

A Mechanism for Multiple Goods and Interdependent Valuations

Rajdeep K Dash, Alex Rogers, and Nicholas R Jennings

School of Electronics and Computer Science,
University of Southampton,
Southampton, SO17 1BJ, UK.
Fax: +44 (0) 23 8059 2865
{rkd02r, ar2, nrj}@ecs.soton.ac.uk

Abstract. This paper reports on the design of an auction mechanism for allocating multiple goods when the buyers have interdependent valuations. We cast the problem as a multi-agent system consisting of selfish, rational agents and develop an auction mechanism which is efficient, incentive compatible and individually rational. We first discuss the necessary assumptions that *any* mechanism developed for this scenario should satisfy so as to achieve the aforementioned properties. We then present our mechanism and show how it is a generalisation of the Vickrey-Clarke-Groves mechanism.

1 Introduction

Auction mechanisms have been proposed as a solution for a variety of task and resource allocation problems that occur in multi-agent systems (MAS) [2, 11, 13]. A common requirement of these systems is that agents of limited complexity can participate fairly without the risk of being exploited by more complex agents indulging in strategic behaviour. As such, *direct* mechanisms which are *incentive compatible* are most often considered, as under these mechanisms, the equilibrium strategy of all agents is simply to truthfully report their type to the auctioneer¹. Of these mechanisms, the Vickrey-Clarke-Groves (VCG) mechanism is the most widely studied because in addition to the above properties, it ensures that the resulting allocation is efficient (i.e. it maximises the global welfare) and that the mechanism is *individually-rational* (i.e. it guarantees any agent joining the mechanism derives a non-negative utility) [5].

However, a key shortcoming of the VCG mechanism is that it relies on private independent valuations to achieve these desirable properties. Such private valuations arise when an agent forms its valuation of the goods or services based solely on its own observation or signal (e.g. the value of a particular car to an agent depends solely on the agent's own perception of the car's use and is not dependent on the valuations of other bidders). However, the more general case is that valuations are actually *interdependent* (e.g. if the agents' valuations were to consider not only the car's use, but also the potential re-sale value of the car in the future, the valuation would clearly be dependent

¹ The *revelation principle*, which states that any mechanism can be transformed into an *incentive compatible* and *direct* revelation mechanism (ICDRM), thereby guarantees that if a more complex mechanism achieves some desirable properties then there is a corresponding ICDRM that can also achieve them.

on the valuations of other bidders). Now, in auctions with interdependent valuations, the desirable properties of the VCG mechanism no longer hold and the auction is not guaranteed to be efficient.

Interdependent valuations occur most commonly within multi-agent systems when agents have noisy or uncertain estimates of the true value of a good. For example, consider the case of agents bidding for a service in some form of computational economy (as is found, with web services or grid computing). In such cases, the value of a service to an agent is often dependent on the time of response between submitting a request and receiving the desired service. However, in many such cases, the dynamic and open nature of most of these systems means that each agent is only likely to have limited previous experience of a given service and thus it will only have an imprecise estimate of its expected response time. Now if the agent knew the response time of other agents who have used this service (e.g. by asking them about their previous experience or by deducing it from their bidding behaviour), it would be able to form a more accurate estimate of the future response time (by cross-correlating from a broader set of experiences). Hence each agent's valuation is dependent on the signals (in this case, the response time) observed by the other agents bidding for the service and thus we again have interdependent valuations.

Another instance where interdependent valuations have been documented is in the FCC spectrum auctions [3] where it was found that bidders formed their valuations based around the beliefs and actions of other bidders. In these auctions, each bidder wanted to infer from the bidding actions of the other bidders how much they valued the spectrum licenses that were being offered. Thus, whilst each bidder had carried out independent research to gauge the market profitability of these spectrum licenses (i.e. how much money can an agent potentially make by using the license if it wins it), they wanted to use the information gained by the other bidders as well.

To overcome the independent valuation limitation, a number of researchers have developed efficient auctions for interdependent valuation scenarios where a single item is allocated (see section 2 for more details). However, in this work we are interested in the case of multiple items being allocated (i.e. where agents may be interested in combinations of items such as a bundle of services). This extension also allows us to consider the important case of combinatorial allocations. These allocations deal with items exhibiting complementarities and substitutabilities and are shown to be more efficient than multiple concurrent auctions of single goods [11, 13]. Such allocations occur in many real world scenarios such as the grid services and FCC spectrum auction.

To this end, we develop a novel direct mechanism that can allocate multiple items in an interdependent valuation scenario where each agent receives a single-dimensional signal (for example, a time of response in the computational economy or market profitability in the case of the FCC spectrum). We restrict our attention to single-dimensional signals because in an interdependent valuation scenario it is not possible to develop an efficient auction for multi-dimensional signals [6]². Moreover, the single-dimensionality of the signal is not overly restrictive because in many cases the necessary information

² However, Mezzetti [8] shows that if we adopt a two-stage approach to the auction design, we can then achieve efficiency and incentive-compatibility.

can be encompassed into a representative single-dimensional signal. In developing this mechanism, we advance the state of the art in the following ways:

1. We extend the standard VCG mechanism to deal with interdependent valuations in the case of multiple goods in which agents receive a single-dimensional signal.
2. We show that our mechanism is general and demonstrate that it reduces to the VCG mechanism for multiple goods in the case of private values.
3. We prove the economic properties of our mechanism. In particular, we show that it is incentive-compatible, individually rational and efficient. We also analyse its computational properties and show that the mechanism does not impose any additional computational load on the agents, but does so in the case of the centre (as compared to an independent valuation scenario).

The remainder of the paper is organized as follows: section 2 presents related work. Section 3 then develops our auction mechanism for the interdependent valuation scenario. We then provide an explanatory example that highlights how the mechanism works in section 4. In section 5 we prove the economic and computational properties of the mechanism. Finally, we conclude and suggest areas of future work in section 6.

2 Related Work

The VCG mechanism and its various extensions have been used in a variety of resource and task allocation scenarios that occur in MAS [11, 13, 7, 5]. However, in these scenarios, work has invariably concentrated on private valuation situations. Specifically, in the case where an agent observes a single-dimensional signal about the objects it wishes to bid on and this signal determines its valuation. This single-dimensional signal is often referred to as the *type* of the agent.

Recently, however, a number of researchers have started to consider interdependent auctions [7, 4, 6]. In particular, there are currently two main approaches to finding an efficient mechanism for the allocation of items with interdependent valuations. Krishna considers a direct mechanism for efficient allocations for multi-unit single items with single-dimensional signals [7]. In this case, agents submit their interdependent valuation functions, as well as their signals, to a central auctioneer who then decides on the efficient allocation. The payment scheme was then devised so that the agents are incentivised to reveal their signals truthfully.

On the other hand, Dasgupta and Maskin developed an indirect efficient mechanism for the case of two non-identical items, again with single-dimensional signals [4]. In their case, agents make contingent bids rather than submitting their valuation functions and observed signals (i.e. agent 1 submits a range of bids which describes its bid when agent 2 bids a particular value and vice versa). Thus the bidding is more complex than in Krishna's mechanism because the agents have to submit bids based on what other agents might bid, rather than just revealing their valuation function and signals. This bidding becomes even more complex in the indirect mechanism they have developed for the case where multiple items need to be allocated.

Given this, in this paper, we adopt the approach by Krishna, since the bidding is more straightforward for the agents. Specifically we develop a direct mechanism in

order to deal with the allocation of multiple items where each agent receives a single-dimensional signal. A naive extension of the VCG mechanism is known not to work in this case [7] and given this we show how to change the payment scheme in order to achieve the desirable economic properties of the VCG. We should note here that we do not concern ourselves with the problem of multi-dimensionality of these signals since it is known that allowing for multi-dimensionality of signals leads to inefficient allocations [6] in direct mechanisms. If the agents can observe the outcome of their reports, then an efficient allocation with multi-dimensional types is possible [8]. However, we believe that this is impractical in many cases because an agent might not be able to observe the outcome from a report (see [8] for an example). Thus, in this paper we consider direct mechanism where the agents can report on their types only once.

3 The Multiple Good Interdependent Mechanism

In this section, we extend Krishna's approach to develop a mechanism that is incentive-compatible, efficient and individually-rational for the case of multiple goods with single-dimensional signals. In this scenario, there is a set of agents \mathcal{I} . Each agent i , $i \in \mathcal{I}$, observes a signal $x_i \in \mathbb{R}_+$ and forms its valuation $v_i(\cdot)$ based on the vector of signals $\mathbf{x} = [x_1, \dots, x_{\mathcal{I}}]$ (where each element in the vector is observed by one agent and is correspondingly indexed) and the particular allocation $f \in \mathcal{F}$ being implemented (\mathcal{F} denotes the set of all possible allocations). Thus, $v_i : \mathbb{R}_+^{|\mathcal{I}|} \times \mathcal{F} \rightarrow \mathbb{R}_+$. For ease of presentation, we shall denote the set $\mathcal{I} \setminus i$ as $-i$. Furthermore, we shall at times denote $v_i(f, \mathbf{x})$ as $v_i(\cdot)$. Our mechanism, $(\mathcal{M}, \mathbf{r})$, then consists of an allocation rule $\mathcal{M} : \mathbb{R}_+^{|\mathcal{I}|} \rightarrow \mathcal{F}$ which chooses the allocations and a payment rule $\mathbf{r} : \mathbb{R}_+^{|\mathcal{I}|} \rightarrow \mathbb{R}_+^{|\mathcal{I}|}$ which determines the payments r_i to each agent, both being based on the reports of the signal values \mathbf{x} . Finally, we shall denote allocations induced by the true report of x_i (x_{-i} being truthful) as f_0^* . As x_i is decreased, it is quite natural to expect that the allocation which is deemed efficient will change because the valuations of each allocation by the agents would also change. These allocations will be denoted by f_l^* with l being the index of each successive induced allocation as x_i is decreased. Mirroring this, as x_i is increased, the successive efficient allocations are denoted by f_{-l}^* . Now, before presenting our mechanism, we shall discuss the assumptions that are critical for the auctions to be efficient.

Assumption 1 $\frac{\partial v_i}{\partial x_j} > 0 \quad \forall i, j \in \mathcal{I}$

This implies that higher values of the signal lead to higher valuations for the agent. This restricts the signal of the agent to vary in one direction only, thereby making it impossible for an agent to have the same valuation of an allocation for two different signal values. For example, in the case of a computational economy, this would imply that the valuation always increases with rapidity of service (which is x_i).

Assumption 2 $\frac{\partial v_i}{\partial x_i} > \frac{\partial v_j}{\partial x_i} \quad \forall i, j \in \mathcal{I}, i \neq j$.

This implies that an agent's signal affects its own valuation more than it affects the valuation of any other agent. This assumption is the single-crossing condition

analogue in the interdependent scenario [7, 9]. Without this condition, no efficient mechanism can exist. In the case of a computational economy, this implies that the agent puts more credence on the rapidity of service it measured as opposed to the one observed by other agents.

Assumption 3 $\frac{\partial v_i}{\partial x_i}(\cdot, f_p^*) \geq \frac{\partial v_i}{\partial x_i}(\cdot, f_q^*)$ if $p < q$

This implies that if a higher value of x_i induces an allocation f_p^* , then agent i 's value changes more rapidly in this new allocation than in the previous allocation f_q^* . This implies that on receiving a higher x_i , the centre allocates a set of goods to i in the new allocation f_p^* where i 's valuation changes more rapidly, than in the previous set f_q^* . To better explain this assumption, consider a situation where there are two services to be allocated and an agent has a complementary valuation of those services. Suppose that the agent is allocated a particular service when $x_i = \alpha$. Now, if x_i is increased, there will come a point $x_i = \beta > \alpha$ when it will be efficient to allocate both services to the agent (since from assumption 2, its valuations will increase more rapidly than that of other agents). This assumption then implies that the rate of change of the valuation with respect to x_i is greater in this new allocation than in the previous one. Consider, for example, two agents bidding for two services being in a grid service economy. Then suppose that as x_i is increased, it first becomes more efficient to allocate one good (denote this allocation as f_{-1}^*) and then both goods to agent i (denote this allocation as f_{-2}^*). Then this assumption implies that $\frac{\partial v_i}{\partial x_i}(x_i, x_{-i}, f_{-2}^*) \geq \frac{\partial v_i}{\partial x_i}(x_i, x_{-i}, f_{-1}^*)$ i.e. agent i 's valuation increases more rapidly with x_i when it is allocated both goods rather than only one.

Given these assumptions, our mechanism then proceeds as follows:

1. Each agent i transmits to the centre its valuation function $v_i(f, \mathbf{x})$ for all the possible allocations $f \in \mathcal{F}$. This function is also over all possible values of \mathbf{x}
2. Each agent i also transmits its observed signal \hat{x}_i .³
3. The centre then computes the optimal allocation f_0^* which is calculated as:

$$f_0^* = \arg \max_{f \in \mathcal{F}} \left(\sum_{i \in \mathcal{I}} v_i(f, \hat{\mathbf{x}}) \right) \quad (1)$$

4. The centre also calculates the payment r_i made by each agent i . To do this, the centre first finds the m next best allocations as the reported signal \hat{x}_i is decreased successively, until the presence of i makes no difference to the allocations. That is, find allocations $f_1^* \dots f_m^*$ and the signal values z_i^l such that:

$$z_i^l = \inf \left\{ y_i : \sum_{i \in \mathcal{I}} v_i(f_l^*, y_i, \mathbf{x}_{-i}) = \sum_{i \in \mathcal{I}} v_i(f_{l+1}^*, y_i, \mathbf{x}_{-i}) \right\} \quad (2)$$

(where each allocation f_l^* is different) until:

$$z_i^m = \inf \left\{ y_i : \sum_{i \in \mathcal{I}} v_i(f_{m-1}^*, y_i, \mathbf{x}_{-i}) = \sum_{i \in \mathcal{I}} v_i(f_m^*, y_i, \mathbf{x}_{-i}) \right\} \quad (3)$$

³ Of course, \hat{x}_i may not be equal to x_i . However, we prove in section 5 that it is a best strategy for the agent to set $\hat{x}_i = x_i$.

where the allocation f_m^* is the optimal allocation when i does not exist i.e.

$$f_m^* = \arg \max_{f \in \mathcal{F}} \sum_{j \in -i} v_j(f, \mathbf{x})$$

Then the transfer ⁴ to buyer i is:

$$r_i = \sum_{l=0}^{m-1} \left[\sum_{j \in -i} v_j(f_l^*, z_i^l, \mathbf{x}_{-i}) - \sum_{j \in -i} v_j(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right] \quad (4)$$

The above scheme rests upon making an agent derive a utility equal to the marginal contribution that its presence makes to the whole system of agents (which is the same intuition as used in the VCG). Thus the additional part of this mechanism is to take into account the effect that an agent's signal x_i has on the overall utility of the system.

This mechanism is general and is shown (below) to reduce to the well-known multiple-good private value model if we take the case of independent valuations i.e when $v_i(\mathbf{x}, \cdot) = v_i(x_i)$. Then the optimal allocation (from equation 1) is:

$$f_o^* = \arg \max_{f \in \mathcal{F}} \left(\sum_{i \in \mathcal{I}} v_i(f, \hat{x}_i) \right)$$

To calculate the payment scheme, we first note that with independent valuations x_i only affects $v_i(\cdot)$. Thus repeatedly decreasing x_i , until the stopping condition on equation 3, does not change the valuation of the other agents $-i$ on the different allocations. This then implies that in the payment (as computed by equation 4) all the terms cancel each other, except for the first and last, leading to a payment of:

$$r_i = \sum_{j \in \mathcal{I} \setminus i} v_j(f_o^*, \hat{x}_j) - \sum_{j \in \mathcal{I} \setminus i} v_j(f_m^*, \hat{x}_j) \quad (5)$$

This is exactly the payment scheme for the multiple-good private values model. Thus, this shows that the classical VCG mechanism is an instance of the generalised mechanism developed here. Furthermore, notice that assumption 2 is automatically satisfied in this independent valuation scenario, since $\frac{\partial v_j}{\partial x_i} = 0$ in such a scenario. Also, since an increase in x_i would only increase $v_i(\cdot, x_i)$, any increase in x_i that induces a new allocation would imply that the rate of change of $v_i(\cdot, x_i)$ with respect to x_i is higher in the new allocation than in the previous allocation. Thus, assumption 3 is also automatically satisfied in the independent valuation scenario.

4 Example of an Interdependent Valuation Scenario

In order to better explain how the mechanism operates to achieve efficiency and incentive-compatibility, in this section, we present an example that demonstrates how it computes the efficient allocation and the payments. We will also consider the assumptions which we made in section 3 and show how the mechanism fails when these do not hold.

We consider a very simple case, namely that with two agents 1 and 2 bidding for two different spectrum licenses A and B . The set of possible allocations consists of

⁴ If the transfer is negative it implies that buyer i pays to the centre.

four members, which are $\mathcal{F} = \{(AB, \emptyset), (A, B), (B, A), (\emptyset, AB)\}$. In this case, each agent perceives a particular signal x_i that determines the market profitability of the spectrum licenses. Table 1 shows the valuations of player 1 and 2 for each allocation as well as the sum of their valuations.

Table 1. Valuations of the players with each allocation

Allocation	$v_1(f, \mathbf{x})$	$v_2(f, \mathbf{x})$	$v_{\mathcal{I}}(f, \mathbf{x})$
(AB, \emptyset)	$4x_1 + 2x_2$	0	$4x_1 + 2x_2$
(A, B)	$2x_1 + x_2$	$x_1 + 2x_2$	$3x_1 + 3x_2$
(B, A)	$x_1 + x_2$	$0.5x_1 + 2x_2$	$1.5x_1 + 3x_2$
(\emptyset, AB)	0	$x_1 + 4x_2$	$x_1 + 4x_2$

We shall now consider how agent 1 views the mechanism as it reports its signal x_1 . The explanation for agent 2 is the same and is therefore omitted. Figure 1 shows how the value of each allocation varies for agents 1, 2 and the set of agents \mathcal{I} , as agent 1's reported signal x_1 is increased. We denote agent 1 by i and agent 2 by $-i$ to demonstrate how this works in cases of more than two agents. Suppose that agent 1 has observed $x_1 = 1.5$ and agent 2 has observed a value of $x_2 = 2$. Then from the figure, we see that the efficient allocation in this case is $f_0^* = (A, B)$ (the efficient allocation is the one that maximises the value of \mathcal{I}). Furthermore, the values of x_i at which it becomes more efficient to implement allocations $f_1^* = (\emptyset, AB)$ and $f_{-1}^* = (AB, \emptyset)$ are $z_i^0 = 1$ and $z_i^{-1} = 2$ respectively (shown in figure 1). Hence we can calculate the overall utility that agent 1 derives from reporting truthfully, which from equation 4, is $v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^0 \mathbf{x}_{-i}) - v_{-i}(f_1^*, z_i^0 \mathbf{x}_{-i}) = 5 + 5 - 9 = 1$. Now, any report in the range $1 \leq x_i \leq 2$ will induce the same allocation and transfer and thus agent 1 has no incentive to report x_i in this range different from the truthful value. If agent 1 reports $x_i > 2$, it will then derive a utility of $v_i(f_{-1}^*, \mathbf{x}) + v_{-i}(f_{-1}^*, z_i^{-1} \mathbf{x}_{-i}) - v_{-i}(f_0^*, z_i^{-1} \mathbf{x}_{-i}) + v_{-i}(f_0^*, z_i^0 \mathbf{x}_{-i}) - v_{-i}(f_1^*, z_i^0 \mathbf{x}_{-i}) = 10 + 0 - 6 + 5 - 9 = 0$, which is less than what it would derive from truthful reporting. Thus agent 1 would not over-report its observed value. The reason why this occurs is because, as shown in figure 1, $v_i(f_{-1}^*, \mathbf{x}) - v_i(f_0^*, \mathbf{x})$ is always less than $v_{-i}(f_0^*, z_i^{-1} \mathbf{x}_{-i}) - v_{-i}(f_1^*, z_i^{-1} \mathbf{x}_{-i})$ when the true value of x_i is in the range $1 \leq x_i \leq 2$. If, on the other hand, the agent reports $x_i < 1$, it would then derive a utility of $v_i(f_1^*, \mathbf{x}) = 0$ which is again less than what it would derive from truthful reporting. We have thus demonstrated how an agent finds it in its best interest to report truthfully (see section 5 for a more general proof).

The mechanism is guaranteed to work in the above example because the valuations satisfy the assumptions presented in section 3. We will now show how this mechanism would fail if ever, one of these assumptions does not hold.

In order to show what happens when assumption 1 fails, consider only the single good A . Suppose that agent 1 has a valuation of $(x_1 - 2)^2 + x_2$ for good A and agent 2 still has the same valuation of $0.5x_1 + 2x_2$. Then the auctioneer in this case has to decide only between two allocations, namely $\mathcal{F} = \{(A, \emptyset), (\emptyset, A)\}$. With these valuations, it is efficient to allocate good A to agent 2 when $2.25 - \sqrt{[(2.25)^2 - (4 - x_2)]} \leq x_1 \leq 2.25 + \sqrt{[(2.25)^2 - (4 - x_2)]}$. If $x_1 \leq 2.25 - \sqrt{[(2.25)^2 - (4 - x_2)]}$ agent 1 obtains




Fig. 1. Valuations of 1, 2 and \mathcal{I} for each bundle as x_1 is increased

the good and pays $2x_2$ according to equation 4. If $x_1 \geq 2.25 - \sqrt{[(2.25)^2 - (4 - x_2)]}$, then agent 1 again obtains good A, but this time, it pays 6 (again using equation 4). Thus, it is always in the interest of agent 1 to state that its signal is in the lower range if its signal happens to occur in either of these ranges. Although assumption 1 may seem to be required only for our mechanism to work, this is not so, as it is required for *any* efficient, incentive-compatible mechanism [9].

Now consider that the valuations of the good A are such that $v_1((A, \emptyset), \mathbf{x}) = 2x_1 + x_2$ and $v_2((\emptyset, A), \mathbf{x}) = 3x_1 + x_2 - 6$ (thus assumption 2 is not satisfied). In this case, it is efficient to allocate A to agent 1 when $x_1 < 6$ and to agent 2 otherwise. However, it is not possible to achieve an efficient mechanism in this case, since agent 1 will always state $x_1 < 6$ no matter what the real value of x_1 is. In the case of our mechanism, agent 1 pays $x_2 - 6$ if it allocated the good. Since $v_1(A, \emptyset)$ is always higher than this, agent 1 will thus lie and always state a value of $x_1 < 6$. This problem can again be shown to extend to be symptomatic of any mechanism rather than our mechanism [4]. Notice that with the original valuations in table 4, such a situation would not arise.

We next consider valuations that break assumption 3. Here the valuations of agents 1 and 2 for the allocation $f = (AB, \emptyset)$ are $v'_1((AB, \emptyset), \mathbf{x}) = 0.5x_1 + 2x_2$ and $v'_2((AB, \emptyset), \mathbf{x}) = 3.5x_1$ as shown in figure 2⁵. Since $v_{\mathcal{I}}$ remains the same for all the allocations, then z_i^{-1} is still the same as shown in figure 2. Using these modified valuations, agent 1 derives a higher utility of 1.75 (using equation 4 and the valua-

⁵ Of course, in practice, agent 2 having a valuation for nothing is highly unlikely to occur. However, we need to use this particular valuation in this case due to the simplicity of our example in order to demonstrate what happens when one of the assumptions fails.




Fig. 2. Modified valuations of 1, 2 and \mathcal{I} for allocations (AB, \emptyset) and (A, B) as x_1 is increased

tion function) if it reports $x_i > 2$ thereby leading to the mechanism no longer being incentive-compatible. The reason this occurs is because if assumption 3 is broken we then have that $v_i(f_{-1}^*, \mathbf{x}) - v_i(f_0^*, \mathbf{x}) > v_{-i}(f_{-1}^*, z_i^{-1} \mathbf{x}_{-1}) - v_{-i}(f_0^*, z_i^{-1} \mathbf{x}_{-1})$ as shown in figure 2. As a result, the agent has an incentive to lie and quote a higher value than z_i^{-1} . Notice that this did not occur with the original valuations. Again this assumption is required in order to find an efficient, incentive-compatible mechanism and is thus not idiosyncratic to our mechanism [4].

Having thus illustrated the working of our mechanism and the necessity of the assumptions via the use of an example, we now turn to formally proving the properties of our mechanism.

5 Properties of the Mechanism

We next prove the properties of our mechanism. We first consider the economic properties; namely that it is incentive-compatible, efficient and strategy proof, whilst intuitively explaining why the mechanism has the aforementioned properties. We then consider the computational properties of the mechanism, showing that the mechanism does not impose any added complexity on the agents' bidding process compared to what it would already face in an independent value scenario. However, it does increase the complexity of calculating the payment, a computational load borne by the centre.

5.1 Economic Properties

Proposition 1. *The mechanism is incentive-compatible in ex-post Nash Equilibrium.*

A mechanism is *incentive-compatible* in ex-post Nash Equilibrium if it is a best response strategy for the players to reveal their types truthfully even after they have complete information about the signal values \mathbf{x} .

Proof. Let $v_{-i}(\cdot) = \sum_{j \in -i} (v_j(\cdot))$ and $v_{\mathcal{I}}(\cdot) = \sum_{i \in \mathcal{I}} (v_i(\cdot))$. Suppose now that all players except i report their signals truthfully (i.e. $\hat{\mathbf{x}}_{-i} = \mathbf{x}_{-i}$). Let the optimal allocation when i reports truthfully be f_0^* . We can then analyse the utility $u_i(\cdot)$ that agent i derives by reporting a certain \hat{x}_i . There are two cases that should be analysed namely when $\hat{x}_i < x_i$ and $\hat{x}_i > x_i$. The utility of an agent on reporting $\hat{x}_i = x_i$ is:

$$u_i(f_0^*, \mathbf{x}) = v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \quad (6)$$

Now suppose an agent reports $\hat{x}_i \neq x_i$ but this does not change the optimal allocation f_0^* implemented. Then, $u_i(f_0^*, \mathbf{x}) = u_i(f_0^*, \hat{x}_i, \mathbf{x}_{-i})$. This is because if the allocation does not change then the agent derives the same value $v_i(f_0^*, \mathbf{x})$ and payment as the signals $z_i^0 \dots z_i^m$ are computed by the centre. Now consider the case that an agent reports $\hat{x}_i < x_i$ such that this changes the allocation. Then some other optimal allocation, which is necessarily one of the allocations f_1^*, \dots, f_m^* , is implemented. Denote the resulting allocation when $\hat{x}_i < x_i$ as f_n^* (i.e. $z_i^n < \hat{x}_i \leq z_i^{n-1}$).

The utility that the agent gets from this new allocation is then:

$$u_i(f_n^*, \mathbf{x}) = v_i(f_n^*, \mathbf{x}) + \sum_{l=n}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \quad (7)$$

The difference, $D_n = u_i(f_0^*, \mathbf{x}) - u_i(f_n^*, \mathbf{x})$ between truthful reporting and under reporting (as given by equations 6 and 7 respectively) is:

$$\begin{aligned} D_n &= v_i(f_0^*, \mathbf{x}) - v_i(f_n^*, \mathbf{x}) + \sum_{l=0}^{n-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^0, \mathbf{x}_{-i}) - v_{-i}(f_n^*, z_i^n, \mathbf{x}_{-i}) - v_i(f_n^*, \mathbf{x}) \\ &\quad + \sum_{l=1}^n \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_l^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) \end{aligned}$$

Since $\frac{\partial v_{-i}(f_l^*, \mathbf{x})}{\partial x_i} \geq 0$, we thus have:

$$\begin{aligned} D_n &> v_i(f_0^*, \mathbf{x}) + v_{-i}(f_0^*, z_i^0, \mathbf{x}_{-i}) - v_{-i}(f_n^*, z_i^n, \mathbf{x}_{-i}) - v_i(f_n^*, \mathbf{x}) \\ &> v_i(f_0^*, \mathbf{x}) - v_i(f_0^*, z_i^0, \mathbf{x}_{-i}) - v_i(f_n^*, \mathbf{x}) + v_i(f_n^*, z_i^n, \mathbf{x}_{-i}) + v_{\mathcal{I}}(f_0^*, z_i^0, \mathbf{x}_{-i}) \\ &\quad - v_{\mathcal{I}}(f_n^*, z_i^n, \mathbf{x}_{-i}) \end{aligned}$$

However, by construction we know that $v_{\mathcal{I}}(f_0^*, z_i^0, \mathbf{x}_{-i}) > v_{\mathcal{I}}(f_n^*, z_i^n, \mathbf{x}_{-i})$ and from assumption 3 we also know that $v_i(f_0^*, \mathbf{x}) - v_i(f_0^*, z_i^0, \mathbf{x}_{-i}) > v_i(f_n^*, \mathbf{x}) - v_i(f_n^*, z_i^n, \mathbf{x}_{-i})$. We thus have $D_n \geq 0$. On the other hand, if an agent reports $\hat{x}_i > x_i$

and this induces an allocation f_{-n}^* , then the utility it derives is:

$$u_i(f_{-n}^*, \mathbf{x},) = v_i(f_{-n}^*, \mathbf{x},) + \sum_{l=-n}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i},) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i},) \right) \quad (8)$$

The difference, $D_{-n} = u_i(f_0^*, \mathbf{x}) - u_i(f_{-n}^*, \mathbf{x})$ between truthful reporting and under reporting (as given by equations 6 and 7 respectively) is:

$$\begin{aligned} D_{-n} &= v_i(f_0^*, \mathbf{x}) - v_i(f_{-n}^*, \mathbf{x}) - \sum_{l=-n}^{-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) - v_i(f_{-n}^*, \mathbf{x}) - \sum_{l=-n}^{-1} \left(v_{\mathcal{I}}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{\mathcal{I}}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &\quad + \sum_{l=-n}^{-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_{-n}^*, z_i^{-n}, \mathbf{x}_{-i}) - v_i(f_{-n}^*, \mathbf{x}) - v_i(f_0^*, z_i^{-1}, \mathbf{x}_{-i}) + v_i(f_0^*, \mathbf{x}) \\ &\quad - \sum_{l=-n+1}^{-1} \left(v_i(f_l^*, z_i^{l-1}, \mathbf{x}_{-i}) - v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) \right) \end{aligned}$$

Using assumption 3 implies that $D_{-n} \geq 0$. We thus see that i derives highest utility when reporting $\hat{x}_i = x_i$.

Proposition 2. *The mechanism is efficient.*

This implies that the centre finds the outcome such that $f^* = \arg \max_f \sum_{i \in \mathcal{I}} v_i(f, \mathbf{x})$.

Proof. The above is a result of the incentive-compatibility of the mechanism. Since the goal of the centre is to achieve efficiency, then given truthful reports, the centre will achieve efficiency.

Proposition 3. *The mechanism is individually rational.*

A mechanism is *individually rational* if there is an incentive for agents to join it rather than opting out of it. We begin by assuming that the utility an agent derives from not joining the mechanism is 0. Then, we need to prove that the utility an agent derives in the mechanism is always ≥ 0 .

Proof. Given that the agents are incentivized to report truthfully, agent i derives utility:

$$\begin{aligned} u_i(f_0^*, \mathbf{x}) &= v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{-i}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{-i}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) + \sum_{l=0}^{m-1} \left(v_{\mathcal{I}}(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_{\mathcal{I}}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &\quad - \sum_{l=0}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \end{aligned}$$

Since $v_{\mathcal{I}}(f_l^*, z_i^l, \mathbf{x}_{-i}) = v_{\mathcal{I}}(f_{l+1}^*, z_i^l, \mathbf{x}_{-i})$ (from equation 2):

$$\begin{aligned} u_i(f_0^*, \mathbf{x}) &= v_i(f_0^*, \mathbf{x}) - \sum_{l=0}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_{l+1}^*, z_i^l, \mathbf{x}_{-i}) \right) \\ &= v_i(f_0^*, \mathbf{x}) - v_i(f_0^*, z_i^0, \mathbf{x}_{-i}) + v_i(f_m^*, z_i^m, \mathbf{x}_{-i}) \\ &\quad + \sum_{l=1}^{m-1} \left(v_i(f_l^*, z_i^l, \mathbf{x}_{-i}) - v_i(f_l^*, z_i^{l+1}, \mathbf{x}_{-i}) \right) \end{aligned}$$

From equation 3, $v_i(f_m^*, z_i^m, \mathbf{x}_{-i}) = 0$. Now, since $\frac{\partial v_i(K, \mathbf{x})}{\partial x_i} \geq 0$, thus $u_i(f_0^*, \mathbf{x}) > 0$.

5.2 Computational Properties

In order for a mechanism to be of use in real world scenarios, we must not only consider its economic properties but also its computational complexity. An important distinction is to differentiate between the computational load which is imposed on the agents within the auction and that imposed on the auctioneer or centre. Specifically, we will analyse the computational properties of the mechanism as opposed to that faced by agents in a standard VCG mechanism. In so doing, we aim to quantify the computational cost that the added richness of this mechanism (namely the ability to express interdependent valuation) imposes.

Outcome Determination. In our mechanism, the centre will need to solve equation 1, which is similar to the winner determination equation in the VCG mechanism, in order to determine the efficient allocation. In both cases the computation involves solving a combinatorial allocation problem which is, in the general case, NP-hard [12]. In fact, the size of the set over which the optimisation is carried out is the same in both cases since this is determined by the number of items $|M|$. Thus our mechanism imposes no additional computational load in terms of the centre calculating the allocation. However, in terms of calculating the payments to the agents, our mechanism does impose a larger computational load. In the case of the VCG mechanism, calculating the payment involves performing the winner determination problem $|\mathcal{I}|$ times over the reduced set of agents $\mathcal{I} \setminus i$ (see [5] for more details). However in our case, the centre needs to successively reduce the value of the report from each agent (and calculate the optimal allocation at each stage) until it reaches an allocation which is the optimal one for the reduced set of agents $\mathcal{I} \setminus i$ (see equations 2 and 3). In the worst case scenario, we have to traverse through all possible allocations (except the efficient one) when calculating the different z_i^l for each agent $i \in \mathcal{I}$. For m goods in a combinatorial auction, this requires $2^m - 1$ calculations and is thus exponential in complexity. However, typically, the number of allocations that need to be traversed (i.e the K_i^l) will be much less than 2^m and there is some redundancy between the calculation of the K_i^l in between the agents in \mathcal{I} . We will exploit this redundancy in future work so as to reduce the computational load on the centre.

Preference Formulation. In the case of a direct mechanism such as the VCG mechanism or our mechanism, the agents do not have additional computational load in formulating their preferences over all possible outcomes. This is because the agents transmit their observed signal θ_i to the centre and thus do not actually compute $v_i(K, \theta)$ over all $K \in \mathcal{K}$. Rather it is the centre which performs this calculation for each agent when solving the winner determination problem. Thus, our mechanism in this case does not add any computational load on the agents.

Strategy Selection. In the VCG mechanism the agent knows *a priori* that it has a dominant strategy, and thus this computational problem does not arise. In our case, an agent has an *ex-post* Nash strategy. Thus if all the agents are behaving rationally, there is no computational load on the agent in this particular case. However, if it becomes common knowledge that some agent is not playing its best-response strategy (i.e. some agent is not rational) then the agents will have to search through their space of strategies again to find their best-response.

Thus, we can observe that there is no additional computational load on the agents when compared with a standard VCG mechanism. Thus we can use the computationally efficient bidding languages developed for VCG mechanisms [11, 10]. This is important since in many proposed applications, whilst the centre may have significant computational power, the agents will be represented by distributed devices of limited computational power.

6 Conclusions and Future Work

In this paper we considered an important class of auctions in which the bidders have interdependent valuations (based on a single dimensional signal measured by each bidder) and bid for multiple goods. In this context, we have significantly extended the standard VCG mechanism and proved that the ensuing mechanism has the ideal economic properties of being efficient, incentive compatible and individually rational. Our mechanism is general and reduces to the VCG mechanism whenever there are independent valuations (as seen in section 3). Thus, we can visualise our mechanism being used even in MAS where the designer is unsure whether the valuations are interdependent or not.

Whilst we have presented our mechanism in terms of resource allocation, it can be easily converted into a task allocation scenario. In such a scenario, agents will first submit cost functions instead of valuation functions. Then, we need to perform a minimisation instead of a maximisation in equations 1, 2 and 3 and take supremums instead of infimums in equations 2 and 3. With these changes, the mechanism still conserves both its computational and economic properties in the task allocation scenario.

Our future work in this area concerns two issues. The first issue will concentrate on how to design mechanisms which take into consideration multi-dimensional signals. Such signals are known to better characterise the preferences of agents in certain MAS such as in a procurement auction where both the price and date of delivery are important [1]. These mechanisms are known not to be efficient in an interdependent scenario as a result of the impossibility result due to Jehiel and Moldovanu [6]. However, we aim to calculate the loss in efficiency when taking into consideration multi-dimensionality. The second issue is concerned with the question of the computational complexity of the

resulting mechanism. We have seen that allowing for interdependent valuations comes at the cost of additional computational complexity on the centre. We intend to investigate methods to reduce this load, by reducing the space of allocations that need to be considered when computing the payments to the agents (as discussed in section 5). Our aim is to achieve a mechanism whose complexity is no greater than that of performing the task of winner determination in the underlying auction.

Acknowledgements

The authors would like to thank R.Mason and anonymous reviewers for helpful comments. Rajdeep K. Dash is funded by a BAE Systems studentship and an ORS scholarship. This research was undertaken as part of the ARGUS II DARP (Defence and Aerospace Research Partnership). This is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University and Southampton University, and is funded by the industrial partners together with the EPSRC, Ministry of Defence (MoD) and Department of Trade and Industry (DTI).

References

1. M. Bichler and J. Kalagnanam. Bidding languages and winner determination in multi-attribute auctions. *European Journal of Operational Research*. to appear.
2. S.H. Clearwater, editor. *Market-Based Control- A paradigm for distributed resource allocation*. World Scientific, 1996.
3. P. Cramton. The fcc spectrum auctions: An early assessment. *Journal of Economics and Management Strategy*, 6(3):431–495, 1997.
4. P. Dasgupta and E. Maskin. Efficient auctions. *Quarterly Journal of Economics*, 115:341–388, 2000.
5. R. K. Dash, D. C. Parkes, and N. R. Jennings. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
6. P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–59, 2001.
7. V. Krishna. *Auction Theory*. Academic Press, 2002.
8. C. Mezzetti. Mechanism design with interdependent valuations: Efficiency and full surplus extraction. Technical report, University of North Carolina, February 2003.
9. R. Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38:175–208, 1971.
10. N. Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 1–12, 2000.
11. D. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, University of Pennsylvania, May 2001.
12. M. H. Rothkopf, A. Pekec, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.
13. T. Sandholm. Making markets and democracy work: A story of incentives and computing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1649–1671, 2003.