Incentive Compatible Two-Tiered Resource Allocation Without Money

Ruggiero Cavallo*
Yahoo Research
111 West 40 St.
New York, NY, 10018
cavallo@yahoo-inc.com

ABSTRACT

We consider a resource allocation problem with two types of goods: a plentiful good that all agents have approximately the same value for, and a scarce good that agents value differently (imagine, e.g., job requests on an ordinary computing cluster versus a restricted high-performance cluster). A social planner seeks to allocate the scarce resource to the agent who values it most. We depart from the usual mechanism design approach by assuming monetary payments are infeasible, and instead use lotteries and the threat of non-allocation to elicit truthful value reporting. Adapting ideas developed in the context of *revenue redistribution*, we examine whether there exist allocation rules yielding expected welfare that—in ex post equilibrium—exceeds that of a baseline that randomly assigns the scarce resource, and find that for i.i.d. values the answer is yes only if the value distribution is heavy-tailed. For a variant of the problem where there is a residual claimant for the plentiful good, we identify a mechanism that obtains welfare converging to that of perfectly efficient allocation as the population size grows.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory

Keywords

Mechanism design; welfare; resource allocation; budgets

1. INTRODUCTION

Monetary payments are often employed to achieve efficient outcomes in group decision settings with individuals that are selfish and hold private information. In the standard setup, agents are assumed to have quasilinear utility, and transfer payments are defined that transform an inherent game of competing interests into one where each agent's utility is maximized when utilitarian social welfare is maximized. This approach has been extremely fruitful, perhaps most manifestly in *auctions*—of highly valuable singular goods such as works of art, large-quantity small-value goods like online advertising impressions, and practically everything in between. However, in certain cases this solution paradigm does not fit the nature of the setting; i.e., *money* is an inapplicable medium for achieving the required incentives, either because agent utilities are idiosyncratically non-quasilinear or, more fundamentally, payments are deemed unacceptable, e.g., for ethical or political reasons. We will consider the problem of optimally allocating a good amongst a group of agents when monetary payments are off the table.

What does it mean to make a decision "optimally" or "efficiently"? When utility is quasilinear and transferable via payments, agents ascribe a dollar-equivalent value to each outcome, and Pareto efficiency amounts to choosing an outcome that maximizes the sum of agent values. But when money cannot be transferred, Pareto efficiency requires only that no alternative to the chosen decision is preferred by some agent without being less preferred by another. In certain settings this remains a meaningful objective, such as when there are multiple heterogeneous goods to be allocated and each agent can consume a limited number of them. But in a *single*-item setting (or even with multiple goods if agents' values strictly increase as they acquire more goods), without transfers *all* allocations are Pareto efficient, and so the Pareto criterion becomes contentless as an objective. Yet, the absence of monetary payments as an available tool does not preclude real distinctions between different agents' values; it just complicates their quantification.

We will be motivated by settings where there is no money and no transfers can take place, but where there is nonetheless an alternate *numéraire* good that allows us to establish a quantitative measure of social welfare. Specifically, we will consider a resource allocation setting where there is a single instance of a high-value good which we'll call "type-$\mathcal{A}$" and an unlimited number of identical lower-value "type-$\mathcal{B}$" goods. Each agent has use for only a single good, whether of type-$\mathcal{A}$ or type-$\mathcal{B}$, so each agent prefers to obtain the type-$\mathcal{A}$ good, but would rather receive a type-$\mathcal{B}$ good than nothing. We will use type-$\mathcal{B}$ goods as the numéraire and seek to maximize utilitarian social welfare—aggregate value to the agents—in those terms, with each agent privately observing his value for the type-$\mathcal{A}$ good with respect to this numéraire.

Since there is no scarcity of type-$\mathcal{B}$ goods, our problem is at root about how to "efficiently" allocate a single good (the type-$\mathcal{A}$ good); the presence of type-$\mathcal{B}$ goods serves the dual

*This work was mostly completed while the author was employed at Microsoft Research, New York.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

purpose of providing an inter-agent welfare evaluation metric, and also a means of incentivizing truthful reporting of private values. To give a sense of our approach, intuitively one can imagine agents each being granted a type-$\mathcal{B}$ good as an initial endowment, and then competing for the opportunity to trade-in this good for a precisely-defined lottery for the type-$\mathcal{A}$ good. The presence of type-$\mathcal{B}$ goods will thus be an imperfect but useful substitute for money, allowing us to apply (some) theory from mechanism design with limited payments, but not, as we will see, providing nearly the flexibility and power of general monetary transfers.

After describing related work, in Section 1.2 we formally introduce the details of the model. In Section 2 we show that it is impossible to do better than random allocation when values are i.i.d. from a distribution with monotone increasing hazard rate, but we propose and evaluate two incentive compatible mechanisms that are superior in other cases. In Section 3 we consider a variant setting in which there is a residual claimant for the type-$\mathcal{B}$ goods, and propose a third mechanism that converges to perfect efficiency in that context as the population size grows. Due to space limitations, several proofs are omitted in this version of the paper.

1.1 Relation to previous work

In settings where monetary payments are possible and utility is quasilinear, efficient solutions to resource allocation problems with private values exist very generally [Clarke, 1971, Groves, 1973, Vickrey, 1961]. However, classic solutions such as the Vickrey–Clarke–Groves (VCG) mechanism often involve very large payments made to an entity outside the group. Relatively recent work on so-called "redistribution mechanisms" seeks to minimize such payments (see, e.g., [Bailey, 1997, Cavallo, 2006, Guo and Conitzer, 2007]), with very positive results (i.e., low "revenue" mechanisms) for allocation settings. However, no allocatively efficient mechanism exists that never imposes *any* aggregate transfer outside the group, even in single-item allocation [Hurwicz and Walker, 1990], and redistribution mechanisms still involve significant payments *between* agents.

On the other hand, in a paper rich with results, Hartline and Roughgarden [2008] consider the problem of maximizing aggregate agent welfare (value of the allocation minus payments) when no payments can be made *between the agents*, a scenario the authors describe as "money burning" since any required payments are viewed as pure loss. Their motivation and setup is related to the one we adopt here, with a few critical differences. Hartline and Roughgarden take as motivation settings where *service degradation* could form the basis of money burning "payments", and assume that such payments can be precisely defined and imposed independent of agents' private information. However, it seems much more natural to instead expect the rate at which degraded service diminishes experienced utility to inherently depend on the agents' private values.[1] In the model we adopt in this paper, each agent's disutility for the kind of "degradation" we impose (i.e., probabilistic non-allocation) flows naturally from his single *private* value for the high-value good to be allocated. We will call attention to other ways in which our setting differs from that of Hartline and Roughgarden in the next section.

[1] In the conclusion of their paper, Hartline and Roughgarden call attention to this issue, observing that when "burnt payments" correspond to degraded service quality, "the designer may not know each agent's relative disutility for such payments" and that this issue "motivates considering the more general setting where agents have a private value for burnt money in addition to their private value for service."

Other work explicitly addresses allocation with private values when monetary payments are *completely* disallowed, as in our setting. Hylland and Zeckhauser [1979] use a market mechanism to find Pareto optimal probabilistic allocations in an assignment problem with multiple heterogeneous items; however, the procedure is vulnerable to strategic manipulation except in the limit as the number of agents becomes large. Harrenstein et al. [2009] seek to map some of the attributes of the VCG mechanism to a setting where there is no numéraire good, proposing a "qualitative VCG" mechanism that obtains a Pareto efficient outcome in allocation and other settings. Guo et al. [2009] seek approximations to optimal welfare in a repeated allocation setting using an artificial currency but no real monetary payments. Dughmi and Ghosh [2010] consider assignment problems (with multiple goods and multiple agents), and achieve results without payments for the restricted case where each agent's value is either zero or some known value, but it is not known which. Procaccia and Tennenholtz [2009] design non-monetary mechanisms that approximate optimal welfare in a facility-location problem, where agents bear costs proportional to their distance from the location selected at which to build the facility. Guo and Conitzer [2010] derive clever mechanisms for two agents with normalized values; they do not posit a prior over values and instead seek mechanisms that approximate the optimal allocation value even in the worst-case. However, this and two related subsequent papers (Han et al. [2011] and Cole et al. [2013b]) impose a structure on agent valuations that renders the problem meaningless in settings where private valuation information regards only a single item, as in the current setting.

A very different line of research considers the so-called "King Solomon dilemma", wherein each agent *knows* whether or not she is the one with highest value for a good to be allocated [Glazer and Ma, 1989, Olszewski, 2003, Perry and Reny, 1999, Qin and Yang, 2009].[2] In this work efficient equilibrium outcomes without payments are achieved through a multi-step process in which the *threat* of payments and the common knowledge assumption play critical roles. Such mechanisms provide compelling solutions when payments are merely undesirable, but they require that payments could be made in principle and in fact *are* made if agents play off-equilibrium strategies.

Finally, the work of Cole et al. [2013a] is of a similar spirit to the approach we apply, though the objectives are quite different. As in our approach, the authors achieve incentive compatibility by probabilistically forgoing opportunities for welfare-improving allocation (discarding fractions of resources, in their case).[3] They are concerned with optimizing a fairness measure rather than welfare.

1.2 Preliminaries

There is a set of agents $I = \{1, \ldots, n\}$ and there are two types of goods: a single *type-$\mathcal{A}$* good, and an unlimited supply of *type-$\mathcal{B}$* goods. Each agent has a cardinal utility for

[2] See Moore [1992] for a very clear explication of early research in this vein.

[3] See also Feige and Tennenholtz [2010], where lotteries are used to obtain truthful reporting of utility functions in a single-agent setting.

each good, is capable of utilizing only a single good (of either type), and is risk neutral. We normalize each agent's utility for receiving a type-$\mathcal{B}$ good to 1, and then each i's utility for receiving the type-$\mathcal{A}$ good is denoted v_i. For each $i \in I$, v_i is private and in $[1, U]$, for some $U \in [1, \infty]$; v_i can be thought of as i's proportional value for the type-$\mathcal{A}$ good versus the type-$\mathcal{B}$ good.[4] Letting $u_i(q_i, p_i)$ denote i's expected utility for obtaining the type-$\mathcal{A}$ good with probability q_i, a type-$\mathcal{B}$ good (and no type-$\mathcal{A}$ good) with probability $p_i \in [0, 1-q_i]$, and nothing with probability $1-q_i-p_i$, note that $u_i(q_i, p_i) = q_i v_i + p_i$. Importantly, this parallels the standard quasilinear utility model for single-item allocation, with p_i here in the place of monetary payments.

We will use the term *welfare* to denote the sum of agent utilities—denominated by type-$\mathcal{B}$ good value, as described above—for a given allocation. For any given vector of values v and any $i \in I$, we use v_{-i} to denote the vector with v_i excluded and, for any $k \in \{1, \ldots, n\}$, $v^{(k)}$ to denote the k^{th}-largest element of v (with $v_{-i}^{(k)}$ defined analogously). We will frequently consider *expectations* over outcomes given a prior distribution over values for the type-$\mathcal{A}$ good; we will only consider i.i.d. settings, and denote the p.d.f. and c.d.f. of the common distribution as f and F, respectively.

From the perspective of a social planner (henceforth, "the center") seeking to maximize welfare, there are only n outcomes that could possibly be optimal, corresponding to the allocation of the single type-$\mathcal{A}$ good to one of the agents and type-$\mathcal{B}$ goods to all others; these are depicted in Table 1. Ideally the center would be able to ascertain which agent's value for the type-$\mathcal{A}$ good is highest and choose the corresponding outcome, which would maximize welfare.

	u_1	u_2	...	u_n
o_1	v_1	1	1	1
o_2	1	v_2	1	1
...	1	1	v_i	1
o_n	1	1	1	v_n

Table 1: Depiction of agent utilities over the space of possibly optimal outcomes (allocations), given unknown agent values $v_1, \ldots, v_n$ for the type-$\mathcal{A}$ good. o_i represents allocation of the type-$\mathcal{A}$ good to agent i and allocation of a type-$\mathcal{B}$ good to all other agents.

However, if the center's only power is to choose amongst these n specific outcomes—making no payments as is the supposition of this paper—he can do no better than to choose independent of the value information held privately by the agents. In other words, assuming agent values are i.i.d., the center can do no better than to pick amongst the n outcomes arbitrarily, as in the following "mechanism".

Mechanism 1. (Arbitrary Allocation Mechanism (AAM)) *An agent is chosen arbitrarily and allocated the type-$\mathcal{A}$ good. All others receive a type-$\mathcal{B}$ good.*

Remark 1. *Expected welfare under AAM equals* $\int_1^U f(x)x\,dx + n - 1$.

[4]Obtaining a type-$\mathcal{A}$ good *and* a type-$\mathcal{B}$ good still yields only utility v_i for i; obtaining more than one type-$\mathcal{B}$ good (and no type-$\mathcal{A}$ good) still yields utility only 1.

But this description of the problem elides the fact that the center has the power to *not* allocate goods. In other words, there are more alternatives besides the n enumerated in Table 1; leveraging them will be the focus of our approach.

2. INCENTIVE COMPATIBLE ALLOCATION MECHANISMS

The scenario pictured in Table 1 leaves out the fact that each depicted outcome o has a distinct corresponding outcome o' in which the agent chosen to receive the type-$\mathcal{A}$ good in o instead receives nothing. We can extend this reasoning also to the type-$\mathcal{B}$ goods, and so the full outcome space actually includes all modifications of the outcomes in Table 1 produced by eliminating a subset of the agents' allocations (yielding 0 value for those agents).

Considering the structure of the setting, if we allow for probabilistic selection of outcomes, the mechanism space can then be described as follows: each agent i reports a claim about his value v_i for the type-$\mathcal{A}$ good, and then for each i the center chooses a probability q_i with which i will obtain the type-$\mathcal{A}$ good and a probability p_i with which i will obtain a type-$\mathcal{B}$ good. Formally, a mechanism $M_{q,p}$ defines a vector of functions $(q_1, \ldots, q_n, p_1, \ldots, p_n)$, where, $\forall i \in I$, $q_i : [1, U]^n \to [0, 1]$ and $p_i : [1, U]^n \to [0, 1]$. A mechanism is *feasible* if it respects the laws of probability and agents' individual "capacity constraints":

Definition 1 (Feasible mechanism). *A feasible mechanism is a vector of functions $M_{q,p} = (q_1, \ldots, q_n, p_1, \ldots, p_n)$ with, $\forall i \in I$, $q_i : [1, U]^n \to [0, 1]$ and $p_i : [1, U]^n \to [0, 1]$, satisfying: $\forall v \in [1, U]^n$, $\sum_{i \in I} q_i(v) \leq 1$ and, $\forall i \in I$, $q_i(v) + p_i(v) \leq 1$.*[5]

We will be concerned with outcomes that occur in *ex post* equilibrium. By the revelation principle, it will be sufficient to consider mechanisms under which truthful value reporting is a dominant strategy.

Definition 2. (Dominant strategy incentive compatibility (DSIC)) *A mechanism $M_{q,p}$ is dominant strategy incentive compatible if and only if, $\forall i \in I$, $\forall v \in [1, U]^n$, $\forall \hat{v}_i \in [1, U]$, $q_i(v_i, v_{-i})v_i + p_i(v_i, v_{-i}) \geq q_i(\hat{v}_i, v_{-i})v_i + p_i(\hat{v}_i, v_{-i})$.*

To reiterate: type-$\mathcal{B}$ goods serve as numéraire and provide the basis for evaluation of welfare, a role typically played by money. Unlike money, though, type-$\mathcal{B}$ goods are indivisible and an agent can receive a limited amount. But indivisibility is not a practically important distinction in a setting that allows probabilistic allocation—receiving a type-$\mathcal{B}$ good with 50% probability is identical to receiving half of a type-$\mathcal{B}$ good, from an incentives and expected welfare perspective, given risk neutral agents. And the "capacity constraint" ($q_i + p_i \leq 1$) is analytically similar to a budget constraint.[6] The relatively close relationship between this setting and that of single-item allocation (of the type-$\mathcal{A}$

[5]More precisely, any mechanism meeting this definition can be feasibly implemented, as follows: the type-$\mathcal{A}$ good lottery is executed first, according to q, and then each agent i receives a type-$\mathcal{B}$ good with probability $\frac{p_i}{1-q_i}$ if he did not get the type-$\mathcal{A}$ good and with probability 0 if he did.

[6]The important difference is that our capacity constraint is more restrictive, since—if mapped to a monetary setting—

good) with money and budget constraints allows us to fairly directly apply Myerson's foundational result characterizing DSIC auctions [Myerson, 1981], giving us the following:

PROPOSITION 1. *A mechanism $M_{q,p}$ is DSIC if and only if, $\forall v \in [1,U]^n$, $\forall i \in I$, $\forall v_i' \in [v_i, U]$, $q_i(v_i, v_{-i}) \le q_i(v_i', v_{-i})$ and:*

$$p_i(v_i, v_{-i}) = p_i(1, v_{-i}) + q_i(1, v_{-i}) - \left[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\, dz\right] \quad (1)$$

Clearly AAM is DSIC (and feasible) since $q_i(v) = \frac{1}{n}$ and $p_i(v) = \frac{n-1}{n}$, $\forall i \in I$, $\forall v \in [1,U]^n$. From a social welfare perspective, we will now see, the only mechanisms worth considering allocate *some* good with certainty to any agent that reports the lowest possible value (i.e., 1).[7]

LEMMA 1. *For any DSIC and feasible mechanism $M_{q,p}$ where, for some $i \in I$ and $v_{-i} \in [1,U]^{n-1}$, $p_i(1, v_{-i}) + q_i(1, v_{-i}) < 1$, there exists an alternate DSIC and feasible mechanism $M_{q',p'}$ that yields at least as much welfare as $M_{q,p}$ on every instance, with, $\forall i \in I$, $\forall v_{-i} \in [1,U]^{n-1}$, $p_i'(1, v_{-i}) + q_i'(1, v_{-i}) = 1$.*

Proof. Consider arbitrary DSIC and feasible mechanism $M_{q,p}$ that, for some $i \in I$ and $v_{-i} \in [1,U]^{n-1}$, defines $p_i(1, v_{-i}) + q_i(1, v_{-i}) < 1$. Note that for arbitrary $v_i \in [1,U]$,

$$\begin{aligned} q_i(v_i, v_{-i}) &= v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(v_i, v_{-i})\, dz \\ &\le v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\, dz \end{aligned}$$

And from Proposition 1 we know that, for arbitrary $v_i \in [1,U]$,

$$\begin{aligned} &p_i(v_i, v_{-i}) + \left[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\, dz\right] \\ &= p_i(1, v_{-i}) + q_i(1, v_{-i}) \end{aligned}$$

Putting these facts together, we have, $\forall v_i \in [1,U]$:

$$p_i(v_i, v_{-i}) + q_i(v_i, v_{-i}) \le p_i(1, v_{-i}) + q_i(1, v_{-i}) < 1$$

Then, consider an alternate mechanism $M_{q,p'}$ that is identical to $M_{q,p}$ except that it defines $p_i'(v) = p_i(v) + \big[1 - p_i(1, v_{-i}) - q_i(1, v_{-i})\big]$, $\forall v \in [1,U]^n$. This mechanism is also DSIC and feasible and yields strictly greater expected efficiency than $M_{q,p}$. □

Since we are concerned with maximizing welfare in dominant strategies, given Proposition 1 and Lemma 1 it is sufficient to consider only mechanisms that define $p_i(1, v_{-i}) = 1 - q_i(1, v_{-i})$ and $p_i(v) = 1 - \big[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\, dz\big]$, $\forall v \in [1,U]^n$, $\forall i \in I$. Let $\mathcal{M}$ denote the set of all such mechanisms. Then, for arbitrary type-$\mathcal{A}$ good allocation function vector q, let $\mathcal{M}_q$ denote the member of $\mathcal{M}$ with the specified q, i.e., $\mathcal{M}_q = M_{q,p}$ for p as defined above.

it would also imply a *minimum* on how much an agent is required to pay as a function of the probability of allocation; this combined with our adoption of DSIC rather than Bayesian incentive compatibility makes the results of Maskin [2000] and Pai and Vohra [2008] on efficient allocation with budget constraints inapplicable here.

[7] See footnote 5 on how lotteries can be implemented in an interdependent way to achieve this.

THEOREM 1. *For arbitrary $q = (q_1, \dots, q_n)$, $\mathcal{M}_q$ is DSIC and feasible if and only if, $\forall v \in [1,U]^n$, $\forall i \in I$, $\forall \hat{v}_i \in [v_i, U]$, $q_i(v_i, v_{-i}) \le q_i(\hat{v}_i, v_{-i})$ and the following two conditions are satisfied:*

(C1) $\forall v \in [1,U]^n$, $\sum_{i \in I} q_i(v) \le 1$

(C2) $\forall v \in [1,U]^n, \forall i \in I, v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\, dz \le 1$

No DSIC mechanism that violates (C1) or (C2) is feasible.

Theorem 1 establishes that evaluating the space of possible (welfare maximizing) mechanisms for this problem reduces to evaluating type-$\mathcal{A}$ good allocation functions (q) that satisfy *(C1)* and *(C2)*. Now, the following lemma establishes that the expected welfare of any mechanism in the class $\mathcal{M}$ can be described very concisely in terms of the type-$\mathcal{A}$ good allocation rule and the value distribution.

LEMMA 2. *The expected welfare of a mechanism $\mathcal{M}_q \in \mathcal{M}$ equals:*

$$\sum_{i \in I} \mathbb{E}_{\tilde{v}}\left[\frac{1 - F(\tilde{v}_i)}{f(\tilde{v}_i)} q_i(\tilde{v})\right] + n$$

With these facts in hand, we will now explore the design of DSIC mechanisms. First, we will see that given the three preceding results—Lemma 1, Theorem 1, and Lemma 2—there is a relatively short path to proving that when the value distribution satisfies a monotone hazard rate condition, it is impossible to achieve better expected welfare than arbitrary allocation does.

2.1 Optimality of arbitrary allocation in MIHR settings

Hartline and Roughgarden [2008], considering a single-item allocation environment with money, demonstrate that the expected aggregate utility of the agents in any DSIC mechanism with allocation rule x equals $\sum_{i \in I} \mathbb{E}_{\tilde{v}}\big[\frac{1-F(\tilde{v}_i)}{f(\tilde{v}_i)} x_i(\tilde{v})\big]$. Mapping x to our type-$\mathcal{A}$ good allocation function q, Lemma 2 demonstrates that expected welfare in our setting is simply a constant plus the same expression derived by Hartline and Roughgarden. In considering welfare-optimal mechanisms, then, the distinction between the results they obtain and what we can expect here will be driven by two factors: 1) Hartline and Roughgarden adopt a *no positive transfers* restriction (which would map to a constraint that $p_i(v) = 0$, $\forall v$, here) that is not at play in our environment, and 2) we are bound by the capacity constraint: $\forall i \in I$, $\forall v \in [1,U]^n$, $q_i(v) + p_i(v) \le 1$.

However, when a hazard rate condition on the value distribution holds, we will show that those distinctions do not lead to a different outcome. We now adapt a lemma from the Hartline and Roughgarden paper; it applies to our setting with essentially no modification. The *hazard rate* of a distribution F, evaluated at $x \in [1,U]$ is defined as: $\frac{f(x)}{1-F(x)}$. F has monotone increasing hazard rate (MIHR) if and only if, $\forall x, x' \in [1,U]$ with $x \le x'$, $\frac{f(x)}{1-F(x)} \le \frac{f(x')}{1-F(x')}$.

PROPOSITION 2. *[Hartline and Roughgarden, 2008][8] If the value distribution has MIHR, then, $\forall i \in I$, $\mathbb{E}_{\tilde{v}}\big[\frac{1-F(\tilde{v}_i)}{f(\tilde{v}_i)} q_i(\tilde{v})\big] \le \big(\int_1^U f(x) x\, dx - 1\big)\, \mathbb{E}_{\tilde{v}}\big[q_i(\tilde{v})\big]$.*

[8] In fact Lemmas 2.8 and 2.10 of Hartline and Roughgarden [2008] together constitute the claim that $\mathbb{E}_{\tilde{v}}\big[\frac{1-F(\tilde{v}_i)}{f(\tilde{v}_i)} q_i(\tilde{v})\big] \le$

This result, which was derived for somewhat different ends in a different context, when combined with Lemma 2 and dominance of the $\mathcal{M}$ class of mechanisms, yields the following fact bounding possible expected welfare in our setting.

COROLLARY 1. *If the value distribution has MIHR, then no DSIC and feasible mechanism yields expected welfare exceeding* $\int_1^U f(x)x\,dx + n - 1$ *in equilibrium.*

Now, the combination of Remark 1, Lemma 1, Lemma 2, and Corollary 1 establishes that when the value distribution has MIHR, no mechanism yields expected welfare that exceeds that yielded by arbitrary allocation.

THEOREM 2. *If the value distribution has MIHR, then no DSIC and feasible mechanism yields greater expected welfare than AAM in equilibrium.*

2.2 A substantive DSIC mechanism without payments

The preceding analysis demonstrates that the case is closed for i.i.d. environments with MIHR distributions—disappointingly, there is no hope of doing better than arbitrary allocation. But many real-world settings will fall outside this class. For instance in the case of i.i.d. values with monotonically *decreasing* hazard rate (MDHR), from Lemma 2 we can see that if we could allocate the good to agent $i \in \arg\max_{j\in J} \frac{1-F(v_j)}{f(v_j)}$ with probability 1, such a mechanism would yield the maximum possible expected efficiency.[9] Such a mechanism would not be feasible because it would violate *(C2)* of Theorem 1, but the following mechanism is a feasible (and DSIC) modification of this idea. Intuitively, the appropriate incentives are created by—instead of directly allocating the type-$\mathcal{A}$ good—allocating the high-bidder a *lottery* for it, defined such that only the agent with highest value would prefer the lottery to the type-$\mathcal{B}$ good he would get (with certainty) as a loser.

MECHANISM 2. (LOTTERY INCENTIVE MECHANISM (LIM)) *Each agent $i \in I$ submits a bid $\hat{v}_i$. Letting $q^{(k)}(\hat{v})$ and $p^{(k)}(\hat{v})$ respectively denote the type-$\mathcal{A}$ and type-$\mathcal{B}$ good allocation probabilities for the agent submitting the k^{th}-highest bid, breaking ties arbitrarily,*

$$q^{(1)}(\hat{v}) = \frac{1}{\hat{v}^{(2)}} \quad \text{and} \quad p^{(1)}(\hat{v}) = 0$$
$$q^{(k)}(\hat{v}) = 0 \quad \text{and} \quad p^{(k)}(\hat{v}) = 1, \quad \forall k \in \{2,\ldots,n\}$$

REMARK 2. *In the truthful equilibrium, given value profile v, expected welfare under LIM equals $n - 1 + \frac{v^{(1)}}{v^{(2)}}$.*

Consider an example with 4 agents whose values are 2, 1.8, 1.5, and 1.2, respectively. Under LIM, agents 2, 3, and 4 receive a type-$\mathcal{B}$ good; agent 1 receives a type-$\mathcal{A}$ good with probability $\frac{1}{1.8}$ and receives nothing with probability $\frac{0.8}{1.8}$. Expected social welfare equals $3 + \frac{2}{1.8}$.

$\int_1^U f(x)x\,dx\,\mathbb{E}_{\tilde{v}}\left[q_i(\tilde{v})\right]$, but this is apparently a minor oversight; the claim as given would hold true only for value distributions that have 0 as the lowest possible value (note that ours have 1 as the lowest possible value).

[9] Hartline and Roughgarden [2008] make this very conclusion for their setting with money.

Besides being DSIC and feasible, this mechanism is optimal among mechanisms that never allocate the type-$\mathcal{A}$ good to anyone other than the high-bidder.

THEOREM 3. *The Lottery Incentive Mechanism is DSIC and feasible. For arbitrary value distribution, among all DSIC and feasible mechanisms under which only the high-bidder has positive probability of receiving the type-$\mathcal{A}$ good, LIM maximizes expected welfare.*

Proof. Consider arbitrary agent $i \in I$ and arbitrary value profile v. First, we show that LIM is in $\mathcal{M}$, which requires that, $\forall v$, $p_i(1, v_{-i}) + q_i(1, v_{-i}) = 1$ and $p_i(v) = 1 - \left[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\,dz\right]$. $p_i(1, v_{-i}) = 1$ and $q_i(1, v_{-i}) = 0$ unless $v_{-i}^{(1)} = 1$ and i wins the tie-breaker, in which case $p_i(1, v_{-i}) = 0$ and $q_i(1, v_{-i}) = 1$. $1 - \left[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\,dz\right] = 0$ if $v_i > v_{-i}^{(1)}$ and equals 1 otherwise, so the conditions for membership in $\mathcal{M}$ are both satisfied. Therefore, by Theorem 1, to show DSIC and feasibility it is sufficient to show that, $\forall v$, q_i is monotone in v_i, $\sum_{i\in I} q_i(v) \le 1$ and $v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\,dz \le 1$. Satisfaction of each of these is obvious from LIM's definition.

Now consider arbitrary DSIC and feasible mechanism $M_{q,p}$ that allocates the type-$\mathcal{A}$ good with positive probability only to the high-bidder. Consider arbitrary value profile v with $v^{(2)} > 1$ (when $v^{(2)} = 1$ LIM is clearly optimal), and let $i \in \arg\max_{j\in I} v_j$. Assume ties are broken in favor of i (the proof extends easily otherwise). By DSIC and Proposition 1, $p_i(v) = p_i(1, v_{-i}) + q_i(1, v_{-i}) - \left[v_i q_i(v_i, v_{-i}) - \int_1^{v_i} q_i(z, v_{-i})\,dz\right]$. By the assumption that only the high-bidder receives the type-$\mathcal{A}$ good with positive probability,

$$p_i(v) = p_i(1, v_{-i}) - \left[v_i q_i(v_i, v_{-i}) - \int_{v_{-i}^{(1)}}^{v_i} q_i(z, v_{-i})\,dz\right]$$

Now assume for contradiction that for some $v_i' \ge v_{-i}^{(1)}$, $q_i(v_i', v_{-i}) = 1/v_{-i}^{(1)} + \delta$ for some $\delta > 0$. Then:

$$\begin{aligned} p_i(v_i', v_{-i}) &= p_i(1, v_{-i}) - \left[v_i'\Big(\frac{1}{v_{-i}^{(1)}} + \delta\Big) - \int_{v_{-i}^{(1)}}^{v_i'} q_i(z, v_{-i})\,dz\right] \\ &\le p_i(1, v_{-i}) - \left[v_i'\Big(\frac{1}{v_{-i}^{(1)}} + \delta\Big) - \int_{v_{-i}^{(1)}}^{v_i'} \Big(\frac{1}{v_{-i}^{(1)}} + \delta\Big)\,dz\right] \\ &= p_i(1, v_{-i}) - 1 - v_{-i}^{(1)}\delta < 0 \end{aligned}$$

The first inequality holds by monotonicity of q_i with respect to v_i, and the second since $p_i(1, v_{-i}) \le 1$. We've reached a contradiction, since p_i cannot be negative. Thus, among all DSIC and feasible mechanisms that only ever allocate the type-$\mathcal{A}$ good to the highest bidder, LIM minimizes the probability that the type-$\mathcal{A}$ good remains unallocated *on every instance*. This combined with Lemmas 1 and 2 entails that LIM maximizes expected efficiency among all such mechanisms. □

In cases where there is more than one bidder who values the type-$\mathcal{A}$ good vastly more than he values type-$\mathcal{B}$ goods, under LIM the type-$\mathcal{A}$ good will unfortunately remain unallocated with very high probability. However, as the proof of Theorem 3 establishes, this is a necessary characteristic of any DSIC mechanism that only allocates the type-$\mathcal{A}$ good to the high-bidder.

But for some natural non-MIHR distributions LIM will in fact yield greater expected welfare than arbitrary allocation. To take one example, consider the Weibull distribution, bounded below at 1, with scale parameter $\lambda = 1$ and shape parameter $k = 0.5$: $\forall x \in [1, \infty]$, $f(x) = \frac{1}{2e^{\sqrt{x-1}}\sqrt{x-1}}$ and $F(x) = 1 - \frac{1}{e^{\sqrt{x-1}}}$. This distribution is not at all unreasonable in our setting—intuitively, values for the type-$\mathcal{A}$ good close to the value for the type-$\mathcal{B}$ good are most probable, with density falling off sharply as the increase in value over that for type-$\mathcal{B}$ goods increases. Expected social welfare under AAM equals: $\int_1^\infty \frac{x}{2e^{\sqrt{x-1}}\sqrt{x-1}}\,dx = 3$. Expected social welfare under LIM when there are 2 agents equals:

$$2\int_1^\infty \frac{x}{2e^{\sqrt{x-1}}\sqrt{x-1}} \int_1^x \frac{1}{2ye^{\sqrt{y-1}}\sqrt{y-1}}\,dy\,dx \approx 3.174$$

And so with this value distribution, when there are 2 agents, in expectation LIM obtains greater equilibrium welfare than the baseline AAM.

2.3 A mechanism that redistributes allocation probability

Theorem 3 demonstrated that LIM is optimal under the constraint that only the high-bidder may receive the type-$\mathcal{A}$ good, but such a constraint is not particularly well-motivated; if allocating the good to lower bidders improves welfare, why shouldn't we do so? In this section we take inspiration from Bailey [1997] and Cavallo [2006], wherein mechanisms with monetary payments are derived with an eye towards minimizing the *revenue* (i.e., aggregate payments made by the agents to the center) in order to maximize the agents' aggregate welfare. The project in these papers is framed as "redistribution" of the revenue generated by the VCG mechanism. The setting we are concerned with here—and the inadequacy of LIM—is quite different, but an analogy can be drawn between revenue in the VCG mechanism and probability of non-allocation in LIM.

Revenue redistribution mechanisms succeed by establishing, for each agent i, a lower-bound on the VCG revenue that would result independent of the type i announces; this provides a basis for returning revenue to i that neither alters his incentives nor runs the risk of a budget deficit. Here, we will instead seek to establish, for each agent i, a lower-bound on the probability with which the type-$\mathcal{A}$ good remains unallocated that is independent of i's bid; we will then use this as a basis for allocating the good to i with a probability that neither changes i's incentives nor runs a risk of making the mechanism infeasible (by specifying a total probability of allocation greater than 1 across all the agents).

Mechanism 3. (Lottery Redistribution Mechanism (LRM)) *Each agent $i \in I$ submits a bid $\hat{v}_i$. Letting $q^{(k)}(\hat{v})$ and $p^{(k)}(\hat{v})$ respectively denote the type-$\mathcal{A}$ and type-$\mathcal{B}$ good allocation probabilities for the agent submitting the k^{th}-highest bid, breaking ties arbitrarily,*

$$q^{(1)}(\hat{v}) = \frac{1}{\hat{v}^{(2)}}\left(1 - \frac{1-\frac{1}{\hat{v}^{(3)}}}{n}\right) + \frac{1-\frac{1}{\hat{v}^{(3)}}}{n} \quad \textit{and} \quad p^{(1)}(\hat{v}) = 0,$$

$$q^{(2)}(\hat{v}) = \frac{1-\frac{1}{\hat{v}^{(3)}}}{n} \quad \textit{and} \quad p^{(2)}(\hat{v}) = 1 - \frac{1-\frac{1}{\hat{v}^{(3)}}}{n},$$

$$q^{(k)}(\hat{v}) = \frac{1-\frac{1}{\hat{v}^{(2)}}}{n} \quad \textit{and} \quad p^{(k)}(\hat{v}) = 1 - \frac{1-\frac{1}{\hat{v}^{(2)}}}{n}, \ \forall k \in \{3, \ldots, n\}$$

To describe the mechanism in a more concise way that highlights the "redistribution" character of it, we can let $r_i = \frac{1}{n}\big(1 - 1/\hat{v}_{-i}^{(2)}\big)$, $\forall i \in I$, $\forall \hat{v} \in [1, U]^n$; then the high-bidder h receives the type-$\mathcal{A}$ good with probability $\frac{1}{\hat{v}_{-h}^{(1)}}(1 - r_h) + r_h$ and never receives the type-$\mathcal{B}$ good, while each other agent j receives the type-$\mathcal{A}$ good with probability r_j and the type-$\mathcal{B}$ good with probability $1 - r_j$. It is not hard to show that as n grows, the probability that the type-$\mathcal{A}$ good remains unallocated converges to 0, regardless of the value distribution. LRM is feasible, provides the proper incentives, and outperforms LIM on every possible value profile.

Theorem 4. *The Lottery Redistribution Mechanism is DSIC and feasible, and yields weakly greater expected welfare than LIM on every value profile.*

Remark 3. *In the truthful equilibrium, given value profile v, expected welfare under LRM equals:*

$$n - 1 + \left(\frac{1 - r^{(1)}(v)}{v^{(2)}} + r^{(1)}(v)\right)v^{(1)} + \sum_{2 \le j \le n} r^{(j)}(v)(v^{(j)} - 1),$$

where, $r^{(1)}(v) = r^{(2)}(v) = \frac{1}{n}\big(1 - \frac{1}{v^{(3)}}\big)$, and $\forall k \in \{3, \ldots, n\}$, $r^{(k)}(v) = \frac{1}{n}\big(1 - \frac{1}{v^{(2)}}\big)$.

In most cases LRM will yield much greater welfare than LIM. However, despite its superiority, we know from Theorem 2 that for MIHR value distributions it will still not outperform AAM (nothing does). As with LIM, the situation is different when the value distribution is heavy-tailed. We will forgo a specific illustration of these welfare differences, but in the next section will provide a comparative analysis across a range of distributions under a somewhat different evaluation metric.

3. RESIDUAL CLAIMANT FOR THE NUMÉRAIRE GOOD

We were able to identify scenarios under which the expected welfare under LIM or LRM significantly exceeds the baseline of random allocation. But for i.i.d. values and many "typical" distributions (uniform, normal, etc.), random allocation cannot be beat (Theorem 2). Intuitively, the impossibility of exceeding that benchmark stems from the fact that, in order to provide the proper incentives for truthtelling, the agent allocated the high-value good must bear a significant probability of receiving nothing—not even a type-$\mathcal{B}$ good. LRM mitigates this issue by largely "redistributing" the type-$\mathcal{A}$ good non-allocation probability to the other agents (so that the type-$\mathcal{A}$ good is very unlikely to *not* be allocated). But this does not change the fact that there is a pure loss whenever the winning agent receives no goods rather than the numéraire (type-$\mathcal{B}$) good.

In this section we will move away from the assumption that type-$\mathcal{B}$ goods contribute to social welfare only if consumed by the agents, instead considering their allocation to be essentially inconsequential to welfare. There are several ways to motivate this: for instance, the center may be able to consume excess type-$\mathcal{B}$ goods himself, or there may be a pool of other agents with no interest in the type-$\mathcal{A}$ good who would obtain value for a type-$\mathcal{B}$ good.

Here the connection to mechanisms with a very constrained space of possible monetary payments becomes stronger; analogically, if type-$\mathcal{B}$ goods were money, the problem would amount to achieving high-welfare allocations of a single good under the following conditions: each agent has a fixed budget of \$1, values for the good are greater than \$1, no positive transfers to the agents are allowed, and each agent must pay at least his allocation probability (to satisfy the capacity constraint).

In this model our definition of *mechanism* as a vector of functions $(q_1, \ldots, q_n, p_1, \ldots, p_n)$ remains, but now the social welfare objective is to maximize $\sum_{i \in I} q_i v_i$ subject to the constraint that $\sum_{i \in I} q_i \leq 1$ and, $\forall i \in I$, $p_i \in [0, q_i]$. Proposition 1 and Theorem 1 characterizing DSIC and feasible mechanisms continue to apply. As in the case without a residual claimant, there exists no DSIC and feasible mechanism outside of $\mathcal{M}$ that yields greater expected welfare than the best DSIC and feasible mechanism in $\mathcal{M}$. But now, expected welfare equals $\sum_{i \in I} \mathbb{E}_{\tilde{v}}[\tilde{v}_i q_i(\tilde{v})]$ and even in MIHR settings arbitrary allocation may no longer be optimal. This is in fact the case; now for large enough populations of agents LRM dominates AAM *for all bounded value distributions*. We bypass formal statement and proof of this because, in fact, a much stronger asymptotic result is available.

We propose the following mechanism that, intuitively, can be thought of as granting each agent an initial endowment of a type-$\mathcal{B}$ good and then auctioning off lottery tickets for the type-$\mathcal{A}$ good; the number of lottery tickets is selected such that the "Vickrey price" of each (in terms of type-$\mathcal{B}$ good allocation probability) is less than each agent's initial endowment (i.e., 1). Recall that U is the maximum value any agent could possibly have.

MECHANISM 4. (MULTI-LOTTERY VICKREY MECHANISM (MVM)) *Each agent $i \in I$ submits a bid $\hat{v}_i$. Let $C = \lceil U \rceil$, and let $q^{(k)}(\hat{v})$ and $p^{(k)}(\hat{v})$ respectively denote the type-$\mathcal{A}$ and type-$\mathcal{B}$ good allocation probabilities for the agent submitting the k^{th}-highest bid, breaking ties arbitrarily. If $C \geq n$ then $q^{(k)}(\hat{v}) = \frac{1}{n}$ and $p^{(k)}(\hat{v}) = \frac{n-1}{n}$, $\forall k \in \{1, \ldots, n\}$. If $C < n$, then:*

$$q^{(k)}(\hat{v}) = \frac{1}{C} \quad \text{and} \quad p^{(k)}(\hat{v}) = 1 - \frac{\hat{v}^{C+1}}{C}, \quad \forall k \in \{1, \ldots, C\}$$

$$q^{(k)}(\hat{v}) = 0 \quad \text{and} \quad p^{(k)}(\hat{v}) = 1, \quad \forall k \in \{C+1, \ldots, n\}$$

THEOREM 5. *For arbitrary value distribution with support $[1, U]$, for arbitrary finite $U > 1$: MVM is DSIC and feasible; for arbitrary $\epsilon > 0$, $\exists n \in Z^+$ such that if there are at least n agents and a residual claimant for type-$\mathcal{B}$ goods, expected welfare under MVM is within ϵ of the expected welfare of optimal allocation.*

Theorem 5 is very good news when the number of agents is large. But to get a sense of how quickly MVM converges to optimal welfare,[10] and also of how the various mechanisms we've proposed compare with each other for smaller numbers of agents, we numerically computed expected allocation values under each across a range of different value distributions. For many distributions and small population sizes, LRM is actually superior to MVM, though only the latter will converge to optimality as the population size grows. Figure 1 illustrates these findings for three distributions (with other tested examples omitted due to lack of space).

[10]Note that convergence time will increase as a function of the upper-bound of the value distribution, and MVM reduces to AAM for unbounded distributions.

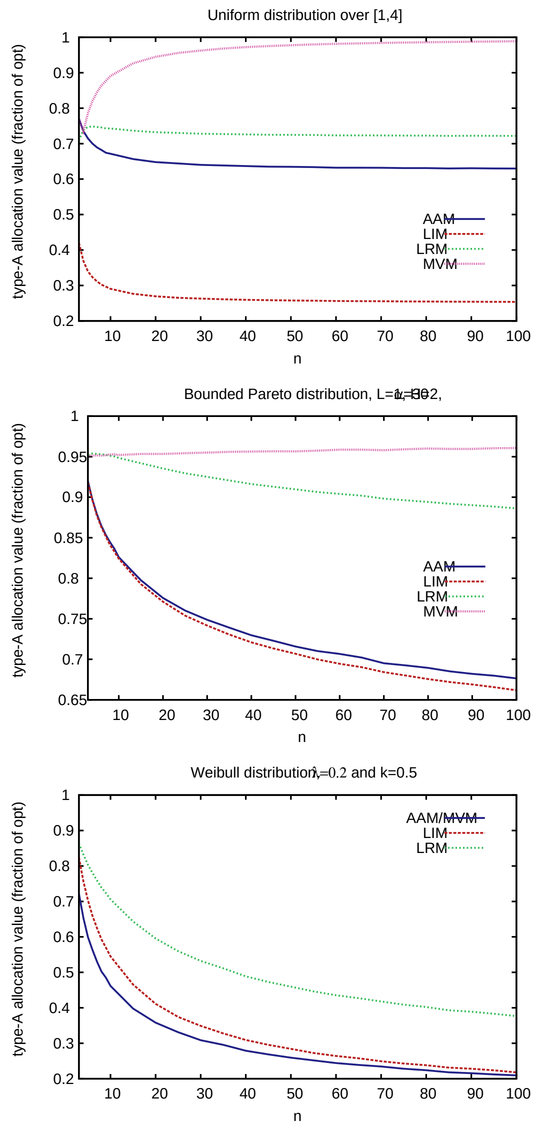

Figure 1: Expected value yielded from allocation of the type-$\mathcal{A}$ good under AAM, LIM, LRM, and MVM as a fraction of the no-private-information optimum (i.e., the expected highest value for the good), under 3 different value distributions. MVM coincides with AAM in the case of unbounded value distributions.

4. CONCLUSION

In this paper we explored non-monetary mechanism design for settings where there are two kinds of goods to be allocated, one that is singular and of high (but variable) value, and another that is plentiful and of common value. We found that if values for the singular good are i.i.d. according to a distribution that has monotone increasing hazard rate, no mechanism yields more welfare than allocating the good arbitrarily, oblivious to the agents' values. But in other cases better mechanisms do exist. We adapted techniques of *redistribution* from a line of research that seeks efficient mechanisms with minimal revenue, and derived a mechanism that outperforms the baseline in a variety of settings. We also proposed a mechanism that asymptotes to perfect efficiency as the population size grows when there is a residual claimant for the plentiful (numéraire) good.

Though space limitations prevented us from including an analysis of settings with non-identical value distributions, it is in such settings that these mechanisms really shine—all three of the mechanisms we introduced very often perform significantly better than arbitrary allocation. Intuitively, the type-$\mathcal{A}$ good allocation-value under LIM equals the ratio of highest to second highest value; in the i.i.d. case this discrepancy will be large in expectation only for heavy-tailed distributions, but if value distributions are asymmetric across agents then it will be large in a broad range of cases.

From a design perspective, we should be curious not just about mechanisms that outperform a baseline, but further about *welfare-optimal* mechanisms. It would be nice to pinpoint, for any given value distribution, a mechanism that maximizes expected welfare among all feasible and DSIC mechanisms. But this is more daunting than it may first seem; it is a more constrained variant of the welfare-optimal mechanism design problem for a setting with payments and budget-constraints, where known results are relatively modest. Recent work there has either identified optimal mechanisms in restricted classes or sought distribution-independent worst-case approximations—not even to the welfare of the optimal mechanism, but to a *proxy* for that welfare (see Devanur et al. [2013]).

In an environment where monetary transfers are impossible, the mere existence of a mechanism that reaches better-than-random allocations is significant; and in some settings the mechanisms we proposed go far beyond that. The two-tiered resource allocation problem we've addressed is one of many where mechanisms without payments could be useful; future work may explore other compelling examples.

References

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

Ruggiero Cavallo. 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)*, pages 882–889, 2006.

Edward Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*, pages 251–268, 2013a.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Positive results for mechanism design without money. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'13)*, pages 1165–1166, 2013b.

Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. Prior-free auctions for budgeted agents. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*, pages 287–304, 2013.

Shaddin Dughmi and Arpita Ghosh. Truthful assignment without money. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC'09)*, pages 325–334. ACM, 2010.

Uriel Feige and Moshe Tennenholtz. Responsive lotteries. *Algorithmic Game Theory*, pages 150–161, 2010.

Jacob Glazer and Albert Ma. Efficient allocation of a "prize"–King Solomon's dilemma. *Games and Economic Behavior*, 1(3): 222–233, 1989.

Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07), San Diego, CA*, pages 30–39, 2007.

Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, pages 881–888, 2010.

Mingyu Guo, Vincent Conitzer, and Daniel M Reeves. Competitive repeated allocation without payments. In *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE'09)*, pages 244–255, 2009.

Li Han, Chunzhi Su, Linpeng Tang, and Hongyang Zhang. On strategy-proof allocation without payments or priors. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE'11)*, pages 182–193, 2011.

B. Paul Harrenstein, Mathijs M. de Weerdt, and Vincent Conitzer. A qualitative vickrey auction. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC'09)*, pages 197–206. ACM, 2009.

Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing (STOC'08)*, pages 75–84, 2008.

Leonid Hurwicz and Mark Walker. On the generic nonoptimality of dominant-strategy allocation mechanisms: A general theorem that includes pure exchange economies. *Econometrica*, 58(3):683–704, 1990.

Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979.

Eric S. Maskin. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44:667–681, 2000.

John Moore. Implementation, contracts, and renegotiation in environments with complete information. *Advances in Economic Theory*, 1:182–281, 1992.

Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

Wojciech Olszewski. A simple and general solution to King Solomon's problem. *Games and Economic Behavior*, 42(2):315–318, 2003.

Mallesh Pai and Rakesh Vohra. Optimal auctions with financially constrained bidders. Technical report, Discussion paper, Center for Mathematical Studies in Economics and Management Science, 2008.

Motty Perry and Philip J Reny. A general solution to King Solomon's dilemma. *Games and Economic Behavior*, 26(2):279–285, 1999.

Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proc. of the 10th ACM Conf. on Electronic Commerce (EC'09)*, pages 177–186. ACM, 2009.

Cheng-Zhong Qin and Chun-Lei Yang. Make a guess: a robust mechanism for King Solomon's dilemma. *Economic Theory*, 39 (2):259–268, 2009.

William Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.