Incentive Compatible Allocation Without Money

RUGGIERO CAVALLO
Yahoo Research

Mechanism design typically involves the use of monetary payments to incentivize truthful reporting of private information. But what if payments are impossible? This letter summarizes recent work on how lotteries and the threat of non-allocation can be leveraged to similar ends in single-item allocation settings with a secondary good that can serve a limited numéraire role.

Categories and Subject Descriptors: J.4 [**Social and behavioral Sciences**]: Economics

General Terms: Economics, Theory, Design

Additional Key Words and Phrases: Mechanism Design, Social Welfare, Lotteries

1. INTRODUCTION

Let's start with a couple stories.

Story 1: Imagine you work at a large tech firm and you're tasked with designing an internal allocation policy for the company's vast internal computational resources. Imagine, further, that these resources take the following form: there is a virtually unlimited amount of "standard issue" computational capacity available, and then there is an elite high-performance computing cluster that is to be made available to only one job at a time. All potential users (researchers, product groups, etc.) have approximately the same value for the standard resource, but values vary widely (and privately) for the elite resource. And, users are selfish. Your task is to determine which user gets access to the high-performance cluster at any given time.

Story 2: Imagine you're at this year's iteration of the Economics and Computation conference, and after one of the sessions, on your way out of the room you notice that one of your colleagues has lost his wedding ring, which you pick up. You want to return it to its rightful owner, but you don't know who that is. And, wary that this crowd may take their standard models of behavior a little too much to heart, you fear that if you were to simply make an announcement asking whose it is, many "self-interested agents" would come forward as claimants.

These two very different scenarios have some important common features: an allocation decision must be made; the decision-maker would like to choose the allocation that maximizes "welfare" (if that can be suitably defined), but faces the challenge of privately held values; and finally, common allocation approaches like auctions seem untenable, since monetary payments would likely be deemed inappropriate. I doubt there are any perfect solutions to this type of problem, but in this letter I'll describe some recent work [Cavallo 2014] that at least suggests possible approaches to allocation problems where social welfare is the goal and payments are disallowed or highly restricted, either for normative reasons or due

Author's address: 111 W 40 St., 9 FL, New York, NY 10018. Email: `cavallo@yahoo-inc.com`

to constraints on agents' *capacity* to pay.[1]

1.1 The roles of money in mechanism design

If mechanism design is the study of how to *engineer* outcomes in settings with selfish actors and private information, the usual engineering *lever* is money. E.g., scenarios with competing interests are transformed into ones with common interests via specific payments. This is one important role of money.

But more fundamentally, money often plays an important role as a measure of utility, i.e., as numéraire. If everyone can assign a dollar-equivalent value for, say, obtaining a particular good, then there is a metric for comparison; and if dollars can be transferred from one person to another and utility functions are quasilinear, seeking Pareto efficiency entails a very specific policy: allocate the good to the person with highest value for it. But what if there's no money with which to transfer utility? While in some settings we can stipulate a sensible optimization criterion that still differentiates the space of possible outcomes,[2] when allocating a *single* item *any* allocation is Pareto efficient. There may well be real differences in "value" across the population (e.g., a wedding ring is uniquely important to its true owner), but we lack a way of measuring them. If we want to discriminate among different possible allocations of a single good, the presence of an additional feature in the outcome description—be it money or something else—seems unavoidable.

2. INCENTIVE COMPATIBLE ALLOCATION

2.1 Measuring utility

In [Cavallo 2014], I consider a setting that maps directly to *Story 1*, in which there is a singular high-value good (which we'll call *type-A*) and an unlimited number of identical lower-value goods (which we'll call *type-B*). Agents have need for only one good, so allocating a *type-B* good to an agent who receives the *type-A* good yields no additional value.[3] An agent i's value v_i for the *type-A* good is defined in terms of how much he prefers it over the lower-value *type-B* good. In other words, the *type-B* good is used as numéraire, and i's utility is: v_i if he receives the *type-A* good, 1 if he instead receives a *type-B* good, and 0 otherwise. Given this set-up, utilitarian social welfare—denominated by *type-B* good value—can be measured as v_i (for the agent i that receives the *type-A* good, if anyone does) plus m, where m is the number of agents allocated a *type-B* good. This is what we seek to maximize.

To allocate the *type-A* good optimally, we need agents to reveal their values. There is a trivial non-monetary mechanism wherein misreporting types can never be beneficial: randomly choose an agent i, give him the *type-A* good, and give all others a *type-B* good. The truthfulness of this mechanism is wasted—reports

[1]On the normative side, Alvin Roth [2007] describes three types of concerns about monetary transactions: the *objectification* of goods or services, the *exploitation* of involved parties, and contribution to a *slippery slope* of over-monetization. The approach I suggest may mitigate some of these issues, but will not be completely immune to related critiques.
[2]Examples include: optimizing matches in an exchange network [Ashlagi and Roth 2013], choosing where to build a public facility [Procaccia and Tennenholtz 2013], fairly splitting up divisible items [Cole et al. 2013], and scheduling jobs to minimize makespan [Koutsoupias 2011].
[3]In *Story 1* terms, any given user only needs to run a single job.

are ignored. The question is: can we do anything *else* while maintaining incentive compatibility? Can we do anything better? The short answers are: yes, we can do something else; and in some cases that "something" is better, but in many cases random allocation cannot be beat.

2.2 Lottery-based incentive mechanisms

In the standard single-item setting with money, the Vickrey auction achieves an optimal allocation by, in essence, making a proposition to each agent: you can have the item if you pay $\$x$. The trick is to define the x's in a way such that only the highest-value agent will want to accept. In the non-monetary setting described above, something similar is possible: we can specify a proposition for each agent of the form "you can have either a *type-B* good or a lottery for the *type-A* good", where the lottery odds are defined such that only the agent with highest value for the *type-A* good will want to take it (now, besides serving as numéraire, the *type-B* good is playing the incentive-shaping role of money). Specifically, we give the agent reporting the highest value a $1/x$ chance of receiving the *type-A* good, where x is the *second*-highest reported value. This mechanism is dominant strategy incentive compatible (DSIC). It may leave the *type-A* good unallocated with high probability; however, somewhat surprisingly, ideas from work on strategyproof revenue redistribution [Bailey 1997; Cavallo 2006] can be adapted here to yield variant lottery mechanisms that go a large way towards remedying this.

Do these mechanisms help overall? That is, given a prior over agents' values for the *type-A* good, do they yield greater expected utilitarian social welfare than random allocation? The answer is mixed: building off results of Hartline and Roughgarden [2008] derived in a very different context, I show that for any i.i.d. value distribution that has monotonically increasing hazard rate, *no DSIC mechanism can beat random allocation.* However, in other cases these mechanisms can in fact do far better. This connects back to *Story 2* from the introduction: it is easy to imagine scenarios where it is likely that one agent (e.g., the true owner of a lost wedding ring) will have a far higher value than all others for the good to be allocated—perhaps because the value distribution is heavy-tailed, or values are not i.i.d.—and it is exactly in such cases that the proposed lottery-based DSIC mechanisms yield major gains over random allocation.[4]

2.3 Relation to settings with restricted monetary payments

In the work described above, the plentiful *type-B* good is used to measure utility and create incentives for truthful value reporting, two key roles traditionally played by money. So, one might ask, how are *type-B* goods any different from money? Are we in fact just doing a kind of money-based mechanism design by another name?

No, in fact what we're doing is more general: the *type-B* good could be dollar bills or anything else. But something useful may be gleaned by the connection this question suggests. Money has properties that make it flexible in ways that the

[4]In the lost wedding ring story, there was no *type-B* good described, but one can conjure many possibilities. For instance, each conference attendee may be given a banquet ticket; then the "bidding" can be defined in these terms: what are the minimum odds for receiving the ring you'd be willing to trade your ticket for?

type-B good of scenarios like *Story 1* will not be. If the *type-B* good were money, to fit snugly into the model we would have to impose restrictions that are highly unnatural in the case of an actual currency. For one, we would have to assume that agents can *receive* (very limited) payments but are unable to make any. The negative results are neutered when these restrictions do not apply.

However, the positive dimension of [Cavallo 2014] continues to apply, and may be a useful starting point for money-based mechanism design in a setting with highly budget-constrained agents, where the center has a modest budget he can put towards inducing truthful value reporting. It would not surprise me if further progress on the types of problems discussed here closely tracks breakthroughs in efficient mechanism design with budget constraints.

3. CONCLUSION

On the one hand, the negative side of this analysis may be the most remarkable: in a setting with a good that can stand in for money as numéraire, we can't beat random allocation when values are i.i.d. and the distribution has monotonically increasing hazard rate. But on the other hand, it is interesting that we can *ever* do better than random allocation without money. And moreover, it is perhaps in settings that defy the monetary approach that we should be most skeptical about the relevance of any specific quantitative model of utility, especially one characterized by pure self-interest. Thus it would be prudent to take an extra grain of salt with negative results in this domain. Maybe people just don't go around claiming other people's rings. I'd be happy if that's the least far-fetched premise considered in this letter.

REFERENCES

Ashlagi, I. and Roth, A. E. 2013. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*. Forthcoming.

Bailey, M. J. 1997. The demand revealing process: To distribute the surplus. *Public Choice 91*, 107–126.

Cavallo, R. 2006. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06)*. 882–889.

Cavallo, R. 2014. Incentive compatible two-tiered resource allocation without money. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'14)*. 1313–1320.

Cole, R., Gkatzelis, V., and Goel, G. 2013. Mechanism design for fair division. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC'13)*. 251–268.

Hartline, J. D. and Roughgarden, T. 2008. Optimal mechanism design and money burning. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing (STOC'08)*. 75–84.

Koutsoupias, E. 2011. Scheduling without payments. In *Proceedings of the 4th International Symposium on Algorithmic Game Theory (SAGT'11)*. 143–153.

Procaccia, A. D. and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation 1,* 4, 18:1–18:26.

Roth, A. E. 2007. Repugnance as a constraint on markets. *Journal of Economic Perspectives 21,* 3, 37–58.