# An Efficient Auction with Variable Reserve Prices for Ridesourcing⋆ ⋆⋆

Chaoli Zhang[1], Fan Wu[1], and Xiaohui Bei[2]

[1] Shanghai Jiao Tong University, Shanghai, China
{chaoli_zhang, fwu}@sjtu.edu.cn
[2] Nanyang Technological University, Singapore
xhbei@ntu.edu.sg

**Abstract.** Ridesourcing refers to the service that matches passengers who need a car to personal drivers. In this work, we study an auction model for ridesourcing that sells multiple items to unit-demand single-parameter agents with variable reserve price constraints. In this model, there is an externally imposed reserve price set for every item, and the price is both item- and bidder-dependent. Such auctions can also find applications in a number of other traditional and online markets, such as ad auction or online laboring market.
Our main result is a truthful, individually rational, and computationally efficient mechanism that respects the reserve price constraints and always achieves at least half of the optimal social benefit (*i.e.*, the sum of the valuations of the winning agents). Furthermore, we show such efficiency approximation is tight by proving that even without any computational constraints, no truthful and individually rational mechanism can achieve better than 2-approximation for social benefit maximization. Finally, we evaluate the performance of our mechanism based on real taxi-trace data. The empirical results show that our mechanism outperforms other benchmark mechanisms in terms of both social benefit and revenue.

**Keywords:** Ridesourcing· Reserve Prices Auction.

## 1 Introduction

Ridesourcing provides a wealth of efficiency and flexibility to urban transportation and has been injected new vigor in recent years due to the popularization of smartphones[3]. More specifically, an online transportation network company takes ride requests from passengers, and dispatches the available cars in the nearby location to serve the passengers. There are two fundamental functions of

⋆ This work was supported in part by the State Key Development Program for Basic Research of China (973 project 2014CB340303), in part by China NSF grant 61672348, 61672353, and 61472252, and in part by Shanghai Science and Technology fund 15220721300. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

⋆⋆ F. Wu is the corresponding author.

[3] Example companies in this domain include Lyft, Uber, Didi, etc.

such platforms: ride-matching (*i.e.*, assigning the drivers to the passengers) and pricing [7].

The above ridesourcing scenario can be modeled as a single-parameter auction design problem: each passenger has a single-parameter valuation for getting a ride. As one of the classical mechanism design problems, auctions have been widely studied as a powerful instrument for resource allocation due to their perceived fairness and allocation efficiency. In addition, each potential ride has a *reserve price*, usually set by the government as a regulation, so as to guarantee the minimum profit from each ride and to prevent destructive competition, which is determined mainly by the driver's cost (*i.e.*, fuel consumption and labor time, etc) or some other factors of the driver (*i.e.*, different drivers may have different critical values for a ride picking). Hence these reserve prices would depend on not only the trip length, but the locations of the driver and the passenger, as the driver needs to pick up the passenger before continuing their ride. It is reasonable and necessary to set different reserve prices for different passengers when they are to be allocated to different drivers.

Reserve price has been a salient feature in many single-parameter auctions and many practical scenarios, such as online auction markets like eBay [3], either to boost the revenue [16], or to favor a particular group of agents [15]. However, in most previous works, the reserve price is set by the auctioneer and only simple and symmetric price rules are considered. That is, either the auctioned item has a common reserve price for all agents, or in certain cases, though the price can be set differently for different agents, different units of the same item always have the same reserve price.

In this paper, we study a simple auction model that captures the most salient aspect of the ridesourcing problem above. Our goal is to design mechanisms that is truthful and individually rational, with the objective of maximizing the *social benefit* of the allocation, defined as the sum of valuations of all agents who win an item individually.

### 1.1 Our Results

Our results in this paper are listed as follows.

- We propose a truthful, individually rational, and 2-approximation mechanism in terms of social benefit for the auctions with variable reserve prices.
- We also exploit limitations of auctions with variable reserve prices, and show that when restricted to truthful and individually rational mechanisms, the 2-approximation social benefit guarantee of our mechanism is optimal.
- Finally, we demonstrate the performance of our mechanism in a case study that considers taxi matching using real taxi-trace data. Experiment results show that our mechanism outperforms other benchmark mechanisms by a clear margin in terms of both social benefit and revenue.

The rest of the paper is organized as follows. In Section 2, we review related works. In Section 3, we show the preliminaries for our work. In Section 4, we give our truthful approximation mechanism and show that the 2-approximation achieved by our mechanism is actually tight. In Section 5, we report the evaluation results. We make a conclusion in Section 6.

## 2 Related Work

Many works have been done for ridesourcing (sometimes also called ridesharing) in recent years [12, 11, 1, 10, 20, 5, 17, 13, 21, 6, 9]. Due to the complexity of the ridesourcing problem, they focused on many different issues, scheduling, pricing, market analysis and so on.

Some of them which are based on market-based mechanisms design are most related with our work. [13] modified the VCG payment scheme proposed in [17] to adapt it to the dynamic requirements of the open-world ridesourcing problem. [21] showed that the deficit generated by the VCG mechanism can be arbitrarily large and proposed less efficient alternatives based on fixed prices and two-sided reserve prices that have deficit control. [6] proposed a coalition game for ridesourcing problem. [9] focused on sharing the resulting cost among passengers which is also different from the pricing problem in ridesourcing.

A majority of the literature on auction design have restricted their study to the single-parameter framework. Such a framework includes many practical scenarios, such as single-item auctions [4], digital goods auctions [8], multicast auctions [19], to name a few.

In many single-parameter auctions, reserve price are set to boost the revenue [16], or to favor a particular group of agents [15]. Although the reserve price constraints are widely used and play an important role to achieve certain design goal, they also bring challenges to the mechanism design field, as the reserve price constraints restrict the space of the possible allocations. Different from many existing works, variable reserve prices are considered in this paper.

We combine the classical greedy weighted maximum matching design technique [18] with the Myerson's Lemma [16, 2] in a creative way to design the novel truthful and individually rational mechanism for social benefit maximization with the variable reserve prices constraints and the approximation ratio achieved by the mechanism we design is actually tight in theory.

## 3 Preliminary

We present the ridesharing problem model in a general auction framework. Consider a multi-unit auction with $m$ items and $n$ unit-demand, single-parameter agents. Each agent $i$ has a private value $v_i$ for getting an item. The auction collects agent bids $(b_1, ..., b_n)$, and outputs an allocation $A$ of the items to the agents as well as payment $p_i$ for each agent $i$. We assume agents have *quasi-linear* utilities defined as $u_i = v_i - p_i$ if $i$ wins an item, and $u_i = 0$ otherwise. Further, there is an externally imposed reserve price $\bar{p}_{ij}$ for agent $i$ to buy item $j$. That is, when item $j$ is assigned to agent $i$, the payment $p_i$ of agent $i$ should not be less than $\bar{p}_{ij}$. We assume that reserve prices $\bar{p}_{ij}$ are public information, but the valuation $v_i$ is private information to agent $i$.

We aim to design auctions with the following properties.

- *Truthfulness* (also known as *incentive compatibility* or *strategy-proofness*): for every agent $i$ and fixed bids of other agents $b_{-i}$, agent $i$ can always maximize her utility by reporting $b_i = v_i$.
- *Individual Rationality*: for each agent $i$ whose value is $v_i$, assume that she is charged $p_i$ ($p_i = 0$ if she loses the auction) when bidding $v_i$, then $v_i - p_i$ must be non-negative.

– *Computational Efficiency*: the allocation and payments can be computed in polynomial time.

*Objectives.* We study the following objectives in this paper. Let $\mathbb{W}(A)$ be the set of agents who are individually assigned to an item in allocation $A$.

– The *social benefit* of an allocation $A$ is defined as the sum of valuations of all winning agents $SB(A) = \sum_{i \in \mathbb{W}(A)} v_i$.
– The *revenue* of the mechanism is defined as $REV(A) = \sum_{i \in \mathbb{W}(A)} p_i$. In this paper we focus mainly on the social benefit objective.

Due to the reserve price constraints, the notion of "optimal social benefit" needs to be defined carefully. Because an allocation that simply maximizes the social benefit may assign a certain item to an agent whose valuation is smaller than its corresponding reserve price. Such an allocation is unfavored because it does not admit any pricing scheme that satisfies both reserve price constraints and individual rationality. Thus, in this paper we define the optimal social benefit as

$$SB_{\text{OPT}} = \max\{SB(A) \mid v_i(j) \geq \bar{p}_{ij} \text{ if } A \text{ assigns item } j \text{ to } i\}.$$

A mechanism is said to have $c$-approximation if it always outputs an allocation with social benefit no less than $\frac{SB_{\text{OPT}}}{c}$.

### 3.1 Discussion on VCG-type Mechanisms

The Vickrey-Clarke-Groves (VCG) mechanism is a cornerstone method in mechanism design when it comes to the optimization of global objective. It is a truthful mechanism that achieves optimal global objective and can be applied to a variety of domains. However, it does not take reserve prices into consideration, hence the resulting payments might violate these constraints.

One might also be tempted to consider variants of the VCG mechanism. One possibility is to restrict the domain of feasible allocations to the ones in which every winning agent has valuation not less than the corresponding reserve price. The reserve price constraints can then be satisfied. Such change, however, might break the truthfulness of the mechanism.

*Example 1.* Consider two agents $\{1, 2\}$ and two items $\{a, b\}$, with variable reserve prices $\bar{p}_{1a} = \bar{p}_{2a} = 1$, $\bar{p}_{1b} = \infty$, $\bar{p}_{2b} = 2$ and $v_1 = 1$, $v_2 = 2$. When reporting truthfully, the VCG-type mechanism will assign item $a$ to agent 1 and item $b$ to agent 2, and charge them 1 and 2 respectively. However, agent 2 can get item $a$ with price 1 if she misreports her valuation as 1.5.

Another variant is to view the reserve price as the "cost" of allocating an item to an agent, and use the allocation that maximizes the total valuation minus the allocation costs as the resulting allocation. Such mechanism could preserve truthfulness. However, it no longer provides any social benefit guarantee.

*Example 2.* Consider two agents $\{1, 2\}$ and one item $a$, which reserve prices $\bar{p}_{1a} = 0, \bar{p}_{2a} = 1$ and valuations $v_1 = \epsilon, v_2 = 1$. The mechanism that maximizes the sum of valuation minus the reserve prices would allocate the item to agent 1 because $\epsilon - 0 > 1 - 1$. While the optimal social benefit, when defined as only the sum of valuations, is obviously 1.

### 3.2 Interpretation in Ridesourcing

The above model describes a general framework that can be applied to many applications. Below we demonstrate its interpretation in the ridesourcing scenario.

In the basic online ridesourcing problem [7], passengers requiring rides are the agents and available drivers offering rides can be considered as the items. Passenger $i$ reports her request $\theta_i = \{l_i^o, l_i^d, b_i\}$ to the platform where $l_i^o$ and $l_i^d$ are the origin and destination of her required ride respectively, $b_i$ is agent $i$'s valuation for the ride.

The platform collects the location information $l_j$ of each driver $j$ through GPS, and calculates the cost of assigning driver $j$ to passenger $i$ as the reserve price $\bar{p}_{ij}$. For instance, one way to define $\bar{p}_{ij}$ is to let $\bar{p}_{ij} = p_u * (dis(l_j, l_i^o) + dis(l_i^o, l_i^d))$, where $dis(m, n)$ denotes the distance between two points $m$ and $n$, and $p_u$ denotes the specific fuel consumption of the car. Then the platform needs to produce an assignment of the drivers to the passengers and the corresponding payment that each passenger pays to her driver.

The goal of the platform is to design a mechanism that motivates passengers to report their true valuations, and ensures that the payment of each passenger always covers the corresponding reserve price, and simultaneously maximizing the social benefit or revenue of the allocation.

## 4 A Truthful Approximation Mechanism

In this section, we present and analyze a truthful, individually rational and 2-approximation mechanism for social benefit maximization in the variable reserve price setting.

In the following, we consider a weighted bipartite graph $G = (V_A, V_I, E)$, where $V_A$ is a subset of the agents, $V_I$ is a subset of the items. Every edge $(i, j) \in E$ is associated with the corresponding reserve price $\bar{p}_{ij}$. We will add or remove vertices and edges of $G$ during the process of the mechanism.

We define a left-perfect matching in $G$ as follows.

**Definition 1.** *A matching $M$ with size $|V_A|$ in $G$ is called a* left-perfect matching *in $G$.*

Clearly, a left-perfect matching is also a maximum matching in $G$.

The intuition behind our mechanism is that, at any moment, we only consider the agents with bids higher than $\bar{p}$ and edges with reserve prices below or equal $\bar{p}$, for some threshold value $\bar{p}$. We continuously decrease the value of $\bar{p}$. Such process will add agents to the graph while removing edges. We will maintain the following invariants throughout the mechanism.

**Invariants:**

---

1. Graph $G$ always has a left-perfect matching.
2. There exists a threshold value $\bar{p}$, such that every agent in $G$ has valuation not less than $\bar{p}$, and every edge in $G$ has reserve price below or equal $\bar{p}$.

---

Each agent added to the graph will pay the smallest threshold value of $\bar{p}$ with which the graph still has a left-perfect matching.

The details of the mechanism are shown below.

**Mechanism 1.** EROS: a truthful, individually rational, 2-approximation mechanism.

**Input**: reported bids $(b_1, \ldots, b_n)$;

Let $C$ be the union set of all agents and edges, the value $v(e)$ for element $e$ in $C$ is defined as its bid if $e$ is an agent, and its reserve price if $e$ is an edge.<br>
Remove all edges $(i, j)$ with $b_i < \bar{p}_{ij}$ in $C$.<br>
Let $V_A = \emptyset$, $V_I$ be the set of all items.<br>
Let bipartite graph $G = (V_A, V_I, E = \emptyset)$.<br>
**for** *each element $e$ in $C$ in decreasing order of $v(e)$* **do**<br>
    **if** *$e$ is an edge $(i, j)$ in $G$ and $i \in V_A$* **then**<br>
        Remove $(i, j)$ from $G$.<br>
        **if** *$G$ no longer has a left-perfect matching* **then**<br>
            Assign item $j$ to agent $i$, let $i$ pay $\bar{p}_{ij}$.<br>
            Remove $i$ and $j$ from $G$.<br>
    **if** *$e$ is an agent $i$* **then**<br>
        **if** *$G \cup \{i\}$ has a left-perfect matching* **then**<br>
            Add vertex $i$ and its incident edges to $G$.<br>
        **else**<br>
            **for** *each agent $i'$ in $G$* **do**<br>
                Let $M$ be a left-perfect matching in $G$<br>
                **if** *$G \setminus \{i'\} \cup \{i\}$ has a left-perfect matching* **then**<br>
                    Assign agent $i'$ her matched item $j$ in $M$, let $i'$ pay $b_i$.<br>
                    Remove $i'$ and $j$ from $G$.<br>
                    $i$ remains unallocated.

**Lemma 1.** *The invariants hold throughout the run of the mechanism.*

**Theorem 1.** EROS *is a truthful, individually rational, and computationally efficient mechanism with 2-approximation for social benefit maximization.*

We prove each part of Theorem 1 separately in the following subsections.

### 4.1 Truthfulness

Our mechanism design problem belongs to the single-parameter domain (*i.e.*, the private information of every agent is a single value). By the well known characterization of truthful mechanisms in single-parameter domains [16], it suffices to show that our mechanism satisfies the following two properties.

- *Allocation Monotonicity.* For every agent $i$ and fixed bids $b_{-i}$ by other agents, $i$ can win an item when bidding $b_i$, then she can still win if she bids any value $b'_i > b_i$.
- *Critical Payment Rule.* The payment charged to each winning agent equals to the minimal bid for this agent to win. This value is also called the *critical bid* of this agent.

*Allocation Monotonicity.*

**Lemma 2.** *The allocation rule in* EROS *is monotone.*

*Proof.* First, it is easy to see from the mechanism that, every agent that is added to graph $G$ in line 11 of the mechanism will end up being assigned an item in the output allocation. Thus the allocation monotonicity rule is equivalent of saying that, if agent $i$ can be added to $G$ with bid $b_i$, $i$ will still be added to $G$ if she bids any value $b'_i > b_i$, assuming all other bids remain the same.

We focus on the agent adding process in EROS. Without loss of generality, assume that $k(k')$ agents, denoted as $A_k(A_{k'})$, have been added to $G$ before agent $i$ bidding $b_i(b'_i)$ is considered at moment $t_i(t'_i)$. Note that $A_k(A_{k'})$ does not only contain agents in $G$ at $t_i(t'_i)$, but also the ones that have been matched and removed from $G$ before $t_i(t'_i)$.

When processing elements (*i.e.*, agents and edges) in the union set $C$ in decreasing order of their values, agent $i$ will be processed earlier when bidding $b'_i$ than bidding $b_i$. Thus we have $t'_i \leq t_i$. With all the other information remaining unchanged, it is easy to see that $A_{k'} \subseteq A_k$. Furthermore, if graph $(A_k \cup \{i\}, I, E)$ has a left-perfect matching, graph $(A_{k'} \cup \{i\}, I, E)$ will have a left-perfect matching as well. Therefore when bidding $b'_i$, agent $i$ will still be added to graph $G$, and thus be assigned an item in the final mechanism outcome.

*Critical-bid Payment.* Next we show the payment specified by our mechanism coincides with the critical bid of the monotone allocation rule.

**Lemma 3.** *The payment of a winning agent $i$ is her critical bid in* EROS.

*Proof.* Here we show that the payment of each assigned agent is the minimum bid that she needs to report to be added in $G$. That is, the agent will lose the item if she bids anything less than that payment.

There are two places in the mechanism where items are assigned to agents. We consider them separately.

1. When the element $e$ being processed is an edge $(i, j)$ (line 8,9).
   Here agent $i$ is assigned item $j$ and pays $\bar{p}_{ij}$. Let $G$ be the bipartite graph at this moment. Based on line 7 condition, we know that $G \setminus \{(i, j)\}$ does not have a left-perfect matching.
   Now assume that agent $i$ bids $b'_i = \bar{p}_{ij} - \varepsilon$ for some small $\varepsilon > 0$. Note that in this case, edge $(i, j)$ will be removed from $C$ at line 2 of the mechanism. Therefore when agent $i$ is being processed, adding it will give us exactly the graph $G \setminus \{(i, j)\}$. Because it does not have a left-perfect matching, we know in this case agent $i$ will not be added to $G$ and hence will be left unassigned.
2. When the element $e$ being processed is an agent $i$ (line 15,16).
   Here agent $i'$ is assigned her matched item in matching $M$ and pays $b_i$, where $i$ is the agent failing to be added to $G$.
   Assume $i'$ bids $b_i - \varepsilon$. According to the mechanism, the following events will happen.
   (a) Agent $i$ will be picked in $C$ before agent $i'$.
   (b) When processing agent $i$, she will be added to the graph since $G \setminus \{i'\} \cup \{i\}$ has a left-perfect matching.

(c) After $i$ is added, agent $i'$ will not be able to added to $G$ any more when processed. Because otherwise it would violate the condition that agent $i$ is not added to $G$ with the original bidding.

We conclude that in this case, agent $i'$ will also not be assigned if she bids below the payment.

Lemma 2 and Lemma 3 together imply the truthfulness of Eros.

### 4.2 Individual Rationality and Reserve Price Constraints

Individual rationality and reserve price constraints can be easily proved by the invariant properties of the mechanism.

**Lemma 4.** Eros *respects the reserve price constraints and is individually rational.*

*Proof.* Note that at each iteration of the mechanism, the value of element $e$ in $C$ is exactly the threshold value mentioned in Invariant 2. Furthermore, whenever we assign an item $j$ to an agent $i$ in the iteration (at line 9 or line 18), the payment charged to the agent is exactly this threshold value. Hence we always have $v_i \geq p_i \geq \bar{p}_{ij}$. This proves both reserve price constraints and individual rationality of the mechanism.

### 4.3 Approximation Ratio

To show the social benefit approximation of Eros, we relate the winning set of agents to the outcome of the following simple greedy allocation algorithm.

---

**Greedy Algorithm**

Consider each agent $i$ in decreasing order of $v_i$, if there is an unassigned item $j$ with $v_i \geq \bar{p}_{ij}$, assign item $j$ to agent $i$ (if there are more than one such items, pick an arbitrary one).

---

This is a variant of the greedy matching algorithm, which is known to achieve 2-approximation for social benefit maximization. The proof is omitted due to space constraints.

**Lemma 5.** *The Greedy Algorithm is a 2-approximation to the optimal social benefit solution.*

Next we will show that the allocation produced by Eros coincides with a particular run of this greedy algorithm.

**Lemma 6.** Eros *always achieves at least half of the optimal social benefit.*

*Proof.* Let $M$ be the matching output by our mechanism. We assume that this matching is produced by the following process:
For each agent $i$ in decreasing order of $v_i$, if $i \in M$, assign agent $i$ her matched item in $M$. Otherwise discard this agent. To prove this allocation is 2-approximation, it suffices to show that the above process is also a valid running process of the greedy algorithm. That is, we want to show that for each agent $i \notin M$, when

$i$ is considered in the greedy algorithm, there is no unassigned item $j$ left with $v_i \geq \bar{p}_{ij}$.

Assume by contradiction that we can find an unassigned item $j$ with $v_i \geq \bar{p}_{ij}$ when processing agent $i$. This means that when processing agent $i$ in line 11 of EROS, the graph $G \cup \{i\}$ will have a left-perfect matching $M \cup (i, j)$. Hence according to the mechanism, agent $i$ will be added to graph $G$ at this step. However, we also observe that every agent added to the graph $G$ in the mechanism will all end up being assigned to an item. This contradicts the assumption that $i \notin M$ and proves the lemma.

### 4.4 Lower Bounds

In this subsection we focus on the lower bounds for the approximation ratio of any truthful and individually rational mechanism. We show that the 2-approximation achieved by our mechanism is actually tight in this model.

**Theorem 2.** *There is no truthful and individually rational mechanism that can achieve an approximate ratio better than 2 for social benefit maximization, even with only two items.*

*Proof.* We prove this theorem through a concrete counter example. Assume otherwise that such a mechanism exists. We focus on the following 2 scenarios.

First, consider two agents $\{1, 2\}$ and two items $\{a, b\}$, with variable reserve prices $\bar{p}_{1a} = \bar{p}_{2a} = 1$, $\bar{p}_{1b} = \infty$, $\bar{p}_{2b} = 1 + \varepsilon$ and $v_1 = 1$, $v_2 = 1 + \varepsilon/2$ where $\varepsilon > 0$. Since neither agent can afford item $b$, and agent 2 has a higher value than agent 1, the truthful mechanism must assign item $a$ to agent 2 with price $p_2 \in [1, 1 + \varepsilon/2]$ and leave item $b$ unassigned.

Next, consider a similar case where from the above example we change the value $v_2$ from $1 + \varepsilon/2$ to $1 + \varepsilon$. Note that although agent 2 can now afford to have item $b$, a truthful mechanism must still assign item $a$ to agent 2 with price $p_2 \in [1, 1 + \varepsilon/2]$ when they are bidding truthfully. This is becasue agent 2 can always get this allocation by reporting $b_2 = 1 + \varepsilon/2$, and this is always a better result for agent 2 than assigning item $b$ to her.

Note that in the later instance, the optimal social benefit is $2 + \varepsilon$ from assigning item $a$ to agent 1 and item $b$ to agent 2. Hence any truthful and individually rational mechanism cannot have approximation ratio better than $\frac{2+\varepsilon}{1+\varepsilon}$, whose limit is 2 when $\varepsilon \to 0$.

## 5 Evaluation

We conduct experiments in a ridesourcing scenario on a real taxi-trace dataset to evaluate the performance of our mechanism. Please refer to Section 3 "Example in ridesourcing" paragraph for the model details.

### 5.1 Simulation Setup

Our experimental evaluation is based on a taxi-trace dataset of city Shanghai in 2015. The data contains the locations of about 100 taxis in that city at a particular time.

We then randomly sample the information of passengers [14]. For each passenger $i$, her origin $l_i^o$ and destination $l_i^d$ are randomly sampled from a uniform

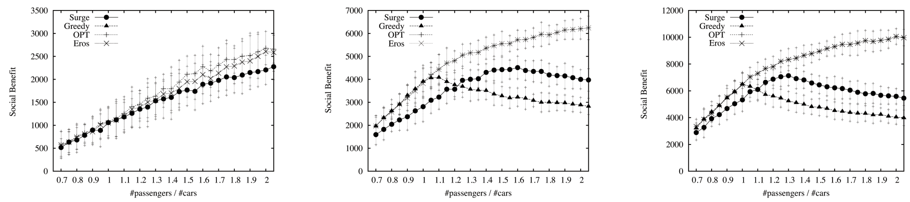


(a) Social Benefit for $k = 1$ (b) Social Benefit for $k = 3$ (c) Social Benefit for $k = 5$

Fig. 1: Social Benefit Comparison

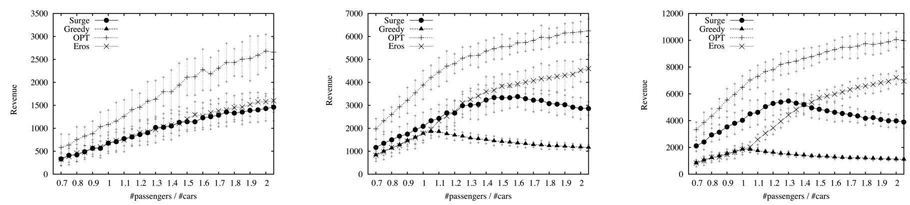


(a) Revenue for $k = 1$ (b) Revenue for $k = 3$ (c) Revenue for $k = 5$

Fig. 2: Revenue Comparison

distribution over the main city area. Recall that the reserve price of assigning taxi $j$ to passenger $i$ is $\bar{p}_{ij} = p_u * (dis(l_j, l_i^o) + dis(l_i^o, l_i^d))$, where $l_j$ is the location of taxi $j$, and $p_u$ is the specific fuel consumption of a car in that city. Note the actual value of $p_u$ is not important since it is a common factor in all costs and passenger valuations.

The individual valuation $v_i$ of each passenger $i$ is sampled from a Gaussian distribution over the interval $2\mu \geq v_i \geq 0$ with mean $\mu = k * p_u * dis(l_i^o, l_i^d)$ and variance $\delta = 20$. The idea is to let $\mu$ represent the standard taxi fare for the trip, and use Gaussian distribution to model that with high probability a passenger would be willing to pay a price close to that taxi fare. The taxi price is on average 5 times the gas or oil cost of the trip in the targeted city. In order to examine the effects of different valuation distributions on the mechanism performance, we run the entire set of experiments with three different values $k = 1, 3, 5$.

To investigate the influence of different levels of supply and demand on the performance of each mechanism, we run all the aforementioned mechanisms with different passenger quantities $\mathcal{P}$, ranging from 70 to 200. For every value of $\mathcal{P}$, we simulate the mechanisms on 50 sets of passenger samples and take the average result. Each data point is also plotted together with the 95% confidence interval.

### 5.2 Benchmark Mechanisms

To evaluate the performance of Eros, we compared Eros with several other mechanisms.

– Greedy. We use this mechanism to model the simplest strategy of assigning each passenger to its nearest car in a greedy fashion. The mechanism runs as follows: Repeatedly pick passenger $i$ and car $j$ with the smallest distance

$dis(l_j, l_i^o)$, offer this car to this passenger at price $p_i = \bar{p}_{ij}$. If passenger $i$ accepts this offer (*i.e.*, her bid is not less than this price), make the match and remove this pair from the system. Otherwise mark passenger $i$ as unassigned and remove her from the system (because under individual budget balance condition this passenger cannot be assigned to any car). Clearly this mechanism is truthful and individually rational.

– Surge. Surge pricing is a dynamic pricing strategy used by many online ridesourcing companies. It is an effective way to extract more revenue from the discrepancy between demand and supply. Here we consider a very simple version of this pricing strategy. Let $\alpha \geq 1$ be the surge factor. We simply employ the Greedy mechanism described above, but replace the offered price $\bar{p}_{ij}$ by $\alpha \cdot \bar{p}_{ij}$. We vary the value of $\alpha$ from 1 to 5 and pick the allocation with the largest revenue as our final allocation. Note that this is *not* a truthful mechanism.

Finally, we use OPT to denote the optimal social benefit in each problem instance. Note that as shown in Section 4.4, no truthful and individually rational mechanism can always achieve optimal social benefit. Thus it only serves as an upper bound of all mechanisms' performances. Also note that when considering revenue as the performance metric, the optimal social benefit still serves as a (perhaps loose) upper bound on the possible revenue that any mechanism can achieve.

### 5.3 Results

Figures 1(a), 1(b), 1(c) show the social benefit achieved by different mechanisms with $k = 1, 3, 5$ respectively[4]. We can see from the results that our mechanism outperforms all other mechanisms by a clear margin. In particular, when the ratio between the number of passengers and the number of taxis is large, *i.e.*, the market is in high demand, our mechanism achieves near optimal social benefit most of the time, while the performances of Greedy and Surge start to drop down.

Figures 2(a), 2(b), 2(c) show the revenue comparisons between each mechanism. Note that in this work we did not provide any theoretical analysis on the revenue aspect of our mechanism. Yet empirical results show that our mechanism is able to generate competitive revenue performance. In particular, when the demand is high compared with the supply, our mechanism again outperforms Surge, which is a non-truthful mechanism designed specifically for the purpose of revenue extraction.

### 5.4 Discussion

From the experiments we can also form the following observations and discussions.

*Low Demand vs. High Demand.* The discrepancy between supply and demand plays an important role in different mechanisms. Note that as the number of passengers increases, OPT also increases. However, such correlation is not reflected in Greedy and Surge. The reason is that these two mechanisms assign cars to passengers based on the distances between them, hence they do not take

[4] In Figure 1(a) Surge mechanism achieves optimal performance with surge factor $\alpha = 1$, *i.e.*, it coincides with the Greedy mechanism. Thus its plot is not displayed in the figure. The same is true also for Figure 2(a).

advantage of the increasing number of high-valuation passengers. On the other hand, our mechanism is able to exploit such discrepancy and picks passengers selectively based on their valuations. This partially explains why our mechanism always gives good performance in the cases with high demand.

*Theory vs. Practice.* In Section 4.4 we established a lower bound of 2 on the social benefit approximation ratio for any truthful mechanisms in this model. However, as the empirical results suggest, in practice one can usually go beyond this lower bound and achieve much better social benefit. This is because practical data usually have special structures and properties. How to formally analyze these properties from a theoretical point of view, and obtain better lower/upper bounds for auctions in specific application domains is an interesting open question in this area.

## 6 Conclusion and Future Work

In this paper, we study an mechanism design model for ridesourcing that allocates cars to unit-demand single-parameter agents with variable reserve price constraints. We give a truthful, individually rational, and computationally efficient mechanism that respects the reserve price constraints and achieves 2-approximation for social benefit maximization. We further evaluate the performance of our mechanism based on real taxi-trace data, and empirical results show that our mechanism outperforms other benchmark mechanisms in terms of both social benefit and revenue.

There are a number of future working directions that worth pursuing. First, given the online nature of the ridesourcing application, an important challenge is to design a dynamic online mechanism. Second, given the simplicity of our mechanism, it would be interesting to know how it can be generalized to be applied in other related domains.Third, appropriate design of the formulation of the social benefit can make our mechanism be adapted to practical applications with more interesting social objectives.

## References

1. Agatz, N., Erera, A., Savelsbergh, M., Wang, X.: Optimization for dynamic ride-sharing: A review. European Journal of Operational Research **223**(2), 295–303 (2012)
2. Archer, A., Papadimitriou, C., Talwar, K., Tardos, É.: An approximate truthful mechanism for combinatorial auctions with single parameter agents. Internet Mathematics **1**(2), 129–150 (2004)
3. Bajari, P., Hortacsu, A.: The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. RAND Journal of Economics pp. 329–355 (2003)
4. Baker, J., Song, J.: A review of single-item internet auction literature and a model for future research. Journal of Electronic Commerce in Organizations **5**(1) (2007)
5. Banerjee, S., Johari, R., Riquelme, C.: Pricing in ride-sharing platforms: A queueing-theoretic approach. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, 2015. EC'15. pp. 639–639. ACM (2015)
6. Bistaffa, F., Farinelli, A., Ramchurn, S.D.: Sharing rides with friends: a coalition formation algorithm for ridesharing. In: Proceedings of the 2015 AAAI Conference on Artificial Intelligence, 2015. AAAI'15. AAAI (2015)

7. Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M.E., Wang, X., Koenig, S.: Ridesharing: The state-of-the-art and future directions. Transportation Research Part B: Methodological **57**, 28–46 (2013)
8. Goldberg, A.V., Hartline, J.D., Wright, A.: Competitive auctions and digital goods. In: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms, 2001. SODA'01. pp. 735–744. Society for Industrial and Applied Mathematics (2001)
9. Gopalakrishnan, R., Mukherjee, K., Tulabandhula, T.: The costs and benefits of ridesharing: Sequential individual rationality and sequential fairness. In: Proceedings of the 8th ACM Conference on Electronic Commerce, 2016. EC'16. ACM (2016)
10. Herbawi, W., Weber, M.: The ridematching problem with time windows in dynamic ridesharing: A model and a genetic algorithm. In: Evolutionary Computation (CEC), 2012 IEEE Congress on. pp. 1–8. IEEE (2012)
11. Herbawi, W.M., Weber, M.: A genetic and insertion heuristic algorithm for solving the dynamic ridematching problem with time windows. In: Proceedings of the 14th annual conference on Genetic and evolutionary computation. pp. 385–392. ACM (2012)
12. Jacob, J., Roet-Green, R.: Ride solo or pool: The impact of sharing on optimal pricing of ride-sharing services (2017)
13. Kamar, E., Horvitz, E.: Collaboration and shared plans in the open world: Studies of ridesharing. In: International Joint Conference on Artificial Intelligence, 2009. IJCAI'09. vol. 9, p. 187 (2009)
14. Kleiner, A., Nebel, B., Ziparo, V.: A mechanism for dynamic ride sharing based on parallel auctions. In: International Joint Conference on Artificial Intelligence, 2011. IJCAI'11. pp. 266–272 (2011)
15. Kotowski, M.H.: On asymmetric reserve prices. Tech. rep., Mimeo (2015)
16. Myerson, R.B.: Optimal auction design. Mathematics of operations research **6**(1), 58–73 (1981)
17. Parkes, D.C., Kalagnanam, J.R., Eso, M.: Achieving budget-balance with vickrey-based payment schemes in exchanges. In: International Joint Conference on Artificial Intelligence, 2001. IJCAI'01. pp. 1161–1168 (2001)
18. Preis, R.: Linear time 1/2-approximation algorithm for maximum weighted matching in general graphs. In: Annual Symposium on Theoretical Aspects of Computer Science, 1999. STACS'99. pp. 259–269. Springer (1999)
19. Varshney, U.: Multicast over wireless networks. Communications of the ACM **45**(12), 31–37 (2002)
20. Zhang, D., Li, Y., Zhang, F., Lu, M., Liu, Y., He, T.: coride: Carpool service with a win-win fare model for large-scale taxicab networks. In: Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems. p. 9. ACM (2013)
21. Zhao, D., Zhang, D., Gerding, E.H., Sakurai, Y., Yokoo, M.: Incentives in ridesharing with deficit control. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, 2014. AAMAS'14. pp. 1021–1028. International Foundation for Autonomous Agents and Multiagent Systems (2014)