

# Incentive-compatible Social Choice

Boi Faltings

Artificial Intelligence Laboratory (LIA)  
Swiss Federal Institute of Technology (EPFL)  
IN-Ecublens, CH-1015 Ecublens, Switzerland

## Abstract

*Many situations present a social choice problem where different self-interested agents have to agree on joint, coordinated decisions. For example, power companies have to agree on how to use the power grid, and airlines have to agree on how to schedule takeoffs and landings.*

*Mechanisms for social choice are called incentive-compatible when cooperative behavior is optimal for all parties. The most well-known examples of incentive-compatible mechanisms are auctions. However, the party that receives the auction revenue has an incentive to manipulate the outcome to increase the revenue. For example, a power grid operator has an interest to reduce capacity and drive up prices. Conversely, if it provides sufficient capacity to every user it derives no revenue to cover its costs.*

*We present a novel mechanism for social choice that is incentive-compatible without generating a payment surplus. We give several examples of applications where it solves the social choice problem without unwanted incentives, and provides significantly better overall utility than any other known mechanism.*

## 1. Social Choice Problems

Agents in the modern world increasingly have to coordinate their actions, typically because they involve sharing resources or cooperating to achieve a task. The *social choice problem* models such scenarios: a group of agents has to choose an outcome that best fits their combined preferences.

Social choice problems can be solved by simple protocols. Among them, the most commonly used is *first-come first-served*: the agent that first presents a request will be able to determine the outcome that is most favorable to it. Theater tickets, hotel rooms, and other scarce resources are distributed in this way. Another mechanism is *voting*: pick the outcome that most agents

vote for. Another possibility is to randomly flip a coin to choose the outcome.

All these mechanisms are designed to make it impossible for agents to manipulate the outcome by their actions. However, the drawback is that the outcomes are usually far from optimal: agents could achieve much higher combined utility through cooperation. For example, people who bought theater tickets months in advance may receive other invitations in the meantime and no longer be interested in using the tickets. People who live in different areas may vote for an unsuitable option in a city planning decision.

Intelligent agent technology offers the possibility to use more complex protocols to obtain a better solution. Coordination, cooperation and optimization algorithms can serve to obtain decisions that better achieve the agents' objectives.

## 2. Social Choice as Constraint Optimization

A social choice problem is a choice among a set of possible outcomes. In order to express preferences on outcomes with manageable complexity, it is 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. Constraint satisfaction has been used with great success for a wide variety of problems, including design, resource allocation, scheduling and planning.

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. The solution of the MCOP is the solution of the underlying CSP where

the sum of agents' utilities is highest. Formally, we define the following notation:

**Definition 1** A 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  $A_i$ .

**Definition 2** An assignment  $V$  is a combination of values  $x_1 = v_1 \in d_1, \dots, x_n = v_n \in d_n$ . It is consistent if all constraints are satisfied, i.e.  $(\forall c_i \in C)c_i(V) = 1$ .

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  $V$ . 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 write  $v_R^*(x_i)$  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*.

An *MCOP* is *solvable* if there is at least one consistent assignment. In this paper, we only consider solvable *MCOP*. When an *MCOP* is unsolvable, the social choice problem itself has no solution.

For *MCOP* of practically interesting size, it will in general not be feasible to compute the optimal assignment  $V_R^*$ . Instead, the mechanism and agents together will only compute an approximation to this optimal solution. We use the notation  $\overline{V}_R$  to denote the assignment returned by this mechanism when optimizing for the set of relations  $R$ , and  $\overline{v}_R(x_i)$  for the value of  $x_i$  in this assignment.

The optimization could be carried out by a central agent that computes the optimal solution using constraint optimization packages such as ILOG Solver or CPLEX. Open constraint optimization techniques ([4]) can be used to limit the amount of information that agents have to communicate to this central solver. The computation could also be distributed among different agents using distributed constraint satisfaction techniques ([5]).

### 3. Example

As an example, consider a city where each neighbourhood has a separate natural gas supply network. The social choice problem is to decide what company gets to supply gas in which neighbourhood. Thus, there is one variable for each neighbourhood whose domain is the set of companies that could serve it. Each company is represented by an agent. Its relations show its utilities, given as the expected revenue minus the cost of providing the service.

Note that this problem cannot be solved by auctioning the rights to serve each neighbourhood, since this would allow the winning company the option of not providing the service by not using the right. In the social choice problem, this option is ruled out by the fact that every variable must be given a value.

For the example, let there be three companies represented by agents  $A_1$ ,  $A_2$  and  $A_3$ . Each company has a capacity to serve at most two neighbourhoods, and can provide the service to any of the neighbourhoods at the following costs:

Agent	$A_1$	$A_2$	$A_3$
Cost	3	4	5

We consider a scenario with three neighbourhoods  $n_1$ ,  $n_2$  and  $n_3$  to be served. We assume that the expected revenue from customers for each neighbourhood is 10, and furthermore serving  $n_1$  and  $n_2$  together creates additional opportunities and results in extra revenue of 10 to whoever serves this combination. This gives rise to the following social choice problem:

variables:  $x_1, x_2$  and  $x_3$

domains:  $\{1, 2, 3\}$

constraints: at most 2 variables have the same value

relations: each  $R_i$  contains 4 relations  $r_{ij}, j = 1..4$ :

$$\left\{ \begin{array}{l} j=1..3: \quad r_{ij}(x_j) = 10 \text{ if } x_j = i, 0 \text{ otherwise;} \\ \quad \quad r_{i4}(x_1, x_2) = 10 \text{ if } x_1 = x_2 = i, 0 \text{ otherwise} \end{array} \right\}$$

The optimal solution is  $x_1 = 1, x_2 = 1, x_3 = 2$ .

### 4. Incentive-compatibility

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 a good, and the winner has to make a payment that depends on how strongly he claims to value the good. A

method for social choice that includes selection of a solution and a set of side (tax) payments is called a *mechanism*.

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. Such mechanisms rule out the possibility that an agent can obtain unfair gains at the expense of other agents. However, an agent could still hurt the other agents through incompetent actions that are not in its own best interest.

When utility optimization is left to the social choice mechanism, IC often corresponds to each agent being best off declaring its preferences truthfully; this property is called *truthful* or *strategyproof*. Such a mechanism makes life easy for the agents since they do not have to choose an optimal strategy themselves, but can leave this to the mechanism. However, it requires that the mechanism finds the true optimal solution, which may be a problem in practice.

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 ([13, 2, 7]). Its application for multi-agent decision making has already been proposed in [3].

In the VCG mechanism, each agent pays the difference in other agents' utilities of the optimal assignment when it is not present,  $V_{R \setminus R_i}^*$ , and the optimal assignment for all agents,  $V_R^*$ :

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

In the example, we would obtain the following taxes:

Agent	$A_1$	$A_2$	$A_3$
Utility	24	6	0
VCGtax	$27-6 = 21$	$29-24 = 5$	0

The VCG tax has the effect of making the objectives of each individual agent that of optimizing the sum of all agent's utilities. This can be seen by considering the utility  $U_t(A_i)$  that agent  $A_i$  achieves from the result of a social choice using a VCG tax:

$$\begin{aligned} U_t(A_i) &= \left[ \sum_{r_l \in R_i} r_l(V_R^*) \right] - VCGtax(A_i) \\ &= \sum_{r_l \in R_i} r_l(V_R^*) + \sum_{r_l \in R \setminus R_i} r_l(V_R^*) - r_l(V_{R \setminus R_i}^*) \\ &= \sum_{r_l \in R} r_l(V_R^*) - \sum_{r_l \in R \setminus R_i} r_l(V_{R \setminus R_i}^*) \end{aligned}$$

Since  $A_i$  has no influence on  $\sum_{r_l \in R \setminus R_i} r_l(V_{R \setminus R_i}^*)$ , it maximizes its own utility by maximizing the utility of the group,  $\sum_{r_l \in R} r_l(V_R^*)$ . Furthermore, we can see that the utility an agent derives from participating in the mechanism,  $U_t(A_i)$ , is never smaller than the utility it would get if it would let the other agents choose their best solution. If it was, then the agent  $A_i$  could show a better overall solution than  $V_R^*$  by choosing the assignment for the other agents as in  $V_{R \setminus R_i}^*$  and selecting its own best matching value with it. Thus, a mechanism based on VCG taxes is *individually rational*: every agent has an interest to participate in it. Finally, since  $\sum_{r_l \in R \setminus R_i} r_l(V_{R \setminus R_i}^*) \geq \sum_{r_l \in R \setminus R_i} r_l(V_R^*)$ , the VCG tax is always non-negative.

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). This creates incentives for whoever receives the surplus to manipulate the outcome. For example, a power grid operator has incentives to create bottlenecks to drive up fees. Conversely, if it provides ample capacity it derives no revenue to maintain its network. In game theory, it has been shown that all incentive-compatible mechanisms that apply to general social choice problems and always generate an optimal outcome must use a tax similar to the VCG tax ([6]), and that such mechanisms cannot be budget-balanced ([6, 9]).

If the social choice mechanism computes the true optimal solution, the VCG tax makes it a dominant strategy equilibrium for each agent to declare its utilities truthfully to this mechanism. When optimization is not perfect, the VCG tax is still incentive-compatible, as each agent's individual utility is still maximized by actions that make the mechanism obtain the optimal solution for all agents. However, as shown in [10], these actions may include non-truthful utility declarations if these would help the optimization algorithm to converge to a better overall solution. In a distributed social choice setting, in general incentive-compatibility is more important than truthfulness.

However, with an imperfect optimization, two complications may arise:

- in rare cases, due to shortcomings of the optimization method, it could happen that for some  $A_i$ ,

$$\sum_{r_l \in R \setminus R_i} r_l(\overline{V_{R \setminus R_i}}) \geq \sum_{r_l \in R \setminus R_i} r_l(\overline{V_R})$$

In this case, the tax paid by  $A_i$  would become negative, since its presence actually helps the optimization algorithm find a better solution. One solution is to replace the solution  $\overline{V_{R \setminus R_i}}$  by the

better one  $\overline{V}_R$ , but this would eliminate the incentive for  $A_i$  to find a better solution. Another solution is to make the optimization mechanism pay this negative tax.

- some  $A_i$  could obtain a better utility than in  $\overline{V}_R$  by choosing the best compatible values with  $\overline{V}_{R \setminus R_i}$ . In this case, for all variables of agents other than  $A_i$ , we set  $\overline{V}_R \leftarrow \overline{V}_{R \setminus R_i}$ , and set variables belonging to  $A_i$  to the best compatible values with that assignment. If there are several such agents, we can select the first one by coin flip and repeat the process until no further such agents remain. This procedure does not affect incentive-compatibility since it only improves  $\overline{V}_R$  and all agents influence this solution anyway.

## 5. An Incentive-compatible, Budget-balanced Social Choice Mechanism

The example shows that the truth-inducing property of the tax can come at a high cost: while the overall utility obtained by the system is 30, the agents must pay 26 units of tax to an outside entity. Thus, their effective total utility is only 4 units!

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.

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.

More specifically, consider the following mechanism:

### Mechanism 1 .

1. Ask each agent  $A_i \in \mathcal{A}, i = 1..k$  to state its relations.
2. Choose an excluded group  $E$  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 group 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. In most cases, it will be best to let it consist of only a single agent, but it can be used to create different expected returns for the participants of the mechanism.

Mechanism 1 is obviously budget-balanced, since all taxes are paid to the agents in the excluded group. For incentive-compatibility and individual rationality, consider an agent  $A_i$ . If  $A_i$  is not in  $E$ , then the mechanism looks just like a VCG tax mechanism with the subset of agents that excludes  $E$ . This is known to be incentive-compatible and individually rational. If  $A_i \in E$ , then the mechanism is incentive-compatible as the result does not depend on  $A_i$ 's actions at all. It is also individually rational, as  $A_i$  receives some of the taxes and thus is at least as well off as if it had not participated at all. Furthermore, all these properties hold *ex-post*, i.e. even if  $A_i$  knew all the relations declared by other agents.

When an optimal solution can not be computed, a suboptimal mechanism can be used with the same considerations as outlined above.

Now consider using Mechanism 1 on the gas utilities example given earlier. Assume that it forms coalitions of single agents and that each of these coalitions is chosen to be excluded with the same probability 1/3. The mechanism chooses between the following three solutions:

Solution	Assignments	Utilities	Payments
$S_1$	$x_1 = 2$	$U(A_1) = 0$	$pay(A_1) = -7$
	$x_2 = 2$	$U(A_2) = 11$	$pay(A_2) = 7$
	$x_3 = 3$	$U(A_3) = 2$	$pay(A_3) = 0$
$S_2$	$x_1 = 1$	$U(A_1) = 13$	$pay(A_1) = 7$
	$x_2 = 1$	$U(A_2) = 0$	$pay(A_2) = -7$
	$x_3 = 3$	$U(A_3) = 2$	$pay(A_3) = 0$
$S_3$	$x_1 = 1$	$U(A_1) = 13$	$pay(A_1) = 8$
	$x_2 = 1$	$U(A_2) = 3$	$pay(A_2) = 0$
	$x_3 = 2$	$U(A_3) = 0$	$pay(A_3) = -8$

Considering that each solution is chosen with probability 1/3, the agents have the following expected returns:

Agent	$A_1$	$A_2$	$A_3$	total
E[Utility]	8.66	4.66	1.33	14.66
E[Payment]	2.66	0	-2.66	0
E[Net Utility]	6	4.66	4	14.66

Thus, in spite of the fact that the allocation produces a suboptimal utility of only 14.66, there are no losses through taxes, and overall the agents are much better off than had they used a VCG scheme where their net utility was only 4!

It is a well-known fact that VCG tax schemes are vulnerable to *collusion*: groups 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. Obviously, the first-price tax is not incentive-compatible, as agents have an interest to claim a lower utility for the optimal solution. However, an agent that overclaims its utilities 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, similar to the Ausubel auction, that allow first-price schemes to become incentive-compatible in certain cases.

## 6. 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:

$x =$	1	2	3
$r_1$	$a$	0	0
$r_2$	0	$b$	0
$r_3$	0	0	$c$

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

Solution	$S_1$	$S_2$	$S_3$
$x =$	3	3	2

which gives us for the expected taxes and utilities:

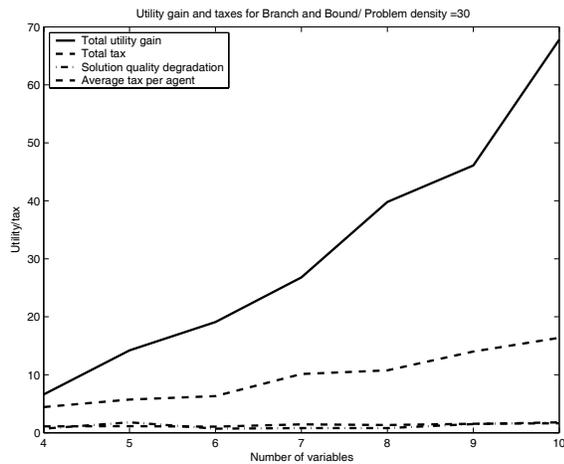
$A_i$	$E[tax]$	$pr(x = i)$	$E[u(A_i)]$
$A_1$	1/3 (-b)	0	b/3
$A_2$	0	1/3	b/3
$A_3$	1/3 b	2/3	2c/3 - b/3

In comparison, in 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 arises when leaving one agent out of the optimization would give only a slightly better result for the remaining agents, but a significantly worse result for the agent that was excluded.

To determine the average performance of the method, we have generated random coloring problems where the optimization consists of assigning variables with colors that are as far apart as possible. Figure 1 shows the following quantities:

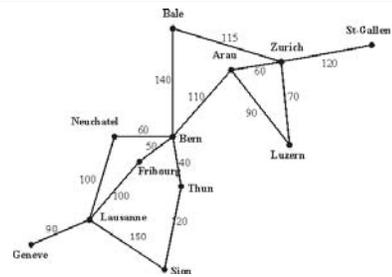


**Figure 1.** Experimental results for randomly generated coloring problems.

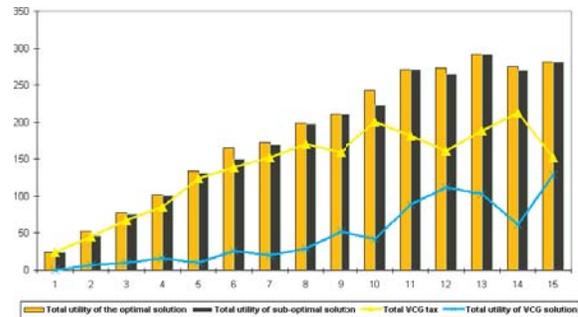
- the solid black curve shows the total gain in utility in an optimized solution versus a randomly chosen assignment.
- the dashed curve shows the total VCG tax that would have to be paid by the agents.
- the other two curves (almost indistinguishable) show the degradation of the total utility when one agent is excluded from the optimization, and the average VCG tax to be paid per agent.

The experiment shows that the VCG tax wipes out a large part of the utility gain that the agents derive from cooperation. 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. Meanwhile, 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 ex-



**Figure 2.** Network used for resource allocation experiment.



**Figure 3.** Utilities achieved by different mechanisms in the experiment.

ecuting the task at all. Constraints specify that no two tasks be assigned to paths that share an arc.

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. Even as the number of tasks increases, the utility of the suboptimal solution remains close to that of the optimal solution; the difference is much less than the amount of tax that would be wasted in a VCG mechanism.

## 7. Related Work

The dAGVA mechanism ([1]) makes each agent pay the VCGtax minus the expected value of that tax. It is optimal, incentive-compatible, 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.

Some earlier approaches have attempted to achieve budget balance by sacrificing optimality. For double

auctions, a special case of social choice, the *trade reduction rule* ([8]) leaves one possible exchange undone to achieve budget balance or possibly surplus. However, it does not apply to general social choice, and even on auctions tends to produce lower overall utility than the scheme presented here.

Parkes et al. ([11]) 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.

Another 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 ([12]) 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 small set of types rather than continuous valuations.

## 8. Conclusions

Intelligent agent technology allows replacing simple protocols for social choice with sophisticated optimization and coordination mechanisms that promise much improved utility to all participants. However, so far most approaches for multi-agent optimization and coordination have neglected the presence of self-interest.

The best existing solutions for taking self-interest into account 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. Since economists have shown many impossibility results (for example, [6, 9]) that prove that it is impossible to combine incentive-compatibility, optimality and budget balance, it is unlikely that this can be overcome in general. Furthermore, auctions do not apply to many instances of social choice because they assume free disposal of the items or rights received.

We have presented a mechanism based on the idea of choosing a solution that is optimal for all but a group of excluded agents. In this way, most tax and auction schemes can be made budget-balanced by returning the surplus to the excluded agents. In contrast to earlier proposals, the mechanism we presented is generally applicable to any social choice problem with a quasilinear utility function, and does not require any a-priori knowledge of agent's preferences.

The numerous impossibility results in game theory have long kept researchers from investigating the many possibilities for incentive-compatible social choice in restricted yet practically useful settings. Our results show that intelligent agent technology enables social choice mechanisms that are both incentive-compatible and provide significantly better utility to the involved parties than the simple protocols in use now. They allow new forms of computer-mediated collective decision-making that can make society more democratic.

## 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] C. d'Aspremont and L.-A. Gerard-Varet: "Incentives and Incomplete Information," *Journal of Public Economics* **11**, pp. 25-45, 1979
- [2] E.H. Clarke: "Multipart pricing of public goods," *Public Choice* **11**, pp. 17-33, 1971
- [3] E. Ephrati, 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.
- [4] B. Faltings, S. Macho-Gonzalez: "Open Constraint Optimization," *Proceedings of CP-2003*, LNCS 2833, pp. 303-317, 2003
- [5] B. Faltings, M. Yokoo (eds.): Special Issue on Distributed Constraint Satisfaction, *Artificial Intelligence*, forthcoming, 2005
- [6] J. Green, J.J. Laffont: "Incentives in public decision making," *Studies in Public Economics* **1**, North-Holland, 1979
- [7] T. Groves: "Incentives in Teams," *Econometrica* **41**, pp. 617-663, 1973
- [8] R. McAfee: "A dominant strategy double auction," *Journal of Economic Theory* **56**, pp. 434-450, 1992
- [9] R.B. Myerson, M.A. Satterthwaite: "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* **29**, pp. 265-281, 1983
- [10] N. Nisan and A. Ronen: "Computationally feasible VCG Mechanisms," *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 2000
- [11] D.C. Parkes, J. Kalagnanam, M. Eso: "Achieving Budget-Balance with Vickrey-based Payment Schemes in Exchanges," *IJCAI-01*, pp. 1161-1168, 2001
- [12] T. Sandholm: "Automated mechanism design: A New Application Area for Search Algorithms," *Proceedings of CP-2003*, LNCS 2833, pp. 19-36, 2003
- [13] W. Vickrey: "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance* **16**, pp. 8-37, 1961