Efficiency of Sequential English Auctions with Dynamic Arrivals

Olivier Compte* Ron Lavi† Ella Segev‡

Abstract

We study a setting of online auctions with expiring/perishable items: K items are sold sequentially, buyers arrive over time and have unit-demand and private values. The goal is to maximize the social welfare – the sum of winners' values. This model was previously studied in several different papers, suggesting several different solutions for the problem. All previous solutions are direct-revelation mechanisms, while in real-life we usually see open mechanisms, most often a sequence of English auctions. We ask whether the previous optimal approximation bounds can be achieved using the more popular/realistic mechanism, or a small variant of it. We observe that a sequence of English auctions (the exact original format) does not guarantee any constant approximation of the welfare, and describe two variants that bring back the approximation guarantee.

In the first variant, the price ascent is stopped when the number of active bidders is equal to the number of remaining items. The winner is chosen from this set of active bidders using some tie-breaking rule. This yields a truthful deterministic 2-approximation. Moreover we show that this ratio is the best possible for any deterministic mechanism that must charge payments at the time of the sale. If the winner is chosen uniformly at random from the set of active bidders, the approximation ratio decreases to $\frac{e}{e-1} \simeq 1.582$ (the currently best approximation ratio for this problem).

The second variant is to disqualify bidders that quit the auction when the number of active bidders is larger than the number of remaining items, i.e. to disallow their participation in future auctions. Under the assumption that the true arrival times are observed by the auctioneer (i.e. values are the only private information) this activity rule again ensures a 2-approximation of the social welfare when players play undominated strategies.

*Paris School of Economics. Email: compte@enpc.fr.

†Faculty of Industrial Engineering and Management, The Technion – Israel Institute of Technology. Email: ronlavi@ie.technion.ac.il. Supported by grants from the Israeli Science Foundation, the Bi-national Science Foundation, the Israeli ministry of Science, and Google.

‡Department of Industrial Engineering and Management, Ben-Gurion University of the Negev. Email: ellasgv@bgu.ac.il.

1 Introduction

1.1 Background and Motivation

Auctions over the Internet are more dynamic than classic auctions, since items and bidders often arrive and depart over time. Several theoretical models of online auctions were studied in the last decade, trying to adjust classic theory to the new electronic settings. One such simple model of online auctions with expiring/perishable items was studied in several papers: There are K time units, and in each time unit there is an item that must be sold at that time unit. Buyers arrive and depart over time and have a private value for exactly one of the items that are being sold during their stay. The goal is to maximize the social welfare – sum of winners' values. The common motivating scenario is the allocation of computational resources like CPU time or network bandwidth, but the same model captures more classic economic settings, for example selling movie tickets or other types of items that are time-dependent.

For this model, Lavi and Nisan (2005) show that any deterministic truthful mechanism must have an approximation ratio of at least K to the social welfare (and suggest a weaker game-theoretic notion to solve the problem). Hajiaghayi, Kleinberg, Mahdian and Parkes (2005) give a truthful 2-approximation in the same model, assuming that departure times cannot be manipulated, and that payments need not be collected at the time of sale but only after all items were sold. Cole, Dobzinski and Fleischer (2008) study "prompt" mechanisms, requiring payments to be made at the time of sale, and give a truthful 2-approximation, assuming that both arrival and departure times cannot be manipulated.

All these are direct-revelation mechanisms, while in real-life we usually see open (indirect, iterative) mechanisms, in particular some variant of an English auction is usually being employed. On top of the psychological preference for open mechanisms, there are several other advantages to such mechanisms, for example they reduce information revelation of (the private) types, and they are more robust to changes in the utility model.[1]

In this paper we are therefore interested in the following question: can we obtain the same approximation bounds by using a sequence of English auctions, which are the most widely acceptable format in reality? We study a special case, in which once a player arrives she remains for all subsequent auctions (i.e. no departures), and start by showing a lower bound of 2 on the approximation ratio of any truthful deterministic mechanism with prompt payments, even for this more restricted setup. Thus the previously suggested mechanisms remain optimal for our special case.

We then observe that the exact original format of the English auction does not guarantee any constant approximation of the welfare. As an example to the difficulties that emerge, consider the following scenario. Two items are sold via two consecutive ascending English auctions, one after the other. There are two players that participate in the first auction; each player desires one of the two items, and is indifferent between the two items. There is a certain probability that a high-value bidder will arrive for the second auction, and in this case the loser in the first auction will also lose the second auction. Clearly, if a player assigns a high probability to this event, she will be willing to compete (almost) up to her value in the first auction while if she assigns a low probability to this event, she will stop competing in the first auction at a low price. If the two players have significantly different beliefs regarding this event, one will retire early and the other

[1] E.g, Dobzinski, Lavi and Nisan (2008) show how Ausubel's clinching auction naturally extends to a utility model with a hard budget constraint. In fact, all our results here immediately extend to a model where each player has a private budget constraint. It is not clear if previous results hold under such an extension.

one will win. This has a negative effect on the social welfare when the player with the higher value incorrectly underestimates the probability of the new second-period arrival. One can easily construct such examples for larger numbers of items, showing that without any modification to the original English auction format there is no hope to extract a constant fraction of the optimal social welfare.[2]

1.2 Main Results

We show how small but meaningful modifications to the English auction bring back the approximation guarantee. We give two different modifications for two different informational assumptions: The first is that arrival times can be manipulated, i.e. a player can arrive *after* her true arrival time, or arrive at the true arrival time and conceal her presence, if this may increase her utility. The second is that arrival times can be observed by the designer and thus cannot be manipulated. In both cases the resulting mechanism yields a deterministic 2-approximation, and we additionally show how distributional/randomized considerations can further improve the bound and beat the deterministic lower bound.

1.2.1 Unobserved Arrivals

For unobserved arrival times, we suggest the following variant of the English auction: when there remain t items for sale (i.e. the current auction is the $K - t$ auction), stop the price ascent when there remain exactly t active bidders. The winner is one of these active bidders, chosen using any arbitrary deterministic or randomized tie-breaking rule. Her payment is the price that was reached. We show that if this seemingly small modification is employed, it is an ex-post equilibrium to play truthfully, i.e. to arrive in the true arrival time and quit each auction exactly when the price reaches the player's value. We then show that, regardless of the tie-breaking rule, this guarantees that at least half of the optimal social welfare will always be obtained.

The deterministic lower bound of 2 can be beaten by using a random tie-breaking rule in this auction. In particular, we show that if we break the tie uniformly at random (each of the active bidders is declared winner with equal probability), the approximation ratio becomes $\frac{e}{e-1} \simeq 1.58$. This is the best known bound even for *non-truthful* randomized algorithms (Bartal, Chin, Chrobak, Fung, Jawor, Lavi, Sgall and Tichý, 2004), and we are able to achieve it even with the presence of strategic players. Furthermore no randomized truthful mechanism for this problem that beats the 2-approximation guarantee was previously known.

The analysis of the welfare loss first shows that the worst-case is when exactly one new player with value 1 arrives in every auction and in addition there are many players with value 0 that are present from the first auction. The expected welfare in this case is simply the expected number of 1-players that win. In the first few auctions, the winners are 0-players with very high probability, and as the number of remaining 1-players increases the probability that a 1-player will win increases. On the other hand, if at some auction t^* the number of remaining 1-players is at least the number of remaining items we are guaranteed that only 1-players win from this auction onwards. Furthermore, one can observe that the overall number of 1-players that win roughly equals t^*. We show that t^*

[2] Said (2008) shows that if one assumes a common-prior then this problem disappears and efficiency is regained, while our paper (as well as all the previous literature mentioned above) studies the "robust" setting, in which no common-prior is assumed.

is the stopping time of a certain super-martingale process, and use this to show that its expected value is about $\frac{e-1}{e}K$, implying the claimed bound.

From a conceptual point of view, the worst-case scenario reveals the fact that (sometimes) the auction awards almost for free the first several items, to artificially create supply shortage and high competition for remaining items. From the strategic aspect, this guarantees that players will not misreport true value and will find it in their best interest to arrive as early as possible. From the efficiency aspect, this guarantees the optimal approximation ratio that can be achieved by any deterministic truthful mechanism, and enables the use of randomization to further improve the bound. It remains an interesting open question whether the specific bound we obtain is the best possible via a truthful randomized mechanism, or perhaps there is a way to obtain truthfulness with less "supply reduction" which will imply larger overall welfare.

1.2.2 Observed Arrivals

Sellers may find it unattractive to stop the price ascent that early, when competition is still strong. We give a second possible modification to the English auction under the assumption that arrival times cannot be manipulated (i.e. the only private parameter of the players is their value). Arrivals are still dynamic and are not known a-priori, but the player's true arrival time is observed by the auctioneer. In this case we suggest the following modification: when there remain t items for sale (i.e. the current auction is the $K-t$ auction), all players that quit the current auction when there are more than t active bidders are disqualified from participating in subsequent auctions. In addition, the price ascent of the next auction $K-t+1$ starts from the price point at auction $K-t$ at which there remain exactly t active bidders. Since players cannot drop too early from the auction as they will be disqualified, this added "activity rule" has the effect of increased competition, and therefore may be attractive to sellers.

With this modification we show that the auction again always obtains at least half of the optimal social welfare, *whenever players play any tuple of undominated strategies*. This last point deserves some more attention. As demonstrated above in the context of the original English auction, players may wish to drop early in the auction, if they believe that future auctions will exhibit less competition than current auction, and alternatively may compete in current auction until price becomes very close to their value if they believe that future auctions will exhibit higher competition. The additional activity rule does not eliminate this state of affairs, it only limits it: regardless of the player's beliefs, it is a dominated strategy to drop in the disqualifying range if the price is lower than the value, and thus we assume that players will not do it. But when only t or less players remain in auction $K-t$, players may still decide differently where to drop, depending on their beliefs, and so there is no unique equilibrium choice. The analysis therefore does not assume a certain drop point for all players, instead it hold for all possible drop points in the valid range. Our activity rule carefully chooses the valid range, so that the efficiency loss will be minimized.

While the worst-case bound here is again 2, it makes sense that "usually" the obtained welfare will be much higher than half of the optimum. However it is not clear how to formally justify such a statement. The usual average-case analysis requires making specific assumptions on the exact nature of the underlying distribution, assumptions that on the other hand make the analysis very specific, so the generality of the conclusions is questionable. We choose here the following middle-ground method that we feel does not restrict generality too much. While a worst-case analysis assumes an adversary that is allowed to determine the number of players, their arrival times, their values, and their (undominated) strategies, we conduct (on top of the worst-case analysis that shows the 2-approximation) an analysis that assumes a slightly weaker adversary: it has powers almost as before, except that instead of assigning arbitrary worst case values to the players, the adversary now has to choose (any) probability distribution, and draw the values independently from that distribution. The adversary can still set all other parameters freely, as before.

Under these very weak average-case assumptions – in fact the only additional assumption is that players' values are drawn i.i.d. from *some* arbitrary distribution – we analyze the ratio between the expected welfare of the auction and the expected optimal welfare, for the special case of two items. We first show that the worst possible probability distribution is a Bernoulli distribution (i.e. with some probability p the value is 0 and with probability $1-p$ the value is 1). Second, we find the exact worst p, and as a result obtain that the ratio between the expected welfare of the sequential English auction with activity rule and the expected welfare of the optimal allocation is at least $\sqrt{2}/2 \simeq 0.707$. Of-course, this bound is achieved only for the *worst possible distribution*. For example, if we take the uniform distribution over some interval then the ratio will increase to 80%. This exercise strengthens our basic intuition that a worst case bound of 2 implies much better bounds even with very mild distributional assumptions. The analysis itself seems interesting, and may potentially be applied to other relevant models.

1.3 Paper Organization

Section 2 describes the model and the lower bound of 2 for deterministic truthful mechanisms. Section 3 describes our results for the case of unobserved arrival times, and section 4 describes our results for the case of observed arrival times.

2 The Setting

A seller sells K identical items using a sequence of K single-item ascending auctions. There are n unit-demand bidders with private values and quasi-linear utilities: a bidder has value v_i for receiving an item; her utility is $v_i - p_i$ if she wins an item and pays p_i, and 0 if she loses. We study a dynamic setting where bidders arrive over time. Formally, bidder i's type includes, besides her value v_i, an arrival time r_i which is an integer between 1 and K, indicating that bidder i may participate only in the auctions for items $r_i, ..., K$. Thus, a bidder's type is a pair $\theta_i = (r_i, v_i)$, and the set of possible types for bidder i is Θ_i. We denote $\Theta = \Theta_1 \times \cdots \times \Theta_n$. We study two variants of the model: in the unobserved arrival time version, the value and the arrival time of a player is her private information, and therefore a player may choose to arrive *strictly after* her true arrival time if she finds it strategically useful. In the observed arrival time version the auctioneer knows the true arrival time of the player when the player truly arrives. We wish to maximize the social welfare – sum of winners' values – and evaluate a given mechanism according to its approximation ratio: the worst-case ratio over all tuples of types between the optimal (i.e. maximal possible) social welfare and the social welfare that the mechanism obtains. As we next show, any deterministic mechanism has an approximation ratio of at least 2 in any ex-post equilibrium.

2.1 Deterministic Lower Bound

By the direct-revelation principle, we focus on direct mechanisms in which truthful reporting of the type is a dominant-strategy. We also assume ex-post Individual Rationality: a winner pays at most her declared value and a loser pays at most zero. We show that there is no deterministic truthful mechanism with approximation ratio strictly smaller than 2, even for the very restrictive setting of two items and three players, where it is common knowledge that players 1 and 2 arrive for the first auction and player 3 arrives for the second auction. This restriction only strengthens the impossibility.

Definition 1 (A limited direct mechanism)**.** *A direct mechanism for two items and three players is a set of four functions:* $w_1(v_1, v_2)$ *determines the winner (either* 1 *or* 2*) of the first item, and she pays a price* $p_1(v_1, v_2)$*, where* $p_1(v_1, v_2) \le v_{w_1(v_1,v_2)}$*.* $w_2(v_1, v_2, v_3)$ *determines the winner of the second item (either* 1*,* 2*, or* 3*, but not* $w_1(v_1, v_2)$*), and she pays a price* $p_2(v_1, v_2, v_3)$*, where* $p_2(v_1, v_2, v_3) \le v_{w_2(v_1,v_2,v_3)}$*. Such a mechanism is called "truthful" if it is a dominant-strategy of every player to report her true type.*

Theorem 1. *Every truthful and individually rational limited direct mechanism obtains in the worst-case at most half of the optimal social welfare.*

Proof. Fix any $\frac{1}{2} \ge \epsilon > 0$. Suppose by contradiction that their exists a truthful mechanism $M = (w_1(\cdot,\cdot), p_1(\cdot,\cdot),\ w_2(\cdot,\cdot,\cdot), p_2(\cdot,\cdot,\cdot))$ that always obtains at least $\frac{1}{2} + \epsilon$ of the optimal social welfare. We first show three short claims.

Claim 1. *If* $v_2 > \frac{v_1}{2\epsilon}$ *then* $w_1(v_1, v_2) = 2$*, i.e. player* 2 *must be the winner of the first auction.*

Proof. Suppose by contradiction that there exists an instance (v_1, v_2, v_3) such that $v_2 > \frac{v_1}{2\epsilon}$ and $w_1(v_1, v_2) = 1$. Consider another instance $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3)$, where $\tilde{v}_1 = v_1, \tilde{v}_2 = v_2$, and $\tilde{v}_3 = v_2$. The optimal social welfare in this instance is $2 \cdot v_2$. We have $w_1(\tilde{v}_1, \tilde{v}_2) = w_1(v_1, v_2) = 1$, and therefore the social welfare that the mechanism obtains is $v_1 + v_2$. But $\frac{v_1+v_2}{2v_2} < \frac{1}{2} + \epsilon$ which contradicts the fact that the mechanism always obtains at least $\frac{1}{2} + \epsilon$ of the optimal social welfare. □

Claim 2. *If* $v_1 = 1, v_2 > \frac{1-2\epsilon}{1+2\epsilon}$*, and* $w_1(v_1, v_2) = 2$*, then* $p_1(v_1, v_2) \le \frac{1-2\epsilon}{1+2\epsilon}$ *(note that* $\frac{1-2\epsilon}{1+2\epsilon} < 1$*).*

Proof. First note that in the instance $(v_1 = 1, v_2 > \frac{1-2\epsilon}{1+2\epsilon}, v_3 = 0)$ players 1 and 2 must win, since any other set of winners has welfare strictly less than a fraction of $\frac{1}{2}+\epsilon$ of the optimal social welfare of this instance.

Now, suppose a contradicting instance (v_1, v_2, v_3) where $p_1(v_1, v_2) > \frac{1-2\epsilon}{1+2\epsilon} + \delta$ for some $\delta > 0$. Note that player 2 wins item 1 and pays the same price in the instance $(v_1, v_2, 0)$ (call this "instance 2"). Consider a third instance $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3)$, where $\tilde{v}_1 = 1, \tilde{v}_2 = \frac{1-2\epsilon}{1+2\epsilon} + \delta$, and $\tilde{v}_3 = 0$. By the first paragraph, player 2 must be a winner in instance 3, and by individual rationality she pays at most $\frac{1-2\epsilon}{1+2\epsilon} + \delta$. Therefore, in instance 2, player 2 has a false announcement ($\frac{1-2\epsilon}{1+2\epsilon} + \delta$ instead of v_2) that strictly increases her utility, a contradiction to truthfulness. □

To reach a contradiction and conclude the proof, consider the instance $(1, 1, 5)$. Suppose w.l.o.g. that $w_1(1, 1) = 1$. To obtain at least half of the optimal welfare we must have $w_2(1, 1, 5) = 3$. Thus player 2 loses and has zero utility. However if she declares some $\tilde{v}_2 > \frac{1}{2\epsilon}$ instead of her true type $v_2 = 1$ then by claim 1 she will win the first item and by claim 2 she will pay a price of at most $\frac{1-2\epsilon}{1+2\epsilon} < 1$. Thus she is able to strictly increase her utility by some false declaration, contradicting truthfulness.

□

3 Unobserved Arrivals

3.1 A Modification to the English Auction

We suggest the following modification to each English auction. In the $(K-t)$'th auction ($t = 0, ..., K-1$), the price ascent stops when there remain exactly t bidders that have not dropped yet (note that t is equal to the number of unsold items). The winner is chosen by some arbitrary tie-breaking rule, and she pays the price that was reached. To summarize, the auction is as follows:

- All players that are present may participate, and a single price ascends continuously from zero.
- Each player continuously decides whether to quit the current auction.
- When the number of remaining players exactly equals the number of remaining items the price ascent stops. Let this price be p_t, and denote the set of players that are active in this price by B_t.
- The winner in the auction is chosen from B_t by some tie-breaking rule. Her payment is p_t.

Two additional technicalities need to be discussed. First, to avoid pathological situations the auction artificially adds K "dummy" players with value zero that arrive for the first auction. This way, the number of players in each auction is not less than the number of remaining items. The optimal allocation is obviously not affected by the dummy players, and our efficiency bounds hold even if some of the dummy players are chosen to win some of the items (meaning that that item is left unallocated). Second, if several players quit at the same price in some auction, and this makes the number of remaining players strictly below t, then the auctioneer includes some of the players that have dropped at p_t in B_t, to make its size exactly equal to t. The choice which of the players who dropped at p_t to include is made by some arbitrary "B completion rule".

We first argue that the truthful strategy of arriving at the true arrival time and in each auction quitting exactly when the price is equal to one's value, forms a symmetric ex-post equilibrium. We prove this by induction on K. For $K = 1$ the auction is the regular English auction which implies the claim. Assume correctness for $K-1$ and let us prove for K. It is enough to prove only that a player i that arrives at time 1 maximizes her utility by staying in the first auction until the price reaches her value, since (i) if the player first deviates from the truthful strategy in some later auction and as a result increases her utility we can construct an instance with fewer items for which deviating in the first auction will increase utility, a contradiction to the inductive assumption (and thus we can assume that in all subsequent auctions the player stays until her value in the auction), and (ii) delaying arrival to a later auction is strategically equivalent to quitting at price zero in the first auction.

We assume that all players besides i play truthfully and wish to prove that i maximizes her utility by remaining in the first auction until the price equals her value. Let p_1 be the price of the first item when player i is truthful. If $v_i < p_1$ then player 1 cannot obtain a positive utility regardless of her actions. Thus assume $v_i \geq p_1$ and $i \in B_1$. If player i wins the first auction when truthful then there is no beneficial deviation since prices of subsequent items must be at least p_1 as all other players in B_1 will be present in all subsequent auctions. If player i does not win the auction and the same player j wins the auction when i is truthful and when i deviates then i obtains the same utility since in subsequent auctions she remains until her value in both cases. If player i

does not win the auction and the deviation causes a different player $j' \notin B_1$ to win the first auction then, since all players in B_1 continue to subsequent auctions, subsequent prices can only increase, which implies that i's utility can only decrease. We get:

Proposition 1. *The truthful strategy is a symmetric ex-post equilibrium.*

Our main goal is to analyze the efficiency of this equilibrium outcome. Let $W_A(\theta)$ denote the resulting welfare of our auction (with any fixed tie-breaking rule), when players' types are θ, and they play the equilibrium strategy. Similarly let $W_{A(U)}(\theta)$ denote the welfare of our auction when the winner is chosen uniformly at random among all players in B_t. Note that $W_{A(U)}(\theta)$ is a random variable. Let $W_{OPT}(\theta)$ denote the optimal welfare for the scenario θ. This variable is constructed deterministically by taking in each auction a player with the maximal value among all participating players.

Theorem 2. *For any tuple of players' types θ, the expected social welfare of the sequential auction, when players play the equilibrium strategies, is at least half of the optimal welfare for θ. Moreover, if the winner is chosen uniformly at random among all players in B_t, then,*

$$\inf_{\theta\in\Theta} \frac{E[W_{A(U)}(\theta)]}{W_{OPT}(\theta)} \geq 1 - \frac{1}{e}.$$

where $e = 2.718...$ is "Euler's number".

The rest of this section proves this theorem, in two parts. We first identify a concrete set of tuples of types ("scenarios"), $\{\theta^*_K\}_{K=1,2,...}$, that are worst-case (in an exact sense to be defined below). We then analyze the auction-to-optimal ratio over this set of types.

3.2 A family of worst-case scenarios

Definition 2. *For a given number of items K, θ^*_K is a tuple of at least $2K$ players. Players $1, ..., K$ (the "1-players") all have a value of 1, and the arrival time of player $i = 1, ..., K$ is $r_i = i$. All other players (the "0-players") have value zero. At least K 0-players arrive for the first auction.*

Although θ^*_K is not uniquely defined, all tuples of types that satisfy the definition of θ^*_K are equivalent for our purposes since the number of zero players does not change the optimal welfare, nor the outcome of the auction (besides perhaps the identity of the zero-players that win). Clearly, $W_{OPT}(\theta^*_K) = K$ for any K. We first show:

Lemma 1. *For any scenario $\theta \in \Theta$ over K items,*

$$\frac{E[W_A(\theta)]}{W_{OPT}(\theta)} \geq \min_{K'\in\{1,...,K\}} \frac{E[W_A(\theta^*_{K'})]}{K'}$$

We prove this in two steps: We first give a procedure to change the players' values so that the new resulting values are either 0 or 1, and the ratio $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)}$ only decreases. We then add more zero-players, and remove 1-players or change their arrival times, until we reach a scenario $\theta^*_{K'}$ for some $K' \leq K$. We again show that these operations can only decrease the ratio $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)}$.

First Step. For any given scenario we construct a different scenario for which the values are either 0 or 1 and the ratio $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)}$ is weakly lower. We rely on three claims whose proofs are given in the appendix.

Claim 3. *Fix any two scenarios θ, θ' with the same set of players, such that $r_i = r'_i$ for every player i, and for any two players i, j, $v_i \geq v_j$ if and only if $v'_i \geq v'_j$. Then there exists a specific B completion rule such that, when using this rule, the probability that a given player wins is the same in both scenarios.*

Claim 4. *Changing the B completion rule does not change the expected welfare of the auction.*

Claim 5. *Fix arbitrary non-negative real values $x_1, ..., x_L$ and arbitrary non-negative real weights $\alpha_1, ..., \alpha_L$ and $\beta_1, ..., \beta_L$. Let $l = argmax_{j=1,...,L} \frac{\alpha_j}{\beta_j}$, and fix any real number $x'_l < x_l$. Then $\frac{\sum_{j=1}^{L} \alpha_j x_j}{\sum_{j=1}^{L} \beta_j x_j} \geq \frac{\sum_{j \neq l} \alpha_j x_j + \alpha_l x'_l}{\sum_{j \neq l} \beta_j x_j + \beta_l x'_l}$.*

Proof. Follows since the partial derivative of the left-hand-side with respect to x_l is non-negative, regardless of the values $x_1, ..., x_L$. □

With these three claims we prove the first step by performing the following procedure iteratively, starting with the original scenario:

1. Denote the current scenario by θ. Let p_i denote the probability that player i wins one of the auctions, and let y_i be an indicator that denotes whether i wins an item in the optimal allocation (both are for the current scenario θ).

2. Let $x_1, ..., x_L$ denote the set of distinct positive values of players in θ in decreasing order. For any $1 \leq j \leq L$, define $\alpha_j = \sum_{i \,:\, v_i = x_j} p_i$ and $\beta_j = \sum_{i \,:\, v_i = x_j} y_i$.

3. If $L = 1$ then stop. Otherwise choose an index l with maximal $\frac{\alpha_l}{\beta_l}$. Form a new scenario θ' by decreasing the values of all players i with $v_i = x_l$ to have value $v'_i = x_{l+1}$ (where $x_{L+1} \equiv 0$). All other players' types remain unchanged.

We argue that in every iteration of this process the auction-to-optimal welfare ratio cannot increase, that is $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)} \geq \frac{E[W_A(\theta')]}{W_{OPT}(\theta')}$. By claim 3 we have that $p_i = p'_i$ where p'_i denotes the probability that player i wins in θ' if the auction uses the specific B completion rule of claim 3 (denote this auction as A'). Thus $E[W_{A'}(\theta')] = \sum_{j \neq l} \alpha_j x_j + \alpha_l x'_l$. We also have $W_{OPT}(\theta') = \sum_{j \neq l} \beta_j x_j + \beta_l x'_l$ since the relative order among the players' values is the same in θ and θ'. Claim 5 implies $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)} \geq \frac{E[W_{A'}(\theta')]}{W_{OPT}(\theta')}$, and claim 4 implies $\frac{E[W_{A'}(\theta')]}{W_{OPT}(\theta')} = \frac{E[W_A(\theta')]}{W_{OPT}(\theta')}$. By repeating this procedure we end up with a scenario in which all positive values are identical. Changing all positive values to 1 will keep the same ratio. This concludes the first step.

Second Step. Let θ be some scenario such that $v_i \in \{0, 1\}$ for any player i. We now wish to prove that

$$\frac{E[W_A(\theta)]}{W_{OPT}(\theta)} \geq \min_{K' \leq K} \frac{E[W_A(\theta^*_{K'})]}{K'}$$

which will conclude the proof of Lemma 1. We rely on two claims whose proofs are given in the appendices.

Claim 6. *Fix a scenario θ such that $v_j \in \{0, 1\}$ for any player j. Additionally let i be some player with $r_i = 1$ and $v_i = 1$. Then $E[W_A(\theta)] \leq 1 + E[W_A(\theta_{-i})]$.*

Claim 7. *Fix two scenarios θ, θ' such that $v_j \in \{0, 1\}$ for any player j in θ, and θ' is a modification of θ such that one of the following holds:*

- *one player arrives one auction later, i.e. there exists a player i such that $\theta'_j = \theta_j$ for any $j \neq i$ and $\theta'_i = (v_i, r_i + 1)$ (if $r_i = K$ then $\theta' = \theta_{-i}$), or*
- *the set of types in θ' includes θ plus some additional zero players that arrive for the first auction.*

Then $E[W_A(\theta)] \geq E[W_A(\theta')]$.

We now prove step 2. We start from a scenario θ such that $v_i \in \{0, 1\}$ for any player i in θ. By claim 7 we can assume that θ contains at least K zero players that arrive for the first auction (otherwise we add them and the auction-to-optimal welfare ratio does not increase). Let a_t for $t = 1, ..., K$ denote the number of players with positive value in θ that arrive in auction t. If there exists an auction t with $a_t \geq 2$ then by claim 7 we move to θ' in which some player with $r_i = t$ now has $r'_i = t + 1$ (and if $t = K$ then $\theta' = \theta_{-i}$) and have $E[W_A(\theta)] \geq E[W_A(\theta')]$. We also have $OPT(\theta) = OPT(\theta')$ since the optimal allocation chooses in each auction a participating player with maximal value and one can verify that in θ and θ' the maximal value in each auction is the same. Thus $\frac{E[W_A(\theta)]}{W_{OPT}(\theta)} \geq \frac{E[W_A(\theta')]}{W_{OPT}(\theta')}$. By repeatedly performing this step we reach a scenario where $a_t \leq 1$ for any auction t. If $a_t = 1$ for any auction t then this is θ^*_K, and we are done. Otherwise if $a_1 = 0$ then we actually have a scenario over $K - 1$ items and by an inductive argument we conclude the claim. Otherwise $a_1 = 1$ and there exists some $t > 1$ with $a_t = 0$. By claim 7 we move to θ' in which the 1-player with $r_i = 1$ now has $r'_i = t$, and $E[W_A(\theta)] \geq E[W_A(\theta')]$. We also have $OPT(\theta) = OPT(\theta')$ since the optimal allocation chooses all players with positive values in both scenarios. Since θ' is a scenario over $K - 1$ items the claim follows by induction.

3.3 Bounding the inefficiency of θ^*_K

We now wish to prove the approximation ratio. Fix any $K \in \mathbb{N}$. Recall that in θ^*_K there are K zero-value players that arrive for the first auction, and in every auction $t = 1, ..., K$ arrives one additional player with value 1. To analyze $E[W_A(\theta^*_K)]$, define a variable Z_t for $t = 1, ..., K$ whose value is the number of zero players that were chosen as winners in auctions $1, ..., t$ in θ^*_K. Z_t may be a random variable. The number of 1-players that participate in auction $t + 1$ is exactly $Z_t + 1$: there are Z_t 1-players that arrived in auctions $1, ..., t$ and were not winners in those auctions since zero-player that is a winner there is a corresponding 1-player which is a loser. An additional 1-player arrives for auction $t + 1$. Let $P_t = \frac{Z_t+1}{K-t}$. If $P_t \leq 1$ then it exactly equals the fraction of 1-player in B_{t+1}, since there are $Z_t + 1$ 1-players at time $t + 1$ and $|B_{t+1}| = K - t$ (t items were already sold so the number of remaining items is $K - t$).

Define an additional variable t^* to be the first auction t for which $P_t \geq 1$, $t^* = \text{argmin}_{1 \leq t \leq K}(P_t \geq 1)$. In other words, in auctions $t^* + 1, ..., K$ the winners must all be 1-players. We also know that the number of 1-players that are winners in auctions $1, ...t^*$ is $t^* - Z_{t^*}$. Thus the total number of winners with value of 1 is $W_A(\theta^*_K) = (K - t^*) + (t^* - Z_{t^*}) = K - Z_{t^*}$. Note that it must be that $t^* \geq K/2$ since before this time not enough 1-players have arrived. Since $1 > P_{t^*-1} = \frac{Z_{t^*-1}+1}{K-(t^*-1)}$, we get $Z_{t^*-1} < K - t^*$ and therefore $Z_{t^*} \leq K - t^*$. Thus $W_A(\theta^*_K) = K - Z_{t^*} \geq t^* \geq K/2$. This implies the first part of theorem 2.

To prove the second part we show that, for any $K \in \mathbb{N}$, $\frac{E[W_{A(U)}(\theta^*_K)]}{K} \geq 1 - \frac{1}{e}$. Since $W_{A(U)}(\theta^*_K) = K - Z_{t^*}$ it is enough to show that $E[Z_{t^*}] \leq \frac{K}{e}$.

With the uniformly-at-random tie-breaking rule, when $P_t \leq 1$ it exactly equals the probability that a 1-player will be the winner in auction $t+1$, since there are $Z_t + 1$ 1-players and the winner is chosen uniformly at random from a set of players of size $K - t$ that contains all the participating 1-players. If $P_t \geq 1$ then as before the winner of auction $t+1$ will be a 1-player with probability 1. Thus

$$Z_{t+1} = \begin{cases} Z_t & \text{with prob. } \min(P_t, 1) \\ Z_t + 1 & \text{with prob. } 1 - \min(P_t, 1) \end{cases}$$

and $Z_0 = 0$. To bound $E[Z_{t^*}]$ we super-martingales. A stochastic process $X_1, ..., X_K$ is a "super-martingale" if, for any $1 \leq t < K$, $E[X_{t+1} \mid X_t = l] \leq l$ for any possible value l of X_t. A "stopping time" for $X_1, ..., X_K$ is a random variable T with the property that for each $1 \leq t \leq K$, the occurrence or non-occurrence of the event $T = t$ depends only on the values of $X_1, ..., X_t$. Therefore t^* is a stopping time for $Z_1, ..., Z_K$. If $X_1, ..., X_K$ is a super-martingale, and t^* is a stopping time, then $E[X_{t^*}] \leq E[X_1]$ (Williams, 1991). Back to our setting, we construct an auxiliary process: for any $t \leq t^*$,

$$X_t = Z_t - \sum_{j=0}^{t-1} \frac{K-t}{K-j},$$

and for $t > t^*$ set $X_t = X_{t-1}$. For $t \geq t^*$ we have $E[X_{t+1} \mid X_t = l] = l$ since $X_{t+1} = X_t$. To show that X_t is a super-martingale we compute, for any $t < t^*$, $E[Z_{t+1} \mid Z_t = l < K - t - 1] = P_t \cdot l + (1 - P_t)(l+1) = (l+1)\frac{K-t-1}{K-t}$.

$$\begin{aligned}
&\text{Now, } E[X_{t+1} \mid X_t = l] \\
&= E\left[Z_{t+1} \mid Z_t = l + \sum_{j=0}^{t-1} \frac{K-t}{K-j}\right] - \sum_{j=0}^{t} \frac{K-t-1}{K-j} \\
&= \frac{K-t-1}{K-t}\left(l + \sum_{j=0}^{t-1} \frac{K-t}{K-j} + 1\right) - \sum_{j=0}^{t} \frac{K-t-1}{K-j} \\
&= \frac{K-t-1}{K-t} \cdot l \leq l,
\end{aligned}$$

implying that $X_1, ..., X_K$ is a super-martingale. We have that $E[Z_1] = 1 \cdot \frac{K-1}{K} + 0 \cdot \frac{1}{K}$ and therefore $E[X_1] = E[Z_1] - \frac{K-1}{K} = 0$. Hence

$$0 = E[X_1] \geq E[X_{t^*}] = E\left[Z_{t^*} - \sum_{j=0}^{t^*-1} \frac{K-t^*}{K-j}\right],$$

or, in other words, $E[Z_{t^*}] \leq E\left[(K - t^*)\sum_{j=0}^{t^*-1} \frac{1}{k-j}\right]$. Since

$$E\left[(K-t^*)\sum_{j=0}^{t^*-1} \frac{1}{K-j}\right] \leq \max_{t=1,...K} (K-t)\sum_{j=0}^{t-1} \frac{1}{K-j}$$

we upper bound the RHS by $\frac{K}{e}$. We use the fact that, for any integer n, if we set $\epsilon_n = \sum_{j=1}^{n} \frac{1}{j} - \ln(n)$, then $0 < \epsilon_{n+1} < \epsilon_n$. As a result, for any t,

$$\sum_{j=0}^{t-1} \frac{1}{K-j} = \sum_{j=1}^{K} \frac{1}{j} - \sum_{j=1}^{K-t} \frac{1}{j}$$
$$= \ln(K) + \epsilon_K - \ln(K-t) - \epsilon_{K-t} < \ln(\frac{K}{K-t}),$$

$$\text{and} \max_{t=1,\ldots K} (K-t) \sum_{j=0}^{t-1} \frac{1}{K-j} \leq \max_{0 \leq t < K} (K-t) \ln(\frac{K}{K-t})$$
$$= \max_{0 < y \leq K} y \ln(\frac{K}{y}) = \frac{K}{e}.$$

where the last equality follows since the function $y \ln(\frac{K}{y})$ obtains its maximum at $y^* = \frac{K}{e}$ and $y^* \ln(\frac{K}{y^*}) = \frac{K}{e}$. Therefore we have shown that $E[Z_{t^*}] \leq \frac{K}{e}$, which implies the second part of theorem 2.

4 Observed Arrivals

4.1 A Second Modified English Auction

We now return to the original sequential English auction format, and suggest a second modification. While it may seem that bidders who do not belong to the top $K-t+1$ in round t have no chance to win in a later round even in the original format if players play undominated strategies, in fact this is not true and very strange strategies are in fact undominated. This is true even in the special case of two items and three bidders, and in particular we show in Appendix A that for this special case, the strategy of quitting at price 0 in the first auction, and remaining until price equals value in the second auction, is *not* dominated. Clearly without ruling out such strategies we cannot hope to obtain good efficiency bounds, therefore we suggest the following "activity rule" that will force players to compete at least until a carefully-chosen cut-off point, where the number of remaining players is equal to the number of remaining items:

An English auction with an activity rule: At auction t (for $t = 1, ..., K$), let p_t be the price point at which there remain exactly $K-t+1$ bidders who have not dropped yet from the current auction (this is the point where the number of remaining bidders is equal to the number of unsold items). If several bidders drop together at p_t, so that more than $K-t+1$ bidders are active before p_t and less than $K-t+1$ bidders are active after p_t, the auction orders them so that it has a set of players of size exactly $K-t+1$ that dropped last. We refer to p_t as the cutoff price at auction t (note that p_t can be strictly below the end price of auction t, as at p_t there are $K-t+1 > 1$ players that are still active). Then,

1. Bidders that do not belong to the set of $K-t+1$ bidders that dropped last are not allowed to participate in subsequent auctions. This implies that $K-t$ of the bidders that participate in this auction are qualified to participate in the next auction, and one additional bidder wins this auction.

2. The next auction $t+1$ starts from price p_t.

As was in the previous auction, we add K dummy players with value zero that arrive for the first auction, to ensure that the number of players in each auction is not less than the number of remaining items. Here as well, our efficiency bounds hold even if some of the dummy players are chosen to win some of the items (meaning that that item is left unallocated).

Since player i obtains zero utility if she drops before the cutoff price, it is clearly a dominated strategy to do so. Thus, we assume throughout that in every auction $t = r_i, ..., K$, bidder i does not drop before the cutoff price p_t, unless her value is lower than p_t. If indeed the price reaches the player's value, and this point is lower than p_t, the player drops at this point since subsequent prices will not be lower than p_t, and thus the player cannot obtain positive utility regardless of her actions.

We should remark at this point that the solution concept of undominated strategies is much stronger than the more standard ex-post equilibrium notion. For example, a player that follows an equilibrium strategy assumes that the other players are rational enough to also follow the same equilibrium, while a player that follows an undominated strategy does not need to assume *anything* about the other players. Thus the additional assumption of known arrivals enables us to significantly strengthen the solution concept.

To analyze the resulting efficiency, we first show that, as a direct result of this activity rule, the bidders with the $K-t+1$ highest values *among all bidders that arrive up to time t and have not won yet* are qualified for auction $t+1$. Formally, let Λ_t be the set of bidders that arrive up to time t and does not win any item $1, ..., t$. Let X_t denote the set of bidders that participate in auction t, and let $Q_t \subseteq \Lambda_t$ be the set of players at auction t that are qualified for auction $t+1$. We prove in the appendix that,

Proposition 2. *If $|X_t| < K-t+1$ then no player was disqualified at any auction $s \leq t$. If $|X_t| \geq K-t+1$ then the $K-t+1$ highest-value bidders in Λ_t have the same set of values as the bidders in Q_t.*

4.2 Worst-case Efficiency

We now turn to analyze the social efficiency of the sequential auction with the activity rule, under the assumption that players may choose to play any tuple of undominated strategies. We use the same notation of section 3: $W_A(\theta)$ denotes the resulting welfare of the auction, when players' types are θ, and they play some tuple of undominated strategies, and $W_{OPT}(\theta)$ is the optimal welfare for θ.

We wish to prove that $\frac{W_A(\theta)}{W_{OPT}(\theta)} \geq 1/2$ for any tuple of types θ. For two items the proof is simple, since the player with the highest value in θ must win an item: if this player arrived for the first auction but was not a winner in the first auction then by proposition 2 she was qualified for the second auction. In the second auction, since all players remain up to their values, the highest player wins. This immediately implies the claimed ratio. When we consider more items the appropriate generalization of this fact is the following lemma.

Fix any tuple of types θ. For simplicity of notation we omit repeating θ throughout. Let OPT be a valid assignment with maximal welfare, and let A be an assignment that results from the sequential auction with the activity rule, when all players play some tuple of undominated strategies. Let $v_1^{OPT}, ..., v_K^{OPT}$ be the values of the winners of OPT, ordered in a non-increasing order (i.e. $v_1^{OPT} \geq v_2^{OPT} \geq \cdots \geq v_K^{OPT}$). (we also set $v_{K+1}^{OPT} = 0$ for notational purposes). Similarly, let $v_1^A, ..., v_K^A$ be the values of the winners of A, again in a non-increasing order.

Lemma 2. *For any index $0 \leq l \leq \lfloor \frac{K}{2} \rfloor$, $v_{l+1}^A \geq v_{2l+1}^{OPT}$.*

Proof. Assume by contradiction that there are at most l winners in A with values that are larger or equal to v_{2l+1}^{OPT}. Let $K-t$ be the last auction at which the winner in A has value strictly smaller than v_{2l+1}^{OPT}. After this auction there remain exactly t more auctions, hence there are at least t players in A with value at least v_{2l+1}^{OPT}. Thus, by the contradiction assumption, $t \leq l$.

Let X be the set of players in OPT with the $2l+1$ highest values. Let $Y = \{i \in X \mid r_i \leq K-t\}$. Note that $|Y| \geq (2l+1)-t$: there are only t auctions after time $K-t$, so there are at least $(2l+1)-t$ players in X that receive an item in OPT at or before time $K-t$, and these must have an arrival time smaller or equal to $K-t$. Let Z be the set of players in Y that win in A before auction $K-t$. Thus $Y \setminus Z \subseteq \Lambda_{K-t}$.

Note that $|Z| \leq l-t$, since after auction $K-t$ all winners in A have values at least v_{2l+1}^{OPT} and all players in Z win in A an item from $1, ..., K-t$. Therefore by the contradiction assumption $|Z| + t \leq l$. This implies that $|Y \setminus Z| \geq (2l+1-t) - (l-t) = l+1 \geq t+1$.

Since $Y \setminus Z \subseteq \Lambda_{K-t}$ and $|Y \setminus Z| \geq t+1$, then the $(t+1)$-highest-value in Λ_{K-t} is larger or equal than the minimal value in $Y \setminus Z$. By proposition 2, the winner in A at auction $K-t$ must have value at least as large as the $(t+1)$ highest value in Λ_{K-t}. Thus the winner in A at auction $K-t$ has value at least as large as the minimal value in $Y \setminus Z$. But all players in $Y \setminus Z$ have values at least v_{2l+1}^{OPT}, and this contradicts our assumption that the winner of auction $K-t$ has value strictly smaller than v_{2l+1}^{OPT}. □

This lemma implies that $W_{OPT}(\theta)$ is at most twice $W_A(\theta)$, since $v_1^{OPT}, v_2^{OPT} \leq v_1^A$, $v_3^{OPT}, v_4^{OPT} \leq v_2^A$, and so on, and thus $W_{OPT}(\theta) \leq 2\sum_{k=1}^{\lfloor \frac{K}{2} \rfloor + 1} v_k^A \leq 2W_A(\theta)$.

Theorem 3. *$W_{OPT}(\theta) \leq 2W_A(\theta)$ for any θ.*

The following simple example shows that the analysis is tight. Suppose two items and two players that arrive at time 1, with values $v_1 = 0, v_2 = 1$. The cutoff price at the first auction is therefore zero, and therefore suppose that player 1 wins the first auction, and player 2 continues to the second auction. In the second auction arrives a third player with $v_3 = 1$. Regardless of the winner in the second auction, the resulting welfare of the auction is 1, while the optimal welfare is 2.

4.3 Average-case Efficiency for Two Items

The traditional worst-case analysis is very pessimistic, and it would be more reasonable to assume that the input is not completely adversarial. In this section we will concentrate on the special case where there are only two items for sale (for which the worst-case bound is also 2), and demonstrate that even a minor shift from the worst-case setting towards an average-case setting will improve the efficiency guarantee quite significantly.

Formally, we assume an adversary that is allowed to choose the number of players, n, and their arrival times. Thus, the adversary determines a number $r \leq n$, such that the first r players arrive for the first auction, and the remaining $n-r$ players arrive for the second auction. The adversary then chooses a cumulative probability distribution F with some support in $[0, \infty]$, and draws the values of the players from this distribution, i.e. the values are i.i.d. The adversary then determines the undominated strategy of each player (the choice of the strategy may depend on the random result of the players' values, as to "fail" the auction). Comparing this setup to the setup of the previous section, we can see that the only change is that now the adversary must draw the players' values from some fixed distribution (but the adversary can choose what distribution to use). We will show that this modification towards an average-case setup implies a significant increase in the efficiency of the sequential auction: the auction will obtain at least $\sqrt{2}/2 \simeq 0.7$ of the optimal efficiency, for any number of players, their arrival times, the chosen distribution of the players' values, and any choice of undominated strategies. Moreover, we show that this bound is tight, i.e. there exists a sequence of distributions that approach this efficiency guarantee in the limit.

The analysis is carried out in the following way. Fixing the number of players, n, the number of players $r \leq n$ that arrive for the first auction, and the cumulative distribution F, we define two random variables: $OPT_{r,n}$ is the highest value among all players that arrive at time 1 plus the highest value among all other players (including those that arrive at time 2). Note that $OPT_{r,n}$ is indeed equal to the optimal welfare, given a specific realization of the values. The second random variable, $\tilde{A}_{r,n}$, is equal to the second highest value among all players that arrive at time 1 plus the highest value among all the remaining players (including those that arrive at time 2). By proposition 2 the winner in the first auction has a value larger or equal to the second highest value among all players present in the first auction, and the winner in the second auction has the largest value among all remaining players. Thus, A's welfare is at least the value of $\tilde{A}$.[3] We show:

Theorem 4. *For any choice of the parameters n, r, F,*

$$\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \frac{\sqrt{2}}{2} \simeq 0.707.$$

The proof proceeds in two parts. We first analyze Bernoulli distributions over $\{0, 1\}$, and bound the ratio of expectations over all such distributions. We then show in a formal way that the worst-case over all distributions is lower bounded by the worst-case over all Bernoulli distributions.

4.3.1 A bound on any Bernoulli distribution

We have n players with i.i.d. values such that $\Pr(v_i = 0) = p$ and $\Pr(v_i = 1) = 1 - p$ for some $0 \leq p < 1$. Players $1, ..., r$ arrive for the first auction (at time 1), and players $r + 1, ..., n$ arrive for the second auction, at time 2, where p, n, r are parameters. We ask what values of p, n, r will minimize the ratio $\frac{E_{F_p}[\tilde{A}_{n,r}]}{E_{F_p}[OPT_{n,r}]}$, where F_p denotes the above-mentioned Bernoulli distribution.

Observe that, since a player's value is either zero or one, the random variables $OPT_{n,r}$ and $\tilde{A}_{n,r}$

[3] There always exists a tuple of undominated strategies such that A's value exactly equals $\tilde{A}$'s value. Recall that the adversary here may choose the strategy after she knows the values of the bidders. Therefore she can choose the following strategies: at the first auction, the player with the highest value quits when the number of remaining bidders is equal to the number of remaining items and all other players quit when the price reaches their value. (In the second auction the only allowed strategy is to quit when price equals value). These strategies ensure that the winner of the first auction is the player with the second highest value.

can take only the values $0, 1, 2$. We calculate:

$$\Pr(\tilde{A}_{n,r} = 0) = p^n,$$
$$\Pr(\tilde{A}_{n,r} = 1) = p^r(1 - p^{n-r}) + r(1 - p)p^{r-1},$$
$$\Pr(\tilde{A}_{n,r} = 2) = 1 - p^r - r(1 - p)p^{r-1}.$$

For example, $\tilde{A} = 1$ if all values at the first auction are 0 and at least one value at auction 2 is 1 (this happens with probability $p^r(1 - p^{n-r})$), or if there exists exactly one value that is equal to 1 at the first auction, and then it does not matter what are the values at the second auction (this happens with probability $r(1 - p)p^{r-1}$). Similarly, we also have that $\Pr(OPT_{n,r} = 0) = p^n$, $\Pr(OPT_{n,r} = 1) = p^r(1-p^{n-r})+r(1-p)p^{n-1}$, and $\Pr(OPT_{n,r} = 2) = 1-p^r-r(1-p)p^{n-1}$. Using this, we show in appendix G that

Proposition 3. *For any n, r, and $0 \leq p < 1$,*

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]} = \frac{2 - p^r - p^n - r(1 - p)p^{r-1}}{2 - p^r - p^n - r(1 - p)p^{n-1}} \geq \frac{\sqrt{2}}{2}$$

The calculations first show that this ratio decreases with n (for any r, p), so it suffices to compute a lower bound on the limit of the ratio of expectations when $n \to \infty$. In that case, a minimum is achieved for $r = 2$ and $p = 2 - \sqrt{2}$. Note that for $p = 1$, the two expectations become zero and the ratio is undefined.

The worst-case scenario of section 4.2 requires three players, but here we need the infinitely many players to approach the worst ratio. These additional players arrive for the *second* auction, which may seem counterintuitive at first, as in the second auction the player with the highest value wins. What is the effect of adding more players to the second auction? To explain this, note that the only events that differentiate OPT and $\tilde{A}$ are those in which, at the first auction, exactly one player has value 1 and the other players have value 0, and at the second auction there exists at least one additional player with value 1. In these events OPT $= 2$ and $\tilde{A} = 1$. As the number of players increases (while keeping r constant), these events get more probability, hence the above-mentioned effect.

4.3.2 Generalizing to any other distribution

To explore the case of a general distribution F with a support in $[0, \infty)$, we must take a closer look at the expression for the expectation of OPT and $\tilde{A}$. We denote by $X_{n-j:n}$ the j'th order statistic of the random variables $v_1, ..., v_n$ (the players' values), which denotes the $(j+1)$'th highest value of the players, i.e., $X_{n:n}$ is a random variable that takes the maximal value among $v_1, ..., v_n$; $X_{n-1:n}$ is a random variable that takes the second largest value among $v_1, ..., v_n$, and so on. If the player with the highest value at time 1 has the $j+1$'th highest value among all the n players[4], then $OPT_{n,r} = X_{n-j:n} + X_{n:n}$. Hence

$$E[OPT_{n,r}|\text{ highest at time 1 is (j+1)-highest overall }]$$
$$= E[X_{n-j:n} + X_{n:n}].$$

[4]If the distribution is discrete we use an arbitrary deterministic tie-breaking rule to ensure that the events (indexed by j) "highest player at time 1 has the (j+1)-highest value overall" are mutually exclusive.

Denote by $q_j^{n,r}$ the probability that the highest value at time 1 is the $j+1$'th highest value among all the n players. It follows that:

$$E_F[OPT_{n,r}] = q_0^{n,r}(E_F[X_{n-1:n}] + E_F[X_{n:n}]) \\ + \sum_{j=1}^{n} q_j^{n,r}(E_F[X_{n-j:n}] + E_F[X_{n:n}]). \tag{1}$$

We remark that the highest player among the players that arrive at time 1 is at least the $n-r+1$ highest player among all players; therefore $q_j^{n,r} = 0$ when $j > n-r$. It will be important for the sequel to verify that the probability $q_j^{n,r}$ does not depend on the distribution F. First, note that since the values are drawn i.i.d. then each value-ordering of the players has equal probability. Thus, the probability of any specific order of all the players is $1/n!$, and the probability that the order of values will satisfy any specific property is simply the number of orderings that satisfy this property, divided by $n!$. To find $q_j^{n,r}$, we thus ask in how many orderings, the highest player among the first r players is exactly the $j+1$ highest among all players. To get one such ordering, one needs to choose one player (say i) out of the r players of time 1 (this is the highest player at time 1), to choose j players out of the $n-r$ players of time 2 (these are the players that are higher than i), to order them in one of the $j!$ orderings, then to place i, and then to order the remaining $n-j-1$ players. Thus, for any $0 \le j \le n-r$, $q_j^{n,r} = \frac{1}{n!} \cdot r \cdot \binom{n-r}{j} \cdot j! \cdot (n-j-1)!$, and $q_j^{n,r} = 0$ for any $n-r+1 \le j \le n$.

Similarly, given that the second-highest player at time 1 is the $j+1$'th highest player among all the n players, the expected welfare of $\tilde{A}$ is $E[X_{n-j:n} + X_{n:n}]$. Denoting by $p_j^{n,r}$ the probability that the second-highest player at time 1 is the $j+1$'th highest player among all the n players (where again this probability does not depend on F), it follows that:

$$E_F[\tilde{A}_{n,r}] = \sum_{j=1}^{n} p_j^{n,r}(E_F[X_{n-j:n}] + E_F[X_{n:n}]) \tag{2}$$

(and we set $p_j^{n,r} = 0$ for any $n-r+2 \le j \le n$).

We now consider the terms $E[X_{n-j:n}]$. Let $F_{n-j:n}(x)$ be the probability distribution of $X_{n-j:n}$. The probability that $X_{n-j:n} \le x$ is the probability that at most j values will be higher than x, and the remaining at least $n-j$ values will be smaller than x, or, in other words,

$$F_{n-j:n}(x) = \Pr(X_{n-j:n} \le x) = \sum_{k=0}^{j} \binom{n}{k} (1-F(x))^k (F(x))^{n-k}$$

Therefore, $F_{n-j:n}(x)$ is a polynomial in $F(x)$, where the coefficients of the polynomial do not depend on the distribution F. A well-known formula for the expectation of an arbitrary non-negative random variable Y with cumulative distribution G is $E[Y] = \int_0^\infty (1-G(y))dy$. In particular, $E[X_{n-j:n}] = \int_0^\infty (1-F_{n-j:n}(x))dx$. In other words, the expectation of the j'th order statistic is an integration over a polynomial in $F(x)$, i.e. there exist coefficients $w_l^{(j)}$ for $l = 0, ..., n$ and $j = 1, ..., n$ (that does not depend on the distribution F) such that $E_F[X_{n-j:n}] = \int_0^1 [\sum_{l=0}^n w_l^{(j)} (F(x))^l] dx$. Combining this equation with equations (2) and (1), we get that both $E_F[OPT_{n,r}]$ and $E_F[\tilde{A}_{n,r}]$ are an integration over a polynomial in $F(x)$, i.e. there exist coefficients $\beta_0^{(n,r)}, ..., \beta_n^{(n,r)}$ and $\gamma_0^{(n,r)}, ..., \gamma_n^{(n,r)}$, that do not depend on the distribution F, such that $E_F[OPT_{n,r}] = \int_0^\infty [\sum_{l=0}^n \beta_l^{(n,r)} (F(x))^l] dx$ and $E_F[\tilde{A}_{n,r}] = \int_0^\infty [\sum_{l=0}^n \gamma_l^{(n,r)} (F(x))^l] dx$.

One additional important observation is that $\sum_{l=0}^{n} \beta_l^{(n,r)} = \sum_{l=0}^{n} \gamma_l^{(n,r)} = 0$. To see this, take some distribution F with a bounded support, say $[0,1]$. The above equality implies that $E_F[OPT_{n,r}] > \int_1^{\infty} [\sum_{l=0}^{n} \beta_l^{(n,r)}]dx$, which is unbounded if $\sum_{l=0}^{n} \beta_l^{(n,r)} \neq 0$. But clearly $E_F[OPT_{n,r}]$ is a finite number since the support is bounded; hence it must be that $\sum_{l=0}^{n} \beta_l^{(n,r)} = 0$. The same argument implies that $\sum_{l=0}^{n} \gamma_l^{(n,r)} = 0$.

The Bernoulli distribution F_p ($0 \leq p < 1$) gives a fixed function over the interval $[0,1)$, specifically $F_p(x) = p$ for any $0 \leq x < 1$, and $F_p(x) = 1$ for $x \geq 1$. Thus for this distribution the integration cancels out, and,

$$E_{F_p}[OPT_{n,r}] = \sum_{l=0}^{n} \beta_l^{(n,r)} p^l, E_{F_p}[\tilde{A}_{n,r}] = \sum_{l=0}^{n} \gamma_l^{(n,r)} p^l. \tag{3}$$

We next show how all the above implies:

Proposition 4. *Fix any α such that $\frac{E_{F_p}[\tilde{A}_{n,r}]}{E_{F_p}[OPT_{n,r}]} \geq \alpha$, for any n, r and $0 \leq p < 1$. Then, for any other cumulative distribution F with $E_F[OPT_{n,r}] > 0$, it must be that $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \alpha$.*

Proof. We need to show that $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \alpha$, or, equivalently, that $E_F[\tilde{A}_{n,r}] - \alpha E_F[OPT_{n,r}] \geq 0$. Using the above equations, this term becomes

$$\int_0^{\infty} [\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l] dx.$$

We will show that, for every $x \geq 0$, $\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l \geq 0$, which implies the above inequality. Fix some $x \geq 0$, if $F(x) = 1$ then indeed

$$\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l = 0 - \alpha \cdot 0 = 0.$$

Otherwise, denote $p = F(x) < 1$. Thus

$$\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l = \\ E_{F_p}[\tilde{A}_{n,r}] - \alpha E_{F_p}[OPT_{n,r}] \geq 0,$$

where the equality follows from Eq. (3), and the inequality follows from the assumption in the claim. □

Corollary 1. *For any cumulative probability distribution F, and any n, r, $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \frac{\sqrt{2}}{2} \simeq 0.707$.*

This completes the proof of theorem 4.

References

Bartal, Y., F. Chin, M. Chrobak, S. Fung, W. Jawor, R. Lavi, J. Sgall and T. Tichỳ (2004). "Online competitive algorithms for maximizing weighted throughput of unit jobs." In *Proc. of the 21st Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 187–198.

Cole, R., S. Dobzinski and L. Fleischer (2008). "Prompt mechanisms for online auctions." In *Proc. of the 1st International Symposium on Algorithmic Game Theory (SAGT'08).*

Dobzinski, S., R. Lavi and N. Nisan (2008). "Multi-unit auctions with budget limits." In *Proc. of the 49th annual Symposium on Foundations of Computer Science (FOCS).*

Hajiaghayi, M., R. Kleinberg, M. Mahdian and D. Parkes (2005). "Online auctions with re-usable goods." In *Proc. of the 6th ACM Conf. on Electronic Commerce (ACM-EC'05).*

Lavi, R. and N. Nisan (2005). "Online ascending auctions for gradually expiring items." In *Proc. of the 16th Symposium on Discrete Algorithms (SODA).*

Said, M. (2008). "Information revelation in sequential ascending auctions." In *Proc. of the 9th ACM Conf. on Electronic Commerce (ACM-EC'08).*

Williams, D. (1991). *Probability with martingales.* Cambridge Univ Press.

A The need for an activity rule

To exemplify the need to add an activity rule to the auction, consider a setting of two items and three bidders that arrive in time 1. In the second auction, it is rather immediate that the strategy of "remaining until price reaches value" weakly dominates all other strategies. In the first auction, it is not hard to verify that there are at least two weakly undominated strategies:

1. In both auctions, remain until price reaches value (call this b_i^{value}).
2. In the first auction, remain until exactly one other bidder remains, and in the second auction, remain until your value (call this b_i^{EPD}, where EPD stands for Earliest Possible Dropping).

With our activity rule, these are the only undominated strategies, when there are two items. Without the activity rule, these two strategies do not weakly dominate all other strategies, as we next show. Suppose three bidders $1, 2, 3$ arrive at time 1, with $v_1 > v_2 > v_3$, and let us consider the strategy b_1^0 for bidder 1, in which she drops at price 0 in the first auction, and remains until her value in the second auction. None of the above two strategies dominates b_1^0, due to the following reasoning.

Consider first the strategy b_1^{value}. This strategy performs strictly worse than b_1^0 in case both 2 and 3 choose to do the same (remain until their value in both auctions, i.e. play b_2^{value}, b_3^{value}), due to the following. By playing b_1^{value}, bidder 1 will win the first auction and will pay v_2. By playing b_1^0, bidder 1 will lose the first auction and will win the second auction for a lower price of v_3.

Consider next the strategy b_1^{EPD}. This strategy performs strictly worse than b_1^0 when bidder 3 plays the strategy b_3^{value} and 2 uses the following strategy: In the first auction, if bidder 1 drops at price 0 then bidder 2 continues until her value, and if bidder 1 does not drop at price 0 then bidder 2 drops immediately after that. In the second auction, bidder 2 remains until her value[5]. In this case, if bidder 1 follows b_1^0 and drops at 0 then bidder 2 will win the first auction, bidder 1 will win the second auction, and will pay v_3. If bidder 1 follows b_1^{EPD} then bidder 2 drops, and bidder 1 drops immediately after that (since now only bidder 3 remains besides 1). Thus, bidder 3 wins the first auction, bidder 1 again wins the second auction, but this time pays v_2 which is larger than v_3.

B Proof of Claim 3

Throughout the following appendices we use the following notation. For a scenario θ, denote $\theta|_{2,\ldots,K} = \{ \ (v_i, \max(r_i, 2) \mid (v_i, r_i) \in \theta \ \}$, i.e. it is a scenario on $K-1$ items with the same players' types, except that players that previously arrived for the first auction now arrive for the second auction, and the first auction is virtually deleted. Note that when j is the winner of the first auction then $\theta_{-j}|_{2,\ldots,K}$ describes the scenario for the remaining K-1 items auction.

Fix any two scenarios θ, θ' with the same set of players, such that $r_i = r'_i$ for every player i, and for any two players i, j, $v_i \geq v_j$ if and only if $v'_i \geq v'_j$. We show that there exists a specific B completion rule such that, when using this rule, the probability that a given player wins is the same in both scenarios.

Suppose the following B completion rule: if either $v_i > v_j$ or $v'_i > v'_j$ then i is preferred over j (if $v_i = v_j$ and $v'_i = v'_j$ the preference is arbitrary). We prove the claim by induction on the number of items K. For $K = 1$ the claim is immediate. Now assume correctness for $K-1$. Let B_1, B'_1 denote the set of possible winners in the first auction for θ, θ', respectively. Because of the specific B completion rule, $B_1 = B'_1$, and the induction assumption implies $Pr(i \text{ wins in } \theta_{-j}|_{2,\ldots,K}) = Pr(i \text{ wins in } \theta'_{-j}|_{2,\ldots,K})$. Thus,

$$\begin{aligned}
&Pr(i \text{ wins in } \theta) = Pr(i \text{ wins 1st auction in } \theta) + \\
&\textstyle\sum_{j \in B_1 \setminus \{i\}} Pr(j \text{ wins 1st auction in } \theta) \cdot Pr(i \text{ wins in } \theta_{-j}|_{2,\ldots,K}) \\
&= Pr(i \text{ wins 1st auction in } \theta') + \\
&\textstyle\sum_{j \in B'_1 \setminus \{i\}} Pr(j \text{ wins 1st auction in } \theta') \cdot Pr(i \text{ wins in } \theta'_{-j}|_{2,\ldots,K}) \\
&= Pr(i \text{ wins in } \theta')
\end{aligned}$$

C Proof of Claim 4

We prove that changing the B completion rule does not change the expected welfare of the auction by induction. For $K = 1$ the claim is immediate. For $K > 1$, the expected welfare obtained in the first auction is the same, regardless of the B completion rule, since any player chosen by any B completion rule has value that is equal to p_1. The B completion rule also does not change the set of types that continue to the second auction (i.e. there exists a one-to-one mapping π from players that continue in one B completion to players that continue in the other B completion such that $v_i = v_{\pi(i)}$ for any player i), hence by the inductive assumption the expected welfare of the auction is the same.

[5] Formally, $b_2((1, v_2), h(t = 1, p, k)) = R$ iff $[p = 0 \text{ or } (0 < p < v_3 \text{ and } 1 \notin I_1(0, s_0^1))]$.

D Proof of Claim 6

Fix a scenario θ such that $v_j \in \{0,1\}$ for any player j. Additionally let i be some player with $r_i = 1$ and $v_i = 1$. We show that $E[W_A(\theta)] \leq 1 + E[W_A(\theta_{-i})]$.

We prove by induction on the number of items K. For $K = 1$, $E[W_A(\theta)] \leq 1$ which implies the claim. Now assume correctness for $K-1$ and let us prove for K. Let p be the probability that a 1-player will be chosen at the first auction. We have

$$\begin{aligned} E[W_A(\theta)] \leq & p[1 + E[W_A(\theta_{-i}|_{2,\ldots,K})]] \\ & + (1-p)E[W_A(\theta|_{2,\ldots,K})]. \end{aligned}$$

By the induction assumption

$$E[W_A(\theta|_{2,\ldots,K})] \leq 1 + E[W_A(\theta_{-i}|_{2,\ldots,K})],$$

and the claim follows.

E Proof of Claim 7

Fix two scenarios θ, θ' such that $v_j \in \{0,1\}$ for any player j in θ, and θ' is a modification of θ such that one of the following holds:

- one player arrives one auction later, i.e. there exists a player i such that $\theta'_j = \theta_j$ for any $j \neq i$ and $\theta'_i = (v_i, r_i + 1)$ (if $r_i = K$ then $\theta' = \theta_{-i}$), or
- the set of types in θ' includes θ plus some additional zero players that arrive for the first auction.

We show that $E[W_A(\theta)] \geq E[W_A(\theta')]$. We prove by induction on the number of items K. For $K = 1$ the claim is immediate. Assume correctness for $K-1$ and let us prove for K. Let p, p' be the probability that a 1-player will be chosen at the first auction in θ, θ', respectively. If all players that arrive for the first auction have value 0 the claim clearly holds. Otherwise fix some 1-player j with $r_j = 1$. We have

$$\begin{aligned} E[W_A(\theta)] =& p[1 + E[W_A(\theta_{-j}|_{2,\ldots,K})]] \\ & + (1-p)E[W_A(\theta|_{2,\ldots,K})], \text{ and} \\ E[W_A(\theta')] =& p'[1 + E[W_A(\theta'_{-j}|_{2,\ldots,K})]] \\ & + (1-p')E[W_A(\theta'|_{2,\ldots,K})]. \end{aligned}$$

To prove that $E[W_A(\theta)] \geq E[W_A(\theta')]$ we use the following fact: given two real intervals $[a,b]$ and $[a',b']$ such that $a' \leq a$ and $b' \leq b$, and $0 \leq p' \leq p \leq 1$, one can verify that $pb + (1-p)a \geq p'b' + (1-p')a'$. We set $a = E[W_A(\theta|_{2,\ldots,K})]$, $b = 1 + E[W_A(\theta_{-j}|_{2,\ldots,K})]$. Note that claim 6 implies $a = E[W_A(\theta|_{2,\ldots,K})] \leq 1 + E[W_A(\theta_{-j}|_{2,\ldots,K})] = b$. Similarly we set $a' = E[W_A(\theta'|_{2,\ldots,K})]$, $b' = 1 + E[W_A(\theta'_{-j}|_{2,\ldots,K})]$ and again note that $a' \leq b'$. If the player i from the statement of the claim is player j then $\theta_{-j}|_{2,\ldots K} = \theta'_{-j}|_{2,..,K}$ and $\theta|_{2,\ldots,K} = \theta'|_{2,\ldots,K}$. Otherwise the induction assumption implies

$$E[W_A(\theta_{-j}|_{2,\ldots,K})] \geq E[W_A(\theta'_{-j}|_{2,\ldots,K})] \text{ and}$$
$$E[W_A(\theta|_{2,\ldots,K})] \geq E[W_A(\theta'|_{2,\ldots,K})].$$

Thus $a' \leq a$ and $b' \leq b$. Since $p' \leq p$ the claim follows.

F Proof of Proposition 2

We first show that if $|X_t| < K - t + 1$ then no player was disqualified at any auction $s \leq t$. We prove that if a player is disqualified at auction s then for every $t \geq s$, $|X_t| \geq K - t + 1$, by induction on t. For $t = s$, since some player was disqualified, then by definition $|X_s| \geq K - s + 1$. Assume the claim is true for t, and let us verify it for $t + 1$. Since $|X_t| \geq K - t + 1$ then by definition $|Q_t| = K - t + 1$; hence $|X_{t+1}| \geq |Q_t| - 1 = K - t = K - (t + 1) + 1$, and the claim follows.

We now show that if $|X_t| \geq K - t + 1$ then the $K - t + 1$ highest-value bidders in Λ_t have the same set of values as the bidders in Q_t. We show that for every player $i \in \Lambda_t \setminus Q_t$ and player $j \in Q_t$ we have $v_j \geq v_i$, which implies the claim.

First observe that this holds for $i \in X_t \setminus Q_t$ and $j \in Q_t$: all players in X_t but not in Q_t have values smaller or equal to the cutoff price p_t, and all players in Q_t have values greater or equal to p_t, hence $v_j \geq v_i$. This also implies that the $K - t + 1$ highest-value bidders in X_t belong to Q_t.

We now prove the claim by induction on t. For $t = 1$, $\Lambda_t = X_t$ and the claim follows from the above argument. Assume the claim is correct for any $t' < t$, and let us prove it for t. If $i \in X_t$ then again the above argument holds. Otherwise $i \in \Lambda_t \setminus X_t$, which implies that i arrived strictly before time t and was disqualified at or before time $t - 1$. Note that by the first part of this proposition we have that $|Q_{t-1}| = K - (t - 1) + 1$. Since player i was disqualified, we have $i \in \Lambda_{t-1} \setminus Q_{t-1}$. Let j' be the player with minimal value in Q_{t-1}. By the induction assumption we have that $v_{j'} \geq v_i$. There are $K - t + 2$ players in Q_{t-1}, out of them $K - t + 1$ continue to auction t (i.e. belong to X_t), all with values larger or equal to $v_{j'}$. Thus, again by the argument in the previous paragraph, any player $j \in Q_t$ has $v_j \geq v_{j'} \geq v_i$, and the claim follows.

G Proof of Proposition 3

We need to show that, for any n, r, and $0 \leq p < 1$,

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]} = \frac{2 - p^r - p^n - r(1-p)p^{r-1}}{2 - p^r - p^n - r(1-p)p^{n-1}} \geq \frac{\sqrt{2}}{2}. \tag{4}$$

We first differentiate this expression with respect to n, to show that it decreases as n increases.

$$\frac{d}{dn}\left(\frac{\left(2 - p^r - p^n - r\left(1-p\right)p^{r-1}\right)}{\left(2 - p^r - p^n - r\left(1-p\right)p^{n-1}\right)}\right) =$$
$$rp^{n-2}\left(\ln p\right)\left(1-p\right)\frac{\left(2p - rp^r + p^{r+1}\left(r-2\right)\right)}{\left(-rp^{n-1} - p^r + \left(r-1\right)p^n + 2\right)^2}.$$

We concentrate on the term

$$G\left(p,r\right)=2p-rp^{r}+p^{r+1}(r-2)$$

and claim that it is nonnegative for every $2\leq r\leq n$ and $p\in[0,1]$. For $p=0$ we have $G(0,r)=0$ and for $p=1$ we have $G(1,r)=0$. Moreover

$$\frac{d^{2}}{dp^{2}}G\left(p,r\right)=rp^{r-2}\left(\left(r+1\right)\left(r-2\right)p-r\left(r-1\right)\right)$$

and since

$$\frac{r\left(r-1\right)}{\left(r+1\right)\left(r-2\right)}>1$$

for every $r\geq 2$ we know that $\frac{d^{2}}{dp^{2}}G\left(p,r\right)\leq 0$ and $G\left(p,r\right)$ is concave in p. We thus conclude that $G\left(p,r\right)\geq 0$ for every $2\leq r\leq n$ and $p\in[0,1]$ and consequently that $\frac{d}{dn}\left(\frac{\left(2-p^{r}-p^{n}-r(1-p)p^{r-1}\right)}{\left(2-p^{r}-p^{n}-r(1-p)p^{n-1}\right)}\right)\leq 0$. We take n to infinity:

$$\lim_{n\to\infty}\left(\frac{\left(2-p^{r}-p^{n}-r\left(1-p\right)p^{r-1}\right)}{\left(2-p^{r}-p^{n}-r\left(1-p\right)p^{n-1}\right)}\right)=1-\frac{r\left(1-p\right)p^{r-1}}{\left(2-p^{r}\right)}$$

and therefore the minimum of $1-\frac{r(1-p)p^{r-1}}{(2-p^{r})}$ will give us a lower bound for (4), for every n, since we obtained that (4) decreases towards the limit as n increases. Equivalently, we look for the maximum of $H\left(p,r\right)=\frac{r(1-p)p^{r-1}}{(2-p^{r})}$. Now,

$$\begin{aligned}\frac{d}{dr}H\left(p,r\right)&=-\frac{p^{r-1}\left(1-p\right)\left(-2r\ln p+p^{r}-2\right)}{\left(2-p^{r}\right)^{2}}\\ \frac{d}{dp}H\left(p,r\right)&=rp^{r-2}\frac{\left(2r\left(1-p\right)+p^{r}-2\right)}{\left(2-p^{r}\right)^{2}}.\end{aligned}$$

Therefore if there exists a global maximum at $0<p<1$ and $2<r$ then we must have $-2r\ln p=2-p^{r}$ and $2r\left(1-p\right)=2-p^{r}$ but this is not possible since for every $0<p<1$ we have $-\ln p>1-p$. We thus conclude that the maximum is achieved on the boundary. For $p=0$, we have $H\left(0,r\right)=0$ and for $p=1$, we have $H\left(1,r\right)=0$; therefore we conclude that the maximum is achieved on the boundary where $r=2$. We find p that solves

$$\max_{p\in(0,1)}H\left(p,2\right)=\max_{p}\frac{2\left(1-p\right)p}{\left(2-p^{2}\right)}$$

and the solution is $p^{*}=2-\sqrt{2}$. Finally, for $r=2,p^{*}=2-\sqrt{2}$ and $n\to\infty$ we have

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]}=\frac{\sqrt{2}}{2}\simeq 0.707\,11.$$