

Dynamic Efficient Auctions^{*}

Dirk Bergemann[†] and Juuso Valimäki[‡]

ABSTRACT

(Extended Abstract) We consider the truthful implementation of the socially efficient allocation in a dynamic private value environment in which agents receive private information over time. We show that a suitable generalization of the Vickrey-Clark-Groves mechanism, based on the marginal contribution of each agent, leads to truth-telling in every period.

A leading example of a dynamic allocation model is the sequential auction of a single good in which the current winner of the object receives additional information about her valuation. We show that a modified sequential second price auction in which only the current winner makes a positive payment leads to truth-telling. In general allocations problems, the marginal contribution mechanism continues to induce truth-telling in every period but may now include positive transfers for many agents.

Keywords

ACM proceedings, L^AT_EX, text tagging

Categories and Subject Descriptors

H.4.m [Information Systems]: Auction Theory

General Terms

Auction theory

1. INTRODUCTION

The seminal analysis of second price auctions by Vickrey (1961) established that single or multiple unit discriminatory auctions can be used to implement the socially efficient allocation in private value models in (weakly) dominant strategies. The subsequent contributions by Clarke (1971) and Groves (1973) showed that the insight of Vickrey extends to more general allocation problems in private value environments. By requiring that the transfer payment of agent i match her externality cost on the remaining agents, agent i internalizes the social objective and is led to report her type truthfully. The resulting net utility for agent i corresponds to her marginal contribution to the social value.

In this paper, we generalize the idea of a marginal contribution mechanism to dynamic environments with private information. We design an intertemporal sequence of transfer payments which allow each agent to receive her flow marginal contribution in every period. In other words, each agent will pay her externality cost in a time consistent manner. In consequence, each agent is willing to truthfully report her information in every period.

The basic idea of the dynamic mechanism is first explored in the context of the sequential allocation of an indivisible object with initially uncertain value to the bidders. We assume that the initial estimate of the value is private information to the bidder. In subsequent periods, a bidder receives additional information only in those periods in which the object is allocated to her. The structure of the payoffs in the model, and in particular the resolution of uncertainty, therefore resembles the multi-armed bandit problem.

The first result reports the construction of a dynamically efficient auction that allocates the object in each period according to the utilitarian welfare criterion under symmetric but imperfect information. We show that a dynamic second price auction truthfully implements the socially efficient allocation period by period subject to Bayesian (and in fact even subject to ex post) incentive constraints. The bandit framework constitutes a natural setting to analyze the repeated allocation of an object or a license over time. The key assumption in the multi-armed bandit setting is that only the current user gains more information about her valuation of the object. If we think about the object as a license to use a facility or to explore a resource for a limited time, it is natural to assume that the current insider gains information relative to the outsiders. A conceptual advantage of the sequential allocation problem is that the structure of the socially efficient program is well understood. As the monetary transfers allow each agent to capture her marginal

^{*}(Produces the WWW2007-specific release, location and copyright information). For use with www2007-submission.cls V1.4. Supported by ACM.

[†]Department of Economics, Yale University, New Haven, CT 06520-8268, U.S.A., dirk.bergemann@yale.edu.

[‡]Department of Economics, Helsinki School of Economics and University of Southampton, 00100 Helsinki, Finland, juuso.valimaki@hse.fi.

contribution, the properties of the social program translate into properties of the marginal program. In the case of the dynamic auction, we therefore obtain surprisingly explicit and informative expressions for the intertemporal transfer prices.

The second result is the description of a dynamic Vickrey-Clark-Groves mechanism in which each agent receives in every period her flow marginal contribution to the social value. We obtain the second result for a general specification of the utility of each agent and the arrival of private information over time. Throughout the paper we maintain the assumptions of quasi-linear utility and of a private value environment.

The objective of the dynamic mechanism is to implement the socially efficient policy. With transferable utilities, the social objective is simply to maximize the expected discounted sum of the individual utilities. The solution to this dynamic optimization problem is by necessity time consistent. In consequence, the dynamic Vickrey-Clark-Groves mechanism is time consistent and the social choice function can be implemented by a sequential mechanism without any ex ante commitment by the designer. In contrast, in revenue maximizing problems, the “ratchet effect” leads to very distinct solutions for mechanisms with and without intertemporal commitment ability (see Freixas, Guesnerie, and Tirole (1985)).

In contrast to the static environment, the truth-telling strategy in the dynamic setting forms an ex-post equilibrium rather than an equilibrium in weakly dominant strategies. The weakening of the equilibrium notion is due to the dynamic nature of the game. If the connection between other agents’ current announcements and their implications on the future continuation payoffs is broken, then truth-telling is not necessarily individually optimal.

In recent years, a number of papers have been written with the aim to explore various issues arising in dynamic allocation problems. Athey and Segal (2006) consider a finite time horizon model with transferable utilities and private values. Their main result is the construction of a balanced budget mechanism in the finite horizon allocation model. Their construction of a rebalancing mechanism is based on a “team mechanism” in which the monetary transfers are paid only at the terminal period and are equal to the sum of the other agents’ terminal utilities. In contrast, we design a sequence of transfers which support the flow marginal contribution as the net utility of each agent in every period. In consequence we do not need a finite terminal time to establish the transfers. Bapna and Weber (2005) consider a sequential allocation problem for a single, indivisible object by a dynamic auction. The basic optimization problem is a multi-armed bandit problem as in the auction we discuss here. Their analysis attempts to use the Gittins index of each alternative allocation as a sufficient statistic for the determination of the transfer price. While the Gittins index is sufficient to determine the efficient allocation in each period, the indices, in particular the second highest index is typically not a sufficient statistic for the incentive compatible transfer price. Bapna and Weber (2005) present necessary and sufficient conditions when an affine but *report-contingent* combination of indices can represent the externality cost. In contrast, we consider a direct mechanism and determine the transfers from general principles of the incentive problem. In particular we do not require any assumptions beyond the

private value environment and transferable utility. In symmetric information environments, Bergemann and Välimäki (2003), (2006) use the notion of marginal contribution to construct efficient equilibria in dynamic first price auctions. In this paper, we emphasize the role of a time-consistent utility flow, namely the flow marginal contribution, to encompass environments with private information.

This paper is organized as follows. Section ?? sets up the basic auction model. Section ?? contains the construction of the efficient dynamic auction. Section ?? extends the construction to general private value environments. Section ?? concludes.

2. MODEL

2.0.0.1 Setting.

We consider a dynamic auction model in discrete time with an infinite horizon. In every period t , a single indivisible object can be allocated to a bidder $i \in \{1, \dots, N\}$. The true valuation of bidder i is given by $\omega_i \in \Omega_i = [0, 1]$. The prior distribution about the valuation ω_i is given by $F_i(\omega_i)$ and the distributions are independent across bidders. In period 0, bidder i does not know the realization of ω_i , instead she receives an informative signal $s_i^0 \in S_i = [0, 1]$ about her true value of the object. The signal s_i is generated by a conditional distribution function $G_i(s_i | \omega_i)$. In each subsequent period t , only the winning bidder in period $t-1$ receives additional information about her valuation ω_i in the form of an additional and conditionally independent signal $s_i^t \in S_i$ from the conditional distribution $G_i(s_i | \omega_i)$. Each signal s_i^t is private information to bidder i and is not observed by any other agent.¹

We denote the private history of bidder i by $h_i^t = (s_i^0, \dots, s_i^{t-1})$. The posterior belief of agent i about ω_i can be calculated by Bayes’ rule using h_i^t . The expected value of the object for bidder i given his private history is denoted by:

$$v_i(h_i^t) = \mathbb{E}[\omega_i | h_i^t].$$

Each agent i has quasi-linear utility and the net value of getting the object in period t is

$$v_i(h_i^t) - p_i^t,$$

where p_i^t is the transfer price paid in period t . Each agent discounts the future with a common discount factor δ , $0 < \delta < 1$.

2.0.0.2 Mechanism.

A dynamic direct mechanism asks the bidders to report their signals in every period t . The report \hat{s}_i^t may or may not be truthful. We define the initial reports by

$$\hat{h}^0 = (\hat{s}_1^0, \dots, \hat{s}_N^0),$$

and inductively the history of reports by:

$$\hat{h}^t = (\hat{h}^{t-1}, \hat{s}_1^t, \dots, \hat{s}_N^t).$$

¹We describe the arrival of new information as a Bayesian sampling process. The equilibrium characterization in Theorem 1 would continue to hold for any stochastic process, possibly non-Markovian, provided that the signal realizations are independent across agents and that signals only arrive for winning bidders.

The set of possible histories of reports in period t is denoted by \widehat{H}^t . The allocation rule for a dynamic direct revelation mechanism is

$$x^t : \widehat{H}^t \rightarrow [0, 1]^N.$$

The allocation in period t is a vector $x^t = (x_1^t, \dots, x_N^t)$, where x_i^t denotes the probability of assigning the object to i in t with

$$\sum_{i=1}^N x_i^t = 1.$$

The transfer (or pricing) rule is given by:

$$p^t : \widehat{H}^t \rightarrow \mathbb{R}^N.$$

A dynamic mechanism $\mathcal{M} = \langle \mathbf{x}, \mathbf{p}, \widehat{\mathbf{H}} \rangle$ is a triple where

$$\mathbf{x} = \{x^t\}_{t=0}^\infty, \quad \mathbf{p} = \{p^t\}_{t=0}^\infty \quad \text{and} \quad \widehat{\mathbf{H}} = \{\widehat{H}^t\}_{t=0}^\infty,$$

are the sequences of public decisions and public reports (histories).

2.0.0.3 Equilibrium.

The bidders evaluate payoffs according to the discounted expected payoff criterion. A reporting strategy for agent i is a mapping

$$m_i^t : S_i \rightarrow S_i.$$

For a given mechanism \mathcal{M} , the expected payoff for bidder i from reporting a sequence $\widehat{\mathbf{s}}_i = \{\widehat{s}_i^t\}$ of signals given that the others are reporting $\widehat{\mathbf{s}}_{-i} = \{\widehat{s}_{-i}^t\}$ is given by

$$\mathbb{E} \sum_{t=0}^\infty \delta^t \left[x_i^t \left(\widehat{h}^{t-1}, \widehat{s}_i^t, \widehat{s}_{-i}^t \right) v_i \left(h_i^t \right) - p_i^t \left(\widehat{h}^{t-1}, \widehat{s}_i^t, \widehat{s}_{-i}^t \right) \right].$$

Given the mechanism \mathcal{M} and the reporting strategies $\widehat{\mathbf{s}}_{-i}$, the optimal reporting strategy of bidder i solves a sequential optimization problem which can be phrased recursively in terms of value functions, or

$$V_i(\widehat{h}^{t-1}, h_i^t) = \max_{\widehat{s}_i^t \in S_i} \mathbb{E} \left\{ \begin{array}{l} x_i^t \left(\widehat{h}^{t-1}, \widehat{s}_i^t, \widehat{s}_{-i}^t \right) v_i \left(h_i^t \right) \\ - p_i^t \left(\widehat{h}^{t-1}, \widehat{s}_i^t, \widehat{s}_{-i}^t \right) + \delta V_i \left(\widehat{h}^t, h_i^{t+1} \right) \end{array} \right\}.$$

We say that the dynamic direct mechanism \mathcal{M} is *Bayesian incentive compatible*, if for every agent i , in every period t , truthtelling is a best response given that all other agents report truthfully. In terms of the value function, it means that for all i and all t , the solution to the dynamic programming equation:

$$V_i(h^{t-1}) =$$

$$\max_{\widehat{s}_i^t \in S_i} \mathbb{E} \left\{ \begin{array}{l} x_i^t \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) v_i \left(h_i^t \right) \\ - p_i^t \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) + \delta V_i \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) \end{array} \right\}.$$

is to report truthfully, i.e. to choose $\widehat{s}_i^t = s_i^t$. Finally, we say that the mechanism \mathcal{M} is *ex post incentive compatible* if truthtelling is a best response for agent i regardless of the distribution of signals of the other agents, or

$$s_i \in \arg \max_{\widehat{s}_i^t \in S_i} \left\{ \begin{array}{l} x_i^t \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) v_i \left(h_i^t \right) \\ - p_i^t \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) + \delta V_i \left(h^{t-1}, \widehat{s}_i^t, s_{-i}^t \right) \end{array} \right\},$$

for all $s_{-i}^t \in S_{-i}$. In the dynamic context, ex post incentive compatibility has to be qualified in the sense that is ex post with respect to all signals received in period t , but not ex post with respect to signals arriving after period t . Consequently, the value function $V_i(h^{t-1}, \widehat{s}_i^t, s_{-i}^t)$ is still the future expected value conditional on $h^{t-1}, \widehat{s}_i^t, s_{-i}^t$.

3. DYNAMIC AUCTION

We start with the single good allocation problem and show that it is possible to implement the socially efficient allocation in ex post equilibrium (and hence in Bayesian Nash equilibrium). The construction resembles to some extent a second price auction in each period. The transfer price of the winning bidder is calculated in each period by comparison to the optimal allocation policy within the set of bidders where the current winner is excluded. As a result, the winning bidder internalizes her effect on the welfare of other bidders. The transfer price of the losing bidders will be equal to zero provided that only the winning bidder receives additional information. The exact construction of the transfer prices follows the spirit of the Vickrey pricing, but the intertemporal trade-offs are fully taken into account.

3.0.0.4 Social Efficiency.

The socially efficient assignment policy is obtained by maximizing the utilitarian welfare criterion, namely the expected discounted sum of utilities. Given a history of signals h^s in period s , the socially optimal program can be written simply as

$$W(h^s) = \max_{\{x^t(h^t)\}_{t=s}^\infty} \mathbb{E} \sum_{t=s}^\infty \sum_{i=1}^N \delta^{t-s} x_i^t(h^t) v_i(h_i^t).$$

Alternatively, we can represent the social program in its recursive form:

$$W(h^s) = \max_{x^s(h^s)} \mathbb{E} \left\{ \sum_{i=1}^N x_i^s(h^s) v_i(h_i^s) + \delta W(h^s, x^s) \right\}.$$

The expected value $\mathbb{E}W(h^s, x^s)$ represents the optimal continuation value conditional upon the state h^s and the allocation x^s today. The socially optimal assignment problem is a standard multi-armed bandit problem and the optimal policy is characterized by an index policy (see Gittins (1989) and Whittle (1982) for a textbook introduction). In particular, we compute for every bidder i the Gittins index based exclusively on the information about bidder i . The index of bidder i in state h_i^t is the solution to the following optimal stopping problem:

$$\gamma_i(h_i^t) = \max_{\tau} \mathbb{E} \left\{ \frac{\sum_{s=0}^{\tau} \delta^s v_i(h_i^{t+s})}{\sum_{s=0}^{\tau} \delta^s} \right\}.$$

The socially efficient allocation policy $\mathbf{x}^* = \{x^{t*}\}_{t=0}^\infty$ is to choose in every period a bidder i with the maximal index:

$$x_i^{t*} > 0 \text{ if } \gamma_i(h_i^t) \geq \gamma_j(h_j^t) \text{ for all } j.$$

3.0.0.5 Marginal Contribution.

In the static Vickrey auction, the price of the winning bidder is equal to the highest valuation among the losing bidders. The highest value among the remaining bidders

represents the social opportunity cost of assigning the object to the winning bidder. In a dynamic framework, the social opportunity cost is determined by the optimal continuation plan in the absence of the current winner. It is therefore useful to define the value of the social program after removing bidder i from the set of agents:

$$W_{-i}(h^s) = \max_{\{x_{-i}^t(h^t)\}_{t=s}^{\infty}} \mathbb{E} \sum_{t=s}^{\infty} \sum_{j \neq i} \delta^{t-s} x_j^t(h^t) v_j(h_j^t).$$

The *marginal contribution* $M_i(h^t)$ of bidder i at history h^t is then naturally defined by:

$$M_i(h^t) = W(h^t) - W_{-i}(h^t). \quad (1)$$

The marginal contribution is the change in social value due to the addition of agent i and hence the possibility of assigning the object to i . The marginal contribution of agent i may be thought of as the information rent that agent i may be able to secure for herself in the direct mechanism. If bidder i can secure her marginal contribution in a time consistent manner, she should be able to receive the *flow marginal contribution* $m_i(h^t)$ in every period. The flow marginal contribution accrues incrementally over each period:

$$M_i(h^t) = m_i(h^t) + \delta M_i(h^t, x^{t*}).$$

As in the notations of the value functions above, $M_i(h^t, x^t)$ represents the marginal contribution of agent i in the continuation problem conditional on the history h^t and the allocation x^t today. The flow marginal contribution can be expressed more directly using the definition of the marginal contribution (1) as

$$m_i(h^t) = (W(h^t) - W_{-i}(h^t)) - \delta (W(h^t, x^{t*}) - W_{-i}(h^t, x^{t*})). \quad (2)$$

3.0.0.6 Dynamic Second Price Auction.

The flow marginal contribution is a natural candidate for the net utility that each bidder should receive in each period t . We now construct a transfer price such that under the efficient allocation, each bidder's net payoff coincides with her flow marginal contribution. We then show that this pricing rule makes truth-telling incentive compatible in the dynamic mechanism.

The winning bidder i receives the object in period t . To match her net payoff to her flow marginal contribution, we must have:

$$m_i(h^t) = v_i(h^t) - p_i(h^t). \quad (3)$$

The remaining bidders, $j \neq i$, do not receive the object in period t and their transfer price must offset the flow marginal contribution:

$$m_j(h^t) = -p_j(h^t).$$

Consider first the efficient bidder i in period t . We expand the flow marginal contribution in (2) by noting that i is the efficient assignment and that another bidder, say k , would constitute the efficient assignment in the absence of bidder i :

$$m_i(h^t) = v_i(h_i^t) - v_k(h_k^t) - \delta (W_{-i}(h^t, i) - W_{-i}(h^t, k)). \quad (4)$$

The optimal assignment policy is without loss of generality a deterministic policy as a function of the history. We therefore replace the vector x^t by the assignment decision which determines the identity of the winning bidder. Thus, in (4), $W_{-i}(h^t, i)$ and $W_{-i}(h^t, k)$ represent the continuation value of the social program without i , conditional on the history h^t and the current assignment being i or k_{-i} respectively. We notice that with private values, the continuation value of the social program without i and conditional on h^t and giving the object to agent i in period t is simply equal to the value of the program conditional on h^t alone, or

$$W_{-i}(h^t, i) = W_{-i}(h^t).$$

The additional information generated by the assignment to agent i only pertains to agent i and hence has no value for the allocation problem once i is removed. We can therefore rewrite the flow marginal contribution of the winning agent i as:

$$m_i(h^t) = v_i(h_i^t) - (1 - \delta) W_{-i}(h^t).$$

The flow marginal contribution of i is therefore her expected flow value minus the delay in the accrual of the social benefit arising from the optimal assignment among all agents excluding agent i . It follows that the transfer price should simply be given by:

$$p_i^*(h^t) = (1 - \delta) W_{-i}(h^t), \quad (5)$$

which is the flow social opportunity cost of assigning the object today to agent i .

A similar analysis, based on the flow marginal contribution (4) leads to the determination of the transfer price for the losing bidders. Consider a bidder j who should not get the object in period t . Her flow utility is clearly zero in period t . Moreover, by the optimality of the index policy, the removal of alternative j from the set of possible allocations does not change the optimal assignment today. In consequence, the identity of the winning bidder does not depend on the presence of alternative j . In other words the efficient assignment to i will remain efficient after we remove j . As a result the flow marginal contribution of the losing bidder is zero, and we have:

$$p_j^*(h^t) = -m_j(h^t) = 0.$$

THEOREM 1 (DYNAMIC SECOND PRICE AUCTION).

The socially efficient allocation rule \mathbf{x}^ is ex post incentive compatible in the dynamic direct mechanism with the payment rule \mathbf{p}^* where:*

$$p_j^*(h^t) = \begin{cases} (1 - \delta) W_{-j}(h^t) & \text{if } x_j^{t*} = 1, \\ 0 & \text{if } x_j^{t*} = 0. \end{cases}$$

PROOF. By the unimprovability principle, it is sufficient to prove that if an agent receives in all future periods her marginal contribution as her continuation value, then truth-telling is incentive compatible for an agent in period t . Suppose then that at h^t , it is socially efficient to assign the object to agent i and suppose that all agents except i report truthful. The incentive constraint for agent i is then given by:

$$v_i(h_i^t) - p_i^*(h^t) + \delta M_i(h^t, i) \geq \delta M_i(h^t, j) \quad (6)$$

for some $j \neq i$. By the determination of the transfer price p_i^* , it follows that (6) can be written as follows

$$M_i(h^t) \geq \delta M_i(h^t, j) \quad (7)$$

and by definition of the marginal contribution, we can rewrite (7) in terms of the social value functions:

$$W(h^t) - W_{-i}(h^t) \geq \delta(W(h^t, j) - W_{-i}(h^t, j)),$$

and expanding by $v_i(h_i^t)$, we have

$$\begin{aligned} & W(h^t) - W_{-i}(h^t) \\ & \geq \\ & v_i(h_i^t) + \delta W(h^t, j) - v_i(h_i^t) - \delta W_{-i}(h^t, j), \end{aligned}$$

but then the result is:

$$W(h^t) - W(h^t, j) \geq W_{-i}(h^t) - W_{-i}(h^t, j). \quad (8)$$

The inequality (8) follows from the fact that the size of the loss due to a suboptimal choice j (weakly) increases in the number of alternatives present.

For the case of an inefficient agent j in period t , we have

$$M_j(h^t) \geq v_j(h_j^t) - p_j(h^t) + \delta M_j(h^t, j). \quad (9)$$

As the transfer price is independent of the report of agent j , and given by (5), we can rewrite (9) as follows

$$M_j(h^t) \geq v_j(h_j^t) - (1 - \delta)W_{-j}(h^t) + \delta M_j(h^t, j).$$

After replacing the marginal contributions by the social value functions, we have

$$\begin{aligned} & W(h^t) - W_{-j}(h^t) \\ & \geq \\ & v_j(h_j^t) - (1 - \delta)W_{-j}(h^t) + \delta(W(h^t, j) - W_{-j}(h^t, j)). \end{aligned}$$

But as $W_{-j}(h^t, j) = W_{-j}(h^t)$, the terms involving the value functions of $-j$ all drop out and we are left with

$$W(h^t) \geq v_j(h_j^t) + \delta W(h^t, j), \quad (10)$$

which is a valid inequality since j is by hypothesis not the efficient choice in period t . \square

The incentive compatible pricing rule has a few interesting implications. First, we observe that in the case of two bidders, the formula for the dynamic second price reduces to the static solution. If we remove one bidder, the social program has no other choice but to always assign it to the remaining bidder. But then, the expected value of that assignment policy is simply equal to the expected value of the object for bidder j in period t by the martingale probability of the Bayesian posterior. In other words, the transfer is equal to the current expected value of the next best competitor. With more than two bidders, the social program without bidder i will contain an option value due to the possibility of assigning the object to the more favorable bidder. In consequence the social opportunity cost is higher than the highest expected valuation among the remaining bidders.

Second, we observe that the transfer price of the winning bidder is independent of her own information about the object. This means, that for any number of periods in which the ownership of the object does not change, the transfer price will stay constant as well, even though the valuation of the object by the winning bidder may undergo substantial change.

The design of the transfer price pursued the objective to match the flow marginal contribution of every agent in every period. The determination of the transfer price is based

exclusively on the reported signals of the other agents, rather than their true signals. For this reason, truthtelling is not only Bayesian incentive compatible, but ex post incentive compatible, if we qualify ex post to mean conditional on all signals received up to and including period t .

An important insight from the static analysis of the private value environment is the fact that incentive compatibility can be guaranteed in weakly dominant strategies. This strong result does not carry over into the dynamic setting due to the interaction of the strategies. In a dynamic setting, each agent can condition her strategy on the past reports of the other agents. In particular, the strategy of truthtelling after all histories fails to be a weakly dominant strategy as it removes the ability to respond to past announcements. Yet our argument shows that the weaker condition of ex post incentive compatibility can be satisfied.

The vital assumption in the dynamic auction model pertained to the flow of information: Each bidder receives additional private information in period $t + 1$ if and only if she received the object in period t . This is the essential informational hypothesis in multi-armed bandit framework. Yet we might be interested in a setting in which each bidder may learn more about the value of the object even in periods in which she does not control the object. The incentive analysis is again based on the flow marginal contribution. But once we leave the bandit framework, then some losing bidders may have to pay a positive price even in periods in which they do not receive the object. Consider a losing bidder j and suppose that the removal of bidder j would change the efficient assignment policy from agent i to agent k . The flow contribution of the losing bidder j would now be equal to:

$$m_j(h^t) = v_i(h_i^t) - v_k(h_k^t) + \delta(W_{-j}(h^t, i) - W_{-j}(h^t, k)) < 0.$$

In other words, if the presence of j changes the efficient assignment policy, then this leads to an externality cost created by agent j and hence strictly positive transfer prices even in periods in which agent j does not receive the object.

4. GENERAL PRIVATE VALUE ENVIRONMENT

In this section we extend the private value environment from a single unit auction to a general allocation model. In addition, we substantially generalize the statistical model of information. The net expected flow utility of agent i in period t is now determined by the (flow) allocation $a^t \in A$, the private history h_i^t and the transfer price p_i^t :

$$v_i(a^t, h_i^t) - p_i^t.$$

The utility function v_i represents the expected utility to agent i from an allocation a^t given the private information h_i^t . The set of available allocations is given by a compact and time invariant set A . The private signal of agent i in period $t + 1$ is generated according to a conditional distribution function:

$$s_i^{t+1} \sim G_i(s_i^{t+1} | a^t, h_i^t).$$

We generalize the information flow by allowing the signal s_i^{t+1} of agent i in period $t + 1$ to be dependent on the current allocative decision a^t and the entire past history of private signals received by agent i . The allocation rule for the direct

mechanism is now given by

$$x^t : \widehat{H}^t \rightarrow \Delta(A),$$

and the transfer rules are given by:

$$p^t : \widehat{H}^t \rightarrow \mathbb{R}^N.$$

As before, we denote the socially efficient policy by $\mathbf{x}^* = \{x^{t*}\}_{t=0}^\infty$. The direct dynamic mechanism $\mathcal{M} = \langle \mathbf{x}^*, \mathbf{p}^*, \widehat{\mathbf{H}} \rangle$ extends the Vickrey-Clark-Groves mechanism to general intertemporal environments by the marginal contribution argument as developed earlier in the context of the single unit allocation problem.

THEOREM 2 (DYNAMIC VCG MECHANISM).

The socially efficient allocation rule $\{x^\}$ is ex post incentive compatible with the payment rule p^* :*

$$p_i^{t*}(x^*(h^t), h_{-i}^t) = m_i(h^t) - v_i(x^*(h^t), h_i^t). \quad (11)$$

PROOF. The basic idea of the proof generalizes the marginal contribution argument in Theorem 1. By the unimprovability principle, it suffices to prove that if agent i will receive as her continuation value her marginal contribution, then truthtelling is incentive compatible for agent i in period t , or:

$$\begin{aligned} & v_i(x^*(h^t), h_i^t) - p_i^t(x^*(h^t), h_{-i}^t) + \delta M_i(x^*(h^t), (h_{-i}^t)) \\ & \geq \\ & v_i(a, h_i^t) - p_i^t(a, h_{-i}^t) + \delta M_i(h^t, a), \end{aligned}$$

for all i, t and $a \in A$. By construction of the transfer price, the lhs of (12) represents the marginal contribution of agent i . Similarly, we can express the continuation marginal contribution $M_i(h^t, a)$ in terms of the values of the different social programs:

$$\begin{aligned} & W(h^t) - W_{-i}(h^t) \\ & \geq \\ & v_i(a, h_i^t) - p_i^t(a, h_{-i}^t) + \delta (W(h^t, a) - W_{-i}(h_{-i}^t, a)). \end{aligned} \quad (13)$$

By construction of the transfer price, we can represent the price that agent i would have to pay if allocation a were to be chosen in terms of the marginal contribution if the reported history h_i^t were the true signal received by agent i . By construction, we have as in (11):

$$p_i^{t*}(x^*(h^t), h^t) = m_i(h^t) - v_i(x^*(h^t), h_i^t).$$

The flow marginal contribution of agent i is given by

$$\begin{aligned} m_i(h_i^t, h_{-i}^t) &= \sum_{j=1}^I v_j(a, h_i^t, h_{-i}^t) - \sum_{j \neq i} v_j(x_{-i}^*, h_{-i}^t) \\ & \quad + \delta (W_{-i}(h_{-i}^t, a) - W_{-i}(h_{-i}^t, x_{-i}^*)). \end{aligned} \quad (14)$$

so that the price is given by:

$$\begin{aligned} p_i^t(h^t) &= \sum_{j \neq i} v_j(x_{-i}^*, h_{-i}^t) - \sum_{j \neq i} v_j(a, h_i^t, h_{-i}^t) \\ & \quad + \delta (W_{-i}(h_{-i}^t, x_{-i}^*) - W_{-i}(h_{-i}^t, a)). \end{aligned} \quad (15)$$

We can now insert the prices into (13) to obtain:

$$\begin{aligned} & W(h^t) - W_{-i}(h_{-i}^t) \geq \\ & v_i(a, h_i^t) - \\ & \sum_{j \neq i} v_j(x_{-i}^*, h_{-i}^t) - \sum_{j \neq i} v_j(a, h_i^t, h_{-i}^t) + \delta W_{-i}(h_{-i}^t, x_{-i}^*) \\ & \quad + \delta ((W(h^t, a))). \end{aligned}$$

But now we can reconstitute the entire expression in terms of the social value of the program with and without agent i and we are lead to the final inequality:

$$W(h^t) - W_{-i}(h_{-i}^t) \geq W(h^t, a) - W_{-i}(h_{-i}^t),$$

where the later is true by the optimality of x^* at h^t . \square

We observe that the pricing rule (11) for agent i depends on the report of agent i only through the determination of the social allocation which already appeared as a prominent feature in the static environment. Theorem 2 gives a general characterization of the transfer prices. In specific environment (such as a public good provision model), we can then gain additional insights into the structure of the efficient transfer prices by analyzing how the policies would change with the addition or removal of an arbitrary agent i .

5. CONCLUSION

This paper suggest the construction of a direct dynamic mechanism in private value environments with transferable utility. The design of the monetary transfers relies on the notions of marginal contribution and flow marginal contribution. These notions allow us to transfer the insights of the Vickrey-Clark-Groves mechanism from a static environment to general dynamic settings. In the case of the sequential allocation of a single indivisible object, we show that the notion of marginal contribution and its relationship to the social program allow us to give explicit solutions of the monetary transfers in each period.

Many interesting questions are left open. The dynamic mechanism considered here satisfies the incentive compatibility and individual participation constraints in every period. In particular, we do not require that the monetary transfer satisfy a balanced budget condition in every period. The recent analysis of Athey and Segal (2006) suggests that a sequential version of AGV mechanism might be able to achieve budget balancing in every period as well. This paper is silent on the issue of revenue maximizing mechanisms. In order to make progress in that direction, a characterization of implementable allocations in dynamic setting will first be necessary. Finally, we restricted our attention to private value environments. A recent literature, beginning with Maskin (1992) and Dasgupta and Maskin (2000) showed how to extend the VCG mechanism to interdependent value environments. In dynamic settings, the single crossing condition will then typically involve a dynamic element which will introduce some complications. These tasks are left for future research.

6. ADDITIONAL AUTHORS

Additional authors: Juuso Valimaki (Department of Economics, Helsinki School of Business and University of Southampton) email: juuso.valimaki@hse.fi)

7. REFERENCES

- ATHEY, S., AND I. SEGAL (2006): “An Efficient Dynamic Mechanism,” Harvard University and Stanford University.
- BAPNA, A., AND T. WEBER (2005): “Efficient Dynamic Allocation with Uncertain Valuations,” Stanford University.
- BERGEMANN, D., AND J. VÄLIMÄKI (2003): “Dynamic Common Agency,” *Journal of Economic Theory*, Forthcoming.
- (2006): “Dynamic Price Competition,” *Journal of Economic Theory*, 127, 232–263.
- CLARKE, E. (1971): “Multipart Pricing of Public Goods,” *Public Choice*, 8, 19–33.
- DASGUPTA, P., AND E. MASKIN (2000): “Efficient Auctions,” *Quarterly Journal of Economics*, 115, 341–388.
- FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): “Planning under Incomplete Information and the Ratchet Effect,” *Review of Economic Studies*, 52, 173–191.
- GITTINS, J. (1989): *Allocation Indices for Multi-Armed Bandits*. London, Wiley.
- GROVES, T. (1973): “Incentives in teams,” *Econometrica*, 41, 617–631.
- MASKIN, E. (1992): “Auctions and Privatization,” in *Privatization: Symposium in Honor of Herbert Giersch*, ed. by H. Siebert, pp. 115–136. J.C.B. Mohr, Tuebingen.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- WHITTLE, P. (1982): *Optimization Over Time*, vol. 1. Wiley, Chichester.