

Equilibrium Prices in Bundle Auctions

Peter R. Wurman and Michael P. Wellman

University of Michigan
Artificial Intelligence Laboratory
{ pwurman, wellman }@umich.edu

Abstract

The allocation of discrete, complementary resources is a fundamental problem in economics and of direct interest to e-commerce applications. In this paper we establish that competitive equilibrium bundle prices always exist that support the efficient allocation in discrete resource allocation problems with free disposal. We believe that this is an important step in the quest for a mechanism that performs well in the face of complementary preferences. We present a family of auctions that use this bundle pricing policy, and make some initial observations on several of its members, including the new Ascending k -Bundle auction.

Introduction

In many potential e-commerce applications, agents have complementary preferences for objects in the marketplace. Consider an agent trying to construct a computer system by purchasing components. Among other things, the agent needs to buy a cpu, harddrive, and monitor, and may have a choice of several models for each component. The agent's valuation of a package depends on the combination of components that it can accumulate. This example is an instance of the general allocation problem characterized by heterogeneous, discrete resources and complementarities in agent preferences.

Discrete resource allocation has been the focus of auction theory since Vickrey's seminal work on the second-price auction (Vickrey 1961). The allocation of single objects, or multiple homogeneous objects, is well understood (Demange, Gale, & Sotomayor 1986; McAfee & McMillan 1987; Milgrom 1989), however a general solution to the allocation of heterogeneous objects has proven elusive. The landscape seems composed of various mechanisms that make different trade-offs among desirable properties.

One branch of research is focused on identifying conditions under which price equilibria exist. For instance, it has been shown that price equilibria exist if a *gross substitutability* condition holds (Kelso & Crawford 1982), or if utility functions satisfy the *no-complementarities* condition (Gul & Stachetti 1997). Bikhchandani and Mamer (Bikhchandani & Mamer

	A	B	AB
Agent 1	0	0	3
Agent 2	2	2	2

Table 1: An example without equilibrium prices for individual goods.

1997) demonstrated that equilibria exist iff the social value of the optimal allocation is equal to the value of the corresponding relaxed linear program.

However, it is easy to construct examples in which price equilibria fail to exist. A simple example presented by McAfee and McMillan (McAfee & McMillan 1996), reproduced in Table 1, illustrates the problem. One agent values the pair of goods at 3 and gets no benefit from the goods singly. The second agent values either good at 2, but gets no added benefit from getting both goods. The *efficient* solution, from a social perspective, assigns both goods to agent 1.¹ For any assignment that includes agent 2 receiving one or more goods, both agents can be made better off by a transfer of money from agent 1 to agent 2 in exchange for agent 2's allocation. In this example, however, no equilibrium prices exist. In order to exclude the first agent from the allocation, the price of each good individually must be greater than or equal to 2. However, these prices put the cost of the pair at 4—above the second agent's valuation.

Another subject of research in this area is the generalization of Vickrey's original results and related results by Clarke (Clarke 1971) and Groves (Groves 1973). This mechanism, which alternately goes by the moniker Vickrey-Clarke-Groves mechanism or the Generalized Vickrey Auction (GVA), is an incentive compatible, individually rational, efficient, direct revelation mechanism. The mechanism's attractive properties come from the fact that an agent pays (receives) the difference in utility that other agents lose (gain) by its presence. Thus, the agent's payment is not dependent upon its own stated preferences, which removes the incentive to misrepresent those preferences.

¹Throughout the examples in this paper we indicate the socially efficient allocation in boldface.

However, the GVA does have several undesirable properties. Foremost among these is that, in the two-sided market, it is not budget-balanced. Even when we restrict our attention to the single-sided case, the mechanism has very high computational costs for the auctioneer, and requires complete information revelation by the agents. We suspect that in many cases it is not necessary for agents to fully disclose their utility function in order to reach the efficient allocation.

Under some conditions, the payments computed by the GVA can be interpreted as prices (Leonard 1983). However, in the general case, we run into two problems. First, the GVA computes payments per agent, not prices per object. When a single agent is allocated a combination of objects, there is no way in which to determine the relative price of each object in the bundle. Moreover, the GVA may not compute *anonymous* prices, where anonymity implies that every agent has the opportunity to purchase the same object at the same price.² Consider an example with two buyers and two identical units of a good. Agent 1 values one unit of the good at 5, and two units at 7. Agent 2 values one and two units at 5 and 8, respectively. The efficient allocation is to give one unit to each agent. However, the GVA payment for agent 1 is $8 - 5 = 3$, whereas the payment by agent 2, for a unit of the same good, is 2. Thus, if participants care that prices be uniform, the GVA will not be an acceptable mechanism.

There are a few notable alternatives to the GVA that attempt to allocate heterogeneous goods under complementarities. The most well-known mechanism is the *Simultaneous Ascending Auction* (SAA) used by the FCC to allocate spectrum rights (McAfee & McMillan 1996; McMillan 1994). Two other promising mechanisms have been studied in laboratory experiments: the *Adaptive User Selection Mechanism* (AUSM) introduced by Banks, et al. (Banks, Ledyard, & Porter 1989), and a combination of the previous two called the *Resource Allocation Design* (RAD) auction (DeMartini et al. 1998).

These three mechanisms differ in many subtle ways. Two dimensions of particular relevance to this discussion are the semantics of bids and the scope of objects being priced. SAA accepts bids on individual goods, and sets prices on the individual goods. RAD accepts bids on bundles, but still computes prices exclusively for the individual goods. AUSM accepts bids on bundles and makes those bids public, essentially announcing the prices of the bundles.³

All three auctions use ascending rules to assure termination, however none of them have been analyzed for their equilibrium properties. In particular, we do

²The antithesis of anonymous pricing is discriminatory pricing.

³This is an oversimplification. The mechanism announces the prices of the winning bundles, and provides a secondary queue with which agents bidding on smaller bundles can coordinate their bids to displace larger bundles. The prices listed in this queue imply prices for their complements with respect to supersets that have bids.

not know whether, given the final prices and the final allocation, any agent would prefer some other bundle to its allocation. In other words, we do not know if *competitive equilibria* exist for the underlying allocation models.

In this paper we present a method for setting prices on bundles that ensures that equilibrium bundle prices always exist that support the efficient allocation when agents act competitively. We introduce a family of auctions, the *k*-bundle auctions, based on this price setting scheme. Finally, we present some observations about several members of this family.

Model

We consider a benevolent auctioneer with an objective to allocate a set of heterogeneous objects in the most socially efficient manner. There are N agents, indexed by i , and M objects, indexed by j . There are $2^M - 1$ different combinations of objects (excluding the empty set). Let $b \in \{0, 1\}^M$ where $b^j = 1$ implies that object j is an element of the bundle b . For two bundles, b and c , we use the superset notation $b \supset c$ to indicate that $\forall j, b^j \geq c^j$.

The utility that an agent derives from a bundle is given by $v_i(b)$. We assume the existence of a currency, denoted m , and that agents have utility quasilinear in m .

$$U(b) = v_i(b) + m.$$

Further, we assume free disposal, which allows us to assume valuations increase monotonically as items are added to a bundle. Formally, $b \supset c$ implies $v_i(b) \geq v_i(c)$.

A solution to the allocation problem is an assignment of bundles to agents that satisfies the material balance constraint. Let $x_{ib} \in \{0, 1\}$ take the unit value only when b is assigned to i .

The *socially efficient* allocation satisfies the following maximization problem:

$$\begin{aligned} \max \quad & \sum_i \sum_b v_i(b) x_{ib} & (1) \\ \text{s.t.} \quad & \sum_b x_{ib} \leq 1, \forall i \\ & \sum_i \sum_b x_{ib} b^j \leq 1, \forall j \\ & x_{ib} \in \{0, 1\}. \end{aligned}$$

The first constraint implies that each agent receives at most one bundle. The second constraint ensures that no object is allocated more than once.

Let X^* be a solution to this integer program. The value of the solution is simply

$$V(X^*) = \sum_i \sum_b v_i(b) x_{ib}^*.$$

A vector of bundle prices, P , specifies a non-negative real number, p_b , for each b . We admit *nonlinear* prices.

That is, for two disjoint bundles b and c , $p_{b \cup c} = p_b + p_c$ does not necessarily hold.

When adopting nonlinear pricing, it may be necessary to have some regulator enforce the allocation. Otherwise, if $p_{b \cup c} > p_b + p_c$, an agent may wish to purchase b and c separately to assemble the bundle. Similarly, if $p_{b \cup c} < p_b + p_c$, two agents may have an incentive to collude in order to purchase the combined set at a discount, and reallocate the objects later.

We require that the price of a bundle be at least as great as any of its subsets, and that the prices be anonymous—a single price vector is reported to all agents.

A competitive equilibrium is one in which each agent treats prices as exogenous and demands the objects that maximize its surplus. In effect, the agent ignores the impact of its actions on prices. Agent i is in equilibrium at prices P with respect to an allocation X if

$$\sum_b x_{ib} [v_i(b) - p_b] = \max_b [v_i(b) - p_b]. \quad (2)$$

The allocation X^* and the price vector P support a competitive equilibrium if (2) holds for all agents.

Equilibrium Prices

Construction

We start with the solution to equation 2, and wish to construct equilibrium prices to support the optimal allocation X^* . The construction proceeds in two steps. First, we compute prices on the bundles that are assigned in X^* , then we compute prices on the rest of the bundles.

Let \mathcal{B}_+ be the subset of bundles which are assigned in X^* , and \mathcal{B}_- be the unassigned bundles. Let A be the set of all agents, \mathcal{A}_+ the agents which are assigned non-empty bundles, and \mathcal{A}_- be the agents that are excluded. For each agent $i \in \mathcal{A}_-$, we introduce a dummy good ϕ_i which represents the agent's null allocation. Let $\Phi \equiv \{\phi_i | i \in \mathcal{A}_-\}$. Every agent has a valuation of zero for every dummy good. Let $G \equiv \mathcal{B}_+ \cup \Phi$, and $g \in G$.

We have now constructed a pseudo-assignment problem for which we already have the solution. Our goal is to compute prices that support the mapping $A \rightarrow G$. To accomplish this, we adopt the dual program used by Leonard (Leonard 1983) to compute minimal prices for the assignment problem. First, we solve LP_{\min} :

$$\begin{aligned} \min \quad & \sum_g p_g \\ \text{s.t.} \quad & s_i + p_g \geq v_i(g), \forall i, g \\ & s_i, p_g \geq 0, \\ & \sum_i s_i + \sum_g p_g = V(X^*). \end{aligned} \quad (3)$$

The s_i term represents the surplus achieved by agent i . LP_{\min} maximizes the agent's surplus within the range

of equilibrium prices that support the efficient allocation. Note that the introduction of the objects in Φ is simply a trick to include the agents in \mathcal{A}_- in the pseudo-assignment problem. In the solution, $s_i = p_{\phi_i} = 0$ for all $i \in \mathcal{A}_-$.

LP_{\min} has a complementary problem, LP_{\max} , which computes upper bound prices:

$$\begin{aligned} \min \quad & \sum_i s_i \\ \text{s.t.} \quad & s_i + p_g \geq v_i(g), \forall i, g \\ & s_i, p_g \geq 0, \\ & \sum_i s_i + \sum_g p_g = V(X^*). \end{aligned} \quad (4)$$

Clearly the price vector produced in finding a solution to either of these linear programs is a competitive equilibrium for the pseudo-assignment problem—the first constraint requires that an agent cannot gain any more surplus from another bundle than it receives from the one allocated to it.

Given a solution to either LP_{\min} or LP_{\max} we must now set prices on the bundles in \mathcal{B}_- . This can be done quite simply. For all $b \in \mathcal{B}_-$,

$$p_b = \max_i [v_i(b) - s_i]. \quad (5)$$

By construction, these prices cannot distract any agent from the allocation it receives in X^* . Moreover, these prices satisfy the same monotonicity constraint we imposed on utility functions.

Let \overline{P} be a vector of prices where the p_b s are calculated by LP_{\max} and (5). Similarly, \underline{P} is the price vector resulting from solving LP_{\min} and applying (5). We can now consider a range of equilibrium prices that support the optimal allocation.

Lemma 1 For any $k \in [0, 1]$, $k\underline{P} + (1-k)\overline{P}$ is a competitive equilibrium.

Proof: Consider two bundles, b and c , where agent i prefers b to c at \overline{P} . Because \overline{P} and \underline{P} both support the same allocation, i must prefer b to c at \underline{P} . Thus we have

$$v_i(b) - \overline{P}_b \geq v_i(c) - \overline{P}_c,$$

and

$$v_i(b) - \underline{P}_b \geq v_i(c) - \underline{P}_c.$$

In both cases, the relation is invariant to positive scalar transformations. Thus, for $k \in [0, 1]$,

$$k[v_i(b) - \overline{P}_b] \geq k[v_i(c) - \overline{P}_c],$$

and

$$(1-k)[v_i(b) - \underline{P}_b] \geq (1-k)[v_i(c) - \underline{P}_c].$$

Adding the two equations and simplifying gives

$$v_i(b) - k\overline{P}_b - (1-k)\underline{P}_b \geq v_i(c) - k\overline{P}_c - (1-k)\underline{P}_c.$$

	A	B	C	AB	AC	BC	ABC
Agent 1	6	6	5	10	7	8	12
Agent 2	3	3	2	8	5	6	11
Agent 3	4	2	1	7	6	5	10

Table 2: An example with three agents.

	AB	C	ϕ_3
Agent 1	10	5	0
Agent 2	8	2	0
Agent 3	7	1	0

Table 3: The pseudo-assignment problem.

□.

The k parameter here is analogous to the parameter used in the k -double auction (Satterthwaite & Williams 1989). \overline{P}^* and \underline{P}^* represent the range of prices for which supply equals demand. Reducing any price in \underline{P}^* would create excess demand for some of the objects. There is a slightly weaker analogy for the upper bound side. We cannot raise the price of any $b \in \mathcal{B}_+$ without disequilibrating supply and demand. However, in some cases we can increase the prices of bundles in \mathcal{B}_- without adverse affects.

Example

Consider the example in Table 2. The efficient allocation is to assign AB to agent 2 and C to agent 1. Agent 3 gets nothing. This allocation has a social welfare of 13.

In order to construct equilibrium prices, we introduce the dummy good ϕ_3 . We then solve the linear programs LP_{\min} and LP_{\max} for the pseudo-assignment problem in Table 3.

The solution to LP_{\min} for this problem is $p_{AB} = 7$, $p_C = 1$, and of course $p_{\phi_3} = 0$. This leaves $s_1 = 4$, $s_2 = 1$, and $s_3 = 0$. We leave the computation of the solution to LP_{\max} as an exercise for the reader. The final upper and lower equilibrium price vectors are given in Table 4.

Discussion

Unfortunately, by expanding the pricing space we admit inefficient equilibria. For example, the price vector $\langle 2, 2, 2.5 \rangle$ supports a competitive equilibrium for both the problem in Table 1 and the problem in Table 5. However, in the latter case the allocation it supports is inefficient—agent 1 gets both goods. For this example, the price vector $\langle 2, 2, 3 \rangle$ supports the efficient

	A	B	C	AB	AC	BC	ABC
\overline{P}^*	4	4	3	8	6	6	11
\underline{P}^*	4	2	1	7	6	5	10

Table 4: Equilibrium prices.

	A	B	AB
Agent 1	0	0	3
Agent 2	2	2	2
Agent 3	2	2	2

Table 5: An example with multiple equilibrium bundle price vectors.

allocation.⁴ Clearly, one of the challenges in developing an efficient allocation mechanism will be avoiding these degenerate equilibria.

Recently, Bikhchandani and Ostroy (Bikhchandani & Ostroy 1998) (B&O) thoroughly analyzed the equilibrium properties of a particular *package assignment* model. They demonstrated that discriminatory, non-linear pricing is required to ensure that competitive equilibria exist in the single-seller case. However, their allocation model differs from the one presented herein in a crucial respect. In B&O’s model, the seller must desire to sell the packages in exactly the bundles that buyers demand. In our model, we have ignored the seller and assumed that all of the goods are available for reallocation by the auctioneer. More importantly, we assume that the auctioneer is interested in maximizing efficiency, and not revenue.

It is also appropriate to question whether the competitive assumption holds for such small problems. We argue that in this complex domain, competitive equilibrium analysis is a good starting point for analysis. A more thorough treatment of agent incentives and strategies will need to be done in the context of specific auction mechanisms based on the k -bundle prices. It is worth noting that some recent algorithms for computing the optimal allocation from bundle bids have performed surprisingly well when the bids are sparse and exhibit complementarities (Fujishima, Leyton-Brown, & Shoham 1999; Sandholm 1999).

The Family of k -Bundle Auctions

Although we have presented the price setting policy under the assumption of an omniscient mediator, the algorithms can be used directly on bids. Thus, we can combine the k -bundle price policy with the full set of auction parameters (Wurman, Wellman, & Walsh 1998; Wurman, Walsh, & Wellman 1999). This gives potentially hundreds of new auction types to explore. We have not completed an in-depth analysis of any member of this auction family, but have a few preliminary observations.

Throughout this section we assume that bids are of the form $r_i(b)$, where $r_i(b) \in \mathcal{R}_+$, but is not necessarily equal to $v_i(b)$. An agent can bid on any or all bundles. There is an important subtlety in the interpretation of

⁴We present the upperbound price vector that is computed by our algorithm, but in this problem, the efficient solution actually has an equilibrium that is supported by prices on individual goods.

bids by the auctioneer. When an agent does not bid on all bundles, the bids on the other superset bundles must be inferred. For example, if an agent bid 2 on A, and has not bid on AB, the mechanism must assume that it “values” AB at 2.

Two problems plague auctions that allocate complementary preferences. The first is the *exposure* problem: an agent desires a bundle but can only purchase elements of the bundle individually, and in order to avoid getting stuck with only a subset of its desires, does not express its true valuations. The second is the *free rider* problem: one agent is winning a bundle and two agents value the components of the bundle at a greater combined value, but do not to raise their bids in the hopes that the other agent will shoulder more of the burden of displacing the winner. We expect these two problems to be arise to varying degrees in the analysis of particular k -bundle auctions.

Sealed-Bid k -Bundle Auctions

The sealed-bid k -bundle auction accepts bids for all bundles. Then, at a specified time, it computes the prices and allocation *from the bids* using the above method.

First, we consider the sealed-bid bundle auction where $k = 0$. From the following example, it is clear that the incentive compatibility results of the assignment problem (Leonard 1983) do not hold for our generalized allocation problem.

Consider again the example in Table 5. Assume that agents 2 and 3 reveal their true valuations. Can agent 1 be better off by lying about its valuation? The answer is yes. By reporting its true valuation, agent 1 is allocated nothing and receives a surplus of zero. Suppose, instead, that agent 1 reports $r_1(AB) = 5$. The mechanism will calculate the optimal allocation to be the one in which agent 1 receives AB, and a supporting price vector $\underline{P}^* = \langle 2, 2, 2 \rangle$. Thus, by lying, agent 1 receives a surplus of 1.⁵

Now we turn our attention to the sealed-bid bundle auction where $k = 1$. This pricing strategy would discourage manipulations like the one above, because agent 1 would be forced to pay 5. However, like the standard sealed-bid first-price auction, it suffers from potential efficiency losses due to strategic behavior. This can be seen in the example in Table 1 in which agent 1 has an incentive to bid $2 + \epsilon$. If the agent’s information is imperfect, it risks underbidding its opponents and losing the good.

Ascending k -Bundle Auction

We now introduce a mechanism that performs quite well in our early tests. We call it the Ascending k -Bundle (AKB) Auction. In particular, we focus on the situation where $k = 1$. The auction accepts bids on bundles, but requires that an agent’s new bid, $r'_i(b)$, beat $\max(p_b, r_i(b))$ by some fixed ϵ . The auction clears

⁵In fact, this holds for any $r_1(AB) \geq 4$.

the market and closes after no new bids are received within a specified period.

After each bid, the auction calculates a new price vector and announces it as the *price quote*. Unlike in RAD, an agent can compare its bid to the price quote and tell whether it is winning. Because of the manner in which the prices are determined, the agent knows it is winning the good that maximizes its *revealed surplus*, $r_i(b) - p_b$. However, if more than one bundle provides it with the maximal revealed surplus, the agent may not be able to tell from among these which one it is winning. This difficulty could be addressed by posting the tentative winning combinations (though not the winner’s identity).

Because agents can bid on bundles in AKB, they do not suffer from the *exposure* problem in which they are reluctant to express their true valuation for a bundle when they need to purchase the elements individually. However, there is still an incentive for an agent to *free ride*. It may be necessary to adopt activity rules like those in SAA and RAD to encourage participation.

We have not yet performed large scale experimentation with AKB, nor empirically compared it with RAD or AUSM. However in our early experiments with myopic, best response agents, the results are promising.

Conclusion

The allocation of discrete, heterogeneous resources when agents have complementarities in preferences is a very general problem that is likely to arise in many e-commerce applications. In this paper we have established that competitive equilibrium bundle prices always exist to support the efficient allocation. We accomplish this by solving a pseudo-assignment problem for the bundles that are assigned in the optimal allocation, and then constructing prices on the rest of the bundles using a simple rule. We believe that this is an important step in the quest for a mechanism that performs well in the face of complementary preferences. We intend to continue our investigation of the family of k -bundle auctions in general, and the Ascending k -Bundle auction in particular.

Acknowledgments

This research is funded by the DARPA information survivability project and the IBM Institute for Advanced Commerce. Helpful comments were provided by Ennio Stacchetti, Jeff MacKie-Mason, William Walsh, Wolfram Conen, David Parkes, and the participants of the Santa Fe Institute 1997 Graduate Workshop in Computational Economics.

References

- Banks, J. S.; Ledyard, J. O.; and Porter, D. P. 1989. Allocating uncertain and unresponsive resources: An experimental approach. *Rand Journal of Economics* 20:1–22.

- Bikhchandani, S., and Mamer, J. W. 1997. Competitive equilibrium in an economy with indivisibilities. *Journal of Economic Theory* 74:385–413.
- Bikhchandani, S., and Ostroy, J. M. 1998. The package assignment model. Technical report, University of California at Los Angeles.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11:17–33.
- Demange, G.; Gale, D.; and Sotomayor, M. 1986. Multi-item auctions. *Journal of Political Economy* 94:863–72.
- DeMartini, C.; Kwasnica, A. M.; Ledyard, J. O.; and Porter, D. 1998. A new and improved design for multi-object iterative auctions. Technical report, California Institute of Technology.
- Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Gul, F., and Stacchetti, E. 1997. Walrasian equilibrium without complementarities. Technical report, Princeton University and University of Michigan.
- Kelso, A. S., and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50:1483–1504.
- Leonard, H. B. 1983. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy* 91(3):461–479.
- McAfee, R. P., and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature* 25:699–738.
- McAfee, R. P., and McMillan, J. 1996. Analyzing the airwaves auction. *Journal of Economic Perspectives* 10(1):159–175.
- McMillan, J. 1994. Selling spectrum rights. *Journal of Economic Perspectives* 8(3):145–162.
- Milgrom, P. 1989. Auctions and bidding: A primer. *Journal of Economic Perspectives* 3(3):3–22.
- Sandholm, T. 1999. An algorithm for optimal winner determination in combinatorial auctions. In *Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*.
- Satterthwaite, M. A., and Williams, S. R. 1989. Bilateral trade with the sealed bid k -double auction: Existence and efficiency. *Journal of Economic Theory* 48:107–33.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.
- Wurman, P. R.; Walsh, W. E.; and Wellman, M. P. 1999. A parameterization of the auction design space. Submitted for publication.
- Wurman, P. R.; Wellman, M. P.; and Walsh, W. E. 1998. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Second International Conference on Autonomous Agents*, 301–8.