

Reservation Exchange Markets for Internet Advertising*

Gagan Goel¹, Stefano Leonardi^{†2}, Vahab Mirrokni¹, Afshin Nikzad^{‡3}, and Renato Paes-Leme¹

1 Google Research, New York

2 Sapienza University of Rome

3 Stanford University

Abstract

Internet display advertising industry follows two main business models. One model is based on *direct deals* between publishers and advertisers where they sign legal contracts containing terms of fulfillment for a future inventory. The second model is a spot market based on auctioning page views in real-time on advertising exchange (AdX) platforms such as DoubleClick's Ad Exchange, RightMedia, or AppNexus. These exchanges play the role of intermediaries who sell items (e.g. page-views) on behalf of a seller (e.g. a publisher) to buyers (e.g., advertisers) on the opposite side of the market. The computational and economics issues arising in this second model have been extensively investigated in recent times.

In this work, we consider a third emerging model called *reservation exchange market*. A reservation exchange is a two-sided market between buyer orders for blocks of advertisers' impressions and seller orders for blocks of publishers' page views. The goal is to match seller orders to buyer orders while providing the right incentives to both sides. In this work we first describe the important features of mechanisms for efficient reservation exchange markets. We then address the algorithmic problems of designing revenue sharing schemes to provide a fair division between sellers of the revenue collected from buyers.

A major conceptual contribution of this work is in showing that even though both clinching ascending auctions and VCG mechanisms achieve the same outcome from a buyer perspective, however, from the perspective of revenue sharing among sellers, clinching ascending auctions are much more informative than VCG auctions.

1998 ACM Subject Classification F.2 ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY

Keywords and phrases Reservation Markets, Internet Advertising, Two-sided Markets, Clinching Auction, Envy-free allocations.

1 Introduction

The universe of internet advertisement is divided in two big worlds: *search ads* and *display ads*. At first glance, they look very similar: both sell impressions using variants of the second price auction with reserves. A closer look, however, reveals that they are very different: in *search ads*, the platform (Google or Bing, for example) is both the *auctioneer* and the

* A preliminary version of this work was presented at the 11th Ad Auctions Workshop, Portland, June 2015.

† This work is supported by the Google Focused Award "Algorithms for Large-scale Data Analysis" and by EU FET project MULTIPLEX no. 317532. Work partially done while the author was a visiting scientist at Google Research, New York

‡ Work partially done while the author was on internship at Google Research, New York



seller, i.e., it sells inventory in their own properties. This makes it a one-sided market design problem, or in other words, the designer needs to reason only about the incentives of the buyers. In *display ads*, the platform is auctioning inventory not owned by them, turning it into a two-sided market design problem, where incentives for buyers (advertisers) and sellers (publishers, such as websites, blogs and news portals) need to be balanced.

Designing practical markets for display ads is challenging, since the theory of market design is much more developed for one-sided markets (there are tools like VCG, Myerson's optimal auction, ...) while for two-sided markets, what classic auction theory offers are mainly impossibility results, such as the Myerson-Satterthwaite Impossibility Theorem [23].

One complicating factor in display ads is that while we are used to think of internet advertisement in the form of auctions, auctions are only the tip of the iceberg. The most premium inventory is sold via the *reservations market* (also called *direct deals* or *guaranteed contracts*.) In this method, a publisher and an advertiser make a deal to allocate a certain number of impressions over a certain period, for a pre-specified price per impression. This deal is made offline in advance for a future inventory. Direct deals are known to suffer from inefficiencies for two reasons: first is the manual nature of formation of these contracts which allows a publisher to sign deals with a small number of advertisers, thus creating allocation inefficiencies. The second reason is the manual negotiation between buyers and publishers which incurs a huge cost and lowers the overall efficiency.


Auctions are fully automated and don't suffer from any of these inefficiencies. On the other hand, they can't guarantee to buyers and sellers the certainty that reservations can. For example, a brand launching a large campaign to advertise a new product certainly benefits from the certainty (both in terms of cost and volume of impressions) provided by reservation contracts.

The idea of *automated reservation market* is to overcome the shortcomings of both auctions (real-time spot market) and traditional reservations (offline negotiation). Such markets would allow sellers and buyers to transact for a bulk inventory in advance. This is inspired by a number of recent two-sided markets for online advertising, e.g., an exchange for future contracts¹. In particular, we study a two-sided market, which we call a *reservation market*, where publishers can post offers for blocks of ad slots characterized by parameters like supply level, reserve price, and their preference for a set of advertisers. Advertisers post requests for ad slots defined by parameters like valuation, demand, and targeting constraints specifying where and when they want to show their ads. The role of the reservation mechanism is to match the seller and buyer orders that attain some economic objectives.

Note that, unlike ad exchange markets that offer impressions available on the spot, the reservation market offers guaranteed deals for an inventory available in the future. Moreover, the reservation market brings together multiple publishers and advertisers with the goal of reducing the intermediation costs and the underlying inefficiencies of one-to-one deals, also by selling bundles of inventories from different publishers.

In this work, we start the investigation of the economics and algorithmic principles that are central to these reservation markets. The major questions we address in this work are: What features and incentive properties form the basis of a successful reservation market? What are the economic objectives of a reservation market? What are the revenue sharing policies that we can employ? What are the algorithmic problems we need to address in the design of reservation mechanisms?

¹ <http://www.massexchange.com/>



1.1 Our contribution

Our main contribution is to propose a formal model of reservation exchange market and discuss what are desirable properties (referred to as *axioms*) for this market. Secondly, we propose two specific mechanisms that help us understand the extent to which some of the aforementioned properties can be simultaneously achieved. We also provide several algorithmic results for the two mechanisms that we study.

Axioms for a reservation exchange market. In Section 2, we provide a simple and clean abstraction of reservation exchange markets for display ads as a two-sided matching market with buyer orders on one side and seller orders on the other side and in Section 3 we identify a list of axioms that we wish any mechanism for these markets to approximately satisfy.

The first axiom for the reservation mechanism that we discuss is the *efficiency of the market*, i.e., the social welfare of all agents of the market. The agents of the market are sellers and buyers, both with quasi-linear utilities. Buyers aim at maximizing their utility, i.e., the total value of the inventory received minus price. Sellers aim at maximizing revenue minus reservation price. The mechanism decides on the allocation to buyer orders of the inventory supplied from each seller, a payment to be charged to each buyer and a distribution of the revenue among sellers.

Individual rationality (IR) requires that participating in the mechanism is beneficial to all agents. *Incentive compatibility* (IC) requires that truthfully reporting one's preferences to the mechanism is the best strategy for each agent, independently from what the other agents report.

Budget Balance (BB) states the payments of the advertisers must entirely and exclusively be transferred to the publishers, i.e., the buyers and the sellers are allowed to trade without leaving to the mechanism any share of the payments, and without the mechanism adding money into the market. This axiom might appear strange at first glance, but it reflects a business practice common to most exchanges, which is of the exchange to get a fixed cut (typically called *revenue sharing margin*) of the seller's revenue. The reasoning behind this rule is that sellers have the option to send their inventory to different exchanges and keeping the revenue sharing margin fixed helps the exchange to be perceived as fair and hence attract more seller's inventory. Since fixed margins are an industrial standard in the ads world, any practical mechanism must implement some of that. Fixed margins are equivalent to budget-balance applied to bids rescaled by the revenue sharing margin up to rescaling bids.

An ideal goal in a reservation exchange is to design IR, IC, BB mechanisms that maximize the social welfare of all agents in the market. Unfortunately, Myerson and Satterthwaite [23] proved impossibility for an IR, IC, and BB mechanism that maximizes social welfare in such a market. The direction we pursue in this work is to ensure full efficiency and incentive compatibility for buyers as the most desirable goal for advertiser. This can be achieved for one-side markets by the family of Vickrey-Clarke-Groves (VCG) mechanisms [6, 29]. As extensively discussed in this work, the main problem with the vanilla VCG allocation is that it does not offer any good incentives to sellers, e.g., VCG can match fungible inventories from different sellers to buyers that offer very different payments thus producing the feeling that revenue is unfairly distributed among sellers.

An alternative to enforcing incentive compatibility for sellers is to design a mechanism that leads to a fair distribution of mechanism's revenue among sellers. Envy-free allocations [28] and other notions of market equilibria have often been considered in markets that cannot achieve full efficiency with truthful allocations. This leads to our third axiom as follows:



sellers should not envy each other with respect to the revenue that is received from the mechanism.

Most of the first part of the paper will be devoted to the discussion of which definition of *envy-freeness among sellers* is meaningful for reservation exchange markets. Our conclusion is that such a definition should crucially rely on the buyer-seller transactions that can arise in an efficient allocation, while it should disregard revenue that can only be obtained from allocations with suboptimal social welfare. With this aim, when introducing our notion of envy-freeness between sellers, we define the concept of *clinchng graph* as the collection of buyer-seller transactions that can arise in a VCG allocation. We conveniently define and compute the clinching graph by resorting to the implementation of VCG through an ascending clinching auction.

A major insight of this paper is that while the usual description of VCG payments as externalities imposed by agents on others offers little clue on how to split the proceeds of the auction among the sellers, the alternative description of VCG as an ascending auction (Ausubel's clinching framework [1]) provides additional structure obtained from the execution of the auction that can be exploited to design revenue sharing schemes. The clinching auction returns not only bundle prices, but the order in which each item was sold and the price at which the sale occurred. In an ideal case, whenever a *clinch* happens, if the clinching auction points to a unique item to be clinched, then there is a unique way to split the revenue among the sellers (and in this case clinching auctions capture the full information in how to split the revenue). However, sometimes, when a clinch happens, there are multiple items that can be used for that clinch. This is precisely the case when the clinching auction, even though it provides more information than VCG, it doesn't lead to a unique revenue sharing scheme, and we rely on a notion of envy-freeness for the revenue sharing scheme.

Finally, we discuss further desirable properties of reservation markets as additional axioms, and study our proposed mechanisms for their satisfaction of these axioms. These axioms are stability properties that prevent the market to be manipulated from buyers or sellers. We define the concept of *buyer monotonicity (BM)* as the property that the revenue of all sellers does not decrease when new buyer orders are presented. A second property called *seller monotonicity (SM)* states that the increase of the reservation price of a seller is not responsible of the decrease of the revenue of another seller.

Algorithmic results. We restrict our attention to buyer incentive-compatible efficient mechanisms based on VCG allocations. For all these mechanisms truth-telling is a dominant strategy for buyers. The major issue we face is to complement the VCG mechanism with a suitable envy-free revenue sharing scheme between sellers. Our first result is actually a negative result:

- There exists no efficient revenue sharing scheme that is both envy-free and budget balance. We actually demonstrate that any envy-free revenue sharing scheme cannot distribute more than a $\sqrt{3} - 1 + \delta$ share of the total revenue, for any $\delta > 0$.

Given the impossibility result above, we investigate the possibility of finding good trade-offs between budget balance or envy freeness. Relaxing one of these two constraints imply that either the mechanism is able to distribute a guaranteed share of the total revenue or that any seller has only limited envy of any other seller. With this goal in mind, we propose two revenue sharing schemes: i) a *revenue sharing by the clinching auction (CA)*, and ii) a *revenue sharing by the eating mechanism (EM)*.

For the CA revenue sharing scheme, we prove the following desirable properties:

- CA is budget balance.



- CA is 1/2-envy free and this bound is tight.
- CA is budget monotone and seller monotone.

Finally, for the EM revenue sharing scheme, we prove that following three results:

- EM is envy-free.
- EM is at most $\frac{11+\epsilon}{12+\epsilon}$ budget balance, for any $\epsilon > 0$.
- EM is at least $\frac{e-1}{e}$ -budget balance.

Most of the technical proofs of the results of the paper are given in the full version of the paper.

1.2 Related work

Double Auctions. Double auctions are special cases of two-sided markets with unit-supply buyers and sellers. Myerson and Satterthwaite [23] proved that it is impossible to obtain an IR, Bayesian IC², and weak BB³ mechanism to maximize social welfare in double auctions. Since then, much of the literature has focused on trading off social welfare for buyers and sellers, incentive compatibility and budget balance for double auctions [21, 25, 26, 11, 7]. The seminal work on double auctions [21] shows that efficiency for both buyers and sellers can actually be achieved asymptotically in large markets. In the context of one-shot auctions, optimal auctions for two-sided settings has been studied by [23], following the Nobel-prize winning work of [22]. The problem of finding the right trade-offs between social welfare, IC and BB is largely open for two-sided markets that model reservation exchanges. In this work we investigate two-sided markets that achieve IC for buyers and envy-freeness for sellers. This follows a line of work that looks at envy-freeness and other market equilibria if social welfare cannot be optimized truthfully [14, 19, 28]. Recently, this literature has also been adopted to design the optimal revenue sharing double auctions in the context of advertising exchanges [18]. In this paper, we focus on two-sided markets where multiple buyers are allocated to multiple sellers and the allocation and pricing are done differently. Other than online advertising systems, optimal two-sided markets can be applied to online and offline retailers and e-commerce websites like Amazon and eBay. A very recent paper by [20] studies EBay’s double auction problem, but their setting is different from this paper as they consider one buyer and multiple sellers, and explore approximately optimal pricing schemes for this setting.

Clinching Ascending Auctions. One fundamental component of this work is the use of the structure that can be obtained from the execution of Ausubel’s clinching auction [1] in designing revenue sharing schemes for the sellers. The clinching auction has been very successful in a variety of scenarios: designing auctions with budget constraints [10, 13, 8, 16], designing online auctions [17], extracting revenue in settings with budgets [3, 9]. The current paper adds to this line of work by showcasing another application of the clinching framework.

Cooperative games. Cooperative game theory may provide insights for modeling the fair sharing of revenue in the ad reservation exchange market. Shapley value [27] is a widely adopted notion of fair division between agents of the value of a game. It is however hard to extend this concept to our model since a crucial axiom of Shapley value (summability)

² Bayesian incentive compatibility is a less restrictive form of incentive compatibility.

³ Weak budget balance allows the mechanism to retain a share of the payments while not subsidizing the market.



does not hold in our case. Similar difficulties can also be found while trying to design a revenue sharing scheme that results in an attribution that lies in the core of a game [15]. On the positive side, we mention that the revenue sharing scheme by the clinching auction we present resembles Shapley values since it is defined as the expected revenue obtained over all possible seller permutations.

Market Equilibria. Several notions of equilibrium in markets have been studied. In a Walrasian equilibrium, we have item prices such that every agent receives a bundle of items that maximizes her utility, the market clears, and the corresponding outcome is efficient. However, except from very special cases (e.g. unit demand buyers), it can't be converted to a mechanism that is incentive-compatible for the buyer [19, 12]. Envy-freeness for buyers is also a concept widely used to characterize the stability of allocations. We do not survey here the extensive literature on this topic. However, we notice that we instead adopt the notion of envy-freeness to characterize fair revenue sharing schemes between sellers.

Reservation-based Internet advertising has also previously considered with more optimization-related questions than mechanism design questions. Examples of this line of work that is quite unrelated to the scope of this paper can be found in [30, 4, 24]. Markets that combine characteristics of the spot market and of direct deals between publishers and sellers have been also considered in [5] with the goal of maximizing the revenue of one single publisher.

2 Preliminaries

We consider a two-sided market, referred to as a *reservation exchange market*, consisting of a set B of buyers and a set S of sellers. Each seller $s_i \in S$ holds a supply of ℓ_i units of an *indivisible* good and has a reserve price ρ_i for each unit of those goods. Each buyer is interested in purchasing at most d_i units and has a value v_i per unit. The structure of the matching market is captured by a bipartite graph $G = (B \cup S, F)$ which indicates which buyer is interested in buying goods from which seller.

For example, in the case of internet advertisement, each buyer $b_j \in B$ corresponds to an advertiser and a seller $s_i \in S$ corresponds to a publisher. An edge $(b_j, s_i) \in F$ indicates that buyer b_j is interested in purchasing inventory from the publisher s_i 's website. We define $B_i = \{b_j \in B; (b_j, s_i) \in F\}$ as the set of buyers who target seller s_i inventory. Similarly, we define $S_j = \{s_i \in S; (b_j, s_i) \in F\}$ as the set of sellers that are targeted by buyer b_j .

We are interested in designing *reservation exchange mechanisms* (or simply reservation mechanisms) which associate for any given matching market described by (B, S, v, d, ℓ, ρ) an outcome composed of:

- an allocation $x_i[j] \in \mathbb{Z}$, indicating how many goods from seller s_i are sold to buyer b_j , respecting demands $a_j := \sum_i x_i[j] \leq d_j$ and supply $c_i := \sum_j x_i[j] \leq \ell_i$.
- a total amount P_j paid by each buyer b_j , such that the payment per unit doesn't exceed buyer j 's value per unit: $P_j \leq a_j \cdot v_j$
- a revenue sharing scheme which allocates for each buyer b_j and seller s_i , a portion $R_i[j]$ of the buyer's payment P_j to seller s_i , such that $\sum_i R_i[j] \leq P_j$. We define $R_i := \sum_j R_i[j]$ to be the total revenue obtained by seller i .

An outcome satisfying the properties above is said to be a *feasible outcome*. Given a feasible outcome, the utility of involved agents are as follows:

- buyers have quasi-linear utilities, i.e, $u_j = v_j \cdot a_j - P_j$.
- sellers have the revenue minus the reservation price $R_i - \rho_i \cdot c_i$ as their utility.



In the next section, we discuss a set of desirable properties for a reservation mechanism and discuss which subsets of those properties can be simultaneously satisfied.

3 Axioms for Reservation Exchange Markets

In this section, we develop an axiomatic approach to reservation markets. First, we define a set of desirable properties, referred to as axioms, for a well-designed market. As it is the case with axiomatic approaches, some seemingly innocuous axioms might generate impossibility results and some seemingly powerful axioms might not prevent the pitfalls we intended. Here, we define a family of axioms (many with different variations) and discuss, using examples, their strengths and weaknesses.

3.1 Fundamental Axioms: efficiency and budget-balance

We establish as our first and most important goal the maximization of market efficiency, which is the sum of the utilities of all agents involved:

- **Efficiency [Eff]:** The implemented outcome maximizes $SW(B \cup S) = \sum_{j \in B} v_j \cdot a_j + \sum_{i \in S} \rho_i \cdot (\ell_i - c_i)$ among all feasible outcomes, assuming the seller derives utility ρ_i for unsold items.

In order to simplify notation, for each seller that has a reserve price $\rho_i > 0$ we add a *proxy buyer* $j(i)$ with demand $d_{j(i)} = \ell_i$ and value $v_{j(i)} = \rho_i$. Let also $j(i)$ be the endpoint of a single edge connecting it to seller i . Notice that there is a social-welfare-preserving one-to-one map between outcomes for sellers with reserve prices and sellers without reserves but with proxy buyers. This reduction allows us to ignore from this point on the reserve prices ρ_i and focus on maximizing $\sum_i v_j \cdot a_j$ as the [Eff] goal.

A second goal of the mechanism is to distribute the revenue between sellers. A budget balance mechanism should distribute the entire revenue collected from the buyers to the sellers. A β -budget balance mechanism should distribute at least a β -fraction of the revenue.

- **β -Budget Balance [β -BB]:** The implemented outcome is β -budget balance for a constant $\beta \in [0, 1]$ if $\sum_{i \in S} R_i \geq \beta \sum_{j \in B} P_j$. We say that the reservation mechanism is exact budget balance if $\beta = 1$.

3.2 Stability properties

A second set of properties describes the stability of the allocation and resilience to manipulation via adding or removing buyers or sellers:

- **Buyer Monotonicity [BM]:** If a new buyer order b_j is added, the revenue of all sellers in S_j does not decrease.
- **Seller monotonicity [SM]:** If a seller increases his reserve price, the revenue of all other sellers does not decrease.

3.3 Incentive compatibility

We next define a set of desirable incentives properties for buyers and sellers.

- **[B-IC] Buyer incentive compatibility:** Buyers maximize their utility by reporting their true valuations to the mechanism.



- **[S-IC] Seller incentive compatibility:** Sellers maximize their utility by declaring true reserve prices and supply levels.

Unfortunately, [B-IC] and [S-IC] cannot be simultaneously achieved in a two-sided market [23, 21] if not at the expense of efficiency. Here, we choose to relax seller incentive compatibility and instead, enforce a fairness constraint among sellers while keeping buyer incentive compatibility.

If we enforce [Eff] and [B-IC], the only mechanism available to decide on the allocation and buyer payments is the VCG mechanism. VCG, however, treats all the sellers as one and therefore doesn't prescribe how the revenue of the auctions should be distributed among the sellers. The central issue in the design of reservation exchange mechanisms is how to distribute the revenue from the VCG auction in a manner that is fair to the sellers. As we will see next, defining a precise notion of fairness that matches our intuition is a quite non-trivial task. First, we show how the most natural definitions fail to capture important situations.

3.4 Seller Fairness and Envy-Freeness

We start by identifying a set of properties that we believe a fair revenue sharing scheme should satisfy. The challenge here is in identifying, when a buyer gets some item at price p , if a seller can stake a claim on this revenue or not. Firstly, a seller s_i can claim revenue only from buyers that are interested in the inventory owned by seller s_i , i.e., $R_i[j] = 0$ for $j \notin B_i$. Also, if a certain buyer b_j never receives goods from seller s_i under any efficient allocation, seller s_i shouldn't be able to claim stake on the revenue of buyer b_j . This is so because even if this seller drops this connection to the buyer, it won't change the set of efficient outcomes. Moreover, the seller may end up getting a lower revenue because of the reduced competition after dropping such a connection.

In order to capture the above notions, we define for each seller s_i , the set $A_i \subseteq B_i$ as the set of buyers that are allocated at least one good from s_i in *some* efficient allocation. We are now able to define the concept of envy-free revenue sharing: roughly speaking, we say that a revenue sharing scheme is envy-free if each seller s_i extracts from A_i more revenue than any other seller with at most the same supply and proportionally more revenue than any seller with higher supply. More specifically, this concept is defined as follows:


- **Envy-free Revenue Sharing [ERS]:** $\forall s_i, s_{i'}, \sum_{j \in A_i} R_i[j] \geq \min(1, \frac{\ell_{s_i}}{\ell_{s_{i'}}}) \cdot \sum_{j \in A_{i'}} R_{i'}[j]$.

If all sellers have the same supply, this boils down to $\sum_{j \in A_i} R_i[j] \geq \sum_{j \in A_{i'}} R_{i'}[j]$. We note that if one is able to design an envy-free mechanism for sellers with unit-supply, this automatically extends to sellers with non-unit supply by the following reduction: transform each seller of supply ℓ_i in ℓ_i unit supply sellers. An envy free allocation in the transformed setting naturally translates to an envy-free allocation in the original setting. For this reason we will assume for the remainder of the paper that sellers are unit supply.

However, as shown in the following example, the above notion of envy-freeness doesn't fully capture the notion of a fair allocation among the sellers:

► **Example 1.** Consider two buyers with valuation $v_1 = 2$, $v_2 = 1$, and demands $d_1 = 2$, $d_2 = 1$. There are two unit supply sellers s_1, s_2 with preference constraints shown in Figure 1. For any mechanism satisfying [Eff], [B-IC] and [ERS], the allocation and payments charged to the buyers must be the one of the VCG mechanism. So, the mechanism sells both items to buyer b_1 at total price 1.





■ **Figure 1** The two sellers receive different revenue.

Since $A_1 = A_2 = \{b_1\}$, ERS imposes to share the revenue equally between s_1 and s_2 . This way to partition of the revenue can be hardly called fair, since the 1 dollar in revenue is caused by the competition with buyer b_2 that is brought to the market by seller s_2 . So a natural intuition is that a ‘fair’ scheme should attribute the 1 dollar in revenue to seller 2.

The above example implies the need to refine the envy-freeness property to incorporate some notion of *which buyers are responsible for putting the price pressure*. To get a handle on such a notion, we consider ascending auctions to refine our envy-freeness property.


3.5 Fairness via ascending auction

The traditional definition of the VCG mechanism is that it allocates according to an efficient allocation and charges each agent according to the externality it imposes on other agents. One problem with this way of defining VCG is that it returns a bundle of items to each agent and a total price but does not specify how much of the payment is attributed to each item. An alternative way to define VCG is via an ascending auction [2, 1], in which there is a price clock p that gradually ascends, and as the price increases items are allocated to buyers. The total payment of the buyer in such a case is the sum over the prices of all individual items, where the price of each item is the value of the price clock when the buyer acquired the item.

The ascending auction description of VCG returns for each buyer b_j his allocation, say $x_j \in \mathbb{Z}_+$ together with x_j prices $p_1 \leq p_2 \leq \dots \leq p_{x_j}$ corresponding to the value of the price clock when he acquires each of those items. In other words, we can describe the outcome of VCG as an ascending auction as a set of n *buying events*, where $n = \sum_i \ell_i$ and each buying event is a pair (b_j, p) indicating that one item was sold to buyer b_j at price p . Note that we assume all items are sold by VCG, which is always the case when we consider proxy buyers as discussed in Section 3.1.

Let E be the set of buying events that represents the outcome of the auction. During its execution, the ascending auction maintains in each step a tentative assignment of buying events to sellers. This allows us to define a bipartite graph between sellers and buying events called the *clinging graph*. We say that a seller s_i is connected to buying event e_j if this seller is tentatively allocated to that buying event in some point of the auction execution. If the auction execution is not unique (because of ties, caused for example by two identical sellers) we consider an edge to be in the clinging graph if for some execution of the auction its corresponding buying event is connected to the corresponding seller.





■ **Figure 2** Each seller is linked in the clinching graph to only one buying event.

We refer to [1, 10, 13, 16] and to the full version of the paper for a formal definition of the ascending auction, but we illustrate its execution for the instance in Example 1. The price clock starts at zero, and at that price the auction is already able to allocate the item in s_1 to buyer b_1 , since he faces no competition on that item. This generates a buying event $(b_1, 0)$ that is associated with seller s_1 . For prices between 0 and 1, both buyers compete for the remaining item. When the price clock reaches 1, buyer b_2 is no longer interested in the remaining item and buyer b_1 can acquire it at price 1, generating a buying event $(b_1, 1)$, which is associated with seller s_2 . There are no ties, so this is the unique execution of the auction, which generates the clinching graph represented in Figure 2.

Given the clinching graph, we are now able to define a stronger notion of envy-freeness based on it. We denote by E_i the set of buying events associated to buyer b_i and by $r_i[j]$ the revenue obtained by seller s_i from buying event $e_j \in E_i$. Let us denote the revenue of unit supply seller s_i by $r_i = \sum_{j \in E_i} r_i[j]$. We now state a new version of Envy-free Revenue Sharing, that we denote by ERSCG, as follows:

- **Envy-free Revenue Sharing from Clinching Graph [ERSCG]:** $\forall s_i, s_{i'}, \sum_{j \in E_i} r_i[j] \geq \sum_{j \in E_{i'}} r_{i'}[j]$.

According to the definition of clinching graph, we obtain for the example of Figure 2 that $E_1 = \{e_1\}$ and $E_2 = \{e_2\}$. The revenue sharing scheme that attributes $r_1[1] = 0$, $r_1[2] = 0$, $r_2[1] = 0$ and $r_2[2] = 1$ is therefore ERSCG. We conclude that the new definition of envy-freeness is able to characterize a fair sharing of the revenue between sellers.

We also define an approximate version of the previous property:

- **α -Envy-free Revenue Sharing from Clinching Graph [α -ERSCG]:** $\forall s_i, s_{i'}, \sum_{j \in E_i} r_i[j] \geq \alpha \sum_{j \in E_{i'}} r_{i'}[j]$.

We conclude by observing that the notion of envy-freeness we introduce can easily be adapted to the original non-unit supply sellers.

4 Impossibility of Envy-freeness and Budget Balance

Before presenting two revenue sharing schemes based on the definition of clinching graph in Sections 5 and 6, we show that envy-freeness and budget balance are indeed contradicting objectives for any revenue sharing scheme based on an efficient allocation.



► **Theorem 2.** *There does not exist any revenue sharing efficient mechanism for the reservation exchange market which is BB and α -ERSCG for $\alpha \geq \sqrt{3} - 1 + \delta \approx 0.732$, for an arbitrary small value $\delta > 0$.*

5 Revenue sharing by the Clinching Ascending Auction.

The first revenue sharing scheme is based on the allocations computed by the clinching ascending auction (CA). We denote by *CA* this revenue sharing scheme.

A detailed description of the use of the clinching ascending auction [2, 1] to compute efficient VCG allocations is given in the full version of the paper. We specifically use a version for matching markets given in [13].

Crucial to the definition of revenue sharing scheme is the notion of *priority order* among sellers that is used in the execution of the CA. Whenever the CA is indifferent about buying between a set of sellers, it decides in "favor" of the seller with lower priority in the precedence order. Intuitively, the seller of higher priority will enjoy a payment that is at least as good as the lower priority seller since the price in the ascending auction is non-decreasing. A priority order between sellers is simply represented by a permutation $\pi \in \Pi(S)$ where $\Pi(S)$ defines the set of all permutations of set S .

We set $r_i^\pi[j] = p_j$ if the execution of CA on permutation π matches e_j to s_i .

Revenue share of seller s_i from buying event e_j is defined as

$$r_i[j] = \mathbb{E}_{\pi \in \Pi(S)}[r_i^\pi[j]]. \quad (1)$$

The revenue of seller s_i is defined as $r_i = \sum_{e_j \in E_i} r_i[j]$. Since the total revenue of the mechanism $REV = \sum_{j \in E} p_j$ is shared between sellers, we state a first property of the revenue sharing scheme CA:

► **Claim 3.** The revenue sharing scheme CA is BB.

We next prove that the revenue sharing scheme CA is not exact ERSCG.

► **Theorem 4.** *The revenue sharing scheme CA is at most 1/2-ERSCG.*

We next prove the approximate envy-freeness of revenue sharing scheme CA.

► **Theorem 5.** *The revenue sharing scheme CA is 1/2-ERSCG.*

We conclude with the properties of buyer monotonicity and seller monotonicity for revenue sharing scheme CA.

► **Theorem 6.** *BM and SM hold for revenue sharing scheme CA.*

6 Revenue sharing by the Eating mechanism.

The eating mechanism is defined as a fractional process in time on the clinching graph $CG = (E \cup S, H)$, $H = \{(e_j, s_i) : e_j \in E_i\}$. At each time $x \in [0, 1]$, the unit supply seller s_i "eats" from the non-exhausted buying event $e_j \in E_i$ of highest payment p_j . Each seller will eat at most up to a fraction of 1. A buying event is exhausted when it has been eaten for 1 unit. The result of the eating mechanism is a fractional assignment $x_i[j] \in [0, 1]$ such that $\sum_{j \in E_i} x_i[j] \leq 1$ for each seller $s_i \in S$ and $\sum_{s_i: e_j \in E_i} x_i[j] \leq 1$ for each buying event $e_j \in E$.



Revenue share of seller s_i from buying event e_j is defined as

$$r_i[j] = x_i[j] \times p_j. \quad (2)$$

The total revenue of seller s_i is also equal to $r_i = \sum_{j \in E_i} r_i[j]$.

It is easy to observe that the revenue shares by the eating mechanism can be computed in polynomial time. We first show that revenue share mechanism EM is envy-free.

► **Theorem 7.** *The revenue sharing scheme EM is ERSCG.*

Proof. Consider any two sellers s_i and $s_{i'}$. At any time $x \in [0, 1]$ of execution of the eating mechanism, seller s_i eats from the non-exhausted buying event $e_j \in E_i$ of highest payment p_j . At the same time x , $s_{i'}$ eats from a buying event $e_{j'}$. Buying event $e_{j'}$ is either outside E_i or it offers payment $p_{j'} \leq p_j$. When s_i stops eating, all $e_j \in E_i$ are exhausted. We conclude that the revenue of s_i on E_i is higher than the revenue of $s_{i'}$ on E_i . ■

It is not clear that in the EM sharing scheme all sellers eat up to 1. We show in the following that EM is not exact budget balance.

► **Theorem 8.** *For any $\epsilon > 0$, the revenue sharing scheme EM is at most $\frac{11+\epsilon}{12+\epsilon}$ -BB.*

► **Theorem 9.** *The revenue sharing scheme EM is $1 - 1/e$ -BB.*

7 Conclusions

The reservation exchange market is an emerging model for internet advertising that brings together multiple publishers and advertisers interested in trading inventories of impressions available in the future. In this work, we present the axioms and the design principles at the basis of mechanisms for reservation exchange markets. The goal we define for these markets is the design of mechanisms that are incentive compatible for buyers, envy-free for sellers, efficient and budget balance. We show that this is possible if one of the requirements of budget balance or envy-freeness is slightly relaxed. Our efficient revenue sharing mechanisms are based on the notion of clinching graph that is a convenient representation of the trades of efficient VCG allocations.

We leave several open problems in the context of reservation exchange markets. First of all, it would be interesting to close some of the gaps on approximate envy-freeness and budget balance of the revenue sharing mechanisms we propose. It is also unknown whether the eating mechanism holds some of the monotonicity properties we define in this paper. Moreover, since the clinching graph seems to provide fundamental insights for designing fair revenue sharing mechanisms, it would be helpful to derive its structure from basic properties of VCG mechanisms.

Acknowledgements. We thank Marek Adamczyk for suggesting the analysis of budget balance of the eating mechanism.



References

- 1 Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, December 2004.
- 2 Lawrence M. Ausubel and Paul R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1:1019–1019, 2002.
- 3 Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. Incentive compatible budget elicitation in multi-unit auctions. In *SODA*, pages 554–572, 2010.
- 4 Craig Boutilier, David C. Parkes, Tuomas Sandholm, and William E. Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 1, AAAI’08*, pages 30–37. AAAI Press, 2008.
- 5 Bowei Chen, Shuai Yuan, and Jun Wang. A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising, ADKDD’14*, pages 1:1–1:9, New York, NY, USA, 2014. ACM.
- 6 E.H. Clarke. Multipart pricing of public goods. *Public Choice*. *Public Choice*, pages 17–33, 1971.
- 7 Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Fifteenth ACM-SIAM Symposium on Discrete Algorithms, SODA ’16*. ACM, 2016. To appear.
- 8 Riccardo Colini-Baldeschi, Monika Henzinger, Stefano Leonardi, and Martin Starnberger. On multiple keyword sponsored search auctions with budgets. In *ICALP (2)*, pages 1–12, 2012.
- 9 Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. Prior-free auctions for budgeted agents. *CoRR*, abs/1212.5766, 2012.
- 10 Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.
- 11 Paul Dütting, Inbal Talgam-Cohen, and Tim Roughgarden. Modularity and greed in double auctions. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC ’14*, pages 241–258. ACM, 2014.
- 12 Michal Feldman and John Lai. Mechanisms and impossibilities for truthful, envy-free allocations. In *Proceedings of the 5th International Conference on Algorithmic Game Theory, SAGT’12*, pages 120–131, Berlin, Heidelberg, 2012. Springer-Verlag.
- 13 Amos Fiat, Stefano Leonardi, Jared Saia, and Piotr Sankowski. Single valued combinatorial auctions with budgets. In *ACM Conference on Electronic Commerce*, pages 223–232, 2011.
- 14 D. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–98, 1967.
- 15 D. B. Gillies. Solutions to general non-zero-sum games. In *In Tucker, A. W.; Luce, R. D.. Contributions to the Theory of Games IV.*, (Annals of Mathematics Studies 40), pages 47–85, 1953.
- 16 Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *STOC*, pages 107–122, 2012.
- 17 Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Clinching auctions with online supply. In *SODA*, 2013.
- 18 Renato Gomes and Vahab S. Mirrokni. Optimal revenue-sharing double auctions with applications to ad exchanges. In *23rd International World Wide Web Conference, WWW ’14, Seoul, Republic of Korea, April 7-11, 2014*, pages 19–28, 2014.
- 19 F. Gul and E Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 56:95–124, 1999.



- 20 Kamal Jain and Christopher A. Wilkens. ebay’s market intermediation problem. *CoRR*, abs/1209.5348, 2012.
- 21 R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, April 1992.
- 22 R. Myerson. Optimal auction design. *Mathematics of operations research*, 6:58–73, 1981.
- 23 R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economics Theory (JET)*, 29:265–281, 1983.
- 24 David C. Parkes and Tuomas Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *In Proceedings of First Workshop on Sponsored Search Auctions*, 2005.
- 25 Mark A. Satterthwaite and Steven R. Williams. The rate of convergence to efficiency in the buyer’s bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):pp. 477–498, 1989.
- 26 Mark A. Satterwhite and Steven R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.
- 27 L.S. Shapley. A value for n-person games. In *H. Kuhn and A.W. Tucker, editors. Proceedings of the 5th International Conference on Algorithmic Game Theory.*, Contributions to the theory of games, pages 120–131, 1953.
- 28 Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- 29 W Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
- 30 William E. Walsh, Craig Boutilier, Tuomas Sandholm, Rob Shields, George L. Nemhauser, and David C. Parkes. Automated channel abstraction for advertising auctions. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*, 2010.

