Auctions with Online Supply

Moshe Babaioff
Microsoft Research
Herzliya, Israel
moshe@microsoft.com

Liad Blumrosen
Department of Economics
The Hebrew University
Jerusalem, Israel 91905
blumrosen@huji.ac.il

Aaron Roth
Computer and Information Sciences
University of Pennsylvania
Philadelphia, PA 19104
aaroth@cis.upenn.edu

Abstract

Online advertising auctions present settings in which there is uncertainty about the number of items for sale. We study mechanisms for selling identical items when the total supply is unknown but is drawn from a known distribution. Items arrive dynamically, and the seller must make immediate allocation and payment decisions with the goal of maximizing social welfare. We devise a simple incentive-compatible mechanism that guarantees some constant fraction of the first-best solution. A surprising feature of our mechanism is that it artificially limits supply, and we show that limiting the supply is essential for obtaining high social welfare. Although common when maximizing revenue, commitment to limit the supply is less intuitive when maximizing social welfare. The performance guarantee of our mechanism is in expectation over the supply distribution; We show that obtaining similar performance guarantee for every realization of supply is impossible.

keywords: Dynamic auctions, unknown supply, online auctions, approximation, dynamic mechanism design, stochastic supply.

JEL codes: D44, D82, D83.

1 Introduction

Consider the problem that an online advertising marketplace faces when selling banner advertisements. The items being sold are "page impressions" of a particular banner-ad slot, and they arrive whenever a new user navigates to the website in question. The items being sold are identical,[1] and in a first cut approximation, the set of bidders in this auction is fixed.[2] On the other hand, the *supply* of items is dynamic: new page impressions are constantly arriving, and although the market maker might have distributional information about the number of impressions he expects to receive in a day, he does not know this quantity with certainty ahead of time. Moreover, these items are perishable: page views cannot be stored and saved for later allocation. Therefore, items must be allocated immediately when they arrive, without waiting for information about future supply. Finally, since the market is essentially ongoing, and it is not possible to wait indefinitely for payments, bidder payments must also be computed before the final supply is fully realized.

This paper considers a natural abstraction of dynamic mechanism design with the goal of capturing the key properties of the advertisement market described above, along with other kinds of markets with similar considerations.[3] In contrast to previous work on dynamic auctions which have focused on dynamic bidder arrival (see e.g. the recent surveys on dynamic auctions by Parkes (2007) and Bergemann and Said (2010)), we have a fixed set of n bidders known to the mechanism, who each have a private value for obtaining a single item. There is an unknown supply of ℓ identical items that arrive one at a time (i.e. ℓ is not known to the mechanism). Finally, we require that the mechanism be *prompt* – that it make allocation and payment decisions immediately when each item arrives, without waiting to see if another item is coming. In this model, we investigate when it is possible for mechanisms to obtain a significant fraction of the social welfare obtainable when the supply is known in advance (the "optimal social welfare").

We study this model with an eye towards one of the "holy grails" of the mechanism-design literature – the design of "detail-free" mechanisms, along the lines of the *Wilson Doctrine* (Wilson (1987)). The ideal goal is to design robust mechanisms which need not be aware of specific characteristics of the environment, and in particular, should not rely on common-prior assumptions. Specifically, we wish to avoid assuming the existence of a known prior on bidder valuations, since it is not clear that an advertising platform can get good information on the valuations of bidders, who may be heterogeneous and anonymous. Specifically, we ask whether any mechanism exists which can guarantee a reasonable fraction of the social welfare when the mechanism has no distributional information about bidder valuations. To formalize the notion of "a reasonable fraction of the social welfare", we borrow a robust asymptotic dichotomy that has proven extremely useful in the computer science literature. We think of social welfare as being approximable in a given setting if there exists some ex-post truthful mechanism which guarantees to always achieve some constant fraction of the social welfare. In contrast, we think of social welfare as being in-approximable if every ex-post truthful mechanism achieves only a diminishing fraction of the optimal social welfare in the worst case. We show that even in the absence of any distributional information about bidder valuations, it is possible to approximate the optimal social welfare when supply is uncertain. Our algorithms do rely on distributional information about the supply. This is more reasonable in our setting, because the auctioneer might have extensive history selling page views of a given website. Moreover, we show that distributional information about the supply is necessary: no mechanism can guarantee a non-diminishing fraction of the optimal social welfare in the worst-case over the realization of the supply.

[1] Each ad impression is identical from the advertiser's perspective in the absence of knowledge about who is viewing the website. It is true that advertisers sometimes have coarse information about users by buying cookie-data from providers such as Acxiom, but in this case we can simply restrict our attention to ad impressions viewed by a particular fixed demographic.

[2] Different advertisers can of course enter and leave the market for a particular banner ad slot, but this happens at a human time scale, whereas ad impressions arrive at a time scale of milliseconds.

[3] 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.

Our positive results highlight a novel phenomenon: in the presence of uncertainty about the supply, it is necessary to artificially *restrict* supply in order to approximate social welfare. Artificially restricting supply is a common technique when the goal is to maximize revenue, but in our setting, it is a surprising necessity, given that our goal is welfare maximization.

1.1 Our Results

We 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. Throughout this paper, we study mechanisms which are ex-post dominant strategy truthful (i.e. truthful for every realization of the unknown supply)[4]. We make the assumption (standard in mechanism design in the context of bidder valuations, but also natural here) that D has a non-decreasing hazard rate[5]. We again stress that we make *no* assumptions about bidder valuations at all. We obtain a positive result:

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 it's analysis is surprisingly subtle. As noted, the incentive properties of our mechanisms do not rely on any distributional information and truthful bidding is a dominant strategy for every set of bids, for *every* realized supply and not only in expectation. Surprisingly, our mechanism relies on restricting supply, which is unusual in a welfare-maximization setting.

We then characterize the set of truthful mechanisms that are constrained to collect payments as items are allocated, and prove a surprising lemma (see Lemma 3.4, Section 3): in the stochastic setting, we can (almost) without loss of generality consider mechanisms that determine an upper bound on the number of items to be sold without considering the bids. More specifically, we observe that in a setting with uncertainty over supply, a mechanism that does not limit the supply can obtain an arbitrarily bad approximation to social welfare. We might a priori think that in order to achieve a good approximation, a mechanism may have to limit the supply in a way that depends on some arbitrarily complicated function of the bids. However, Lemma 3.4 shows that this is not the case: for any truthful mechanism, there is another truthful mechanism obtaining almost the same welfare approximation, which deterministically limits supply in a simple, bid-independent banner.

Using this characterization, we ask whether or not our assumption that the mechanism knows a distribution on the supply is necessary. We show that this is indeed necessary. In the adversarial supply setting in which welfare guarantees are required to hold for any realization of supply, we show the following negative result:

Theorem: *Every truthful mechanism achieves a diminishing fraction (in the number of bidders) of the optimal social welfare in the worst case over realization of the supply.*

[4]The weaker solution concept of truthfulness in expectation over the realization of the supply would require that we assume that the bidders are risk neutral. We note that if we only required truthfulness for risk-neutral bidders, then in the stochastic setting we could obtain optimal social welfare by merely charging every agent their "expected" VCG price – but in addition to risk neutrality, this would require an assumption that all bidders agree with the mechanism on the prior distribution over supply. We view these as needlessly strong assumptions, especially in light of our positive results.

[5]A 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.

1.2 Related Work

Dynamic settings with immediate payment decisions have been studied in several recent papers. Gershkov and Moldovanu (2009, 2012) study dynamic models in which the values of the bidders are drawn from a distribution that is initially unknown to the seller. The seller learns the distribution as bidders arrive, but cannot charge them expected VCG prices while the distribution is unknown. Like in our model, a VCG solution would be easy to implement if payment decisions could be deferred until the end of the auction. Cole et al. (2008) also required that bidders learn their payment immediately upon winning an item (they called such mechanisms *prompt mechanisms*). They study 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,[6] 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 (2005) and by Hajiaghayi et al. (2005). 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, yet bidders arrive and depart over time. Bounds on the efficiency loss in dynamic auctions were also shown, e.g., in Lavi and Nisan (2004) and Compte et al. (2012).

Mahdian and Saberi (2006) studied mechanisms in which the supply is unknown and arrives online. They consider revenue maximization in the adversarial supply setting when payments may be deferred until the end of the auction. They achieve a constant approximation to the optimal revenue. Devanur and Hartline (2009) studied the Bayesian optimal mechanism for the same setting. Our focus on the other hand is on welfare approximation and not on revenue, and we require immediate payments.

A recent line of papers studied dynamic mechanism design in Bayesian settings (e.g., Cavallo et al. (2006); Athey and Segal (2013); Bergemann and Valimaki (2010)), 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 impossibility results show, socially-efficient outcomes cannot be truthfully implemented. In the economics literature, stochastic supply has only been studied in a handful of papers. Most of this work (see, for example, Jeitschko (1999); Neugebauer and Pezanis-Christou (2007)) studied a Bayesian model and focused on characterizing equilibrium prices. Uncertain supply models can be viewed as more complicated versions of the classic sequential-auction model, which is technically hard to analyze even without uncertainty on the supply (see, e.g., Milgrom and Weber (1982); McAfee and Daniel (1993)). Finally, two recent papers with algorithmic results in the same spirit as ours is the work by Feldman et al. (2009), who showed that the $e/(e-1)$-approximation barrier can actually be improved when the input to the matching problem is stochastic, and the work of Devenur and Hayes (2009), who studied an online keyword matching problem where bidders are constrained by a budget; Devenur and Hayes (2009) showed that when statistical information is used (they assumed a random permutation of bidders), then a similar $e/(e-1)$-approximation can be significantly improved. Note that unlike the work of Feldman et al. (2009) and Devenur and Hayes (2009), the main constraint in our paper is incentive compatibility.

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. 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

[6] An α competitive (approximation) mechanism always guarantees $1/\alpha$ fraction of the optimum. See Hartline (2012) for some background on approximation in auctions.

i some price p_i. When ℓ items arrive and bidders have submitted bids $v'_1, \ldots, v'_n$, we denote the outcome of the mechanism by $\mathcal{M}_\ell((v'_1, \ldots, v'_n), r)$ where r is a random bitstring which may be used by randomized mechanisms (think about a set of coin-flip results). We note that the mechanism is unaware of ℓ, as it only encounters the items one at a time as they arrive. To simplify notation we sometimes do not explicitly mention the dependance on 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 $u_i(v_i; \mathcal{M}_\ell(v'_1, \ldots, v'_n)) = v_i - p_i$, where p_i is the payment that the mechanism $\mathcal{M}_\ell$ determines for bidder i for the given input. Bidders who do not receive an item pay 0 and obtain utility 0 (which make our mechanisms ex-post individually rational). We assume that bidders aim to maximize utility and cannot collude.

As we do not assume any distribution on the preferences of the bidders, we consider dominant strategy equilibria. Due to the revelation principle (Myerson (1979)), we can focus on *truthful* mechanisms: mechanisms where bidders are incentivized to report their true valuations, regardless of the bids of others or the realizations of the supply. Following the literature (e.g., Goldberg and Hartline (2003), Guruswami et al. (2005)) 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: $u_i(v_i; \mathcal{M}_\ell((v_i, v'_{-i}), r)) \geq u_i(v_i; \mathcal{M}_\ell((v'_i, v'_{-i}), r))$*

In other words, a mechanism (which may employ randomness) is ex-post truthful if it is a probability distribution over deterministic mechanisms which are each dominant strategy truthful for every realization of the supply. This is a stronger condition than asking that the mechanism be merely truthful in expectation – i.e. truthful only for risk-neutral expectation maximizers. We will sometime use the term *truthful mechanisms* to denote mechanisms with truthful dominant-strategy equilibrium in the above sense. Note that the dominant-strategy solution concept enables us not to assume anything regarding the beliefs of the bidders on the rationality of the other players. Restricting attention to direct-revelation truthful mechanisms also allows us to use the same notations for the bidders' actual types and their bids in the auction.

Without loss of generality, we imagine that $v_1, \ldots, v_n$ 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 $W(\mathcal{M}_\ell((v_1, \ldots, v_n), r))$.[7] 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 range $\{1, ..., n\}$, we define $\mathbf{OPT} = \mathrm{E}_\ell[\mathbf{OPT}_\ell] = \sum_{i=1}^{n} \mathbf{OPT}_i \cdot \Pr[l = i]$. (Having n items is w.l.o.g. assuming free disposal of items if there are more than n items, and assigning zero probability for items if there are less than n items.)

We measure performance by the approximation guarantees to social welfare. The expectation of the social welfare is taken over the distribution of the supply and possibly also on the randomization made by the mechanism. We will call this model, where the supply distribution is known, the *stochastic-supply model*. We will also consider a model without a distribution on the supply, where the expected welfare is taken only over the randomization of the mechanism. In this case, we would like that our mechanism will achieve the desired approximation for every realization of the supply. In other words, this model takes the worst-case approach, as if the realization is determined by some adversary of our approximation mechanism. We will thus sometime refer to this worst-case model by the term *the adversarial-supply model*.

Definition 2.2. *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 for every profile of values $v_1, ..., v_n$:*

$$\frac{\mathrm{E}_\ell[\mathbf{OPT}_\ell]}{\mathrm{E}_{\ell,r}[W(\mathcal{M}_\ell((v_1, \ldots, v_n), r))]} \leq \alpha \tag{1}$$

[7]Note that the social welfare is defined as the sum of bidders utilities and the seller's surplus (sum of payments), which is equal in models with quasi-linear utilities to the sum of the values of the winning bidders.

A mechanism $\mathcal{M}$ achieves an α- approximation to social welfare in the adversarial supply *setting if for* every *supply ℓ and for every profile of values $v_1, ..., v_n$:*

$$\frac{\mathbf{OPT}_\ell}{\mathrm{E}_r[W(\mathcal{M}_\ell((v_1, \ldots, v_n), r))]} \leq \alpha \tag{2}$$

If for some constant c, which is independent of the parameters of the problem, a mechanism achieves a c-approximation, we say that this mechanism achieves a *constant approximation*. We use the asymptotic notations of $\Omega(f(n))$- and $O(f(n))$-approximation to denote approximation that is asymptotically worse or better, respectively, than $f(n)$-approximation. Using these notations, a *lower bound* on the approximation ratio relates to an impossibility result and an *upper bound* describes a positive result. More on asymptotic and approximation notations can be found in Cormen et al. (2001).

In computing our approximation ratios, we are comparing the performance of our online mechanism with the expected value of the mechanism that knows the supply level in advance - the *offline* optimal mechanism. We let $\mathbf{OPT}$ denote the social welfare achieved by the efficient auction that knows the supply in advance (running VCG with fixed supply). Comparing the mechanism's performance to $\mathbf{OPT}$ is natural in the adversarial-supply setting. In the stochastic-supply setting, we might alternatively seek to compare our mechanisms to the optimal *online* policy, a weaker benchmark. Yet, we adopt the offline benchmark for the stochastic setting as even with respect to this stronger benchmark we are *still* able to achieve constant factor approximations, which makes our results stronger. However, this strong benchmark may make the notation $\mathbf{OPT}$ somewhat misleading: as we show, no online mechanism can achieve $\mathbf{OPT}$ exactly. The benchmark $\mathbf{OPT}$ can also be viewed as the *first-best* outcome, that is, the optimal result had the agents been non-strategic. We remark that our setting, even in the stochastic case, is not a Bayesian setting, and we do not have any probabilistic assumptions on the preferences of the bidders; Thus, the concept of "the efficient mechanism" is not well-defined in our model. This is one reason that we use a worst-case approach in our paper.

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

Definition 2.3. ***(Non-Decreasing Hazard Rate.)*** *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 (usually for the prior on bidder valuations, but also a natural assumption here), and is satisfied by many natural distributions, including the exponential, uniform, and binomial distributions. There are many natural distribution for which this condition does not hold, bimodal distributions are a prominent example.

3 Characterization of Truthful Mechanisms

In this section, we provide a characterization of (dominant-strategy) truthful mechanisms in our setting, and some useful lemmas that allow us to consider mechanisms from a particularly simple class in the stochastic setting. The characterization follows from two simple lemmas. These lemmas claim that in truthful mechanisms a bid of a winning bidder cannot affect the price she pays and what item (out of the series of arriving items) she receives.

Lemma 3.1. *For every truthful mechanism and for any realization of the supply, 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 the supply 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. □

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

Proof. Note first that although the sold items are identical, they do differ in the order they arrive to the market. Suppose for some realization of the supply, and for some fixed set of bids of the other bidders, bidder b win the j-th item when bidding v_b, and wins the i-th item for $i < j$ when bidding v'_b. 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 $v_b - p_b > 0$. Therefore, the mechanism is not truthful. □

The converse of the characterization is immediate: any mechanism that assigns winner order and prices in a bid independent way is truthful.

It turns out that in the stochastic-supply setting, we can (almost) without loss of generality consider a particularly simple kind of mechanisms. Such mechanisms choose some supply level, and commit not to sell more than this number of units. We show in Lemma 3.4 that the approximation factor achieved by any truthful mechanism can be approximated by such mechanisms.

Definition 3.3. ***(Mechanisms with Bid-Independent Supply Limit)*** *A* mechanism with bid-independent supply limit *chooses an ordering on the bidders π and a bound g on supply independently of the bids. It fixes a list of the g highest bidders, ordered according to π. When an item arrives it is sold to the next bidder on the list, and this bidder pays the $g + 1$st highest bid.*

We remark that the more straightforward mechanism that assigns items to bidders by the order of their bids and charges everyone the $g+1$st highest price is *not* truthful, exactly because of the supply uncertainty. For instance, if less than g items arrive, and the price charged is the $g + 1$st price, then bidder g who does not receive an item can pretend to have a very high value and win, gaining a positive surplus this way. Our equilibrium concept, however, requires bidders to have a truthful dominant strategy even after the supply had been realized. Fixing the order of the players in a way that is independent of the bids, as suggested in Definition 3.3, solves this problem.

We now show that the social welfare of any truthful mechanism can be approximated by a mechanism with bid-independent supply limit.

Lemma 3.4. *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 mechanism M' with bid-independent supply limit 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 sells the items at uniform prices 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. □

We use Lemma 3.4 in our lower bounds proofs, as it enables us to restrict attention to mechanisms with bid-independent supply limit for ruling out constant approximations in the adversarial-supply setting. As mentioned, it may also serve as a rule of thumb to mechanism designers under uncertain supply, saying that they can focus on this family of simple, intuitive mechanisms without losing much efficiency. This will hold although committing for not selling the entire supply may sound non-intuitive.

We now show via simple examples that mechanisms that do not commit to limit supply perform very badly. We say that a mechanism *sells every item* if for every realized supply, the mechanism always sells every available item. We observe that any truthful mechanism that sells every item performs very poorly.

Observation 3.5. *There exists a non-decreasing hazard rate supply distribution such that every deterministic truthful mechanism that sells every item has an approximation which is exponential in n.*

Proof. Consider the exponential distribution with probability of supply realized to l being proportional to e^{-l}. By Lemma 3.1 the price p_b that bidder b is charged upon winning (any) item is independent of his bid. In case that the supply is n, the mechanism sells all n items and every bidder wins (even if bidding arbitrarily low), thus it must be the case that $p_b = 0$ for every bidder b. Moreover, by Lemma 3.2 if bidder b wins an item, which item bidder b wins must be independent of his bid whenever $p_b = 0 < v_b$, which is always true. Thus, any truthful mechanism must fix some order that is independent of the bids, and allocate according to that order. Now, fix any small $\epsilon > 0$ and consider the input with all bidders having value $\epsilon > 0$, while the last bidder according to the mechanism's order having value 1. The social welfare is maximized by always allocating the first item to this bidder, achieving welfare of at least 1. On the other hand, the mechanism allocates to the high value bidder with probability proportional to e^{-n}, thus achieving expected value of at most $(n-1)\epsilon + e^{-n}$. □

We also show (see Appendix A.1), that randomized mechanism that always sell the entire supply achieve no better than n approximation.

4 Stochastic Supply

In this section, we give our main result, a truthful mechanism that achieves a constant-approximation to social welfare for any distribution with non-decreasing hazard rate. We also show that a better than a constant approximation (that is, a $(1+\epsilon)$-approximation for every ϵ) is impossible. We do that by proving that no truthful deterministic mechanism can achieve better than a $\phi \approx 1.63$ approximation to social welfare, and no randomized mechanism can achieve better than a $4/(2+\sqrt{2}) \approx 1.17$ approximation, for distributions that satisfy the monotone hazard rate condition.

We consider the following mechanism (we call "HazardGuess") 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*, thus it belongs to the family of mechanisms bid-independent supply limit. Although it seems somewhat surprising, it still achieves good approximation when the non-decreasing hazard rate condition holds. We later show (Theorem 5.2) that it is impossible to construct a mechanism that achieve a constant approximation for every realization of the supply (rather than on average).

HazardGuess(D):

1. Fix an arbitrary permutation π on the bidders.
2. Solicit bids, and denote them $v_1, \ldots, v_n$ 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 $g = n$).

[a] Alternatively, 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 constant approximation to social welfare in expectation over D, for any distribution D with non-decreasing hazard rate.*

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 the mechanism 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 $\mathbf{OPT}$.

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 $s \in \{0, \ldots, n\text{-}1\}$ we have the following bound on α in terms of D, which we denote 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:

$$\mathbf{OPT}_i < \frac{\mathbf{OPT}}{\beta \cdot \Pr[\ell \geq i]} \tag{3}$$

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 $v_1, \ldots, v_n$ 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 3 and multiplying both sides by $\beta/\mathbf{OPT}$ 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. □

Remark 4.3. *We must now show that for every distribution D, there exists an s such that 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, 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.2. □

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

Proof. Given a distribution D, we wish to find the value of s such that 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 $s^* \geq \frac{\Pr[\ell \geq s^*]}{\Pr[\ell = s^*]}$. If no such s^* exists, we choose $s^* = n$. We now show that Bound(s^*) gives $\alpha \leq 5$. We bound the two terms of 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:

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

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

$$\Pr[\ell \geq s^* + 1](3s^* + 1)$$

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

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

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

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.

Let s^* be the smallest integer such that $s^* \geq \frac{\Pr[\ell \geq s^*]}{\Pr[\ell = s^*]}$. Whenever $s^* > 3$, HazardGuess(D) achieves welfare at least $\mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*]$. When $s^* \leq 3$, 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, 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:

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

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

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

which is a bound on the performance of 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 $s^* = 3$ it achieves at least $\mathbf{OPT}_i/3 \cdot \Pr[\ell \geq i]$. Combining the $\frac{27}{8}$ factor and the 5 factor from Lemma 4.5 we get at least a $16\frac{7}{8}$ approximation, which concludes the proof. □

4.1 Better Bounds for Specific Distributions

In our analysis of the mechanism HazardGuess we take a worst-case approach over the possible distributions of the supply, and we get a guarantee of $16\frac{7}{8}$ approximation which appear to be low for practical purposes. However, we note that this exact mechanism can be shown to achieve a better constant approximation for specific distributions of interest. Proofs from this section are found in Appendix A.2.1. For example:

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

Proposition 4.7. ***(Geometric distributions.[8])*** *HazardGuess(D) achieves an $(e+1)/(e-1) \approx 2.16$-approximation to social welfare in expectation over D when D is the geometric distribution defined as $\Pr_D[\ell = k] = (1-p)^{k-1}p$, for any value of the parameter $p \in (0,1)$.*

(Binomial distributions.) We also compute bounds on the approximation ratio of our mechanism for the binomial distribution for specific values of n, and find that it achieves at least 64% of the optimum in the range of n that we check. This approximation factor quickly converges to 100% as n grows. See Appendix A.2.1 for details of the analysis and a graph of the bound we prove on the approximation factor for specific values of n.

4.2 Impossibility Results for the Stochastic Model

In considering these bounds, one should note that it is impossible to get perfect efficiency, since we are comparing ourselves to a very strong adversary: the *offline* optimum. We show that it is impossible to get close to perfect efficiency. We prove constant lower bounds, bounded away from 1, both for deterministic and randomized truthful mechanisms:

Proposition 4.8. *No deterministic truthful mechanism guarantees better than a ϕ-approximation to social welfare in expectation over distributions that satisfy the monotone hazard rate condition, where $\phi = (1+\sqrt{5})/2$ is the golden ratio. No randomized truthful mechanism can guarantee better than a $4/(2+\sqrt{2})$ approximation.*

Proof. In Appendix A.2.2. □

The Necessity of a Condition on the Distribution: We now show the necessity of our non-decreasing hazard rate assumption. We show that there exists a regular distribution that does not satisfy the non-decreasing hazard rate assumption, for which for *no* deterministic mechanism can achieve constant approximation to social welfare.

Theorem 4.9. *There exists a regular distribution (that does not satisfy the non-decreasing hazard rate condition) such that no deterministic truthful mechanism can achieve an $o(\sqrt{\log n/\log\log n})$ approximation to social welfare when faced with stochastic supply drawn from this distribution.*

Proof can be found in Appendix A.2.2. We prove the theorem in two stages. First, we consider only mechanisms with bid-independent supply limit (see Definition 3.3), Such mechanisms sells the first g items at the $(g+1)$-st highest price to the first g bidders, ordered according to some permutation π. We note that HazardGuess is such a mechanism. 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 using Lemma 3.4 which shows that we can restrict our attention to such mechanisms almost without loss of generality: If a deterministic truthful mechanism could achieve a $o(\sqrt{\log n/\log\log n})$-approximation, then by Lemma 3.4 a mechanism with bid-independent supply limit could achieve a $o(\log n/\log\log n)$ approximation, which we show is impossible.

5 Adversarial Supply

In this section we consider a more demanding model than the stochastic-supply model – the adversarial-supply model. In this model, we do not have a distribution over the supply and we require a good approximation to social welfare for *any* number of items that arrive. We first show that deterministic truthful

[8]We remark that the geometric distribution is the extremal MHR distribution: the hazard rate of the geometric distribution is constant.

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. We also show that $\log n$ approximation is achievable by a randomized truthful mechanism, and that, when assuming that all winners are charged the same price, this is essentially optimal.

5.1 Deterministic Mechanisms

We begin by proving that deterministic mechanisms can only achieve a trivial approximation. The proof follows easily from our characterization of truthful mechanisms, and a simple lemma:

Lemma 5.1. *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. It could happen than a bidder's valuation, and therefore his bid b, is more than n times the value of 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. □

Theorem 5.2. *No deterministic truthful mechanism can achieve better than an n approximation to social welfare.*

Proof of Theorem. By Lemma 5.1, any bidder can win the first item with an appropriately high bid. The bidder who wins the first item will pay a payment p_b which is independent of her bid due to Lemma 3.1. But by Lemma 3.2, the item that the bidder wins is independent of his bid, and therefore of the bidder wins the first item for some $p_b < v_b$ she 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 bidders 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* bidders 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. □

5.2 Randomized Mechanisms

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

We next present our first impossibility result for randomized truthful mechanisms.

Theorem 5.3. *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, see Yao (1977), 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 5.4. *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:* $\mathrm{E}[\mathbf{OPT}_k] \geq H_{k+1} - 1$ *where H_{k+1} denotes the $k+1$st harmonic number. In particular,* $\mathrm{E}[\mathbf{OPT}_k] > (\log k)/2$.

Proof. We defer the proof to Appendix A.3. □

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.[9] 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 $c_b \in \{0, \ldots, \log n - 1\}$.

When all n items arrive, the expected welfare achieved by a mechanism is:

$$\sum_{b=1}^{n} \Pr[v_b \geq \frac{1}{2^{c_b}}] \cdot E[v_b | v_b \geq \frac{1}{2^{c_b}}] = \frac{1}{n-1} \sum_{b=1}^{n} (c_b + 1). \tag{4}$$

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 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 $1 \leq k \leq n$:

$$\sum_{b=1}^{n} (c_b + 1) \Pr[N_{b-1} < k] \geq \frac{(n-1)\mathbf{OPT}_k}{\alpha} \geq \frac{(n-1)\log k}{2\alpha} \tag{5}$$

where the last inequality follows from Lemma 5.4. After offering the item to b bidders, the expected number of sales is $E[N_b] = 1/(n-1) \cdot \sum_{i=1}^{b} (2^{c_b+1} - 1)$.

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 $\sum_{i=1}^{b_k} 2^{c_i} \geq (n-1) \cdot k$. By plugging our bound into Constraint (5), 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 5.5. *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.3. □

[9] An 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.

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:

$$\sum_{i=1}^{b_k} c_i \geq \frac{n \log k}{4\alpha} \tag{6}$$

$$\sum_{i=1}^{b_k-1} 2^{c_i} < n \cdot k \tag{7}$$

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 7 continues to hold, and so this is without loss of generality.

We let $k = 2^{15\alpha}$ and consider the sequence of integers $k, 2k, 4k, \ldots, 2^t k$ 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 6 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:

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

where the last inequality follows from inequality 7. 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 $\alpha \geq \Theta(\log \log n)$. □

5.2.2 A truthful $\log n$-approximation mechanism

Here we show a simple randomized mechanism that achieves a $\log n$ approximation to social welfare. In Section 5.3 we show that this is nearly optimal for the natural class of mechanisms that must charge all winners the same price.

Let **RandomGuess** be the mechanism that selects a supply $g \in \{2, 4, 8, \ldots, 2^i, \ldots, n\}$ 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}.[10]

Proposition 5.6. *RandomGuess is truthful and achieves a $\log n$ approximation to social welfare.*

Proof. We defer the (quite standard) proof of this proposition to Appendix A.3. □

[10] We thank Andrew Goldberg for suggesting this mechanism, which is a significant simplification of our original mechanism.

We leave open the problem of closing the gap between the above $\log n$ factor and the $\Omega(\log \log n)$ lower bound of Theorem 5.3. In Section 5.3 we strengthen this lower bound to $\Omega(\log n / \log \log n)$ for the class of mechanisms that must charge all winners the same price. We conjecture that RandomGuess is asymptotically optimal.

5.3 Uniform Charging Mechanisms

All our mechanisms (RandomGuess and HazardGuess) charge all winners the same price (uniform charging), which seems like a desired property. For mechanisms for adversarial supply that are allowed to charge winners different amounts, Theorem 5.3 shows that no truthful randomized mechanism can achieve an $o(\log \log n)$ approximation to social welfare. Here, we present a stronger impossibility result for truthful mechanisms that must charge all winners the same price.

Proposition 5.7. *No truthful mechanism (even randomized) that charges all winners the same price can achieve an* $o(\log n / \log \log n)$ *approximation to social welfare when faced with adversarial supply.*

Proof. We defer the proof of this proposition to Appendix A.4. □

Note that Proposition 5.7 is nearly tight, since RandomGuess achieves a $\log n$ approximation factor.

We remark that while mechanisms that charge all winners the same price seem more fair than arbitrary mechanisms that charge different winners different prices, they are still *not* envy free in the usual sense. An offline mechanism is envy-free if no agent prefers another agent's allocation and payment to his own (see, for example, Goldberg and Hartline (2003); Guruswami et al. (2005)). [11] In our online setting some bidders with value above the price paid by the winners might still be losing, due to shortage of supply, and thus envy the winners. Yet, this kind of envy is inevitable in the online setting if the mechanism sometimes sells more than a single item (which is essential to get sublinear approximation).

6 Discussion

In this paper we discussed the effect of uncertain supply on auction design, an issue that has not received much attention. We presented a simple model that tries to isolate the effect of supply uncertainty. Our goal was to understand what kinds of mechanisms can achieve good approximations to the first-best social welfare (which is impossible for any truthful mechanism to obtain exactly). It is already known from the literature that if the seller has sufficient Bayesian knowledge, and if all payments can be deferred to the end of the auction, then full efficiency can be achieved using variants of dynamic Vickrey-Clarke-Groves mechanisms. However, in many environments payments should be determined when the allocation is made, or at least periodically, and this restriction makes the design more challenging. Our main result of the paper (Theorem 4.1) is positive, showing that obtaining a constant fraction of the first-best social welfare is possible even with prompt payments. We show this by devising a simple mechanism that gains the above guarantee without assuming anything about how the preferences of the bidders are distributed, and for any number of bidders. The mechanism does require knowing how the supply is distributed, and we show that such positive results cannot be achieved without knowledge of the supply distribution.

A striking feature of our results is the role of *commitment*. Our auction ("HazardGuess") relies on the fact that the seller can reliably commit not to sell items above a certain quota. We show that due to the incentive constraints in our model, commitments to limit the supply are necessary for any mechanism that is to obtain a non-diminishing fraction of the first-best social welfare. Without commitment, any approximation guarantee

[11] In traditional economics, envy-free prices often correspond to Walrasian- or market equilibria, which are prices under which no customer prefers the allocation and payment of another customer over his own.

diminishes exponentially with the number of players (Observation 3.5); Furthermore, mechanisms with bid-independent supply limits can achieve approximately the same social welfare as any other incentive-compatible mechanism (Lemma 3.4). Withholding the sale of items is a common practice for monopoly sellers attempting to maximize revenue: see for example the celebrated auction by Myerson (1981); In our setting, surprisingly, committing to limit supply is essential to approximate efficiency.

In this paper, we made several simplifying assumptions in order to focus on the effect of unknown supply. One important assumption is that all items are identical. Heterogenous items can raise issues of substitutability and complementarity that would make the incentive and algorithmic aspects in our analysis more complex. One future direction is studying online supply with more general substitutes valuations (e.g., submodular valuations for which there is a 2 approximation algorithm that already has an online nature, see Lehmann et al. (2006)). Another simplifying assumption we have made is that bidders' preferences do not change over time and thus the mechanism can solicit bids only once, at the beginning of the auction. A richer model will allow bidders' preferences to change over time and this will make the incentive constraints harder to satisfy and might harm our positive results; Future work with dynamic bidding might therefore consider other solution concepts rather our strong (ex-post) dominant-strategy equilibrium. Our main positive result is given for the monotone hazard-rate condition, and another future direction is to identify other properties of distributions that may give better approximation guarantees. Such properties may be taken from a more explicit probabilistic model for the dynamic arrival of units.

Acknowledgements. Aaron Roth was supported in part by an NSF Graduate Research Fellowship, and portions of his work were done while visiting Microsoft Research - Silicon Valley. Portions of this work were done while Liad Blumrosen was a post-doc researcher in Microsoft Research - Silicon Valley, and Liad Blumrosen was also supported by the Israeli Science Foundation (grant no. 230/10) and by the Google Inter-university center for Electronic Markets and Auctions. We thank Andrew Goldberg and Ilya Segal for helpful discussions.

References

Athey, S., Segal, I., 2013. An efficient dynamic mechanism. Econometrica 81 (6), 2463–2485.

Bergemann, D., Said, M., 2010. Dynamic auctions: A survey. Cowles Foundation Discussion Paper 1757.

Bergemann, D., Valimaki, J., 2010. The dynamic pivot mechanism. Econometrica 78 (2), 771–789.

Cavallo, R., Parkes, D. C., Singh., S., 2006. Optimal coordinated planning amongst self-interested agents with private state. In: Conference on Uncertainty in Artificial Intelligence (UAI),. pp. 55–62.

Cole, R., Dobzinski, S., Fleischer, L., 2008. Prompt Mechanisms for Online Auctions. Symposium on Algorithmic Game Theory (SAGT): LNCS 4997, 170.

Compte, O., Lavi, R., Segev, E., 2012. Efficiency of sequential english auctions with dynamic arrivals. In: ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012. p. 369.
URL `http://doi.acm.org/10.1145/2229012.2229041`

Cormen, T., Leiserson, C., Rivest, R., Stein, C., 2001. Introduction to Algorithms, second edition Edition. MIT Press,MA.

Devanur, N. R., Hartline, J. D., 2009. Limited and online supply and the bayesian foundations of prior-free mechanism design. In: ACM conference on Electronic Commerce (EC). pp. 41–50.

Devenur, N. R., Hayes, T. P., 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In: ACM conference on Electronic Commerce (EC). pp. 71–78.

Feldman, J., Mehta, A., Mirrokni, V., Muthukrishnan, S., 2009. Online stochastic matching: Beating 1-1/e. In: IEEE Symposium on Foundations of Computer Science (FOCS). pp. 117–126.

Gershkov, A., Moldovanu, B., 2009. Learning about the future and dynamic efficiency. American Economic Review. 99 (4), 1576–1587.

Gershkov, A., Moldovanu, B., 2012. Optimal search, learning, and implementation. Journal of Economic Theory. 147, 881–909.

Goldberg, A. V., Hartline, J. D., 2003. Envy-free auctions for digital goods. In: ACM conference on Electronic commerce (EC). pp. 29–35.

Guruswami, V., Hartline, J. D., Karlin, A. R., Kempe, D., Kenyon, C., McSherry, F., 2005. On profit-maximizing envy-free pricing. In: ACM-SIAM symposium on Discrete algorithms (SODA). pp. 1164–1173.

Hajiaghayi, M., Kleinberg, R., Mahdian, M., Parkes, D. C., 2005. Online auctions with re-usable goods. In: ACM Conf. on Electronic Commerce (EC). pp. 165–174.

Hartline, J. D., 2012. Approximation in mechanism design. The American Economic Review 102 (3), 330–336.

Jeitschko, T. D., 1999. Equilibrium price paths in sequential auctions with stochastic supply. Economics Letters 64 (1), 67–72.

Lavi, R., Nisan, N., 2004. Competitive analysis of incentive compatible on-line auctions. Theoretical Computer Science. 310 (1), 159–180.

Lavi, R., Nisan, N., 2005. Online ascending auctions for gradually expiring items. In: ACM-SIAM Symposium On Discrete Algorithms (SODA). pp. 1146–1155.

Lehmann, B., Lehmann, D., Nisan, N., 2006. Combinatorial auctions with decreasing marginal utilities. Games and Economic Behaviour (GEB) 55 (2), 270–296.

Mahdian, M., Saberi, A., 2006. Multi-unit Auctions with Unknown Supply. ACM conference on Electronic Commerce (EC), 243 – 249.

McAfee, R. P., Daniel, V., 1993. The declining price anomaly. Journal of Economic Theory (JET) 60 (1), 191–212.

Milgrom, P. R., Weber, R. J., 1982. A theory of auctions and competitive bidding. Econometrica 50 (5), 1089–1122.

Myerson, R. B., 1979. Incentive compatibility and the bargaining problem. Econometrica 47 (1), 61–73.

Myerson, R. B., 1981. Optimal Auction Design. Mathematics of Operations Research 6 (1), 58–73.

Neugebauer, T., Pezanis-Christou, P., 2007. Equilibrium price paths in sequential auctions with stochastic supply. Journal of Economic Behavior and Organization 63 (1), 55–72.

Parkes, D. C., 2007. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani (eds.) Chapter 16, Online Mechanisms. Cambridge University Press.

Wilson, R., 1987. Game-Theoretic Approaches to Trading Processes. In Advances in Economic Theory: Fifth World Congress, ed. by T. Bewley, chap. 2,. Cambridge University Press.

Yao, A. C.-C., 1977. Probabilistic computations: Toward a unified measure of complexity. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science. IEEE, pp. 222–227.

A Proofs

A.1 Proofs from Section 3

We also present an impossibility for randomized mechanisms that always sell all the realized supply, showing that the best performance is achieved by allocating using a uniform random order, still achieving poor approximation of only n.

Observation A.1. *Every truthful randomized mechanism that sells every item has approximation which is no better than n on non-decreasing hazard rate supply distributions.*

Proof. Fix $\epsilon > 0$ and large $K > n^2$. Consider the exponential distribution with probability of supply realized to l being proportional to K^{-l}. The randomized truthful mechanism is a randomization over deterministic truthful mechanisms that sells every item. By the arguments presented in Observation 3.5, each of these mechanisms fixes some order that is independent of the bids, and allocates according to that order, charging each bidder 0. Let i be a bidder with the smallest probability of being first to get an item. That probability is at most $1/n$. Consider the input with bidder i having a value of 1 while all other bidders having value $\epsilon > 0$. The social welfare is maximized by always allocating the first item to bidder i, achieving welfare of at least 1. On the other hand, the randomized mechanism achieve expected welfare of at most $1/n + Pr[l \geq 2]$. The result follows from $Pr[l \geq 2]$ tending to 0 as K grows to infinity. □

A.2 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 $f(h_{s+1}, \ldots, h_n) \equiv \sum_{i=s+1}^{n} (i \cdot h_i \cdot \prod_{j=s+1}^{i-1} (1 - h_j))$ and consider the partial derivative at h_k:

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

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

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

Fix some assignment to the h_i that maximizes $f(h_{s+1},\ldots,h_n)$ 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:

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

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 $\sum_{i=k}^{\infty} i\cdot r^{i-k} = (k+r-kr)/(r-1)^2$. □

A.2.1 Improved Bounds for Specific Distributions

In this section, we prove improved bounds on the approximation ratio of HazardGuess for specific distributions of interest. We make use of the following lemma in all of our proofs:

Lemma A.2. *For any distribution D, we have the following bounds on $\mathbf{OPT}$ and ALG, the welfare achieved by HazardGuess(D):*

$$\mathbf{OPT} \leq \mathbf{OPT}_{s^*}\cdot\Pr[\ell < s^*]+\sum_{i=s^*}^{\infty}\Pr[\ell = i]\cdot i/s^*\cdot\mathbf{OPT}_{s^*}$$

$$ALG \geq \mathbf{OPT}_{s^*}\cdot\Pr[\ell \geq s^*]+\sum_{i=1}^{s^*-1}\Pr[\ell = i]\cdot i/s^*\cdot\mathbf{OPT}_{s^*}$$

Proof. This follows from the definition of $\mathbf{OPT}_{s^*} = \sum_{i=1}^{s^*} v_i$, and the fact that v_i are ordered in decreasing order of value. Thus, for $i < s^*$, $\mathbf{OPT}_i \leq \mathbf{OPT}_{s^*}$, and for $i > s^*$, $\mathbf{OPT}_i \leq \frac{i}{s^*}\mathbf{OPT}_{s^*}$. □

Note that the ratio of these bounds for $\mathbf{OPT}$ and ALG always constitute an upper bound on the approximation ratio of HazardGuess.

Proposition 4.6: *HazardGuess(D) achieves a $\frac{5}{3}$-approximation to social welfare in expectation over D when D is the uniform distribution over $\{1, \ldots, n\}$. Moreover, there are values for which HazardGuess(D) cannot get better than a $\frac{4}{3}$-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 $\{1, \ldots, n\}$ (we note that if the range starts from a number larger than 1 the problem becomes easier and the algorithm achieves better approximation.) For the uniform distribution, HazardGuess selects $s^* = n/2$.[12] We prove that the algorithm achieves at least 60% of the optimum by applying Lemma A.2.

$$OPT \le \frac{OPT_{\frac{n}{2}}}{2} + \frac{1}{n} \cdot \sum_{l=\frac{n}{2}+1}^{n} \frac{l}{n/2} OPT_{\frac{n}{2}} \tag{8}$$

$$= OPT_{\frac{n}{2}} \left(\frac{1}{2} + \frac{2}{n^2} \sum_{l=\frac{n}{2}+1}^{n} l \right) \tag{9}$$

$$= OPT_{\frac{n}{2}} \left(\frac{1}{2} + \frac{2}{n^2} \left(\frac{n(n+1)}{2} - \frac{\frac{n}{2}(\frac{n}{2}+1)}{2} \right) \right) \tag{10}$$

$$= OPT_{\frac{n}{2}} \left(\frac{5}{4} + \frac{1}{2n} \right) \tag{11}$$

Our algorithm achieves expected welfare of

$$ALG = \frac{1}{n} \cdot \sum_{l=1}^{\frac{n}{2}} \frac{l}{n/2} OPT_{\frac{n}{2}} + \frac{1}{n} \cdot \sum_{l=\frac{n}{2}+1}^{n} OPT_{\frac{n}{2}} \tag{12}$$

$$= OPT_{\frac{n}{2}} \left(\frac{2}{n^2} \sum_{l=1}^{\frac{n}{2}} l + \frac{1}{2} \right) \tag{13}$$

$$= OPT_{\frac{n}{2}} \left(\frac{2}{n^2} \frac{\frac{n}{2}(\frac{n}{2}+1)}{2} + \frac{1}{2} \right) \tag{14}$$

$$= OPT_{\frac{n}{2}} \left(\frac{3}{4} + \frac{1}{2n} \right) \tag{15}$$

$$\ge \frac{OPT}{\frac{5}{4} + \frac{1}{2n}} \cdot \left(\frac{3}{4} + \frac{1}{2n} \right) \quad \ge \quad \frac{3}{5} OPT \tag{16}$$

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. ☐

Proposition 4.7: *HazardGuess(D) achieves an $(e+1)/(e-1) \approx 2.16$-approximation to social welfare in expectation over D when D is the geometric distribution defined as $\Pr_D[\ell = k] = (1-p)^{k-1}p$, for any value of the parameter p.*

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

Proof. We recall that $\Pr[\ell \geq k] = (1-p)^{(k-1)}$. For simplicity, assume $1/p$ is an integer (the proof is similar even if not). Given a geometric distribution with parameter p, HazardGuess selects a supply $s^* = 1/p$ if $1/p > 3$, and $s^* = 1$ otherwise. (If $1/p > n$, we select $s^* = n$, and our mechanism only does better in this case) We can again apply Lemma A.2:

$$\mathbf{OPT} \leq \mathbf{OPT}_{s^*} \cdot \Pr[\ell < s^*] + \sum_{i=s^*}^{\infty} \Pr[\ell = i] \cdot i/s^* \cdot \mathbf{OPT}_{s^*}$$

Similarly:

$$ALG \geq \mathbf{OPT}_{s^*} \cdot \Pr[\ell \geq s^*] + \sum_{i=1}^{s^*-1} \Pr[\ell = i] \cdot i/s^* \cdot \mathbf{OPT}_{s^*}$$

Simplifying, we find that for $s^* = 1/p$:

$$\mathbf{OPT}/ALG \leq \frac{2}{1-(1-p)^{1/p}} - 1$$

Taking the limit as $p \to 0$ we get that the approximation factor is always at most $(e+1)/(e-1)$. For the natural special case of $p = 1/2$, $s^* = 1$. In this case:

$$\mathbf{OPT}/ALG \leq \frac{\mathbf{OPT}_1 \sum_{i=1}^{n} 2^{-i} \cdot i}{\mathbf{OPT}_1} = 2$$

□

(Binomial distributions.) We also analyze HazardGuess on the binomial distribution, defined by the probability mass function $\Pr_D[\ell = k] = \frac{\binom{n}{k}}{2^n}$. It appears to be difficult to compute an analytic expression for s^* for the binomial distribution, so rather than proving a general bound, as we do for the uniform and geometric distributions, we numerically compute s^* for various values of n, and prove a bound for each value of n separately by applying Lemma A.2. We find that HazardGuess achieves an approximation factor of at least 64% for $n = 3$, and quickly converges to 100% as n grows. See Figure A.2.1, which plots the approximation factor we prove for HazardGuess on the binomial distribution by n. For each value of n, we derive the expression for the inverse hazard rate of the binomial distribution on n elements, and find the exact value of s^* by using binary search. Given s^* for a particular value of n, we apply the bounds on $\mathbf{OPT}$ and ALG given in Lemma A.2, and compute their ratio. Figure A.2.1 plots the ratio of these bounds on $\mathbf{OPT}/ALG$ for various values of n.

A.2.2 Lower Bounds for the Stochastic Model

Proposition 4.8: *No deterministic truthful mechanism guarantees better than a ϕ approximation to social welfare in expectation over distributions that satisfy the monotone hazard rate condition, where $\phi = (1+\sqrt{5})/2$ is the golden ratio. No randomized truthful mechanism can guarantee better than a $4/(2+\sqrt{2})$ approximation.*

Before proving the theorem, we prove a lemma regarding the deterministic case. The lemma shows that if the supply consists of at most two potential items, we may without loss of generality consider mechanisms with bid-independent supply limit (see Definition 3.3).

Lemma A.3. *For any distribution D over at most 2 items, if mechanism M guarantees an $1/\alpha$ approximation to social welfare in expectation over D, there exists a mechanism M' with bid-independent supply limit that also achieves a $1/\alpha$ approximation to social welfare in expectation over D.*

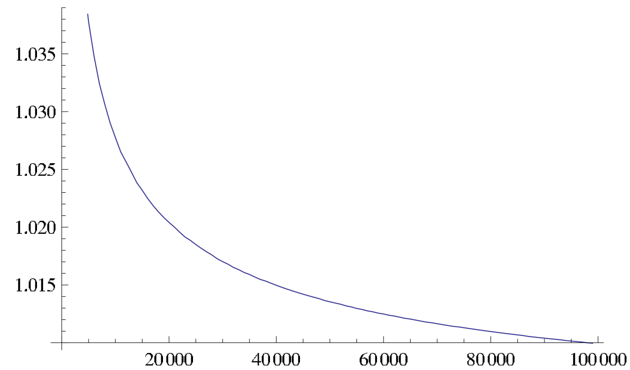

Figure 1: The approximation factor of HazardGuess on the binomial distribution. The horizontal axis represents n, and the vertical axis represents the corresponding bound on the approximation factor of HazardGuess that we prove. We plot values of n up to 100000, with one data point for every multiple of 1000.

Proof. Let D be a distribution over 2 items such that $\Pr_D[\ell = 1] = 1 - \beta$ and $\Pr_D[\ell = 2] = \beta$. Note that D satisfies the monotone hazard rate condition, as do all 2 item distributions.

We can assume that for some bid profiles, M sells 1 item, and for others, M sells 2 items. (Otherwise, M is clearly dominated by a mechanism with bid-independent supply limit, and we are done). Let M' be one of the two mechanisms with bid-independent supply limit that picks $g = 2$, that is the mechanism sells at most 2 items, to the 2 highest bidders at some arbitrary order (charging the third highest bid). Assume that there are only 2 bidders. We first note that since M guarantees a $1/\alpha$ approximation to social welfare over D, we have: $\alpha \leq \beta$ (see the proof of lemma 3.4 for the argument). After scaling, say that the bids are as follows: $v_1 = 1 \geq v_2 = x$. Therefore, we have $\mathbf{OPT} = 1 + \beta \cdot x$. On the other hand, M' achieves welfare of at least

$$x + \beta \cdot 1 \geq x + \alpha.$$

Therefore, the approximation achieved by M' is at most:

$$\frac{(1 + \beta \cdot x)}{x + \alpha} \leq \frac{1}{\alpha}.$$

□

Proof of Proposition 4.8. We first prove that no deterministic mechanism can achieve better than a ϕ approximation to social welfare by exhibiting a hard distribution over two items. By the above lemma, we can restrict our attention to mechanisms with bid-independent supply limit. Let $t = 1/\phi$. We construct an instance with two bidders with values $v_1 = 1 \geq v_2 = x$. Let D be the distribution over 2 items such that $\Pr_D[\ell = 1] = 1 - t$ and $\Pr_D[\ell = 2] = t$. Note that $\mathbf{OPT} = 1 + t \cdot x$. Let M_1 be the mechanism which always sells 1 item (without loss of generality, to the highest bidder). Then we have that the expected welfare of M_1 is 1 (since it always sells an item to bidder 1). Let M_2 be one of the mechanisms with bid-independent supply limit with $g = 2$ (a mechanism which always sells as many items that arrive). Without loss of generality, M_2 always sells the first item to the lowest bidder. This is because we can rescale bids, letting $v_1 = 1$, $v_2 = 1/x$, which gives us an identical instance in which the only change is the bid of bidder 2: but by lemma 3.2, this does not change which item the second bidder receives. Therefore, the expected welfare of M_2 is $(1-t)x + t(1+x) = t + x$. If we face M_1, we let $x = 1$ and achieve approximation ratio $(1 + t) = \phi$. If we face M_2 we let $x = 0$ and achieve approximation ratio $1/t = \phi$.

We now prove that no randomized mechanism can achieve better than a $4/(2+\sqrt{2})$ approximation to social welfare. We again consider a distribution D over two items. Let $\Pr[\ell = 1] = 1-\beta$ and $\Pr[\ell = 2] = \beta$. We exhibit a distribution over vectors of values such that no truthful deterministic mechanism can achieve better than the stated approximation ratio in expectation. By the min-max principle, this gives a lower bound for truthful randomized mechanisms. Our distribution over vectors of values is as follows: With probability α, we have $v_1 = 1, v_2 = x$ with $x \leq 1$. With probability $1-\alpha$ we have $v_1 = 1, v_2 = 1/x$. As argued above, because these two bid instances differ only in the bid of a single bidder, any deterministic truthful mechanism must allocate items the same way given both instances. We have $\mathbf{OPT} = \alpha(1+\beta x) + (1-\alpha)(1/x + \beta)$. Let M_1 be a mechanism that happens only to sell a single item (without loss of generality to the highest bidder). Then M_1 achieves expected welfare $R_1 \leq \alpha + (1-\alpha)/x$. As before, let M_2 be one of mechanisms with bid-independent supply limit with $g = 2$. Since M_2 will sell the first item to the lowest bidder in at least one instantiation of the bid values, we have that M_2 achieves expected welfare at most:

$$R_2 \leq \max[\alpha(1+\beta x) + (1-\alpha)(1+\beta/x), \alpha(x+\beta) + (1-\alpha)(1/x+\beta)]$$

To prove a lower bound, we wish to maximize $\mathbf{OPT}/\max[R_1, R_2]$. Optimizing this quantity, we get $\alpha = 1/\sqrt{2}, \beta = x = \sqrt{2}-1$ which achieves the stated bound. □

Theorem 4.9: *There exists a regular distribution (that does not satisfy the non-decreasing hazard rate condition) such that no deterministic truthful mechanism can achieve an $o(\sqrt{\log n/\log\log n})$ approximation to social welfare when faced with stochastic supply drawn from this distribution.*

Theorem 4.9 follows directly from the following lemma and Lemma 3.4.

Lemma A.4. *There exists a regular distribution (that does not satisfy the non-decreasing hazard rate condition) such that no deterministic mechanism with bid-independent supply limit can achieve an $o(\log n/\log\log n)$ approximation to social welfare when faced with stochastic supply drawn from this distribution.*

Proof. We give a regular 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 $\Pr[\ell = i] = 1/(i+i^2)$ for $i < n$ and $\Pr[\ell = n] = 1/n$. 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 mechanism with bid-independent supply limit 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. □

A.3 Proofs of Lemmas from Section 5

Lemma 5.4: 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:

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

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

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 $i \in \{0, 1, \ldots, \log n - 1\}$. We consider the inverse CDF function $F^{-1}(x) : [0,1] \rightarrow \{1, 1/2, 1/4, \ldots, 2/n\}$. 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 $[0, 1]$: $E[X_{i,n}] = 1 - i/(n+1)$. This standard fact follows from a simple symmetry argument. We are now ready to complete the proof of the lemma:

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

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

Lemma 5.5: 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 $f(c_{b_k+1},\ldots,c_n)) \equiv \sum_{i=b_k+1}^{n} \frac{(c_i+1)}{\exp(\frac{\sum_{j=b_k+1}^{i-1} 2^{c_j}}{n-1})}$. We consider the partial derivative at the i'th offer price:

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

But this is negative unless

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

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:

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

□

Proposition A.5 (Proposition 5.6)**.** *RandomGuess is truthful, charges all winners the same price, and achieves a* $\log n$ *approximation to social welfare.*

Proof. Truthfulness and the fact that the mechanism charges all winners the same price 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 $\mathbf{OPT}_g/2 \geq OPT_I/2$. 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 of at least $(1/\log n)(\mathbf{OPT}_I/2 + \mathbf{OPT}_I/2) = \mathbf{OPT}_I/\log n$. □

A.4 Proofs from Section 5.3

Proposition 5.7: *No truthful mechanism that charges all winners the same price can achieve an* $o(\log n/\log\log n)$ *approximation to social welfare when faced with adversarial supply.*

Proof. For a mechanism that charges all winners the same price, we may assume that all offered prices $c_1, \ldots, c_n$ are equal: for all i, $c_i = c$. We apply inequality 5 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:

$$n \cdot c \geq \frac{(n-1)\log n}{2\alpha} - n \tag{17}$$

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:

$$(c+1) \cdot \sum_{i=1}^{n} \left(\frac{n - 2^{c+1}}{n-1}\right)^{i-1} \geq \frac{n-1}{2\alpha} \tag{18}$$

Setting $\alpha = o(\log n/\log\log n)$, we see that Constraint (17) requires $c = \omega(\log\log n)$. It is simple to verify that the left hand side of Constraint (18) is decreasing in c in the range $[\log\log n, \log(n) - 1]$, and that setting $c = \omega(\log\log n)$ fails to satisfy Constraint (18), which proves the claim. □