Competitive Repeated Allocation Without Payments*

Mingyu Guo[1], Vincent Conitzer[1], and Daniel M. Reeves[2]

[1] Duke University, Department of Computer Science, Durham, NC, USA
{mingyu, conitzer}@cs.duke.edu
[2] Yahoo! Research, 111 W. 40th St. 17th Floor, New York, NY, USA
dreeves@yahoo-inc.com

Abstract. We study the problem of allocating a single item repeatedly among multiple competing agents, in an environment where monetary transfers are not possible. We design (Bayes-Nash) incentive compatible mechanisms that do not rely on payments, with the goal of maximizing expected social welfare. We first focus on the case of two agents. We introduce an artificial payment system, which enables us to construct *repeated* allocation mechanisms *without payments* based on *one-shot* allocation mechanisms *with payments*. Under certain restrictions on the discount factor, we propose several repeated allocation mechanisms based on artificial payments. For the simple model in which the agents' valuations are either high or low, the mechanism we propose is 0.94-competitive against the optimal allocation mechanism with payments. For the general case of any prior distribution, the mechanism we propose is 0.85-competitive. We generalize the mechanism to cases of three or more agents. For any number of agents, the mechanism we obtain is at least 0.75-competitive. The obtained competitive ratios imply that for repeated allocation, artificial payments may be used to replace real monetary payments, without incurring too much loss in social welfare.

1 Introduction

An important class of problems at the intersection of computer science and economics deals with allocating resources among multiple competing agents. For example, an operating system allocates CPU time slots to different applications. The resources in this example are the CPU time slots and the agents are the applications. Another example scenario, closer to daily life, is "who gets the TV remote control." Here the resource is the remote control and the agents are the members of the household. In both scenarios the resources are allocated repeatedly among the agents, and monetary transfers are infeasible (or at least inconvenient). In this paper, we investigate problems like the above. That is, we study how to allocate resources in a repeated setting, without relying on payments. Our objective is to maximize social welfare, i.e., allocative efficiency.

The problem of allocating resources among multiple competing agents when monetary transfers are possible has been studied extensively in both the one-shot mechanism design setting [9, 6, 20, 16, 19, 15] and the repeated setting [11, 7, 10, 5]. A question that

* Guo and Conitzer thank the National Science Foundation and the Alfred P. Sloan Foundation for support (through award number IIS-0812113 and a Research Fellowship, respectively).

has recently been drawing the attention of computer scientists is how to design mechanisms without payments to achieve competitive performance against mechanisms with payments [21, 13]. [3] This paper falls into this category. We consider mechanisms without payments in repeated settings. A paper that lays out many of the foundations for repeated games is due to Abreu *et al.* [2], in which the authors investigate the problem of finding pure-strategy sequential equilibria of repeated games with imperfect monitoring. Their key contribution is the state-based approach for solving repeated games, where in equilibrium, the game is always in a *state* which specifies the players' long-run utilities, and on which the current period's payoffs are based. There are many papers that rely on the same or a similar state-based approach [22, 18, 17, 8].

The following papers are more related to our work: Fudenberg *et al.* [14] give a folk theorem for repeated games with imperfect public information. Both [14] and our paper are built on the (dynamic programming style) *self-generating* technique in [2] (it is called *self-decomposable* in [14]). However, [14] considers self-generation based on a certain supporting hyperplane, which is guaranteed to exist only when the discount factor goes to 1. [4] Therefore, their technique does not apply to our problem because we are dealing with non-limit discount factors. [5] Another difference between [14] and our paper is that we are designing specific mechanisms in this paper, instead of trying to prove the existence of a certain class of mechanisms. With non-limit discount factors, it is generally difficult to precisely characterize the set of feasible utility vectors (optimal frontier) for the agents. Several papers have already proposed different ways of approximation (for cases of non-limit discount factors). Athey *et al.* [4] study approximation by requiring that the payoffs of the agents must be symmetric. In what, from a technical perspective, appears to be the paper closest to the work in this paper, Athey and Bagwell [3] investigate collusion in a repeated game by approximating the optimal frontier by a line segment (the same technique also appears in the work of Abdulkadiroğlu and Bagwell [1]). One of their main results is that if the discount factor reaches a certain threshold (still strictly less than 1), then the approximation comes at no cost. That is, the optimal (first-best) performance can be obtained. However, their technique only works for finite type spaces, as it builds on uneven tie-breaking.

The main contribution of this paper can be summarized as follows. First, we introduce a new technique for approximating the optimal frontier for repeated allocation problem. Our technique works for non-limit discount factors and is not restricted to symmetric payoffs or finite type spaces. The technique we propose is presented in the form of an artificial payment system, which corresponds to approximating the optimal frontier by triangles. The artificial payment system enables us to construct repeated al-

[3] In the previous work, as well as in this paper, the first-best result can be achieved by mechanisms with payments.

[4] In [14], it is shown that any feasible and individually rational equilibrium payoff vector v can be achieved in a perfect public equilibrium (self-generated based on certain supporting hyperplanes), as long as the discount factor reaches a threshold $\underline{\beta}$. However, the threshold $\underline{\beta}$ depends on v. If we consider all possible values of v, then we essentially require that the discount factor/threshold approach 1, since any discount factor that is strictly less than 1 does not work (for some v).

[5] In this paper, we also require that the discount factor reaches a threshold, but here the threshold is a constant that works for all possible priors.

location mechanisms without payments based on one-shot allocation mechanisms with payments. We analytically characterize several repeated allocation mechanisms that do not rely on payments, and prove that they are competitive against the optimal mechanism with payments.

This paper also contributes to the line of research on designing competitive mechanisms without payments. The proposed artificial payment system provides a link between mechanisms with payments and mechanisms without payments. By proposing specific competitive mechanisms that do not rely on payments, our paper also provides an answer to the question: *Are monetary payments necessary for designing good mechanisms?* Our results imply that in repeated settings, artificial payments are "good enough" for designing allocation mechanisms with high social welfare. Conversely, it is easy to see that for one-shot settings, artificial payments are completely useless in the problem we study (single-item allocation).

The idea of designing mechanisms without payments to achieve competitive performance against mechanisms with payments was explicitly framed by Procaccia and Tennenholtz [21], in their paper titled *Approximate Mechanism Design Without Money*. That paper carries out a case study on locating a public facility for agents with single-peaked valuations. (The general idea of approximate mechanism design without payments dates back further, at least to work by Dekel *et al.* [13] in a machine learning framework.) To our knowledge, along this line of research, we are the first to to study allocation of private goods. Unlike the models studied in the above two papers [13, 21], where agents may have consensus agreement, when we are considering the allocation of private goods, the agents are fundamentally in conflict. Nevertheless, it turns out that even here, some positive results can be obtained if the allocation is carried out repeatedly. Thus, we believe that our results provide additional insights to this line of research.

2 Model Description

We study the problem of allocating a single item repeatedly between two (and later in the paper, more than two) competing agents. Before each allocation period, the agents learn their (private) valuations for having the item in that period (but not for any future periods). These preferences are independent and identically distributed, across agents as well as periods, according to a distribution F. We assume that these valuations are nonnegative and have finite expectations. F does not change over time. There are infinitely many periods, and agents' valuations are discounted according to a discount factor β. Our objective is to design a mechanism that maximizes expected social welfare under the following constraints (we allow randomized mechanisms):

- *(Bayes-Nash) Incentive Compatibility:* Truthful reporting is a Bayes-Nash equilibrium.
- *No Payments:* No monetary transfers are ever made.

In the one-shot mechanism design setting, incentive compatibility is usually achieved through payments. This ensures that agents have no incentive to overbid, because they may have to make large payments. In the repeated allocation setting, there are other ways to achieve incentive compatibility: for example, if an agent strongly prefers to obtain the item in the current period, the mechanism can ensure that she is less likely to obtain it in future periods. In a sense, this is an artificial form of payment. Such payments introduce some new issues that do not always occur with monetary payments, including that each agent effectively has a limited budget (corresponding to a limited amount of future utility that can be given up); and if one agent makes a payment to another agent by sacrificing some amount of future utility, the corresponding increase in the latter agent's utility may be different from the decrease in the former agent's utility.

3 State-Based Approach

Throughout the paper, we adopt the state-based approach introduced in Abreu *et al.* [2]. In their paper, the authors investigated the problem of finding pure-strategy sequential equilibria of repeated games with imperfect monitoring. Their problem can be rephrased as follows: Given a game, what are the possible pure-strategy sequential equilibria? Even though in our paper we are considering a different problem (we are *designing* the game), the underlying ideas still apply. In their paper, states correspond to possible equilibria, while in our paper, states correspond to feasible mechanisms. In this section, we review a list of basic results and observations on the state-based approach, specifically in the context of repeated allocation.

Let M be an incentive compatible mechanism without payments for a particular (fixed) repeated allocation problem, defined by a particular type distribution and a discount factor. If, under M, the expected long-term utilities of agents 1 and 2 (at the beginning) are x and y respectively, then we denote mechanism M by state (x, y). All mechanisms that can be denoted by (x, y) are considered equivalent. If we are about to apply mechanism M, then we say the agents are in state (x, y). In the first period, based on the agents' reported values, the mechanism specifies both how to allocate the item in this period, and what to do in the future periods. The rule for the future is itself a mechanism. Hence, a mechanism specifies how to allocate the item within the first period, as well as the state (mechanism) that the agents will be in in the second period. We have that $(x, y) = E_{v_1,v_2}[(r_1(v_1, v_2), r_2(v_1, v_2)) + \beta(s_1(v_1, v_2), s_2(v_1, v_2))]$, where v_1, v_2 are the first-period valuations, r_1, r_2 are the immediate rewards obtained from the first-period *allocation rule*, and (s_1, s_2) gives the second-period state, representing the *transition rule*.

State (x, y) is called a *feasible* state if there is a feasible mechanism (that is, an incentive compatible mechanism without payments) corresponding to it. We denote the set of feasible states by S^*. Let e be an agent's expected valuation for the item in a single period. $E = \frac{e}{1-\beta}$ is the maximal expected long-term utility an agent can receive (corresponding to the case where she receives the item in every period). Let O be the set of states $\{(x, y)|0 \leq x \leq E, 0 \leq y \leq E\}$. We have that $S^* \subseteq O - \{(E, E)\} \subsetneq O$.

S^* is convex, for the following reason. If (x_1, y_1) and (x_2, y_2) are both feasible, then $(\frac{x_1+x_2}{2}, \frac{y_1+y_2}{2})$ is also feasible (it corresponds to the randomized mechanism where we flip a coin to decide which of the two mechanisms to apply). S^* is symmetric with respect to the diagonal $y = x$: if (x, y) is feasible, then so is (y, x) (by switching the roles of the two agents).

The approximate shape of S^* is illustrated in Figure 1. There are three noticeable extreme states: $(0, 0)$ (nobody ever gets anything), $(E, 0)$ (agent 1 always gets the item), and $(0, E)$ (agent 2 always gets the item). S^* is confined by the x-axis (from $(0, 0)$ to $(E, 0)$), the y-axis (from $(0, 0)$ to $(0, E)$), and, most importantly, the bold curve, which corresponds to the optimal frontier. The square specified by the dotted lines represents O.

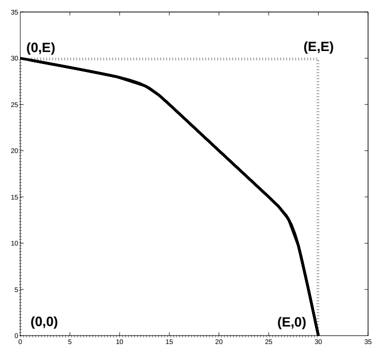

Fig. 1. The shape of S^*.

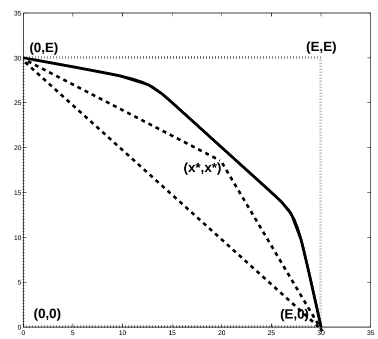

Fig. 2. Bow shape approximated by triangle.

Our objective is to find the state $(x^*, y^*) \in S^*$ that maximizes $x^* + y^*$ (expected social welfare). By convexity and symmetry, it does not hurt to consider only cases where $x^* = y^*$.

We now define a notion of when one set of states is *generated by* another. Recall that a mechanism specifies how to allocate the item within the first period, as well as which state the agents transition to for the second period. Let S be any set of states with $S \subset O$. Let us assume that, in the second period, exactly the states in S are feasible. That is, we assume that, if and only if $(x, y) \in S$, starting at the second period, there exists a feasible mechanism under which the expected utilities of agent 1 and 2 are x and y, respectively. Based on this assumption, we can construct incentive compatible mechanisms starting at the first period, by specifying an *allocation rule* for the first period, as well as a *transition rule* that specifies the states in S to which the agents will transition for the beginning of the second period. Now, we only need to make sure that the first period is incentive compatible. That is, the allocation rule in the first period, combined with the rule for selecting the state at the start of the second period, must incentivize the agents to report their true valuations in the first period. We say the set of resulting feasible states for the first period is *generated by* S, and is denoted by $Gen(S)$.

The following claim provides a general guideline for designing feasible mechanisms.

Claim 1 *For any $S \subseteq O$, if $S \subseteq Gen(S)$, then $S \subseteq S^*$. That is, if S is* self-generating, *then all the states in S are feasible.*

We now consider starting with the square O that contains S^* and iteratively generating sets. Let $O^0 = O$ and $O^{i+1} = Gen(O^i)$ for all i. The following claim, together with Claim 1, provide a general approach for computing S^*.

Claim 2 *The O^i form a sequence of (weakly) decreasing sets that converges to S^* if it converges at all. That is, $S^* = Gen(S^*)$. $S^* \subseteq O^i$ for all i. $O^{i+1} \subseteq O^i$ for all i. If $O^i = O^{i+1}$, then $O^i = S^*$.*

The above guideline leads to a numerical solution technique for finite valuation spaces. With a properly chosen numerical discretization scheme, we are able to compute an underestimation of O^i for all i, by solving a series of linear programs. The underestimations of the O^i always converge to an underestimation of S^* (a subset of S^*). That is, we end up with a set of feasible mechanisms. We are also able to show that as the discretization step size goes to 0, the obtained feasible set approaches S^*. That is, the numerical solution technique produces an optimal mechanism in the limit as the discretization becomes finer. Details of the numerical solution technique are omitted due to space constraint.

One drawback of the numerical approach is that the obtained mechanism does not have an elegant form. This makes it harder to analyze. From the agents' perspective, it is difficult to comprehend what the mechanism is trying to do, which may lead to irrational behavior. Another drawback of the numerical approach is that it only applies to cases of finite valuation spaces. For the rest of the paper, we take a more analytical approach. We aim to design mechanisms that can be more simply and elegantly described, work for any valuation space, and are (hopefully) close to optimality.

At the end of Section 4.2, we will compare the performances of the mechanisms obtained numerically and the mechanisms obtained by the analytical approach.

4 Competitive Analytical Mechanism

In this section, we propose the idea of an artificial payment system. Based on this, we propose several mechanisms that can be elegantly described, and we can prove that these mechanisms are close to optimality.

4.1 Artificial Payment System

Let us recall the approximate shape of S^* (Figure 2). The area covered by S^* consists of two parts. The lower left part is a triangle whose vertices are $(0, 0)$, $(E, 0)$, and $(0, E)$. These three states are always feasible, and so are their convex combinations. The upper right part is a bow shape confined by the straight line and the bow curve from $(0, E)$ to $(E, 0)$. To solve for S^*, we are essentially solving for the largest bow shape satisfying that the union of the bow shape and the lower-left triangle is self-generating. Here, we consider an easier problem. Instead of solving for the largest bow shape, we solve for the largest triangle (whose vertices are $(0, E)$,$(E, 0)$, and (x^*, x^*)) so that the union of the two triangles is self-generating (illustrated in Figure 2). That is, we want to find the largest value of x^* that satisfies that the set of convex combinations of $(0, 0)$, $(E, 0)$, $(0, E)$, and (x^*, x^*) is self-generating.

The triangle approximation corresponds to an *artificial payment system*. Let (x^*, x^*) be any feasible state satisfying $x^* \geq \frac{E}{2}$. Such a feasible state always exists (*e.g.*, $(\frac{E}{2}, \frac{E}{2})$). We can implement an artificial payment system based on (x^*, x^*), $(E, 0)$, and $(0, E)$, as follows. At the beginning of a period, the agents are told that the default option is that they move to state (x^*, x^*) at the beginning of the next period. However, if agent 1 wishes to pay v_1 ($v_1 \leq \beta x^*$) units of artificial currency to agent 2 (and agent 2 is not paying), then the agents will move to $(x^* - \frac{v_1}{\beta}, x^* + \frac{E-x^*}{x^*}\frac{v_1}{\beta})$. That is, the future state is moved $\frac{v_1}{\beta}$ units to the left along the straight line connecting $(0, E)$ and (x^*, x^*). (This corresponds to going to each of these two states with a certain probability.) By paying v_1 units of artificial currency, agent 1's expected utility is decreased by v_1 (the expected utility is decreased by $\frac{v_1}{\beta}$ at the start of the next period). When agent 1 pays v_1 units of artificial currency, agent 2 receives only $\frac{E-x^*}{x^*}v_1$ (also as a result of future utility). In effect, a fraction of the payment is lost in transmission. Similarly, if agent 2 wishes to pay v_2 ($v_2 \leq \beta x^*$) units of artificial currency to agent 1 (and agent 1 is not paying), then the agents will move to $(x^* + \frac{E-x^*}{x^*}\frac{v_2}{\beta}, x^* - \frac{v_2}{\beta})$. That is, the future state is moved $\frac{v_2}{\beta}$ units towards the bottom along the straight line connecting (x^*, x^*) and $(E, 0)$. If both agents wish to pay, then the agents will move to $(x^* - \frac{v_1}{\beta} + \frac{E-x^*}{x^*}\frac{v_2}{\beta}, x^* - \frac{v_2}{\beta} + \frac{E-x^*}{x^*}\frac{v_1}{\beta})$, which is a convex combination of $(0, 0)$, $(0, E)$, $(E, 0)$, and (x^*, x^*).

Effectively, both agents have a *budget* of βx^*, and when an agent pays the other agent, there is a *gift tax* with rate $1 - \frac{E-x^*}{x^*}$.

Based on the above artificial payment system, our approach is to design repeated allocation mechanisms without payments, based on one-shot allocation mechanisms with payments. In order for this to work, the one-shot allocation mechanisms need to take the gift tax into account, and an agent's payment should never exceed the budget limit.

The budget constraint is difficult from a mechanism design perspective. We circumvent this based on the following observation. An agent's budget is at least $\beta\frac{E}{2} = \frac{e\beta}{2-2\beta}$, which goes to infinity as β goes to 1. As a result, for sufficiently large discount factors, the budget constraint will not be binding. For the remainder of this paper, we ignore the budget limit when we design the mechanisms. Then, for each obtained mechanism, we specify how large the discount factor has to be for the mechanism to be well defined (that is, the budget constraint is not violated). This allows us to work around the budget constraint. The drawback is obvious: our proposed mechanisms only work for discount factors reaching a (constant) threshold (though it is not as restrictive as studying the limit case as $\beta \to 1$).

4.2 High/Low Types

We start with the simple model in which the agents' valuations are either H (high) with probability p or L (low) with probability $1 - p$. Without loss of generality, we assume that $L = 1$. We will construct a repeated allocation mechanism without payments based on the following *pay-only* one-shot allocation mechanism:

Allocation: If the reported types are the same, we determine the winner by flipping a (fair) coin. If one agent's reported value is high and the other agent's reported value is low, then we allocate the item to the agent reporting high.

Payment: An agent pays 0 if its reported type is low. An agent pays $\frac{1}{2}$ if its reported type is high (whether she wins or not); this payment does not go to the other agent.

Claim 3 *The above pay-only mechanism is (Bayes-Nash) incentive compatible.*

Now we return to repeated allocation settings. Suppose (x^*, x^*) is a feasible state. That is, we have an artificial payment system with gift tax rate $1 - \frac{E-x^*}{x^*}$. We apply the above one-shot mechanism, with the modifications that when an agent pays $\frac{1}{2}$, it is paying artificial currency instead of real currency, and the other agent receives $\frac{1}{2}\frac{E-x^*}{x^*}$. We note that the amount an agent receives is only based on the other agent's reported value. Therefore, the above modifications do not affect the incentives.

Under the modified mechanism, an agent's expected utility equals $\frac{T}{2} - P + P\frac{E-x^*}{x^*} + \beta x^*$. In the above expression, $T = 2p(1-p)H + p^2 H + (1-p)^2$ is the expected value of the higher reported value. $\frac{T}{2}$ is then the ex ante expected utility received by an agent as a result of the allocation. $P = \frac{p}{2}$ is the expected amount of artificial payment an agent pays. $P\frac{E-x^*}{x^*}$ is the expected amount of artificial payment an agent receives. βx^* is the expected future utility by default (if no payments are made).

If both agents report low, then, at the beginning of the next period, the agents go to (x^*, x^*) by default. If agent 1 reports high and agent 2 reports low, then the agents go to $(x^* - \frac{1}{2\beta}, x^* + \frac{E-x^*}{2\beta x^*})$, which is a convex combination of (x^*, x^*) and $(0, E)$. If agent 1 reports low and agent 2 reports high, then the agents go to $(x^* + \frac{E-x^*}{2\beta x^*}, x^* - \frac{1}{2\beta})$, which is a convex combination of (x^*, x^*) and $(E, 0)$. If both agents report high, then the agents go to $(x^* - \frac{1}{2\beta} + \frac{E-x^*}{2\beta x^*}, x^* - \frac{1}{2\beta} + \frac{E-x^*}{2\beta x^*})$, which is a convex combination of (x^*, x^*) and $(0, 0)$. Let S be the set of all convex combinations of $(0, 0)$, $(E, 0)$, $(0, E)$, and (x^*, x^*). The future states given by the above mechanism are always in S. If an agent's expected utility under this mechanism is greater than or equal to x^*, then S is self-generating. That is, (x^*, x^*) is feasible as long as x^* satisfies $x^* \leq \frac{T}{2} - P + P\frac{E-x^*}{x^*} + \beta x^*$.

We rewrite it as $ax^{*2} + bx^* + c \leq 0$, where $a = 1 - \beta$, $b = 2P - \frac{T}{2}$, and $c = -EP$. The largest x^* satisfying the above inequality is simply the larger solution of $ax^{*2} + bx^* + c = 0$, which is $\frac{\frac{T}{2} - 2P + \sqrt{(2P - \frac{T}{2})^2 + 4(1-\beta)EP}}{2(1-\beta)}$.

This leads to a feasible mechanism M^* (corresponding to state (x^*, x^*)). The expected social welfare under M^* is $2x^*$, where x^* equals the above solution.

We have not considered the budget limit. For the above M^* to be well-defined (satisfying the budget constraint), we need $\beta x^* \geq \frac{1}{2}$. Since $x^* \geq \frac{E}{2} = \frac{e}{2-2\beta} \geq \frac{1}{2-2\beta}$, we only need to make sure that $\frac{\beta}{2-2\beta} \geq \frac{1}{2}$. Therefore, if $\beta \geq \frac{1}{2}$, then M^* is well-defined. For specific priors, M^* could be well-defined even for smaller β.

Next, we show that (whenever M^* is well-defined) M^* is very close to optimality. Consider the *first-best allocation mechanism*: the mechanism that always successfully identifies the agent with the higher valuation and allocates the item to this agent (for free). This mechanism is not incentive compatible, and hence not feasible. The expected social welfare achieved by the first-best allocation mechanism is $\frac{T}{1-\beta}$, which is an upper bound on the expected social welfare that can be achieved by any mechanism with (or without) payments (it is a strict upper bound, as the dAGVA mechanism [12] is efficient, incentive compatible, and budget balanced).

Definition 1. *When the agents' valuations are either high or low, the prior distribution over the agents' valuations is completely characterized by the values of H and p. Let W be the expected social welfare under a feasible mechanism M. Let W^F be the expected social welfare under the first-best allocation mechanism. If $W \geq \alpha W^F$ for all H and p, then we say M is α-competitive. We call α a* competitive ratio *of M.*

Claim 4 *Whenever M^* is well-defined for all H and p, (*e.g., *$\beta \geq \frac{1}{2}$), M^* is 0.94-competitive.*

As a comparison, the lottery mechanism that always chooses the winner by flipping a fair coin has competitive ratio (exactly) 0.5 (if H is much larger than L and unlikely to occur).

In the following table, for different values of H, p, and β, we compare M^* to the near-optimal *feasible* mechanism obtained with the numerical solution technique. The table elements are the expected social welfare under M^*, the near-optimal feasible mechanism, the first-best allocation mechanism, and the lottery mechanism.

	M^*	Optimal	First-best	Lottery
$H = 2, p = 0.2, \beta = 0.5$	2.6457	2.6725	2.7200	2.4000
$H = 4, p = 0.4, \beta = 0.5$	5.5162	5.7765	5.8400	4.4000
$H = 16, p = 0.8, \beta = 0.5$	30.3421	30.8000	30.8000	26.0000
$H = 2, p = 0.2, \beta = 0.8$	6.6143	6.7966	6.8000	6.0000
$H = 2, p = 0.8, \beta = 0.8$	9.4329	9.8000	9.8000	9.0000
$H = 16, p = 0.8, \beta = 0.8$	75.8552	77.0000	77.0000	65.0000

4.3 General Valuation Space

In this section, we generalize the earlier approach to general valuation spaces. We let f denote the probability density function of the prior distribution. (A discrete prior distribution can always be smoothed to a continuous distribution that is arbitrarily close.)

We will construct a repeated allocation mechanism without payments based on the following *pay-only* one-shot allocation mechanism:

Allocation: The agent with the higher reported value wins the item.

Payment: An agent pays $\int_0^v tf(t)dt$ if it reports v.

This mechanism is actually a[6] dAGVA mechanism [12], which is known to be (Bayes-Nash) incentive compatible.

[6] "The" dAGVA mechanism often refers to a specific mechanism in a class of Bayes-Nash incentive compatible mechanisms, namely one that satisfies budget balance. In this paper, we will use "dAGVA mechanisms" to refer to the entire class, including ones that are not budget-balanced. Specifically, we will only use dAGVA mechanisms in which payments are always nonnegative.

The process is similar to that in the previous section. Due to space constraints, we omit the details. At the end, we obtain a feasible mechanism M^*. The expected social welfare under M^* is $2x^*$, where x^* equals $\frac{\frac{T}{2}-2P+\sqrt{(2P-\frac{T}{2})^2+4(1-\beta)EP}}{2(1-\beta)}$. Here, $T = \int_0^\infty \int_0^\infty \max\{t,v\} f(t)f(v)dtdv$ is the expected value of the higher valuation. $P = \int_0^\infty \int_0^v tf(t)dtf(v)dv$ is the expected amount an agent pays.

For the above M^* to be well-defined, we need the budget βx^* to be greater than or equal to $\int_0^\infty tf(t)dt = e$ (the largest possible amount an agent pays). Since $x^* \geq \frac{E}{2} = \frac{e}{2-2\beta}$, we only need to make sure $\frac{\beta e}{2-2\beta} \geq e$. Therefore, if $\beta \geq \frac{2}{3}$, then M^* is well-defined. For specific priors, M^* may be well-defined for smaller β.

Next, we show that (whenever M^* is well-defined) M^* is competitive against the first-best allocation mechanism for *all* prior distribution f. Naturally, the competitive ratio is slightly worse than the one obtained previously for high/low valuations. We first generalize the definition of competitiveness appropriately.

Definition 2. *Let W be the expected social welfare under a feasible mechanism M. Let W^F be the expected social welfare under the first-best allocation mechanism. If $W \geq \alpha W^F$ for all prior distributions, then we say that M is α-competitive. We call α a* competitive ratio *of M.*

Claim 5 *Whenever M^* is well-defined for all prior distributions (*e.g., $\beta \geq \frac{2}{3}$*), M^* is 0.85-competitive.*

5 Three or More Agents

We have focused on allocation problems with two agents. In this section, we generalize our analytical approach to cases of three or more agents.

Let n be the number of agents. We will continue with the state-based approach. That is, a mechanism (state) is denoted by a vector of n nonnegative real values. For example, if under mechanism M, agent i's long-term expected utility equals x_i, then mechanism M is denoted by $(x_1, x_2, \ldots, x_n)$. If we are about to apply mechanism M, then we say the agents are in state $(x_1, x_2, \ldots, x_n)$.

For any n, it is easy to see that the set of feasible states is convex and symmetric with respect to permutations of the agents. A state is called *fair* if all its elements are equal. For example, $(1, 1, 1)$ is a fair state ($n = 3$). When there is no ambiguity about the number of agents, the fair state $(x, x, \ldots, x)$ is denoted simply by x.

An artificial payment system can be constructed in a way that is similar to the case of two agents. Let μ_{n-1} be any feasible fair state for the case of $n-1$ agents. Then, the following n states are also feasible for the case of n agents:

$$(0, \underbrace{\mu_{n-1}, \ldots, \mu_{n-1}}_{n-1}), (\mu_{n-1}, 0, \underbrace{\mu_{n-1}, \ldots, \mu_{n-1}}_{n-2}), \ldots, (\underbrace{\mu_{n-1}, \ldots, \mu_{n-1}}_{n-1}, 0).$$

We denote the above n states by s_i for $i = 1, 2, \ldots, n$. Let $\hat{S}$ be the set of all feasible states with at least one element that equals 0. $\hat{S}$ is self-generating. Suppose we have a fair state μ_n for the case of n agents. Let S be the smallest convex set containing μ_n and all the states in $\hat{S}$. The s_i are in both $\hat{S}$ and S. An artificial payment system can be implemented as follows (for the case of n agents): The agents will go to state μ_n by default. If for all i, agent i chooses to pay v_i units of artificial currency, then we move to a new state whose i^{th} element equals $\mu_n - \frac{v_i}{\beta} + \gamma \sum_{j \neq i} \frac{v_j}{\beta}$. Here $\gamma = \frac{\mu_{n-1}-\mu_n}{\mu_n}$.[7] The new state M is in S. (The reason is the following. If only agent i is paying, and it is paying nv_i instead of v_i, then the new state M_i is $(\underbrace{\mu_n + \gamma\frac{nv_i}{\beta}, \ldots, \mu_n + \gamma\frac{nv_i}{\beta}}_{i-1}, \mu_n - \frac{nv_i}{\beta}, \underbrace{\mu_n + \gamma\frac{nv_i}{\beta}, \ldots, \mu_n + \gamma\frac{nv_i}{\beta}}_{n-i})$, which is a convex combination of μ_n and s_i. The average of the M_i over all i is just M. Thus M is a convex combination of μ_n and the s_i, which implies $M \in S$. [8])

With the above artificial payment system, by allocating the item to the agent with the highest reported value and charging the agents dAGVA payments, we get an incentive compatible mechanism. We denote agent i's reported value by v_i for all i. The dAGVA payment for agent i equals $E_{v_{-i}}(I(v_i \geq \max\{v_{-i}\}) \max\{v_{-i}\})$, where I is the characteristic function (which evaluates to 1 on *true* and to 0 otherwise) and v_{-i} is the set of reported values from agents other than i.

We still use P to denote the expected amount of payment from an agent. We use T to denote the expected value of the highest reported value. The expected utility for an agent is then $\frac{T}{n} - P + (n-1)\frac{\mu_{n-1}-\mu_n}{\mu_n}P + \beta\mu_n$.

To show S is self-generating, we only need to show μ_n is in $Gen(S)$. That is, μ_n is a feasible fair state as long as μ_n satisfies the following inequality: $\mu_n \leq \frac{T}{n} - P + (n-1)\frac{\mu_{n-1}-\mu_n}{\mu_n}P + \beta\mu_n$.

The largest solution of μ_n equals $\frac{\frac{T}{n} - nP + \sqrt{(nP - \frac{T}{n})^2 + 4(1-\beta)(n-1)\mu_{n-1}P}}{2(1-\beta)}$.

The above expression increases when the value of μ_{n-1} increases. The highest value for μ_1 is E (when there is only one agent, we can simply give the item to the agent for free). A natural way of solving for a good fair state μ_n is to start with $\mu_1 = E$, then apply the above technique to solve for μ_2, then μ_3, *etc.*

Next, we present a claim that is similar to Claim 5.

Claim 6 *Let n be the number of agents. Let M_n^* be the mechanism obtained by the technique proposed in this section. Whenever $\beta \geq \frac{n^2}{n^2+\frac{3}{4}}$, M_n^* is well defined for all priors, and is α_n-competitive, where $\alpha_1 = 1$, and for $n > 1$,*

$$\alpha_n = \min_{\{1 \leq u \leq \frac{n}{n-1}\}} n\frac{\frac{u}{n} - n + nu - u + \sqrt{(n - nu + u - \frac{u}{n})^2 + 4\alpha_{n-1}\frac{n-nu+u}{n}}}{2u}.$$

For all i, $\alpha_i \geq \frac{3}{4}$ holds.

As a comparison, the lottery mechanism that always chooses the winner uniformly at random has competitive ratio (exactly) $\frac{1}{n}$, which goes to 0 as n goes to infinity.

---

[7] It should be noted that when one agent pays 1, then *every* other agent receives γ. In a sense, γ already incorporates the fact that the payment must be divided among multiple agents.

[8] The above argument assumes that the available budget is at least n times the maximum amount an agent pays.

References

1. A. Abdulkadiroğlu and K. Bagwell. Trust, reciprocity and favors in cooperative relationships. Working paper, 2005.
2. D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–63, September 1990.
3. S. Athey and K. Bagwell. Optimal collusion with private information. Working papers 99-17, Massachusetts Institute of Technology (MIT), Department of Economics, Sept. 1999.
4. S. Athey, K. Bagwell, and C. Sanchirico. Collusion and price rigidity. *Review of Economic Studies*, 71(2):317–349, 04 2004.
5. S. Athey and I. Segal. An efficient dynamic mechanism. Technical report, 2007. Working Paper.
6. M. J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.
7. D. Bergemann and J. Välimäki. Efficient dynamic auctions. Cowles Foundation Discussion Papers 1584, Cowles Foundation, Yale University, Oct. 2006.
8. A. Casella. Storable votes. *Games and Economic Behavior*, 51:391–419, 2005.
9. R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.
10. R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 220–229, Chicago, IL, USA, 2008.
11. R. Cavallo, D. C. Parkes, and S. Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, Cambridge, MA, USA, 2006.
12. C. d'Aspremont and L.-A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.
13. O. Dekel, F. Fischer, and A. D. Procaccia. Incentive compatible regression learning. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 884–893, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
14. D. Fudenberg, D. I. Levine, and E. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, September 1994.
15. M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 210–219, Chicago, IL, USA, 2008.
16. M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009. Special Section of Games and Economic Behavior Dedicated to the 8th ACM Conference on Electronic Commerce.
17. H. A. Hopenhayn and C. Hauser. Trading favors: Optimal exchange and forgiveness. 2004 Meeting Papers 125, Society for Economic Dynamics, 2004.
18. M. M. Möbius. Trading favors. May 2001. Manuscript.
19. H. Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
20. R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118:209–228, 2004. Early version appeared in IJCAI-01.
21. A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 177–186, Stanford, CA, USA, 2009.
22. Y. Sannikov. Games with imperfectly observable actions in continuous time. *Econometrica*, 75(5):1285–1329, 09 2007.