

Auctions with Online Supply

Moshe Babaioff
Microsoft Research
Mountain View, CA 94043
moshe@microsoft.com

Liad Blumrosen
Microsoft Research
Mountain View, CA 94043
liadbl@microsoft.com

Aaron L. Roth
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15217
alroth@cs.cmu.edu

May 21, 2009

Abstract

We study the problem of selling identical goods to n unit-demand bidders in a setting in which the total *supply* of goods is unknown to the mechanism. Items arrive dynamically, and the seller must make the allocation and payment decisions online with the goal of maximizing social welfare. We consider two models of unknown supply: the adversarial supply model, in which the mechanism must produce a welfare guarantee for any arbitrary supply, and the stochastic supply model, in which supply is drawn from a distribution known to the mechanism, and the mechanism need only provide a welfare guarantee in expectation.

Our main result is a separation between these two models. We show that all *truthful* mechanisms, even randomized, achieve a diminishing fraction of the optimal social welfare (namely, no better than a $\Omega(\log \log n)$ approximation) in the adversarial setting. In sharp contrast, in the stochastic model, under a standard *monotone hazard-rate* condition, we present a truthful mechanism that achieves a constant approximation. We show that the monotone hazard rate condition is necessary, and also characterize a natural subclass of truthful mechanisms in our setting, the set of *online-envy-free* mechanisms. All of the mechanisms we present fall into this class, and we prove almost optimal lower bounds for such mechanisms. Since auctions with unknown supply are regularly run in many online-advertising settings, our main results emphasize the importance of considering distributional information in the design of auctions in such environments.

1 Introduction

Auctions have recently received attention in computer science because they crystalize many of the incentive issues in algorithmic game theory, and have direct application to the fast-growing market for online advertising. This paper belongs to a line of research that studies *online mechanism design*, which focuses on markets in which decisions are made dynamically before information regarding the state of the world has been fully revealed. Previous work in online mechanism design mainly concerned settings where *customers* that arrive dynamically compete for buying a known set of items (see a recent survey [25]). However, in many real-world settings the *supply* arrives dynamically and the exact number of items for sale is uncertain. This, for example, is the case in the sale of clicks on banner ads, where the number of clicks is not known in advance to the seller; Such a seller must decide which advertisement to show in a fraction of a second after the item arrives, while the future supply is uncertain.¹

In this work, we investigate a natural online setting, in which a mechanism must allocate items to a fixed set of bidders when the supply of items is unknown, and arrives online. We require that the mechanism allocates items and extracts payment for them as they arrive. The restriction that the mechanism extract payment at the time of sale is a natural practical constraint, and is satisfied by most real-world markets. Even in markets in which customers are able to defer their payments (such as auctions for search ads), the seller typically calculates payments immediately, which allows customers to better keep track of their spending. We introduce a stochastic model where the seller knows how the supply is distributed, but we do not assume any prior distribution on the bidders' valuations, nor do we require that the bidders know how the supply is distributed. One of the conceptual contributions of our paper is this *hybrid stochastic model*, in which the supply is drawn from some prior distribution, but no distributional assumptions are made on the preferences of the bidders. This captures scenarios such as online advertising, in which sellers can easily collect statistics on the supply (e.g., number of ad impressions per day) but obtaining statistics on the actual valuations of the bidders is harder and may require modeling, for example, their equilibrium behavior. Most of the recent work in computer science on online mechanism design has been in the fully adversarial setting, when in actuality, mechanism designers have a wealth of distributional information at their disposal. In economics, at the other extreme, dynamic mechanism design has been recently studied in a full Bayesian setting that assumes the existence of prior distributions on the bidders' preferences.

We wish to maximize social welfare, which is a desirable goal even from the perspective of a for-profit seller that does not have the luxury of operating under monopoly conditions. An economically efficient market (one that maximizes the *combined* welfare of the customers *and* the seller) will be more attractive to customers, and avoids harming the seller in the long term at the expense of short-term profits. In fact, the generalized second price auction currently used to sell search advertisements has social welfare, rather than revenue guarantees [10].

We explore the cost of ignoring distributional information. We produce a strong separation: Our main results are *lower bounds* in the adversarial setting, and *truthful approximation mechanisms* in the stochastic setting.

Notably, the algorithmic problem that we face is simple. If bidder valuations were known, then the greedy algorithm which simply allocated each arriving item to the unsatisfied bidder with the highest value would achieve optimal social welfare even in the adversarial supply setting. The difficulty of the problem stems from the fact that bidders may misrepresent their valuations for personal gain. Any allocation rule that we design must be associated with a corresponding payment rule which incentivizes bidders to truthfully report their valuations. As we shall show, the incentive constraint proves to be an insurmountable barrier to developing mechanisms guaranteeing a constant approximation to social welfare in the adversarial supply setting, but can be overcome in the stochastic supply setting.

1.1 Our Results

We first consider the adversarial supply setting in which welfare guarantees are required to hold for any realization of supply. Our first main result are lower bounds on the approximation obtainable by truthful

¹ Uncertainty on the supply appears in various environments. More examples include markets for computing resources and also traditional markets, like agricultural markets, where produce and fish continue to arrive after markets has been opened.

mechanisms:

Theorem: *Every truthful mechanism achieves a diminishing fraction (in the number of bidders) of the optimal social welfare. Specifically, no deterministic truthful mechanism achieves better than n -approximation and no randomized truthful mechanism achieves better than $\Omega(\log \log n)$ -approximation.*

The linear lower bound is simple, and is in the spirit of the lower bound given by Lavi and Nisan [19] for a model in which bidders that arrive online bid for a fixed set of expiring items. We note that an n -approximation to social welfare can be achieved by the trivial mechanism which simply allocates the first item to the highest bidder at the second highest price, and does not allocate any additional items. The randomized lower bound is more technically challenging. To prove it we give a characterization of truthful mechanisms in our setting, and a distribution over bidder values. From this, we derive a system of equations that can be simultaneously satisfied only if there exists a mechanism which achieves a strong welfare guarantee when given this distribution over bidders. We show that no such satisfying assignment exists, which gives the lower bound.

If we further require that our mechanisms be *online envy-free* (a desirable fairness property that we define in section 5), we can strengthen the above lower bound to show that no randomized truthful mechanism can achieve better than an $\Omega(\log n / \log \log n)$ approximation to social welfare. We show that this last result is almost tight by giving a truthful, online-envy-free mechanism which achieves a $\log n$ approximation to social welfare. We leave open the problem of closing the gap between our upper and lower bounds for non-envy-free randomized mechanisms, which seems to require different techniques. All our lower bounds hold even for algorithms that are not computationally restricted, while our upper bounds follow from computationally efficient mechanisms.

Given the impossibility in the adversarial model, we then consider the stochastic supply setting in which supply is drawn from a distribution D known to the mechanism, and welfare guarantees are required to hold in expectation over D . We make the assumption (standard in mechanism design in other contexts) that D has a non-decreasing hazard rate². Our second main result is a positive one:

Theorem: *There exists a truthful mechanism that achieves a constant approximation to social welfare when supply is drawn from a known distribution with non-decreasing hazard rate.*

This mechanism is simple, deterministic, computationally efficient, and easy to implement, but its analysis is surprisingly subtle. We stress that the incentive properties of the mechanisms we give do not rely on any distributional information. In particular, truthful bidding is a dominant strategy for every set of bids, for every supply, and for any realization of the coin flips of the mechanism (truthful "in the universal sense", see [24, 9]), not only in expectation. Truthfulness in expectation over supply realization would require that all the bidders and the seller share the same beliefs on how the supply is distributed. This is unlikely either because bidders do not have the resources needed for estimating these priors, or, because they may have private information that creates heterogeneity in their beliefs (see, e.g., [2]).³

We also show that the non-decreasing hazard rate assumption is necessary: no deterministic mechanism can achieve a constant approximation (or, in particular, better than an $\Omega(\sqrt{\log n / \log \log n})$ approximation) to social welfare over arbitrary distributions. As mentioned, our mechanism is deterministic, and does not involve randomization techniques used in previous papers for obtaining truthful approximations (like random sampling, see [14, 9]).

Finally, we also consider the setting in which the bidders preferences may exhibit complementarities for multiple items (*increasing* marginal utilities). We study the the extreme case of *knapsack valuations* (or single-minded bidders) and show strong lower bounds (even in the stochastic supply setting) on the competitive ratio that any algorithm can achieve, even without incentive constraints. We provide an algorithm with an exactly matching competitive ratio to prove that our lower bound is tight.

²A cumulative distribution F with density f has *non-decreasing hazard rate* (sometimes called *monotone hazard rate*) if $\frac{f(x)}{1-F(x)}$ is non-decreasing with x .

³We note that in the stochastic setting, we can achieve optimal welfare using expected VCG prices if we were to require only truthfulness *in expectation* over the supply ℓ . However, this seems to be a weak solution concept, since bidders may be motivated to misrepresent their valuations if their understanding of the supply distribution D differs from the mechanism's, or if they are not risk-neutral. In this paper, we show that positive results can be achieved even with this stronger solution concept.

1.2 Related Work

The works most related to ours are Mahdian and Saberi [20], Cole, Dobzinski, and Fleischer [7] and Lavi and Nisan [19]. Mahdian and Saberi [20] is the only other work that we are aware of to study mechanisms in which the supply is unknown and arrives online. They study the sale of multiple types of goods to bidders who desire only a single item, and wish to design mechanisms to maximize revenue. They consider only the adversarial supply setting, and allow extracting all payments when the entire supply has been exhausted. In this model, they give a truthful mechanism that is constant competitive with respect to the optimal auction that is restricted to selling all items at a single price, and show a lower bound of $(e+1)/e$. Their mechanism is randomized, and is based on random-sampling techniques to achieve truthfulness.

Cole, Dobzinski, and Fleischer [7] introduce the concept of *prompt mechanisms*, which impose the natural condition that bidders learn their payment immediately upon winning an item. They observe that mechanisms which are not prompt are often unusable, because, e.g., they tie up bidders to the auction for too long, they make debt collection difficult, and they require a high level of trust in the auctioneer. They study prompt mechanisms for a problem in which the supply of m expiring items is fixed and known to the mechanism, but the bidders arrive and depart online. They wish to maximize social welfare, and give a truthful $\log m$ competitive mechanism, and show a lower bound of 2 even for randomized mechanisms. Similar models of online auctions with expiring goods were studied earlier by Lavi and Nisan [19] and by Hajiaghayi et al. [13]. These models relate to ours since the allocation decisions for items with expiration date (airline tickets, for instance) must be made online. In these papers, however, there is no uncertainty on the supply and bidders arrive and depart over time. More on online auctions, which were first discussed by Lavi and Nisan [18], can be found in the survey [25].

A recent line of papers studies online mechanism design in a Bayesian setting ([5, 3, 4]), where welfare-maximizing, and even budget balanced, generalizations of VCG mechanisms are presented for online settings. Our paper does not assume a Bayesian preference model and, as our lower bounds show, socially-efficient outcomes cannot be truthfully implemented. In the economics literature, stochastic supply has not been studied in many papers. Most of this work (see, for example, [16, 23]) studied a Bayesian model, and focused on the characterization of equilibrium prices. Uncertain supply models can be viewed as more complicated versions of the classic sequential auctions model, which is technically hard to analyze even without uncertainty on the supply (see, e.g., [21, 26]).

While our paper focuses on auctions for identical goods with bidders that are interested in a single item, we briefly discuss a more general domain in which single minded bidders are interested in multiple items in Section 6. Knapsack auctions (or auctions for single-minded bidders) were studied by [1, 8] for static settings with known supply.

We proceed as follows. After presenting our formal model in the next section, we present our main results in Sections 3 (adversarial supply) and 4 (stochastic supply). We then discuss online-envy-free mechanisms in Section 5 and strengthen our lower bounds, and consider Knapsack valuations in Section 6.

2 Model and Definitions

We consider a set of n bidders $\{1, \dots, n\}$, each desires a single item from a set of identical items (except in Section 6 in which we expand our model to agents interested in multiple items.) Each bidder has a non-negative valuation v_i for an item. A *mechanism* \mathcal{M} is a (possibly randomized) allocation rule paired with a payment rule. Bidders report their valuations to the mechanism before any item arrives, and the mechanism assigns items as they arrive to bidders, and simultaneously charges each bidder i some price p_i . When ℓ items arrive and bidders have submitted bids v'_1, \dots, v'_n , we denote the outcome of the mechanism by $\mathcal{M}_\ell((v'_1, \dots, v'_n), r)$ where r is a random bitstring which may be used by randomized mechanisms. We note that the mechanism is unaware of ℓ , as it only encounters the items one at a time as they arrive. We will leave out the r when it is clear from context. We adopt standard notation and write v'_{-i} to denote the set of valuations reported by all bidders other than bidder i . A bidder i who receives an item obtains utility $u_i(v_i; \mathcal{M}_\ell(v'_1, \dots, v'_n)) = v_i - p_i$. Bidders who do not receive an item obtain utility 0. Bidders wish to maximize their own utility, and may misrepresent their valuations to the mechanism in order to do so.

We require that our mechanisms be *truthful*: that bidders should be incentivized to report their true valuations, regardless of the bids of others or the realizations of the supply. Following the literature (e.g. Goldberg et al. [11], Guruswami et al. [12]) we define a randomized truthful mechanism to be a probability distribution over deterministic truthful mechanisms.

Definition 2.1. A mechanism \mathcal{M} is (*ex-post*) *truthful* if for every bidder i with value v_i , for every set of bids v'_{-i} , for every alternative bid v'_i and for every r and ℓ : $u_i(v_i; \mathcal{M}_\ell((v_i, v'_{-i}), r)) \geq u_i(v'_i; \mathcal{M}_\ell((v'_i, v'_{-i}), r))$

We will assume that bidders submit their true valuations to truthful mechanisms, since it is a dominant strategy for them to do so.

Without loss of generality, we imagine that v_1, \dots, v_n are written in non-increasing order. The social welfare achieved by a mechanism is the sum of the values of the bidders to whom it has allocated items, which we denote by $W(\mathcal{M}_\ell((v_1, \dots, v_n), r))$. When ℓ items arrive, we will denote the optimal social welfare by $\mathbf{OPT}_\ell = \sum_{i=1}^{\ell} v_i$. When ℓ is drawn from a distribution D over the support (w.l.o.g.) $\{1, \dots, n\}$, we define $\mathbf{OPT} = \mathbb{E}_\ell[\mathbf{OPT}_\ell] = \sum_{i=1}^n \mathbf{OPT}_i \cdot \Pr[l = i]$.

We will be concerned with approximation guarantees to social welfare in both the *adversarial supply* setting and the *stochastic supply* setting.

Definition 2.2. A mechanism \mathcal{M} achieves an α -approximation to social welfare in the *adversarial supply* setting if for every supply ℓ : $\frac{\mathbf{OPT}_\ell}{\mathbb{E}_r[W(\mathcal{M}_\ell((v_1, \dots, v_n), r))]} \leq \alpha$

When ℓ is drawn from a distribution D , a mechanism \mathcal{M} achieves an α -approximation to social welfare in the *stochastic supply* setting if: $\frac{\mathbb{E}_\ell[\mathbf{OPT}_\ell]}{\mathbb{E}_{\ell, r}[W(\mathcal{M}_\ell((v_1, \dots, v_n), r))]} \leq \alpha$

In the stochastic setting, we will assume unless otherwise specified that D satisfies the *non-decreasing hazard rate* condition:

Definition 2.3. The hazard rate of a distribution D at i is: $h_i(D) = \frac{\Pr[\ell=i]}{\Pr[\ell \geq i]}$. We write simply h_i when the distribution is clear from context.

D satisfies the *non-decreasing hazard rate* condition if $h_i(D)$ is a non-decreasing sequence in i .

The non-decreasing hazard rate condition is standard in mechanism design (see, for example, [22, 17] and recent computer-science work [6, 15]), and is satisfied by many natural distributions, including the exponential, uniform, and binomial distributions.

One might also consider an intermediate model in which supply is drawn from a distribution satisfying the non-decreasing hazard rate condition, but the distribution is unknown to the mechanism. However, we note that since point distributions satisfy the hazard rate condition, adversarial supply is a special case of this model, and so our lower bounds apply.

3 Adversarial Supply

In this section we consider the adversarial model in which we do not have a distribution over supply and we require a good approximation to social welfare for any number of items that arrive. We first show that deterministic truthful mechanisms cannot achieve any approximation better than the trivial n -approximation. We then consider randomized mechanisms, and give a lower bound of $\Omega(\log \log n)$, proving in particular that no constant approximation is possible.

3.1 Deterministic Mechanisms

We begin by proving that deterministic mechanisms can only achieve a trivial approximation. We present a sketch of the proof and defer the details to Appendix A.1. First, we characterize deterministic truthful mechanisms by two useful observations:

Lemma 3.1. *For every truthful mechanism and for any realization of items, the price p_b that bidder b is charged upon winning (any) item is independent of his bid.*

Lemma 3.2. *For every truthful mechanism and for any realization of items, if bidder b wins an item, which item bidder b wins is independent of his bid whenever $p_b < v_b$.*

Theorem 3.3. *No deterministic truthful mechanism can achieve better than an n approximation to social welfare.*

Proof. (Sketch) We show that if the mechanism achieves any finite approximation to social welfare, every bidder has a bid such that he is allocated the first item. Applying lemmas 3.1 and 3.2, we conclude that any deterministic truthful mechanism that achieves a finite approximation to social welfare can only sell a single item, which implies that it cannot achieve better than an n approximation when all bidders have the same value for an item. See the appendix for further details. \square

3.2 Randomized Mechanisms

3.2.1 An $\Omega(\log \log n)$ lower bound

We next present our first main result, a lower bound for randomized truthful mechanisms.

Theorem 3.4. *No truthful randomized mechanism can achieve an $o(\log \log n)$ approximation to social welfare when faced with adversarial supply.*

Proof. A truthful randomized mechanism is simply a probability distribution over deterministic truthful mechanisms. To prove our randomized lower bound, we will exhibit a distribution over bidder values such that no deterministic truthful mechanism achieves a good approximation to welfare in expectation over this random instance. By Yao’s min-max principle, this is sufficient to prove a lower bound on randomized mechanisms.

We define a distribution V with support over values $1/2^i$ for $0 \leq i \leq \log n - 1$. For each realization $v \in V$, we let: $\Pr[v = 1/2^i] = 2^i/(n - 1)$. Therefore, we have $\Pr[v \geq 1/2^i] = (2^{i+1} - 1)/(n - 1)$ and $E[v|v \geq 1/2^i] = (i + 1)/(2^{i+1} - 1)$.

Lemma 3.5. *Consider a set of n valuations drawn from V and let OPT_k denote the sum of the k highest valuations from the set. Then: $E[\text{OPT}_k] \geq H_{k+1} - 1$ where H_{k+1} denotes the $k + 1$ st harmonic number. In particular, $E[\text{OPT}_k] > (\log k)/2$.*

Proof. We defer this proof to Appendix A.2. \square

By Lemma 3.1 and Lemma 3.2, we may characterize deterministic truthful mechanisms as follows: The mechanism assigns to each bidder b a bin i_b and a threshold t_b . i_b and t_b are independent of b ’s bid v_b , but are assigned such that at most one bidder in each bin can have a bid above his threshold.⁴ If $v_b > t_b$, b wins item i (if it arrives) at price t_b . Equivalently, we may imagine the mechanism operating by ordering bidders in some permutation π such that for all i , every bidder in bucket i is ordered before every bidder in bucket $j > i$. When the first item arrives, the mechanism offers it to each bidder at their threshold price, in order of π until some bidder b accepts. We continue in this manner, offering the next item to bidders starting at $b + 1$ until one accepts, etc.

We construct a distribution over instances by drawing each bidder’s valuation independently from the distribution V described above. Since bidder’s thresholds and buckets are independent of their own bids, each value encountered by the mechanism when making offers in order of π is distributed randomly according to V (note that although the values are distributed randomly, they need not be independent of each other). We may assume without loss of generality that each threshold $t_b = 1/2^{c_b}$ for some $c_b \in 0, \dots, \log n - 1$.

When all n items arrive, the expected welfare achieved by a mechanism is: $\sum_{b=1}^n \Pr[v_b \geq \frac{1}{2^{c_b}}] \cdot E[v_b|v_b \geq \frac{1}{2^{c_b}}] = \frac{1}{n-1} \sum_{b=1}^n (c_b + 1)$. Let N_b denote the number of items sold by a mechanism after making offers to b bidders. Then we have more generally, when k items arrive, the expected welfare achieved by a mechanism is: $\sum_{b=1}^n \Pr[v_b \geq \frac{1}{2^{c_b}}] \cdot E[v_b|v_b \geq \frac{1}{2^{c_b}}] \cdot \Pr[N_{b-1} < k] = \frac{1}{n-1} \sum_{b=1}^n (c_b + 1) \Pr[N_{b-1} < k]$. If

⁴An example of such a function is for each bidder’s threshold to be the highest bid of any other bidder in his bin. This results in exactly one bidder (the highest) having a bid above his threshold, while maintaining the property that each bidders threshold is independent of his bid.

our mechanism achieves an α approximation to social welfare, we therefore have the following n constraints on the values of c_b chosen by the mechanism. For all $1 \leq k \leq n$:

$$\sum_{b=1}^n (c_b + 1) \Pr[N_{b-1} < k] \geq \frac{(n-1)\mathbf{OPT}_k}{\alpha} \geq \frac{(n-1) \log k}{2\alpha} \quad (1)$$

where the last inequality follows from Lemma 3.5. After offering the item to b bidders, the expected number of sales is $E[N_b] = 1/(n-1) \cdot \sum_{i=1}^b (2^{c_i+1} - 1)$.

By a Chernoff bound: $\Pr[N_{b-1} < k] \leq \exp(-(\frac{E[N_{b-1}]}{2} - k + 1)) \leq \exp(-\frac{\sum_{i=1}^{b-1} 2^{c_i}}{n-1} + k)$. Let b_k be the first index such that $\sum_{i=1}^{b_k} 2^{c_i} \geq (n-1) \cdot k$. Then by plugging our bound into constraint 1, we have for all k :

$$\sum_{i=1}^{b_k} (c_i + 1) + \sum_{i=b_k+1}^n \frac{(c_i + 1)}{\exp(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1})} \geq \frac{(n-1) \log k}{2\alpha}$$

Lemma 3.6. For $c_i \in [0, \log n - 1]$:

$$\sum_{i=b_k+1}^n \frac{(c_i + 1)}{\exp(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1})} < 2.5 \cdot n$$

Proof. We defer the proof of this technical lemma to Appendix A.2. \square

So, for all k , there must exist an integer b_k such that simultaneously the two equations hold: $\sum_{i=1}^{b_k} c_i \geq \frac{(n-1) \log k}{2\alpha} - (2.5 \cdot n + b_k)$, and $\sum_{i=1}^{b_k-1} 2^{c_i} < (n-1) \cdot k$. In particular, if $k \geq 2^{15\alpha}$ and $n \geq 30$, then $\frac{n \log k}{4\alpha} \leq \frac{(n-1) \log k}{2\alpha} - 3.5n$. Therefore, there must exist integers b_k to satisfy the equations:

$$\sum_{i=1}^{b_k} c_i \geq \frac{n \log k}{4\alpha} \quad (2)$$

$$\sum_{i=1}^{b_k-1} 2^{c_i} < n \cdot k \quad (3)$$

We will consider the smallest such set of b_k : For all k , we will have that $\sum_{i=1}^{b_k} c_i \geq \frac{n \log k}{4\alpha}$, but $\sum_{i=1}^{b_k-1} c_i < \frac{n \log k}{4\alpha}$. Note that if we reduce a larger b_k in this manner, inequality 3 continues to hold, and so this is without loss of generality.

We let $k = 2^{15\alpha}$ and consider the sequence of integers $k, 2k, 4k, \dots, 2^t k$ such that $n \geq 2^t k > n/2$. For $j \geq 1$ we write $\Delta_k^j = (b_{2^j k} - b_{2^{j-1} k})$, and $\Delta_k^0 = b_k$. We note that from inequality 2 and our assumption on the b_k , we have: $\sum_{i=b_{2^{j-1} k}}^{b_{2^j k}} c_i \geq \frac{n(\log k + j)}{4\alpha} - \sum_{i=1}^{b_{(k \cdot 2^{j-1})-1}} c_i \geq \frac{n}{4\alpha}$.

Exponentiating both sides and applying the AM-GM inequality we have:

$$2^{n/(4\alpha \Delta_k^j)} \leq \left(\prod_{i=b_{2^{j-1} k}}^{b_{2^j k}} 2^{c_i} \right)^{1/\Delta_k^j} \leq \frac{\sum_{i=b_{2^{j-1} k}}^{b_{2^j k}} 2^{c_i}}{\Delta_k^j} \leq \frac{n(2^j k + 1)}{\Delta_k^j}$$

where the last inequality follows from inequality 3. This gives us: $\Delta_k^j \geq \frac{n}{4\alpha(\log n + \log(2^{j+1} k) - \log \Delta_k^j)}$. We can expand the above recursive bound to isolate Δ_k^j and find $\Delta_k^j = \Omega(n/(\alpha(j + \alpha)))$.

We recall that $n > b_{2^t k} = \sum_{i=0}^t \Delta_k^i$. Using the above bound, we see that n is at least $\sum_{i=0}^t \Omega(n/(\alpha(i + \alpha))) = \Omega(\frac{n \log(t/\alpha)}{\alpha})$. Therefore, we have $\alpha \geq \Theta(\log(t/\alpha))$ and so $\alpha \geq \Theta(\log t)$. We recall that $k = 2^{15\alpha}$ and $2^t k = 2^{15\alpha+t} \leq n$. t is therefore constrained such that: $\log n \geq 15\alpha + t \geq \Theta(t)$. And so we may take t to be as large as $\Theta(\log n)$, giving us a lower bound of $\alpha \geq \Theta(\log \log n)$. \square

3.2.2 A truthful $\log n$ -approximation mechanism

Here we show a simple randomized mechanism that achieves a $\log n$ approximation to social welfare. In Section 5 we show that this is nearly optimal for the natural class of "online envy-free" mechanisms.

Let **RandomGuess** be the mechanism that selects a supply $g \in \{2, 4, 8, \dots, 2^i, \dots, n\}$ uniformly at random, and considers only the highest g bidders according to permutation order. When an item arrives the mechanism sells it to the first of the remaining such bidders and charges him v_{g+1} .⁵

Proposition 3.7. *RandomGuess is truthful and achieves a $\log n$ approximation to social welfare.*

Proof. We defer this proof to Section A.2 in the Appendix. □

We leave open the problem of closing the gap between the $\log n$ factor achieved by RandomGuess and the $\Omega(\log \log n)$ lower bound of Theorem 3.4. In section 5 we strengthen this lower bound to $\Omega(\log n / \log \log n)$ for the class of *online-envy-free* mechanisms, also defined in section 5. We conjecture that RandomGuess is optimal.

4 Stochastic Supply

Given the strong lower bounds we have shown in the adversarial setting, we now consider the stochastic setting in which supply is drawn from some distribution D known to the mechanism. In this section, we give our second main result, a deterministic truthful mechanism that achieves an $O(1)$ -approximation to social welfare for any distribution with non-decreasing hazard rate. At the end of this section we show that the monotone hazard rate condition is actually necessary to achieve constant approximation.

We consider the following mechanism that takes as input a distribution D . The mechanism is deterministic, so all probabilities are over the distribution D . We note that the mechanism decides on a maximal number of items it is going to sell *without looking at the bids*. Although it seems somewhat surprising it still achieves good approximation when the non-decreasing hazard rate condition holds.

HazardGuess(D):

1. Fix an arbitrary permutation π on the bidders.
2. Solicit bids, and denote them v_1, \dots, v_n in non-increasing order.
3. Let s^* be the smallest integer such that $s^* \geq \frac{\Pr[\ell \geq s^*]}{\Pr[\ell = s^*]}$. If $s^* > 3$ let $g = s^*$. Otherwise let $g = 1$.^a
4. Consider only the highest g bidders ordered according to π .
When an item arrives sell it to the first of the remaining such bidders and charge him v_{g+1} (or 0 if $g = n$).

^aAlternatively, we can pick $g = s^*$ always, but then we must pick a random permutation in step 1 of HazardGuess. We choose to present a deterministic mechanism.

Theorem 4.1. *HazardGuess(D) is truthful, and achieves a $16\frac{7}{8}$ -approximation to social welfare in expectation over D , for any distribution D such that the hazard rate $h_i(D)$ is non-decreasing.*

Truthfulness is immediate: Every bidder with bid higher than v_{g+1} faces a single take-it-or-leave-it offer at the same price (v_{g+1}). The offer and the order in which they receive the offer is independent of their own bids. To prove the approximation guarantee, we will need a series of lemmas.

The following lemmas, 4.2, 4.4 and 4.5 will show that for any distribution with non-decreasing hazard rate, $\max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i] \geq \mathbf{OPT}/5$. To complete the proof, we will then prove that HazardGuess achieves welfare at least $(8/27) \cdot \max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i]$, and thus achieves a $16\frac{7}{8}$ approximation to \mathbf{OPT} .

⁵The authors thank Andrew Goldberg for suggesting this mechanism, which is a significant simplification of our original mechanism.

Lemma 4.2. Let α be the smallest value such that for any set of bids, $\mathbf{OPT}/(\max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i]) \leq \alpha$. Then for each integer $0 \leq s \leq n-1$ we have the following bound on α in terms of D , which we denote $\text{Bound}(s)$:

$$\alpha \leq \sum_{i=1}^s \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} + \frac{\sum_{i=s+1}^n \Pr[\ell = i] \cdot i}{(s+1) \cdot \Pr[\ell \geq s+1]}$$

Proof. Suppose $\alpha > \beta$. That is, there exists a set of bids such that for all i we have $\mathbf{OPT}_i \cdot \Pr[\ell \geq i] < \mathbf{OPT}/\beta$, or equivalently:

$$\mathbf{OPT}_i < \frac{\mathbf{OPT}}{\beta \cdot \Pr[\ell \geq i]} \quad (4)$$

Recall that by definition, we have $\mathbf{OPT} = \sum_{i=1}^n \mathbf{OPT}_i \cdot \Pr[\ell = i]$. Observe that for all $1 \leq i \leq n-1$: $\mathbf{OPT}_{i+1} \leq \frac{i+1}{i} \mathbf{OPT}_i$ since v_1, \dots, v_n is a non-increasing sequence. By repeated application of this observation, we get the following n upper-bounds on \mathbf{OPT} indexed by $0 \leq s \leq n-1$:

$$\mathbf{OPT} \leq \sum_{i=1}^s \mathbf{OPT}_i \cdot \Pr[\ell = i] + \mathbf{OPT}_{s+1} \cdot \left(\sum_{i=s+1}^n \frac{i}{s+1} \Pr[\ell = i] \right)$$

Applying inequality 4 and multiplying both sides by β/\mathbf{OPT} we obtain:

$$\beta < \left(\sum_{i=1}^s \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} + \frac{\sum_{i=s+1}^n \Pr[\ell = i] \cdot i}{(s+1) \cdot \Pr[\ell \geq s+1]} \right).$$

If α is the optimal approximation factor, there is some input such that for every $\epsilon > 0$, $\max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i]$ achieves an α approximation but does not achieve a $\beta = \alpha - \epsilon$ approximation, and the above bound on β holds. Since $\alpha = \beta + \epsilon$, letting ϵ tend to zero, we obtain the lemma. \square

Remark 4.3. We must now show that for every distribution D , there exists an s such that $\text{Bound}(s)$ gives $\alpha \leq 5$. Note that the order of quantifiers is important! It is not the case that there exists an s such that for every distribution, $\text{Bound}(s)$ gives $\alpha \leq O(1)$.

Lemma 4.4. For any $s \geq 1$ and $h_i \in [1/s, 1]$: $\sum_{i=s+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1 - h_j) \right) \leq 3s + 1$.

Proof. We defer the proof of this technical lemma to Appendix A.3. \square

Lemma 4.5. For any set of bids, and for any distribution D with non-decreasing hazard rate, $\frac{\mathbf{OPT}}{\max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i]} \leq 5$.

Proof. Given a distribution D , we wish to find the value of s such that $\text{Bound}(s)$ gives the sharpest bound on α (the approximation factor from lemma 4.2). We choose $s^* \leq n$ to be the smallest integer such that $s^* \geq \Pr[\ell \geq s^*]/\Pr[\ell = s^*]$. If no such s^* exists, we choose $s^* = n$. We now show that $\text{Bound}(s^*)$ gives $\alpha \leq 5$. We bound the two terms of $\text{Bound}(s^*)$ separately. Consider the first term:

$$\sum_{i=1}^{s^*} \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} \leq (s^* - 1) \cdot \frac{\Pr[\ell = s^* - 1]}{\Pr[\ell \geq s^* - 1]} + \frac{\Pr[\ell = s^*]}{\Pr[\ell \geq s^*]} \leq 1 + \frac{\Pr[\ell = s^*]}{\Pr[\ell \geq s^*]} \leq 2$$

since the hazard rate is non-decreasing and by definition of s . We now consider the second term: $\frac{\sum_{i=s^*+1}^n \Pr[\ell = i] \cdot i}{(s^*+1) \cdot \Pr[\ell \geq s^*+1]}$. Since D has a non-decreasing hazard rate, we know that for all $i \geq s^*$, $h_i \equiv \Pr[\ell = i]/\Pr[\ell \geq i] \geq 1/s^*$. Therefore, we have:

$$\begin{aligned} \sum_{i=s^*+1}^n \Pr[\ell = i] \cdot i &= \sum_{i=s^*+1}^n \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} \cdot \Pr[\ell \geq i] \cdot i \\ &= \sum_{i=s^*+1}^n \left(i \cdot h_i \cdot \Pr[\ell \geq s^* + 1] \cdot \prod_{j=s^*+1}^{i-1} (1 - h_j) \right) \\ &\leq \Pr[\ell \geq s^* + 1] (3s^* + 1) \end{aligned}$$

where the inequality follows from Lemma 4.4. Therefore, finally we have for all s^* :

$$\frac{\sum_{i=s^*+1}^n \Pr[\ell = i] \cdot i}{(s^* + 1) \cdot \Pr[\ell \geq s^* + 1]} \leq \frac{\Pr[\ell \geq s^* + 1](3s^* + 1)}{(s^* + 1) \cdot \Pr[\ell \geq s^* + 1]} \leq 3$$

Combining these two bounds, we finally get that $\text{Bound}(s^*)$ gives $\alpha \leq 5$. \square

Now we are ready to complete the proof of our theorem:

Proof of Theorem 4.1. We show that HazardGuess achieves welfare at least $(8/27) \cdot (\max_i \mathbf{OPT}_i \cdot \Pr[\ell \geq i])$. Together with lemma 4.5, this proves that HazardGuess achieves at least a $16\frac{7}{8}$ approximation to social welfare.

Let s^* be the smallest integer such that $s^* \geq \frac{\Pr[\ell \geq s^*]}{\Pr[\ell = s^*]}$. Whenever $s^* > 3$, $\text{HazardGuess}(D)$ achieves welfare at least $\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*]$. When $s^* \leq 3$, $\text{HazardGuess}(D)$ achieves welfare at least $\mathbf{OPT}_{s^*}/3$ (since it sells a single item to the highest bidder, and $\mathbf{OPT}_1 \geq \mathbf{OPT}_3/3$). First consider the case in which $i > s^* \geq 1$. In this case, we know $\Pr[\ell \geq i] \leq \Pr[\ell \geq s^*] \cdot (1 - \frac{1}{s^*})^{i-s^*}$, since the hazard rate h_i is non-decreasing, and $h_{s^*} \geq 1/s^*$. Therefore, we have:

$$\begin{aligned} \mathbf{OPT}_i \cdot \Pr[\ell \geq i] &\leq \frac{i}{s^*} \cdot \mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq i] \\ &\leq \frac{i}{s^*} \cdot \mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*] \cdot (1 - \frac{1}{s^*})^{i-s^*} \\ &\leq (\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*]) \cdot \left(\frac{i}{s^*} \cdot \frac{1}{e^{i/s^*-1}} \right) \\ &\leq (\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*]) \end{aligned}$$

Therefore, in this case, $\text{HazardGuess}(D)$ achieves welfare at least $\mathbf{OPT}_i \cdot \Pr[\ell \geq i]/3$. Now consider the case in which $1 \leq i < s^*$: By definition of s^* : $\frac{\Pr[\ell \geq s^*-1]}{\Pr[\ell = s^*-1]} > s^* - 1$. Alternatively, we may write the hazard rate at $s^* - 1$: $h_{s^*-1} < 1/(s^* - 1)$. Since the hazard rate is non-decreasing, we have that for all $i \leq s^* - 1$, $h_i < 1/(s^* - 1)$. Therefore we have:

$$\Pr[\ell \geq s^*] = \prod_{i=1}^{s^*-1} (1 - h_i) > \prod_{i=1}^{s^*-1} (1 - \frac{1}{s^*-1}) = \left(\frac{s^*-2}{s^*-1} \right)^{s^*-1}$$

If $s^* \geq 4$, then this gives $\Pr[\ell \geq s^*] \geq 8/27$. Therefore:

$$\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*] \geq \mathbf{OPT}_i \cdot \Pr[\ell \geq 4] \geq \frac{8}{27} \mathbf{OPT}_i$$

which is a bound on the performance of $\text{HazardGuess}(D)$, since $s^* > 3$. Finally we consider the special case of $s^* \in \{2, 3\}$. If $s^* = 2$, then $i \in \{1, 2\}$ achieves welfare $\mathbf{OPT}_i/2 \cdot \Pr[\ell \geq i]$ since HazardGuess sells one item. Similarly, if $s^* = 3$ HazardGuess achieves welfare at least $\mathbf{OPT}_i/3 \cdot \Pr[\ell \geq i]$. This concludes the proof. \square

We note that our analysis is worst-case, and that this mechanism can be shown to achieve a better constant approximation for specific distributions of interest. For example:

Theorem 4.6. *$\text{HazardGuess}(D)$ achieves a $\frac{3}{5}$ -approximation to social welfare in expectation over D when D is the uniform distribution over $\{1, \dots, n\}$. Moreover, there are values for which $\text{HazardGuess}(D)$ cannot get better than a $\frac{3}{4}$ -approximation when D is the uniform distribution.*

The proof is deferred to the appendix.

4.1 The Necessity of the Monotone Hazard Rate Condition

We next show that the monotone hazard rate condition is necessary: for arbitrary distributions *no* deterministic mechanism can achieve constant approximation to social welfare.

Theorem 4.7. *No deterministic truthful mechanism can achieve an $o(\sqrt{\log n / \log \log n})$ approximation to social welfare when faced with arbitrary stochastic supply (without the non-decreasing hazard rate condition).*

We prove this theorem in Appendix A.3. The proof proceeds in two stages. First, we consider a class of truthful mechanisms that fix -independently of the bids- an ordering π on the bidders, and a supply g . Such a mechanism sells the first g items that arrive at the $(g + 1)$ -st highest price to the g highest bidders, ordered according to π . We note that HazardGuess is such a mechanism, and all such mechanisms satisfy a notion of envy-freeness which we define in the next section. We show that such mechanisms cannot achieve an $o(\log n / \log \log n)$ approximation to social welfare when faced with arbitrary stochastic supply. We then complete the proof by showing that we can restrict our attention to such mechanisms almost without loss of generality: for *any* deterministic mechanism, there exists a mechanism that chooses its supply independently of the bids that loses only a quadratic factor in its approximation to social welfare.

5 Envy-Free Mechanisms

All our mechanisms satisfy a notion of fairness which is our adaptation of envy-freeness to the online setting. An offline mechanism is envy-free if no agent prefers another agent's allocation and payment to his own (see, for example, [11, 12]). In the case of unit demand bidders and identical goods this means that there is a price p such that any winner pays the same price p and has value at least p , and any loser has value at most p . This is clearly not possible to achieve for online supply, except by trivial mechanisms (for example, the mechanism that only sells a single item to the highest bidder at the second highest price). Informally, in an online envy-free mechanism, the only source of envy is a shortage of supply, not price discrimination on the part of the mechanism.

Definition 5.1. A deterministic mechanism is *online-envy-free* if it is envy-free (in the offline sense) when the supply is enough to satisfy the demand of all of the bidders (that is, when $l = n$). A randomized mechanism is *online-envy free* if it is a distribution over deterministic online-envy-free mechanisms.

Note that this definition ensures that all sold items are sold for the same price, even when the supply is smaller than n . Also note that both our mechanisms RandomGuess and HazardGuess are online-envy-free.

In Theorem 3.4 we showed that no truthful randomized mechanism can achieve an $o(\log \log n)$ approximation to social welfare when faced with adversarial supply. Here, we present an improved lower bound for truthful online-envy free mechanisms.

Proposition 5.2. *No truthful online-envy-free mechanism (even randomized) can achieve an $o(\log n / \log \log n)$ approximation to social welfare when faced with adversarial supply.*

We defer the proof to appendix Section A.4. Note that proposition 5.2 is nearly tight, since RandomGuess achieves a $\log n$ approximation factor.

6 Valuations with complementarities: Knapsack Valuations

So far we have discussed bidders that are interested in a single item out of a set of identical items. It is natural to consider the case of bidders with increasing-marginal utility valuations, corresponding to *complements* valuations. In the extreme case, we get *knapsack valuations*.

We say that a bidder i has a *knapsack valuation* if he has a value c_i and a desired quantity k_i : For all $k < k_i$, $v_i(k) = 0$, and for all $k \geq k_i$, $v_i(k) = c_i$. That is, bidder i desires at least k_i units of the good, is not satisfied with fewer, and has no value for more than k_i units.

Knapsack valuations can be seen as modeling advertising campaigns: a buyer wishes to build brand name recognition through banner-advertisements, and so has little value for a small number of advertisements; A campaign is worth c_i to the advertiser, but additional advertising saturation has little added benefit.

Unfortunately, the online nature of the problem makes knapsack valuations difficult to handle for any algorithm, even without truthfulness (and computational) constraints. Here, we present an algorithm in the stochastic setting, and show that its (poor) competitive ratio is optimal over the class of all (not necessarily truthful) algorithms. Without loss of generality, we can assume that D has finite support over $[1, m]$ for $m = \sum_{i=1}^n k_i$.

Our lower bound for Knapsack valuations shows that with online supply, no algorithm can guarantee a better approximation ratio than the cumulative hazard rate. This welfare guarantee is quite poor. For the uniform distribution, this gives $\alpha = \Theta(\log m)$. For the binomial distribution, $\alpha = \Theta(m)$. We also present a matching upper bound showing that our lower bound is tight. Both proofs are in Appendix A.5.

Proposition 6.1. *No algorithm can have better than a $\sum_{i=1}^m h_i$ approximation to optimal social welfare.*

Proposition 6.2. *For any distribution D with (arbitrary) hazard rate h_i there exists an algorithm that achieves at least a $\sum_{i=1}^m h_i$ approximation to optimal social welfare.*

References

- [1] G. Aggarwal and J. D. Hartline. Knapsack auctions. In *ACM-SIAM Symposium On Discrete Algorithm (SODA)*, pages 1083–1092, 2006.
- [2] Susan Athey and Jonathan Levin. Information and competition in u.s. forest service timber auctions. *Journal of Political Economy*, 109(2):375–417, 2001.
- [3] Susan Athey and Ilya Segal. An efficient dynamic mechanism, 2007. Working paper.
- [4] D. Bergemann and J. Valimaki. Efficient dynamic auctions, 2006. Cowles Foundation Discussion Paper 1584.
- [5] R. Cavallo, D. C. Parkes, and S. Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 55–62, 2006.
- [6] S. Chawla, J.D. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *ACM conference on Electronic Commerce (EC)*, pages 243–251, 2007.
- [7] R. Cole, S. Dobzinski, and L. Fleischer. Prompt Mechanisms for Online Auctions. *Symposium on Algorithmic Game Theory (SAGT): LNCS*, 4997:170, 2008.
- [8] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *ACM conference on Electronic Commerce (EC)*, pages 346–351, 2007.
- [9] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC*, 2006.
- [10] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [11] A.V. Goldberg and J.D. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 29–35. ACM New York, NY, USA, 2003.
- [12] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics Philadelphia, PA, USA, 2005.

- [13] M. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. Parkes. Online auctions with re-usable goods. In *ACM Conf. on Electronic Commerce (EC)*, pages 165–174, 2005.
- [14] Jason D. Hartline and Anna R. Karlin. In *N.Nisan, T.Roughgarden, E.Tardos, and V.Vazirani (eds.) Chapter 13, Profit Maximization in Mechanism Design*. Cambridge University Press., 2007.
- [15] J.D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *ACM Symposium on Theory Of Computing (STOC)*, pages 75–84, 2008.
- [16] Thomas D. Jeitschko. Equilibrium price paths in sequential auctions with stochastic supply. *Economics Letters*, 64(1):67–72, 1999.
- [17] V. Krishna. *Auction theory*. Academic press, 2002.
- [18] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science.*, 310(1):159–180, 2004.
- [19] R. Lavi and N. Nisan. Online ascending auctions for gradually expiring items. In *ACM-SIAM Symposium On Discrete Algorithms (SODA)*, pages 1146–1155, 2005.
- [20] M. Mahdian and A. Saberi. Multi-unit Auctions with Unknown Supply. *ACM conference on Electronic Commerce (EC)*, pages 243 – 249, 2006.
- [21] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.
- [22] R.B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [23] Tibor Neugebauer and Paul Pezanis-Christou. Equilibrium price paths in sequential auctions with stochastic supply. *Journal of Economic Behavior and Organization*, 63(1):55–72, 2007.
- [24] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior (GEB)*, 35(1-2):166–196, 2001.
- [25] David Parkes. In *N.Nisan, T.Roughgarden, E.Tardos, and V.Vazirani (eds.) Chapter 16, Online Mechanisms*. Cambridge University Press., 2007.
- [26] McAfee R. Preston and Vincent Daniel. The declining price anomaly. *Journal of Economic Theory (JET)*, 60(1):191–212, 1993.

A Proofs

A.1 Proof: Lower bound for deterministic mechanism with adversarial supply

In this section we prove Theorem 3.3.

Theorem A.1. *No deterministic truthful mechanism can achieve better than an n approximation to social welfare.*

The theorem will follow from three simple lemmas.

Lemma A.2. *For every truthful mechanism and for any realization of items, the price p_b that bidder b is charged upon winning (any) item is independent of his bid.*

Proof. This is a standard fact characterizing truthful auctions; If there is some realization of items for which bidder b has two distinct bids which result in bidder b winning an item, but at a different price, then in the case in which his valuation is equal to the bid that yields an item at the higher price, he will report falsely that his valuation is equal to the bid that yields an item at the lower price. \square

Lemma A.3. For every truthful mechanism and for any realization of items, if bidder b wins an item, which item bidder b wins is independent of his bid whenever $p_b < v_b$.

Proof. Suppose for some realization of items, and for some fixed set of bids of the other bidders, bidder b can change his bid to v_b or v'_b , and win one of two items, item i or item j , and that if he bids his true valuation v_b , he wins item $j > i$. Now consider a realization in which only i items arrive; If bidder b bids v_b , he wins no item and receives utility 0. If he bids v'_b , he wins item i at his (bid independent) price p_b , and achieves higher utility $v_b - p_b$. Therefore, the mechanism is not truthful. \square

Lemma A.4. For any deterministic mechanism that achieves an n -approximation to social welfare, every bidder has a bid such that they are allocated the first item.

Proof. Any bidder b can set his bid to more than n times the second highest bidder. If the mechanism does not allocate the first item to b , then if there are no further items, the mechanism has not achieved an n -approximation to social welfare. \square

Proof of Theorem. By Lemma A.4, any bidder can win the first item with an appropriately high bid. But by Lemma 3.2, any bidder such that $p_b < v_b$ who has a bid for which he can win the first item cannot win any other item with any bid. Therefore, for any set of bidders b_i such that for all b_i , $p_{b_i} \neq v_{b_i}$, then any deterministic truthful mechanism that achieves an n -approximation can only sell the first item. If all bidders have value $1 \leq v_{b_i} \leq 1 + \epsilon$, this achieves no better than an n -approximation when all items arrive. It remains to demonstrate such a set of bidders: Consider an arbitrary set of $n + 1$ distinct values between 1 and $1 + \epsilon$. For each bidder, choose a value from this set independently at random. Since each bidder's price p_{b_i} is independent of his bid, by Lemma 3.1, the probability that $v_{b_i} = p_{b_i}$ is at most $1/(n + 1)$, and by the union bound, the probability that *any* bidder's bid equals its price threshold is at most $n/(n + 1) \leq 1$. Therefore, there exists a set of bids sampled from this set with the desired property, which completes the proof. \square

A.2 Proofs of Lemmas from Section 3

Lemma 3.5: Consider a set of n valuations drawn from V and let \mathbf{OPT}_k denote the sum of the k highest valuations from the set. Then:

$$\mathbb{E}[\mathbf{OPT}_k] \geq H_{k+1} - 1.$$

where H_{k+1} denotes the $k + 1$ st harmonic number. In particular, $\mathbb{E}[\mathbf{OPT}_k] > (\log k)/2$.

Proof. Let $F(y)$ denote the cumulative distribution function of V . We note that $F(y)$ is a step function taking values $F(y) = (n - 1/y)/(n - 1)$ for all y of the form $y = 1/2^i$ for $i \in \{0, 1, \dots, \log n - 1\}$. We consider the inverse CDF function $F^{-1}(x) : [0, 1] \rightarrow \{1, 1/2, 1/4, \dots, 2/n\}$. It is simple to verify the following pointwise lower bound on $F^{-1}(x)$:

$$F^{-1}(x) \geq \frac{1}{n - x(n - 1)}$$

which follows from inverting the discrete CDF. We denote the quantity in this bound $A(x) = 1/(n - x(n - 1))$, and observe that $A(x)$ is convex in the range $[0, 1]$.

Let $v_{i,n}$ denote the i 'th largest value out of n draws from V , and let $X_{i,n}$ denote the i 'th largest value out of n draws from the uniform distribution over $[0, 1]$. We consider the following method of drawing a value v from V : we draw x uniformly from $[0, 1]$ and let $v = F^{-1}(x)$. Since F^{-1} is monotone, the i 'th largest draw from the uniform distribution corresponds to the i 'th largest draw from V : $v_i = F^{-1}(x_i)$.

Recall the expected value of the i 'th largest of n draws from the uniform distribution over $[0, 1]$: $\mathbb{E}[X_{i,n}] = 1 - i/(n + 1)$. This standard fact follows from a simple symmetry argument. We are now

ready to complete the proof of the lemma:

$$\begin{aligned}
\mathbb{E}[\mathbf{OPT}_k] &= \sum_{i=1}^k \mathbb{E}[v_{i,n}] \\
&= \sum_{i=1}^k \mathbb{E}[F^{-1}(X_{i,n})] \\
&\geq \sum_{i=1}^k \mathbb{E}[A(X_{i,n})] \\
&\geq \sum_{i=1}^k A(\mathbb{E}[X_{i,n}]) \\
&= \sum_{i=1}^k \frac{1}{1 + i(n-1)/(n+1)} \\
&\geq \sum_{i=1}^k \frac{1}{1+i} \\
&= H(k+1) - 1
\end{aligned}$$

where the second inequality is an application of Jensen's inequality, which follows since $A(x)$ is convex. \square

Lemma 3.6: For $c_i \in [0, \log n - 1]$:

$$\sum_{i=b_k+1}^n \frac{(c_i + 1)}{\exp\left(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1}\right)} < 2.5 \cdot n$$

Proof. Let $f(c_{b_k+1}, \dots, c_n) \equiv \sum_{i=b_k+1}^n \frac{(c_i+1)}{\exp\left(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1}\right)}$. We consider the partial derivative at the i 'th offer price:

$$\begin{aligned}
\frac{\partial}{\partial c_i} f(c_{b_k+1}, \dots, c_n) &= \frac{1}{e^{\sum_{j=b_k+1}^{i-1} 2^{c_j}/(n-1)}} - \left(\frac{2^{c_i} \ln 2}{(n-1)e^{2^{c_i}/(n-1)}} \right) \cdot \sum_{j=i+1}^n \frac{c_j + 1}{\exp\left(\sum_{\ell=b_k+1}^{j-1} 2^{c_\ell}/(n-1)\right)} \\
&\leq 1 - \frac{\ln 2}{n-1} \cdot \left(\sum_{j=i+1}^n \frac{c_j + 1}{\exp\left(\sum_{\ell=b_k+1}^{j-1} 2^{c_\ell}/(n-1)\right)} \right)
\end{aligned}$$

But this is negative unless

$$R_i \equiv \sum_{j=i+1}^n \frac{c_j + 1}{\exp\left(\sum_{\ell=b_k+1}^{j-1} 2^{c_\ell}/(n-1)\right)} \leq \frac{n-1}{\ln 2}$$

Fixing any maximal assignment to the c_i variables, let i' be the largest index for which the above condition on $R_{i'}$ fails to hold. We know that for all $i \leq i'$, $c_i = 0$, since the partial derivative at i is negative, and so if we could reduce c_i further this would contradict the fact that we selected a maximal assignment.

Therefore, we have:

$$\begin{aligned}
f(c_{b_k+1}, \dots, c_n) &= \sum_{i=b_k+1}^{i'} \frac{c_i + 1}{\exp\left(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1}\right)} + \sum_{i=i'+1}^n \frac{c_i + 1}{\exp\left(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1}\right)} \\
&\leq i' + \frac{1}{e^{2^{c_{i'}}/(n-1)}} \cdot R_{i'} \\
&\leq n + \frac{n-1}{\ln 2} \\
&< 2.5n
\end{aligned}$$

□

Proposition A.5 (Proposition 3.7). *RandomGuess is truthful, online-envy-free, and achieves a $\log n$ approximation to social welfare.*

Proof. Truthfulness and envy-freeness are immediate: every winning bidder faces a single take-it-or-leave-it offer independent of their bid, in an order independent of their bid. All items are sold at the same price, v_{g+1} . When n items arrive, all bidders with valuations higher than the offer price have been allocated items. We now prove the approximation guarantee.

Suppose that I items arrive, and $\mathbf{OPT}_I = \sum_{i=1}^I v_i$, the sum of the I highest bids. With probability $1/\log n$, $I < g \leq 2I$, and with probability $1/\log n$, $I/2 < g \leq I$. In the first case, RandomGuess allocates the I items to at least half of the top g bidders in random order, and so achieves welfare in expectation at least $\mathbf{OPT}_g/2 \geq \mathbf{OPT}_I/2$. In the second case, RandomGuess allocates at least half of the I items to all of the top g bidders, and achieves welfare $\mathbf{OPT}_g = \sum_{i=1}^g v_i$. Since $g > I/2$, $\mathbf{OPT}_g > \mathbf{OPT}_I/2$ because $\{v_i\}$ is a non-increasing sequence. Our mechanism therefore achieves in expectation welfare at least $(1/\log n)(\mathbf{OPT}_I/2 + \mathbf{OPT}_I/2) = \mathbf{OPT}_I/\log n$. □

A.3 Proofs from Section 4

Lemma 4.4: For any $s \geq 1$ and $h_i \in [1/s, 1]$:

$$\sum_{i=s+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1 - h_j) \right) \leq 3s + 1$$

Proof. Let $f(h_{s+1}, \dots, h_n) \equiv \sum_{i=s+1}^n (i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1 - h_j))$ and consider the partial derivative at h_k :

$$\begin{aligned}
\frac{\partial}{\partial h_k} f(h_{s+1}, \dots, h_n) &= k \cdot \prod_{j=s+1}^{k-1} (1 - h_j) - \sum_{i=k+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1, j \neq k}^{i-1} (1 - h_j) \right) \\
&\leq k \cdot \left(1 - \frac{1}{s}\right)^{k-s-1} - \sum_{i=k+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1, j \neq k}^{i-1} (1 - h_j) \right)
\end{aligned}$$

where the inequality follows from $h_i \geq 1/s$ for all i . But this is negative unless

$$R_k \equiv \sum_{i=k+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1, j \neq k}^{i-1} (1 - h_j) \right) \leq k \cdot \left(1 - \frac{1}{s}\right)^{k-s-1}$$

Fix some assignment to the h_i that maximizes $f(h_{s+1}, \dots, h_n)$ and let k' be the first index at which the above condition holds. Then for all $i < k'$, $h_i = 1/s$, since otherwise this would contradict the fact that the assignment maximizes f . Therefore, we have:

$$\begin{aligned}
\sum_{i=s+1}^n \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1-h_j) \right) &= \sum_{i=s+1}^{k'-1} \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1-h_j) \right) + \sum_{i=k'}^n \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1-h_j) \right) \\
&\leq \sum_{i=s+1}^{k'-1} \left(\frac{i}{s} \left(1 - \frac{1}{s}\right)^{i-s-1} \right) + \sum_{i=k'}^n \left(i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1-h_j) \right) \\
&= \sum_{i=s+1}^{k'-1} \left(\frac{i}{s} \left(1 - \frac{1}{s}\right)^{i-s-1} \right) + k' \cdot h_{k'} \cdot \prod_{j=s+1}^{k'-1} (1-h_j) + (1-h_{k'}) \cdot R_{k'} \\
&\leq \sum_{i=s+1}^{k'-1} \left(\frac{i}{s} \left(1 - \frac{1}{s}\right)^{i-s-1} \right) + h_{k'} \cdot (k' \cdot (1 - \frac{1}{s})^{k'-s-1}) + (1-h_{k'}) \cdot (k' \cdot (1 - \frac{1}{s})^{k'-s-1}) \\
&= \frac{1}{s} \sum_{i=s+1}^{k'-1} \left(i \left(1 - \frac{1}{s}\right)^{i-s-1} \right) + k' \cdot (1 - \frac{1}{s})^{k'-s-1} \\
&\leq \frac{1}{s} \sum_{i=s+1}^{\infty} \left(i \left(1 - \frac{1}{s}\right)^{i-s-1} \right) + (s+1) \\
&= 3s+1
\end{aligned}$$

where the second inequality follows from the fact that for all i , $h_i \geq 1/s$, the third inequality follows from the fact that $k \geq s+1$ and so $k' \cdot (1 - \frac{1}{s})^{k'-s-1}$ is decreasing in k' , and the last equality follows from the identity $\sum_{i=k}^{\infty} i \cdot r^{i-k} = (k+r-kr)/(r-1)^2$. \square

Theorem 4.6: HazardGuess(D) achieves a $\frac{3}{5}$ -approximation to social welfare in expectation over D when D is the uniform distribution over $\{1, \dots, n\}$. Moreover, there are values for which HazardGuess(D) cannot get better than a $\frac{3}{4}$ -approximation when D is the uniform distribution.

Proof. Consider the case that there are n agents and the supply is chosen uniformly at random from $\{1, n\}$ (we note that if the range starts from a number larger than 1 the problem becomes easier and the algorithm achieves better approximation.) We analyze the approximation achieved by picking the supply $k = n/2$ and selling at most k items,⁶ in a random order over the top k values. We prove that the algorithm achieves at least 60% of the optimum.

Assume the values are sorted $v_1 \geq v_2 \geq \dots \geq v_n$. Define $OPT_l = \sum_{i=1}^l v_i$. The expected welfare of the optimal algorithms is $OPT = 1/n \cdot \sum_{l=1}^n OPT_l$. Splitting the sum to two parts we get the following.

$$\begin{aligned}
OPT &= \frac{1}{n} \cdot \sum_{l=1}^{\frac{n}{2}} OPT_l + \frac{1}{n} \cdot \sum_{l=\frac{n}{2}+1}^n OPT_l \leq \frac{OPT_{\frac{n}{2}}}{2} + \frac{1}{n} \cdot \sum_{l=\frac{n}{2}+1}^n \frac{l}{n/2} OPT_{\frac{n}{2}} = OPT_{\frac{n}{2}} \left(\frac{1}{2} + \frac{2}{n^2} \sum_{l=\frac{n}{2}+1}^n l \right) = \\
&OPT_{\frac{n}{2}} \left(\frac{1}{2} + \frac{2}{n^2} \left(\frac{n(n+1)}{2} - \frac{\frac{n}{2}(\frac{n}{2}+1)}{2} \right) \right) = OPT_{\frac{n}{2}} \left(\frac{5}{4} + \frac{1}{2n} \right)
\end{aligned}$$

Our algorithm achieves expected welfare of

$$ALG = \frac{1}{n} \cdot \sum_{l=1}^{\frac{n}{2}} \frac{l}{n/2} OPT_{\frac{n}{2}} + \frac{1}{n} \cdot \sum_{l=\frac{n}{2}+1}^n OPT_{\frac{n}{2}} = OPT_{\frac{n}{2}} \left(\frac{2}{n^2} \sum_{l=1}^{\frac{n}{2}} l + \frac{1}{2} \right) =$$

⁶For simplicity we assume that n is even. Essentially the same argument will work for the case that n is odd.

$$OPT_{\frac{n}{2}} \left(\frac{2}{n^2} \frac{\frac{n}{2}(\frac{n}{2} + 1)}{2} + \frac{1}{2} \right) = OPT_{\frac{n}{2}} \left(\frac{3}{4} + \frac{1}{2n} \right) \geq \frac{OPT}{\frac{5}{4} + \frac{1}{2n}} \cdot \left(\frac{3}{4} + \frac{1}{2n} \right) \geq \frac{3}{5} OPT$$

Finally we observe that this algorithm gets at most 75% of the optimum. Consider the input with one value of 1 and all the rest of the values are 0. The optimal algorithm will always get welfare of 1. Our algorithm will get the 1 with probability

$$\sum_{l=1}^{n/2} \frac{1}{n} \cdot \frac{l}{n/2} + \frac{1}{2} = \frac{n+1}{4n} + \frac{1}{2} < \alpha$$

for any constant $\alpha > 3/4$ when n is large enough. \square

Theorem 4.7: No deterministic truthful mechanism can achieve an $o(\sqrt{\log n / \log \log n})$ approximation to social welfare when faced with arbitrary stochastic supply (without the non-decreasing hazard rate condition).

The theorem follows directly from two lemmas.

Definition A.6. A *bid-independent supply mechanism* chooses an ordering on the bidders π and a supply g independently of the bids. It then sells items as they arrive to the g highest bidders, ordered according to π , at the $g + 1$ st highest price.

Note that all mechanisms presented in this paper are bid-independent supply mechanisms.

Lemma A.7. No deterministic bid-independent supply mechanism can achieve an $o(\log n / \log \log n)$ approximation to social welfare when faced with arbitrary stochastic supply (without the non-decreasing hazard rate condition).

Proof. We give a distribution with a decreasing hazard rate such that no mechanism that determines a maximum supply g independent of the bids v_i can achieve an $o(\log n / \log \log n)$ approximation to social welfare.

We define D such that $\Pr[\ell = i] = 1/(i + i^2)$. Note that $\Pr[\ell \geq i] = 1/i$, and the hazard rate at i is decreasing: $h_i(D) = 1/(1 + i)$. Consider the welfare achieved by a bid-independent mechanism that chooses supply g . If at least g items arrive, it achieves welfare exactly \mathbf{OPT}_g . Otherwise, if $j < g$ items arrive, it achieves expected welfare at most $(j/g)\mathbf{OPT}_g$. Therefore, the welfare it achieves is at most:

$$\begin{aligned} \mathbf{OPT}_g \cdot \Pr[\ell \geq g] + \frac{1}{g} \cdot \sum_{j=1}^{g-1} j \cdot \Pr[\ell = j] &= \mathbf{OPT}_g \cdot \left(\frac{1}{g} + \frac{H_g - 1}{g} \right) \\ &= \Theta \left(\mathbf{OPT}_g \cdot \left(\frac{\log g}{g} \right) \right) \end{aligned}$$

We consider two possible sets of bidder values: In the Single Bidder case, we have $v_1 = 1$ and $v_j = 0$ for all $j > 1$. In the All Bidder case, we have $v_j = 1$ for all j . Note that in the Single Bidder case, we have $\mathbf{OPT} = 1$ and $\mathbf{OPT}_i = 1$ for all i . In the All Bidder case we have $\mathbf{OPT} = H_{n+1} - 1 = \Theta(\log n)$ and $\mathbf{OPT}_i = i$. Therefore, in the Single Bidder case, a mechanism that achieved an $o(\log n / \log \log n)$ approximation to social welfare would have $(\log g)/g = \omega(\log \log n / \log n)$, and in the All Bidder case would have $\log g = \omega(\log \log n)$. There is no $g \in [1, n]$ that satisfies both of these equations simultaneously. Since g is chosen independently of the bids, the two cases are indistinguishable, and any such mechanism much achieve an approximation ratio no better than $\Omega(\log n / \log \log n)$ in at least one of them. \square

Lemma A.8. For any distribution D and any deterministic truthful mechanism M that achieves an α approximation to social welfare over D , there is a truthful deterministic online-envy-free bid-independent supply mechanism M' that achieves an α^2 approximation to social welfare.

Proof. Let g_{\max} be the maximum number of items M sells when full supply is realized, where the maximum is taken over all possible bid profiles. Let M' be the mechanism that always sells the first g_{\max} items to the g_{\max} highest bidders in some predetermined order at the $g_{\max} + 1$ st highest price, and sells no further items. Note that M' is online-envy-free and has bid-independent sell sequence. First observe that $\mathbf{OPT}_{g_{\max}} \geq \mathbf{OPT}/\alpha$. This follows because by definition, M can never achieve welfare beyond $\mathbf{OPT}_{g_{\max}}$, but by assumption, M achieves an α approximation to the optimal social welfare. Next, observe that $\Pr_D[\ell \geq g_{\max}] \geq 1/\alpha$. To see this, consider some bid profile which causes M to produce a supply g_{\max} . Let b_i be the bidder who receives item g_{\max} , and consider raising his valuation v_i until it constitutes all but a negligible fraction of the total possible social welfare. By lemmas 3.1 and 3.2, raising b_i 's bid does not affect either the supply offered by the mechanism, or the order in which b_i receives an item: that is, it continues to be the case that b_i receives an item if and only if at least g_{\max} items arrive. However, since b_i now constitutes an arbitrarily large fraction of the total social welfare, and M is an α -approximation mechanism, it must be that $\Pr[\ell \geq g_{\max}] \geq 1/\alpha$.

Finally, we observe that our mechanism achieves welfare at least $\mathbf{OPT}_{g_{\max}} \cdot \Pr[\ell \geq g_{\max}] \geq \mathbf{OPT}/\alpha^2$, which completes the proof. \square

A.4 Proofs from Section 5

Proposition A.9 (Proposition 5.2). *No truthful online-envy-free mechanism can achieve an $o(\log n / \log \log n)$ approximation to social welfare when faced with adversarial supply.*

Proof. For an envy-free mechanism, we may assume that all offered prices c_1, \dots, c_n are equal: for all i , $c_i = c$. We apply inequality 1 to obtain constraints for the case in which n items arrive, and the case in which 1 item arrives. When n items arrive, we have for all i $\Pr[N_{i-1} < n] = 1$, and obtain the constraint:

$$n \cdot c \geq \frac{(n-1) \log n}{2\alpha} - n \quad (5)$$

When a single item arrives, we have $\Pr[N_{i-1} < 1] = ((n - 2^{c+1})/(n-1))^{i-1}$, since each bidder independently accepts the offer price $1/2^c$ with probability $(2^{c+1} - 1)/(n-1)$. Also, $\mathbf{OPT}_1 \geq 1/2$. We obtain the constraint:

$$(c+1) \cdot \sum_{i=1}^n \left(\frac{n - 2^{c+1}}{n-1} \right)^{i-1} \geq \frac{n-1}{2\alpha} \quad (6)$$

Setting $\alpha = o(\log n / \log \log n)$, we see that constraint 5 requires $c = \omega(\log \log n)$. It is simple to verify that the left hand side of constraint 6 is decreasing in c in the range $[\log \log n, \log(n) - 1]$, and that setting $c = \omega(\log \log n)$ fails to satisfy 6, which proves the claim. \square

A.5 Proofs from Section 6

We begin by presenting a lower bound for Knapsack utilities.

Proposition A.10 (Proposition 6.1). *No algorithm can guarantee better than a $\sum_{i=1}^m h_i$ approximation to optimal social welfare.*

Proof. Consider any arbitrary distribution D and scale it so that it has positive support on $[m+1, 2m]$. Alternately, imagine it has positive support on $[1, m]$, and that m items are guaranteed to arrive; the distribution is on how many additional items will arrive. We construct a set of $n = m$ bidders $1, \dots, m$. Bidder i has $k_i = m + i$ and $c_i = 1/\Pr[\ell \geq i]$. By construction, at most one bidder can have his demand satisfied by any knapsack size. Since bidder values are non-decreasing, we have

$$\mathbf{OPT} = \sum_{i=1}^m c_i \cdot \Pr[\ell = i] = \sum_{i=1}^m \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} = \sum_{i=1}^m h_i$$

However, since at most one bidder can be satisfied by any knapsack size, no algorithm can do better than picking some bidder i and assigning all items that arrive to bidder i . Such an algorithm achieves welfare c_i

in the case that k_i items arrive. By construction, this yields expected welfare $(1/\Pr[\ell \geq i]) \cdot \Pr[\ell \geq i] = 1$, which completes the proof. \square

KnapsackGuess(D):

1. Solicit bids. For each bidder i , create a knapsack instance with one item corresponding to each bidder i , with size k_i and value c_i . For each $s \in [1, m]$ let \mathbf{OPT}_s be the value of the optimal solution to this knapsack instance when the knapsack has size s .
2. Let $s^* = \arg \max_s \Pr[\ell \geq s] \cdot \mathbf{OPT}_s$.
3. Assign items as they arrive to bidders corresponding to the optimal solution for a knapsack of size s^* in an arbitrary order, until each bidder i in the solution has received his demand, k_i items.

Remark A.11. Rather than solving the knapsack problem exactly to find \mathbf{OPT}_s , we can use the greedy-by-density algorithm to find a 2-approximation.⁷ It is simple to see that the greedy knapsack algorithm can only ever output at most $2n$ distinct solutions, regardless of knapsack size. Therefore, at the cost of a factor of 2, our algorithm only has to consider $2n$ solutions, each of which can be computed in polynomial time.

Proposition A.12 (Proposition 6.2). *For any distribution D with (arbitrary) hazard rate h_i KnapsackGuess(D) achieves at least a $\sum_{i=1}^m h_i$ approximation to optimal social welfare.*

Proof. KnapsackGuess(D) achieves welfare \mathbf{OPT}_{s^*} whenever s^* items arrive, which occurs with probability $\Pr[\ell \geq s^*]$. Therefore, KnapsackGuess achieves welfare at least $\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*] \geq \mathbf{OPT}_{s'} \cdot \Pr[\ell \geq s']$ for all s' . Let \mathbf{OPT} denote the expected optimal welfare when the number of items to be sold is drawn from D . If KnapsackGuess achieves no better than an α approximation to social welfare, then for all $s' \in [1, m]$: $\mathbf{OPT}_{s'} \cdot \Pr[\ell \geq s'] \leq \mathbf{OPT}/\alpha$, or equivalently:

$$\mathbf{OPT}_{s'} \leq \frac{\mathbf{OPT}}{\alpha \Pr[\ell \geq s']}.$$

By definition:

$$\mathbf{OPT} = \sum_{i=1}^m \mathbf{OPT}_i \cdot \Pr[\ell = i].$$

Using our above bound on \mathbf{OPT}_i :

$$\mathbf{OPT} \leq \sum_{i=1}^m \mathbf{OPT} \cdot \frac{\Pr[\ell = i]}{\alpha \Pr[\ell \geq i]}.$$

Therefore:

$$\alpha \leq \sum_{i=1}^m \frac{\Pr[\ell = i]}{\Pr[\ell \geq i]} = \sum_{i=1}^m h_i$$

which completes the proof. \square

⁷The greedy-by-density algorithm first discard all items of size larger than the knapsack size and then picks the best of the following two allocations: the greedy-by-density allocation that picks requests in decreasing ratio of value to size until the next element does not fit, and the allocation that gives all the items to the request of highest value.