Strategy-proof implementation and second-best efficient mechanisms: beyond Groves mechanisms.

Research project proposal for academic year 17-18
Timos Athanasiou
eathanasiou@nes.ru
http://www.nes.ru/en/people/catalog/a/efthymios-athanasiou

April 26, 2017

1 The problem in broad strokes

The Friday morning flight from Rome to Paris has been overbooked. Three individuals lay claim to the last two available seats. Two of them will be later boarding an overseas flight operated by a different airline. The other needs to attend to a pressing business engagement later in the day. They have all made their case to the representative of the airline and they are awaiting her decision. To whom should the last seats be assigned?

The individuals living in an old building in Paris have decided to install an elevator. Each individual agrees this is a necessary improvement and once the elevator will be installed everyone will be able to use it. However they all attach a different value to this public project and they have all reported their willingness to contribute to the estate manager. How much each of them should be asked?

The head of the department of economics of some university needs to allocate some courses to the new members of the department. They are all asked to rank they available courses as a function of their teaching interests. Who should teach what course?

These examples capture the core features of the problems associated with this research project. A benevolent planner wishes to:

- assign m indivisible private goods to n claimants. Each individual attaches a different value to any of the available goods.
- provide a public good whose cost should be covered by the community. Each individual attaches a different value to the public good.

- assign a task to certain members of a community. Performing a task entails a different cost for any of the members of the community.

Our travelers to Paris have incentives to misreport their valuation of the good. The representative of the airline may legitimately wonder whether, for instance, all the individuals are telling the truth about the predicament they will be in if they miss their flight. Similarly the people who need to contribute to the installation of the elevator may lie about the value they attach to the project because they will be eventually asked to contribute less but they would still be able to use the elevator. This also applies to the department members who may strategically provide a false ranking to the head of the department in order to avoid a certain course.

In any of these cases, a mechanism that solves the problem needs to overcome this difficulty. A *strategy-proof* mechanism induces, for each individual, if asked, a dominant strategy to tell the truth.

Strategy-proofness constitutes a particularly robust form of implementation. Indeed, one can reasonably expect that a rational agent who has a dominant strategy will actually play it regardless of what the other agents are claiming. Unlike the equilibrium strategies in Nash related equilibrium concepts, for each agent, telling the truth is the optimal strategy without any presumption on how the other claimants are acting. Moreover, strategy-proof implementation is prior-free. It does not require knowledge of the distribution of individual valuations. Essentially, strategy-proof implementation is robust even if some of the participants have wrong beliefs about the distribution of other people's evaluations. In a similar fashion, the designer does not need have any belief about the participants' evaluations for the mechanism to work.

2 The problem in the literature

The three assignment problems mentioned above have received considerable attention in the literature. In general, most of the existing work on strategy-proof implementation focuses on assignment-optimal mechanisms. Roughly speaking, assuming transferable utility, assignment-optimality means that the assignment should maximize the sum of individual valuations (for example, for the first problem I have mentioned before, assignment-optimality simply means that the seats never remain unallocated and always go to the agents with highest evaluation).

The solution that has received most attention is represented by the family of Vickrey-Clarke-Groves (VCG) mechanisms (see Vickrey [**?**]; Clarke [**?**] ; Groves [**?**]). Such family of mechanisms is characterized by strategy-proofness and assignment optimality (see Holmstrom [**?**]). In general, VCG mechanisms fail to balance the budget (see Green and Laffont

[?]). As a consequence, complying with strategy-proofness and assignment optimality entails a welfare loss. The waste takes the form of a budget deficit.

Intuitively, relaxing the constraint of assignment optimality helps to reduce the amount of money wasted (for example, leaving the seats on flight from Rome to Paris unallocated and performing no transfers is a strategy-proof and feasible mechanism that wastes no money). Of course, dropping assignment optimality may also entails a welfare loss.

In the presence of informational constraints, a loss of welfare can manifest in two different ways: either in the form of a deficit, as it happens with the VCG mechanisms, or directly in the form of a suboptimal assignment. A key question is wether looking beyond VCG mechanisms, one might find mechanisms that do not commit to assignment optimality but are welfare superior to them.

Very few papers have tried to tackle this question so far. For example de Clippel et al., [?], in a model where m identical goods have to be assigned to n agents, show that the welfare superior mechanism among those satisfying strategy-proofness and other desirable properties involves destroying the good. The conclusion is that abandoning assignment optimality may determine a welfare improvement. In a simpler framework (one indivisible good to be assigned to n agents), Moulin [?] considers the set of mechanism satisfying feasibility, strategy-proof and other desirable properties. He studies the problem of minimizing the worst relative surplus loss within this set of mechanisms. He allows for non-VCG mechanisms but the solution he describes is a VCG mechanism.

3 The key question

A common feature of the papers just mentioned is that they base their conclusions on a specific metric that captures the welfare loss. Mechanisms are compared with each other on the basis of the sum of utilities the allocations they prescribe produce.

Consider an alternative approach. A mechanism Pareto dominates another mechanism when all individuals agree that the outcome under the first is better than the outcome under the second. A mechanism is *second-best efficient* if it is anonymous, feasible and strategy-proof and there does not exist another anonymous, feasible and strategy-proof mechanism that Pareto dominates it. So, the main question this project intend to pursue, for different economic domains, is:

What is the set of second-best efficient mechanisms?

There are two immediate advantages that this approach brings about. First, if a mechanism is Pareto dominated then it is certainly not a good solution. Second, knowing the shape of the Pareto frontier provides a truly solid foundation to the exercise of choosing a *fair* allocation.

Quite surprisingly, when it comes to strategy-proof implementation, very little is known about the shape of the Pareto frontier. Certain prominent examples of VCG mechanisms, like the Vickrey Mechanism, have been widely studied and used in real life applications (i.e., auctions) in spite of the fact that it was not clear wether there exists another mechanism, not necessarily belonging to the family of VCG mechanisms, that Pareto dominates it. The few papers that try to perform a welfare comparison of strategy-proof mechanisms still focus only on VCG mechanisms. For example Ohseto [**?**] and Guo and Conitzer [**?**] identify Pareto undominated Groves mechanisms. Their approach involves assignment optimal mechanisms. The former paper describes the Pareto frontier of the class of envy-free Groves mechanisms. The latter one offers, in a fairly general setting, a set of necessary and sufficient conditions for Pareto-optimality within the class of Groves mechanisms.

4 Examples of particular projects

1 - Private Goods

A number m private goods need to be allocated among n agents, with $m < n$. In Athanasiou, Dey and Valletta [**?**] we study Groves mechanisms in a context of private goods that accounts for the effect of an externality. In economies comprising two agents, we single out the only feasible and anonymous Groves mechanism that is not Pareto dominated by another strategy-proof, feasible and anonymous mechanism. However, the two contributions on which I intend to build more directly are and Sprumont [**?**].

Sprumont [**?**], in a framework with one indivisible good and n agents proposes a parametrized family of mechanisms, the *maxmed* family of mechanisms, including the canonical Vickrey mechanism. The members of such family are the only ones satisfying strategy-proofness, anonymity, no-envy[1] and voluntary participation that are not Pareto dominated by another strategy-proof, anonymous and envy-free mechanism.

Athanasiou [**?**] identifies three conditions that are necessary and, together with voluntary participation, sufficient for a mechanism to be second-best efficient. For domains comprising two individuals the paper provides an explicit characterization of the family of second-best efficient mechanisms. Interestingly, this family coincides, for $n = 2$, with the *maxmed* family. The problem with Athanasiou's characterization is that it relies on three abstract conditions which cannot be easily related with a particular class of mechanisms. On the other hand, the *maxmed* mechanisms are not second-best efficient (i.e.,

[1] No-envy is a fairness requirement largely studied in the literature: at an envy-free allocation no agents prefers the bundle assigned to some other agent to the bundle she has been assigned.

they are generally not Pareto-optimal within the larger class of feasible, strategy-proof and anonymous mechanisms). For example, Cavallo [**?**], Bailey [**?**] and Porter, Shoham and Tennenholtz [**?**] offer examples of mechanisms that Pareto dominate the Vickrey mechanism when $n \geq 3$.

One objective that may be pursued is the characterization of the class of second-best mechanisms when only one object has to be allocated. In doing so one may build on the methodology used both by Athanasiou [**?**] and Sprumont [**?**]. Both papers rely on a result that can be found in Nisan [**?**]. He shows that all strategy-proof mechanisms for this problem have a precise structure.

A more ambitious objective, is to try to obtain the same characterization for the more general model comprising m identical goods and n individuals, the so called rationing problem (Moulin [**?**]).

2 - Public Goods

The simplest pure public good provision model one can think of takes the following form. There is a set of n agents. There is single indivisible public good project about which a yes or no implementation decision has to be made. If the public good is built everyone can use it and agents have a different valuation of the project. Pecuniary transfers are possible. In such a simple framework, one of the most studied solutions is the Pivotal Mechanism proposed by Clarke [**?**]. This mechanism is a particular Groves mechanism, so it is strategy-proof and assignment optimal (the project is realized only if the sum of the willingness to pay of the N agents is greater than its cost). Moreover it is feasible, the sum of the transfers covers for the cost of the project, actually the sum of transfers is typically greater than the cost so the mechanism generates a waste. In an early contribution Moulin [**?**] studies the welfare properties of the Pivotal Mechanism within the class of Groves mechanisms. It turns out that, if there are only two agents, then the Pivotal Mechanism Pareto dominates any other Groves Mechanism. This conclusion however does not carry over to three agents and more. In a more recent contribution, Guo and Conitzer [**?**] show that, within the class of pay-only Groves mechanisms (the class of Groves mechanisms that only allows for negative tranfers), the pivotal mechanism is the only Pareto-undominated mechanism. However, more in general, even if one just considers the class of Groves mechanisms it is not clear at all what is the shape of the Pareto frontier. Needless to say this issue becomes even more obscure if one considers the whole class of strategy-proof, feasible and anonymous mechanisms.

Some papers have shown that this problem becomes more manageable if one considers public good provision models where it is possible to exclude agents from using the public goods. A prominent example is the Serial Cost Sharing mechanism (Moulin and Shenker

[**?**]). A series of papers (Moulin [**?**]; Olszewski [**?**] and Maniquet and Sprumont [**?**] among others) discuss the properties of such a mechanism. As far as this project is concerned it is interesting to note that these papers insist on the fact that exclusion is not detrimental and, under certain circumstances, necessary for identifying the class of second-best efficient mechanisms.

A fundamental objective in this context is to characterizing the class of second-best mechanisms when a single *pure* public good project has to be implemented.

A further, certainly more challenging, extension is to study the Pareto frontier in a more realistic pure public good model where the project has a variable (discrete) size. In this case it is particularly interesting to understand how the trade-off between feasibility and assignment optimality affects the choice of the size of the project.

3 - Tasks

A finite set of indivisible tasks is to be allocated among a n agents. All tasks must be allocated. An agent can be assigned either no task, a single task, or more than one task. Each task is assigned to only one agent. Performing one or more tasks entails a cost for each agent. Each agent has a cost function which measures her capacity to perform such tasks. Pecuniary transfers are allowed and are meant to compensate agents for the tasks she needs to perform. Preferences are quasi-linear and, at a given allocation, the final utility of each agent is equal to the cost associated with the task she has been assigned (a negative number) plus the transfer she receives (or pays). This simple model can describe many real life situations. One can simply think of the job-task assignment problem mentioned earlier or, even more interestingly, the problem of choosing the locations of noxious facilities (prisons, hazardous materials facilities such as chemical process facilities, waste disposal sites, nuclear facilities, etc.), also called the "not-in-my-backyard" (NIMBY) problem. Although all the members of the society have the same rights and responsibilities over the public bads, only the hosting localities will experience the costs. Hence, localities have an incentive to misrepresent their preferences to manipulate the allocation rule in their favor.

Most of existing literature concerning this model has focused on the implementation of the assignment optimal allocation. Assignment optimality here means that the tasks must be assigned in such a way that the sum of the cost faced by agents is minimized. Some papers then focus either on implementing the efficient outcome in Nash equilibrium (like, for example, Lurent-Lucchetti and Leroux [**?**]) while others focus on assignment optimal and budget balanced mechanisms (like, for example, Sakay [**?**]). The former approach requires all agents to know the cost function of the other agents present in the economy, which may be too strong of an informational requirement for certain practical applications, while the second approach might be too narrow since rules out strategy-proof implementation. As far this kind of implementation is concerned, the existing papers never look beyond assignment optimal mechanisms (like, for example, [**?**])moreover, in certain cases the emphasis is put mechanisms that are financed by the social planner i.e., they run a deficit (like, for example, Porter, Shoham and Tennenholtz, [**?**]).

Again, the idea of identifying the shape of the Pareto frontier within the class of strategy-proof, anonymous and feasible mechanism has been completely neglected. This issue seems particularly sensible here since often the assignment optimal allocation cannot be implemented because of social or political reason (for example, a certain community fiercely rejects to host a noxious facility even if that would be the most efficient thing to do). In such cases, knowing the set of second-best mechanisms simply allows to choose an alternative that is politically practicable and that determines a distribution of welfare sufficiently close to the one that would been have obtained under the assignment optimal allocation.

References

[1] Athanasiou, E., 2013. A Solomonic Solution to the Problem of Assigning a Private Indivisible Good. Games and Economic Behavior 82, 369-387.

[2] Athanasiou, E., Dey, S., Valletta, G., 2012. Groves mechanisms and communication externalities. GSBE Research Memorandum 12/016.

[3] Bailey, M.J., 1997. The demand revealing process: to distribute the surplus. Public Choice 91, 107-126.

[4] Cavallo, R., 2006. Optimal decision-making with minimal waste: Strategy-proof redistribution of VCG payments. In: AAMAS 2006. Hakodate, Japan.

[5] Clarke, E.H., 1977. Multipart pricing of public goods. Public Choice 11,17-33.

[6] de Clippel, G., Naroditskiy, V., Polukarov, P., Greenwald, A., Jennings, N., 2014. Destroy to Save. Games and Economic Behavior 88, 392-404.

[7] Green, J., Laffont, J.-J., 1979. Incentives in Public Decision Making. North-Holland, Amsterdam.

[8] Groves, T., 1973. Incentives in teams. Econometrica 41, 617-631.

[9] Guo, M., Conitzer, V., 2008. Undominated VCG redistribution mechanisms. In: Proceedings of the 7th International Conference of Autonomous Agents and Multi-Agent Systems. Estoril, Portugal, pp. 1039-1046.

[10] Holmstrom, B., 1979. Groves' scheme on restricted domains. Econometrica 47, 1137-1144.

[11] Lurent-Lucchetti, J., Leroux , J., 2011. Choosing and sharing. Games and Economic Behaviour 73, 296Ð300.

[12] Maniquet, F., Sprumont, Y., 2010. Sharing the cost of a public good: An incentive-constrained axiomatic approach. Games and Economic Behavior 68, 275-302.

[13] Moulin, H., 1986. Characterizations of the pivotal mechanism, Journal of Public Economics 31, 53-78.

[14] Moulin, H., 1994. Serial cost-sharing of excludable public goods. Review of Economic Studies 61, 305-325.

[15] Moulin, H., 2009. Almost budget-valanced VCG mechanisms to assign multiple objects. Journal of Economic Theory 144, 96-119.

[16] Moulin, H., 2010. Auctioning or assigning an object: some remarkable VCG mechanisms. Social Choice and Welfare 34, 193-216.

[17] Moulin, H., Shenker, S., 1992. Serial cost-sharing. Econometrica 60, 1009-1037.

[18] Nisan, N., 2007. Introduction to mechanism design (for computer scientists), in: Nisan, N., Rougarden, T., Tardos E., V. Vazirani, V., (Eds.), Algorithmic Game Theory, Cambridge University Press, New York, USA, pp. 209-241.

[19] Ohseto, S., 2006. Characterizations of strategy-proof and fair mechanisms for allocating indivisible goods. Economic Theory 29, 111-121.

[20] Olszewski, W., 2004. Coalition strategy-proof mechanisms for provision of excludable public goods. Games and Economic Behavior 46, 88-114.

[21] Porter, R., Shoham, Y., Tennenholtz, M., 2004. Fair imposition. Journal of Economic Theory 118, 209-228.

[22] Sakai, T., 2012. Fair waste pricing: An axiomatic analysis to the NIMBY problem. Economic Theory 50, 499Ð521.

[23] Sprumont,Y., 2013. Constrained-optimal strategy-proof assignment: Beyond the Groves mechanisms. Journal of Economic Theory 148,1102-1121.

[24] Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. Journal of Finance 16, 8-37.

[25] Yengin, D., 2013. Identical preferences lower bound for allocation on heterogeneous tasks and NIMBY problems. Journal of Public Economic Theory, 15, 580-601.