords: search auction; deadline; sample sizes; adaptive search; clock-and-counter; reserve prices

JEL classification: D44; D82; D83

1 Introduction

Time constraint is a common concern in many search problems. A consumer searches for a Christmas gift on Amazon, a seller searches for buyers to allocate a perishable good, an employer interviews candidates to fill a position by a deadline, a football club has to sign new players by the end of a transfer window, the board of directors (BOD) searches for potential acquirers in a complex and multi-step M&A process, and so on. This paper investigates how

*This version: August, 2025. We thank Nejat Anbarci, Henrique Castro-Pires, Andrew Clausen, Makoto Hanazono, Rongzhu Ke, Ehud Lehrer, Takeharu Sogo, Cole Williams, Zaifu Yang, and Andy Zapechelnyuk for helpful discussion and comments, and the participants of IIOC 2024 (Boston), AMES 2024 (Hangzhou), APIOC 2024 (Seoul), YDE Theory Workshop 2025 (Edinburgh), Workshop in honor of Hervé Moulin's 75th Birthday (Glasgow) and the seminar participants at Durham University and Sungkyunkwan University for valuable comments. This paper was supported by Samsung Research Fund, SKKU, 2023 and Durham University Seedcorn Funds, 2023.

†Department of Economics, Sungkyunkwan University, Korea. E-mail: joosungx@gmail.com.

‡Durham University Business School. Email: daniel.li@durham.ac.uk.

a seller would optimally search for potential bidders to allocate a product by a deadline. The seller can invite the bidders batch by batch *adaptively* across different periods, but she incurs a search cost for each bidder she invites. Bidders also anticipate the seller's and other bidders' optimal decisions when participating and submitting their bids. If the seller's objective is to maximize the expected revenue net of gross search costs, what would be the optimal selling mechanism? One may observe that, due to the presence of a deadline, a one-by-one sequential search may not be optimal, and due to the presence of a search cost, a full-scale simultaneous transaction, e.g., inviting all bidders in one period, may not be optimal either, as the seller would value the intertemporal release of bidder information.

Our setting encompasses many important situations. A well-known puzzle in M&As is that the dominant selling process is non-competitive negotiation – for instance, when selling a company, the BOD typically contacts only one bidder at the outset (Andrade et al., 2001; Betton et al., 2008). This is against conventional wisdom that a seller benefits from competition among bidders, as it raises bid premiums. Some researchers therefore model the M&A process as a multi-period search for potential bidders. For example, if negotiations fail at one stage, the seller invites additional bidders to join the competition in subsequent stages. As a result, the pressure of following auctions will drive up the premium in the early negotiation stage (Aktas et al., 2010). This search argument is supported by empirical evidence that there is no significant difference in bid premiums across negotiations and auctions in M&As (Boone and Mulherin, 2007, 2008).

There are many other real-world examples. In academic hiring, UK universities commonly run several rounds of recruitment campaigns and increase their search intensities when the Research Excellence Framework (REF) deadline approaches. In public procurement, a government agency needs to screen and contact potential contractors and complete the procurement auction by a deadline. In dating markets, men and women regularly attend speed-dating events to search for potential partners, arguably under pressure from age. When exploring price quotes, a customer sequentially searches for quotes on goods or services offered by potential suppliers, typically under a time constraint.

In this paper, we develop a formal model of optimal search auctions, where a seller allocates an indivisible product among potential bidders within T discrete periods. Bidders are ex-ante homogeneous, and their product values are independently drawn from a common distribution.[1] Bidders cannot bid unless they are invited. The seller makes adaptive decisions about how many bidders to invite in each period, and her objective is to maximize the expected search profit. A constant search cost c is incurred for every bidder the seller

[1]For tractability, we assume an infinite pool of ex-ante homogeneous bidders. Absent this homogeneity, the seller would confront an NP-hard combinatorial problem, having to evaluate every possible subset of bidders across periods. The assumption also allows us to isolate how the seller's search intensity and stopping rule drive the dynamics of discriminatory reserve prices, free from confounding bidder-specific heterogeneity.

invites.[2]

We fully characterise the optimal search rule and profit in both cases of long-lived and short-lived bidders.[3] For long-lived bidders (Theorem 1), the optimal search rule is featured by a time-invariant cutoff revenue for optimal stopping, and a sequence of optimal sample sizes (i.e., the number of bidders to invite in each period), which exhibits two monotonic properties: other things being equal, the sequence of optimal sample sizes increases over time (the deadline effect) and decreases in the seller's fallback revenue (the fallback revenue effect). Since realizations of fallback revenues are random, the sequence of optimal sample sizes generally forms a stochastic process. For short-lived bidders (Theorem 2), the sequence of cutoff revenues for optimal stopping is decreasing over time and that of optimal sample sizes is increasing. Both sequences are deterministic and predetermined. This is because the fallback value effect is absent in the short-lived case. For example, if the seller declines the current highest bid and continues to search, her fallback revenue drops to zero when bidders are short-lived. In this case, the cutoff revenues for optimal stopping and the optimal sample sizes are determined solely by the deadline effect, i.e., the number of remaining periods.

To implement the optimal selling mechanisms, we further construct sequential second-price auctions with appropriately chosen reserve prices. The implementation problem is simpler for short-lived bidders, as they participate in the transaction only once and bidding their true value remains a dominant strategy in standard static second-price auctions. Therefore, by following the optimal sampling rule and setting a reserve price whose virtual value equals to the cutoff revenue for optimal stopping, the seller can implement the optimal selling outcomes (Proposition 2). The implementation is not straightforward for long-lived bidders, as bidders remain active and their bids can affect the seller's sampling decisions (i.e., the number of competitors invited) in subsequent periods. Moreover, as bidders' ability to affect the future competition levels depends not only on their bid, but also on their entry period, incentive compatible reserve prices must account for both dimensions. We show that the seller can ensure truthful bids by setting *discriminatory reserve prices* based on a *"clock"* and a *"counter"*, that is, (i) a bidder's entry period and (ii) the number of subsequent entrant bidders induced by his bid, given that the seller follows the optimal sampling rule. As a result, a bidder faces a lower reserve price if he enters the auction later or if his bid triggers more entrants in subsequent periods (Proposition 1).

A key advantage of our framework is that it enables a direct comparison of the seller's optimal search rules and profits between the cases of long-lived and short-lived bidders. This

[2]The assumption of a constant search cost is made for tractability and is not crucial to our main results. Heterogeneous search costs do not qualitatively alter our conclusions.

[3]A long-lived bidder, once invited, remains in the transaction until the deadline, whereas a short-lived bidder participates only in the stage when invited. The two cases of long-lived and short-lived bidders are analogous to sequential search with *full* and *no* recall, respectively.

comparison reveals four insights. First, the seller's cutoff revenue for optimal stopping is lower with short-lived bidders than with long-lived bidders. This is because the continuation search profit is lower when bidders are short-lived, and therefore, a seller with short-lived bidders is willing to accept a lower bid and stop. We refer to this as the *discouragement effect* of short-lived bidders (Corollary 1). Second, conditional on continuing search, a seller with short-lived bidders searches more intensively (i.e., a larger sample size) than one with long-lived bidders. This is because the seller forgoes her current fallback revenue when searching with short-lived bidders, and hence, a lower fallback revenue (zero indeed) increases marginal search revenue and encourages her to search more intensively. We refer to this as the *encouragement effect* of short-lived bidders (Corollary 2). Third, we define the difference in optimal search profits between long- and short-lived bidders as the *value of bidder longevity*. This value is clearly positive. Interestingly, however, it is single-peaked in the seller's fallback revenue: the value is maximized when the fallback revenue reaches the stopping cutoff of a seller with short-lived bidders. This non-monotonicity has practical implications – for example, as the current highest bid rises beyond a certain point, the value of bidder longevity declines, and the seller may in fact be willing to pay less to retain that bidder in future periods (Corollary 3). Finally, as the deadline T tends to infinity, the optimal search rules and profits in both cases of long- and short-lived bidders converge to the standard results of the stationary and infinite horizon (SIH) search problems. Specifically, the cutoff revenues for optimal stopping converge to the same constant, the optimal search profits converge to the same value, and the optimal sampling rules are both one-by-one sequential search (Section 6.3).

We also characterise efficient selling mechanisms for both long- and short-lived bidders, where the seller's objective is to maximize the expected net welfare, i.e., the winner's expected product value net of gross search costs (Proposition 3 and Proposition 4). We find that optimal selling mechanisms exhibit *over-searching* in the sense that the cutoff revenues for optimal stopping are higher than those of efficient selling mechanisms. This result holds for both cases of long- and short-lived bidders. It also indicates that the inefficiency of an optimal selling mechanism may stem from its inefficient search rule, which echoes the similar results in static search auctions (Szech, 2011; Xu and Li, 2019).

The remainder of this paper is organized as follows. Section 2 reviews the literature. Section 3 sets up the model. Section 4 solves the optimal selling mechanism with long-lived bidders, and Section 5 discusses its implementation. Section 6 solves the optimal selling mechanism with short-lived bidders. Section 7 characterises the efficient selling mechanisms in both cases of long- and short-lived bidders. Section 8 concludes. All missing proofs appear in the Appendix.

2 Related Literature

Our work is mainly related to the following strands of literature: sequential search and search with a deadline, search mechanism, dynamic auctions and revenue management, negotiation versus auction, and finite horizon decision-making problems.

For classic sequential search problems with *full* recall (so-called Pandora's problem), Gittins (1979) and Weitzman (1979) have fully characterised the optimal search rule. To be specific, a searcher faces a number of closed boxes, each containing a random prize; she can open the boxes sequentially, each at a search cost; her objective is to maximize the expected prize value discovered net of gross search costs. They show that a unique cutoff value can be allocated to each box, at which the searcher is indifferent between keeping that value and inspecting that box. The optimal search rule is simple: i) the searcher inspects the boxes in the order of descending cutoff values; ii) she stops searching whenever the value discovered is greater than the highest cutoff value of the remaining unopened boxes.[4]

When the search is bounded by a deadline, a one-by-one sequential search may no longer be optimal. Gal et al. (1981) introduces a deadline into the sequential job search model of Lippman and McCall (1976). With *no* recall, they show that the optimal search rule is featured by decreasing reservation wages and increasing search intensities over time. Morgan (1983) further studies the case of *full* recall and shows that the sequence of optimal search intensities is a stochastic process.[5] Lee and Li (2022) investigate a problem of sequential search with a deadline and continuous search effort to characterise the optimal search rule and value. Our work differs from pure search problems in which the search objectives (e.g., boxes) are non-strategic. In contrast, because bidders in search auctions are strategic agents, the seller must account for incentive compatibility when designing the selling mechanism. Moreover, our model treats sample size as a discrete choice variable is discrete, whereas Lee and Li (2022) assumes a continuous sample size for tractability.

A strand of literature studies search mechanisms, where the targets for search are strategic agents rather than non-strategic objects like boxes, yet focuses on one-by-one search with no deadline. McAfee and McMillan (1988) study a procurement problem where a buyer searches for homogeneous long-lived suppliers one-by-one sequentially. They show that the optimal search mechanism is a sequential auction with a constant reserve price. Crémer et al. (2007) examine a search mechanism where a seller searches for heterogeneous long-lived bidders. Using a mechanism design approach, they prove that the optimal search mechanism problem

[4] Kleinberg et al. (2017) prove that the sequential search problem *à la* Pandora can be reformulated as a static combinatorial optimization problem, and Armstrong (2017) and Choi et al. (2018) further show some important applications of this new approach in consumer search problems.

[5] Morgan and Manning (1985) investigate the problem where a searcher also chooses the number of periods for search, and present some results on the existence and properties of optimal search rules.

can be reformulated as Pandora's problem. They further show that, in the case of private values and no discounting, the mechanisms can be implemented by a sequence of second-price auctions with decreasing reserve prices and one-by-one sequential search.[6] Different from one-by-one sequential search, Lee and Li (2023) study a seller's compound search, where a seller commits to a partition of T bidder-samples and searches across them sequentially. If one takes a bidder-sample as an aggregate bidder, their compound search problem can be recast as Pandora's problem as in Crémer et al. (2007). For long-lived bidders, they show that the optimal stopping cutoff for a bidder-sample is decreasing in its sample size. Therefore, the optimal search rule suggests a search order of decreasing reserve prices and increasing sample sizes over time.

However, the above papers on search mechanisms did not consider time constraints. For example, Lee and Li (2023) study a seller's T-period sequential search across T preset bidder-samples, and there is no time constraint indeed. Moreover, their compound search mechanism is suboptimal in a T-period search problem, as the seller commits to preset bidder-samples. In other words, the seller does not have full discretion to choose a sample size given the information released from a search history. In this paper, we forgo the commitment assumption and characterise an optimal search mechanism where a seller can make fully contingent search decisions in each period. As a result, an entrant bidder has incentives to overbid or underbid so as to influence the seller's future search decisions. This intertemporal incentive issue arises uniquely in the presence of a finite deadline.

This paper is also connected to the literature on dynamic mechanisms and revenue management, where the set of participating buyers changes over time while their product values are persistent (Li, 2009; Pai and Vohra, 2013; Said, 2011, 2012; Board and Skrzypacz, 2016; Chaves and Ichihashi, 2024). These papers consider how to allocate multiple products to potential buyers within T periods, and focus mainly on the periodic cutoffs for allocating products. In general, they take the sequence of buyer arrivals as exogenously given. In contrast, though our paper also studies the allocation of a product to potential bidders within T periods, yet the sequence of bidder arrivals is endogenously determined by the seller's optimal search decisions. In this paper, we not only characterise the optimal periodic cutoffs for allocating a product, but also derive the sequence of optimal sample sizes of bidders.[7]

[6] Lauermann and Wolinsky (2016) study a common-value search auction, where an informed buyer searches for short-lived sellers one-by-one sequentially. An invited seller may partially learn the buyer's type through a noisy signal. They find that information aggregation through price is worse in this search auction than in a standard common-value auction. They attribute this failure of information aggregation to a stronger form of the winner's curse that arises with sequential search.

[7] Zhang (2021) studies the optimal sequence of posted-price and auction in a sequential mechanism, where a population of short-lived bidders enters the market periodically. In each period, the seller chooses between a posted-price and an auction mechanism. He shows that, when there is a deadline and the auction cost is moderate, the optimal mechanism sequence takes the form of posted-prices then auctions.

Our results offer new insights into dynamic auction design, particularly regarding discriminatory reserve prices that depends on a "clock" (the bidder's entry time) and a "counter" (the number of competitors). There is evidence in the literature that the seller's optimal reserve price may decline monotonically as the deadline approaches (Board and Skrzypacz, 2016; Mierendorff, 2016; Kanoria and Nazerzadeh, 2020).[8] This monotonicity also appears in our model with short-lived bidders. As time goes, the market gets thicker and the lower reserve price is charged to the new entrant accordingly. Importantly, this *clock-based reserve schedule* does *not* depend on the actual number of bidders who have participated *prior to* a bidder's entry. This is because the seller's search policy is determined by the current highest bid and the remaining search opportunities, but not by the number of past participants. As a result, the winner's payment is determined by the maximum of the second-highest bid and the clock-based reserve price, driven by the remaining search opportunities.

When bidders are long-lived, however, the reserve must respond not only to the time of a bidder's entry (the clock) but also to the actual number of other bidders who enter *after him* (the counter). A winner's payment must fall further if the seller subsequently invites additional bidders; otherwise he may overbid to discourage further search by the seller.[9] This ex-post adjustment of the reserve price reflects the market thickness that materializes after a bidder's entry. Moreover, the optimal reserve evolves as a stochastic process, as newly observed bids reveal the intensity of future competition. Such ex-post reserve price adjustment in a stochastic auction is akin to Che and Choi (2025)'s queuing model, in which the optimal reserve is determined by the number of buyers in a queue.[10] When the seller strategically searches for bidders over time, we find that the optimal reserve price combines the *time-declining component* based on a bidder's entry period and the *post-entry adjustment* by counting the number of competitors.

Notably, our results address the persistent debate on the choice between negotiation and auction as an optimal selling mechanism. Bulow and Klemperer (2009) argue that simultaneous auctions outperform sequential negotiations in a setting of costly entry, where an early entrant may submit a jump bid to deter the entries of outside bidders, which may harm the seller.[11] But the empirical evidence does not support their results in general. For

[8]The source of the decreasing reserve price varies across models. In Board and Skrzypacz (2016), the optimal reserve price remains constant when only one product is for sale. Mierendorff (2016) analyses a 2-period 2-bidder case in which each bidder may face different deadline, but not the seller. In Kanoria and Nazerzadeh (2020), the reserve price decreases only when the seller learns about the underlying distribution; otherwise, it stays constant.

[9]This is a distinctive feature of auctions with long-lived bidders. The mechanism must prevent any bidder from influencing future search intensity by shading her bid. Che and Klumpp (2016) study such intertemporal incentives in a fixed second-price auction format.

[10]Since Che and Choi (2025) study a multi-item auction with inventory, the optimal reserve depends not only on the number of buyers in a queue but also the number of items in the seller's inventory.

[11]Lu et al. (2021) study how to orchestrate costly information acquisition in an auction with a pre-short-

example, in M&As, the dominant selling process is negotiation, not auction. Our results offer an alternative explanation for this puzzle by modelling the M&A selling process as a seller's adaptive search for bidders by a deadline. That is, sequential negotiation can outperform simultaneous auction if the seller adaptively invites subsequent bidders and mitigates bidders' intertemporal incentives to overbid through an ex-post adjustment of the reserve price that accounts for the actual number of competitors triggered by their bid.

Finally, our work also contributes to the literature on finite-horizon decision-making problems. Baucells and Zhao (2018) study a continuous-time decision-making problem within a finite horizon, and formalize the notion that fatigue accumulates with effort and decays with rest. They show that the optimal effort follows a U-shaped pattern over time, which is supported by the empirical evidence from swimming competitions. Du et al. (2022) consider a seller's optimal effort management when she sells a product over a finite horizon with consumers arriving by a Poisson process. Under an all-or-nothing contract, they show that the optimal sales effort is non-monotonic in both the remaining time and the outstanding sales volume required to reach the target. In contrast, we study a sequential search mechanism in discrete time, where the choice variable (i.e., the sample size) is also discrete.

3 The Model

A seller (she) wants to allocate an indivisible product to a number of potential bidders (he) over T discrete periods, indexed by t. To invite a bidder, the seller needs to pay a search cost $c > 0$, and she can decide how many bidders to invite in each period. Bidders can not bid if not invited. The seller's product value is normalized to 0. A bidder's product value, X_i, is drawn IID from a distribution F on $\mathcal{X} = [0, \bar{x}]$, with its density $f > 0$ on $(0, \bar{x})$. Denote x_i the realization of X_i, which is persistent over time and is bidder i's private information. We assume F is of increasing failure rate, and the virtual value

$$\psi(x) = x - \frac{1 - F(x)}{f(x)}$$

is hence strictly increasing. A bidder's payoff is equal to his product value minus his payment if he is allocated the product, and zero otherwise. The seller's search profit is the product revenue net of her search costs. Both the seller and the bidders are risk neutral, and we disregard time discounting in this paper.[12]

listing stage. Bidders are initially endowed with private signals that are positively correlated to their true values, and a bidder can learn his true value by paying an entry cost. They show that, under a sequential short-listing rule, the seller admits the most efficient remaining bidder in each round, provided that his conditional expected contribution to the virtual surplus is positive.

[12]Introducing time discounting to the model does not change our results qualitatively.

Let M_t denote the sample of bidders the seller invites in period t, and $N_t := \bigcup_{\tau=0}^{t} M_\tau$ is the union of disjoint bidder-samples the seller has invited till the end of period t, with $N_0 \equiv M_0 \equiv \emptyset$. Denote m_t the sample size of M_t, which also measures the seller's search intensity in period t. We consider both long and short-lived bidders. A long-lived bidder, once invited, will stay in the transaction until the deadline. Yet, a short-lived bidder only participates in the stage transaction when invited. Apparently, for long- and short-lived bidders, the sets of participating bidders in period t are N_t and M_t, respectively.

At period 0, the seller selects a selling mechanism that is commonly understood by both the seller and the bidders. Without loss of generality, we restrict our attention to direct revelation mechanisms, in which each invited bidder reports his value to the seller upon arrival. Henceforth, a "mechanism" refers to such a direct revelation mechanism. A mechanism consists of a sequence of search rules (i.e., given a history of bidders' reports, the seller decides whether to stop or how many new bidders to invite in the next period) and a sequence of stage mechanisms (i.e., the payment and allocation rules for each stage transaction).

A mechanism is incentive feasible if it is both incentive compatible (i.e., each bidder wishes to truthfully report his value given that all other bidders are truthful), and is individually rational (i.e., each bidder is willing to participate immediately upon invitation). An optimal mechanism is an incentive feasible mechanism that maximizes the seller's expected search profit (i.e., the expected product revenue net of gross search costs).

4 Optimal Search Rule for Long-lived Bidders

In an incentive feasible mechanism, as each bidder's value is perfectly persistent, the maximum revenue a seller can secure from a truthful bidder is the bidder's virtual value.[13] To be specific, let G denote the distribution of the virtual value $V := \psi(X)$ on $[\underline{v}, \bar{v}]$ with $\underline{v} = \psi(0) < 0$ and $\bar{v} = \psi(\bar{x}) = \bar{x}$. It then follows that

$$G(v) = \Pr[V \leq v] = \Pr[\psi(X) \leq v] = F(\psi^{-1}(v)). \tag{1}$$

Denote V^m and X^m the largest order statistics of m independent draws from G and F respectively, which follow the distributions of G^m and F^m, with $V^0 \equiv \underline{v}$ and $X^0 \equiv 0$.

We next derive the seller's optimal search rule in an incentive feasible mechanism. At the beginning of period t, given that all bidders $i \in N_{t-1}$ reported truthfully, the seller then has a fallback revenue $v := \max_{i \in N_{t-1}}\{\psi(x_i)\} \geq 0$ which she can claim by stopping.[14] If she

[13] See Bergemann and Välimäki (2019) for a recent review of dynamic mechanism design problems.

[14] With the set N_{t-1} of participating bidders, following Myerson (1981), the seller can guarantee a fallback

continues to invite m bidders in period t at a search cost mc, the fallback revenue at the beginning of period $t+1$ is $\max\{v, V^m\}$, as bidders are long-lived.

Let $\Pi_t(v)$ denote the optimal search profit that the seller can achieve by following an optimal search rule from the beginning of period t on. The Bellman equation for the search problem is then, for $t \leq T$,

$$\Pi_t(v) = \max\left\{v, \max_{m\in\mathbb{N}_0}\left\{\mathbb{E}\Pi_{t+1}(\max\{v, V^m\}) - cm\right\}\right\}, \tag{2}$$

with $\mathbb{N}_0 := \{0, 1, \cdots\}$ and $\Pi_{T+1}(v) = v$ as the seller keeps whatever she has after the deadline.

We can solve the Bellman equation (2) by backward induction. Specifically, at the beginning of period T, if the seller stops, she can claim the fallback revenue v. If she continues to invite m bidders, the expected search revenue in period T, denote by $R_T(m; v)$, is

$$R_T(m; v) := \mathbb{E}\max\{v, V^m\} = v + \int_v^{\bar{v}} (1 - G(w)^m)dw, \tag{3}$$

and the marginal search revenue of inviting one more bidder is

$$MR_T(m; v) := R_T(m+1; v) - R_T(m; v) \quad = \quad \int_v^{\bar{v}} G(w)^m(1 - G(w))dw, \tag{4}$$

which is strictly decreasing both in m and v, and converges to 0 when $m \to \infty$.

Let v^0 denote the cutoff revenue for optimal stopping in period T, which is given by

$$MR_T(0; v^0) = \int_{v^0}^{\bar{v}} (1 - G(w))dw = c. \tag{5}$$

Apparently, at the beginning of period T, if the fallback revenue is greater than or equals to the cutoff revenue ($v \geq v^0$), then $MR_T(0; v) \leq c$, and the seller will stop and claim v. If $v < v^0$, the seller will continue inviting bidders until the marginal revenue becomes smaller than the marginal cost. Let $m_T^*(v)$ denote the optimal number (sample size) of new bidders to invite in period T. It then follows that

$$m_T^*(v) = \min\{m \in \mathbb{N}_0 : MR_T(m; v) \leq c\}. \tag{6}$$

The properties of $MR_T(m; v)$ guarantee the existence of a unique $m_T^*(v)$.

Given the period T solution, we can solve the earlier period problems recursively. To be specific, we can define the expected search revenue of inviting m bidders in period $t \leq T$ by

$$R_t(m; v) := \mathbb{E}\Pi_{t+1}\left(\max\left\{v, V^m\right\}\right), \tag{7}$$

and the corresponding marginal revenue by $MR_t(m; v) := R_t(m+1; v) - R_t(m; v)$. Lemma 1 below summarizes some useful properties of the marginal revenue.

---

revenue

$$v = \max\left\{0, \max_{i\in N_{t-1}} \psi(x_i)\right\} \geq 0$$

using the second-price auction with a reserve price of $\psi^{-1}(0)$.

Lemma 1. *For any $t \leq T$, the marginal revenue $MR_t(m; v)$ is*

i) strictly positive for any $v < \bar{v}$ and any $m \geq 0$;

ii) strictly decreasing in v and $\lim_{v \to \bar{v}} MR_t(m; v) = 0$; and

iii) strictly decreasing in m and $\lim_{m \to \infty} MR_t(m; v) = 0$.

These properties are intuitive. For example, part i) implies that having one more bidder is always valuable unless the seller has already obtained the highest possible revenue $\bar{v}$. Part ii) simply suggests that, when the fallback revenue v gets higher, it becomes less likely to achieve a higher revenue by inviting one more bidder. Part iii) states that the expected revenue increases in the number of bidders at a diminishing rate, and in the limit when the number of bidders converges to infinity, the marginal benefit of inviting one more bidder is negligible. Figure 1, as well as Example 1, below illustrates the properties of $MR_t(m; v)$.

With the results of Lemma 1, it is easy to show that, given a fallback revenue v at the beginning of any period t, there exists a unique optimal sample size $m_t^*(v)$ at which the marginal search revenue starts to be smaller than the marginal search cost. We can then characterise the seller's optimal search rule and search profit in the case of long-lived bidders. The formal result is summarized in Theorem 1 below.

Theorem 1 (Optimal search for long-lived bidders)**.** *Let v be the seller's fallback revenue at the beginning of period $t \leq T$. The optimal sample size $m_t^*(v)$ in period t is given by*

$$m_t^*(v) = \min\{m \in \mathbb{N}_0 : MR_t(m; v) \leq c\}, \tag{8}$$

where the marginal revenue $MR_t(m; v)$ is recursively determined by

$$MR_t(m; v) = \int_v^{\bar{v}} G(w)^m (1 - G(w)) \prod_{\tau \geq t+1} G(w)^{m_\tau^*(w)} dw. \tag{9}$$

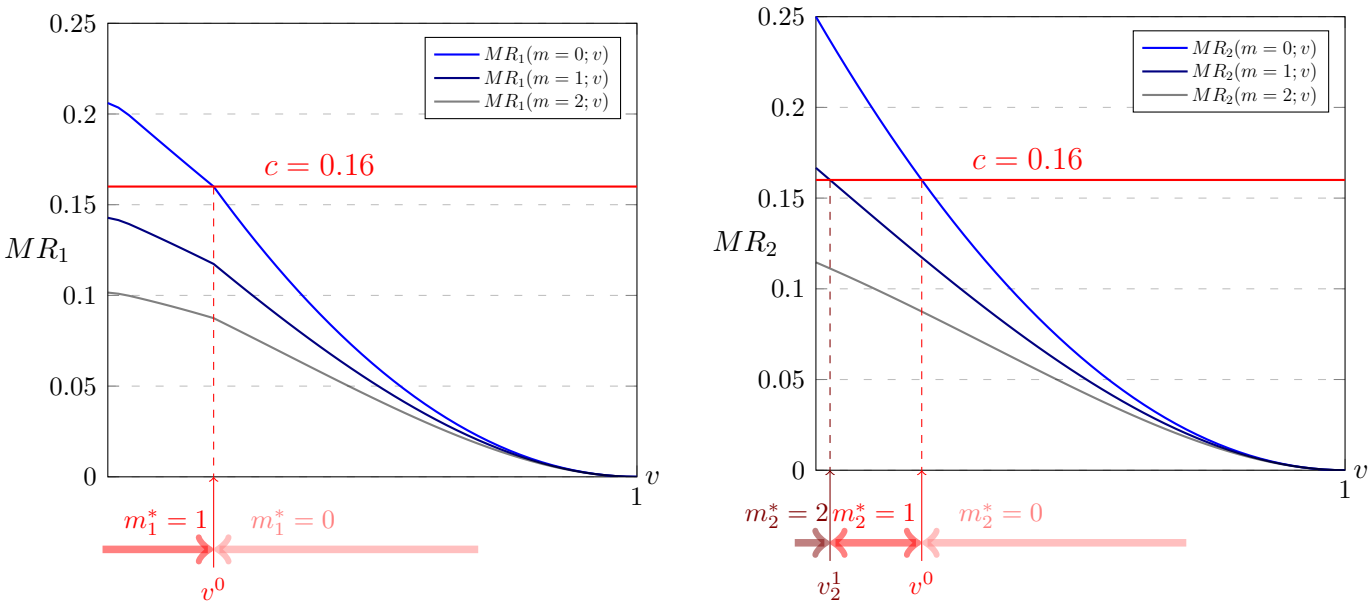
Specifically, it is optimal to stop at the beginning of period t, i.e., $m_t^*(v) = 0$, if and only if $v \geq v^0$. The optimal search profit $\Pi_t(v)$ is given by

$$\Pi_t(v) = \mathbb{E} \max\{v, V_t^*\} = v\mathcal{G}_t(v) + \int_v^{\bar{v}} w d\mathcal{G}_t, \tag{10}$$

where V_t^* is the continuation search profit achieved by following the optimal search rule from period t on, and its distribution function $\mathcal{G}_t(v)$ is recursively defined by

$$\mathcal{G}_t(v) = G(v)^{m_t^*(v)} \mathcal{G}_{t+1}(v) = \prod_{\tau \geq t} G(v)^{m_\tau^*(v)}. \tag{11}$$

Figure 1: Optimal search rule $m_t^*(v)$ in Example 1 ($c = 0.16$, $F(x) = x\mathbb{1}(0 \le x \le 1)$, $T = 2$)



(a) Period 1

(b) Period 2

Note. The optimal search rule $m_t^*(v)$ is determined by the marginal search revenue $MR_t(m, v)$ and the search cost c. In period 1, the optimal number of invitees is $m_1^*(v) = \mathbb{1}(v < v^0)$, where $v^0 = 0.2$. In period 2, the optimal number of invitees is $m_2^*(v) = \mathbb{1}(v < v_2^1) + \mathbb{1}(v < v^0)$, where $v_2^1 \approx 0.0267$.

Remark 1. It is clear from Lemma 1 that the optimal sample size $m_t^*(v)$ given by (8) is unique and is decreasing in v. Therefore, for each $m = 0, 1, \ldots, m_t^*(0) - 1$, there exists a unique cutoff revenue, denoted by v_t^m, that solves $MR_t(m; v) = c$. When the fallback revenue $v \ge v_t^0$, $m_t^*(v) = 0$ and it is optimal to stop in period t. Importantly, Theorem 1 suggests that $v_t^0 = v^0$ for any $t \le T$, that is, the cutoff revenue for optimal stopping is constant over time.[15] Given the decreasing sequence of $v^0 > v_t^1 > v_t^2 > \ldots > v_t^{m_t^*(0)-1}$, we have the following explicit expression of the optimal sample sizes in any period $t \le T$:

$$m_t^*(v) = \sum_{m=1}^{m_t^*(0)} \mathbb{1}(v < v_t^{m-1}) = \begin{cases} 0 & \text{if } v \ge v^0 \\ 1 & \text{if } v_t^1 \le v < v^0 \\ 2 & \text{if } v_t^2 \le v < v_t^1 \\ \vdots & \\ m_t^*(0) & \text{if } v < v_t^{m_t^*(0)-1}. \end{cases} \tag{12}$$

Remark 2. The marginal revenue $MR_t(m, v)$ is decreasing in v by Lemma 1 and increasing

[15] For instance, in period $T - 1$, if the seller's fallback revenue v is greater than or equals to v_T^0, then it would be her last period of search, as she will stop searching in period T given her next period fallback revenue will be no less than v. The proof of Theorem 1 formally shows this property by induction.

in t by (9) as $G(w)^{m_\tau^*(w)} \leq 1$. It follows that the optimal sample size $m_t^*(v)$ is decreasing in v and increasing in t. We call the former the *fallback revenue effect* and the latter the *deadline effect* of an optimal sample size $m_t^*(v)$. That is, when the fallback revenue v increases, the seller will search less intensively, and when the deadline gets closer, the seller will search more intensively. A direct implication is that, all other things being equal, the seller will invite the smallest number of bidders in period $t = 1$. The result may explain why one-on-one negotiation can be a dominant selling process in many important markets, such as M&As.

Remark 3. A time-invariant cutoff revenue v^0 for optimal stopping may also explain why there is no significant difference in bid premiums between negotiation and auction in M&As. This is because a seller has the same cutoff revenue v^0 for optimal stopping regardless of whether it is in the early negotiation stage or the later auction stages. The result of a constant v^0 seems counter-intuitive, yet it is based on the fact that the search is with free recall. When bidders are long-lived, the seller's fallback revenue increases weakly over time. Suppose the cutoff revenue for optimal stopping in period T is v^0, and that the seller has a fallback revenue v^0 at the beginning of period $T-1$. Then period $T-1$ will be the seller's last period of search, as she will definitely stop searching in period T. Therefore, the seller's cutoff revenue for optimal stopping in period $T-1$ is also v^0.

Remark 4. By (10), the optimal search profit $\Pi_t(v)$ is increasing and convex in v and decreasing in t. To be specific:

i) For $v \in [0, v^0)$, $\Pi_t(v)$ is strictly convex, and $\Pi_t(v) > v$ and $\Pi_t(v) > \Pi_{t+1}(v)$

ii) For $v \in [v^0, \bar{v}]$, $\Pi_t(v) = v$.

In other words, the optimal search profit $\Pi_t(v)$ is larger if the fallback revenue v, which she can reclaim in the case of stopping, is larger, or if it is in an earlier period (a smaller t) such that the seller will have more opportunities of search to increase her expected profit. One may observe that the distributions of V_t^* are ordered in terms of first order stochastic dominance, i.e., $\mathcal{G}_{t+1}(v) \geq \mathcal{G}_t(v)$, which implies that $\Pi_t(v) > \Pi_{t+1}(v)$.

At the end of this section, we provide an example of two-period search problem to illustrate the main results of Theorem 1.

Example 1 (Optimal search rule for long-lived bidders)**.** Let $c = 0.16$, $F(x) = x\mathbb{1}(0 \leq x \leq 1)$, and $T = 2$. The virtual value $\psi(x) = 2x - 1$ then has the distribution $G(v) = \frac{1}{2}(v+1)\,\mathbb{1}(-1 \leq v \leq 1)$. We next derive the optimal search rule and the optimal search profit, using the results of Theorem 1.

- In the last period $t = 2$, the marginal search revenue is

$$MR_2(m; v) = \int_v^1 G(w)^m (1 - G(w)) dw = \frac{2^{m+2} - (1+v)^{m+1}(m+3-v(m+1))}{2^{m+1}(m+1)(m+2)},$$

which is illustrated in Figure 1(b) for $m = 0, 1, 2$. As $MR_2(2; 0) \approx 0.1146 < c$, the optimal sample size must be smaller than 2 from (12). For $m = 0, 1$, the solution to $MR_2(m; v) = c$ gives the cutoff values $v^0 = v_2^0 = 0.2$ and $v_2^1 \approx 0.02667$. Given a fallback revenue v, the optimal sample size in period $t = 2$ is then

$$m_2^*(v) = \mathbb{1}(v < v_2^1) + \mathbb{1}(v < v^0). \tag{13}$$

From $\mathcal{G}_2(v) = G(v)^{m_2^*(v)}$, we have the optimal search profit in period $t = 2$

$$\Pi_2(v) = \mathbb{E} \max\{v, V_2^*\} = v\mathcal{G}_2(v) + \int_v^{\bar{v}} w d\mathcal{G}_2(w).$$

- In the first period $t = 1$, the marginal search revenue, by (9), is

$$MR_1(m; v) = \int_v^1 G(w)^{m+m_2^*(w)}(1 - G(w))dw.$$

Using (13), we have the following three possible cases:

i) If $v < v_2^1$, then $m_2^*(w) \in \{0, 1, 2\}$ and

$$\begin{aligned} MR_1(m; v) &= \int_v^{v_2^1} G(w)^{m+2}(1 - G(w))dw + \int_{v_2^1}^{v^0} G(w)^{m+1}(1 - G(w))dw \\ &\quad + \int_{v^0}^1 G(w)^m(1 - G(w))dw. \end{aligned}$$

ii) If $v_2^1 \le v < v^0$, then $m_2^*(w) \in \{0, 1\}$ and

$$MR_1(m; v) = \int_v^{v^0} G(w)^{m+1}(1 - G(w))dw + \int_{v^0}^1 G(w)^m(1 - G(w))dw;$$

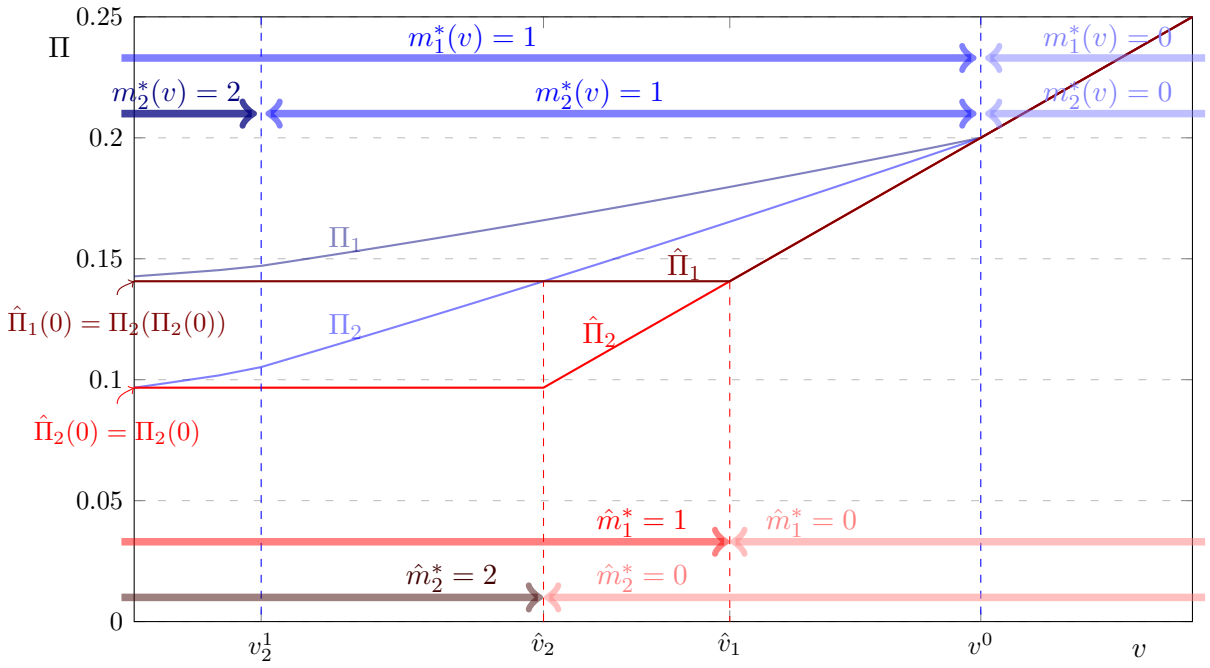
iii) If $v \ge v^0$, then $m_2^*(w) \in \{0\}$ and

$$MR_1(m; v) = \int_v^1 G(w)^m(1 - G(w))dw.$$

The first-period marginal revenue $MR_1(m; v)$ is illustrated in Figure 1(a), for $m = 0, 1, 2$. Observe that for any $v \ge v^0$, $MR_2(m, v) = MR_1(m, v)$, which suggests that the cutoff revenue v^0 for optimal stopping does not change over time. For $v < v^0$, the optimal sample size $m_2^*(v) \ge 1$, which implies that $MR_1(m; v) < MR_2(m; v)$. Therefore, the marginal search revenue is increasing over time, and the seller searches more intensively as the deadline approaches. This deadline effect of $m_1^*(v) \le m_2^*(v)$ is illustrated by Figure 2 (presented in blue) or by comparing Figure 1(a) and Figure 1(b). As $MR_1(1; 0) \approx 0.1428 < c$, the optimal sample size in period $t = 1$ is at most 1, as given by

$$m_1^*(v) = \mathbb{1}(v < v^0).$$

Figure 2: Optimal Search Rules and Profits: Long vs. Short-lived Bidders



Note. $c = 0.16$, $F(x) = x\mathbb{1}(0 \leq x \leq 1)$, and $T = 2$. Π_t is the search profit with long-lived bidders from the optimal search rule m_t^* (presented in blue); whereas $\hat{\Pi}_t$ is that with short-lived bidders from the corresponding optimal search rule $\hat{m}_t^*$ (presented in red).

Given $m_1^*(v)$ and $m_2^*(v)$, the continuation search profit V_1^* has the following distribution

$$\mathcal{G}_1(v) = G(v)^{m_1^*(v)+m_2^*(v)} = \begin{cases} G(v)^3 & \text{if } v < v_2^1 \\ G(v)^2 & \text{if } v_2^1 \leq v < v^0 \\ G(v)^0 = 1 & \text{if } v \geq v^0. \end{cases}$$

Given an initial fallback revenue of 0, the seller's optimal search profit is then

$$\Pi_1(0) = \mathbb{E}\max\{0, V_1^*\} = \int_0^1 w d\mathcal{G}_1 = v^0 - \int_{v_2^1}^{v^0} G(w)^2 dw - \int_0^{v_2^1} G(w)^3 dw \approx 0.1427.$$

The optimal search profits for each period $t = 1, 2$ are plotted in Figure 2 (in blue).

5 Implementation of the Optimal Search Auction

It is not straightforward to implement the optimal mechanism characterised in Theorem 1. Observe that the optimal search rule bears two key features:

- an adaptive sampling rule $m_t^*(v)$ which is decreasing in the fallback revenue v and increasing in time t; and

- a constant cutoff revenue v^0 for optimal stopping at the beginning of any period $t \leq T$.

Given the seller's optimal search rule, bidders may have incentives to overbid or underbid their values, since their bids can influence the seller's stopping decisions and her choices of sample sizes in subsequent periods. We construct a sequence of second-price auctions with discriminatory reserve prices that implements the optimal search rule while ensuring incentive compatibility. Specifically, a winner faces a *clock-and-counter reserve price* that is determined by both his entry period (the clock) and the number of new entrant bidders in subsequent periods (the counter).

In the remaining parts of this section, we first propose a sequential search auction and specify its invitation rule, along with the construction of the *clock-and-counter reserve prices*. Next, we prove that the proposed search auction is incentive feasible and implements the optimal search mechanism. Finally, we continue Example 1 and show how to implement the optimal search outcomes using the proposed search auction with the adaptive discriminatory reserve prices.

5.1 The Sequential Second-Price Auction

The seller runs a sequence of second-price auctions in T discrete periods. At the beginning of any period $t \leq T$, depending on the bid history, the seller decides how many new bidders to invite in that period. Each entrant bidder submits a bid upon his arrival. The seller records each entrant bidder's bid and entry period. A bidder's type is defined by (x, t), where x is his product value and t is his entry period, and only x is his private information. We define a *period-t winning bid*, denoted by $y_t := \max_{i \in N_t}\{x_i\}$, which is the highest bid among the set N_t of entrant bidders. Denote the seller's product value $y_0 \equiv 0$. We may suppress the subscript 't' of y_t sometimes.

Bid-triggered invitation rule Given the period-t winning bid y, at the beginning of period $t+1$, the seller will invite $\tilde{m}^*_{t+1}(y)$ new bidders according to

$$\tilde{m}^*_{t+1}(y) := m^*_{t+1}\big(\psi(y)\big). \tag{14}$$

If bidders bid truthfully, $v = \psi(y)$ is the seller's fallback revenue, and $\tilde{m}^*_{t+1}(y)$ then implements the seller's optimal sampling rule (8) as specified in Theorem 1. Specifically, if $y \geq p^0 := \psi^{-1}(v^0)$, then $\tilde{m}^*_{t+1}(y) = 0$ and it is optimal for the seller to stop.

Competition in subsequent periods If the period-t winning bid y remains being the winning bid in all subsequent periods and the seller follows the invitation rule (14), then

$$\ell_t(y) := \sum_{\tau=t+1}^{T} \tilde{m}_\tau^*(y) \tag{15}$$

is the total number of new entrant bidders the seller will invite from period $t+1$ to the last period T, with $\ell_T(y) \equiv 0$. It is clear that $\ell_t(y)$ is decreasing in both y and t. For an admissible integer $\ell \in \mathbb{N}_0$, we define a *cutoff price*

$$p_t^\ell := \inf\{y \in \mathcal{X} : \ell_t(y) = \ell\}. \tag{16}$$

That is, if the period-t winning bid $y \geq p_t^\ell$, then at most ℓ new bidders will be invited in the subsequent periods. Observe that p_t^ℓ is decreasing in both ℓ and t.

When $\ell = 0$, we have the highest cutoff price $p_t^0 = \psi^{-1}(0)$ for all t, and hence, simply denote the time-invariant cutoff price for stopping $p^0 = p_t^0$. The lowest non-trivial cutoff price is r_T, which corresponds to the largest number of $\bar{\ell}_t := \ell_t(r_T)$. The cutoff prices p_t^ℓ are defined only for certain admissible values of ℓ. For example, if $y \geq p^0$, then the seller stops and $\ell = 0$. If y is slightly lower than p^0, then the seller will invite bidders one by one sequentially, and in total $\ell = T - t$ new entrant bidders will be invited in subsequent periods. A typical set of admissible values of ℓ is[16]

$$L_t := \left\{0,\, T-t,\, T-t+1,\, \ldots,\, \bar{\ell}_t\right\}.$$

Note that the cutoff prices partition the value interval $\mathcal{X}$ into segments that correspond to different total number ℓ of entrant bidders to invite in subsequent periods.

Clock-and-counter reserve prices Now, we construct an incentive-compatible reserve price from the bidders' perspective. Suppose a bidder enters in period $t \leq T$ and wins the auction with a bid x, which induces a total number $\ell = \ell_t(x)$ of new bidders invited after his entry. Since the "clock" and the "counter" correspond to t and ℓ, respectively, the auction sets a discriminatory reserve price r_t^ℓ for the bidder. Specifically, for any $t \leq T$ and $\ell \in L_t$, r_t^ℓ is defined as the unique solution to

$$\int_{r_T}^{p_t^\ell} \mathcal{F}_{t+1}(z)dz = \int_{r_t^\ell}^{p_t^\ell} F(z)^\ell dz, \tag{17}$$

[16]Some integers between $T-t$ and $\bar{\ell}_t$ may not be admissible. Since each function $\tilde{m}_\tau^*(x)$ is generally a step function and different periods may share the same cutoff thresholds, the total number of future invites $\ell_t(x)$ can exhibit upward jumps greater than one. That is, as x decreases, $\ell_t(x)$ may increase from ℓ to $\ell + k$ for some $k \geq 2$, skipping over intermediate values.

where $r_T = \psi^{-1}(0)$ is the Myerson's static optimal reserve and $\mathcal{F}_{t+1}(x)$ is the probability that the bidder's bid x remains the highest in subsequent periods under the optimal invitation rule of (14), which is recursively determined by

$$\mathcal{F}_t(x) = F\big(x\big)^{\tilde{m}_t^*(x)}\, \mathcal{F}_{t+1}(x), \tag{18}$$

with $\mathcal{F}_{T+1}(x) = 1$.[17] The construction of r_t^ℓ is based on the following payoff equivalence result: if the sequential search auction implements the optimal mechanism of Theorem 1, any bidder of the same type (x, t) will receive the same expected payoff in both mechanisms. Clearly, the discriminatory reserve price r_t^ℓ is decreasing both in ℓ and t.

Remark 5. Equation (17) is worth more explanation. Conditional on x being period-t winning bid for a bidder of type (x, t), let $U_t(x'; x)$ denote his expected payoff by reporting his value being x' instead. The expected payoff of truthful reporting is then $U_t(x) := U_t(x, x)$. In the optimal mechanism of Theorem 1, we have $U_t'(x) = \mathcal{F}_{t+1}(x)$ by envelope theorem, where $\mathcal{F}_{t+1}(x)$ is the probability of winning the product when the seller follows the optimal search rule of (8). It then follows that

$$U_t(x) = U_t(r_T) + \int_{r_T}^{x} \mathcal{F}_{t+1}(z)dz = \int_{r_T}^{x} \mathcal{F}_{t+1}(z)dz$$

as $U_t(r_T) = 0$. Therefore, the LHS of (17) is the expected payoff of a bidder of type (p_t^ℓ, t) in the optimal mechanism of Theorem 1. Similarly, the RHS is the expected payoff of the same bidder in a second-price search auction facing a discriminatory reserve price r_t^ℓ.

Since $\ell_T(x) = 0$, all the period-T entrant bidders face the same reserve price of $r_T^0 = r_T$, as in the standard static second-price auction.

The optimal auction With the invitation rule (14) and the clock-and-counter reserve prices given by (17), the sequential search auction is run as follows.

i) At the beginning of any period $t \le T$, given the winning bid $y_{t-1} = \max_{i \in N_{t-1}}\{x_i\}$, the seller invites $\tilde{m}_t^*(y_{t-1}) \ge 1$ new bidders, according to (14), and records their bids and entry period.[18]

ii) At the end of period $t < T$, if the period-t winning bid $y_t \ge p^0$, the seller then stops and allocates the product to the highest bidder among the set N_t of bidders. The payment by the winning bidder is $\max\{y_t^{(2)}, r_t^0\}$, where the reserve price r_t^0 is given by (17) and $y_t^{(2)}$ is the second highest bid among the set N_t of bidders. If $y_t < p^0$, the seller then proceeds to period $t + 1$, following the invitation rule (14).

[17] One may find that $\mathcal{F}_t$ is analogous to $\mathcal{G}_t$ in Theorem 1, yet is now expressed in terms of the product value and bidders' bid, instead of the virtual value and the seller's fallback revenue.

[18] If $\tilde{m}_t^*(y_{t-1}) = 0$, then the product should have been allocated without proceeding to period t.

iii) At the end of period $t = T$, if the period-T winning bid $y_T \geq r_T$, the seller then allocates the product to the highest bidder among the set N_T of bidders. The payment of the highest bidder is $\max\{y_T^{(2)}, r_{t^*}^{\ell}\}$, where the reserve price $r_{t^*}^{\ell}$ is given by (17) with t^* being the highest bidder's entry period and $\ell = \ell_{t^*}(y_T) = \sum_{\tau=t^*+1}^{T} \tilde{m}_\tau^*(y_T)$. If $y_T < r_T$, the seller retains the product.

5.2 Incentive Compatibility

Given a seller's optimal search rule, bidders may have incentives to overbid or underbid so as to affect the seller's search decisions. Such incentives become the strongest when a bidder's value is at the cutoff price p_t^{ℓ}, because an infinitesimal amount of deviation can induce a discrete change in the total number of future entrant bidders. Equation (17) guarantees that this incentive for deviation can be offset by setting the clock-and-counter reserve price r_t^{ℓ}, which drops discretely with increasing ℓ. That is, if a bidder's deviation induces fewer (more) entrant bidders in the subsequent periods, he will face a discontinuous jump (drop) of discriminatory reserve price by the "counter." The formal result is stated as below.

Proposition 1. *A sequence of second-price auctions with the invitation rule $\tilde{m}_t^*$ defined in (14) and the discriminatory reserve price r_t^{ℓ} defined in (17) implements the optimal mechanism of Theorem 1.*

Let us reconsider Example 1 and see how to implement the optimal search outcomes using a sequential search auction with the clock-and-counter reserve prices.

Example 2 (Optimal search auction with long-lived bidders)**.** Reconsider the two period search problem of Example 1. We now show how to implement the optimal search outcomes by a sequence of second-price auctions with the clock-and-counter reserve prices. At the beginning of period $t = 1$, given a fallback revenue $0 = \psi(r_T)$, the seller invites one bidder (denote by B_1) as $\tilde{m}_1^*(r_T) = m_1^*(0) = 1$ from (14).

- If B_1's bid $x \geq p^0 = \psi^{-1}(v^0) = 0.6$, the seller stops and B_1 immediately wins the auction. As $\ell_1(x) = 0$ from (15), the reserve price r_1^0 from (17) is the solution to

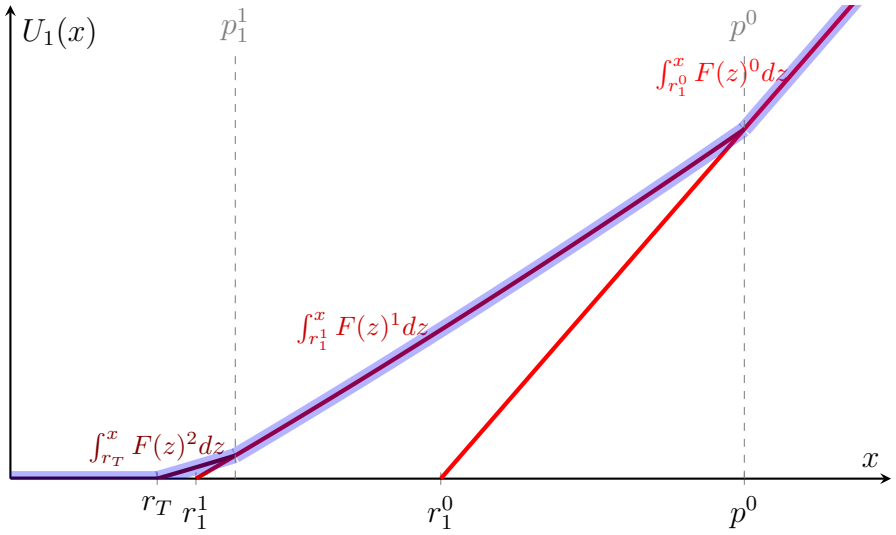
$$\int_{r_T}^{p^0} \mathcal{F}_2(z)dz = \int_{r_1^0}^{p^0} F(z)^0 dz. \tag{19}$$

  As $p_1^1 = \psi^{-1}(v_2^1) \approx 0.5133 > r_T = 0.5$, equation (19) then gives

$$\int_{r_T}^{p_1^1} z^2 dz + \int_{p_1^1}^{p^0} z dz = \int_{r_1^0}^{p^0} 1 dz, \tag{20}$$

  solving which we have the reserve price $r_1^0 \approx 0.5483$.

Figure 3: B_1's Expected Payoff of Truthful Bidding



Note. Bidder B_1's expected payoff $U_1(x)$ of truthful bidding is the envelop curve marked in blue, $U_1(x) = \max\left\{0, \int_{r_T}^x F(z)^2 dz, \int_{r_1^1}^x F(z)^1 dz, \int_{r_1^0}^x F(z)^0 dz\right\}$, as the discriminatory reserve prices r_1^1 and r_1^0 makes it convex to guarantee incentive compatibility.

- If B_1's bid $x < p^0$, the seller continues to invite $\tilde{m}_2^*(x)$ new bidders in period $t = 2$.

  i) If B_1's bid $x \in [p_1^1, p^0)$, the seller invites one more bidder (denoted by B_2). As $\ell_1(x) = 1$, B_1's reserve price is $r_1^1 \approx 0.5066$ which from (17) solves

  $$\int_{r_T}^{p_1^1} z^2 dz = \int_{r_1^1}^{p_1^1} z dz. \tag{21}$$

  It is apparent from (17) that B_2's reserve price is $r_T = 0.5$.

  ii) If B_1's bid $x < p_1^1$, the seller invites two more bidders (denoted by B_2 and B_3). Again, from (17), the solution to $\int_{r_T}^{p_1^2} z^2 dz = \int_{r_1^2}^{p_1^2} z^2 dz$ gives B_1's reserve price $r_1^2 = r_T$. It is clear from (17) that B_2 and B_3 also have the same reserve price r_T.

For the entrants in the period $T = 2$, it is obviously their weakly dominant strategy to bid true value. We need to confirm B_1 will also bid truthfully in period $t = 1$ given the specified reserve prices. Let B_1's true value be x, and Figure 3 plots his expected payoff $U_1(x)$ of truthful bidding. Note that the cutoff prices that induce different numbers of entrant bidders in period $T = 2$ are $r_T < p_1^1 < p^0$. First, B_1 will not deviate by bidding any x' such that $\ell_1(x') = \ell_1(x)$. That is, if x' and x induce the same number of entrants in period $T = 2$ and hence the same reserve price r_1^ℓ, B_1 will always bid his true value x. Next, we show B_1

will not deviate either if $\ell_1(x') \neq \ell_1(x)$. For example, consider $x \in [p_1^1, p^0)$, and a truthful bid x induces one entrant. If he underbids $x' \in [r_T, p_1^1)$ and induces two entrants, as illustrated in Figure 3, his expected payoff is

$$U_1(x'; x) = U_1(x') + \int_{x'}^{x} z^2 dz = U_1(p_1^1) + \int_{p_1^1}^{x} z^2 dz < U_1(p_1^1) + \int_{p_1^1}^{x} z dz = U_1(x).$$

If he overbid $x' \geq p^0$ and induces a stop decision of the seller, his expected payoff is

$$U_1(x'; x) = U_1(x') - \int_{x}^{x'} z^0 dz = U_1(p^0) - \int_{x}^{p^0} z^0 dz < U_1(p^0) - \int_{x}^{p^0} z dz = U_1(x).$$

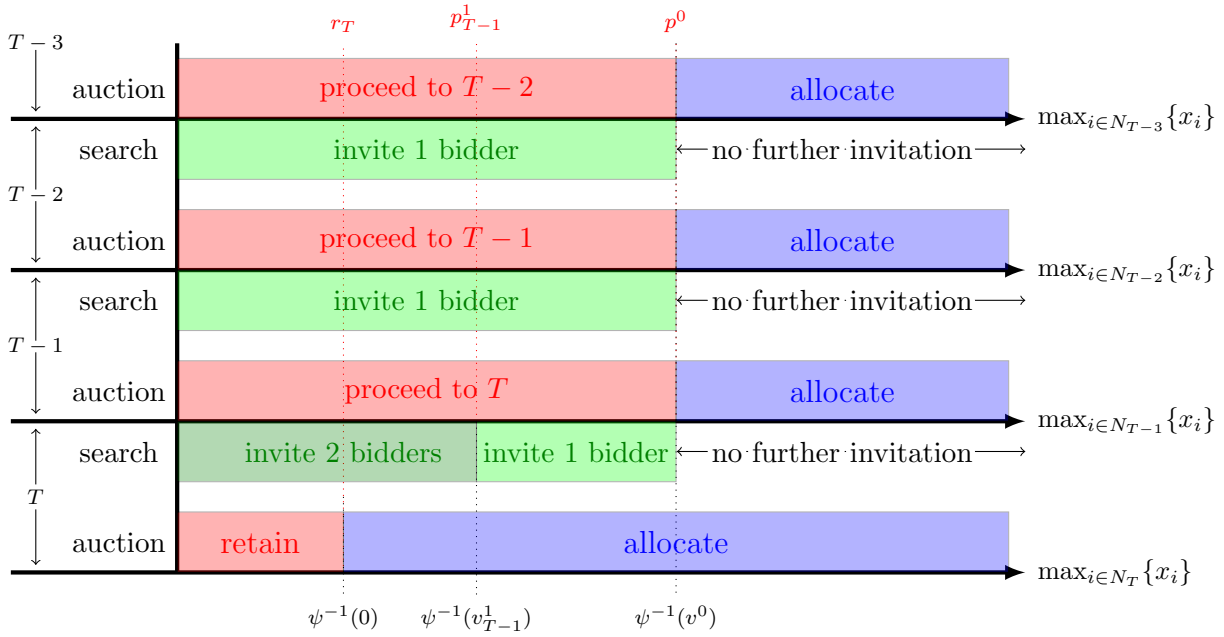
This example can be extended to any finite T-period problem with the same cost parameter c and value distribution function F. Note that in the two-period problem, we have $\tilde{m}_{T-1}^*(y) = \mathbb{1}(y < p^0)$. The deadline effect ($t' < t \implies \tilde{m}_{t'}^*(y) \leq \tilde{m}_t^*(y)$) and the fallback revenue effect ($y' < y \implies \tilde{m}_t^*(y') \geq \tilde{m}_t^*(y)$) together imply that $\tilde{m}_t^*(y) = \mathbb{1}(y < p^0)$ for all $t < T$ in any finite period problem. That is, the optimal sampling rule prescribes one-by-one invitations in every period $t < T$. Specifically, at the beginning of any period $t < T$, if the highest bid observed so far is less than $p^0 = 0.6$, the seller invites one (and only one) additional bidder in period t. In the last period T, however, the seller may invite two bidders, as in the above two-period example.

Figure 4(a) illustrates the implementation of the optimal search outcomes of a T-period problem as an extension of Example 2. The horizontal axes represent the period winning bid $y_t := \max_{i \in N_t}\{x_i\}$ at the end of each period t. There is a constant cutoff price p^0 for allocating the product throughout all periods $t < T$, and it drops to r_T in the last period T. For the T-period problem, the sequence of second price auctions runs as follows. At the beginning of any period $t < T$, if the period $t-1$ winning bid $y_{t-1} \geq p^0$, the seller stops and allocates the product to the highest bidder of the set N_{t-1} of bidders; otherwise, the seller invites one new bidder in period t who will submit a bid upon his arrival. At the beginning of the last period T, however, if $y_{T-1} \geq p^0$, the seller then stops and allocates the product. If $p_{T-1}^1 \leq y_{T-1} < p^0$, the seller invites one more bidder; if $y_{T-1} < p_{T-1}^1$, the seller invites two more bidders in period T. The period-T entrants submit their bids and the highest bidder in N_T wins and his reserve price is determined by (17).

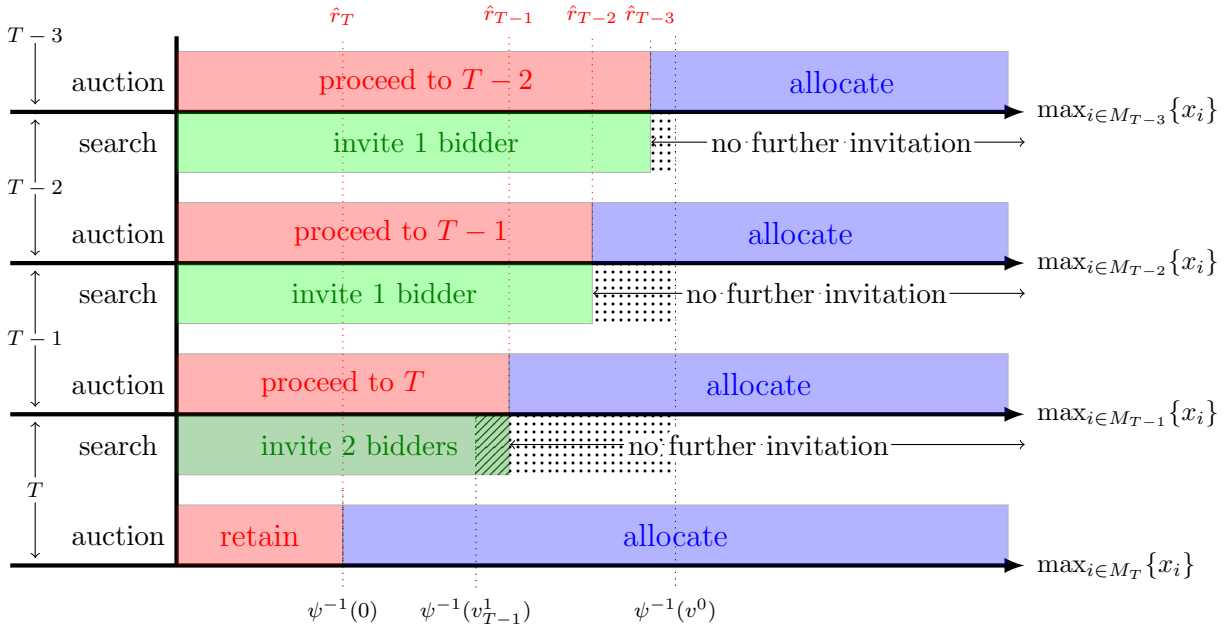
6 Optimal Search for Short-lived Bidders

Short-lived bidders participate in the transaction only for one period, making the implementation problem relatively simple. Again, we focus on direct revelation mechanisms, and characterise the optimal selling mechanism. We then show how to implement it with a sequence of second-price auctions with discriminatory reserve prices that depend solely on the clock (i.e, the bidder's entry period), but *not* the counter.

Figure 4: Optimal Search Auctions: Long vs. Short-lived Bidders



(a) Long-lived Bidders: The reserve price p^0 for allocating the product is constant throughout periods $t < T$ and drops to r_T in the last period T.



(b) Short-lived Bidders: The reserve price $\hat{r}_t$ is determined by the cutoff for optimal stopping in the next period, and is decreasing over time. Compared to (a), the dotted areas capture the discouragement effect of short-lived bidders, and the green slashed area in period T highlights the encouragement effect.

Specifically, at the beginning of period t, given that all active bidders $i \in M_{t-1}$ report truthfully, the seller then has a fallback revenue $v := \max_{i \in M_{t-1}}\{\psi(x_i)\} \geq 0$ which she can claim by stopping. If she continues to invite m bidders in period t, the fallback revenue at the beginning of period $t+1$ becomes $\max\{0, V^m\}$, as bidders are short-lived and she has to forgo the revenue v. Let $\hat{\Pi}_t(v)$ denote the optimal search profit the seller can achieve by following an optimal search rule from the beginning of period t on. For short-lived bidders, the Bellman equation of the search problem is then, for $t \leq T$,

$$\hat{\Pi}_t(v) = \max\left\{v, \max_{m \in \mathbb{N}_0}\{\mathbb{E}\hat{\Pi}_{t+1}\left(\max\{0, V^m\}\right) - cm\}\right\}, \tag{22}$$

and $\hat{\Pi}_{T+1}(v) = v$ obviously. Again, we solve the problem by backward induction.

At the beginning of the last period T, suppose the seller has a fallback revenue v from the set M_{T-1} of bidders. If she stops, she can claim the revenue v. If she continues to search, the expected revenue of inviting m short-lived bidders in period T is

$$\mathbb{E}\max\{0, V^m\} = \int_0^{\bar{v}} (1 - G(z)^m)dz = R_T(m; 0),$$

where the last equality is from (3) in the case of long-lived bidders. Note that $R_T(m; 0)$ is the period-T search revenue for long-lived bidders with a fallback revenue 0, and the marginal revenue of inviting one more bidder is then equal to $MR_T(m; 0)$, as given in (4). Hence, if the seller continues to search for short-lived bidders in period T, the optimal sample size is given by $\hat{m}_T^* := m_T^*(0)$, which is independent of the fallback revenue v from the previously invited bidders.

Let $\hat{v}_T$ be the continuation search profit for short-lived bidders in period T, that is,

$$\hat{v}_T := \max_{m \in \mathbb{N}_0}\{\mathbb{E}\max\{0, V^m\} - cm\} = R_T(m_T(0); 0) - c \cdot m_T(0) = \Pi_T(0).$$

From (22), the optimal search profit in period T is

$$\hat{\Pi}_T(v) = \max\left\{v, \max_{m \in \mathbb{N}_0}\{\mathbb{E}\max\{0, V^m\} - cm\}\right\} = \max\{v, \hat{v}_T\},$$

and it is optimal to stop at the beginning of T if and only if the fallback revenue $v \geq \hat{v}_T$. That is, the cutoff revenue for optimal stopping in period T is equal to the continuation search profit in that period. For short-lived bidders, the equivalence between the cutoff revenue for optimal stopping and the continuation search profit holds in any period as shown in Theorem 2.

Let $\{\hat{v}_t\}_{t=1}^T$ denote the sequence of cutoff revenue for optimal stopping and $\{\hat{m}_t^*\}_{t=1}^T$ the sequence of optimal sample sizes. Theorem 2 characterises the optimal search mechanism for short-lived bidders, which is featured by decreasing cutoff revenues $\hat{v}_t$ for optimal stopping and increasing sample sizes $\hat{m}_t^*$ over time.

Theorem 2 (Optimal search for short-lived bidders)**.** *In an incentive feasible mechanism, let v be the seller's fallback revenue at the beginning of period $t \leq T$. It is optimal to stop searching in period t iff $v \geq \hat{v}_t$, where the cutoff revenue $\hat{v}_t$ for stopping is recursively determined by*

$$\hat{v}_t = \Pi_T(\hat{v}_{t+1}) \tag{23}$$

with $\hat{v}_{T+1} \equiv 0$. Moreover, $\hat{v}_{t+1} < \hat{v}_t < v^0$ and $\lim_{T\to\infty} \hat{v}_1 = v^0$. The optimal sample sizes $\hat{m}_t^*$ in period t is given by

$$\hat{m}_t^* = m_T^*(\hat{v}_{t+1})\mathbb{1}(v < \hat{v}_t), \tag{24}$$

and $\hat{m}_t^*$ is increasing over time. The optimal search profit is

$$\hat{\Pi}_t(v) = \max\{v, \hat{v}_t\}. \tag{25}$$

Remark 6. The recursive relation (23) reveals an interesting link between the optimal search for short-lived and that for long-lived bidders, copied as follows

$$\hat{v}_t = \max_{m\in\mathbb{N}_0} \{\mathbb{E}\max\{\hat{v}_{t+1}, V^m\} - mc\} = \Pi_T(\hat{v}_{t+1}).$$

Note that $\hat{v}_t$ is the continuation search profit in period t, which is achieved by choosing the optimal sample size $\hat{m}_t^*$ in period t and then following the optimal search rule till the deadline. Equation (23) establishes a recursive relation between $\hat{v}_t$ and $\hat{v}_{t+1}$ through the optimal search profit function Π_T for long-lived bidders. To be specific, it states that, when a searcher with short-lived bidders decides to search in period t, she behaves as if she is searching for long-lived bidders in the last period T, yet with a fallback revenue $\hat{v}_{t+1}$. Moreover, the cutoff revenue $\hat{v}_t$ for optimal stopping can be expressed as a repeated composition of optimal one-shot profit $\Pi_T(0)$ for long-lived bidders, that is,

$$\hat{v}_t = \Pi_T(\hat{v}_{t+1}) = \Pi_T(\Pi_T(\hat{v}_{t+2})) = \cdots = \underbrace{\Pi_T \circ \cdots \circ \Pi_T}_{T-t \text{ times}}(\hat{v}_T) = \underbrace{\Pi_T \circ \cdots \circ \Pi_T}_{T-t+1 \text{ times}}(0). \tag{26}$$

Remark 7. It is intuitive that the continuation search profit $\hat{v}_t$ – and thus the cutoff revenue for optimal stopping – decreases over time, as the seller has fewer remaining search opportunities to improve her profit as the deadline approaches. This also implies that the optimal sample size $\hat{m}_t^*$ increases over time, as the marginal search revenue $MR_T(m; \hat{v}_{t+1})$ gets larger when $\hat{v}_{t+1}$ declines over time.

These monotonic properties of the optimal search rule fit many real-world observations. For example, in M&As, if bidders are short-lived, a seller will contact the smallest number of bidders with the highest reserve price at the beginning. Therefore, it may explain why negotiation is the dominant selling process in M&As. In academic recruiting, UK universities generally increase their recruitment intensities when the REF deadline approaches.

Naturally, job candidates on a waiting list may no longer be available with time going on, and therefore can be thought of as short-lived candidates. Our results may also explain the dynamics of recruitment intensities observed in the UK academic markets.

Remark 8. For short-lived bidders, the sequences of $\{\hat{v}_t\}_{t=1}^T$ and $\{\hat{m}_t^*\}_{t=1}^T$ are both deterministic and predetermined. This is because, when the seller continues to search, she has to forgo the current fallback revenue v, and therefore, v will not affect her choices of sample sizes in subsequent periods. In fact, both sequences of $\{\hat{v}_t\}_{t=1}^T$ and $\{\hat{m}_t^*\}_{t=1}^T$ are solely determined by the number of remaining periods. This is different from the case of long-lived bidders, where the sequence of optimal sample sizes is a random process, as the optimal sample size $m_t^*(v)$ depends on the realization of the fallback value v, which is random.

6.1 Implementation: Short-lived Bidders

Unlike long-lived bidders, short-lived bidders do not care about the number of future entrants. As such, the seller can induce truthful bidding simply by setting a reserve price whose virtual value equals to the cutoff revenue for optimal stopping. Since there is no intertemporal incentive issues, each bidder's bid does not influence the seller's search intensity or the level of future competition. Consequently, the seller does not need to use the "counter" to count the number of competitors when setting reserve prices. Instead, bidders' bid affects only the seller's stopping decision. Because the seller's cutoff revenue for stopping depends on the remaining number of search opportunities, it suffices to use the "clock" to track the bidder's entry time when determining reserve prices for short-lived bidders. Specifically, in each period t, the seller sets a discriminatory reserve price

$$\hat{r}_t := \psi^{-1}(\hat{v}_{t+1}) \tag{27}$$

such that its virtual value equals her continuation profit. By (26), the clock-based discriminatory reserve price can be expressed by the number of remaining search periods, that is,

$$\hat{r}_t = \psi^{-1}\big(\underbrace{\Pi_T \circ \cdots \circ \Pi_T}_{T-t \text{ times composition}}(0)\big).$$

Proposition 2 formally describes the implementation of the optimal search outcomes in Theorem 2 for short-lived bidders, using a sequence of second-price auctions. Its proof is standard and is omitted.

Proposition 2. *At the beginning of each period $t \le T$, given the seller's fallback revenue $v = \max\{0, \psi(\max_{i\in M_{t-1}}\{x_i\})\}$ from the outcome of period $t-1$ with the set M_{t-1} of bidders, she invites $\hat{m}_t^* = m_T^*(\hat{v}_{t+1})$ new bidders in period t and runs a second-price auction*

with a reserve price of $\hat{r}_t = \psi^{-1}(\hat{v}_{t+1})$ *as defined in* (27)*. The outcome of the auction proceeds as follows:*

- *If the highest bid among the* $\hat{m}_t^*$ *bidders is greater than or equal to the reserve price* $\hat{r}_t$*, the highest bidder wins the product and pays the maximum of the second-highest bid and the reserve price.*
- *If the highest bid is below* $\hat{r}_t$ *and* $t < T$*, the seller proceeds to period* $t+1$*, continuing the search for new bidders.*
- *If the highest bid is below* $\hat{r}_t$ *and* $t = T$*, the seller retains the product.*

Next, we reconsider Example 1 in the case of short-lived bidders. Example 3 shows how to derive the optimal search rule and profit for short-lived bidders and how to implement the optimal search outcomes by a sequence of second-price auctions.

Example 3 (Optimal search for short-lived bidders)**.** Consider the same setting of Example 1, yet the bidders are now short-lived.

- At the beginning of the last period $T = 2$, if the seller continues to search, she will choose the optimal sample size $\hat{m}_2^* = m_2^*(0) = 2$ by (13) in Example 1. From (23), the continuation search profit $\hat{v}_2 = \Pi_2(0) \approx 0.0967$ using equation (1). Hence, the seller's optimal search profit with a fallback revenue v at the beginning of period T is

$$\hat{\Pi}_2(v) = \max\{v, \hat{v}_2\}.$$

- At the beginning of the first period $t = 1$, if she continues to search, she will choose the optimal sample size $\hat{m}_1^* = m_2^*(\hat{v}_2) = 1$ by (13) again. The continuation search profit $\hat{v}_1 = \Pi_2(\hat{v}_2) \approx 0.1407$ using equation (1). Hence, the seller's optimal search profit with a fallback revenue 0 at the beginning of period 1 is

$$\hat{\Pi}_1(0) = \max\{0, \hat{v}_1\} = \hat{v}_1.$$

The optimal sample sizes $\hat{m}_t^*$ and the search profits $\hat{\Pi}_t$ are plotted in Figure 2 (in red).

The optimal search profit can be obtained by running a sequential second-price auction. At the beginning of the period $t = 1$, the seller's fallback revenue $v = 0$, which is smaller than the continuation search profit $\hat{v}_1 \approx 0.1407$. The seller will then invite $\hat{m}_1^* = m_2^*(\hat{v}_2) = 1$ bidder and run a second-price auction with a reserve price $\hat{r}_1 = \psi^{-1}(\hat{v}_2) \approx 0.5703$. If the bidder's bid $x \geq \hat{r}_1$, the seller then allocates the product to him at price $\hat{r}_1$; otherwise, the seller continues to search in period 2, by inviting $\hat{m}_2^* = m_2^*(0) = 2$ new bidders and running a second-price auction with a reserve price of $\hat{r}_2 = \psi^{-1}(0) = 0.5$.

The two-period example can be easily extended to any finite T-period problems. As shown above, one can recursively derive the optimal sample size $\hat{m}_t^*$ and continuation search profit $\hat{v}_t$ from the last period T. Note that the optimal sample size $\hat{m}_t^*$ is increasing over time, and $\hat{m}_{T-1}^* = 1$ in this example. It is then follows that, at the beginning of any period $t < T$, if the seller decides to search, it is optimal to invite just one bidder. Second, the cutoff revenue for optimal stopping $\hat{v}_t$ is strictly decreasing over time. The monotone properties of optimal sample sizes and the reserve prices are illustrated in Figure 4(b).

Similarly, for short-lived bidders, we can implement the optimal search profit by a sequence of second-price auctions, as illustrated in Figure 4(b). Specifically, the optimal reserve price in any period t is $\hat{r}_t = \psi^{-1}(\hat{v}_{t+1})$, and an invited bidder will bid his true product value. As illustrated in Figure 4(b), at the beginning of any period t, if the highest bid $\max_{i \in M_{t-1}}\{x_i\}$ of the period $t-1$ participating bidders is higher than the reserve price $\hat{r}_{t-1}$, the seller then allocates the product to the highest bidder in M_{t-1}; otherwise, she will continue by inviting $\hat{m}_t^*$ new bidders in period t and run an auction at the reserve price $\hat{r}_t$. The seller runs the auctions in this way till the last period T. Figure 4(b) also illustrates that $\hat{r}_t$ is decreasing and $\hat{m}_t^*$ is increasing over time.

6.2 Long-lived vs. short-lived bidders

We here present some comparative results between long-lived and short-lived bidders. With short-lived bidders, a seller cannot reclaim previously declined offers as those bidders already drop out of the transaction. This induces two seemingly opposite effects on a seller's optimal search decisions. On the one hand, a seller with short-lived bidders is more likely to stop and accept a lower offer than one with long-lived bidders. That is, for short-lived bidders, the cutoff revenue for optimal stopping is lower than that for long-lived bidders, i.e., $\hat{v}_t < v^0$. We term this as the *discouragement effect* of short-lived bidders – it discourages the seller from further searching by accepting a lower fallback revenue. On the other hand, conditional on search, a seller with short-lived bidders will search more intensively than one with long-lived bidders. This is because the seller's fallback revenue turns to 0 if she continues to search, and a lower fallback revenue will increase the marginal search revenue and hence the optimal sample size. We term this effect as the *encouragement effect* of short-lived bidders, which encourages the seller to search more intensively. Corollaries 1 and 2 formally summarize the two effects.

Corollary 1 (Discouragement Effect)**.** *A seller with short-lived bidders stops at a lower cutoff revenue than one with long-lived bidders in any period $t = 1, \cdots, T$. To be specific,*

$$v^0 > \hat{v}_1 > \hat{v}_2 > \cdots > \hat{v}_T > 0.$$

Corollary 2 (Encouragement Effect)**.** *Conditional on searching, a seller with short-lived bidders searches more intensively than one with long-lived bidders in any period* $t = 1, \cdots, T$. *That is, for any* $v < \hat{v}_t$, *we have* $\hat{m}_t^*(v) \geq m_t^*(v)$.

We can also compare the optimal search profit with long-lived and short-lived bidders. Intuitively, as a seller can reclaim a long-lived bidder's previous offer, long-lived bidders are more valuable, and hence, the optimal search profit with long-lived bidders is greater than that with short-lived bidders. We call the difference between them as the *value of bidder longevity*, which is shown to be positive and single-peaked.

Corollary 3 (Value of Bidder Longevity)**.** *The value of* $\Pi_t(v) - \hat{\Pi}_t(v)$ *is*

i) non-negative for all v *and* t*;*

ii) strictly increasing in $0 \leq v \leq \hat{v}_t$ *and strictly decreasing in* $\hat{v}_t \leq v \leq v^0$.

The comparative results between long-lived and short-lived bidders are illustrated in Example 4.

Example 4 (Optimal search: long vs. short-lived bidders)**.** Figure 2 summarizes the optimal search rules and profits for both cases of long and short-lived bidders, as derived in Example 1 and Example 3. The horizontal axes is the fallback revenue v, and the vertical one is the search profit. First, we observe that a seller with short-lived bidders has lower cutoffs for optimal stopping than one with long-lived bidders, i.e., $\hat{v}_2 < \hat{v}_1 < v^0$. This demonstrates the discouragement effect (Corollary 1), that is, a seller with short-lived bidders is willing to accept a lower reserve price and then stop. Second, we also observe that, conditional on searching, a seller with short-lived bidders searches more intensively. Particularly, when $v_2^1 < \hat{v}_2$, $\hat{m}_2^* = 2$ for short-lived bidders is strictly greater than $m_2^*(v) = 1$ for long-lived bidders. This reveals the encouragement effect (Corollary 2). Third, Figure 2 also illustrates that the value of bidder longevity is always positive and single-peaked. Particularly, when $v \geq v^0$, $\Pi_t(v) = \hat{\Pi}_t(v)$ as the seller always stops and keeps the fallback revenue v. When $v < v^0$, we have $\Pi_t(v) > \hat{\Pi}_t(v)$ for any t and any v, except for $\Pi_T(0) = \hat{\Pi}_T(0)$. Moreover, $\Pi_t(v) - \hat{\Pi}_t(v)$ is single-peaked, achieving its maximum at $\hat{v}_t$ (Corollary 3).

Remark 9. The results of Corollary 1 and 2 are also illustrated in Figure 4. By comparing Figure 4(a) and 4(b), first, we find that the optimal auction reserve prices for long-lived bidders are greater than those of short-lived bidders in each period, i.e., $r_t > \hat{r}_t$ for $t < T$ and $r_T = \hat{r}_T$ for $t = T$. Specifically, the dotted areas in Figure 4(b) highlight the discouragement effect (Corollary 1) that a seller with short-lived bidders is willing to accept a lower reserve price and stop in each period. Second, conditional on searching, we find

that a seller with short-lived bidders will search more intensively than one with long-lived bidders, i.e., $\hat{m}_t^* > m_t^*(v)$ if $v < \hat{v}_t$. The green slashed area in period T in Figure 4(b) highlights the encouragement effect (Corollary 2) by showing that $\hat{m}_T^* = 2 > m_T(v)^* = 1$ for $v_{T-1}^1 < v < \hat{v}_T$.

6.3 A Convergence Result

The optimal search rules and values between long- and short-lived bidders differ, as shown in Theorems 1 and 2, and discussed in Section 6.2. However, as $T \to \infty$, the difference between the two cases vanishes, and both converge to the stationary and infinite horizon (SIH) sequential search problem.

We first consider the optimal search auction with short-lived bidders. Theorem 2 establishes that the cutoff revenue for optimal stopping converges to $\lim_{T\to\infty} \hat{v}_1 = v^0$. Therefore, when $T \to \infty$, the recursive formula (26) implies that the cutoff revenue also converges $\hat{v}_t = v^0$ for *any* finite t. Moreover, from (24), the optimal search rule becomes

$$\hat{m}_t^* = \mathbb{1}(0 \leq v < v^0),$$

that is, one-by-one sequential search is optimal. The corresponding optimal search profit is

$$\hat{\Pi}_t(v) = \max\{v, v^0\}.$$

We next consider the optimal search auction with long-lived bidders. When $T \to \infty$, (11) implies $\lim_{T\to\infty} \mathcal{G}_{t+1} = 0$, and hence the optimal search profit with long-lived bidders also becomes

$$\Pi_t(v) = \max\{v, v^0\},$$

which is achieved again by one-by-one sequential search. Therefore, as $T \to \infty$, both cases converge to the optimal search rule and profits of SIH search problems.

7 Efficient Search

In many situations, the objective of a seller is not to maximize profit but social welfare. This section turns to the efficient search for bidders by a deadline, where a seller's objective is to maximize the expected welfare, which is the expected product value of the highest bidder net of gross search costs. Now the seller cares about a bidder's product value x itself rather than its virtual value $v = \psi(x)$ she can secure from a truthful bidder.

We will derive the efficient search rules and values for both long and short-lived bidders. Similarly, the efficient mechanisms can be implemented by a sequence of second-price auctions with appropriately set reserve prices and search rules. We will demonstrate

that a profit-maximizing search is characterised by over-searching compared to a welfare-maximizing search, in terms of either higher cutoff values for stopping or greater sample sizes. This result is robust in both cases of long- and short-lived bidders. Therefore, the inefficiency of an optimal search auction can stem from its inefficient search rule.

7.1 Efficient Search for Long-lived Bidders

When the seller's objective is to maximize welfare, she prioritizes the highest product value over product revenue. Let x be the fallback product value at the beginning of period t, i.e., the highest product value revealed by the N_{t-1} bidders, and $J_t(x)$ be the net search welfare achieved by following an efficient search rule from period t on.Similar to (2), the Bellman equation for an efficient search problem is as follows

$$J_t(x) = \max\left\{x, \max_{m\in\mathbb{N}_0}\left\{\mathbb{E}J_{t+1}\left(\max\left\{x, X^m\right\}\right) - mc\right\}\right\}, \tag{28}$$

where X^m is the highest product value from a sample of m bidders, with $J_{T+1}(x) = x$. We define the expected search welfare of inviting m bidders in period t by

$$W_t(m; x) := \mathbb{E}J_{t+1}\left(\max\left\{x, X^m\right\}\right), \tag{29}$$

and the marginal search welfare $MW_t(m; x) := W_t(m + 1; x) - W_t(m; x)$. It's obvious that $MW_t(m; x)$ is strictly decreasing in m and converges to 0 when $m \to \infty$, and is strictly decreasing in x. Let $m_t^{\mathrm{e}}(x)$ denote the efficient sample size in period t that maximizes the expected net search welfare $W_t(m; x) - mc$, and we have

$$m_t^{\mathrm{e}}(x) = \min\left\{m \in \mathbb{N}_0 : MW_t(m; x) \le c\right\}. \tag{30}$$

As before, for each $0 \le m < m_t^{\mathrm{e}}(0)$, there exists a unique cutoff product value x_t^m that solves $MW_t(m; x) = c$. Moreover, the cutoff value x_t^0 for efficient stopping is constant for all $t \le T$, and we denote $x^0 = x_t^0$ which is the unique solution to

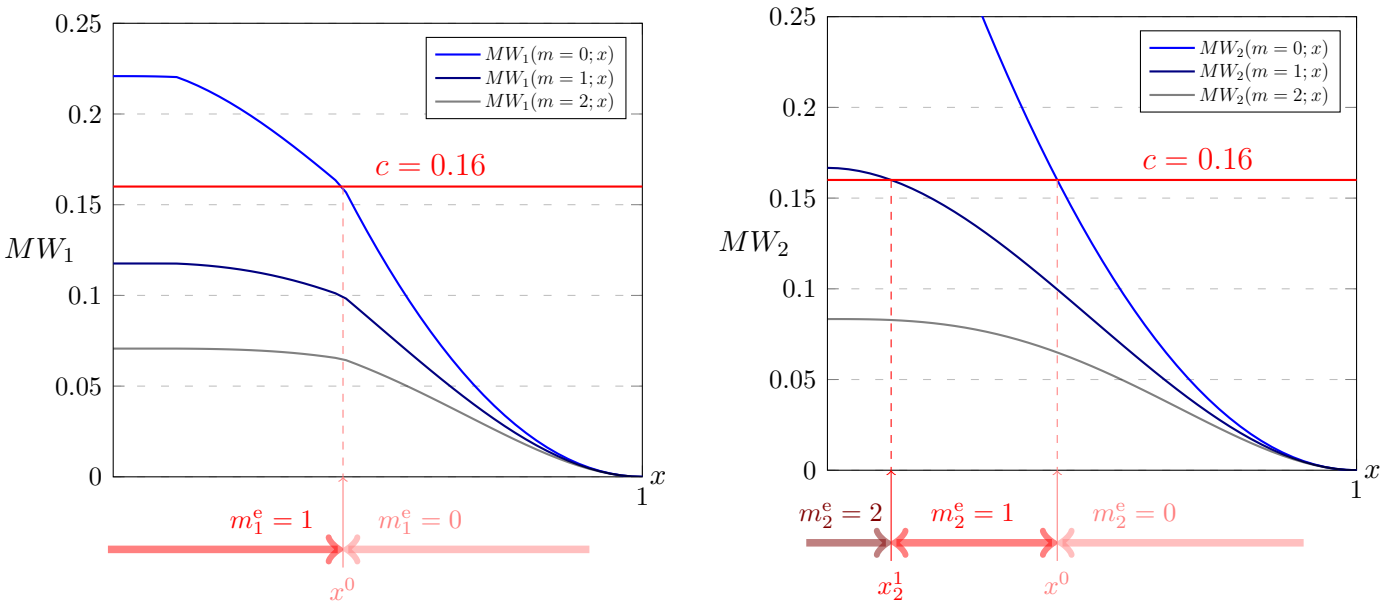
$$MW_T(0; x^0) = \int_{x^0}^{\bar{x}} (1 - F(z))dz = c. \tag{31}$$

Similarly, we have the following efficient search rule and net welfare for long-lived bidders.

Proposition 3 (Efficient Search for Long-lived Bidders)**.** *In an incentive feasible mechanism, let x be the seller's fallback product value at the beginning of period $t \le T$. The efficient sample size $m_t^e(x)$ in period t is given by*

$$m_t^e(x) = \min\{m \in \mathbb{N}_0 : MW_t(m; x) \le c\}, \tag{32}$$

Figure 5: Efficient search rule $m_t^e(x)$ in Example 5



(a) Period 1

(b) Period 2

Note. $c = 0.16$, $F(x) = x\mathbb{1}(0 \le x \le 1)$, and $T = 2$. In period 1, the efficient sample size is $m_1^e(x) = \mathbb{1}(x < x^0)$, where $x^0 = 0.4343$. In period 2, the efficient sample size is $m_2^e(x) = \mathbb{1}(x < x_2^1) + \mathbb{1}(x < x^0)$, where $x_2^1 \approx 0.1204$.

where the marginal search welfare $MW_t(m;x)$ is recursively determined by

$$MW_t(m;x) = \int_x^{\bar{x}} F(z)^m(1-F(z)) \prod_{\tau \ge t+1} F(z)^{m_\tau^e(z)} dz. \tag{33}$$

It is efficient to stop iff $x \ge x^0$. The maximum net search welfare $J_t(x)$ is

$$J_t(x) = \mathbb{E}\max\{x, X_t^e\} = x\mathcal{F}_t^e(x) + \int_x^{\bar{x}} z d\mathcal{F}_t^e(z), \tag{34}$$

where X_t^e is the continuation net welfare achieved by following the efficient search mechanism from period t on, and its distribution function $\mathcal{F}_t^e(x)$ is recursively defined by

$$\mathcal{F}_t^e(x) = F(x)^{m_t^e(x)}\mathcal{F}_{t+1}^e(x) = \prod_{\tau \ge t} F(x)^{m_\tau^e(x)}. \tag{35}$$

Remark 10. The property of $MW_t(m;x)$ guarantees that, for each $m = 0, 1, \ldots, m_t^e(0) - 1$, there exists a unique cutoff product value x_t^m that solves $MW_t(m;x) = c$. When $x \ge x_t^0$, $m_t^e(x) = 0$ and it is efficient to stop. As before, we can show $x_t^0 = x^0$ for all $t = 1, \cdots, T$, and therefore, there is a time-invariant cutoff product value for efficient stopping. Similar

to (12), we provide the following explicit expression of $m_t^e(x)$ in period t:

$$m_t^e(x) = \sum_{m=1}^{m_t^e(0)} \mathbb{1}(x < x_t^{m-1}) = \begin{cases} 0 & \text{if } x \geq x^0 \\ 1 & \text{if } x_t^1 \leq x < x^0 \\ 2 & \text{if } x_t^2 \leq x < x_t^1 \\ \vdots & \\ m_t^e(0) & \text{if } x < x_t^{m_t^e(0)-1}. \end{cases} \tag{36}$$

Remark 11. Similar to Section 5, the efficient search outcomes for long-lived bidders can be implemented by a sequence of second-price auctions with discriminatory reserve prices. For example, let $\ell_t^e(x) := \sum_{\tau=t+1}^{T} m_\tau^e(x)$ be the total number of new entrant bidders from period $t+1$ to T, given that x is the current highest bid and the seller follows the efficient sampling rule (32). For an admissible integer $\ell^e \in \mathbb{N}_0$, the cutoff price $p_t^{\ell^e} := \inf\{x \in \mathcal{X} : \ell_t^e(x) = \ell^e\}$. Suppose a bidder enters the auction in period $t \leq T$ and wins the auction at a bid x eventually, which induces a total number of $\ell^e = \ell_t^e(x)$ entrant bidders from period $t+1$ to T. The seller then set a reserve price $r_t^{\ell^e}$ for the bidder which is the unique solution to

$$\int_0^{p_t^{\ell^e}} \mathcal{F}_{t+1}^e(z)dz = \int_{r_t^{\ell^e}}^{p_t^{\ell^e}} F(z)^{\ell^e} dz,$$

where $\mathcal{F}_{t+1}^e$ is given by (35).

We next reconsider Example 1, yet the seller's objective is now to maximize expected net welfare. Example 5 shows how to construct an efficient search rule using the results of Proposition 3. In particular, we are interested in comparing the stopping cutoffs and sampling rules between the two cases of optimal and efficient search for long-lived bidders.

Example 5 (Efficient search for long-lived bidders)**.** Consider Example 1 again, yet we now examine an efficient, not an optimal, search rule for long-lived bidders.

- In the last period $t = 2$, the marginal search welfare is

$$MW_2(m; x) = \int_x^1 F(z)^m(1 - F(z))dz = \frac{1 - (m+2)x^{m+1} + (1+m)x^{m+2}}{(m+1)(m+2)},$$

which is plotted in Figure 5(b). As $MW_2(2; 0) = 1/16 < c$, the efficient sample size can not be greater than 2. For $m = 0, 1$, solving $MR_2(m; x) = c$ gives the cutoff values $x^0 = x_2^0 \approx 0.4343$ and $x_2^1 \approx 0.1204$. From (32), the efficient sample size is

$$m_2^e(x) = \mathbb{1}(x < x_2^1) + \mathbb{1}(x < x^0). \tag{37}$$

We have $\mathcal{F}_2^e(x) = F(x)^{m_2^e(x)}$, and the maximum net welfare in period $t = 2$ is

$$J_2(x) = \mathbb{E}\max\{x, X_2^e\} = x\mathcal{F}_2^e(x) + \int_x^{\bar{x}} z d\mathcal{F}_2^e(z). \tag{38}$$

- In the first period $t = 1$, the marginal search welfare is

$$MW_1(m; x) = \int_x^1 F(z)^m (1 - F(z)) \mathcal{F}_2^e(z) dz,$$

which is illustrated in Figure 5(a). It's interesting to observe that, for and $x \geq x^0$, $MW_2(m, x) = MW_1(m, x)$ for any m, and therefore, there is a time-invariant cutoff x^0 for efficient stopping. If $x < x^0$, the positive sample size in period $t = 2$ (i.e., $m_2^{\mathrm{e}}(x) \geq 1$), decreases the marginal search welfare in period 1 as $\mathcal{F}_2^e(x) < 1$, and hence, $MW_1(m, x) < MW_2(m, x)$. As $MW_1(1; 0) = \int_0^{x_2^1} z^3(1-z)dz + \int_{x_2^1}^{x^0} z^2(1-z)dz + \int_{x^0}^1 z(1-z)dz \approx 0.1428 < c$, the efficient sample size in period $t = 1$ is at most 1, and

$$m_1^{\mathrm{e}}(x) = \mathbb{1}(x < x^0).$$

Given $m_1^{\mathrm{e}}(x)$ and $m_2^{\mathrm{e}}(x)$, the continuation net search welfare X_1^{e} has the distribution

$$\mathcal{F}_1^e(x) = F(x)^{m_1^{\mathrm{e}}(x) + m_2^{\mathrm{e}}(x)} = \begin{cases} F(x)^3 & \text{if } x < x_2^1 \\ F(x)^2 & \text{if } x_2^1 \leq x < x^0 \\ F(x)^0 = 1 & \text{if } x \geq x^0. \end{cases}$$

The maximum net search welfare at the beginning of period $t = 1$ is

$$J_1(0) = \mathbb{E} \max\{0, X_1^{\mathrm{e}}\} = \int_0^1 z d\mathcal{F}_1^e = x^0 - \int_{x_2^1}^{x^0} F(z)^2 dz - \int_0^{x_2^1} F(z)^3 dz \approx 0.4075.$$

7.2 Efficient Search for Short-lived Bidders

Similar to Theorem 2, we can derive the efficient search rule and welfare for short-lived bidders. Let $\hat{x}_t$, $\hat{m}_t^e$ and $\hat{J}_t(x)$ respectively denote the cutoff value for efficient stopping, the efficient sample size, and the maximum net search welfare in the case of short-lived bidders. Proposition 4 below gives the results.

Proposition 4 (Efficient Search for Short-lived Bidders)**.** *In an incentive feasible mechanism, let x be the seller's fallback product value at the beginning of period $t \leq T$. It is efficient for the seller to stop in period t if and only if $x \geq \hat{x}_t$, where $\hat{x}_t$ is recursively determined by*

$$\hat{x}_t = J_T(\hat{x}_{t+1}) \tag{39}$$

Moreover, $\hat{x}_{t+1} < \hat{x}_t < x^0$ and $\lim_{T\to\infty} \hat{x}_1 = x^0$. The efficient sample sizes $\hat{m}_t^e(x)$ is given by

$$\hat{m}_t^e(x) = m_T^e(\hat{x}_{t+1}) \mathbb{1}(x < \hat{x}_t), \tag{40}$$

and $\hat{m}_t^e(x)$ is increasing over time, i.e., $\hat{m}_t^e \leq \hat{m}_{t+1}^e$. The maximum net search welfare is

$$\hat{J}_t(x) = \max\{x, \hat{x}_t\}. \tag{41}$$

Remark 12. Following the same spirit of Section 6.1, we can implement the efficient search mechanism for short-lived bidders by a sequence of second-price auctions with the sequences of reserve prices $\{\hat{x}_t\}_{t=1}^T$ and sample sizes $\{\hat{m}_t^e\}_{t=1}^T$.

Example 6 derives the efficient search rule and outcome with short-lived bidders. We are interested in comparing the efficient search rules across different cases, i.e., between short and long-lived bidders, and between efficient and optimal searches.

Example 6 (Efficient search for short-lived bidders)**.** Let us reconsider Example 5, yet now with short-lived bidders. Again we solve the problem by backward induction.

- At the beginning of the last period $t = 2$, if the seller continues to search, she will choose a sample size $\hat{m}_2^e = m_2^e(0) = 2$ as given in Example 5 and the continuation net search welfare is $\hat{x}_2 = J_2(0) \approx 0.3467$ from (38). The maximum net search welfare at the beginning of period $t = 2$ is hence $\hat{J}_2(x) = \max\{x, \hat{x}_2\}$.

- At the beginning of period $t = 1$, if the seller continues to search, she will choose a sample size $\hat{m}_1^e = m_2^e(\hat{x}_2) = 1$, and the continuation net search welfare is $\hat{x}_1 = J_2(\hat{x}_2) \approx 0.4001$ again from (38). The maximum net search welfare at the beginning of period $t = 1$ is hence $\hat{J}_1(0) = \max\{0, \hat{x}_1\}$.

The efficient mechanisms can be implemented by the following sequential second-price auction. At the beginning of period $t = 1$, as the fallback value 0 is smaller than $\hat{x}_1 \approx 0.4001$, the seller then invites $\hat{m}_1^{\mathrm{e}} = 1$ bidder (B_1) and runs a second-price auction with a reserve price $\hat{x}_2 \approx 0.3467$. If the bid is higher than $\hat{x}_2$, the seller then allocates the product to B_1 at the price $\hat{x}_2$. Otherwise, the seller continues to invite $\hat{m}_2^{\mathrm{e}} = 2$ bidders in period $t = 2$ and run a second-price auction with a zero reserve price.

7.3 Optimal vs. Efficient Search

The optimal search rules are different from the efficient ones, in both cases of long and short-lived bidders. Particularly, in terms of product value, we can show that the optimal cutoff value for stopping is higher than the efficient one, and therefore, an optimal search mechanism is featured by over-searching. This result echoes the similar result in static auctions, where the product value associated with the optimal reserve, i.e., $r_T = \psi^{-1}(0)$, is greater than one with the efficient reserve, i.e., 0. Our results then extend the static result to the more general case of sequential search mechanisms with a deadline.

Corollary 4 below states that, for long-lived bidders, the cutoff product value for optimal stopping is higher than that for efficient stopping. As a results, a profit-maximizing seller would over search bidders than a welfare-maximizing seller.

Corollary 4. *For long-lived bidders, the cutoff product value for optimal stopping is greater than that for efficient stopping, i.e.,* $v^0 > \psi(x^0)$.

For short-lived bidders, we have the similar result that the optimal sample size in period t is greater than the efficient one, if their continuation search values are equal. Therefore, the over-searching result for an optimal search auction also holds for short-lived bidders.

Corollary 5. *For short-lived bidders, the optimal sample size in period* t *is greater than the efficient one, i.e.,* $\hat{m}_t^* \geq \hat{m}_t^e$*, if their continuation search values are equal in the sense that* $\hat{v}_{t+1} = \psi(\hat{x}_{t+1})$.

8 Concluding Remarks

We have studied sequential search involving strategic agents under a time constraint, focusing on search auctions in which a seller searches for bidders before a finite deadline. Due to the time constraint, the decision-maker may simultaneously search for multiple agents in each period, and the search intensity adaptively evolves over time. In response, agents strategically adjust their behaviour to affect the decision-maker's adaptive search rule. To address the intertemporal incentive issues arising from the time constraint, the decision-maker employs two measures: the "clock" to track each agent's entry time (i.e., the expected market thickness in terms of the number of remaining search opportunities) and the "counter" to record the number of other agents (i.e., the ex-post market thickness realised after the agent's entry).

In addition to analysing the role of time constraints in sequential search for strategic agents, this paper makes a few contributions to the literature. First, our model enables a direct comparison between long-lived and short-lived bidders, uncovering clear links between the two cases. In contrast, existing models of search mechanisms typically focus on one or the other in isolation. Second, this paper contributes to the literature on dynamic mechanisms and revenue management, where the sequence of buyer arrivals is treated as exogenous. In our setting, by contrast, bidder arrivals are endogenously determined by the seller's optimal search decisions. Third, our work contributes to the longstanding debate over the choice between negotiation and auction as selling mechanisms. For instance, our results provide an alternative explanation for the observed dominance of negotiation in M&A transactions.

Our results are applicable to a wide range of real-world problems in which a decision-maker searches for strategic agents by a deadline. Below, we propose several illustrative applications.

Revenue Management The typical literature on revenue management usually makes two assumptions: (i) buyers are short-lived, and (ii) the sequence of buyer arrivals is exogenously given (Talluri and Van Ryzin, 2006). In an influential paper, Board and Skrzypacz (2016) study a revenue management problem with long-lived bidders, yet the sequence of buyer arrivals is still assumed to be exogenously given. With our framework, we can relax both assumptions (i) and (ii) and study the optimal allocation of products to potential buyers within T periods, where the buyers can be either short- or long-lived, and the seller's optimal search decisions endogenously determine the sequence of buyer arrivals.

M&As Our results can obviously be applied to the study of M&A and other sequential selling processes. A typical M&A process lasts weeks, even months, and the inter-temporal interplay between the target and the bidders is therefore of pivotal importance. In a two-bidder model of costly information acquisition, Fishman (1988) argues that an early bidder makes a pre-emptive bid mainly to deter a second bidder from competing. A similar argument appears in Bulow and Klemperer (2009) who state that a simultaneous auction is better than a sequential negotiation in M&As. Also in a two-bidder model of costly information acquisition, Chen and Wang (2023) show that the optimal selling mechanism reveals a favouritism towards the initial bidder who has an endogenous informational disadvantage at the time of entry. Our approach is different. We focus on a seller's costly search for potential bidders, and the seller imposes higher reserve prices on earlier entrant bidders to offset their inter-temporal incentives for overbidding. Importantly, our results provide a theoretical explanation for the dominance of negotiation in M&A selling processes.

Recruiting Campaign Recruitment campaigns are usually under time constraints. For example, a football club needs to sign new players within a transfer window, and the UK academic job markets are usually driven by the REF (Research Excellence Framework) cycles. Our model can be applied to study various recruitment campaigns with a deadline. For example, an employer needs to fill a position by a deadline. She invites job candidates batch by batch for interviews, each at a search cost. Job candidates can be long-lived or short-lived, depending on the job market situation. After each round of interviews, the employer decides whether or not to make an offer. If the employer does not make an offer and continues to search, she decides how many job candidates to invite in the next round.

References

Aktas, N., E. De Bodt, and R. Roll (2010): "Negotiations under the threat of an auction," *Journal of Financial Economics*, 98, 241–255.

Andrade, G., M. Mitchell, and E. Stafford (2001): "New evidence and perspectives on mergers," *Journal of economic perspectives*, 15, 103–120.

Armstrong, M. (2017): "Ordered Consumer Search," *Journal of the European Economic Association*, 15, 989–1024.

Baucells, M. and L. Zhao (2018): "It Is Time to Get Some Rest," *Management Science*, 65, 1717–1734.

Bergemann, D. and J. Välimäki (2019): "Dynamic mechanism design: An introduction," *Journal of Economic Literature*, 57, 235–274.

Betton, S., B. E. Eckbo, and K. S. Thorburn (2008): "Corporate takeovers," *Handbook of corporate finance: Empirical corporate finance*, 2, 291–430.

Board, S. and A. Skrzypacz (2016): "Revenue management with forward-looking buyers," *Journal of Political Economy*, 124, 1046–1087.

Boone, A. L. and J. H. Mulherin (2007): "How are firms sold?" *The Journal of Finance*, 62, 847–875.

——— (2008): "Do auctions induce a winner's curse? New evidence from the corporate takeover market," *Journal of Financial Economics*, 89, 1–19.

Bulow, J. and P. Klemperer (2009): "Why do sellers (usually) prefer auctions?" *American Economic Review*, 99, 1544–75.

Chaves, I. N. and S. Ichihashi (2024): "Auction timing and market thickness," *Games and Economic Behavior*, 143, 161–178.

Che, X. and T. Klumpp (2016): "Entry deterrence in dynamic second-price auctions," *American Economic Journal: Microeconomics*, 8, 168–201.

Che, Y.-K. and A. B. Choi (2025): "Optimal Auction Design for Dynamic Stochastic Environments: Myerson Meets Naor," *arXiv preprint arXiv:2505.22862*.

Chen, Y. and Z. Wang (2023): "Optimal sequential selling mechanism and deal protections in mergers and acquisitions," *The Journal of Finance*, 78, 2139–2188.

Choi, M., A. Y. Dai, and K. Kim (2018): "Consumer search and price competition," *Econometrica*, 86, 1257–1281.

Crémer, J., Y. Spiegel, and C. Z. Zheng (2007): "Optimal search auctions," *Journal of Economic Theory*, 134, 226–248.

Du, L., M. Hu, and J. Wu (2022): "Sales effort management under all-or-nothing constraint," *Management Science*, 68, 5109–5126.

Fishman, M. J. (1988): "A theory of preemptive takeover bidding," *The Rand Journal of Economics*, 88–101.

Gal, S., M. Landsberger, and B. Levykson (1981): "A compound strategy for search in the labor market," *International Economic Review*, 597–608.

Gittins, J. C. (1979): "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 41, 148–164.

Kanoria, Y. and H. Nazerzadeh (2020): "Dynamic reserve prices for repeated auctions: Learning from bids," *arXiv preprint arXiv:2002.07331*.

Kleinberg, B., B. Waggoner, and E. G. Weyl (2017): "Descending price optimally coordinates search," *Working Paper*.

Lauermann, S. and A. Wolinsky (2016): "Search with adverse selection," *Econometrica*, 84, 243–315.

Lee, J. and D. Z. Li (2022): "Sequential search with adaptive intensity," *International Economic Review*, 63, 803–829.

——— (2023): "Seller compound search for bidders," *The Journal of Industrial Economics*, 71, 1004–1037.

Li, A. (2009): "Selling storable goods to a dynamic population of buyers: A mechanism design approach," *Unpublished manuscript, Stanford University*.

Lippman, S. A. and J. J. McCall (1976): "Job search in a dynamic economy," *Journal of Economic Theory*, 12, 365–390.

Lu, J., L. Ye, and X. Feng (2021): "Orchestrating information acquisition," *American Economic Journal: Microeconomics*, 13, 420–465.

McAfee, R. P. and J. McMillan (1988): "Search mechanisms," *Journal of Economic Theory*, 44, 99–123.

Mierendorff, K. (2016): "Optimal dynamic mechanism design with deadlines," *Journal of Economic Theory*, 161, 190–222.

Morgan, P. and R. Manning (1985): "Optimal search," *Econometrica: Journal of the Econometric Society*, 923–944.

Morgan, P. B. (1983): "Search and optimal sample sizes," *The Review of Economic Studies*, 50, 659–675.

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of operations research*, 6, 58–73.

Pai, M. M. and R. Vohra (2013): "Optimal dynamic auctions and simple index rules," *Mathematics of Operations Research*, 38, 682–697.

Said, M. (2011): "Sequential auctions with randomly arriving buyers," *Games and Economic Behavior*, 73, 236–243.

——— (2012): "Auctions with dynamic populations: Efficiency and revenue maximization," *Journal of Economic Theory*, 147, 2419–2438.

Szech, N. (2011): "Optimal advertising of auctions," *Journal of Economic Theory*, 146, 2596–2607.

Talluri, K. T. and G. J. Van Ryzin (2006): *The theory and practice of revenue management*, vol. 68, Springer Science & Business Media.

Weitzman, M. L. (1979): "Optimal search for the best alternative," *Econometrica: Journal of the Econometric Society*, 641–654.

Xu, M. and D. Z. Li (2019): "Equilibrium competition, social welfare and corruption in procurement auctions," *Social Choice and Welfare*, 53, 443–465.

Zhang, H. (2021): "The optimal sequence of prices and auctions," *European Economic Review*, 133, 103681.

Appendix: Omitted Proofs

Proof of Lemma 1. It follows from (7) that

$$MR_t(m;v) = \mathbb{E}\Pi_{t+1}\left(\max\left\{v, V^{m+1}\right\}\right) - \mathbb{E}\Pi_{t+1}\left(\max\left\{v, V^m\right\}\right).$$

Noting that V's are independent draws, we have $\max\{v, V^{m+1}\} = \max\{v, \max\{V^m, V\}\} = \max\{\max\{v, V^m\}, V\}$. As Π_{t+1} is strictly increasing, we have

$$\begin{aligned}&\Pi_{t+1}\left(\max\{\max\{v, V^m\}, V\}\right) - \Pi_{t+1}\left(\max\left\{v, V^m\right\}\right)\\ &= \begin{cases}\Pi_{t+1}(V) - \Pi_{t+1}(\max\{v, V^m\}) & \text{if } V \geq \max\{v, V^m\}\\ 0 & \text{if } V < \max\{v, V^m\}.\end{cases}\end{aligned}$$

As a result, an alternative expression of the marginal revenue is

$$MR_t(m;v) = \mathbb{E}\max\{\Pi_{t+1}(V) - \Pi_{t+1}(\max\{v, V^m\}), 0\},$$

which directly implies the results (i)-(iii) in the Lemma. □

Proof of Theorem 1. Assume $MR_T(0,0) > c$, otherwise the problem becomes trivial as the seller never invites any bidder. Let $\pi_t^*(v) := \max_{m\in\mathbb{N}_0}\{R_t(m;v) - cm\}$. We prove the results by backward induction.

Step 1: For the last period T, substituting $\Pi_{T+1}(v) = v$ into (7), the expected revenue is

$$R_T(m;v) = \mathbb{E}\max\left\{v, V^m\right\} = v + \int_v^{\bar{v}} (1 - G(w)^m)dw, \tag{42}$$

and $MR_T(m;v) = \int_v^{\bar{v}} G(w)^m(1-G(w))dw$. Let $v^0 \equiv v_T^0$, which is the unique value satisfying $MR_T(0,v^0) = \int_{v^0}^{\bar{v}}(1-G(w))dw = c$. We have $0 < v^0 < \bar{v}$ as $MR_T(0,\bar{v}) < c < MR_T(0,0)$. As $\pi_T^*(v)$ is a contraction on $[0,v^0]$, there exists a unique fixed point on $[0,v^0]$ such that $\pi_T^*(v) = v$ due to the Banach fixed point theorem. In fact, this fixed point is v^0 as $m_T^*(v) = 0$ for $v \geq v^0$, and $\pi_T^*(v) = v$. Therefore, the optimal search profit in period T is

$$\Pi_T(v) = \max\{v, \pi_T^*(v)\} = \begin{cases} v & \text{if } v \in [v^0, \bar{v}]\\ \pi_T^*(v) = R_T(m_T^*(v);v) - cm_T^*(v) & \text{if } v \in [0, v^0),\end{cases}$$

which is obtained by the optimal search rule described in (8).

We next show that $\Pi_T(v)$ is convex and can be represented as (10). First, as $\partial R_T(m;v)/\partial v = G(v)^m$ is increasing in v, $R_T(m;v) - cm$ is convex in v for any given m; and $\pi_T^*(v)$, as the maximum of a family of convex functions, is hence also convex in v. Second, the optimal profit $\Pi_T(v)$ is the maximum of two convex functions. It is then convex and absolutely

continuous, and its derivative $\Pi_T'(v)$ is defined almost everywhere.As such, $\Pi_T(v)$ can be represented by a definite integral of its derivative, i.e.,

$$\Pi_T(v) = \Pi_T(\bar{v}) - \int_v^{\bar{v}} \Pi_T'(w)dw = \bar{v} - \int_v^{\bar{v}} \Pi_T'(w)dw,$$

where we apply $\Pi_T(\bar{v}) = \bar{v}$. As $\partial R_T(m;v)/\partial v = G(v)^m$, applying envelop theorem gives

$$\Pi_T'(v) = \begin{cases} 1 & \text{if } v \in [v^0, \bar{v}] \\ G(v)^{m_T^*(v)} & \text{if } v \in [0, v^0), \end{cases}$$

which is equivalent to $\mathcal{G}_T(v)$ as given in (11). The optimal search profit $\Pi_T(v)$ is then represented by (10). Note that $\mathcal{G}_T(x) = \Pi_T'(x)$ is increasing in x as $\Pi_T(x)$ is convex.

Step 2: For $t < T$, as an induction hypothesis, assume

$$\Pi_{t+1}(v) = \bar{v} - \int_v^{\bar{v}} \mathcal{G}_{t+1}(w)dw, \tag{43}$$

where $\mathcal{G}_{t+1}(v)$ is given by (11) and is increasing in x. Integrating by parts then gives

$$\begin{aligned} R_t(m;v) &= \mathbb{E}\Pi_{t+1}(\max\{v, V_m\}) \\ &= \Pi_{t+1}(v)G(v)^m + \int_v^{\bar{v}} \Pi_{t+1}(w)dG(w)^m \\ &= \bar{v} - \int_v^{\bar{v}} G(w)^m \mathcal{G}_{t+1}(w)dw, \end{aligned}$$

and $MR_t(m;v) = \int_v^{\bar{v}} G(w)^m(1-G(w))\mathcal{G}_{t+1}(w)dw$. At the cutoff revenue v^0, we still have $MR_t(0;v^0) = c$ as $\mathcal{G}_{t+1}(v) = 1$ for $v \geq v^0$. Again, $\pi_t^*(v)$ is a contraction on $[0, v^0]$, on which v^0 is the unique fixed point. Therefore, for $t < T$, the optimal search profit is

$$\Pi_t(v) = \max\{v, \pi_t^*(v)\} = \begin{cases} v & \text{if } v \in [v^0, \bar{v}] \\ \pi_t^*(v) = R_{t+1}(m_t^*(v);v) - cm_t^*(v) & \text{if } v \in [0, v^0), \end{cases}$$

which is obtained by the optimal search rule described in (8).

Next, we have $\partial R_{t+1}(m;v)/\partial v = G(v)^m \mathcal{G}_{t+1}(v)$, which is increasing in v. As in **Step 1**, $R_{t+1}(m;v) - cm$ is convex in v for any given m, and hence $\pi_t^*(v)$ is also convex in v. The optimal profit $\Pi_t(v)$, as the maximum of two convex functions, is convex and absolutely continuous. Hence, its derivative $\Pi_t'(v)$ is defined almost everywhere and $\Pi_t(v)$ can be represented by

$$\Pi_t(v) = \Pi_t(\bar{v}) - \int_v^{\bar{v}} \Pi_t'(w)dw = \bar{v} - \int_v^{\bar{v}} \Pi_t'(w)dw,$$

where we apply $\Pi_t(\bar{v}) = \bar{v}$. As $\partial R_{t+1}(m;v)/\partial v = G(v)^m \mathcal{G}_{t+1}(v)$, envelop theorem yields

$$\Pi_t'(v) = \begin{cases} 1 & \text{if } v \in [v^0, \bar{v}] \\ G(v)^{m_t^*(v)} \mathcal{G}_{t+1}(v) & \text{if } v \in [0, v^0), \end{cases}$$

which is equivalent to $\mathcal{G}_t(v)$ given in (11). □

Proof of Theorem 2. We solve the search problem (22) by backward induction. For $t \leq T$, define

$$\hat{v}_t := \max_{m \in \mathbb{N}_0} \left\{ \mathbb{E}\hat{\Pi}_{t+1} \left(\max\{0, V^m\} \right) - mc \right\}, \tag{44}$$

which is the continuation search profit of following an optimal search rule from period t on, and $\hat{v}_{T+1} \equiv 0$. Substituting (44) into (22) then gives (25), i.e., $\hat{\Pi}_t(v) = \max\{v, \hat{v}_t\}$. It's clear that the seller stops searching in period t if and only if $v \geq \hat{v}_t$.

Starting from the last period T, substituting $\hat{\Pi}_{T+1}(v) = v$ into (44) gives

$$\hat{v}_T = \max_{m \in \mathbb{N}_0} \{ \mathbb{E} \max\{0, V^m\} - mc \} = \Pi_T(0) = \Pi_T(\hat{v}_{T+1}) > 0,$$

where we apply $\hat{v}_{T+1} = 0$. The strict inequality is from the fact that, for $v \in [0, v^0)$, $v < \Pi_T(v) < v^0$ (Remark 4). It then confirms (23) and that $v^0 > \hat{v}_t > \hat{v}_{t+1}$ for $t = T$. Suppose the induction hypotheses hold for periods $t+1, \cdots, T$. In period t, we have

$$\hat{v}_t = \max_{m \in \mathbb{N}_0} \left\{ \mathbb{E}\hat{\Pi}_{t+1} \left(\max\{0, V^m\} \right) - mc \right\} = \max_{m \in \mathbb{N}_0} \{ \mathbb{E} \max\{\hat{v}_{t+1}, V^m\} - mc \} = \Pi_T(\hat{v}_{t+1}),$$

where the second equality is from (25) and the fact that $\hat{v}_{t+1} > 0$, e.g.,

$$\hat{\Pi}_{t+1} \left(\max\{0, V^m\} \right) = \max \left\{ \max\{0, V^m\}, \hat{v}_{t+1} \right\} = \max \left\{ \hat{v}_{t+1}, V^m \right\}.$$

Again $\hat{v}_{t+1} < \hat{v}_t < v^0$, and the induction hypotheses hold for any period $t \leq T$. Furthermore, by monotone convergence theorem and that $\Pi_T(v^0) = v^0$, we have $\lim_{T \to \infty} \hat{v}_1 = v^0$.

We next solve for the optimal sample size $\hat{m}_t^*$ when the seller continues to search. Similarly, for the payoff function on the RHS of (44), we define the search revenue

$$\hat{R}_t(m; v) := \mathbb{E}\hat{\Pi}_{t+1}(\max\{0, V^m\}) = R_T(m, \hat{v}_{t+1}).$$

The maximizer $\hat{m}_t^*$ is then given by

$$\hat{m}_t^* = \min \{ m \in \mathbb{N}_0 : MR_T(m, \hat{v}_{t+1}) \leq c \} = m_T^*(\hat{v}_{t+1}).$$

As $\hat{v}_t > \hat{v}_{t+1}$, we have $m_T^*(\hat{v}_t) \leq m_T^*(\hat{v}_{t+1})$, which implies that $\hat{m}_t^* \leq \hat{m}_{t+1}^*$. □

Proof of Proposition 1. By construction, the invitation rule (14) implements the optimal sampling rule (8) if bidders bid truthfully. To complete the proof, we just need to show that the search auction with the discriminatory reserve price set by (17) is incentive compatible.

Consider a bidder of type (x, t), and suppose x is the highest product value among the set N_t of bidders and that all other bidders bid truthfully. Denote $U_t(x'; x)$ the bidder's expected payoff of bidding x', and his expected payoff of truthful bidding is then $U_t(x) := U_t(x; x)$. To show incentive compatibility, it suffices to show $U_t(x) \geq U_t(x'; x)$ for all $x' \in \mathcal{X}$.

Let $\ell_- < \ell < \ell_+$ be three adjacent admissible numbers in L_t and $x \in (p_t^{\ell_+}, p_t^{\ell}]$. If the bidder bids truthfully, then it induces a total number $\ell = \ell_t(x)$ of entrant bidders in subsequent periods and a reserve price r_t^{ℓ}. There are three possible scenarios of deviation.

i) Any deviating bid $x' \in (p_t^{\ell_+}, p_t^{\ell}]$ does not change the total number of entrant bidders in subsequent periods, which is still ℓ, nor the reserve price r_t^{ℓ}. Therefore, bidding true value x is a weakly dominant strategy as in a standard second-price auction.

ii) A downward deviating bid $x' \leq p_t^{\ell_+}$ has two effects: an increase of the total number of future entrant bidders from ℓ to ℓ_+, and a decrease of reserve price from r_t^{ℓ} to $r_t^{\ell_+}$. That is, underbidding may lower his reserve price yet induce more competition in subsequent periods. However, observing that

$$\begin{aligned} U_t(x';x) &= U_t(x') + \int_{y'}^{y} F(x)^{\ell_+} dx = U_t(p_t^{\ell}) + \int_{p_t^{\ell}}^{y} F(x)^{\ell_+} dx \\ &< U_t(p_t^{\ell}) + \int_{p_t^{\ell}}^{y} F(x)^{\ell} dx = U_t(x), \end{aligned}$$

the bidder has no incentive to underbid $x' \leq p_t^{\ell_+}$.

iii) An upward deviating bid $x' > p_t^{\ell}$ has two effects: a decrease of the total number of future entrant bidders from ℓ to ℓ_- and an increase of reserve price from r_t^{ℓ} to $r_t^{\ell_-}$. That is, by overbidding, the bidder may reduce competition in subsequent periods yet face a higher reserve price. As $r_t^{\ell_-}$ is constructed by (17), we have

$$\begin{aligned} U_t(x';x) &= \int_{r_t^{\ell_-}}^{y} F(x)^{\ell_-} dx - (F(x')^{\ell_-} - F(x)^{\ell-1})(y'-y) \\ &< \int_{r_t^{\ell_-}}^{x} F(z)^{\ell_-} dz = \int_{r_T}^{y} \mathcal{F}_{t+1}^{*}(z) dz = U_t(x). \end{aligned}$$

Therefore, the bidder has no incentive to overbid $x' > p_t^{\ell}$ either. □

Proof of Corollary 1. The proof is already given in the proof of Theorem 2. □

Proof of Corollary 2. When $v < \hat{v}_t$, it is optimal to search in both cases of short and long-lived bidders. The result is then clear by comparing Theorem 1 and Theorem 2. □

Proof of Corollary 3. i) It is easy to show that $\Pi_T(v) \geq \hat{\Pi}_T(v)$. Then the result $\Pi_t(v) \geq \hat{\Pi}_t(v)$ can be recursively derived from the last period T, using the Bellman equations of (2) and (22) and that $\Pi_{t+1}(v) \geq \hat{\Pi}_{t+1}(v)$. ii) Note that $\hat{\Pi}_t(v) = \max\{v, \hat{v}_t\}$ by (25), and $\Pi_t(v) = \mathbb{E}\max\{v, V_t^*\}$ where V_t^* has the distribution $\mathcal{G}_t(v) = \prod_{\tau \geq t} G(v)^{m_\tau^*(v)}$ by (11). The single-peak result then comes from comparing the slopes of $\hat{\Pi}_t(v)$ and $\Pi_t(v)$. □

Proof of Proposition 3. Let $w_t^*(x) := \max_{m \in \mathbb{N}_0}\{W_t(m;x) - cm\}$. We prove the results by backward induction.

Step 1: Let $t = T$. Substituting $J_{T+1}(x) = x$ into (29), the expected search welfare is

$$W_T(m;x) = \mathbb{E}\max\{x, X^m\} = x + \int_x^{\bar{x}} (1 - F(z)^m) dz,$$

and $MW_T(m;x) = \int_x^{\bar{x}} F(z)^m(1-F(z))dz$. Denote $x^0 < \bar{x}$ the unique value satisfying $MW_t(0;x^0) = \int_{x^0}^{\bar{x}}(1-F(z))dz = c$. Note $w_T^*(x)$ is a contraction on $[0,x^0]$, and there is a unique fixed point on $[0,x^0]$ such that $w_T^*(x) = x$. In fact this fixed point is x^0, as $m_T^{\mathrm{e}}(x) = 0$ for $z \geq x^0$ and $w_t^*(x) = x$. Therefore, the net search welfare function is

$$J_T(x) = \max\{x, w_T^*(x)\} = \begin{cases} x & \text{if } x \in [x^0, \bar{x}], \\ w_T^*(x) = W_T(m_T^{\mathrm{e}}(x);x) - cm_T^{\mathrm{e}}(x) & \text{if } x \in [0, x^0), \end{cases}$$

which is obtained by the efficient search rule described in (32).

We next show that $J_T(x)$ is convex and can be represented as (34). First, as $\partial W_T(m;x)/\partial x = F(x)^m$ is increasing in x, $W_T(m;x)$ is convex in x for any given m. $w_T^*(x)$, as the maximum of a family of convex functions, is hence convex in x. Second, the net search welfare $J_T(x)$ is the maximum of two convex functions. It is then convex and absolutely continuous, and its derivative $W_T'(x)$ is defined almost everywhere, and $J_T(x)$ can be represented by a definite integral of its derivative, i.e.,

$$J_T(x) = J_T(\bar{x}) - \int_x^{\bar{x}} W_T'(z)dz = \bar{x} - \int_v^{\bar{v}} \Pi_T'(z)dz,$$

where we apply $J_T(\bar{x}) = \bar{x}$. As $\partial w_T(m;x)/\partial x = F(x)^m$, applying envelop theorem gives

$$W_T'(x) = \begin{cases} 1 & \text{if } x \in [x^0, \bar{x}], \\ F(x)^{m_T^{\mathrm{e}}(x)} & \text{if } x \in [0, x^0), \end{cases}$$

which is equivalent to $\mathcal{F}_T^e(x)$ as given in (35). The net search welfare $J_T(x)$ is then represented by (34). Note that $\mathcal{F}_T^e(x) = W_T'(x)$ is increasing in x as $J_T(x)$ is convex.

Step 2: For $t < T$, as an induction hypothesis, assume

$$J_{t+1}(x) = \bar{x} - \int_x^{\bar{x}} \mathcal{F}_{t+1}^e(z)dz,$$

where $\mathcal{F}_{t+1}^e(z)$ is given by (35) and is increasing in x. Integrating by parts then gives

$$W_t(m;x) = \bar{x} - \int_x^{\bar{x}} F(z)^m \mathcal{F}_{t+1}^e(z)dz,$$

and $MW_t(m;x) = \int_x^{\bar{x}} F(z)^m(1-F(z))\mathcal{F}_{t+1}^e(z)dz$. For the cutoff value x^0, we still have $MW_t(0;x^0) = c$ as $\mathcal{F}_{t+1}^e(z) = 1$ for $x \geq x^0$. Again, x^0 is the unique fixed point on $[0,x^0]$ such that $w_t^*(x^0) = x^0$. Therefore, for $t < T$, the net search welfare is

$$J_t(x) = \max\{x, w_t^*(x)\} = \begin{cases} x & \text{if } x \in [x^0, \bar{x}], \\ w_t^*(x) = W_t(m_t^{\mathrm{e}}(x);x) - cm_t^{\mathrm{e}}(x) & \text{if } x \in [0, x^0), \end{cases}$$

which is obtained by the efficient search rule described in (32).

Next, $\partial W_t(m;x)/\partial x = F(x)^m \mathcal{F}^e_{t+1}(x)$, which is increasing in x. As in **Step 1**, $w_t(m;x)$ is convex in x for any given m, and hence $w^*_t(x)$ is also convex in x. The search value $J_t(x)$, as the maximum of two convex functions, is convex and absolutely continuous. Hence, its derivative $W'_t(x)$ is defined almost everywhere and $J_t(x)$ can be represented by

$$J_t(x) = J_t(\bar{x}) - \int_x^{\bar{x}} W'_t(z)dz = \bar{x} - \int_x^{\bar{x}} W'_t(z)dz,$$

where we apply $J_t(\bar{x}) = \bar{x}$. As $\partial w_t(m;x)/\partial x = F(x)^m \mathcal{F}^e_{t+1}(x)$, envelop theorem yields

$$W'_t(x) = \begin{cases} 1 & \text{if } x \in [x^0, \bar{x}], \\ F(x)^{m_t(x)} \mathcal{F}^e_{t+1}(x) & \text{if } x \in [0, x^0), \end{cases}$$

which is equivalent to $\mathcal{F}^e_t(x)$ given in (35). □

Proof of Proposition 4. The proof is analogous to that of Theorem 2 and Proposition 3, and is hence omitted here. □

Proof of Corollary 4. v^0 and x^0 are given by $\int_{v^0}^{\bar{v}}[1-G(z)]dz = \int_{x^0}^{\bar{x}}[1-F(z)]dz = c$, where $\bar{v} = \bar{x}$. Suppose $v^0 \leq \psi(x^0)$. We then have

$$\begin{aligned} \int_{v^0}^{\bar{v}}[1-G(z)]dz &\geq \int_{\psi(x^0)}^{\bar{x}}[1-G(z)]dz = \int_{\psi(x^0)}^{\bar{x}}[1-F(\psi^{-1}(z))]dz \\ &= \int_{x^0}^{\bar{x}}[1-F(x)]\psi'(x)dx > \int_{x^0}^{\bar{x}}[1-F(x)]dx, \end{aligned}$$

where we replace $z = \psi(x)$ and the last inequality is from the fact that $\psi'(x) > 1$, which results in a contradiction. It then must be that $v^0 > \psi(x^0)$. □

Proof of Corollary 5. We know that the optimal and efficient sample sizes are given by

$$\hat{m}^*_t = \min\{m \in \mathbb{N}_0 : MR_T(m, \hat{v}_{t+1}) \leq c\} \text{ and } \hat{m}^e_t = \min\{m \in \mathbb{N}_0 : MW_T(m, \hat{x}_{t+1}) \leq c\},$$

respectively. To prove that $\hat{m}^*_t \geq \hat{m}^e_t$ when $\hat{v}_{t+1} = \psi(\hat{x}_{t+1})$, it suffices to show that $MR_T(m, \hat{v}_{t+1}) \geq MW_T(m, \hat{x}_{t+1})$. We have

$$\begin{aligned} MR_T(m, \hat{v}_{t+1}) &= \int_{\hat{v}_{t+1}}^{\bar{v}} G(z)^m(1-G(z))dz = \int_{\psi(\hat{x}_{t+1})}^{\bar{x}} G(z)^m(1-G(z))dz \\ &= \int_{\hat{x}_{t+1}}^{\bar{x}} F(x)^m(1-F(x))\psi'(x)dx \geq MW_T(m;\hat{x}_{t+1}), \end{aligned}$$

where we substitute $z = \psi(x)$ and the last inequality is from the fact that $\psi'(x) > 1$ under the assumption of increasing failure rate. □