A Field Guide to Personalized Reserve Prices

Renato Paes Leme
Google NYC
renatoppl@google.com

Martin Pál
Google NYC
mpal@google.com

Sergei Vassilvitskii
Google NYC
sergeiv@google.com

ABSTRACT

We study the question of setting and testing reserve prices in single item auctions when the bidders are not identical. At a high level, there are two generalizations of the standard second price auction: in the *lazy* version we first determine the winner, and then apply reserve prices; in the *eager* version we first discard the bidders not meeting their reserves, and then determine the winner among the rest. We show that the two versions have dramatically different properties: lazy reserves are easy to optimize, and A/B test in production, whereas eager reserves always lead to higher welfare, but their optimization is NP-complete, and naive A/B testing will lead to incorrect conclusions. Despite their different characteristics, we show that the overall revenue for the two scenarios is always within a factor of 2 of each other, even in the presence of correlated bids. Moreover, we prove that the eager auction dominates the lazy auction on revenue whenever the bidders are independent or symmetric. We complement our theoretical results with simulations on real world data that show that even suboptimally set eager reserve prices are preferred from a revenue standpoint.

1. INTRODUCTION

A key part of auctions is setting the minimum price at which the seller is willing to part with the item. The so called reservation, or, *reserve*, price is critical to maximizing revenue, as proven by Myerson in his Nobel prize winning work [Mye81]. In the online advertising scenario, setting the reserve price is a non-trivial exercise—the auctions are repeated, hence agents may be adapting their behavior to influence the learning [ARS13, MM14, CGM15], only a glimpse into the buyers' valuations is known [DRY15], and there is a heterogeneity in the sophistication level of the bidders. On the other hand, as observed by Celis et al. [CLMN14], many of the auctions have very few bidders (the median number of bidders is six in their dataset),

thereby exacerbating the need for good reserve prices to maximize revenue.

The fact that reserve prices are a good idea comes from Myerson's seminal work [Mye81] which shows that if the valuations are drawn independently and identically (iid) from a distribution satisfying a certain regularity assumption, then the optimal auction takes the form of a second price auction with a reserve price. Myerson's result generalizes past the iid setting: if agents are independent but not identical, the optimal auction involves sorting the agents by a function of their bid (called the *virtual value*) and assigning the item to the agent with largest non-negative virtual value.

In practice, the optimal auction is complicated to implement, since it involves learning the distributions of valuations in a robust enough manner to allow the computation of virtual values. And even if this computation were feasible, in the spirit of Wilson doctrine [Wil85], simple detail-free mechanisms are often preferred in complex environments in order to mitigate the risks introduced by assumptions of the model. In the spirit of designing simpler auctions for the non-iid setting, Hartline and Roughgarden [HR09] show that if players are independent and their distribution obey the monotone hazard rate condition, then there exist a vector of personalized reserve prices that generates revenue which is at least half of the optimal revenue. Both assumptions (independence and monotone-hazard-rate) are necessary for their result.

In this work, we revisit the topic of setting personalized reserve prices in second price auctions from the perspective of a practitioner: How to compute them? How to make sure they are computed correctly? How to apply them in an efficient manner? While asking this question, we avoid (in most of our results) making any assumptions about the shape of the bid distributions.

Evaluation. In practice computing a reserve price that behaves well in offline simulations is only the first part of the process. A successful experimental evaluation is key in proving that assumptions made in theory are reasonable, and do not lead to unintended consequences.

Testing reserve prices is a separate research challenge in and of itself. We distinguish between short term studies, where the goal is to measure the immediate benefit of the reserve prices, and long term studies that intend to capture the strategic interactions between the bidders and auctioneer. Here we focus on the short term studies, and observe that testing personalized reserve prices, even in this, relatively simple, setting, is non-trivial, and can lead to incorrect conclusions.

1.1 Our Contributions

We show that the problem of computing and applying personalized reserve prices is nuanced, and has paradoxical behavior, particularly when testing the efficacy of reserves in an A/B test. Our empirical evaluations show that the paradoxical behavior is not limited to theory, but does occur in practice as well.

We begin by describing two different approaches to applying personalized reserve prices: *lazy* and *eager*, and show that while sometimes one approach dominates another on revenue, they are always within a factor of two of each other (Theorem 4.1).

We then identify two mildly restrictive settings (those of symmetric bidders, and independent bidders) and show that in these situations optimal eager reserve prices always yield more revenue (Theorems 4.4 and 4.5).

Turning to computational issues in setting the optimal lazy and eager reserve prices, we show that a simple nearly linear time algorithm can compute the optimal lazy reserves from previous history, while computing the optimal eager reserves is NP-hard (Propositions 5.1 and 5.2).

We show that naive A/B testing of eager reserve prices *always* leads to a drop in revenue, even when the reserve prices are set correctly! (Theorem 6.1) We observe that lazy reserves do not suffer from this problem, and behave in an intuitive manner.

Finally, we present an empirical evaluation of our findings on real world data, and show that the performance of the algorithms is much better than the large approximation factors guaranteed by the theory (Section 7).

2. PRELIMINARIES

We consider the standard setting of single item auctions. Let $A = \{1, 2, \ldots, n\}$ be the set of agents interested in the item. Each agent i has a value v_i for the item, and submits a bid b_i to the auctioneer. Given a vector of bids $\mathbf{b} = (b_1, b_2, \ldots, b_n)$, we denote by $b^{(1)}$ the highest bid, and $b^{(2)}$ the second highest bid. We will denote by $\mathbf{r} = (r_1, r_2, \ldots, r_n)$ the vector of personalized reserve prices.

We assume that the valuation of each agent is drawn independently from an unknown distribution with CDF F and PDF f. To ease the exposition we assume that the distributions F are regular, in that the virtual value function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing.

In the case when all of the distributions are identical and known to the seller, $F_1 = F_2 = \ldots = F_n$, Myerson proved the following characterization of the optimal auction:

1. Collect bids $b_1, b_2, \ldots, b_n$.

2. Discard all bids that are below $\phi^{-1}(0)$.

3. Allocate the item to the agent with the highest bid, and charge her the maximum of $\phi^{-1}(0)$ and $b_{(2)}$.

Note that $\phi^{-1}(0)$ acts as a reserve price for the auction: this is the minimum bid any agent must submit to win, and also acts as a lower bound on the revenue to the seller.

In this work our focus is on the non-identical setting. In this case Myerson proved that the optimal auction is:

1. Collect bids $b_1, b_2, \ldots, b_n$.

2. For each agent, discard bid b_i if it is lower than $\phi_i^{-1}(0)$.

3. Allocate the item to the agent with the highest virtual value, $\phi_i(b_i)$, and charge her the maximum of $\phi_i^{-1}(0)$ and $\phi_i^{-1}(\phi_j(b_j))$, where j is the agent with the second highest virtual value.

There are two major differences from the identical setting. First, instead of having a universal reserve price $\phi^{-1}(0)$, each agent now has a personalized reserve price $\phi_i^{-1}(0)$. Second, the winner is determined as the agent whose bid has the highest virtual value, not the one who has the highest bid.

While the first auction is easily implemented in practice, the latter is much more problematic. First, the virtual value functions depend critically on the value distributions, which themselves are not always known and must be estimated. Second, the auction is counterintuitive to outsiders, as the agent with the highest bid does not always win the item.

2.1 Personalized Reserve Prices

To combat the potential bid inversion that comes with ordering by virtual values, Hartline and Roughgarden [HR09] proposed keeping the personalized reserve prices aspect of the optimal auction, but ordering items by bid instead. They show that with monopoly reserves $\phi_i^{-1}(0)$, this auction yields a 2-approximation to the revenue of the optimal auction when agents are independent and follow the monotone hazard rate condition.

This natural approach leads to two possible flavors of second price auctions, which were first introduced by Dhangwatnotai et al. [DRY15]. Informally, in the *eager* regime, we first discard all of the bids that do not meet their personalized reserve prices and then run the second price auction on the remaining bids. In the *lazy* regime, we always try to allocate the item to the agent with highest bid. If her bid is below the reserve price the good goes unallocated, otherwise the agent is charged the maximum between her reserve price and the second highest bid. We describe these two auctions more formally in Section 3.

2.2 Estimating Value Distributions

Myerson's theoretical analysis requires us to know the cumulative density function F, as well as the PDF, f, to compute the virtual value function $\phi(\cdot)$. In practice this is a tall order. In the online advertising context, the auctions are repeated, and thus we can observe a number of draws from the distribution. For example, we observe the sequence of bids for agent i: $b_i^1, b_i^2, b_i^3, \ldots$, from which we can compute empirical estimates $\hat{F}$ and $\hat{f}$ for F and f.

In this work we focus on the computational complexity of computing optimal personalized reserve prices from the previous bids. In particular we show that the optimization question for eager and lazy auctions has very different profiles, one being solvable in polynomial time, and the other being NP-complete.

3. LAZY AND EAGER AUCTIONS

Insisting that advertisers are ranked by bid leads to two flavors of second price auctions. As before, let $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ be the bids submitted to the auctioneer, and assume without loss of generality that $b_1 \geq b_2 \geq b_3 \ldots b_n$. Let $\mathbf{r} = (r_1, r_2, \ldots, r_n)$ be the vector of reserve prices, with reserve price r_i applying to bidder i.

Following the work of Dhangwatnotai, Roughgarden and Yan [DRY15], we define second price auctions with *lazy* and *eager* reserves.

Lazy Reserves:

- If $b_1 < r_1$ then there is no winner and the item goes unsold.

- if $b_1 \geq r_1$, allocate the item to bidder 1 and charge her $\max(r_1, b_2)$.

Eager Reserves:

- Let $S = \{i : b_i \geq r_i\}$ be the set of bidders who bid above reserve.

- Let j be bidder with the highest bid in S (ties broken by the original ordering).

- Allocate the item to bidder j and charge her the maximum of her reserve and second highest bid in S: $\max(r_j, \max_{i \in S - \{j\}} b_i)$.

Observe that when the reserve prices are identical, $r_1 = r_2 = \ldots = r_n$, then both of these approaches implement the standard second price auction. We denote by $\mathrm{REV}_L(\mathbf{b}; \mathbf{r})$ and $\mathrm{REV}_E(\mathbf{b}; \mathbf{r})$ the revenue obtained by running the lazy and eager auctions on the same set of bids and reserve prices.

Both versions of the auction are incentive compatible and individually rational, since, whenever player i wins with bid b_i, she still wins with all bids $b_i' \geq b_i$. Moreover, in either case, the payment of i correspond to her critical bid, i.e., the infimum of the bids for which he wins.

3.1 Example

To demonstrate the difference between the two approaches and show why in general the revenues are incomparable, consider the following two examples.

EXAMPLE 3.1 (EAGER DOMINATES LAZY). *Consider an auction with three bidders A, B, and C, who bid 7, 5, and 3 respectively. Suppose the vector of reserve prices is 8, 1 and 2. The auction with lazy reserves tries to allocate the item to A since she has the highest bid. However, since her bid is lower than her reserve price, the item goes unallocated, and the seller collects no revenue. The auction with eager reserves first filters out A, allocates the item to B and collects revenue of 3.*

EXAMPLE 3.2 (LAZY DOMINATES EAGER). *Let A, B, and C be the three bidders bidding 7, 5, and 3. Suppose the vector of reserve prices is 2, 6 and 1. The auction with lazy reserves allocates the item to A since her bid is above her reserve and charges her 5 (the second highest bid). On the other hand, the auction with eager reserves first removes B from consideration (since his bid is below reserve). A still wins the item, but is charged only 3.*

Intuitively, eager reserve prices lead to higher revenue when the highest bid is priced out; whereas lazy reserve prices lead to higher revenue when the second highest bidder is priced out.

4. COMPARING LAZY AND EAGER RESERVE PRICES

The second price auctions with eager and lazy reserves are identical when all agents are subject to the same reserve price. If the reserve prices are personalized, however, the outcomes of the auctions can be very different. In terms of welfare, it is easy to see that the auction with eager reserves always dominates the auction with lazy reserves. The auction with lazy reserves allocates the item only if the highest bidder is above her reserve price. In such cases, the auction with eager reserves also allocates the item. This fact alone makes a second price auction with eager reserves more attractive to sellers that care about match rate.

As we saw in examples in Section 3, in general the revenue gains due to lazy and eager reserve prices are incomparable. In this section we provide a tighter characterization, and identify broad classes where one mechanism dominates another.

We begin with Theorem 4.1 and prove that for any (possibly correlated) bid distribution no auction generates more than twice the revenue of the other, and give two examples to show this bound is asymptotically tight.

Then, in Section 4.2 we show that the revenue of the optimal eager mechanism dominates that of the optimal lazy mechanism whenever either (i) the bidders are symmetric (the joint bid distribution doesn't change when bidders are permuted), or, (ii) the bidders' bids are drawn from independent (not necessarily identical) distributions.

THEOREM 4.1. *Let $\mathcal{D}$ be any distribution over bid vectors $\mathbf{b}$. Also, let $\mathrm{REV}_L(\mathbf{b}; \mathbf{r})$ and $\mathrm{REV}_E(\mathbf{b}; \mathbf{r})$ denote the revenue of the lazy and eager auctions under bid vectors $\mathbf{b}$ and personalized reserves $\mathbf{r}$. Let:*

$$\mathbf{r}_L^* = \arg\max_{\mathbf{r}} \mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_L(\mathbf{b}, \mathbf{r})],$$

$$\mathbf{r}_E^* = \arg\max_{\mathbf{r}} \mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_L(\mathbf{b}, \mathbf{r})].$$

Then

$$\mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_L(\mathbf{b}, \mathbf{r}_L^*)] \leq 2 \cdot \mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_E(\mathbf{b}, \mathbf{r}_E^*)],$$

$$\mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_E(\mathbf{b}, \mathbf{r}_E^*)] \leq 2 \cdot \mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}_L(\mathbf{b}, \mathbf{r}_L^*)].$$

PROOF. For the first inequality, consider running the eager second price auction with $\mathbf{r}_L^*$. For any bid vector $\mathbf{b}$, if the highest player is below the reserve then: $\mathrm{REV}_L(\mathbf{b}, \mathbf{r}) = 0 \leq \mathrm{REV}_E(\mathbf{b}, \mathbf{r})$. If the highest player is above the reserve and her payment in the lazy auction is her reserve, then both auctions generate the same revenue. The final case is the case where the highest bidder is above the reserve and the payment in the lazy auction is the second highest bid. In this case, the revenue of the eager auction might be lower if the second highest bidder is below her reserve. For that bid vector, however, the revenue of the lazy auction is the second highest bid, which is equal to the revenue of the auction with no reserve prices. Therefore:

$$\mathrm{REV}_L(\mathbf{b}, \mathbf{r}_L^*) \leq \max[\mathrm{REV}_E(\mathbf{b}, \mathbf{r}_L^*), \mathrm{REV}_E(\mathbf{b}, \mathbf{0})]$$
$$\leq \mathrm{REV}_E(\mathbf{b}, \mathbf{r}_L^*) + \mathrm{REV}_E(\mathbf{b}, \mathbf{0})$$

Taking expectations over $\mathbf{b}$ we obtain the first inequality:

$$\mathsf{E}[\mathrm{REV}_L(\mathbf{b}, \mathbf{r}_L^*)] \leq \mathsf{E}[\mathrm{REV}_E(\mathbf{b}, \mathbf{r}_L^*) + \mathrm{REV}_E(\mathbf{b}, \mathbf{0})]$$
$$\leq 2 \cdot \mathsf{E}[\mathrm{REV}_E(\mathbf{b}, \mathbf{r}_E^*)]$$

For the second inequality, consider running a lazy second price auction with $\mathbf{r}_E^*$. If for any bid vector $\mathbf{b}$, the highest player is above her reserve, the lazy auction is guaranteed to generate more revenue than the eager auction. If not, then the eager auction generates at most revenue equal to the bid

of the second highest bidder (since revenue is dominated by welfare), which is the revenue of the second price auction with no reserves. Therefore:

$$\text{REV}_E(\mathbf{b}, \mathbf{r}_E^*) \leq \max[\text{REV}_L(\mathbf{b}, \mathbf{r}_E^*), \text{REV}_L(\mathbf{b}, \mathbf{0})]$$
$$\leq \text{REV}_L(\mathbf{b}, \mathbf{r}_E^*) + \text{REV}_L(\mathbf{b}, \mathbf{0})$$

We obtain the second inequality by again taking expectations over $\mathbf{b}$:

$$\mathsf{E}[\text{REV}_E(\mathbf{b}, \mathbf{r}_E^*)] \leq \mathsf{E}[\text{REV}_L(\mathbf{b}, \mathbf{r}_E^*) + \text{REV}_L(\mathbf{b}, \mathbf{0})]$$
$$\leq 2 \cdot \mathsf{E}[\text{REV}_L(\mathbf{b}, \mathbf{r}_L^*)]$$

$\square$

4.1 Lower Bound Examples

The following examples complement the bounds in Theorem 4.1. First, we show an example where revenue from the eager mechanism is almost twice the revenue from the best lazy mechanism.

EXAMPLE 4.2. *There is an instance with n bidders with valuations drawn independently from the same distribution, where the best eager mechanism generates $2 - o(1)$ times more revenue than the best lazy mechanism.*

PROOF. Consider n identical bidders, with each bidder choosing to bid $b_i = n$ with probability $1/n^2$ and with probability $1 - 1/n^2$ bidding $b_i = 1$. (Each bid is drawn independently. To break ties, each bid is perturbed by adding a noise term drawn independently from $[0, \epsilon]$ for infinitesimally small ϵ)).

With probability $1 - 1/n + O(n^{-2})$, all bidder valuations will be 1, and with probability $1/n - O(n^{-2})$, some bidder will have a valuation of n. Thus, the optimal welfare given these valuations is $2 - O(n^{-2})$.

We claim that a lazy pricing mechanism extracts at most 1 revenue. By symmetry each buyer is equally likely to be the winner. Now consider the reserve price for an individual bidder i.

- If $r_i < 1 + \epsilon$, revenue will be bounded by $1 + \epsilon$ unless both bidder i and another bid bid high (probability of $O(n^{-2})$), in which case revenue is bounded by n. The expected revenue is then $1 + o(1)$.

- If $r_i > 1 + \epsilon$, revenue is only earned if $b_i = n$, and is bounded by n. Since the probability of a high bid is less than n^{-1}, the expected revenue is bounded by 1.

On the other hand, we claim that an eager mechanism can earn $2 - o(1)$ revenue. This can be done by imposing a high reserve price $r_i = n$ for $i = 2, \ldots, n$ (all buyers except one) and a low reserve price $r_1 = 1$ on the remaining bidder. With this setup:

- With probability $1/n - o(n^{-2})$, one of bidders $2, \ldots, n$ will submit a high bid, generating revenue $n - O(1)$.

- With the remaining probability, the auction is guaranteed to clear because $b_1 \geq r_1 = 1$ and will generate at least 1 in revenue.

The overall expected revenue is thus $\frac{1}{n} \cdot (n - o(1)) + (1 - \frac{1}{n}) \cdot 1 = 2 - o(1)$. $\square$

Next, we show an example where the lazy mechanism generates twice the revenue of the best eager mechanism. This example requires both that bidders are not symmetric, and that their bids are correlated.

EXAMPLE 4.3. *There is an instance with two correlated heterogenous bidders where a lazy mechanism generates $2 - o(1)$ times higher revenue than the best eager mechanism.*

PROOF. Let M be a sufficiently large constant, and consider the following joint distribution of bids of two bidders.

- With probability $\frac{\log M}{M}$, bidder 1 bids 0 while bidder 2 bids $b_2 = M$.

- With probability $1 - \frac{\log M}{M}$, first bidder's bid b_1 is drawn from a truncated equal revenue distribution $F(b) = 1 - \frac{1}{b}$ for $b \in [1, M)$ and $F(M) = 1$ and the second bidder's bid is set to $b_2 = (1 - \epsilon)b_1$ for some arbitrarily small ϵ.

Lazy reserve prices $r_1 = 0, r_2 = M$ extract the full surplus as $\epsilon \to 0$, achieving expected revenue of

$$\frac{\log M}{M} \cdot M + \left(1 - \frac{\log M}{M}\right)(1 - \epsilon)\left(\int_1^M bf(b)db + \frac{M}{M}\right)$$
$$= \log M + \left(1 - \frac{\log M}{M}\right)(1 - \epsilon) \cdot (\log M + 1)$$
$$= (2 - o(1))\log M.$$

Suppose the eager mechanism imposes reserves of r_1 and r_2. With probability $\log M/M$ the revenue is r_2. For the remaining case, we condition on whether the second bidder bid above reserve. In the case $b_2 < r_2$, bidder 1 is competing against his reserve price, r_1, generating revenue of at most:

$$r_1(1 - F(r_1)) \geq r_1\left[1 - \left(1 - \frac{1}{r_1}\right)\right] = 1.$$

If, on the other hand, when $b_2 \geq r_2$, the total revenue is bounded by:

$$\int_{r_2}^M bf(b)db = \log M - \log r_2 = \log(M/r_2)$$

Putting these together the overall revenue is at most :

$$r_2 \frac{\log M}{M} + \left(1 - \frac{\log M}{M}\right) \cdot \left(1 + \log \frac{M}{r_2}\right)$$
$$\leq 1 + r_2 \frac{\log M}{M} + \log M - \log r_2$$

This function is convex in r_2, thus achieving its maximum at the endpoints of the interval. It obtains its maximum value of $(1 + o(1))\log M$ at $r_2 = 1$. $\square$

4.2 Restricted Settings

Above we proved that while the revenue from lazy and eager auctions is always within a factor of two of each other, an unconditional bound is impossible in general. Here we consider two restricted settings, first of symmetric bidders, and then of independent bidders. In both cases we show that eager auctions dominate lazy auctions.

Symmetric Bidders. We say that bidders are symmetric (sometimes also called exchangeable) if the bid distribution is invariant under permutations. Formally, for every

permutation $\pi : [n] \to [n]$ let $\mathbf{b}^\pi$ be the vector $(b_{\pi(1)}, b_{\pi(2)}, \ldots, b_{\pi(n)})$. A distribution is symmetric if the distribution of $\mathbf{b}^\pi$ is the same as the distribution of $\mathbf{b}$. Notice that iid implies symmetry, but symmetry is more general—for example, consider the bid distribution obtained by choosing in each time one buyer at random and letting her bid H and letting every other buyer bid L. The distribution is clearly symmetric, but it is not independent (since there is always exactly one buyer bidding H).

THEOREM 4.4. *If bidders are symmetric the optimal eager mechanism yields at least as much revenue as the optimal lazy mechanism.*

PROOF. Proposition 5.1 tells us that the optimal lazy reserve price for bidder i is a function of the joint distribution of winning bids and prices conditioned on bidder i winning. If bidders are symmetric, the optimal lazy reserve prices are the same for all bidders; $r_1 = r_2 = \cdots = r_n$. If all reserve prices are the same, lazy and eager mechanisms behave identically. (It is however possible that a different reserve price vector yields higher revenue for the eager mechanism, as in Example 4.2.) □

Independent Bidders. We show that if each bidder draws her value independently, auctions with eager reserves are always at least as good auctions with lazy reserves.

THEOREM 4.5. *If bidder valuations are drawn independently (not necessarily from identical distributions), the optimal eager mechanism yields at least as much revenue as the optimal lazy mechanism.*

To prove the theorem we will exhibit a method that, given any vector $\mathbf{r}$ of lazy reserve prices, produces a vector $\mathbf{r}_E$ such that $\mathsf{E}[\mathrm{REV}_E(\mathbf{b}; \mathbf{r}_E)] \geq \mathsf{E}[\mathrm{REV}_L(\mathbf{b}; \mathbf{r})]$. The construction will rely heavily on the independence of the bids. A key concept in the proof is that of trimmed distributions:

DEFINITION 4.6 (TRIM). *Let X be a non-negative random variable distributed according to some distribution $\mathcal{D}$, and a real number $r \geq 0$. We use $\mathtt{trim}(\mathcal{D})$ to denote the distribution of the random variable $X' = X \cdot \mathbf{1}\{X \geq r\}$, i.e.. random variable that is equal to X if $X \geq R$ and is zero if $X < r$.*

This definition allows us to define the main lemma. In the following proofs, we will abbreviate $\mathsf{E}_{\mathbf{b} \sim \mathcal{D}}[\mathrm{REV}(\mathbf{b}, \mathbf{r})]$ by $\mathrm{REV}(\mathcal{D}, \mathbf{r})$.

LEMMA 4.7. *Given independent bid distributions $\mathcal{D}_i$, $\mathcal{D} = \mathcal{D}_1 \times \ldots \mathcal{D}_n$ and a reserve price vector $\mathbf{r}$, there exists a vector of reserve prices $\mathbf{r}'$ such that:*

$$\mathrm{REV}_L(\mathcal{D}, \mathbf{r}) \leq \mathrm{REV}_L(\mathcal{D}', \mathbf{r}')$$

where $\mathcal{D}' = \mathcal{D}'_1 \times \ldots \times \mathcal{D}'_n$ and $\mathcal{D}'_i = \mathtt{trim}(\mathcal{D}_i; r'_i)$.

First we show how to use Lemma 4.7 to prove Theorem 4.5:

PROOF OF THEOREM 4.5. Let $\mathbf{r}'$ and $\mathcal{D}'$ be as in Lemma 4.7. Notice that for $\mathbf{b} \sim \mathcal{D}'$, $\Pr[0 < b_i < r_i] = 0$. Since a bidder is never blocked by the reserve, the revenue in both the eager and lazy auctions is the maximum of the second highest bid and the reserve of the highest bidder. This implies that $\mathrm{REV}_E(\mathcal{D}', \mathbf{r}') = \mathrm{REV}_L(\mathcal{D}', \mathbf{r}')$. Since $\mathcal{D}$

stochastically dominates $\mathcal{D}'$, there is a distribution on pairs of vectors $(\mathbf{b}', \mathbf{b})$ such that the marginals are $\mathcal{D}'$ and $\mathcal{D}$ and $b'_i \leq b_i$ for all i for every realization of the random variables. Hence, $\mathrm{REV}_E(\mathbf{b}'; \mathbf{r}')] \leq \mathrm{REV}_E(\mathbf{b}; \mathbf{r}')$. Taking expectations we conclude that

$$\mathrm{REV}_E(\mathcal{D}', \mathbf{r}') \leq \mathrm{REV}_E(\mathcal{D}, \mathbf{r}').$$

Putting this together with the inequality from Lemma 4.7:

$$\mathrm{REV}_L(\mathcal{D}, \mathbf{r}) \leq \mathrm{REV}_L(\mathcal{D}', \mathbf{r}') = \mathrm{REV}_E(\mathcal{D}', \mathbf{r}') \leq \mathrm{REV}_E(\mathcal{D}, \mathbf{r}')$$

□

Now, all is left to do is to prove Lemma 4.7:

PROOF OF LEMMA 4.7. Assume that the bidders are sorted such that $r_1 \geq r_2 \geq \ldots \geq r_n$. We will define an algorithmic procedure that iterates through bidders 1 to n and at each iteration i, trims the distribution of the i-th bidder and possibly increases the reserves of agents $j > i$. It is useful to think of $\mathbf{r}$ and $\mathcal{D} = \mathcal{D}_1 \times \ldots \times D_n$ as variables that are updated in the course of the procedure.

The procedure will keep the following invariants: (i) $\mathrm{REV}_L(\mathcal{D}, \mathbf{r})$ cannot decrease; (ii) for all bidders already processed, their distribution is trimmed at their reserve, i.e., $\Pr[0 < b_i < r_i] = 0$. (iii) the reserve prices will continue to be sorted.

Now, we are ready to describe each iteration. When we process bidder i, we perform the following procedure:

Choose $x < r_i$ such that:

$$\mathsf{E}_{\mathcal{D}}[\mathrm{REV}_L(\mathbf{b}; \mathbf{r}) | b_i = x] \geq \mathsf{E}_{\mathcal{D}}[\mathrm{REV}_L(\mathbf{b}; \mathbf{r}) | b_i < r_i]$$

Set $\mathcal{D}_i = \mathtt{trim}(\mathcal{D}_i, r_i)$ and $r_j = \max(r_j, x)$ for $j > i$.

Clearly we maintain invariants (ii) and (iii). Now, we only need to argue that invariant (i) is also maintained. It is convenient to write $\mathrm{REV}_L(\mathbf{b}, \mathbf{r}) = \sum_j \mathrm{REV}_L^j(\mathbf{b}, \mathbf{r})$ where $\mathrm{REV}_L^j(\mathbf{b}, \mathbf{r})$ is the revenue obtained from bidder j.

First notice that $\mathsf{E}_{\mathcal{D}}[\mathrm{REV}_L^i(\mathbf{b}; \mathbf{r})] = \mathsf{E}_{\mathcal{D}'}[\mathrm{REV}_L^i(\mathbf{b}; \mathbf{r}')]$ since the lazy auction just extracts revenue from i when she is the highest bidder and above her reserve (and those events are unaffected by trimming). Also, in the lazy auction, the reserves on bidders other than the highest bidder do not affect the outcome.

For all other $j \neq i$, notice that $\mathsf{E}_{\mathcal{D}}[\mathrm{REV}_L^j(\mathbf{b}; \mathbf{r}) | b_i \geq r_i] = \mathsf{E}_{\mathcal{D}'}[\mathrm{REV}_L^j(\mathbf{b}; \mathbf{r}') | b_i \geq r_i]$. Conditioned on $b_i > r_i$ the only thing changing in the two scenarios is that the reserve price of j is now $r'_j = \max(r_j, x)$; since $x < r_i$, it cannot be binding as bidder i is bidding $b_i > r_i$. (Notice that for bidders $j < i$, their reserve price was already at least r_i, so $r'_j = r_j = \max(r_j, x)$ since $x < r_i < r_j$.)

Finally, $\mathsf{E}_{\mathcal{D}'}[\mathrm{REV}_L^j(\mathbf{b}; \mathbf{r}') | b_i < r_i] = E_{\mathcal{D}}[\mathrm{REV}_L^j(\mathbf{b}; \mathbf{r}) | b_i = x]$, since in the first case i must be bidding zero and j is subject to reserve $r_j = \max(r_j, x)$, but this reserve can be implemented by having i bid x.

Combining all of the expressions, we get:

$$\mathsf{E}_{\mathcal{D}'}[\mathrm{Rev}_L(\mathbf{b}, \mathbf{r}')]$$

$$= \mathsf{E}_{\mathcal{D}'}[\mathrm{Rev}_L^i(\mathbf{b}, \mathbf{r}')]$$
$$+ \sum_{j \neq i} \mathsf{E}_{\mathcal{D}'}[\mathrm{Rev}_L^j(\mathbf{b}, \mathbf{r}')|b_i \geq r_i] \Pr(b_i \geq r_i)$$
$$+ \sum_{j \neq i} \mathsf{E}_{\mathcal{D}'}[\mathrm{Rev}_L^j(\mathbf{b}, \mathbf{r}')|b_i < r_i] \Pr(b_i < r_i)$$

$$= \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^i(\mathbf{b}, \mathbf{r})]$$
$$+ \sum_{j \neq i} \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^j(\mathbf{b}, \mathbf{r})|b_i \geq r_i] \Pr(b_i \geq r_i)$$
$$+ \sum_{j \neq i} \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^j(\mathbf{b}, \mathbf{r})|b_i = x] \Pr(b_i < r_i)$$

Since $\mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^i(\mathbf{b}; \mathbf{r})|b_i < r_i] = 0$, x was picked such that

$$\sum_{j \neq i} \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^j(\mathbf{b}; \mathbf{r})|b_i = x] \geq \sum_{j \neq i} \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L^j(\mathbf{b}; \mathbf{r})|b_i < r_i]$$

Plugging this into the last expression, we get that:

$$\mathsf{E}_{\mathcal{D}'}[\mathrm{Rev}_L(\mathbf{b}, \mathbf{r}')] \geq \mathsf{E}_{\mathcal{D}}[\mathrm{Rev}_L(\mathbf{b}, \mathbf{r})]$$

□

5. COMPUTING OPTIMAL RESERVES

In this section we investigate the computational complexity of computing the optimum reserve prices in the eager and lazy settings. We assume that the input to the problem is given as a set of bids submitted to previous auctions. Since the number of auctions run daily is extremely large, it is important for this procedure to be linear, or nearly linear in the size of the input.

We prove that from a computational perspective, lazy and eager auctions are vastly different. Computing the optimal vector of lazy reserves $\mathbf{r}_L^*$ can be done in linear time. On the other hand, computing the optimal vector of eager reserves $\mathbf{r}_E^*$ is NP-hard.

PROPOSITION 5.1. *The optimal vector of lazy reserves $\mathbf{r}_L^*$ can be computed in nearly linear time in the size of input logs.*

PROOF. Given bids $\{b_{i,t}\}$, for each query t, let w_t correspond to the agent that would win if no reserve prices were set. For any given vector of reserves $\mathbf{r}$, either w_t wins query t or no agent wins. Let Q_i be the queries for which $w_t = i$. Then we can write the revenue for the vector $\mathbf{r}$ of reserves as:

$$\mathrm{Rev}_L(\mathbf{b}; \mathbf{r}) = \sum_i \sum_{t \in Q_i} \mathbf{1}\{b_t^{(1)} \geq r_i\} \cdot \max(r_i, b_t^{(2)})$$

where $b_t^{(1)}$ $b_t^{(2)}$ are respectively the highest and second highest bid for query t. The previous expression shows that for lazy reserves, the problem of computing reserve prices can be decoupled for every i. Also, the optimal reserve price should be of the form $b_t^{(2)}$ or $b_t^{(1)}$ for some t. If not, we can increase it to the next point and the revenue can only increase.

This observation gives an algorithm for computing the optimal vector of reserves with running time $O(\sum_i |Q_i|^2)$. In order to turn it into a nearly linear time algorithm, notice that if we sort the bids appropriately, we can compute the revenue for setting each $r_i = b_t^{(1)}$ and $r_i = b_t^{(2)}$ in constant time.

In order to do so, construct an array with all the bids $b_t^{(1)}$ and $b_t^{(2)}$ for $t \in Q_i$ and annotate each entry of weather it is a highest bid or a second highest bid. If $n_i = |Q_i|$, this is an array of $2n_i$ numbers. Sort the array by bids in increasing order, which takes $O(n_i \log n_i)$, and call its elements $r_1 \leq r_2 \leq \dots r_{2n_i}$. Assume for simplicity that all elements are distinct. Given a certain r_j, we can write the revenue associated with it $R(r_j) = \sum_{t \in Q_i} \mathbf{1}\{b_t^{(1)} \geq r_j\} \cdot \max(r_j, b_t^{(2)})$ as $r_j \cdot k_j + s_j$ where $s_j = \sum_{t \in Q_i} b_t^{(2)} \cdot \mathbf{1}\{r_j \leq b_t^{(2)}\}$.

If we show how to compute (s_{j+1}, k_{j+1}) from (s_j, k_j) in constant time, we have a nearly-linear time algorithm. Doing it is easy. This can be done in two cases:

- $r_j = b_t^{(1)}$ for some t. Therefore increasing the reserve past $b_t^{(1)}$ will make query t to be unallocated. All other queries are unaffected. Set $k_{j+1} = k_j - 1$ and $s_{j+1} = s_j$.

- $r_j = b_t^{(2)}$ for some t. Therefore increasing the reserve past $b_t^{(2)}$ will cause the reserve price to bind for query t, instead of the second highest bid. Update: $k_{j+1} = k_j + 1$, $s_{j+1} = s_j - r_j$.

Since we can reconstruct the revenue for each value in the array, we can choose the optimal vector of reserves in time $O(\sum_i |Q_i| \cdot \log |Q_i|)$. It is not hard to see that this algorithm still works if values of the array are repeated. □

PROPOSITION 5.2. *Computing the optimal vector of eager reserves $\mathbf{r}$ given the bids from a set of previous auctions is NP-hard.*

PROOF. We give a reduction to the independent set problem. Let L and H be constants that we will chose later. Given a graph $G = (V, E)$, we map the independent set problem on this graph to the following instance of the reserve price problem with eager reserves: consider $|V|$ agents and $|E| + |V|$ queries. For each edge $e = (u, v)$ consider queries where $b_u = L$, $b_v = L$ and all other agents bid zero. And for each node u, consider queries with $b_u = H$ and all other agents bid zero. We select L and H such that $L < H < 2L$.

Clearly, for the optimal vector of reserves, $r_i \in \{L, H\}$. Also, notice that in the optimal solution, the set of nodes $I = \{u \in V; r_u = H\}$ must form an independent set. Indeed, if there is an edge $e = (u, v)$ with $r_u = r_v = H$, then there is zero revenue from the queries corresponding to edges e. If we switch either u or v to have reserve L, then we gain L revenue from edge e and lose $H - L$ from node u. Since $L > H - L$, this is a profitable deviation.

The revenue associated with setting $u \in I$ to H and other nodes to L is given by

$$L \cdot (|E| + |V|) + (H - L) \cdot |I|$$

So the optimal vector of reserve prices would give a solution to the maximum independent set problem. □

Combining the proof of Theorem 4.1 and Proposition 5.1 we get a 2-approximation to the optimal revenue obtained by running an eager second price auction.

We remark that the algorithm proposed in Proposition 5.1 is different from the heuristic proposed by Hartline and Roughgarden in [HR09]: their heuristic consists in choosing the vector of monopoly reserve prices $\mathbf{r}_M$ such that $r_i^M = \arg\max_{r \geq 0} \sum_{t=1}^{T} r \cdot \mathbf{1}\{b_{i,t} \geq r\}$. They observe that while this is a very good choice when the valuations are independent and follow the monotone hazard rate condition, this can be arbitrarily far from the revenue of the optimal auction if either condition is violated. We complement the observation showing that this vector can also be arbitrarily far from the optimal revenue of the second price auction with personalized reserves (which is a weaker benchmark than the optimal auction).

CLAIM 5.3. *For every constant C, there is a distribution of bids such that* $\text{REV}(\mathbf{b}; \mathbf{r}^M) \leq \frac{1}{C}\text{REV}(\mathbf{b}; \mathbf{r}^*)$ *for both the eager and lazy auctions.*

PROOF. Consider 2 bidders and consider a distribution on bids that sets $b_1 = b_2 = 2^{k-1}$ with probability $\frac{1}{2^k}$ probability, for $k = 1, \ldots, K - 1$, and $b_1 = b_2 = 2^{K-1} + \epsilon$ with probability $\frac{2}{2^K}$. The monopoly reserve prices are $r_u^M = 2^{K-1} + \epsilon$ generating revenue $1 + \epsilon \cdot 2^{1-K}$. The optimal vector of reserves is zero for both eager and lazy reserves: $\text{REV}(\mathbf{b}; \mathbf{0}) = K + \epsilon \cdot 2^{1-K}$. □

6. A/B TESTING

In the previous sections we described auctions with eager and lazy reserve prices, and compared the differences in welfare, revenue, and computational complexity of the two approaches. The theoretical models are clean and elegant, and allow us to abstractly reason about the benefits of one approach over another; however they do not capture the messy realities of setting reserve prices in practice.

Before releasing a new model into the wild, a key step is measuring the impact of its change on a small sample of traffic. This too is a non-trivial step, and requires a lot of care both in setting up the experiments [KHS07] and effectively measuring long term, as well as, short term impacts [HOT15],

One way to test reserve prices is to partition all of the auctions into treatment and control, and then only apply reserves in treatment scenarios. However, such a test gives a biased estimate of the revenue lift—since the effect on an individual buyer is small (as the test applies only to a small fraction of the auctions), the buyer is unlikely to react strategically, for example changing her bidding behavior, or looking for alternative places to buy the impressions.

A different approach is to partition the bidders into treatment and control groups, and only apply reserve prices to the bidders in treatment. It is easy to see that applying reserve prices to only a fraction of the bidders will reduce the overall revenue gains, however since a single buyer is now subject to reserve prices on all of her auctions, she is more likely to update her behavior in response.

It is natural to expect that applying a vector of optimal reserves to a subset of the buyers would yield an improvement in revenue that would allow us to evaluate the impact of applying reserves to all of the buyers. Counterintuitively, we show that even in the simplest possible setting (iid buyers with regular distributions), applying the Myerson reserve

price to a subset of the bidders is *worse than not applying any reserve prices at all*! In other words, as we add bidders to the treatment group, the total revenue *decreases*, and it is only when all of the bidders are treated that we realize the revenue gains. As we show in Section 7, this phenomenon is not a purely theoretical construct, but is also not uncommon in practice.

We state the results formally.

THEOREM 6.1. *Assume agents are iid with regular distribution F and let $\text{REV}_E(k)$ be the revenue obtained from applying the Myerson reserve price eagerly to k out of n agents and applying no reserve to the remaining agents. Then:*

$$\text{REV}_E(0) \geq \text{REV}_E(1) \geq \text{REV}_E(2) \geq \ldots \geq \text{REV}_E(n-1)$$

PROOF. Since agents are iid we can sample agents according to the following procedure: draw n iid samples from F and take a random assignment from them to agents. This is equivalent to drawing the valuation v_i for each agent and then choosing k out of n at random players to apply reserve prices. We will denote by $v^{(t)}$ the t-th largest bid.

We consider two scenarios: (i) the highest bidder is above the reserve. In this case, she will be allocated the good. (ii) the highest bidder is below the reserve. In this case, the agent to whom we allocate the good is the highest agent for whom no reserve price is applied.

Let $\text{MAXR}_t(z_1, z_2, \ldots, z_n)$ be the expected maximum of a uniformly random subset of t elements drawn from $\{z_1, z_2, \ldots, z_n\}$. Notice that for any vector $\mathbf{z}$, the function $\text{MAXR}_t(z_1, z_2, \ldots, z_n)$ is monotone non-decreasing in t. Second, observe that if z_i are drawn iid, then:

$$\mathsf{E}[\max\{z_1, \ldots, z_t\}] = \text{MAXR}_t(z_1, z_2, \ldots, z_n).$$

For $k < n$, by the Myerson Lemma:

$$\text{REV}_E(k) = \mathsf{E}[\phi(v^{(1)}) \cdot \mathbf{1}\{v^{(1)} \geq r\}] +$$
$$\mathsf{E}[\text{MAXR}_{n-k}(\phi(v_1), \ldots, \phi(v_n))|v^{(1)} < r] \cdot \Pr(v^{(1)} < r) \tag{$\star$}$$

We finish by noticing that when $v_i < r$ then $\phi(v_i) < 0$. Therefore the second term above is negative and nonincreasing in k by the monotonicity of the MAXR operator. For $k = n$, however, the second term disappears and we recover the optimal auction. □

We now show that the auction with lazy reserve prices does not suffer from this kind of paradoxical behavior.

THEOREM 6.2. *Assume agents are iid with regular distribution F and let $\text{REV}_L(k)$ be the revenue obtained from lazily applying the Myerson reserve price to k out of n agents and applying no reserve to the remaining agents. Then*

$$\text{REV}_L(k) = \tfrac{k}{n}\text{REV}_L(n) + \left(1 - \tfrac{k}{n}\right)\text{REV}_L(0)$$

PROOF. As in the previous theorem, we sample the agents by drawing n iid samples from F and taking a random assignment from them to agents.

In the second price auction with lazy reserves, we will choose the agent with bid $v^{(1)}$ and declare her as the winner if $v^{(1)}$ is at least the reserve price, which is r with probability k/n and 0 with the remaining probability. By Myerson's lemma:

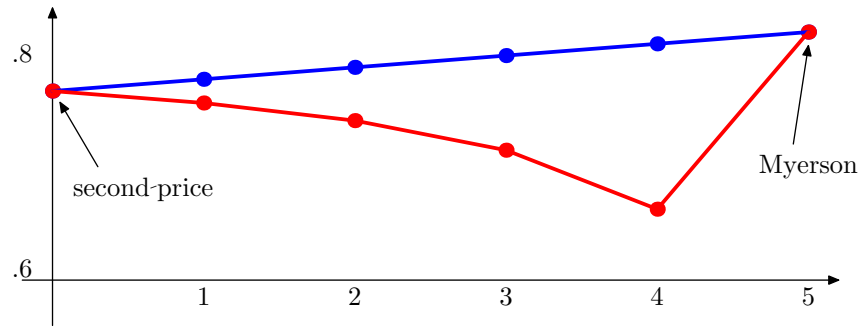$$\text{REV}_L(k) = \mathsf{E}[\phi(v^{(1)}) \cdot \mathbf{1}\{v^{(1)} \geq r^{(1)}\}]$$

Figure 1: $\mathbf{Rev}_E(k)$ in red and $\mathbf{Rev}_L(k)$ in blue for $n = 5$ agents with iid uniform distributions

where $r^{(t)}$ denotes the personalized reserve of agent with the t-th highest bid. Therefore:

$$\text{Rev}_L(k)$$
$$= \left(1 - \tfrac{k}{n}\right)\mathsf{E}[\phi(v^{(1)})] + \left(1 - \tfrac{k}{n}\right)\mathsf{E}[\phi(v^{(1)}) \cdot \mathbf{1}\{v^{(1)} \geq r\}]$$
$$= \left(1 - \tfrac{k}{n}\right)\text{Rev}_L(0) + \left(1 - \tfrac{k}{n}\right)\text{Rev}_L(n)$$

$\square$

To better understand the detrimental effect of eager reserve prices, it's useful to write equation $(\star)$ in explicit form so that we can evaluate the impact of applying reserves to a subset of buyers for particular distributions:

LEMMA 6.3. *In the setting of Theorem 6.1:*

$$\text{Rev}_E(k) = \int_r^1 \phi(x) \cdot nF(x)^{n-1}f(x)dx$$
$$- \mathbf{1}\{k < n\}F^n(r)\int_0^r \left(\frac{F(x)}{F(r)}\right)^{n-k}\phi'(x)dx$$

PROOF. For the first term, notice that the distribution of $v^{(1)}$ is given by $nF(x)^{n-1}f(x)$ since $\Pr(v^{(1)} \leq x) = \prod \Pr(v_i \leq x) = F(x)^n$. For the second term, since the agents are iid, the maximum of $n-k$ randomly chosen agents from among n buyers is identical to the maximum of $n-k$ buyers. By the principle of deferred decisions, we can first sample $n - k$ buyers and then draw their values.

Since in the second expression we condition their value to be at most r, we can simply compute the maximum over $n - k$ buyers with density $f(x)/F(r)$ for $0 \leq x \leq r$. By the same argument as before the density of the maximum is $(n - k) \cdot \frac{f(x)}{F(r)} \cdot \left(\frac{F(x)}{F(r)}\right)^{n-k-1}$. Using this fact, We can write the second term as:

$$F(r)^n \cdot \int_0^r \left[(n - k) \cdot \frac{f(x)}{F(r)} \cdot \left(\frac{F(x)}{F(r)}\right)^{n-k-1}\right]\phi(x)dx$$
$$= \mathbf{1}\{k < n\} \cdot F^n(r)\left\{\phi(x) \cdot \left(\frac{F(x)}{F(r)}\right)^{n-k}\Big|_0^r\right.$$
$$\left. - \int_0^r \left(\frac{F(x)}{F(r)}\right)^{n-k}\phi'(x)dx\right\}$$
$$= -\mathbf{1}\{k < n\} \cdot F^n(r)\int_0^r \left(\frac{F(x)}{F(r)}\right)^{n-k}\phi'(x)dx$$

$\square$

6.1 Case study: Uniform distribution

To get some intuition about Theorem 6.1 we look at the setting with n iid bidders distributed according to the $[0, 1]$-uniform distribution, for which $\phi(x) = 2x - 1$ and $r = 1/2$:

$$\text{Rev}_E(k) = \int_{1/2}^1 (2x - 1)nx^{n-1}dx$$
$$- \mathbf{1}\{k < n\} \cdot \left(\frac{1}{2}\right)^n \cdot \int_0^{1/2} (2x)^{n-k} \cdot 2dx$$
$$= \frac{n + 2^{-n} - 1}{n + 1} - \mathbf{1}\{k < n\} \cdot \frac{2^{-n}}{n - k + 1}$$

In Figure 5 we plot the revenue values for $\text{Rev}_E(k)$ and $\text{Rev}_L(k)$ with $n = 5$.

7. EXPERIMENTAL RESULTS

In the previous sections we discussed versions of the second price auction with personalized reserve prices: the eager (E) and lazy (L) auction, and proved that while the two notions of reserve prices are incomparable, eager reserves dominate lazy reserve, except for non-identical and non-independent settings, and in all of the settings total revenue of one is always within a factor of two of another. Further, we showed that applying the eager reserve prices to a *subset* of the bidders is non monotonic, and applying reserves to more bidders may (in theory) lead to lower revenues. In this section we validate these findings by simulating the effect of reserve prices on real world data.

Data. We collect bids sent to a large advertising exchange over the course of part of a day, and then restrict our attention to five ad slots with the highest traffic volume. Each of the ad slots has the bids submitted for hundreds of thousands of auctions. We report the results for each of the ad slots individually.

7.1 Non-monotonicity

Recall that in Section 6 we considered the standard setting where each bidder i draws an independent bid from bid distribution F_i. In the case of eager reserve prices, applying reserves to only a subset of the bidders leads to a *decrease* in the overall revenue to the auctioneer.

In this section we show that this scenario is not restricted to theory. We consider a set of auctions and compute the optimal lazy reserve price for each bidder. We then apply this reserve price in both lazy and eager fashions to a subset
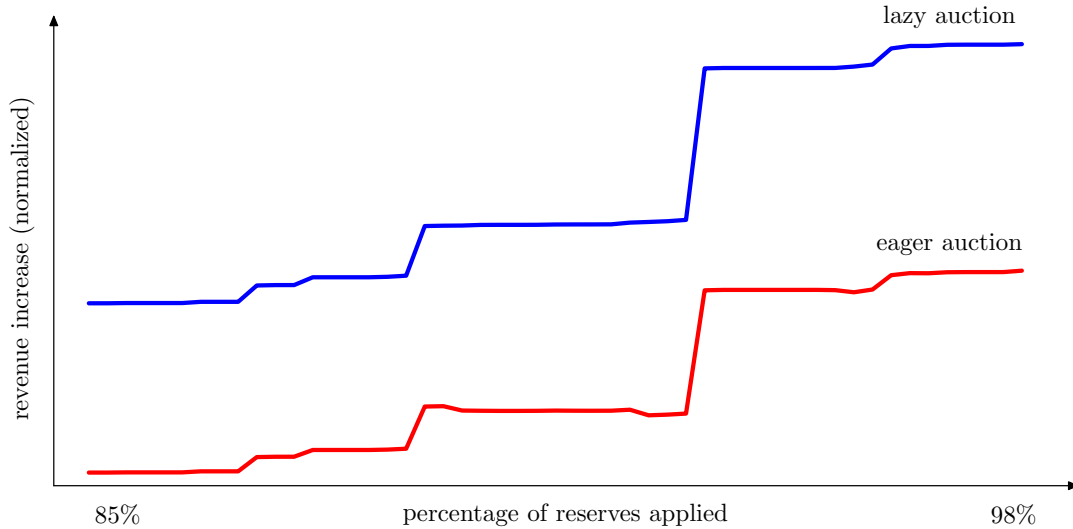
Figure 2: Revenue gain obtained by applying the optimal lazy reserve prices to a percentage of the buyers in the auction. Note that the revenue of the lazy auction is monotone non-decreasing, while the eager auction doesn't have this property.

of the bidders and plot the overall revenue. The results are shown in Figure 6.1.

We confirm our theoretical findings: while the revenue in the auction with the lazy reserve pricing strategy is monotonically increasing in the number of bidders subject to the reserve, the same cannot be said for the greedy reserve pricing strategy. The overall trend is positive, but it is non monotonic, and at times increasing the *number* of bidders subject to the optimal reserve price decreases the overall revenue.

7.2 Revenue Gains

Our second set of experiments addresses the question whether eager or lazy reserve prices lead to higher revenues in practice. For the best comparison, we simulate the auctions under both lazy and eager strategies while applying the optimal lazy reserve prices $\mathbf{r}_L^*$ (see Proposition 5.1) and the monopoly reserve prices $\mathbf{r}_M$ (see the discussion following the proof of Proposition 5.2).

We compute the revenue lifts due to personalized reserves under four strategies:

$$\Delta_L(\mathbf{r}_L^*) = \mathsf{E}[\text{REV}_L(\mathbf{b}; \mathbf{r}_L^*)] - \mathsf{E}[\text{REV}(\mathbf{b}; \mathbf{0})]$$

$$\Delta_E(\mathbf{r}_L^*) = \mathsf{E}[\text{REV}_E(\mathbf{b}; \mathbf{r}_L^*)] - \mathsf{E}[\text{REV}(\mathbf{b}; \mathbf{0})]$$

$$\Delta_L(\mathbf{r}_M) = \mathsf{E}[\text{REV}_L(\mathbf{b}; \mathbf{r}_M)] - \mathsf{E}[\text{REV}(\mathbf{b}; \mathbf{0})]$$

$$\Delta_E(\mathbf{r}_M) = \mathsf{E}[\text{REV}_E(\mathbf{b}; \mathbf{r}_M)] - \mathsf{E}[\text{REV}(\mathbf{b}; \mathbf{0})]$$

where the expectation $\mathsf{E}[\cdot]$ denotes the average over all historical queries.

We normalize the lifts by setting the revenue of the lazy reserve auction with optimally set reserve prices to 1 ($\Delta_L(\mathbf{r}_L^*) = 1$), and report multiplicative improvement over this setting.

Table 1 shows that the eager auction outperforms the lazy auction in practice both when the optimal lazy reserves are used and when monopoly reserves are used. We also observe that there is no clear winner between using optimal lazy

Table 1: Revenue Lift Comparison. The units are normalized so that $\Delta_L(\mathbf{r}_L^*) = 1$

$\Delta_E(\mathbf{r}_L^*)$	$\Delta_L(\mathbf{r}_M)$	$\Delta_E(\mathbf{r}_M)$
1.13204	0.892116	1.18965
1.24867	0.958164	1.28977
1.16233	0.942408	1.09623
1.19286	0.886347	1.11872
1.14805	0.942208	1.08097

reserves or monopoly reserves in the eager auction. For 3 out of 5 slots, the vector of optimal lazy reserves outperforms the monopoly reserves. In practice one may want to start with any of those two vectors and perform local updates to improve the performance of the eager auction.

Another axis along which it makes sense to compare the eager and lazy auctions is the welfare loss of the allocation, i.e., how much welfare is lost due to the application of reserve prices. If $W^L(\mathbf{b}; \mathbf{r})$ and $W^E(\mathbf{b}; \mathbf{r})$ are respectively the welfare of the lazy and eager auctions when reserve price $\mathbf{r}$ are applied, we define the quantities:

$$\tilde{\Delta}_L(\mathbf{r}_L^*) = \mathsf{E}[W(\mathbf{b}; \mathbf{0})] - \mathsf{E}[W_L(\mathbf{b}; \mathbf{r}_L^*)]$$

$$\tilde{\Delta}_E(\mathbf{r}_L^*) = \mathsf{E}[W(\mathbf{b}; \mathbf{0})] - \mathsf{E}[W_E(\mathbf{b}; \mathbf{r}_L^*)]$$

$$\tilde{\Delta}_L(\mathbf{r}_M) = \mathsf{E}[W(\mathbf{b}; \mathbf{0})] - \mathsf{E}[W_L(\mathbf{b}; \mathbf{r}_M)]$$

$$\tilde{\Delta}_E(\mathbf{r}_M) = \mathsf{E}[W(\mathbf{b}; \mathbf{0})] - \mathsf{E}[W_E(\mathbf{b}; \mathbf{r}_M)]$$

which are the analogues for social welfare of the quantities described in Table 1.

As expected, Table 2 shows that the welfare loss in the eager auction is larger than the welfare loss in the lazy auction. Similarly to what occurs for revenue, there is no clear winner between the optimal lazy reserves and the monopoly reserves for the eager auction.

Table 2: Welfare Loss Comparison. The units are normalized so that $\tilde{\Delta}_L(\mathbf{r}_L^*) = 1$

$\tilde{\Delta}_E(\mathbf{r}_L^*)$	$\tilde{\Delta}_L(\mathbf{r}_M)$	$\tilde{\Delta}_E(\mathbf{r}_M)$
0.689856	1.34808	0.896388
0.70264	1.16671	0.814005
0.602512	1.046	0.695923
0.616777	1.0458	0.647172
0.636252	0.995548	0.671452

8. CONCLUSION

The results in this work follow two major themes. The first lies in devising methods for testing new strategies, be they reserve prices, bidding agents, or mechanisms, in competitive environments [CHN14]. As we saw in Section 6 these can have counterintuitive effects, and robust experimental design and analysis is required for making correct decisions.

The second is in foregoing the assumption that bidders' value distributions are known and investigating the computational complexity of optimizing from previously observed data. Whether designing new mechanisms in these environments [CR14], understanding the amount of past data needed [MR15], or, providing approximation algorithms, as we do in this work, this is a rich and exciting open area.

9. REFERENCES

[ARS13]   Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 1169–1177, 2013.

[CGM15]   Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.

[CHN14]   Shuchi Chawla, Jason Hartline, and Denis Nekipelov. Mechanism design for data science. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 711–712, New York, NY, USA, 2014. ACM.

[CLMN14] L. Elisa Celis, Gregory Lewis, Markus M. Mobius, and Hamid Nazerzadeh. Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science*, 60(12):2927–2948, 2014.

[CR14]   Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 243–252, 2014.

[DRY15]   Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[HOT15]   Henning Hohnhold, Deirdre O'Brien, and Diane Tang. Focusing on the long-term: It's good for users and business. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 1849–1858, New York, NY, USA, 2015. ACM.

[HR09]   Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6–10, 2009*, pages 225–234, 2009.

[KHS07]   Ron Kohavi, Randal M. Henne, and Dan Sommerfield. Practical guide to controlled experiments on the web: Listen to your customers not to the hippo. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pages 959–967, New York, NY, USA, 2007. ACM.

[MM14]   Mehryar Mohri and Andres Muñoz Medina. Optimal regret minimization in posted-price auctions with strategic buyers. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1871–1879, 2014.

[MR15]   Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 136–144. Curran Associates, Inc., 2015.

[Mye81]   R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Wil85]   R. Wilson. *Game-Theoretic Analysis of Trading Processes.* Defense Technical Information Center, 1985.