

One-to-Many Matching Auctions in Platforms

Hemant K. Bhargava ^{*} Gergely Csapó[†] Rudolf Müller [‡]

May 15, 2017

Abstract

Many platforms conduct matching in settings where a one-to-many match is possible, but where a one-to-one match could produce higher revenue because participants value exclusive matches higher than shared matches. This paper studies the problem of designing a set of rules (an auction) for allocation (i.e., matching) and pricing of goods or services (e.g., a sales lead) in such settings. We require the auctions to be deterministic, individually rational, and implementable in dominant strategies. Due to the great demand from practitioners for simple and speedy solutions, we focus on heuristics that provide good revenue relative to the optimal auction. Notably, we demonstrate that even the simplest auctions, such as selling always exclusively or always non-exclusively, produce revenue within a constant factor approximation of the optimal revenue. We identify two different single-dimensional relaxations of the problem, for which we determine the optimal auction using well-known techniques. The relaxations provide revenue bounds that can be used to evaluate the quality of heuristic auctions. We also devise a heuristic mechanism from the class of affine maximizers and demonstrate by means of simulation that it yields revenue very close to the upper bounds, and thus very close to optimality.

^{*}Graduate School of Management, University of California Davis, hemantb@ucdavis.edu

[†]Department of Quantitative Economics, Maastricht University, g.csapo@maastrichtuniversity.nl

[‡]Department of Quantitative Economics, Maastricht University, r.muller@maastrichtuniversity.nl

1 Introduction

Technology-enabled platform marketplaces have stormed into business in the last couple of decades. Platforms facilitate multiple groups of trading partners (say, shoppers and merchants) to congregate, discover, and transact with each other (Choudary et al., 2016). Because platforms focus on *enabling* value creation and exchange, rather than value production itself, a vital function is to *match* consumers and providers (Evans and Schmalensee, 2016). One such platform is the lead marketing firm BuyerLink.com (formerly Reply.com) which matches merchants selling specific products (such as automobiles, real estate, insurance, etc.) with shoppers who have expressed interest in purchasing these products. While many platforms seek to support the entire shopper-merchant transaction (e.g., Uber participates in discovery, matching, fulfillment, payment and post-sales issues), others primarily *facilitate* matchmaking (e.g., Craigslist, AirBnB, match.com, and eBay). In the middle, matching platforms such as Beebell, BuyerLink.com, CreditKarma, and Google Search operate a mechanism through which *they* match a shopper and a merchant. This matching role is crucial from a monetization perspective as well, because the platform gets paid for the match, either the connection itself or some metric of success in the interaction between the shopper and merchant.¹

In many platforms, matching is executed as a one-to-one process, where one shopper is matched with one merchant out of several interested ones. For instance, a platform that serves video advertisements against some news or entertainment content can only show one video ad at a time. Many such platforms employ digital online auctions, picking one winning merchant based on winner and price determination algorithms applied to merchants' bids against the shopper's attributes. Advertising auctions are a popular example of this

¹The terms *shopper* and *merchant* are used as a placeholder for two roles, which could be patient-provider, content consumer-content producer, app developer-smartphone operating system, etc. Specifically, the shopper could be a consumer or a buyer firm, and the merchant could be a firm or an individual.

approach.² The theoretical and technological infrastructure for such one-to-one matching auctions is relatively well studied (Varian, 2009; Milgrom, 2004).

This paper extends one-to-one matching and pushes deeper into the matching role of platforms when multiple merchants can potentially be matched with each shopper. One-to-many matches are relevant because matching platforms provide merchants the *possibility* to sell something to a *potential* shopper (rather than an actual sale). A shopper may intrinsically want to connect with multiple merchants in order to find the best price-quality match for herself. This is common in industries served by **BuyerLink**, such as home loans and automobile sales. Conversely, some platforms pick *one* “best” merchant, but only imperfectly. For instance, **CreditKarma**, a platform for matching shoppers and sellers of financial products, matches each shopper to a single merchant by attempting to predict the likelihood that a shopper would be approved for a loan, credit card or other product by each merchant. This imperfection could be addressed by making multiple matches. The emerging standard is to provide a handful of matches. **BuyerLink**’s policy is to match a shopper with up to 3 merchants. This is also increasingly relevant today in the context of location-based advertising on small mobile devices, exemplified by **Beebell** which connects an event visitor with restaurants and other merchants in proximity to the event.

One-to-many matches can improve welfare by increasing competition for the shopper, and they can boost the platform’s profit by leveraging variations’ in merchants’ valuations for shared vs. exclusive matches. Implementing this opportunity motivates the research questions examined in this paper. First, what matching method should the platform deploy: a) only exclusive matches? (b) only shared matches? or c) pick exclusive and shared match in each instance, based on some pre-defined rules, based on dual bids from sellers for exclu-

²Superficially, advertising systems such as *Google Adwords* match a shopper to multiple advertisers by filling multiple ad slots at once. However, these slots are ranked, hence in this case multiple vertically ordered products are sold to multiple buyers, and each product only to one buyer. Moreover, only one ad is clicked at a time, i.e., the search engine connects and monetizes a web user to one single advertiser at a time.

sive and shared access to the potential buyer? While many platforms practice the simpler approaches (a) and (b), some matching platforms such as BuyerLink implement this hybrid allocation strategy. Second, how should this hybrid auction be designed, specifically what rules should be employed to assign winners and prices? We analyze and present a revenue enhancing design for a hybrid auction and quantify its performance gain over revenue-optimal auctions for approaches (a) and (b). Our design finds a clever way of mixing shared and exclusive allocations, with a hybrid auction design that guarantees revenues for the platform that are close to optimal. The following example elaborates the design challenge for a hybrid allocation strategy.

Example 1. *Consider a shopper in zip code 60173, who has expressed interest in a BMW mini. The platform has bids from 5 dealers in this area who are interested in such shoppers, and can connect the shopper simultaneously with up to 3 dealers. The shared purchase (s_i) row of Table 1 lists the reservation values of the 5 dealers. If dealers actually bid these values, then the shopper would be allotted to dealers D3, D4 and D1. However, if dealers had higher values for an exclusive match (as indicated in the exclusive purchase and exclusivity margin (m_i) rows) then, and assuming that dealers indeed bid these values, the platform would make an exclusive allocation to D1 in Scenario 2, while it would prefer the one-to-many match in Scenario 1. If the platform uses pay-your-bid pricing, the dealers may apply some bid shading rule rather than bid their true values. Moreover, knowing that dealers will bid strategically in this manner, the platform might be better off adopting a different set of rules for determining the matches and corresponding prices.*

Table 1: Exclusive and shared valuations for dealers associated with a matching platform, for access to the platform’s customers who have interest in a BMW mini.

Valuation	D1	D2	D3	D4	D5	D1	D2	D3	D4	D5
Shared purchase (s_i)	10	8	14	11	9	10	8	14	11	9
Exclusive purchase	25	20	20	15	9	36	20	20	15	9
Exclusivity margin (m_i)	15	12	6	4	0	26	12	6	4	0
	Scenario 1					Scenario 2				

When the platform has committed to making a one-to-many shared match, the resulting auction design problem is well-understood. But more generally, platforms recognize both

that a single shopper could simultaneously be auctioned off to multiple merchants and that, sometimes, a single merchant might pay a higher price for exclusive access to the shopper. The platform could therefore choose the type of allocation based on the bids. However, it must propose the auction design in advance without knowing the draws of buyer valuations. In this case, intuitively, the platform should adopt a parameterized design where each instance is *configured* with respect to known segment-specific information. For example, a platform which trades many similar “products” in segmented markets (e.g., millions in the automobile lead marketing case, covering one of over 300 models across one of about 40,000 zip codes, and with different demand profiles for shared vs. exclusive allocations) would announce the same general format across all, but customize the auction rules (e.g., setting a reserve price) based on priors for how car sellers value customer leads for a particular class of models within a particular range of neighborhood zip codes. For either one-to-one matches only, or one-to-many matches only, an optimal reserve price is the best way to maximize expected revenue, and the optimization process is well-understood. This paper extends the analysis for a hybrid design that switches between the two types of allocations.

The essential complexity in this setting lies in the fact that merchants’ reservation values for the item depend on whether or not they become the sole winner. An auction that does not allow for bids that express this fact lacks relevant information for allocating efficiently in all cases and is likely to miss revenue opportunities. Therefore, we model this as an auction design problem with two-dimensional private information, where the first parameter represents the base valuation for being served (called *shared* value), while the second one is the *exclusivity margin*. The total value for being served exclusively is the sum of the shared value and the exclusivity margin. Given a prior on private information (for a particular market segment), the goal is to find the expected revenue maximizing mechanism. We aim to identify a deterministic, implementable in dominant strategies, and ex-post individually rational, auction. Expectations are taken with respect to a prior of private information, that

is a distribution of shared values and exclusivity margins which is *a priori* known to the platform. An additional “soft constraint” could be to obtain auction rules which do not vary significantly across segments. The solutions we provide turn out to satisfy this constraint without adding the constraint as a hard constraint to the design problem. Furthermore, the auction rules are equivalent to a posted-price mechanism whenever only a single buyer happens to be interested in the item.

We resolve the challenges posed by this setting through an innovative three-step approach. First, we formulate the search for an optimal mechanism in the space of mechanisms as an optimization program. Second (because of the challenge in identifying the optimal solution within this space), we develop approaches that generate specialized, parametrized mechanisms whose parameters can be efficiently computed, and the mechanism can be easily implemented. Third, we develop techniques for estimating tight bounds for the original problem and use these bounds to demonstrate the revenue performance of the proposed specialized mechanisms. Conversely, we also demonstrate that our proposed auctions outperform alternative and naïve mechanisms. Additionally, we show that the bounds that we develop are tighter than bounds based on standard methods in literature. Our results require mild conditions, that distributions of shared valuations and exclusive margins are independent, and have monotone hazard rate.

The multi-matching auction market examined in this paper is already employed by some matching platforms, as noted above. Moreover, the underlying construct (multiple sales, when buyers have a margin for exclusive purchase) is applicable to many other information goods because of their non-rivalrous property. Higher value for exclusive purchase can be fueled by the threat of competition, by a sense of privilege, or by special customer preferences (e.g., luxury goods). For example, a prospector would derive higher value from exclusive possession of information regarding a natural resource repository. A retailer may perceive greater value from exclusive right to sell a good because it avoids competition with other

shops. A newspaper advertiser could buy exclusive access on a page, or split the audience by purchasing a fraction of the advertising space. A computing platform can entice more users when a marquée software application is solely available on that platform and not on competing platforms. Mobile search advertisers may, due to the device’s limitations in display size and navigation, be willing to pay substantially more for exclusive promotion, and major search engines have examined designs where advertisers can place two-dimensional bids for receiving the click exclusively or shared (Sayedi, 2012; Jerath and Sayedi, 2012). The insights from this paper are therefore broadly applicable to many business settings.

Our proposed mechanisms and results have several implications for practice. From a practical perspective, the mechanism provides the matchmaking platform with an algorithm for computing allocations and prices after receiving participants’ bids as inputs. The platform owner requires only a reasonable prior about the distributions of valuations in order to configure the algorithm (specifically, by setting reserve prices). Once configured, the algorithm is very efficient at computing the actual allocations and prices for any particular configuration of bids. Moreover, as the platform owner obtains improved information about the distribution, the algorithm can efficiently be reconfigured with new reserve prices. Finally, the mechanism is extremely straightforward for bidders because the optimal equilibrium strategy for them is simply to bid their true valuations.

In the next section relevant works from the literature are discussed. Then we introduce the main concepts and notation for our model in §3, where we specify a mathematical program to represent the optimal dominant strategy incentive compatible auction (DSA). In section 4 we set the foundation for deriving an approximately optimal auction by constructing restricted variants of the original problem; optimal solutions for these designs constitute lower bounds on revenue of the optimal DSA. These bounds are used in section 5 to prove that simple mechanisms can approximate the optimal solution. We conclude that section by constructing a heuristic from the affine maximizer family of mechanisms. Finally, section 6

presents a simulation study on the expected revenue of the discussed mechanisms for various settings.

2 Related Literature

The problem of designing the revenue maximizing auction when private information is given by a single parameter has been solved in the seminal (Myerson, 1981) paper. His approach provides a closed form characterization of the optimal mechanism under mild assumptions for most of the cases. In contrast, for settings with multi-dimensional types, the optimal mechanism design problem quickly becomes intractable and there hasn't been a general framework so far to treat these problems. Some researchers use linear programs to compute approximations (e.g., Cai et al., 2011), but the solutions are non-practical as they must be described by an explicit table of inputs (bids) and outputs (allocations and prices) that is exponential in the number of bidders and items. Thus, even if one finds the optimal mechanism, it is generally too complicated to implement in practice. The optimal mechanism is usually tailored in a complex manner to specific details of the distribution on agent preferences.

Some researchers choose to fold the two-dimensional problem into a single dimensional model (e.g., Deng and Pekeč, 2013), by assuming that the exclusivity margin is either public information or a fixed relative mark-up on, or a constant multiple of their value for shared allocation. Using this assumption the techniques of (Myerson, 1981) can be applied. However, this approach is insufficient for matching platforms because it eliminates the case where the ability to dynamically choose between shared and exclusive allocations (vs. predetermined exclusive-only or shared-only) is most relevant (i.e., when exclusivity margins are heterogeneous, distributed differently across problems instances, and not correlated with shared valuations).

These undesirable properties are reflected in the works of (Rochet and Choné, 1998),

(Armstrong, 1996) and (Manelli and Vincent, 2007). For more critiques and thoughts on optimal mechanism design see (Hartline, 2012).

There are different attempts to circumvent the inherent hardness in multi-dimensional mechanism design. One approach is to stay in the single-dimensional domain and assume that, apart from a single parameter of private information, the rest of the parameters are public information. This path was chosen in (Deng and Pekeč, 2013) with a focus on exclusivity contracts. In that paper an agent is allocated exclusively if nobody else from his direct neighbourhood receives it. The neighbourhood is defined on a publicly known network. Our setting covers the case when this neighbourhood graph is a clique. To get their main result they restrict themselves to clique neighbourhood graphs and single-dimensional valuations, where the value for the shared allocation is private information, while the exclusivity margin is public. The rest of the assumptions are similar to ours: private values are distributed according to a monotone hazard rate distribution and they search for deterministic mechanisms. Due to the single-dimensional assumption they can apply the techniques of Myerson and derive the optimal mechanism, which is the one that maximizes the sum of virtual valuations for each type report. In contrast, our model keeps multi-dimensionality where the Myerson framework doesn't have a bite.

Along this line of research, (Salek and Kempe, 2008) and (Pei et al., 2014) study the problem of selling a digital good with unlimited supply of copies to bidders whose value for the good is decreasing in the number of bidders obtaining it. The function according to which the valuation depreciates is public information, therefore their setting is single-dimensional. (Salek and Kempe, 2008) provides the revenue maximizing Bayes-Nash implementable mechanism based on the Myerson techniques, while (Pei et al., 2014) adapts the prior-free “single-sample” mechanism from (Dhangwatnotai et al., 2010) that yields a constant approximation for that setting. Their theorems hinge on the assumption that types are single-dimensional and are independently distributed according to monotone hazard rate

distributions. Moreover, besides that the “single-sample” mechanism is not deterministic, it does not even extend directly to our setting as applying reserve prices combined with the Vickrey-Clark-Groves mechanism (VCG, see (Vickrey, 1961), (Clarke, 1971) and (Groves, 1973)) is not incentive compatible in general, demonstrated by Example 2 on page 24.

Jerath and Sayedi (2012) studies an extension of GSP for sponsored search advertising patented by Yahoo!. This auction takes two bids as input: one for being displayed among other ads and a second bid for exclusive display. As truth-telling is not dominant strategy in that auction the authors aim to identify and analyse bidding strategies that lead to a Bayes-Nash equilibrium. To render their analysis tractable they restrict themselves to a highly stylised setting including only three agents: two having no exclusivity margin and one whose exclusivity margin is a fixed portion of the shared valuation. They find that allowing advertisers to bid for exclusivity usually increases the search engine’s revenue (because of an increased competition effect) but that revenue may fall under certain conditions. Note that for such a single-dimensional setting the optimal auction format is characterised by (Deng and Pekeč, 2013).

Cai et al. (2011) and Cai et al. (2013) weaken the requirement on implementability by looking for multi-dimensional mechanisms that are Bayes-Nash incentive compatible. They also admit randomized mechanisms, which allows them to use linear programming to acquire a polynomial-time approximation scheme. Their method is not applicable for our problem as it doesn’t handle allocational externalities. Moreover, if we go for deterministic and dominant strategy incentive compatible mechanisms, then we have to face an integer linear program with exponential number of variables and constraints, which renders this approach impractical.

Another line of research focuses on heuristic mechanisms, which possess a succinct description and might exhibit a guaranty on its expected revenue. In relation to our problem the work of Devanur et al. (2011) and Dhangwatnotai et al. (2010) bear relevance. They

deal with multi-parameter mechanism design involving unit-demand bidders, regular type distributions and matroid or downward closed feasibility constraints. They derive simple mechanisms that achieve constant approximations of the optimal revenue. The key point of their proof is to employ the solution from a single-dimensional analogue as an upper bound for the optimal revenue. Then the revenue of their simple mechanism is compared to that upper bound. In spite of the similarities, these results cannot be utilized directly for our problem as their setting cannot accommodate allocational externalities. Even if one treats the different allocations as items, the corresponding feasibility constraints are not downward closed, which is a necessary assumption in their model. The way we derive the second upper bound is reminiscent to the technique developed in Devanur et al. (2011) in the sense that we also split the multi-dimensional agents to single-dimensional representatives. The twist in our method is that in order to handle allocational externalities we endow the representatives with interdependent valuations.

Allocational externalities are similar in nature to interdependent valuations and identity dependent externalities in the sense that they all try to catch the impact of the agents on each other. The implications of externalities have been studied in various settings, see, e.g., Jehiel et al. (1996) and Segal (1999) and Figueroa and Skreta (2011). In Aseff and Chade (2008) the optimal auction for selling two identical goods is considered, where bidders impose externalities on each other. Their model can be parametrized such that it coincides with a special case of Deng and Pekeč (2013), but it is limited to selling two items, multiplicative externality and single-dimensional valuations.

The idea of using affine maximizers to improve on the revenue of the VCG mechanism has appeared mostly in connection with combinatorial auctions. Likhodedov and Sandholm (2004) specifies a special class of affine maximizers and tries to fine-tune its parameters. Tang and Sandholm (2012) considers the case of two bidders and two items and searches for the best parameters for a given affine maximizer. One of our heuristics is also an affine maximizer,

but its parametrization idea differs from that of the two previous papers. Moreover, with the help of simulation and the derived upper bound mechanisms, we can demonstrate in a novel way how small the gap is between the heuristic mechanisms and the optimal revenue.

3 Model Preliminaries

Suppose an item can be sold simultaneously to k out of n bidders (in the set $N = \{1, \dots, n\}$) who each have unit demand for the item.³ The value of bidder (or agent) i for obtaining the item is characterized by his *type* $t^i = (s^i, m^i)$, where s^i is the valuation for getting the item shared with some other agent, while m^i is the margin for exclusive possession, i.e., $s^i + m^i$ is the valuation for receiving the item exclusively. The shared valuation s^i is drawn from the set S^i with cumulative distribution F (with density f), while the exclusivity margin m^i is drawn from M^i with cumulative distribution G (density g). The set of possible types for agent i is $T^i = S^i \times M^i$. We make the following assumptions. F and G are independent and strictly increasing. Agent characteristics are i.i.d, i.e., $T^i = T^j = [0, \tilde{s}] \times [0, \tilde{m}]$ for all $i, j \in N$, where $\tilde{s}, \tilde{m} \in \mathbb{R}^+$. Types are distributed independently among agents, so the distribution of type t^i is $F \times G$ and the distribution of type-tuples is $(F \times G)^n$. The basic notation used in the paper is summarised in Table 3 on page 42 in the appendix.

Following the Revelation Principle (Myerson, 1981) we restrict our attention to direct auctions. Each direct auction (x, p) can be characterized by its allocation rule $x : T \rightarrow \{0, 1\}^n$ and its payment scheme $p : T \rightarrow \mathbb{R}^n$. Let $a(t) = \{i : x^i(t) = 1\}$ be the set of agents ($\subseteq N$) who receive the item for bid t . Note that $a(t)$ would be a singleton under exclusive allocation, and $|a(t)| \geq 2$ represents shared allocation. Given an auction (x, p) and report

³The limit k is often set by policy, and a value of 3 or 4 provides sufficient competition and variety without overloading customers with unwanted sales calls. For instance, lead generation platforms for automobile trades typically match a customer with 3 car dealers.

profile $\hat{t} = (\hat{s}, \hat{m})$, the realized valuation v of agent i having type $t^i = (s^i, m^i)$ is

$$v^i(x(\hat{t}), t^i) = \begin{cases} s^i & \text{if } i \in a(\hat{t}) \text{ and } |a(\hat{t})| \geq 2, & \text{(agent } i \text{ receives a shared allocation)} \\ s^i + m^i & \text{if } a(\hat{t}) = \{i\}, & \text{(agent } i \text{ receives exclusive allocation)} \\ 0 & \text{otherwise.} \end{cases}$$

while the net utility (after subtracting payment) is $u^i(x(\hat{t}), p(\hat{t}), t^i) = v^i(x(\hat{t}), t^i) - p^i(\hat{t})$.

Because v depends only on the allocation (a) component of x , we will occasionally write $v^i(a(\hat{t}), t^i)$ for expositional clarity rather than $v^i(x(\hat{t}), t^i)$.

3.1 Design Goals and Choices

The possibility of allotting a single item to multiple bidders raises, at a high level, the following design choices: (a) ignore the possibility for shared allocation and allocate to the bidder with the highest bid for exclusive purchase, (b) ignore possibility of exclusive allocation and always allocate to bidders with highest shared valuations, (c) pick shared vs. exclusive allocation based on priors about distribution of valuations, and (d) pick shared vs. exclusive allocation for given bids. In each case, the auction design specifies allocation and pricing rules, and then agents make utility-maximizing bids given these pre-announced rules. For our problem, design objectives are a combination of profitability and implementability, i.e., maximizing auction revenue subject to practical considerations and costs to all participants. At a high level, the choice is between mechanisms that promote truthful bidding vs. those under which agents may scheme and shade their bids. In general, the best among the latter kind of auctions has the potential to produce higher revenue, but at higher participation costs for agents because agents must compute how to shade their bids while making assumptions and anticipating behavior of other agents. These costs may be substantial (relative to the simpler “bid your true value” rule), and an iterative procedure can impose substantial costs

on the auctioneer as well. If agents cannot reliably identify their optimal bidding strategies, then realized revenue may fall short of predicted revenue. Hence our joint objective of profitability and implementability steers us in the direction of incentive compatible mechanisms, i.e, rules that maximize profits while preserving encouraging truthful bidding, and being easy for agents to participate in and for the seller to execute in real-time. Formally, this defines a feasible space of auction designs in which truth-telling is the optimal strategy for agents and it leads to non-negative utility for every agent.

Definition 1 (DSIC). A direct auction is dominant strategy incentive compatible (DSIC) if truth-telling is a dominant strategy for each agent: given the other agents' bids, every agent's utility is maximized by bidding truthfully. Formally, a direct auction (x, p) is DSIC if for every i , t^{-i} , t^i and \hat{t}^i it holds that

$$u^i(x(t), p(t), t^i) \geq u^i(x(\hat{t}^i, t^{-i}), p(\hat{t}^i, t^{-i}), t^i).$$

Definition 2 (EPIR). A direct auction is ex-post individual rational (EPIR) if a truth-telling agent has non-negative utility for every report of other agents. Formally, a direct mechanism (x, p) is EPIR if for every i and $t = (t^i, t^{-i})$ it holds that

$$u^i(x(t), p(t), t^i) \geq 0.$$

This combination of rules not only is DSIC and EPIR, but is also socially optimal.

3.2 VCG vs. Dominant Strategy Auctions

A potential candidate for auction design given these objectives is to allocate to maximize agent value (i.e., to allot exclusive or shared based on whether the highest exclusive bid exceeds the sum of the k -highest shared bids), and combine this allocation rule with a VCG pricing rule. Formally, the allocation rule assigns to agent i exclusively if $(\hat{s}^i + \hat{m}^i)$ exceeds (or equals) the sum of the highest- k \hat{s} values, or otherwise to the first k agents in descending

Table 2: VCG allocation and prices for the example from Table 1. Allocation type (shared or exclusive) and allocations (x_i) optimize profit given the bids. Prices charged are a difference of column D_i values (for the two rows labeled “Others’ value ...”), representing $\sum_{j \neq i} v_j$ under efficient allocation, when bidder i is not present, and present, respectively.

	D1	D2	D3	D4	D5	D1	D2	D3	D4	D5
Allocation (x_i)	1	0	1	1	0	1	0	0	0	0
Others’ value if i not present	34	35	30	33	35	34	36	36	36	36
Others’ value if i present	25	35	21	24	35	0	36	36	36	36
Prices (p_i)	9	0	9	9	0	34	0	0	0	0
	Scenario 1					Scenario 2				

order of \hat{s}^j values (using some tie-breaking rule when necessary).⁴ The VCG pricing rule charges every bidder the externality he poses on other bidders, thereby eliminates bid-shading incentives that will arise under a pay-your-bid pricing rule. More precisely, the price bidder i has to pay equals the total value of the efficient allocation for other bidders, would he not have participated in the auction, minus the value that all other bidders get when he is participating. Note that for bidders who do not win the item this results in zero payments.

Table 2 gives the VCG allocations and prices corresponding to Scenarios 1 and 2 of our example. While this design maximizes social value while being easy to implement, it is unlikely to have strong revenue performance. It is well-known that VCG auctions may yield very low revenue (*cite the lovely but lonely Vickrey auction from CA book*) and that, due to Revenue Equivalence, there is no way to adjust payments without losing DSIC and EPIR while maintaining social optimality. Because we seek to retain the advantages of truth-telling and then to maximize revenue (§3.1), we therefore need to give up allocation efficiency and include the allocation rule in our search space. This yields the following multi-matching auction problem:

find a direct auction (allocation and pricing rules) which maximizes expected rev-

⁴Shouldn’t the s ’s and m ’s have “hats” indicating they are bids rather than values? It is the same due to DSIC property, but in principle ... ?? Do we even have a notation? We have \hat{t} but not \hat{s} or m . I added it.

enue while preserving individual rationality (EPIR) and truthful bidding (DSIC).

This statement concisely represents our design goal. We call this design a dominant strategy auction (DSA), and model it as a mathematical program where the parameters in the objective are determined by the priors of the distributions, and the constraints represent incentives and allocation feasibility. Once solved, the decision variables tell for each report of types who receives the item and prices to be paid. Because it is DSIC, agents' reports are simply their true values. The model looks as follows.

$$\max_{x(t), p(t)} \mathbb{E}_t \left[\sum_i p^i(t) \right] \quad (\text{DSA}^*)$$

subject to

$$u^i(x(t), p(t), t^i) \geq u^i(x(\hat{t}^i, t^{-i}), p(\hat{t}^i, t^{-i}), t^i) \quad \forall i, \forall t^{-i}, \forall t^i, \forall \hat{t}^i, \quad (\text{DSIC})$$

$$u^i(x(t), p(t), t^i) \geq 0 \quad \forall i, \forall t, \quad (\text{EPIR})$$

$$\sum_i x^i(t) \leq k \quad \forall t, \quad (\text{SUPPLY})$$

$$x^i(t) \in \{0, 1\} \quad \forall i, \forall t. \quad (\text{BI})$$

Constraint (DSIC) is responsible for dominant strategy incentive compatibility, while (EPIR) enforces ex-post individual rationality. (SUPPLY) sets the upper bound on the number of copies that can be sold, and (BI) ensures that the mechanism is deterministic. Note that with *DSA* we refer to the whole mathematical program. The optimal solution (x, p) for *DSA* will be called DSA^* .

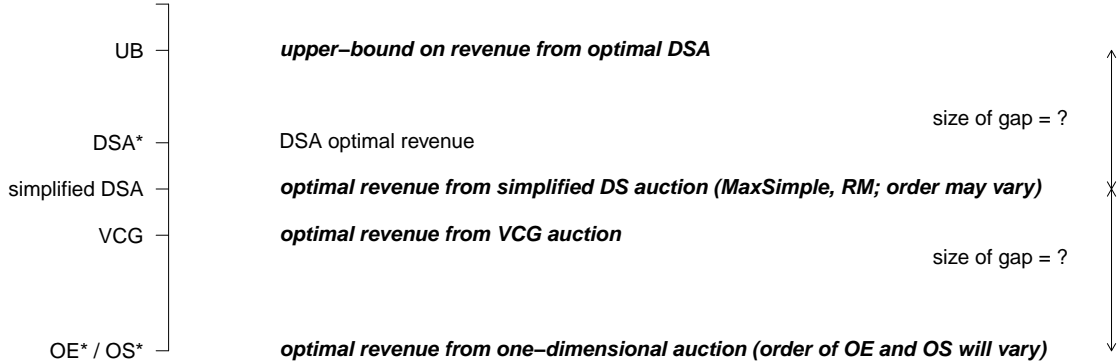


Figure 1: A map of comparative performance for alternative mechanisms. OE and OS (only-exclusive and only-shared) are one-dimensional auctions; MaxSimple is (ex ante) the better of the two. VCG and RM are two-dimensional auctions. Actual comparisons, determined through theoretical and simulation techniques, will be presented in Section 5.

3.3 Optimal DSA and Heuristic Mechanisms

We note that the DSA formulation is, for us, primarily a theoretical construct. It has a huge number of variables and constraints, the latter not even being linear. Even if we could solve it, that solution would not be practical as it would only provide a full list of $x(t)$ and $p(t)$ vectors, one for each t . For us, the program serves a different purpose. On one hand, it enables us to quickly generate alternative practical mechanisms, either by adding certain constraints to this program (see §4.1) or via a special class of DSIC two-dimensional auctions (presented in §4.2). Optimal solutions to these refined programs are therefore feasible solutions to the original *DSA* program, and provide lower bounds to the optimal DSA revenue. On the other hand, we develop relaxations of this program which readily can be optimized, thereby generating upper bounds for the optimal solution of *DSA*. Combining these two features, we show how to achieve an implementable mechanism with provable performance, and indeed we demonstrate in a later section that this mechanism has excellent revenue performance. Fig. 1 visualizes our plan for comparing the alternative mechanisms and bounds.

We end this section with two technical concepts from mechanism design that will play a

role in our construction of heuristics. The first is the so-called monotonicity condition that is imposed solely on allocation rules. The condition is obtained by adding up two (DSIC) constraints such that the first constraint is applied to agent i having type t^i and reporting \hat{t}^i , while the second one is expressed for the same agent having \hat{t}^i and reporting t^i , to yield

$$v^i(x(\hat{t}^i, t^{-i}), \hat{t}^i) - v^i(x(\hat{t}^i, t^{-i}), t^i) \geq v^i(x(t), \hat{t}^i) - v^i(x(t), t^i). \quad (\text{MON})$$

Allocation rules satisfying (MON) will be called monotone. The significance of this property is that designs which fail it cannot have a combination of allocation and pricing rules that together satisfy (DSIC). Hence such designs can be eliminated, and our search can be limited to the space of monotone allocations.

The second is the notion of virtual valuations that plays a central role in the design of revenue maximizing single-item auctions, when valuations for the item are independent and identically distributed. The *virtual valuation* of agent i for real-valued random variable $r^i \in \mathbb{R}$, that is distributed as D , is defined as

$$\phi_D(r^i) = r^i - \frac{1 - D(r^i)}{d(r^i)}. \quad (1)$$

In our context r^i could be s^i , m^i or $s^i + m^i$ for example, so that D would be F , G or $F \times G$. Distributions that have increasing virtual valuation are called regular. An interesting subclass of regular distributions are those having monotone hazard rate (MHR), otherwise known as log-concave probability distributions. In words, these are distributions with tail not fatter than that of the exponential distribution. This class is popular in many areas of economics as its properties usually yield special structure for the problem at hand and hence allow for insightful qualitative implications (see, e.g., Bagnoli and Bergstrom, 2005). Main members of this class are the normal, exponential, some parametrization of gamma, Pareto and the uniform distribution. Regular distributions D have a unique point where

the corresponding virtual valuation hits zero. This value is called the reserve value r_D of D .

Definition 3 (Reserve Value). For a regular distribution D , the reserve value r_D is the unique value for which $\phi_D(r_D) = 0$.

4 Approximate optimal auctions (Revenue Lower Bound)

The DSA problem statement primarily serves the purpose of a theoretical rendering of the optimal DSIC auction. This section develops five different feasible solutions to this problem by constructing five, solvable, restrictions of the original DSA problem. By construction, each method will provide a lower bound for the optimal DSA. Note that the optimal revenue DSA* is not known, hence the performance of these lower bounds will be evaluated by computing *upper* bounds on DSA*, which we do in the next section by relaxing DSA.

The first three designs allow only for one-dimensional bids (OE, OS, MaxSimple, recall Fig. 1). OE and OS commit the platform to either exclusive-only or shared-only allocations for *all* instances or problem categories. Analysis of these two cases is useful in some cases because a platform might prefer a consistent allocation policy, hence this analysis will provide guidance regarding which one to choose. Intuitively, this will depend on the magnitude and distribution of s and m sets across multiple instances. In other cases, a platform might wish to exploit category-specific prior information (e.g., leads for BMW buyers in zip code 93940 during June) and then pick the better of OE and OS based on expected revenue performance given the prior, and then commit to it for that category. This is the MaxSimple mechanism in our framework. For each instance, participants need to provide only one bid (for exclusive or shared allotment, as specified by the platform) but which kind it is can vary across instances. We prove that this auction yields a constant factor approximation of the optimal revenue. The fourth and fifth auctions allow for two-dimensional bids and have allocation and pricing rules that are easy to implement. The first of the two is the VCG mechanism discussed

before. We show that it yields a constant factor approximation of the optimal revenue. For the second two-dimensional auction (RM) we have not proven such an approximation ratio, but show in simulations presented in Section 6 that it performs superior to VCG.

4.1 One-dimensional auctions

The three one-dimensional auctions have an advantage that the revenue optimal mechanism can be found in each case, hence they provide useful lower bounds for the original optimal auction problem (DSA*). For OE (always exclusive allocation) and OS (always shared), the DSA requirement is met by adding respectively, the constraint $\sum_i x^i(t) \leq 1$ (exclusive only) or $\sum_i x^i(t) \geq 2$ (shared only). In the first case we can use the optimal single item auction, i.e., the best out of those mechanisms that allocate to at most one agent. The second is the other extreme: it is the optimal mechanism among those which never allocate exclusively. Notice that these optimal auction problems are single-dimensional as in each of them only one valuation counts (shared or exclusive). Therefore we can utilize for both of them the framework of Myerson, 1981 that gives a characterization of the optimal single-dimensional mechanism. As we need this framework as well for our upper bounds in Section 5, we repeat how these solutions look like in a general setting with unit demand agents and homogeneous, indivisible items. For our purposes, we adapt a generalization of the original result of Myerson, 1981 from Devanur et al., 2011.

Theorem 1 (Myerson, 1981). *Let NU be a set of unit-demand agents, and $A \subseteq 2^{NU}$ the set of feasible allocations, i.e., the set of agents who can be served simultaneously. Assume that agent i 's type t^i is single-dimensional and drawn from set $T^i \subset \mathbb{R}$, with regular distribution D^i . Let $T = \times_i T^i$, and let $D = D^1 \times \dots \times D^n$ be the distribution over T . Then the single-parameter environment (NU, A, T, D) has the following properties.*

I/ For every DSIC and EPIR mechanism (x, p) , where p is the maximal revenue given x and $x^i(t)$ is the probability that agent i is served, the expected value of total payments

can be written as

$$\mathbb{E}_t \left[\sum_i p^i(t) \right] = \mathbb{E}_t \left[\sum_i \phi_D(t^i) x^i(t) \right].$$

II/ The revenue-maximizing DSIC and EPIR allocation x can be characterized as $x^i(t) = 1$ if $i \in a(t)$ and $x^i(t) = 0$ otherwise, where

$$a(t) = \arg \max_{a \in A} \left\{ \sum_{i \in a} \phi_D(t^i) x^i(t) \right\}.$$

With the help of Theorem 1 we can directly define the first two one-dimensional auctions.

Definition 4 (Optimal exclusive auction (OE)). Let

$$a_e(t) = \arg \max_{\substack{a \subseteq N \\ |a| \leq 1}} \left\{ \sum_{i \in a} \phi_C(s^i + m^i) \right\}.$$

Ties are broken arbitrarily in case of multiple optimal arguments. Then OE is given by its allocation and payment rule $(x_{\text{OE}}, p_{\text{OE}})$ as

$$x_{\text{OE}}^i(t) = \begin{cases} 1 & \text{if } i \in a_e(t) \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

$$p_{\text{OE}}^i = \begin{cases} \min\{\hat{s}^i + \hat{m}^i \mid x_{\text{OE}}^i((\hat{s}^i, \hat{m}^i), t^{-i}) = 1\} & \text{if } x_{\text{OE}}^i(t) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Lemma 1. *If F and G have MHR, then OE is feasible for (DSA*) and achieves the highest expected revenue among those mechanisms which allocate to at most one agent. Moreover, its expected revenue can be expressed as*

$$\text{Rev}(\text{OE}) = \mathbb{E}_t \left[\sum_i \phi_C(s^i + m^i) x_{\text{OE}}^i(t) \right]. \quad (4)$$

Proof. It follows directly from Theorem 1 given that C has MHR as the set of MHR distributions is closed under convolution (see Barlow et al., 1963). \square

Definition 5 (Optimal shared auction (OS)). Let

$$a_s(t) = \arg \max_{\substack{a \subseteq N \\ |a| \leq k, |a| \neq 1}} \left\{ \sum_{i \in a} \phi_F(s^i) \right\}$$

Ties are broken arbitrarily in case of multiple optimal arguments. Then OS is given by its allocation and payment rule (x_{OS}, p_{OS}) as

$$x_{OS}^i(t) = \begin{cases} 1 & \text{if } i \in a_s(t) \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

$$p_{OS}^i = \begin{cases} \inf\{\hat{s}^i \mid x_{OS}^i((\hat{s}^i, m^i), t^{-i}) = 1\} & \text{if } x_{OS}^i(t) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Lemma 2. *OS is feasible for (DSA*) and achieves the highest expected revenue among those mechanisms which never allocate exclusively. Moreover, its expected revenue can be expressed as*

$$Rev(OS) = \mathbb{E}_t \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right]. \quad (7)$$

Proof. It follows directly from Theorem 1. Note that Myerson requires only that for each agent the allocation is monotone non-decreasing in its own type given the others', therefore the fact that OS never allocates for singletons is not an issue for invoking Theorem 1. \square

The allocation rule of OS deserves some explanation. By the requirement that we always allocate to at least two bidders, bidders might win an item even if their virtual value is below 0, or in other words, their value is below the reserve price. Due to this bidders will also likely pay prices that are below the reserve prices. While this sounds counter intuitive it is necessary to preserve incentive constraints. Indeed, a bidder who happens to get an exclusive allocation would enjoy extra value from this allocation. In a setting where he and all other bidders have shared value below the reserve price, he would have incentives to report a higher shared value than his true one, securing an exclusive allocation. This demonstrates

that platforms who reduce for the sake of simplicity the expressiveness of bids need to design allocations carefully in order to not invite strategic bidding related to information that is not revealed in the auction.

As noted earlier, OE and OS both provide useful lower bounds for (DSA*). However, as we shall show later, the quality of these bounds varies considerably across the parameter regions. Because the quality variation generally moves in opposite directions, it becomes worthwhile to combine the two auction formats into the MAXSIMPLE defined below.

Definition 6 (MAXSIMPLE mechanism). For any instance of the exclusivity auction problem MAXSIMPLE calculates the expected revenue of OS and OE, then commits the mechanism with the highest value.

Note that the choice in MaxSimple is made upfront for given priors F and G and not for every type realization, therefore the actions of the agents don't influence which of the two single item auctions is executed. Despite the fact that MAXSIMPLE capitalizes only on one type of bid and allocation, a later result (Theorem 4, which can be established only after setting up additional results) reveals that the seller can still capture a constant fraction of the optimal revenue by implementing it. This is useful for practitioners because they can keep the bidding language and the rules of the auction simple for only a small sacrifice in the revenue. Furthermore, there is only one parameter (the reserve price) to change when priors change.

4.2 A two-dimensional auction: RM

Although MAXSIMPLE is intuitive and simple, it is a combination of two different mechanisms. Therefore if a platform provider has multiple similar products to offer, then changing the rules and the type of allocation product by product, in order to derive revenue advantage from the alternate rule,⁵ can create confusion among the buyers. Besides, MAXSIMPLE

⁵Note: Even for the same “product” the demand profiles can vary when the product is repeated at a different date, day of week, time, location etc., potentially motivating switching the auction format.

is only single-dimensional, thus it doesn't take advantage of the extra information and the wider offer menu it is possible to utilize by taking into account both types of valuations, and allowing for both types of allocations. The next natural step accordingly, is to devise mechanisms that consider two types of bids: one for shared and one for exclusive allocations.

One obvious choice for an auction that takes both dimensions of types into account would be VCG, discussed earlier in §3 and illustrated in Table 2. It is DSIC and achieves the highest revenue among those mechanism that maximize welfare. However, the drawback is that its expected revenue can be quite low. We seek here to improve on the revenue performance of VCG. We propose a new mechanism (RM) along the lines developed in the OE and OS methods, while preserving truthful bidding. The combination will also have the advantage that it selects based on actual reports for exclusive or shared allocation in each occurrence, rather than deciding *ex-ante* which mechanism to choose as done in the MAXSIMPLE Mechanism.

In order to define the auction we recall the definition of an affine maximizer, also called Generalized VCG auctions.

Definition 7 (Affine maximizer auctions). Let A_N be the set of feasible allocations for agents in N . A mechanism is an affine maximizer if there are $\gamma^a, \forall a \in A$ and $\lambda^i > 0, \forall i \in N$ such that for each t its allocation is $a(t)$, where

$$a(t) \in \arg \max_{a \in A_N} \left[\gamma^a + \sum_{i \in N} \lambda^i v^i(a, t^i) \right]. \quad (8)$$

Similarly, define

$$a(t^{-i}) \in \arg \max_{a \in A_{N \setminus i}} \left[\gamma^a + \sum_{j \in N \setminus i} \lambda^j v^j(a, t^j) \right].$$

Then the payment rule of the affine maximizer mechanism is equal to

$$p^i(t) = \frac{1}{\lambda^i} \left[\sum_{j \neq i} \lambda^j v^j(a(t^{-i}), t^j) + \gamma^{a(t^{-i})} - \sum_{j \neq i} \lambda^j v^j(a(t), t^j) - \gamma^{a(t)} \right]. \quad (9)$$

Fact 1. *Affine maximizer mechanisms are DSIC and EPIR.*

This means that any affine maximizer is a feasible solution of DSA*. Choosing $\lambda_i = 1$ for all $i \in N$ and $\gamma_a = 0$ for all $a \in A$ yields VCG.

We observe that OS and OE auction are affine maximizers. To see this for OS, choose A to be the set of all allocations that allocate to at least 2 and at most k bidders, and observe that the two rules

$$a(t) = \operatorname{argmax}_{2 \leq |a| \leq k} (-|a|r_F + \sum_{i \in a} s^i)$$

and

$$a(t) = \operatorname{argmax}_{2 \leq |a| \leq k} \sum_{i \in a} \phi_F(s^i)$$

choose the same allocation, when the same tie-breaking rule is used, and when F satisfies MHR. The same applies to OE and the affine maximizer with A being equal to all singletons, and setting $\gamma_a = -r_C$.

Using these observations we define the Reserve Mechanism (RM) as the mechanism that will point wise choose either the allocation as OE or the allocation as OS.

Definition 8 (Reserve mechanism (RM)). Let r_F and r_C denote the shared and exclusivity reserve values respectively (see Definition 3). Then RM is defined as an affine maximizer, where $\lambda^i = 1$ for all i , $A = \{a \subseteq N \mid |a| \leq k\}$, $\gamma^{\{i\}} = -r_C$ for all i , $\gamma^\emptyset = 0$ and $\gamma^a = -r_F|a|$ for all $a \subseteq N \mid 2 \leq |a| \leq k$.

Observe that RM either allocates as OE or as OS, depending which of the ones yields more welfare. Prices paid are also higher because the first term in the price determination is also higher.

We close this section by observing that we cannot simply set reserve prices for exclusive and shared bids, eliminate bids below the reserves, and then allocate efficiently with respect to the remaining bids. This rule would not be incentive compatible in general demonstrated by Example 2.

Example 2. Let $N = \{1, 2\}$, $k = 2$, $t^1 = (3, 3)$, $t^2 = (2, 1)$ and $r_F = 0, r_C = 5.5$. Then the efficient allocation that meets the reserve is that agent 1 receives the item exclusively. Now,

set $\hat{t}^1 = (1, 4)$ and note that the efficient allocation that satisfies the reserve is to share the item between agent 1 and 2. This allocation violates monotonicity condition (MON).

RM is an appealing choice for platform providers because it is simple and fast enough to be implemented, but complex enough to be able to capitalize on the particularity of the multi-dimensional valuations. The reserve prices r_F and r_C are determined uniquely for each “product” based on knowledge about the distribution of reservation prices for that product. For instance, for Reply’s lead marketing auctions for automotive sales, this would mean computing these values for each “make” and “model” combination for groups of geographical locations which are similar in their distribution of reservation prices. Notably this requires no more information than that needed for designing a single-dimensional mechanism. Moreover, as the experimental evaluations in the next section highlight, its expected revenue is very close to the optimal one regardless of the number of agents.

5 Theoretical Upper Bounds on Revenue

The previous section has developed, for the one-to-many matching problem, several potential auction formats that are easy to implement and computationally tractable. As we show later, the revenue performance of these auction formats can be evaluated, thereby establishing conditions under which each format outperforms the others. However, this evaluation does not establish whether the winning format is “good enough” relative to the optimal value provided by the optimal solution to problem *DSA* (the optimal design within the space of DSIC mechanisms). However, there are no known techniques for identifying the optimal design for this two-dimensional problem under general valuations. Therefore, our approach to estimating the revenue performance of the heuristic mechanisms is to establish and compare against upper bounds to (*DSA**). It matters therefore to derive upper bounds that are tight across the parameter space.

A straightforward upper bound is the optimal expected welfare, as it is always larger or equal to the optimal revenue due to the EPIR assumption. The problem is that the gap between the two can be significant, therefore it is a too loose measure most of the time. See for example graph 7 on page 36 for justification, where the optimal revenue is less than 60% of the optimal welfare. In order to find tighter bounds we take the approach of relaxing the conditions of (DSA*) such that the resulting optimization problem admits a closed form optimal solution. The plan is to relax the problem such that the resulting setting is single-dimensional, because in such environments the framework of Myerson, 1981 solves the revenue maximization problem.

In this section we construct two such non-trivial upper bounds. For notational convenience the expected revenue achieved by any mechanism MECH will be symbolized by $Rev(\text{MECH})$.

5.1 Relaxation 1: Public exclusivity margin

We establish the first upper bound by solving (DSA*) under the assumption that the exclusivity margin is public knowledge. Because m^i is known, each agent i needs to report only their shared value s^i . The optimization problem reflecting this change can be written as

$$\max \mathbb{E}_t \left[\sum_i p^i(t) \right] \tag{UBM}$$

subject to

$$u^i(x(t), p(t), t^i) \geq u^i(x((\hat{s}^i, m^i), t^{-i}), p((\hat{s}^i, m^i), t^{-i}), t^i) \quad \forall i, \forall t^{-i}, \forall t^i = (s^i, m^i), \forall \hat{s}^i, \quad (\text{DSIC2})$$

$$u^i(x(t), p(t), t^i) \geq 0 \quad \forall i, \forall t,$$

$$\sum_i x^i(t) \leq k \quad \forall t,$$

$$x^i(t) \in \{0, 1\} \quad \forall i, \forall t.$$

We will refer to the optimal mechanism for (UBM) as UBM^* . Note that the only difference between (DSA^*) and (UBM) is that in (DSIC2) agent i gets only incentives not to lie about valuation s^i . Therefore (UBM) is clearly a relaxation of (DSA^*) , which means that each feasible solution for (DSA^*) is feasible for (UBM), moreover, $\text{Rev}(\text{UBM}^*) \geq \text{Rev}(\text{DSA}^*)$. In order to define UBM^* succinctly some additional notations are introduced. For given type profile $t = (s, m)$ define $a_s(t)$ and $a_e(t)$ to comprise the set agents with the highest contribution to the expected revenue.

$$a_s(t) = \arg \max_{\substack{a \subseteq N \\ |a| \geq 2 \\ |a| \leq k}} \left\{ \sum_{i \in a} \phi_F(s^i) \right\} \quad (\text{set of agents with } k\text{-highest shared virtual values}),$$

$$a_e(t) = \arg \max_i \{ \phi_F(s^i) + m^i \} \quad (\text{agent with highest sum of shared virtual value and exclusivity margin}).$$

Ties are broken arbitrarily in case of multiple optimal arguments. Note that $\sum_{i \in a_s(t)} \phi_F(s^i) \geq 0$ as the empty set is also a solution. The k -highest virtual valuations (for shared possession) add up to $\sum_{i \in a_s(t)} \phi_F(s^i)$, while the highest sum of shared virtual value and exclusivity margin is $\phi_F(s^{a_e(t)}) + m^{a_e(t)}$. The next theorem demonstrates the advantage of this single-

dimensional relaxation by showing that the optimal expected revenue can be expressed in terms of virtual valuations. Moreover, the set of winners in the optimal auction is equal to either $a_s(t)$ or $a_e(t)$, dependent on which one contributes the most to the expected revenue.

For cases when only shared valuations are private information and exclusivity margins are public constants Deng and Pekeč, 2013 characterises the optimal DSIC and EPIR mechanism by building on the framework of Myerson, 1981. We rely on their result to determine the optimal mechanism for (UBM).

Theorem 2. *If F has MHR and the exclusivity margins are public information, then the allocation rule of UBM* given type profile $t = (s, m)$ can be characterized as*

$$x^i(t) = 1 \text{ for } \begin{cases} i \in a_s(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) \geq \phi_F(s^{a_e(t)}) + m^{a_e(t)} \\ i = a_e(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) < \phi_F(s^{a_e(t)}) + m^{a_e(t)} \end{cases} \quad (10)$$

and $x^i(t) = 0$ otherwise. The expected revenue under this allocation rule is equal to

$$\mathbb{E}_t \left[\sum_i p^i(t) \right] = \mathbb{E}_t \left[\sum_i \left(\phi_F(s^i) + m^i \prod_{j \neq i} (1 - x^j(s^i, m^i, t^{-i})) \right) x^i(s^i, m^i, t^{-i}) \right]. \quad (11)$$

The proof can be found in the appendix. In general, the optimal solution of (UBM) might yield strictly higher revenue than (DSA*). This is due to the fact that the optimal allocation in (UBM) is not always feasible for (DSA*) as the following example shows.

Example 3. Let $N = \{1, 2\}$, $s^i \sim U(0, 1)$, $m^i \sim U(0, 1)$ for all $i \in N$. Assume that $t^1 = (0.3, 0.3)$, $t^2 = (0.4, 0.1)$. Then we have that $\phi_{U(0,1)}(s^1) + m^1 = -0.1$, $\phi_{U(0,1)}(s^2) + m^2 = -0.1$, $\phi_{U(0,1)}(s^1) = -0.4$ and $\phi_{U(0,1)}(s^2) = -0.2$, therefore according to UBM* nobody gets anything. Now, change the valuations of agent 1 such that $\hat{t}^1 = (0.5, 0.05)$. As $\phi_{U(0,1)}(\hat{s}^1) + \hat{m}^1 = 0.05$ and $\phi_{U(0,1)}(\hat{s}^1) = 0$ agent 1 receives the item exclusively. This allocation rule violates monotonicity condition (MON), which means that there is no payment scheme such that the incentive compatibility constraints are satisfied. Therefore it cannot be part of a solution for (DSA*).

5.2 Relaxation 2: Bound with representatives

The derivation of the second upper bound mechanism is more convoluted. For every instance of (DSA*) we define an instance of an alternative setting by introducing two representatives for each agent: one for each dimension of his type. To distinguish this setting in the notation, we add a “bar” to each item, as in \bar{t} .

Definition 9 (Representative environment). For any instance of (DSA*) given by the set of agents N , the set of types T and distributions of types F, G , the representative environment is defined as follows:

- For each $i \in N$ introduce i_s and i_e , such that the set of agents is $\bar{N} = \bar{N}_s \cup \bar{N}_e$, where $\bar{N}_s = \{1_s, \dots, i_s, \dots, n_s\}$ and $\bar{N}_e = \{1_e, \dots, i_e, \dots, n_e\}$.
- $\bar{A} = \{a \subseteq \bar{N}_s \mid |a| \leq k\} \cup \{i_e\}_{i_e \in \bar{N}_e}$, the set of feasible allocations
- $\bar{T}^{i_s} = S^i$ for all $i_s \in \bar{N}_s$, $\bar{T}^{i_e} = M^i$ for all $i_e \in \bar{N}_e$, $\bar{T} = \times_{i \in \bar{N}} \bar{T}^i$, the set of type profiles
- Let the valuations for all $t \in \bar{T}$, for all $i \in \bar{N}$ and for allocation rule \bar{x} be:

$$\bar{v}^i(\bar{x}(t^i, t^{-i}), (t^i, t^{-i})) = \begin{cases} s^{i_s} & \forall i_s \in \bar{N}_s, \text{ if } \bar{x}^{i_s} = 1, \\ s^{i_s} + m^{i_e} & \forall i_e \in \bar{N}_e, \text{ if } \bar{x}^{i_e} = 1, \\ 0 & \forall i \in \bar{N} \text{ otherwise.} \end{cases}$$

- $\bar{u}^i(\bar{x}(t^i, t^{-i}), \bar{p}(t^i, t^{-i}), (t^i, t^{-i})) = \bar{v}^i(\bar{x}(t^i, t^{-i}), (t^i, t^{-i})) - \bar{p}^i(t^i, t^{-i})$, the utility function
- $t^i \sim F$ for all $i_s \in \bar{N}_s$, $t^i \sim G$ for all $i_e \in \bar{N}_e$

Note that agent i_e exhibits informational externality in his valuation. In such settings dominant strategy incentive compatibility is too demanding as it requires truthfulness for every possible report of the other agents, leaving small room for non-trivial mechanisms (see for example Roughgarden and Talgam-Cohen, 2013). Therefore we relax the incentive compatibility condition to ex-post incentive compatibility (EPIC), which requires truthfulness only for every true type of the other agents.

Definition 10 (EPIC). A direct mechanism is ex-post incentive compatible (EPIC) if for each agent truth-telling is a dominant strategy given that the others report truthfully. Formally, a direct mechanism (x, p) is EPIC, if for every i , t^{-i} , t^i and \hat{t}^i it holds that

$$u^i(x(t), p(t), (t^i, t^{-i})) \geq u^i(x(\hat{t}^i, t^{-i}), p(\hat{t}^i, t^{-i}), (t^i, t^{-i})).$$

Another difference compared to the principles of (DSA*) is that we allow for allocations where only one shared agent receives the item. This relaxation increases the set of feasible allocations, hence the possibility to generate more revenue.

The revenue optimization problem for the representative environment can be stated as

$$\max \mathbb{E}_t \left[\sum_i \bar{p}^i(t) \right] \quad (\text{UBR})$$

subject to

$$\bar{u}^i(\bar{x}(t), \bar{p}(t), (t^i, t^{-i})) \geq \bar{u}^i(\bar{x}(\hat{t}^i, t^{-i}), \bar{p}(\hat{t}^i, t^{-i}), (t^i, t^{-i})) \quad \forall i, \forall t^i, \forall \hat{t}^i, \forall t^{-i} \quad (\text{EPIC})$$

$$\bar{u}^i(\bar{x}(t), \bar{p}(t), t^i) \geq 0 \quad \forall i, \forall t \quad (\text{EPIR2})$$

$$(1 - \bar{x}^{j_e}(t))k \geq \sum_{i_s \in \bar{N}_s} \bar{x}^{i_s}(t) \quad \forall t, \forall j_e \in \bar{N}_e \quad (\text{FeasA})$$

$$\sum_{i_e \in \bar{N}_e} \bar{x}^{i_e}(t) \leq 1 \quad \forall t \quad (\text{FeasB})$$

$$\bar{x}^i(t) \in \{0, 1\} \quad \forall i, \forall t.$$

(EPIC) is responsible for ex-post incentive compatibility, while (EPIR2) for ex-post individual rationality. (FeasA) ensures that at most k items are allocated to agents in \bar{N}_s , and that there cannot be allocations where agents both from \bar{N}_s and \bar{N}_e receive an item. (FeasB) states that at most one agent from \bar{N}_e can be served. We will refer to the optimal mechanism for (UBR) as UBR*.

There are more ways of defining a single-dimensional variant of (DSA*), but this formu-

lation is particularly useful as the following proposition demonstrates.

Proposition 1. *For any feasible mechanism (x, p) for (DSA*), there is a feasible mechanism (\bar{x}, \bar{p}) for (UBR) such that $Rev(\bar{x}, \bar{p}) \geq Rev(x, p)$.*

Here, only the intuition behind the proof is given, for technical details the reader is referred to the appendix. As each representative corresponds to a particular dimension of the original agents' valuation, given any allocation in (DSA*) it is easy to induce a feasible allocation for (UBR). The incentive constraints (DSIC) give a strong structure to the allocation rule that is carried over to the single-dimensional variant. Utilizing on that we can construct a payment scheme that is not just EPIC together with the induced allocation rule, but also dominates in expectation the payments for (DSA*). The technical challenges in the proof are due to the fact that the desired structure of the allocation rule is lacking on a null set of types.

It is immediate from Proposition 1 that the optimal expected revenue of (UBR) serves as another non-trivial upper bound for $Rev(\text{DSA}^*)$. As (UBR) is single-dimensional, it is possible to provide a closed form solution for UBR^* in a similar way as before. For given type profile t let

$$\bar{a}_s(t) = \arg \max_{\substack{a \subseteq \bar{N}_s \\ |a| \leq k}} \left\{ \sum_{i_s \in a} \phi_F(t^{i_s}) \right\}$$

$$\bar{a}_e(t) = \arg \max_{i_e \in \bar{N}_e} \{t^{i_s} + \phi_G(t^{i_e})\}.$$

Ties are broken arbitrarily in case of multiple optimal arguments. Note that $\sum_{i \in \bar{a}_s(t)} \phi_F(t^{i_s}) \geq 0$ as the empty set is also a solution.

Theorem 3. *When F and G have MHR, the allocation rule of UBR^* represented by \bar{x} , for*

type profile t , is computed as follows,

$$\bar{x}^i(t) = 1 \text{ for } \begin{cases} i \in \bar{a}_s(t) & \text{when } \sum_{j_s \in \bar{a}_s(t)} \phi_F(t^{j_s}) \geq \max_{j_e \in \bar{N}_e} \{t^{j_s} + \phi_G(t^{j_e})\} \\ i \in \bar{a}_e(t) & \text{when } \sum_{j_s \in \bar{a}_s(t)} \phi_F(t^{j_s}) < \max_{j_e \in \bar{N}_e} \{t^{j_s} + \phi_G(t^{j_e})\} \end{cases}$$

and $\bar{x}^j(t) = 0$ otherwise. The optimal revenue under this allocation is

$$Rev(\text{UBR}^*) = \mathbb{E}_t \left[\sum_{i_s \in \bar{N}_s} \phi_F(t^{i_s}) \bar{x}^{i_s}(t) + \sum_{i_e \in \bar{N}_e} (t^{i_s} + \phi_G(t^{i_e})) \bar{x}^{i_e}(t) \right]$$

such that $Rev(\text{UBR}^*) \geq Rev(\text{DSA}^*)$.

5.3 Significance of Relaxations

The two relaxations provide us two alternative upper bounds. To get a sense for the quality of these bounds, we compare them against the welfare-maximizing solution, which also provides an upper bound (albeit a weak one) on optimal revenue. In order to do that we calculated the revenue of UBR^* and UBM^* along with the optimal welfare for different instances by means of simulation (more details on the technique in section 6). The results are depicted in Figure 2 and 3. In any comparison between welfare-maximizing and other solutions, note that for a “neutral” case of optimizing the revenue from a linear demand function with uniform pricing, the optimal monopoly revenue is no more than half the welfare-maximizing value.

UBR denotes the upper bound with representatives, whilst UBM stands for the bound with public exclusivity margin. As the displayed revenue ratios are taken over the optimal welfare, it is apparent that both relaxation bounds are much tighter than the welfare upper bound. In particular, in Figure 2 UBM is around half of the optimal welfare. Furthermore, dependent on the support of the two valuations, the difference between UBR and UBM might be significant. In general, when exclusivity margins are small relative to shared valuations, UBM^* is tighter, while UBR^* is closer to the optimal revenue when exclusivity

margins dominate.

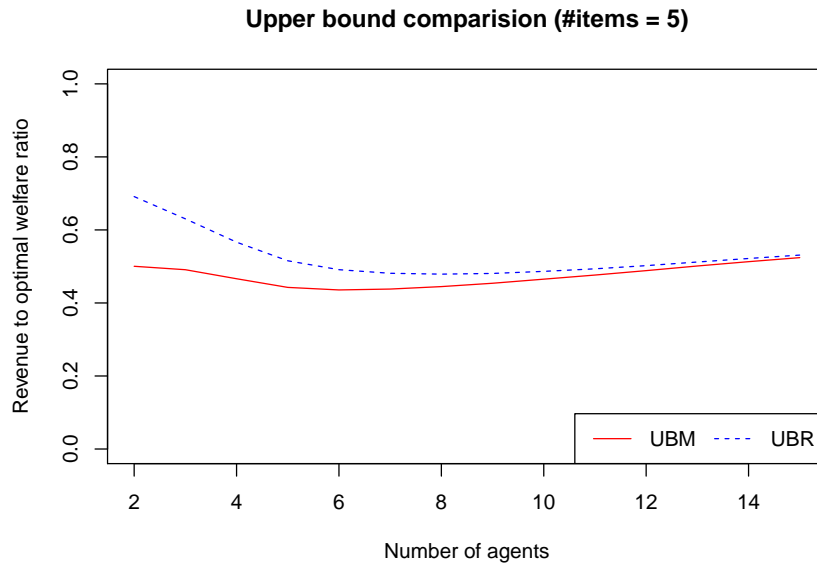


Figure 2: Revenue over optimal welfare of the two upper bounds ($s^i \sim Exp(1)$, $m^i \sim Exp(2)$).

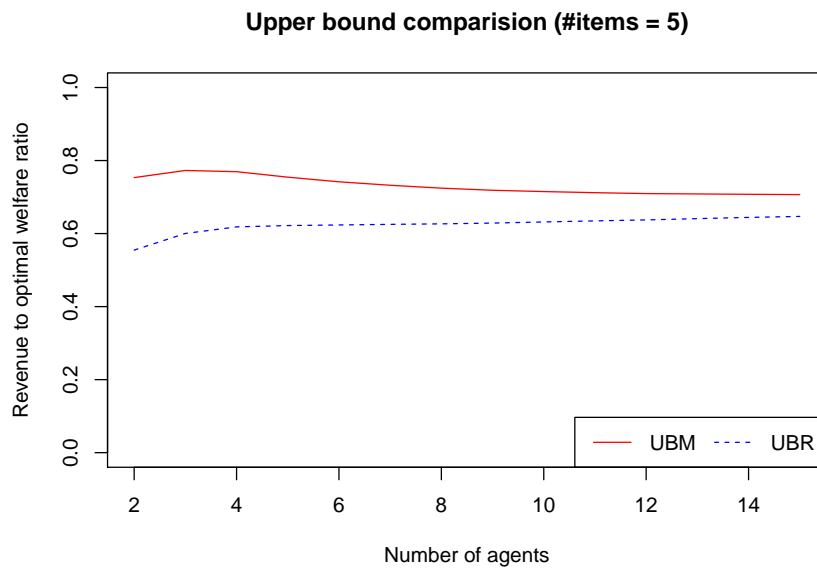


Figure 3: Revenue over optimal welfare of the two upper bounds ($s^i \sim Exp(1)$, $m^i \sim Exp(0.5)$).

UBM can as well be used to prove that MAXSIMPLE, introduced previously as a lower bound, achieves a constant factor approximation of DSA^* (see the appendix for proof). Note that, as stated earlier, the monopoly optimal revenue under a “neutral” demand setting (linear demand, zero marginal cost) is $\frac{1}{2}$ of optimal welfare, hence a heuristic that attains, in its worst case, $\frac{1}{4}$ of optimal welfare actually has a worst-case revenue performance of 50%.

Theorem 4. *Take an instance of the exclusivity auction problem, where F and G are MHR distributions. Then expected revenue of MAXSIMPLE is at least $\left(\frac{1-1/H_n}{2-1/H_n}\right)$ the expected revenue of the optimal mechanism. The minimum of this approximation ratio is 1/4, attained in case of two bidders. Moreover, this approximation ratio holds for the optimal Bayes-Nash implementable mechanism as well.*

6 Performance Analysis

This section describes and compares the performance of the discussed mechanisms across a spectrum of type distributions. The expected revenues are acquired via computational simulations in the following way. We sample from the type distribution many times and calculate the payment for each type report. The expected revenue is then estimated by the average of the simulated payments. To show the robustness of the results we also derive the 99% confidence intervals for each point estimate. The upper (lower) limit of the confidence interval is calculated as

$$\beta + (-)t_\alpha \frac{\sigma}{\sqrt{n^*}}, \quad (12)$$

where n^* is the number of simulations, β is the average and σ is the standard deviation of the simulated values, while t_α is Student’s t -distribution value for the given critical level α and degrees of freedom $n^* - 1$. For our experiments we set $\alpha = 99\%$ and $n^* = 20000$, therefore the corresponding t -value is $t_{99\%} \approx 2.6$. \bar{a} and σ are estimated from the samples for each simulation. For each setting we included the upper bound on the optimal revenue

(UB), which is calculated as the minimum of $Rev(UBM^*)$ and $Rev(UBR^*)$.

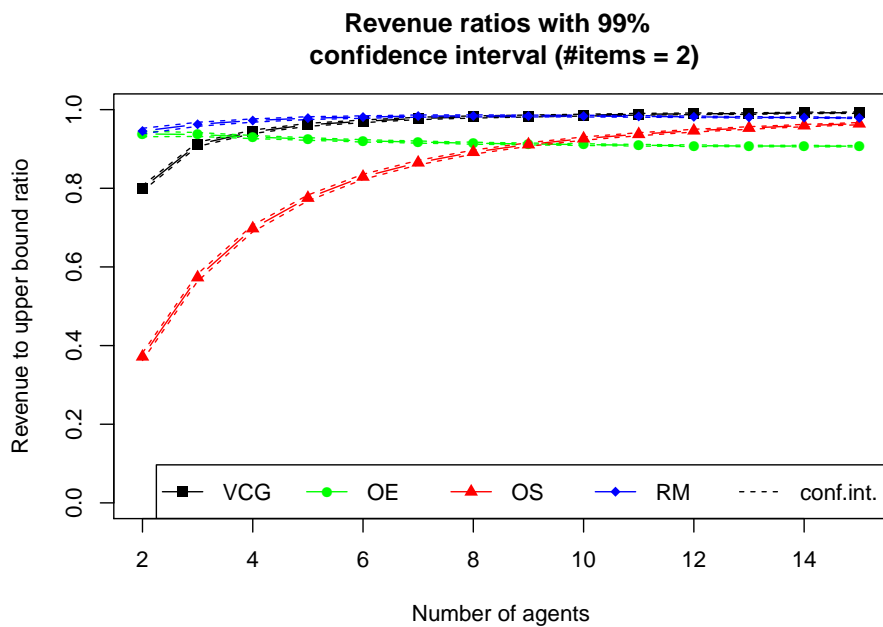


Figure 4: Revenue to upper bound ratios ($s^i \sim U(0, 1)$, $m^i \sim U(0, 1)$).

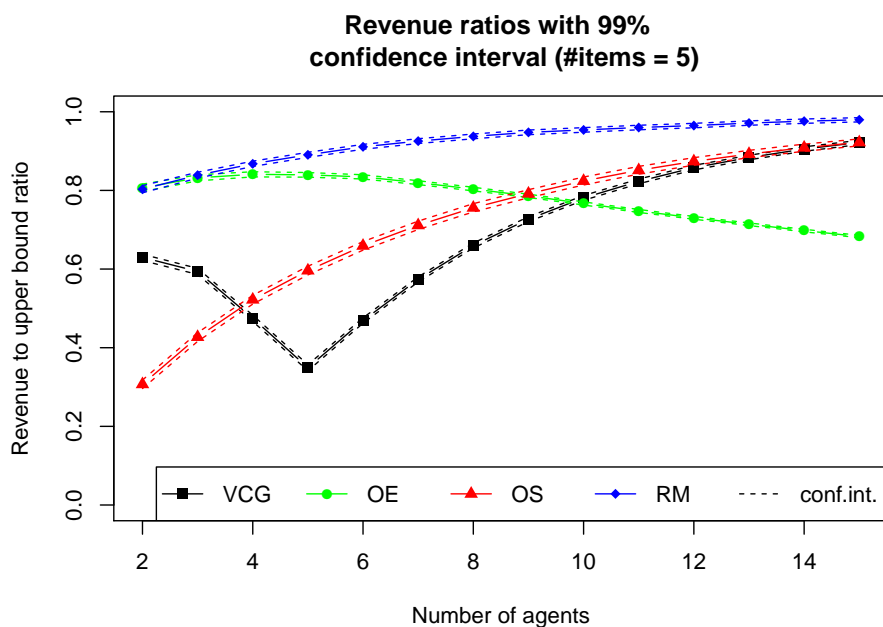


Figure 5: Revenue to upper bound ratios ($s^i \sim Exp(0.5)$, $m^i \sim Exp(0.5)$).

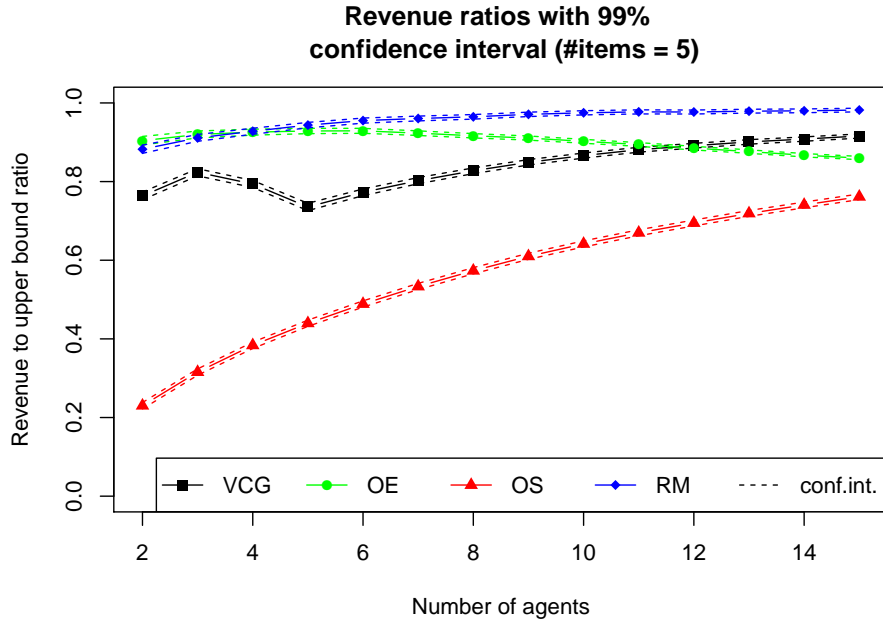


Figure 6: Revenue to upper bound ratios ($s^i \sim Exp(0.5)$, $m^i \sim Exp(0.25)$).

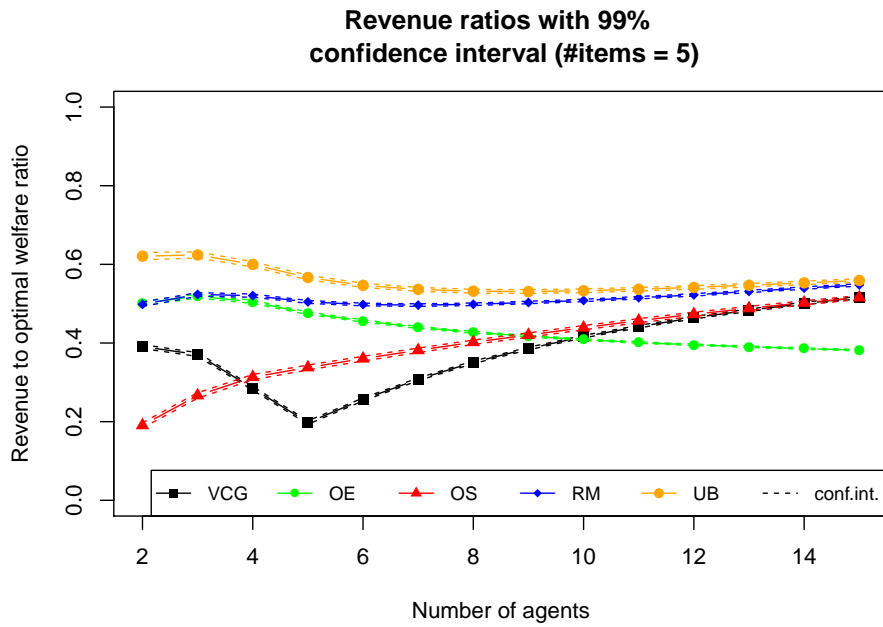


Figure 7: Revenue to optimal welfare ratios ($s^i \sim Exp(0.5)$, $m^i \sim Exp(0.5)$).

7 Conclusion

This paper has developed results for a new kind of one-to-many matching auction format which is relevant for many of today's platforms. Many platform applications already exist in a setting where one-to-many matches are possible (e.g., a single lead can be assigned to multiple interested sellers, or a single web surfer could be shown multiple simultaneous ads) but such matches are either done crudely or not at all because of the complexity in running such auctions. The complexity arises from the fact that some auction participants could potentially have very high value for exclusive purchase, and this causes the platform to have a method for choosing between exclusive vs. shared (i.e., multiple) allocation. The model and results developed in this paper can advance the practice of matching platforms through a proposed auction format that has high revenue performance, is easy for bidders to participate in and is predictable (because truth-telling is the dominant strategy for bidders).

The choice of deterministic, ex-post individual rational and dominant strategy implementable mechanism might seem to be limiting. Indeed, it is folklore knowledge that in multi-dimensional settings randomization improves revenue and that the class of Bayes-Nash implementable mechanisms allows more flexibility for the designer and admits different solving techniques (e.g., Cai et al., 2011). However, dominant truth-telling setting is highly relevant for two reasons: the deep need in practice for speed and simplicity, and (as we demonstrate) little loss in revenue from doing so. Most of the applications, especially lead marketing and sponsored search, involve real-time auctions where participants exhibit varied levels of sophistication. Therefore, both the auctioneer and bidders desire fast, predictable and simple mechanisms that require minimal user interaction and that make the available strategies clear for the users. Moreover, the dominant strategy implementation is prior-free, hence it is more robust from a practical perspective. In contrast, Bayes-Nash implementation requires a lot more of the participants: all of them should have the same prior about

the type distributions, and they have to be able to compute expected utilities based on their actions. The usage of lotteries might be acceptable in certain applications, but it is not yet desired in everyday business. Paying more than your bid is likely to alienate users from a system as they may regret their participation in the auction, therefore ex-post individual rationality can be seen as a natural property of an applied mechanism. Moreover, our main theorem demonstrates that focusing on dominant strategy implementable, ex-post individual rational, simple therefore robust mechanisms costs only a small fraction of the optimal revenue. This means that the trade-off between optimality and simplicity is small.

References

- Armstrong, Mark (1996). “Multiproduct Nonlinear Pricing”. In: *Econometrica* 64.1, pp. 51–75. ISSN: 00129682. URL: <http://www.jstor.org/stable/2171924>.
- Aseff, Jorge and Hector Chade (2008). “An Optimal Auction with Identity-Dependent Externalities”. English. In: *The RAND Journal of Economics* 39.3, pp. 731–746. ISSN: 07416261. URL: <http://www.jstor.org/stable/25474394>.
- Bagnoli, Mark and Ted Bergstrom (2005). “Log-Concave Probability and Its Applications”. English. In: *Economic Theory* 26.2, pp. 445–469. ISSN: 09382259. URL: <http://www.jstor.org/stable/25055959>.
- Barlow, Richard E, Albert W Marshall, and Frank Proschan (1963). “Properties of probability distributions with monotone hazard rate”. In: *The Annals of Mathematical Statistics* 34.2, pp. 375–389.
- Cai, Yang, Constantinos Daskalakis, and S. Matthew Weinberg (2011). “On Optimal Multi-Dimensional Mechanism Design”. In: *SIGecom Exch.* 10.2, pp. 29–33. ISSN: 1551-9031. DOI: 10.1145/1998549.1998555. URL: <http://doi.acm.org/10.1145/1998549.1998555>.

- Cai, Yang, C. Daskalakis, and S.M. Weinberg (2013). “Understanding Incentives: Mechanism Design Becomes Algorithm Design”. In: *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pp. 618–627. DOI: 10.1109/FOCS.2013.72.
- Choudary, Sangeet P., Marshall W. Van Alstyne, and Geoffrey G. Parker (2016). “Platform revolution: How Networked Markets Are Transforming the Economy—And How to Make Them Work for You”. In:
- Clarke, Edward H. (1971). “Multipart pricing of public goods”. English. In: *Public Choice* 11.1, pp. 17–33. ISSN: 0048-5829. DOI: 10.1007/BF01726210. URL: <http://dx.doi.org/10.1007/BF01726210>.
- Deng, Changrong and Saša Pekeč (2013). “Optimal Allocation of Exclusivity Contracts”. In:
- Devanur, Nikhil R., Jason D. Hartline, Anna R. Karlin, and C. Thach Nguyen (2011). “Prior-Independent Multi-parameter Mechanism Design”. In: *WINE*, pp. 122–133.
- Dhangwatnotai, Peerapong, Tim Roughgarden, and Qiqi Yan (2010). “Revenue Maximization with a Single Sample”. In: *Proceedings of the 11th ACM Conference on Electronic Commerce. EC '10*. Cambridge, Massachusetts, USA: ACM, pp. 129–138. ISBN: 978-1-60558-822-3. DOI: 10.1145/1807342.1807364. URL: <http://doi.acm.org/10.1145/1807342.1807364>.
- Evans, David S. and Richard Schmalensee (2016). *Matchmakers: The New Economics of Multisided Platforms*. Harvard Business Review Press.
- Figuroa, Nicolás and Vasiliki Skreta (2011). “Optimal allocation mechanisms with single-dimensional private information”. English. In: *Review of Economic Design* 15.3, pp. 213–243. ISSN: 1434-4742. DOI: 10.1007/s10058-010-0107-6. URL: <http://dx.doi.org/10.1007/s10058-010-0107-6>.
- Groves, Theodore (1973). “Incentives in Teams”. English. In: *Econometrica* 41.4, pp. 617–631. ISSN: 00129682. URL: <http://www.jstor.org/stable/1914085>.

- Hartline, Jason D (2012). “Approximation in mechanism design”. In: *The American Economic Review* 102.3, pp. 330–336.
- Jehiel, Philippe, Benny Moldovanu, and Ennio Stacchetti (1996). “How (Not) to Sell Nuclear Weapons”. English. In: *The American Economic Review* 86.4, pp. 814–829. ISSN: 00028282. URL: <http://www.jstor.org/stable/2118306>.
- Jerath, Kinshuk and Amin Sayedi (2012). “Exclusive Display in Sponsored Search Advertising”. working paper.
- Likhodedov, Anton and Tuomas Sandholm (2004). “Methods for Boosting Revenue in Combinatorial Auctions”. In: *Proceedings of the 19th National Conference on Artificial Intelligence*. AAAI’04. San Jose, California: AAAI Press, pp. 232–237. ISBN: 0-262-51183-5. URL: <http://dl.acm.org/citation.cfm?id=1597148.1597187>.
- Manelli, Alejandro M. and Daniel R. Vincent (2007). “Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly”. In: *Journal of Economic Theory* 137.1, pp. 153–185. ISSN: 0022-0531. DOI: <http://dx.doi.org/10.1016/j.jet.2006.12.007>. URL: <http://www.sciencedirect.com/science/article/pii/S0022053107000348>.
- Milgrom, P R (2004). *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press.
- Myerson, Roger (1981). “Optimal auction design”. In: *Mathematics of Operations Research* 6.1, pp. 58–73.
- Pei, Jinxiang, Diego Klabjan, and Wei Xie (2014). “Approximations to auctions of digital goods with share-averse bidders”. In: *Electronic Commerce Research and Applications* 13.2, pp. 128–138. ISSN: 1567-4223.
- Rochet, Jean-Charles and Philippe Choné (1998). “Ironing, Sweeping, and Multidimensional Screening”. English. In: *Econometrica* 66.4, pp. 783–826. ISSN: 00129682. URL: <http://www.jstor.org/stable/2999574>.

- Roughgarden, Tim and Mukund Sundararajan (2007). “Is efficiency expensive”. In: *In Third Workshop on Sponsored Search Auctions*.
- Roughgarden, Tim and Inbal Talgam-Cohen (2013). “Optimal and Near-optimal Mechanism Design with Interdependent Values”. In: *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. Philadelphia, Pennsylvania, USA: ACM, pp. 767–784. ISBN: 978-1-4503-1962-1. DOI: 10.1145/2482540.2482606. URL: <http://doi.acm.org/10.1145/2482540.2482606>.
- Salek, Mahyar and David Kempe (2008). “Auctions for Share-Averse Bidders.” In: *WINE*. Ed. by Christos H. Papadimitriou and Shuzhong Zhang. Vol. 5385. Lecture Notes in Computer Science. Springer, pp. 609–620. ISBN: 978-3-540-92184-4.
- Sayedi, Amin (2012). “Essays on Sponsored Search Advertising”. AAI3524658. PhD thesis. Pittsburgh, PA, USA. ISBN: 978-1-267-57425-1.
- Segal, Ilya (1999). “Contracting with Externalities”. English. In: *The Quarterly Journal of Economics* 114.2, pp. 337–388. ISSN: 00335533. URL: <http://www.jstor.org/stable/2587012>.
- Tang, Pingzhong and Tuomas Sandholm (2012). “Mixed-bundling Auctions with Reserve Prices”. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. AAMAS '12. Valencia, Spain: International Foundation for Autonomous Agents and Multiagent Systems, pp. 729–736. ISBN: 0-9817381-2-5, 978-0-9817381-2-3. URL: <http://dl.acm.org/citation.cfm?id=2343776.2343801>.
- Varian, Hal R. (2009). “Online Ad Auctions”. In: *The American Economic Review* 99.2, pp. 430–434.
- Vickrey, William (1961). “Counterspeculation, auctions, and competitive sealed tenders”. In: *The Journal of Finance* 16.1, pp. 8–37. ISSN: 1540-6261. DOI: 10.1111/j.1540-6261.1961.tb02789.x. URL: <http://dx.doi.org/10.1111/j.1540-6261.1961.tb02789.x>.

A Technical Details and Proofs

Notation	Meaning
k	Maximum number of sold items
$N = \{1, \dots, n\}$	Set of agents
A	Set of possible allocations
S^i, M^i	Set of shared valuations and exclusivity margins of agent i respectively
$T^i = S^i \times M^i$	Set of types of agent i
$T = T^1 \times \dots \times T^n$	Set of type profiles
$t^i = (s^i, m^i)$	Type of agent i
$t^{-i} = \times_{j \neq i} t^j$	Type profile excluding the type of agent i
F, G	CDF of shared valuations and exclusivity margins respectively
C	CDF of the exclusive value, i.e., the convolution of F and G
$\phi_D(t^i)$	Virtual valuation of agent i for random variable t^i distributed as D
$x : T \rightarrow \{0, 1\}^N$	Allocation rule
$p : T \rightarrow \mathbb{R}^N$	Payment scheme
$v^i(x(\hat{t}), t^i)$	Valuation of agent i for allocation $x(\hat{t})$ having actual type t^i
$u^i(x(\hat{t}), p(\hat{t}), t^i)$	Utility of agent i for allocation $x(\hat{t})$, payment $p(\hat{t})$ having actual type t^i
$SW_N(a t)$	Total welfare of agents in N for allocation a and type profile t
$a(t) = \{i : x^i(t) = 1\}$	Set of agents who receive the item for report t

Table 3: Summary of Notation

Proof of Theorem 2. (UBM) can be decomposed by treating all possible exclusivity margin profiles separately, because none of the constraints involve variables that are related to

different exclusivity margins, furthermore the probabilities are independent among agents and among the dimensions of their type. Therefore it is sufficient to solve the subproblems separately for each fixed m and acquire the expected revenue as the expectation over the optimal objective values of the subproblems. From Proposition 6 in Deng and Pekeč, 2013 it follows that for fixed m the allocation rule x of the optimal DSIC and EPIR mechanism given type profile $t = (s, m)$, is

$$x^i(t) = 1 \text{ for } \begin{cases} i \in a_s(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) \geq \phi_F(s^{a_e(t)}) + m^{a_e(t)} \\ i = a_e(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) < \phi_F(s^{a_e(t)}) + m^{a_e(t)} \end{cases} \quad (13)$$

and $x^i(t) = 0$ otherwise. The expected revenue under this allocation rule is equal to

$$\mathbb{E}_s \left[\sum_i p^i(s, m) \right] = \mathbb{E}_s \left[\sum_i \left(\phi_F(s^i) + m^i \prod_{j \neq i} (1 - x^j(s, m)) \right) x^i(s, m) \right]. \quad (14)$$

Since 13 holds for any fixed m , it concludes 10. Finally, 11 follows from taking expectation over 14 with respect to the exclusivity margins.

Note that Deng and Pekeč, 2013 originally studies ex-post incentive compatible (EPIC) mechanisms, because their exclusivity margin is assumed to be a linear combination of the shared valuations of the other agents. In our case exclusive margins do not depend on the valuation of the other agents, hence the notions DSIC and EPIC coincide (for formal definition of EPIC see Definition 10). \square

Proof of Proposition 1. Before proceeding to the main proof it is useful to get some insights on the allocation rule and the pricing scheme of any DSIC mechanism for DSA*. We start with a lemma for the payments saying that for the same allocation one has to be charged the same amount regardless of his bid, if the bids of the others kept unchanged. Together with ex-post individual rationality this implies an intuitive, but non-trivial upper

bound for each payment.

Lemma 3. *Let (x, p) be a feasible mechanism for (DSA*). For given i , t^{-i} , t^i and \hat{t}^i , having $x(t^i, t^{-i}) = x(\hat{t}^i, t^{-i})$ implies that $p^i(t^i, t^{-i}) = p^i(\hat{t}^i, t^{-i})$. Moreover, for given i , t^{-i} , t^i we have that*

$$p^i(t^i, t^{-i}) \leq \inf_{\hat{t}^i} \{v^i(x(\hat{t}^i, t^{-i}), \hat{t}^i) \mid x(\hat{t}^i, t^{-i}) = x(t^i, t^{-i})\}.$$

Proof. Follows directly from (DSIC) and ex-post individual rationality. \square

Regarding the allocation rule the following lemma states that the higher your valuation is, the more chance you have to receive the item. The statements are grouped based on which dimension of the type is changed. Fix agent i , the type of the others t^{-i} and the allocation rule x . Let $R_e(x, t^{-i}) \subseteq T^i$, $R_s(x, t^{-i}) \subseteq T^i$ be the set of types such that the allocation rule assigns the item to agent i exclusively and shared respectively. Similarly, $R_0(x, t^{-i}) \subseteq T^i$ is the set of types such that agent i gets nothing.

Lemma 4. *Let x be a monotone allocation rule and fix agent i and the report of the others t^{-i} . Then the following two statements hold:*

I/ *For all s^i, m^i, \hat{m}^i such that $\hat{m}^i > m^i$, we have that if $(s^i, m^i) \in R_e(x, t^{-i})$, then $(s^i, \hat{m}^i) \in R_e(x, t^{-i})$.*

II/ *There is an $m^* \in M^i \cup \{\infty\}$ such that for all s^i, m^i, \hat{s}^i , where $\hat{s}^i > s^i$ and $m^i \neq m^*$, we have that if $(s^i, m^i) \in R_s(x, t^{-i})$, then $(\hat{s}^i, m^i) \in R_s(x, t^{-i})$.*

Proof. Assume that the first statement is not true, that is, $(s^i, \hat{m}^i) \in R_0(x, t^{-i})$ or $(s^i, \hat{m}^i) \in R_s(x, t^{-i})$. Applying (MON) for both cases with $t^i = (s^i, m^i)$ and $\hat{t}^i = (s^i, \hat{m}^i)$ results in

$$0 \geq \hat{m}^i - m^i.$$

This is a contradiction as $\hat{m}^i > m^i$ by assumption.

For the second statement define m^* as the infimum of exclusivity margins such that agent i is allocated exclusively, i.e.,

$$m^* = \inf\{\tilde{m}^i \mid (\tilde{s}^i, \tilde{m}^i) \in R_e(x, t^{-i})\}. \quad (15)$$

Now, assume the contrary of the second statement, that is, there are s^i, m^i, \hat{s}^i such that $\hat{s}^i > s^i$, $m^i \neq m^*$, $(s^i, m^i) \in R_s(x, t^{-i})$ and we have that $(\hat{s}^i, m^i) \in R_0(x, t^{-i})$ or $(\hat{s}^i, m^i) \in R_e(x, t^{-i})$. First tackle the case when $(\hat{s}^i, m^i) \in R_0(x, t^{-i})$. Applying (MON) with $t^i = (s^i, m^i)$ and $\hat{t}^i = (\hat{s}^i, m^i)$ results in

$$0 \geq \hat{s}^i - s^i.$$

This is a contradiction as $\hat{s}^i > s^i$ by assumption. Finally, assume that $(\hat{s}^i, m^i) \in R_e(x, t^{-i})$. Note that due to the definition of m^* this assumption implies that $m^i > m^*$ and that there exists a (\bar{s}^i, \bar{m}^i) such that $\bar{m}^i < m^i$ and $(\bar{s}^i, \bar{m}^i) \in R_e(x, t^{-i})$. Applying (MON) with $t^i = (s^i, m^i)$ and $\hat{t}^i = (\bar{s}^i, \bar{m}^i)$ results in

$$\bar{s}^i + \bar{m}^i - s^i - m^i \geq \bar{s}^i - s^i.$$

This is a contradiction as $m^i > \bar{m}^i$ by assumption. □

Now, we are ready to prove directly Proposition 1. By construction we have that for every $i \in N$ there is a pair $(i_s, i_e) \in \bar{N}$, and for every $t \in T$ there is a $\bar{t} \in \bar{T}$ such that $t^i = (s^i, m^i) = (\bar{t}^{i_s}, \bar{t}^{i_e})$. The steps of the proof are the following. First, for every feasible mechanism (x, p) for (DSA*), we define an allocation rule \bar{x} for (UBR) that satisfies (FeasA) and (FeasB). Then we construct a \bar{p} such that (\bar{x}, \bar{p}) satisfies (EPIC) and (EPIR2). Finally, we show that $\mathbb{E}_{\bar{t} \in \bar{T}} [\sum_{i \in \bar{N}} \bar{p}^i(\bar{t})] \geq \mathbb{E}_{t \in T} [\sum_{i \in N} p^i(t)]$.

Fix agent i and the type of the others t^{-i} . Let $t = (t^i, t^{-i})$, where $t^i = (s^i, m^i)$. Define m^*

as in (15). Now, set $\bar{x}^{i_s}(\bar{t}) = 1$ if and only if $m^i \neq m^*$ and $t^i \in R_e(x, t^{-i})$, and set $\bar{x}^{i_e}(\bar{t}) = 1$ if and only if $t^i \in R_e(x, t^{-i})$. In words, whenever x allocates shared to agent i , then \bar{x} gives the item to representative i_s provided that $m^i \neq m^*$. Furthermore, if x allocates exclusively to agent i , then \bar{x} gives the item to representative i_e . It is easy to see that \bar{x} satisfies (FeasA) and (FeasB).

Next, let us define the payment rule \bar{p} for agent $i_s \in \bar{N}_s$ as

$$\bar{p}^{i_s}(\bar{t}^{i_s}, \bar{t}^{-i_s}) = \begin{cases} \inf\{\hat{t}^{i_s} \mid \bar{x}^{i_s}(\hat{t}^{i_s}, \bar{t}^{-i_s}) = 1\} & \text{if } \bar{x}^{i_s}(\bar{t}^{i_s}, \bar{t}^{-i_s}) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

and for agent $i_e \in \bar{N}_e$ as

$$\bar{p}^{i_e}(\bar{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = \begin{cases} \inf\{\hat{t}^{i_e} \mid \bar{x}^{i_e}(\hat{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = 1\} + \bar{t}^{i_s} & \text{if } \bar{x}^{i_e}(\bar{t}^{i_e}, \bar{t}^{-i_e}) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

To see incentive compatibility, consider agent i_e and fix $(\bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)})$. If there is no \bar{t}^{i_e} such that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 1$, then (EPIC) and (EPIR2) trivially hold. Otherwise let $\underline{m} = \inf\{\hat{t}^{i_e} \mid \bar{x}^{i_e}(\hat{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = 1\}$. Note that Lemma 4 ensures that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)})$ is non-decreasing in \bar{t}^{i_e} for all \bar{t}^{-i_e} . This implies that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 1$ for all $\bar{t}^{i_e} > \underline{m}$, and $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 0$ for all $\bar{t}^{i_e} < \underline{m}$. It is easy to see that this allocation rule together with payment $\bar{p}^{i_e}(\bar{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = \underline{m} + \bar{t}^{i_s}$ satisfy (EPIC) and (EPIR2). With respect to the comparison of the payments we can observe the following. When $x(t)$ allocates exclusively to agent i , then

$$\begin{aligned} p^i(s^i, m^i, t^{-i}) &\leq \inf\{v^i(x(\hat{s}^i, \hat{m}^i, t^{-i}), (\hat{s}^i, \hat{m}^i)) \mid x(\hat{s}^i, \hat{m}^i, t^{-i}) = x(s^i, m^i, t^{-i})\} \\ &\leq \inf\{\bar{t}^{i_s} + \hat{t}^{i_e} \mid \bar{x}^{i_e}(\hat{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 1, \bar{t}^{i_s} = s^i\} \\ &= \bar{p}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}). \end{aligned}$$

The inequalities follow from Lemma 3 and from the fact that the second infimum is taken over a smaller set. This means that in case of exclusive allocation $\bar{p}^{ie}(\bar{t}) + \bar{p}^{is}(\bar{t}) \geq p^i(t)$.

Finally, consider agent i_s and fix $(\bar{t}^{ie}, \bar{t}^{-(is,ie)})$. If $\bar{t}^{ie} = m^*$ or there is no \bar{t}^{is} such that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 1$, then $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 0$ for all \bar{t}^{is} , therefore (EPIC) and (EPIR2) trivially hold. Otherwise assume that $\bar{t}^{ie} \neq m^*$ and let $\underline{s} = \inf\{\hat{t}^{is} \mid \bar{x}^{is}(\hat{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 1\}$. Note that Lemma 4 ensures that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)})$ is non-decreasing in \bar{t}^{is} for all $(\bar{t}^{ie}, \bar{t}^{-(is,ie)})$. As a result we have that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 1$ for all $\bar{t}^{is} > \underline{s}$ and $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 0$ for all $\bar{t}^{is} < \underline{s}$. It is easy to see that this allocation rule together with payment $\bar{p}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = \underline{s}$ satisfy (EPIC) and (EPIR2). With respect to the comparison of the payments we can observe the following. When $x(t)$ allocates shared to agent i and $\bar{t}^{ie} \neq m^*$, then

$$\begin{aligned} p^i(s^i, m^i, t^{-i}) &\leq \inf\{v^i(x(\hat{s}^i, \hat{m}^i, t^{-i}), (\hat{s}^i, \hat{m}^i)) \mid x(\hat{s}^i, \hat{m}^i, t^{-i}) = x(s^i, m^i, t^{-i})\} \\ &\leq \inf\{\hat{t}^{is} \mid \bar{x}^{is}(\hat{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}) = 1, \bar{t}^{ie} = m^i\} \\ &= \bar{p}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is,ie)}). \end{aligned}$$

The inequalities follow from Lemma 3 and from the fact that the second infimum is taken over a smaller set. This means that in case of shared allocation $\bar{p}^{ie}(\bar{t}) + \bar{p}^{is}(\bar{t}) \geq p^i(t)$.

Combining the results of all subcases implies that for any feasible solution (x, p) for

(DSA*) there is a feasible (\bar{x}, \bar{p}) for (UBR) such that

$$\begin{aligned}
\mathbb{E}_{t \in T} \left[\sum_{i \in N} p^i(t) \right] &= \sum_{i \in N} \mathbb{E}_{t^{-i} \in T^{-i}} \left[\mathbb{E}_{(s^i, m^i) \in T^i} [p^i(s^i, m^i, t^{-i})] \right] \\
&= \sum_{i \in N} \mathbb{E}_{t^{-i} \in T^{-i}} \left[\mathbb{E}_{(s^i, m^i) \in T^i, m^i \neq m^*} [p^i(s^i, m^i, t^{-i})] \right] \\
&\leq \sum_{i \in N} \mathbb{E}_{\bar{t}^{-i} \in \bar{T}^{-i}} \left[\mathbb{E}_{t^{is} \in \bar{T}^{is}, t^{ie} \in \bar{T}^{ie} \setminus m^*} \left[\bar{p}^{is}(t^{is}, t^{ie}, \bar{t}^{-i, ie}) + \bar{p}^{ie}(t^{is}, t^{ie}, \bar{t}^{-i, ie}) \right] \right] \\
&= \sum_{i \in N} \mathbb{E}_{\bar{t}^{-i} \in \bar{T}^{-i}} \left[\mathbb{E}_{t^{is} \in \bar{T}^{is}, t^{ie} \in \bar{T}^{ie}} \left[\bar{p}^{is}(t^{is}, t^{ie}, \bar{t}^{-i, ie}) + \bar{p}^{ie}(t^{is}, t^{ie}, \bar{t}^{-i, ie}) \right] \right] \\
&= \mathbb{E}_{\bar{t} \in \bar{T}} \left[\sum_{i \in \bar{N}} \bar{p}^i(\bar{t}) \right].
\end{aligned}$$

The second and third equalities hold because type distributions are continuous, hence the measure of the type set where $t^{ie} = m^*$ is zero. The inequality follows from the results of the previous subcases. This concludes the proof of Proposition 1. \square

Proof of Theorem 3. Roughgarden and Talgam-Cohen, 2013 extends the results of Myerson, 1981 to settings with informational externalities and correlated type distributions for mechanism that are ex-post incentive compatible and ex-post individual rational. In particular, it is shown that the expected revenue of any EPIC and EPIR mechanism equals to expected sum of virtual valuations provided that the payments are maximal, i.e., the utility of an agent with zero type is zero. Moreover, the revenue maximizing mechanism allocates such that the sum of virtual valuations is maximal for each type profile, given that the resulting allocation is monotone non-decreasing for each agent in their own type. Note that the virtual valuation of agent i_e having type \bar{t}^{ie} is $\bar{t}^{is} + \phi_G(\bar{t}^{ie})$ due to the informational externality. Having all these in mind the only thing left to show is that all agent's probability of receiving the item is monotone non-decreasing in their reported types.

As F has monotone hazard rate, we have that $\phi_F(\bar{t}^{is})$ is monotone non-decreasing in \bar{t}^{is}

for all $i_s \in \bar{N}_s$. Therefore if $i_s \in \bar{a}_s(\bar{t}^{i_s}, \bar{t}^{-i_s})$, then $i_s \in \bar{a}_s(\hat{t}^{i_s}, \bar{t}^{-i_s})$ for any $\hat{t}^{i_s} > \bar{t}^{i_s}$. Similarly, as G has monotone hazard rate, we have that $\phi_G(\bar{t}^{i_e})$ is also non-decreasing in \bar{t}^{i_e} for all $i_e \in \bar{N}_e$. This ensures that once $i_e \in \bar{a}_e(\bar{t}^{i_e}, \bar{t}^{-i_e})$, then $i_e \in \bar{a}_e(\hat{t}^{i_e}, \bar{t}^{-i_e})$ also for any $\hat{t}^{i_e} > \bar{t}^{i_e}$. The last critical point is that when any agent $i_s \in \bar{N}_s$ increases \bar{t}^{i_s} , because then both $\phi_F(\bar{t}^{i_s})$ and $\bar{t}^{i_s} + \phi_G(\bar{t}^{i_e})$ increase at the same time. To see why it is not a problem observe that the MHR assumption ensures that $\phi_F(\bar{t}^{i_s})$ increases at least as fast as $\bar{t}^{i_s} + \phi_G(\bar{t}^{i_e})$ in \bar{t}^{i_s} . \square

Proof of Theorem 4.

In order to prove this Theorem we need to elaborate first in more detail on the VCG auction and prove a couple of results. Let $SW_N(a | t) = \sum_{i \in N} v^i(a, t^i)$ denote the welfare of agents in N for allocation a and type profile t .

Definition 11 (VCG mechanism). Let A_N be the set of feasible allocations for agents in N and let $a(t) \in A_N$ be an allocation for each type profile t such that

$$a(t) \in \arg \max_{a \in A_N} SW_N(a | t).$$

Similarly, let $A_{N \setminus \{i\}}$ be the set of feasible allocations for agents in $N \setminus \{i\}$ and let $a(t^{-i}) \in A_{N \setminus \{i\}}$ be an allocation for each type profile t^{-i} such that

$$a(t^{-i}) \in \arg \max_{a \in A_{N \setminus \{i\}}} SW_{N \setminus \{i\}}(a | t^{-i}).$$

Then for each t VCG chooses $a(t)$ as allocation and elicits payment

$$p^i(t) = \sum_{j \neq i} v^j(a(t^{-i}), t^j) - \sum_{j \neq i} v^j(a(t), t^j). \quad (18)$$

VCG allocates efficiently, i.e., maximizes the total welfare pointwise, moreover, it is EPIR. The following Lemma helps computing its expected revenue.

Lemma 5. *Let $N = \{1, \dots, n\}$ be the set of agents with i.i.d. types, and let $Q = \{1, \dots, n - 1\}$ denote the same set of agents with one less member. Let $T_N = \times_{i \in N} T^i$ denote the set of possible type profiles of agents in N and, similarly let $T_Q = \times_{i \in Q} T^i$ be the set of possible*

type profiles of $n - 1$ agents. Then

$$Rev(\text{VCG}) = n\mathbb{E}_{t \in T_Q} [SW_Q(a(t) | t)] - (n - 1)\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)] \quad (19)$$

Proof. According to (18) the expected value of the VCG payments can be written as

$$\begin{aligned} Rev(\text{VCG}) &= \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \left(\sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) - \sum_{j \in N \setminus \{i\}} v^j(a(t), t^j) \right) \right] \\ &= \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) \right] - \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \sum_{j \in N \setminus \{i\}} v^j(a(t), t^j) \right] \\ &= \sum_{i \in N} \mathbb{E}_{t^i \in T^i} \left[\mathbb{E}_{t^{-i} \in T_N^{-i}} \left[\sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) \right] \right] - (n - 1)\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)] \\ &= \sum_{i \in N} \mathbb{E}_{t^i \in T^i} \left[\mathbb{E}_{t \in T_Q} \left[\sum_{j \in Q} v^j(a(t^{-i}), t^j) \right] \right] - (n - 1)\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)] \\ &= n\mathbb{E}_{t \in T_Q} \left[\sum_{j \in Q} v^j(a(t), t^j) \right] - (n - 1)\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)] \\ &= n\mathbb{E}_{t \in T_Q} [SW_Q(a(t) | t)] - (n - 1)\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)]. \end{aligned}$$

The equalities are direct consequences of the assumption that types are i.i.d.. □

Corollary 1. *If*

$$\frac{\mathbb{E}_{t \in T_Q} [SW_Q(a(t) | t)]}{\mathbb{E}_{t \in T_N} [SW_N(a(t) | t)]} \geq 1 - \varrho,$$

then the expected revenue of VCG is at least $1 - n\varrho$ times the expected optimal welfare.

The message of Corollary 1 is that if one wants to compare $Rev(\text{VCG})$ to the optimal expected welfare, then it is sufficient to know the added value of an extra agent to the welfare. Let $H_i = \sum_{j=1}^i 1/j$ represent the i^{th} Harmonic number and set $H_0 = 0$. The next lemma is useful for providing lower bounds on the revenue-welfare ratio.

Lemma 6 (Lemma 3 from Roughgarden and Sundararajan, 2007). *Draw independently n times from an MHR distribution. Then the expected value of the l^{th} -largest value of n samples is at least $(H_n - H_{l-1})/(H_{n+j} - H_{l-1})$ times that of the l^{th} -largest value of $n + j$ samples.*

Theorem 5. *Consider the single-item auction problem with $n \geq 2$ agents who have unit demand and single-dimensional valuations i.i.d. according to an MHR distribution. Then VCG extracts at least $1 - 1/H_n$ fraction of the optimal welfare in terms of expected revenue.*

Proof. Let $N = \{1, \dots, n\}$ be the set of agents with i.i.d. types, and let $Q = \{1, \dots, n-1\}$ denote the same set of agents with one less member. Furthermore, denote the l^{th} largest value from n samples by $v_{[l:n]}$. Lemma 6 implies that

$$\mathbb{E} [v_{[l:n-1]}] \geq (H_{n-1} - H_{l-1})/(H_n - H_{l-1})\mathbb{E} [v_{[l:n]}] .$$

Therefore we have that

$$\begin{aligned} \frac{\mathbb{E}_{t \in T_Q} [SW_Q(x(t) | t)]}{\mathbb{E}_{t \in T_N} [SW_N(x(t) | t)]} &= \frac{\mathbb{E} [v_{[1:n-1]}]}{\mathbb{E} [v_{[1:n]}]} \\ &\geq \frac{(H_{n-1}/H_n)\mathbb{E} [v_{[l:n]}]}{\mathbb{E} [v_{[l:n]}]} \\ &= 1 - \frac{1}{nH_n} . \end{aligned}$$

The proof is concluded by invoking Corollary 1 and letting $\varrho = \frac{1}{nH_n}$. □

We note that according to Theorem 4 of Roughgarden and Sundararajan, 2007 the ratio of the VCG revenue to the optimal welfare is at least $1 - 1/n$ for monotone hazard rate distributions. Their result is apparently not precise as our bound is tight for exponential distributions and $1 - 1/H_n$ is generally lower than $1 - 1/n$.

Now we are ready to proof Theorem 4

According to Theorem 2 $Rev(\text{UBM}^*)$ is an upper bound on $Rev(\text{DSA}^*)$, therefore it is sufficient to show that $Rev(\text{MAXSIMPLE})$ approximates $Rev(\text{UBM}^*)$. Let x_{UBM^*} represent the allocation rule of UBM^* . Define $T_{\text{UBM}^*}^e = \{t \in T \mid \exists i : x_{\text{UBM}^*}^i(t) = 1, \forall j \neq i : x_{\text{UBM}^*}^j(t) = 0\}$ and $T_{\text{UBM}^*}^s = \{t \in T \mid \exists i, j \neq i : x_{\text{UBM}^*}^i(t) = x_{\text{UBM}^*}^j(t) = 1\}$. Then according to Theorem 2 the revenue of UBM^* can be split such that $Rev(\text{UBM}^*) = Rev(\text{UBM}^*)_s + Rev(\text{UBM}^*)_e$, where

$$Rev(\text{UBM}^*)_s = \mathbb{E}_{t \in T_{\text{UBM}^*}^s} \left[\sum_i \phi_F(s^i) x_{\text{UBM}^*}^i \right] \quad (20)$$

and

$$Rev(\text{UBM}^*)_e = \mathbb{E}_{t \in T_{\text{UBM}^*}^e} \left[\sum_i (\phi_F(s^i) + m^i) x_{\text{UBM}^*}^i \right]. \quad (21)$$

The idea of the proof is to bound the two terms separately. We start with $Rev(\text{UBM}^*)_e$.

As $\phi_F(s^i) + m^i \leq s^i + m^i$ for all i , we have that $Rev(\text{UBM}^*)_e$ is less than or equal to the optimal welfare of a single-item auction. Note that OE achieves at least as much expected revenue as VCG does for the single-item auction. Therefore due to Theorem 5 $Rev(\text{OE})$ is at least $1 - 1/H_n$ times the optimal welfare of a single-item auction. This leads to the conclusion that $Rev(\text{OE}) \geq (1 - 1/H_n) Rev(\text{UBM}^*)_e$.

To bound $Rev(\text{UBM}^*)_s$ observe that whenever UBM^* allocates shared for type report t , then $x_{\text{UBM}^*}(t) = x_{OS}(t)$. This is due to the fact $a_s(t)$ is defined the same way for both

mechanism. Using this observation together with Lemma 2 we can write

$$\begin{aligned}
Rev(\text{OS}) &= \mathbb{E}_t \left[\sum_i \phi_F(s^i) x_{\text{OS}}^i(t) \right] \\
&= \mathbb{E}_{t \in T_{\text{UBM}}^s} \left[\sum_i \phi_F(s^i) x_{\text{OS}}^i(t) \right] + \mathbb{E}_{t \notin T_{\text{UBM}}^s} \left[\sum_i \phi_F(s^i) x_{\text{OS}}^i(t) \right] \\
&= Rev(\text{UBM}^*)_s + \mathbb{E}_{t \notin T_{\text{UBM}}^s} \left[\sum_i \phi_F(s^i) x_{\text{OS}}^i(t) \right] \\
&\geq Rev(\text{UBM}^*)_s.
\end{aligned}$$

The last inequality holds because under OS allocation only occur when the sum of virtual valuations is non-negative. Putting together the two bounds results in

$$\frac{1}{1 - 1/H_n} Rev(\text{OE}) + Rev(\text{OS}) \geq Rev(\text{UBM}^*)_s + Rev(\text{UBM}^*)_e = Rev(\text{UBM}^*) \geq Rev(\text{DSA}^*).$$

To finish the proof note that $Rev(\text{MAXSIMPLE}) = \max\{Rev(\text{OS}), Rev(\text{OE})\}$, hence

$$\left(\frac{1}{1 - 1/H_n} + 1 \right) Rev(\text{MAXSIMPLE}) \geq Rev(\text{DSA}^*).$$

Finally, note that we can replace DSIC to Bayes-Nash incentive compatibility in (DSA*) and relax the resulting mathematical program by letting the exclusivity margin public information as in (UBM). As we arrive at a single-dimensional setting it is folklore knowledge that the optimal Bayes-Nash mechanism is DSIC. This means that UBM* is optimal even among Bayes-Nash implementable mechanisms, therefore it is also an upper bound for the Bayesian relaxation of (DSA*). \square