

# A Budget-balanced, Incentive-compatible Scheme for Social Choice

Boi Faltings

Artificial Intelligence Laboratory (LIA)  
Swiss Federal Institute of Technology (EPFL)  
IN-Ecublens, CH-1015 Ecublens, Switzerland  
`boi.faltings@epfl.ch`

**Abstract.** Many practical scenarios involve solving a *social choice problem*: a group of self-interested agents have to agree on an outcome that best fits their combined preferences. We assume that each outcome presents a certain utility to an agent and that the best outcome is the one that maximizes the sum of these utilities. We call a mechanism for solving social choice problems *incentive-compatible* if for each agent, the behavior that maximizes its own utility is also the one that maximizes the group's utility.

One way to achieve incentive-compatibility is the Vickrey-Clarke-Groves (VCG) tax ([5]) mechanism. However, it produces a surplus of taxes that cannot be redistributed to the agents and can severely reduce agents' utilities. Game theory has shown that it is not possible to have a general scheme that is incentive-compatible, budget-balanced and guarantees a Pareto-efficient solution.

We present a scheme that sacrifices Pareto-efficiency to achieve budget balance while being both incentive-compatible and individually rational. On randomly generated social choice problems, the scheme results in significantly better overall agent utility than the VCG tax mechanism.

## 1 Social Choice Problems

Many practical situations involve social choice: a group of agents has to choose from a fixed set of choices an outcome that best fits their combined preferences. For example, a group going out to have dinner together has to choose a restaurant that fits everyone's preferences. Tenants of a building have to decide on features of a planned renovation. Spectrum has to be divided up among different mobile telephone providers.

A mechanism for solving a social choice problem takes as inputs declarations of the agents' utilities for each outcome, and outputs as a solution the optimal choice and possibly other information.

Social choice problems become difficult to solve when agents have conflicting preferences, as each agent will exaggerate its preferences to obtain a better outcome for itself. It is possible to counteract this tendency using *tax* schemes where agents have to pay for the preferences they claim. An example of such tax

schemes are auctions: the social choice is to decide who receives the good, and the winner has to make a payment that depends on how strongly he claims to value the good. Note that while auctions are a special case of social choice, reverse and double auctions as the outcome space is not independent of the agents (an agent may not agree to give up an item).

In an *incentive-compatible* (IC) mechanism, the incentives of each agent are aligned with those of the group: the behavior that optimizes the utility of an individual agent also optimizes the utility of the group. When utility optimization is left to the social choice mechanism, this often corresponds to each agent being best off declaring its preferences truthfully; this is called *truthful* or *strategyproof*. Such a mechanism makes life easy for the agents since they do not have to speculate to obtain the best outcome. It also avoids choosing a suboptimal outcome because of such speculation.

A well-known mechanism for achieving IC is the Vickrey-Clarke-Groves (VCG) tax ([5]) mechanism. It assumes that the mechanism chooses an outcome that maximizes the sum of agents' utilities (called the *Pareto-efficient* outcome, PE), and makes each agent pay a tax that is calculated so that the agent cannot gain from misreporting its utility. Furthermore, the VCG tax is *individually rational* (IR) in that the tax paid by an agent never exceeds the utility gain it gets from participating in the optimization as opposed to letting the other agents pick the outcome.

Any tax mechanism produces a surplus of taxes that cannot be redistributed to the agents without losing the incentive-compatible property, i.e. they are not *budget-balanced* (BB). In game theory, it has been shown that all incentive-compatible mechanisms that apply to general social choice problems and always generate a Pareto-efficient outcome must use a tax of a form similar to the VCG tax ([7, 9]). It has further been shown that such a mechanism cannot be budget-balanced ([7, 10]).

In the special case of auctions, the surplus can be used to pay the sellers of the goods; the resulting VCG scheme is called the *Vickrey* auction protocol. However, in many cases, there is no use for this surplus. It reduces agents' utilities, and creates incentives for the receiver of the surplus to manipulate the setting to maximize taxes. For example, in spectrum allocation, governments can obtain huge windfall profits by creating scarcity, but in so doing hurt the public in general.

Since the lack of budget-balance may cause huge losses to the involved agents, we consider IC and IR mechanisms that do not always choose the Pareto-efficient outcome, but achieve budget balance. In particular, we propose a novel randomized scheme that can be applied to any tax mechanism to achieve budget balance. It preserves all IC and IR properties of the underlying tax scheme, but generates suboptimal solutions, i.e. it is not Pareto-efficient. On randomly generated constraint optimization problems, it can be seen that the loss of agent utility due to the lack of PE is in general much smaller than the loss they would incur through the taxes in a VCG mechanism. The mechanism has much better per-

formance than previous proposals, and shows an interesting new direction for solving social choice problems.

We model social choice problems as general constraint optimization problems. In Section 2, we formally define the framework and give an example of a social choice problem formulated in this way. In Section 3, we review the properties of tax schemes. Section 4 introduces our new mechanism, and Section 5 presents examples and experimental results. Section 6 presents extensions of the mechanism. Section 7 presents related work, and Section 8 the conclusions.

## 2 Modeling Social Choice Problems

A social choice problem can be simply formulated as a choice among a set of possible outcomes. However, it is often useful to further structure this outcome space. In particular, we assume that the outcome space is the set of solutions to a *constraint satisfaction problem* (CSP). It is defined by a set of variables that can be assigned values in associated domains. A solution is a combination of value assignments to all variables such that a set of *constraints* is simultaneously satisfied. Note that the space of possible outcomes is independent of the agents' actions; thus, there can only be negative externalities and it is for example not possible to model a double auction as a social choice problem.

Based on the CSP formulation, we model social choice problems as multi-agent constraint optimization problems. These are CSP where a set of *agents* has declared relations that specify the utilities they attach to different value combinations. Formally, they are defined as follows:

**Definition 1.** A *discrete* multi-agent constraint optimization problem (*MCOP*) is a tuple  $\langle A, X, D, C, R \rangle$  where:

- $A = \{A_1, \dots, A_k\}$  is a set of agents.
- $X = \{x_1, \dots, x_n\}$  is a set of variables.
- $D = \{d_1, \dots, d_n\}$  is a set of domains of the variables, each given as a finite set of possible values.
- $C = \{c_1, \dots, c_p\}$  is a set of constraints, where a constraint  $c_i$  is a function  $d_{i1} \times \dots \times d_{il} \rightarrow \{0, 1\}$  that returns 1 if the value combination is allowed and 0 if it is not.
- $R = \{r_1, \dots, r_o\}$  is a set of relations, where a relation  $r_i$  is a function  $d_{i1} \times \dots \times d_{il} \rightarrow \mathbb{R}$  giving the utility of choosing each combination of values.
- $R_i$  is the subset of  $R$  that gives the relations associated with agent  $A_i$ .

A solution to an MCOP is a consistent assignment to the underlying CSP that maximizes the sum of the agents' utilities. Formally, we define:

**Definition 2.** An assignment  $V$  is a combination of values  $x_1 = v_1 \in d_1, \dots, x_n = v_n \in d_n$ .

We write  $r_i(V)$  and  $c_i(V)$  for the result of applying  $r_i$  or  $c_i$ , respectively, to the relevant variables with the assignments in  $V$ .

An assignment  $V$  is consistent if all constraints are satisfied, i.e.  $(\forall c_i \in C)c_i(V) = 1$ .

An MCOP is solvable if there is at least one consistent assignment.

We write  $V_R^*$  for the consistent assignment such that the sum of the utilities obtained by the relations corresponding to  $R$  is maximal; if there are several such assignments it is the one that is lexicographically smallest. We use  $v_R^*(x_i)$  for the value of  $x_i$  in that assignment.

The solution to a MCOP is the assignment  $V_R^*$  where  $R$  is the set of all relations of the MCOP.

In this paper, we only consider solvable MCOP. When an MCOP is unsolvable, the social choice problem itself has no solution.

As an example, consider the following social choice problem. A building has 4 tenants, represented by agents  $A_1$  through  $A_4$ . The tenants have to agree on three architectural features  $X = \{x_1, x_2, x_3\}$  of a planned renovation. For the three features, we have options  $x_1 \in d_1 = \{A, B, C\}$ ,  $x_2 \in d_2 = \{A, C\}$  and  $x_3 \in d_3 = \{B, C\}$ . For structural reasons, the combination of  $x_2 = A$  and  $x_3 = C$  is not allowed, but all other combinations are feasible. This restriction is modelled by a constraint.

Each agent attaches different utilities to a feature or feature combination. Specifically:

- for agent  $A_1$  and any of the three features, choosing  $A$  has a utility of -1,  $B$  a utility of 0 and  $C$  a utility of +1.
- for agent  $A_2$ , the utility is determined by the combination of features  $x_1$  and  $x_2$  according to the following table:

$(x_1, x_2)$	(B,C)	(C,C)	(A,C)	(B,A)	(C,A)	(A,A)
utility	2	2	1	-1	-1	-3

- for agent  $A_3$ , the utility is determined by the combination of features  $x_1$  and  $x_3$  according to the following table:

$(x_1, x_3)$	(A,C)	(C,C)	(A,B)	(C,B)	(B,C)	(B,B)
utility	2	2	1	-1	-1	-3

Finally, agent  $A_4$  attaches the following utilities to combinations of assignments:

$(x_1, x_2, x_3)$	utility	$(x_1, x_2, x_3)$	utility	$(x_1, x_2, x_3)$	utility
(A,A,B)	3	(B,A,C)	1	(B,C,C)	-1
(B,A,B)	3	(B,C,B)	1	(C,A,C)	-1
(A,A,C)	1	(C,A,B)	1	(C,C,B)	-1
(A,C,B)	1	(A,C,C)	-1	(C,C,C)	-3

The optimal solution to this problem is the combination  $x_1 = C, x_2 = C, x_3 = C$ , which provides a total utility of 4. However, this solution can only be found if agents report their true utilities. In this example, agent  $A_4$  could exaggerate its utility for combination  $x_1 = B, x_2 = A, x_3 = B$  to 10. This would give it an overall supposed utility of 5, and make it the one chosen by the mechanism. Agent  $A_4$  has improved its own utility from -3 to +3, but the true overall utility has decreased from +4 to -2. Other agents may follow similar reasoning and perturb the result further.

### 3 Achieving Incentive-Compatibility

The incentives of each agent can be aligned with those of the group by making each agent pay a tax reflecting the cost that their preferences are causing to others. A well-known mechanism is the Vickrey-Clarke-Groves (VCG) tax mechanism ([14, 5, 8]). Its application for multi-agent decision making has already been proposed in [6].

In the VCG mechanism, each agent pays the difference in other agents' utilities of the optimal solution when it is not present and the optimal solution when it is. Recall that  $R_i$  is the set of relations  $R_i$  imposed by agent  $A_i$ . The tax, called the Clarke tax, is then:

$$VCGtax(A_i) = \sum_{r_l \in R \setminus R_i} r_l(V_{R \setminus R_i}^*) - r_l(V_R^*)$$

The VCG tax has the effect of making the objectives of each individual agent that of optimizing the sum of all agent's utilities.

If the optimization is left to the social choice mechanism, this makes it a dominant strategy equilibrium for each agent to declare its utilities truthfully.

In our example, we consider the following solutions:

Solution	$u(A_1)$	$u(A_2)$	$u(A_3)$	$u(A_4)$	total
$v_R^* = (C, C, C)$	3	2	2	-3	4
$v_{R \setminus R_1}^* = (A, C, B)$	0	1	1	1	3
$v_{R \setminus R_2}^* = (C, C, C)$	3	2	2	-3	4
$v_{R \setminus R_3}^* = (B, C, B)$	1	2	-3	1	1
$v_{R \setminus R_4}^* = (C, C, C)$	3	2	2	-3	4

Thus, the VCG tax payments of the agents would be:

Agent	VCG tax
$A_1$	$3-1 = 2$
$A_2$	$2-2 = 0$
$A_3$	$4-2 = 2$
$A_4$	$7-7 = 0$

The example shows that the truth-inducing property of the tax can come at a high cost: in this example the total tax paid, 4 units, completely erases the utility of 4 that the agents jointly get out of their renovation.

### 4 Budget balanced mechanisms for social choice

The main difficulty with applying tax schemes to social choice is that they generate a surplus of taxes that reduces overall agent utility and creates unwanted incentives for whatever party gets this surplus. We now show a simple scheme that is always strictly budget balanced, but produces sub-optimal solutions. We assume that the agents are solving an MCOP whose variables, domains and

constraints are fixed and known. Furthermore, we are going to assume that the MCOP is solvable, i.e. it has at least one consistent assignment.

The basic idea is to randomly select an agent or a group of agents whose relations will receive a lower priority in the optimization. In return, this agent or group of agents will be paid the tax collected from the remaining agents. The scheme is by definition budget balanced since all taxes are paid between the agents themselves. Since the agents receiving the tax have no influence on the declarations and thus the taxes of the remaining agents, the scheme preserves all incentive-compatibility properties of the tax scheme itself. However, it chooses solutions that are not optimal for all agents and is thus not Pareto-efficient.

We now present several budget balanced social choice mechanisms based on this idea. We assume that the tax scheme is that of a VCG tax. However, the mechanisms can be applied with any tax scheme, including the first-price tax where agents pay the declared utilities in the chosen solution.

In this paper, we consider the following mechanism:

- Mechanism 1**
1. Each agent  $A_i \in \mathcal{A}, i = 1..k$  is asked to state its relations.
  2. Choose an excluded coalition  $E$  of one or more agents using a method that does not depend on the relations stated by the agents.
  3. Compute the assignment:

$$S_E = V_{R \setminus R_E}^*$$

where  $R_E = \bigcup_{A_i \in E} R_i$ . Optionally, if there are several equally optimal assignments, choose the one with the best utility according to the relations in  $R_E$ .

4. Make each agent  $A_i$  pay to agents in  $E$  the VCGtax for the solution  $S_E$ :

$$\begin{aligned} \text{pay}(A_i \rightarrow E) &= \text{VCGtax}_{-E}(A_i) \\ &= \sum_{r_m \in R \setminus (R_i \cup R_E)} r_m(V_{R \setminus (R_i \cup R_E)}^*) - r_m(V_{R \setminus R_E}^*) \end{aligned}$$

and distribute the tax among the agents in  $E$  according to some predetermined scheme.

The excluded coalition can be chosen by any mechanism that does not depend on the utility declarations of the agents. In the interest of fairness, it will often be useful to make this choice randomly. The excluded coalition can consist of one or more agents. In most cases, it will be best to choose only a single agent and let the optimization take into account the relations of a maximum number of agents. However, we will see later that larger coalitions may be useful in certain circumstances.

We now show several properties of Mechanism 1.

**Proposition 1.** *Mechanism 1 is incentive-compatible ex-post.*

**Proof.** Consider an agent  $A_i$ .

When  $A_i \in E$ , the agent's declarations have no influence on the outcome nor its tax (which is equal to 0), so it cannot gain by misreporting.

When  $A_i \notin E$ , then the chosen solution is optimal for a social choice problem where  $A_i$  is included. Each agent pays the VCG tax corresponding to that problem. This tax is known to be incentive-compatible *ex-post*.

**Proposition 2.** *Mechanism 1 is individually rational ex-post.*

**Proof.** Consider an agent  $A_i$ .

When  $A_i \in E$ , the mechanism chooses a solution that does not consider the utilities of  $A_i$ . Its utility is not worse than if it had not participated in the mechanism at all, and it pays no tax. Thus, it is individually rational to participate.

When  $A_i \notin E$ ,  $A_i$  is included in the optimization, and it pays the VCG tax. This scheme is known to be individually rational ex-post.

**Proposition 3.** *Mechanism 1 is budget balanced ex-post.*

**Proof.** All taxes are paid to agents in the excluded coalition  $E$ , so no tax surplus or deficit remains to be distributed.

For the example problem, assume that the mechanism chooses to randomly leave out each of the 4 agents individually with probability 1/4. We then have the following payments:

Solution	$\text{tax}(A_1)$	$\text{tax}(A_2)$	$\text{tax}(A_3)$	$\text{tax}(A_4)$
$v_{R \setminus R_1}^* = (A, C, B)$	-5	2	1	2
$v_{R \setminus R_2}^* = (C, C, C)$	5	-7	2	0
$v_{R \setminus R_3}^* = (B, C, B)$	0	0	-2	2
$v_{R \setminus R_4}^* = (C, C, C)$	0	0	0	0
Average tax	0	-5/4	1/4	1

The mechanism is obviously not Pareto-efficient, as it does not always chose the optimal solution as the final result. However, while the VCG tax mechanism chooses the Pareto-efficient solution, the tax payments cause considerable utility loss. As we have seen before, in the optimal solution of this example, the joint utility of the agents is 4 and is completely eaten up by the sum of the VCG taxes which is also 4. Thus, the expected utility to the community of agents is 0. In contrast, in Mechanism 1 the expected utility is

$$1/4 \cdot (3 + 4 + 1 + 4) = 12/4 = 3$$

which is significantly better. In fact, for this example it is better than the VCG mechanism no matter what agent is excluded from the optimization.

The mechanism may have a choice of several solutions  $S_E$  that have equal utility for the agents except  $E$ , but different utilities for agents in the excluded coalition  $E$ . This can lead to variations in the efficiency of the mechanism. For example, the solution  $(A, A, B)$  is also optimal for  $v_{R \setminus R_2}^*$ , but it has a total utility of -1. If this had been chosen as the solution in the scenario given above, the expected utility in Mechanism 1 would be 7/4 instead of 3 (but still better than in the VCG mechanism). This shows the importance of choosing the best

solution for agents in the excluded coalition as well. This optimization could also be done by letting agents in  $E$  choose which of several equivalent solutions  $S_E$  is to be chosen to minimize their costs.

## 5 Examples and experimental results

As another example, consider an auction of a single item among three agents  $A_1, A_2$  and  $A_3$ . It can be represented by a variable  $x$  that represents the final allocation of the good by an integer 1, 2 or 3 indicating which agent gets the good. Let the agents' valuations be expressed by the relations  $r_1, r_2$  and  $r_3$  on  $x$  as follows:

$$\begin{array}{c|ccc} x = & 1 & 2 & 3 \\ \hline r_1 & a & 0 & 0 \\ r_2 & 0 & b & 0 \\ r_3 & 0 & 0 & c \end{array}$$

and assume that  $a < b < c$ , i.e.  $A_3$  values the good the most. We assume that the mechanism chooses as excluded coalition a single agent, where each of the three agents is chosen randomly with probability  $1/3$ . We have the following solutions  $S_i = v_{R \setminus R_i}^*(x)$ :

$$\begin{array}{c|ccc} \text{Solution} & S_1 & S_2 & S_3 \\ \hline x = & 3 & 3 & 2 \end{array}$$

which gives us the following expected taxes and utilities:

$$\begin{array}{c|ccc} A_i & E[\text{tax}] & pr(x = i) & E[u(A_i)] \\ \hline A_1 & 1/3 (-b) & 0 & b/3 \\ A_2 & 0 & 1/3 & b/3 \\ A_3 & 1/3 b & 2/3 & 2c/3 - b/3 \end{array}$$

We can verify that no agent has an incentive to misreport its valuation:

- if agent  $A_1$  overreports a valuation  $a'$  so that  $a' > b > a$ , its true expected utility drops from  $b/3$  to  $a/3$ . Underreporting has no effect.
- if agent  $A_2$  underreports a valuation  $b'$  so that  $b > a > b'$ , then its true expected utility drops from  $b/3$  to  $a/3$ . If it overreports  $b'$  so that  $b' > c > b$ , then its true expected utility drops to  $2b/3 - c/3$ .
- if agent  $A_3$  underreports a valuation  $c'$  so that  $c > b > c'$ , then its true expected utility drops to  $c/3$ . Overreporting has no effect.

In comparison, in a VCG tax scheme, the Vickrey auction, agent  $A_3$  always gets the good and pays tax  $b$ , and both other agents get nothing. Only agent  $A_3$  has an expected utility of  $c - b$ . Thus, agents  $A_1$  and  $A_2$  are always better off, whereas  $A_3$  is better off only as long as  $c \leq 2b$ . This condition is likely to be satisfied in competitive markets where valuations tend to be close to one another.




The major difference with classical auction schemes is that this way of allocating the good does not produce any revenue for a third party. Such a *revenue-free auction* is often desirable for public goods such as airport slots, water or pollution rights, and the use of distribution networks.

It is possible to construct cases where in spite of the wasted tax, the VCG mechanism would still achieve better overall efficiency. This would arise when leaving one agent out of the optimization would give only a marginally better result for the remaining agents, but a significantly worse result for the agent that was excluded. Consider the following example:  $n$  people have to go from Geneva to London. They can get a group ticket for up to  $n - 1$  people on Swiss at a cost of \$100 per person, or a group ticket on British Air for up to  $n$  people at a cost of \$110 per person. They can also buy individual tickets for a business jet at a cost of \$10'000 per person.

The optimal solution for any subset of  $n - 1$  agents is to buy a group ticket on Swiss, forcing the remaining agent to buy an individual ticket. Mechanism 1 will choose one of these solutions, so the total cost to all agents is  $\$100(n - 1) + \$10'000$ , and no agent pays any tax. For  $n = 10$ , this amounts to \$10'900.

A VCG mechanism chooses the overall best solution: buy a group ticket on British Air. The total expense is  $\$110n$ , but on top of this each agent has to pay a tax of  $\$(110 - 100)(n - 1)$ . Thus, the total expense for all agents is  $\$110n + \$10n(n - 1)$ . For  $n = 10$ , this amounts to \$2'000, significantly less than with Mechanism 1. In fact, the VCG scheme will be better as long as  $n < 33$ .



**Fig. 1.** Costs for randomly generated constraint optimization problems where each variable belongs to a different agent. It shows the total cost to all agents of the optimal solution, the total VCG tax paid by all agents, the degradation of the total cost when the suboptimization solution is used, and the average VCG tax paid by one agent.

On randomly generated problems, we have observed that on average the VCG tax is much larger than the degradation in solution quality incurred by using a

suboptimal solution. Figure 1 shows a comparison on randomly generated problems with 5 to 11 variables. The problems involve as many agents as variables. For each agent, we randomly generate several binary relations (between 2 variables) where each value combination is assigned a valuation between 0 and 1 with a uniform random distribution, and the number of constraints is limited so that the problem is solvable. We apply Mechanism 1 where each agent is excluded individually with equal probability.

Figure 1 shows the following quantities:

- the solid line shows the total cost of coexistence in the optimal solution. This is the difference of the total agent utilities if each of them could choose the solution it preferred and the ones obtained in the optimal joint solution.
- the dashed line (lowest in the figure) shows the average degradation in the total agent utility when one agent’s constraints are not taken into account in the optimization.
- the lightly dotted line (second lowest in the figure) shows the average VCG tax that an agent would have to pay in the VCG (Clarke) tax mechanism.
- the densely dotted line (second highest in the figure) shows the total of all VCG taxes that would have to be paid by all agents.

The experiment shows that the VCG tax mechanism makes agents pay an amount of tax that is almost comparable to the utility loss due to their coexistence. While the tax per agent tends to decrease slightly with the number of agents, the total amount of taxes and thus the loss of social welfare continues to increase with problem size. On the other hand, the cost of the degradation incurred by having a single agent excluded of the optimization is much smaller. In fact, it is comparable or even below the average VCG tax for a *single* agent. This means that on average, even the agent that is excluded from the optimization would tend to get a comparable or better utility than it can expect in the VCG mechanism!

Another experiment has been conducted on resource allocation in the transportation (or communication) network shown by the graph in Figure 2, where each agent has different (randomly varied) costs for the arcs. We randomly generate tasks which require using a path between two points in the network. For each task, we calculate up to three shortest disjoint paths and define a decision variable whose domain is the cross product of the path to use and the agent that executes it. A further value is provided that corresponds to not executing the task at all. Constraints specify that no two tasks be assigned to paths that share an arc. Each agent is asked to evaluate the costs for the three different paths and thus state its utility (task payoff - cost) if it were assigned the task and the corresponding path.

Figure 3 shows the performance of Mechanism 1 compared to an optimization with a VCG mechanism. The bars show the average total utility to all agents in the optimal (shaded) and optimal with one agent excluded (black) solutions. The curves show the average total amount of tax in a VCG mechanism, and the resulting net total utility. It can be seen that even as the number of tasks increases, the utility of the suboptimal solution remains close to that of the




Fig. 2. Network used for resource allocation experiment.




Fig. 3. Utilities achieved by different mechanisms in the experiment.

optimal solution. While the total amount of VCG tax levels off as the number of tasks increases, there is still a large amount of tax that needs to be wasted and causes utility loss to the agents.

## 6 Collusion

It is a well-known fact that VCG tax schemes are vulnerable to collusion: coalitions of agents can act together to achieve unfair advantages over the others. In our example, suppose that agent  $A_4$ , who fares particularly badly in the optimal solution, bribes agent  $A_3$  to help it impose solution  $x_1 = A, x_2 = A, x_3 = B$ . The result lowers  $A_3$ 's utility by 1 but increases  $A_4$ 's utility by 6, so  $A_4$  can pay  $A_3$  3 units for its trouble and both will benefit from the manipulation.  $A_3$  and  $A_4$  can impose this solution by each adding a ternary relation between all three variables that would give utility 100 to this value and 0 to all other combinations. Since the solution would remain the same if either  $A_3$  or  $A_4$  was removed, no agent would pay any taxes, so the manipulation comes for free.

In auctions, collusion can be avoided by using mechanisms that make bidders pay their bid rather than the second highest bid. Similarly, we can define a tax mechanism for social choice, which we call the *first-price tax*, where every agent pays as tax the utility gain it gets in the chosen solution, i.e.:

$$Ftax(A_i) = \sum_{r_l \in R_i} r_l(V_R^*)$$

Obviously, the first-price tax is not incentive-compatible, as agents have an interest to claim a lower utility for the optimal solution and a higher utility for the case where they are excluded. However, it does protect against overclaiming true utilities since an agent that does so will be forced to pay a higher tax than its true utility should the particular solution be chosen. Furthermore, there are ascending-price elicitation schemes, such as the Ausubel auction, that allow first-price schemes to become incentive-compatible in certain cases.

Mechanism 1 can be readily adapted to use a first-price tax rather than a VCG tax as the underlying tax scheme. It is instructive to see what the effect of such a modification is in the example of the revenue-free auction. We would obtain the following expected taxes and utilities:

$A_i$	$E[tax]$	$pr(x = i)$	$E[u(A_i)]$
$A_1$	$1/3 (-c)$	0	$c/3$
$A_2$	$b/3 - c/3$	$1/3$	$c/3$
$A_3$	$c/3$	$2/3$	$c/3$

Thus, in this example, if all agents report the truth, the mechanism will make them all have equal expected utility. Note however that it is obviously in the interest of all agents to speculate by underclaiming utilities. This process can be supported using an ascending-price mechanism similar to English auctions that can be readily adapted to this scenario.

## 7 Related Work

The social choice problem has been the subject of considerable interest in economics, game theory and more recently computer science research, and there is therefore a large amount of related work that is impossible to survey accurately and completely here. Examples of surveys are Moulin ([9]) and recently Chung and Ely ([4]), but there are numerous others. Game theory has largely focussed on the feasibility rather than design of actual mechanisms. Fundamental results ([7, 10]) show that it is impossible to have a budget balanced, Pareto-efficient and incentive-compatible mechanism for the general case.

A first approach to design feasible mechanisms is based on relaxing one of the conditions. Vickrey-Clarke-Groves mechanisms ([14, 5, 8]), in particular the Clarke tax ([5]), are Pareto-efficient, incentive-compatible and individually rational, but in general not budget balanced.

One can easily imagine mechanisms which are incentive-compatible, budget balanced but not efficient, such as simply choosing a solution at random or

choosing a solution that is optimized for agents according to a fixed priority sequence. The approach described here is similar, but comes much closer to an optimal result and thus provides much higher utility to the agents.

Parkes et al. ([12]) have investigated VCG mechanisms which are only approximately incentive-compatible in order to achieve budget balance. Since it is not known a priori what benefits manipulation can bring in a particular case, such a mechanism places a burden on agents who need to evaluate potential manipulations for possible gains.

The dAGVA mechanism ([3, 1]) is an example of a mechanism that is Pareto-efficient, incentive-compatible, and budget balanced and individually rational on average. However, it requires a-priori knowledge of the true probability distributions of the agent's preferences, which is rarely available in practice.

A second approach is to relax the requirement of a general mechanism that works for all valuation structures, and design a mechanism specifically for a particular scenario. Recent work on automated mechanism design ([13]) has shown that given the exact valuations for each of the agents' types, it is computationally feasible to search for mechanisms that have all desired properties. However, this process requires that the uncertainty about agent's preferences is limited to a finite set of types rather than continuous valuations.

## 8 Conclusions

The internet has enabled the creation of networked enterprises consisting of multiple agents. So far, most protocols for optimizing their collective behavior have assumed cooperative behavior but neglected the presence of self-interest.

When self-interest has to be taken into account, the best existing solutions for optimal behavior are based on auctions. However, auctions are not budget balanced: they generate a surplus that reduces agents' utilities and creates unwanted incentives for the party that receives it. This is the case in particular for the Clarke tax ([5]), the most well-known mechanism for dealing with self-interest in social choice.

Since economists have shown many impossibility results (for example, [7, 10]) that prove that it is impossible to combine incentive-compatibility, optimality and budget balance, it is unlikely that this can be overcome in general.

This paper has presented a mechanism that chooses a solution that is optimal for all but a group of excluded agents. In this way, any tax or auction scheme can be made budget-balanced by returning the surplus to the excluded agents. The important observation is that the quality of the solution optimized for all but a small group of agents is very close to that of the solution optimized for all agents. In fact, the loss of utility incurred by this suboptimal solution is much smaller than the loss that would be incurred through the wasted taxes. Thus, it provides an attractive scheme for achieving budget-balance that can be applied to any auction scheme.

In contrast to other proposals for budget balanced mechanisms, our proposal is general and applies to any quasilinear utility function. It also does not require any a-priori knowledge of agent's preferences.

## Acknowledgements

I would like to thank Tuomas Sandholm for numerous insights regarding incentives and computing, and Adrian Petcu, Santiago Macho-Gonzalez and Quang-Huy Nguyen for running experiments.

## References

1. K. Arrow: "The property rights doctrine and demand revelation under incomplete information," in M. Boskin (ed.): *Economics and Public Welfare* Academic Press, 1979
2. K.J. Arrow: "Social Choices and Individual Values," Wiley and Sons, New York, 1963
3. C. d'Aspremont and L.-A. Gerard-Varet: "Incentives and Incomplete Information," *Journal of Public Economics* **11**, pp. 25-45, 1979
4. K.-S. Chung and J.C. Ely: "Ex-Post Incentive-Compatible Mechanism Design," *Working Paper*, Department of Economics, Northwestern University
5. E.H. Clarke: "Multipart pricing of public goods," *Public Choice* **11**, pp. 17-33, 1971
6. E. Ephrati and J. S. Rosenschein: "The Clarke tax as a consensus mechanism among automated agents," *Proceedings of the 9th National Conference on Artificial Intelligence*, pp. 173-178, San Jose, California, July 1991.
7. J. Green, J.J. Laffont: "Incentives in public decision making," *Studies in Public Economics* **1**, North-Holland, 1979
8. T. Groves: "Incentives in Teams," *Econometrica* **41**, pp.617-663, 1973
9. H. Moulin: "Axioms of Cooperative Decision-making," *Econometric Society Monographs* **15**, Cambridge University Press, 1988
10. R.B. Myerson, M.A. Satterthwaite: "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* **29**, pp. 265-281, 1983
11. N. Nisan and A. Ronen: "Computationally feasible VCG Mechanisms," *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 2000
12. D.C. Parkes, J. Kalagnanam, M. Eso: "Achieving Budget-Balance with Vickrey-based Payment Schemes in Exchanges," *17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pp. 1161-1168, 2001
13. T. Sandholm: "Automated mechanism design: A New Application Area for Search Algorithms," in F. Rossi (ed.): *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP-2003)*, pp. 19-36, 2003
14. W. Vickrey: "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance* **16**, pp. 8-37, 1961