Dynamic Marginal Contribution Mechanism*

Dirk Bergemann† Juuso Välimäki‡

First Version: September 2006

Current Version: June 2007

*We thank the editor, Eddie Dekel, and two anonymous referees for many helpful comments. The current paper is a major revision and supersedes "Efficient Dynamic Auctions" (2006). We are grateful to Larry Ausubel, Jerry Green, Paul Healy, John Ledyard, Michael Ostrovsky, and David Parkes for many informative conversations. The authors gratefully acknolwedge financial support through the National Science Foundation Grants CNS 0428422 and SES 0518929 and the Yrjö Jahnsson's Foundation, respectively. We thank seminar participants at DIMACS, Ohio State University, University of Iowa, University of Madrid and the University of Maryland for valuable comments.

†Department of Economics, Yale University, New Haven, U.S.A., dirk.bergemann@yale.edu.

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

Abstract

We consider truthful implementation of the socially efficient allocation in a dynamic private value environment in which agents receive private information over time. We propose a suitable generalization of the Vickrey-Clarke-Groves mechanism, based on the marginal contribution of each agent. In the marginal contribution mechanism, the ex post incentive and ex post participations constraints are satisfied for all agents after all histories. It is the unique mechanism satisfying ex post incentive, ex post participation and efficient exit conditions.

We develop the marginal contribution mechanism in detail for a sequential auction of a single object in which each bidders learn over time her true valuation of the object. We show that a modified second price auction leads to truthtelling.

JEL CLASSIFICATION: C72, C73, D43, D83.

KEYWORDS: Vickrey-Clarke-Groves Mechanism, Pivot Mechanism, Ex Post Equilibrium, Marginal Contribution, Multi-Armed Bandit, Bayesian Learning.

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 general allocation problems in private value environments. The central idea behind the Vickrey-Clarke-Groves mechanism is to convert the indirect utility of each agent i into the social objective function - up to a term which is a constant from the point of view of agent i. In the class of transfer payments which accomplish this internalization of the social objective, the pivot mechanism (due to Green and Laffont (1977)) requires the transfer payment of agent i to match her externality cost on the remaining agents . 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 (or the pivot mechanism) to dynamic environments with private information. We design an intertemporal sequence of transfer payments which allows each agent to receive her flow marginal contribution in every period. In other words, after each history, the expected transfer that each player must pay coincides with the dynamic externality cost that she imposes on other agents. In consequence, each agent is willing to truthfully report her information in every period.

We consider a general intertemporal model in discrete time and with a common discount factor. The private information of each agent in each period is her perception of her future payoff path conditional on the realized information and allocations. We assume throughout that the information that the agents have is statistically independent across agents. At the reporting stage of the direct mechanism, each agent reports her information. The planner then calculates the efficient allocation given the reported information. The planner also calculates for each i the optimal allocation when agent i is excluded from the mechanism. The total expected discounted payment of each agent is set equal to the externality cost imposed on the other agents in the model. In this manner, each player receives as her payment her marginal contribution to the social welfare in every conceivable continuation mechanism.

With transferable utilities, the social objective is simply to maximize the expected discounted sum of the individual utilities. Since this is essentially a dynamic programming problem, the solution is by construction time consistent. In consequence, the dynamic marginal contribution 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)). Furthermore since marginal contributions are positive by definition, dynamic marginal contribution mechanism induces all productive agents to participate in the mechanism after all histories.

In contrast to the static environment, the thruthtelling 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. The reports of other agents in period t determine the allocation for that period. In the dynamic game, the agents' intertemporal payoffs depend on the expected future allocations and transfers as well. As a result, the agents' current reports need not maximize their current payoff. Since dishonest reports distort current and future allocations in different ways, agent i's optimal report may depend on the reports of others. Nevertheless, truthful reporting is optimal for all realizations of other players' private information as long as their reports are truthful.

In the intertemporal environment there is a multiplicity in transfer schemes that support the same incentives as the marginal contribution mechanism. In particular, the monetary transfers necessary to induce the efficient action in period t may always become due at some later period s, provided that the transfers maintain a constant net present value. We say that a mechanism supports efficient exit if an agent who ceases to affect the current and future allocations also ceases to receive transfers. Our second characterization result shows that the marginal contribution mechanism is the unique mechanism that satisfies ex post incentive, ex post participation and efficient exit conditions.

The basic idea of the marginal contribution mechanism is first explored in the context of a scheduling problem where a set of privately informed bidders compete for the services of a central facility over time. This class of problems is perhaps the most natural dynamic analogue of static single unit auctions. Besides the direct revelation mechanism, we also show that there is dynamic ascending price auction implements the efficient allocation when each bidder has a single task that can be completed in a single period. Unfortunately in the case of multiple tasks per bidder, the ascending price auction and other standard auction formats fail to be efficient. In contrast, the marginal contribution mechanism continues to support the efficient allocation. This gap calls for a more complete understanding of bidding mechanisms expressible in the willingness to pay in intertemporal environments.

In section 5, we use the marginal contribution mechanism to derive the optimal dynamic auction format for a model where bidders learn their valuations for a single object over time. The Bayesian learning framework constitutes a natural setting to analyze the repeated allocation of an object or a license over time. The key assumption in the learning setting is that only the current winner gains additional 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, often referred to as multi-armed bandit problem, is that the structure of the socially efficient program is well understood. As the monetary transfers allow each agent to capture her marginal 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.

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 (2007) consider a similar model as ours. Their focus is on mechanisms that are budget balanced whereas our paper focuses on mechanisms where the participation constraint is satisfied in each period. In the last section of their paper, Athey and Segal (2007) show that in infinite horizon problems, participation constraint can be satisfied using repeated game strategies if the discount factors are high enough. The same repeated game strategies are employed by **?** with a focus on repeated bilateral trade. 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 our results work equally well for the finite horizon model as for the infinite one. Cavallo, Parkes, and Singh (2006) consider a dynamic Markovian model and derive a sequence of Groves like payments which guarantee interim incentive compatibility but not interim participation constraints. 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. Friedman and Parkes (2003) and Parkes and Singh (2003) consider a specific dynamic environments with randomly arriving and departing agents in a finite horizon model. A dynamic version of the VCG mechanism, termed "delayed VCG" is suggested to guarantee interim incentive compatibility but again does not address interim participation constraints. 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 2 sets up the general model, introduces the notion of a dynamic mechanism and defines the equilibrium concept. Section 3 introduces the main concepts in a simple example. Section 4 analyzes the marginal contribution mechanism in the general environment. We also show that the marginal contribution mechanism is the unique dynamic mechanism which satisfies ex post incentive compatibility, ex post participation and efficient exit condition. Section 5 analyzes the implications of the general model for a licensing auction with learning. Section 6 concludes.

2 Model

Payoffs We consider an environment with private and independent values in a discrete time, infinite horizon model. The flow utility of agent $i \in \{1, 2, ..., I\}$ in period $t \in \{0, 1, 2, ...\}$ is determined by the past and present allocations and a monetary transfer. The allocation space A_t in period t is assumed to be a compact space and an element of the allocation space is denoted by $a_t \in A_t$. An allocation profile until period t is denoted by:

$$a^t = (a_0, a_1, ..., a_t) \in A^t = \bigtimes_{s=0}^{t} A_s.$$

The allocation profile a^t gives rise to a flow utility $\omega_{i,t}$:

$$\omega_{i,t} : A^t \to \mathbb{R}_+,$$

and we assume that the flow payoff in t is quasi-linear in the transfer $p_{i,t}$ and given by:

$$\omega_{i,t}\left(a^t\right) + p_{i,t}.$$

By allowing the flow utility $\omega_{i,t}$ of agent i in period t to depend on the past allocations, the model can encompass learning-by-doing and habit formation.[1]

All agents discount the future with a common discount factor δ, $0 < \delta < 1$.

Information The family of flow payoffs of agent i over time

$$\{\omega_{i,t}(\cdot)\}_{t=0}^{\infty}$$

is a stochastic process which is privately observed by agent i. In an incomplete information environment, the private information of agent i in period t is her information about her current (and future) valuation profile. The type of agent i in period t is

[1] An alternative (and largely equivalent) approach would allow the past consumption to influence the distribution of future random utility.

therefore simply her information about her current (and future) valuation profile. It is convenient to model the private information of agent i in period t about his current and future valuations as being represented by a filtration $\{\mathcal{F}_{i,t}\}_{t=0}^{\infty}$ on a probability space $(\Omega_i, \mathcal{F}_i, \mathbb{P}_i)$. An element ω_i of the sample space Ω_i is the infinite sequence of valuation functions $\omega_i = (\omega_{i,0}, \omega_{i,1}, ...)$. We take Ω_i to be the set of all infinite sequences of uniformly bounded and continuous functions. In other words, there exists $K > 0$ such for all i, all t and all a^t, $\omega_{i,t}(a^t)$ is continuous in a^t and $\omega_{i,t}(a^t) < K$. The σ-algebra $\mathcal{F}_i$ represents the family of measurable events in the sample space Ω_i and $\mathbb{P}_i$ is the probability measure on Ω_i. The filtration $\{\mathcal{F}_{i,t}\}_{t=0}^{\infty}$ is an increasing family of sub $\sigma-$algebras of $\mathcal{F}_i$. Intuitively, the filtration $\mathcal{F}_{i,t}$ is the information about $\omega_i \in \Omega_i$ available to agent i at time t. We follow the usual convention to augment the filtration $\mathcal{F}_{i,t}$ by all subsets of zero probability of $\mathcal{F}_i$. We denote a typical element of the filtration $\mathcal{F}_{i,t}$ in period t by

$$\omega_i^t \in \mathcal{F}_{i,t}.$$

The element $\omega_i^t \in \mathcal{F}_{i,t}$ thus represents the information of agent i about her current and future valuations function $(\omega_{i,t}, \omega_{i,t+1}, ...)$ at time t.[2] We observe that the information model is sufficiently rich to accommodate random entry and exit of the agents over time. In particular, for any $k \in \mathbb{N}$, the first k utility functions *or* all but the first k utility functions can be equal to zero utilities.

Finally, in the dynamic model, the independent value condition is guaranteed by assuming that the prior probabilities $\mathbb{P}_i$ and the filtrations $\mathcal{F}_{i,t}$ are independent across i.

[2] An common alternative model of private values in static (and dynamic models) is to assign each individual a utility function $u_i(a, \omega_i)$ which depends on the allocation and a privately observed random variable ω_i. In our specification, we take the utility function itself to be a random function. This direct approach via random utilities is useful for the characterization results in Theorem 1 and 2.

Histories In the presence of private information we have to distinguish between private and public histories. The private history of agent i in period t is the sequence of private information received by agent i until period t, or $h_{i,t} = \left(\omega_i^0, ..., \omega_i^{t-1}\right)$. The set of possible private histories in period t is denoted by $H_{i,t}$. In the dynamic direct mechanism to be defined shortly, each agent i is asked to report her current information in every period t. The report $r_{i,t}$ of agent i, truthful or not, is an element of the filtration $\mathcal{F}_{i,t}$ for every t. The public history in period t is then a sequence of reports until t and allocative decisions until period $t-1$, or $h_t = (r_0, a_0, r_1, a_1, ...r_{t-1}, a_{t-1}, r_t)$, where each $r_s = (r_{1,s}, ..., r_{I,s})$ is a report profile of the I agents. The set of possible public histories in period t is denoted by H_t. The sequence of reports by the agents is part of the public history and hence the past and current reports of the agents are observable to each one of the agents.

Mechanism A *dynamic direct mechanism* asks every agent i to report her information ω_i^t in every period t. The report $r_{i,t}$, truthful or not, is an element of the filtration $\mathcal{F}_{i,t}$ for every i and every t. A dynamic direct mechanism is then represented by a family of allocative decisions:

$$a_t : H_t \to \Delta\left(A_t\right),$$

and monetary transfer decisions:

$$p_t : H_t \to \mathbb{R}^I,$$

such that the decisions in period t respond to the reported information of all agents in period t. A dynamic direct mechanism $\mathcal{M}$ is then defined by

$$\mathcal{M} = \left\{\{H_t\}_{t=0}^{\infty}, \{a_t\}_{t=0}^{\infty}, \{p_t\}_{t=0}^{\infty}\right\}$$

such that the decisions $\{a_t, p_t\}_{t=0}^{\infty}$ are adapted to the histories $\{H_t\}_{t=0}^{\infty}$.

Social Efficiency In an environment with quasi-linear utility the socially efficient policy is obtained by maximizing the utilitarian welfare criterion, namely the expected discounted sum of valuations. Given a history h_t in period t under truthful reporting, the socially optimal program can be written simply as

$$W(h_t) = \max_{\{a_s\}_{s=t}^{\infty}} \mathbb{E}\left\{\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i=1}^{I} \omega_{i,s}(a^s)\right\}.$$

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

$$W(h_t) = \max_{a_t} \mathbb{E}\left\{\sum_{i=1}^{I} \omega_{i,t}\left(a^t\right) + \delta \mathbb{E} W(h_t, a_t)\right\},$$

where $W(h_t, a_t)$ represents the *optimal continuation value* conditional upon history h_t and allocation a_t. We note that the optimal continuation value $W(h_t, a_t)$ is well defined for all feasible allocations $a_t \in A_t$ The socially efficient policy is denoted by $\mathbf{a}^* = \{a_t^*\}_{t=0}^{\infty}$. In the remainder of the paper we focus attention on direct mechanisms which truthfully implement the socially efficient policy $\mathbf{a}^*$.

The social externality cost of agent i is determined by the optimal continuation plan in the absence of agent i. It is therefore useful to define the value of the social program after removing agent i from the set of agents:

$$W_{-i}(h_t) = \max_{\{a_s\}_{s=t}^{\infty}} \mathbb{E} \sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \neq i} \omega_{i,s}(a^s).$$

The *marginal contribution* $M_i(h_t)$ of agent i at history h_t is naturally defined by:

$$M_i(h_t) \triangleq W(h_t) - W_{-i}(h_t). \tag{1}$$

The marginal contribution is the change in social value due to the addition of agent i.

Equilibrium In a dynamic direct mechanism, a reporting strategy for agent i in period t is a mapping from the private and public history into the filtration $\mathcal{F}_{i,t}$:

$$r_{i,t} : H_{i,t} \times H_{t-1} \to \mathcal{F}_{i,t}.$$

Each agent i reports her information on the current and future valuation process that she has gathered up to period t. In a dynamic direct mechanism, a reporting strategy for agent i in period t is simply a mapping from the private and public history into an element of the filtration $\mathcal{F}_{i,t}$ in period t:

$$r_{i,t} : H_{i,t} \times H_{t-1} \rightarrow \mathcal{F}_{i,t}.$$

In other words, each agent i reports her information on her entire valuation process that she has gathered up to period t. For a given mechanism $\mathcal{M}$, the expected payoff for agent i from reporting strategy $r_{\mathbf{i}} = \{r_{i,t}\}_{t=0}^{\infty}$ given that the others agents are reporting $r_{-\mathbf{i}} = \{r_{-i,t}\}_{t=0}^{\infty}$ is given by

$$\mathbb{E}\sum_{t=0}^{\infty}\delta^t \left[\omega_{i,t}\left(a^t\left(h_{t-1}, r_{i,t}, r_{-i,t}\right)\right) + p_{i,t}\left(h_{t-1}, r_{i,t}, r_{-i,t}\right)\right].$$

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

$$V_i(h_{t-1}, h_{i,t}) = \max_{r_{i,t}\in\mathcal{F}_{i,t}} \mathbb{E}\left\{\omega_{i,t}\left(a^t\left(h_{t-1}, r_{i,t}, r_{-i,t}\right)\right) + p_{i,t}\left(h_{t-1}, r_{i,t}, r_{-i,t}\right) + \delta V_i\left(h_t, a_t, h_{i,t+1}\right)\right\}.$$

The profile of allocative decisions $a^t\left(h_{t-1}, r_{i,t}, r_{-i,t}\right)$ is determined by the past history h_{t-1} as it includes the past choices $(a_0, ...a_{t-1})$ and the current choice a_t is determined by the history h_{t-1} and the current reports r_t. The value function $V_i\left(h_t, a_t, h_{i,t+1}\right)$ represents the continuation value given the current history h_t, the current action a_t and tomorrow's private history $h_{i,t+1}$. We say that a dynamic direct mechanism $\mathcal{M}$ is *interim incentive compatible*, if for every agent and every period, truthtelling is a best response given that all other agents report truthfully. In terms of the value function, it means that a solution to the dynamic programming equations is to report truthfully $r_{i,t} = \omega_i^t$:

$$\omega_i^t \in \arg\max_{r_{i,t}\in\mathcal{F}_{i,t}} \mathbb{E}\left\{\omega_{i,t}\left(a^t\left(h_{t-1}, r_{i,t}, \omega_{-i}^t\right)\right) + p_{i,t}\left(h_{t-1}, r_{i,t}, \omega_{-i}^t\right) + \delta V_i\left(h_t, a_t, h_{i,t+1}\right)\right\}.$$

We say that $\mathcal{M}$ is *periodic ex post incentive compatible* if truthtelling is a best response regardless of the signal realization of the other agents:

$$\omega_i^t \in \underset{r_{i,t} \in \mathcal{F}_{i,t}}{\arg\max} \left\{ \omega_{i,t} \left(a^t \left(h_{t-1}, r_{i,t}, \omega_{-i}^t \right) \right) + p_{i,t} \left(r_{i,t}, \omega_{-i}^t \right) + \delta \mathbb{E} V_i \left(h_t, a_t, h_{i,t+1} \right) \right\},$$

for all $\omega_{-i}^t \in \mathcal{F}_{-i,t}$. In the dynamic context, the notion of ex post incentive compatibility has to be qualified by periodic as it is ex post with respect to all signals received in period t, but not ex post with respect to signals arriving after period t. The periodic qualification is natural in the dynamic environment as agent i may receive information at some later time $s > t$ such that in retrospect she would wish to change the allocation choice in t and hence her report in t.

Finally we consider the interim participation constraint of each agent. If agent i were to irrevocably leave the mechanism in period t, then an efficient mechanism would prescribe the efficient policy a_{-i}^* for the remaining agents. By leaving the mechanism, agent i may still enjoy the value of allocative decisions supported by the remaining agents. We define the value of agent i from being outside the mechanism as:

$$O_i(h_{t-1}, h_{i,t}) = \max_{r_{i,t} \in \mathcal{F}_{i,t}} \mathbb{E} \left\{ \omega_{i,t} \left(a_{-i}^t \left(h_{t-1}, r_{-i,t} \right) \right) + \delta O_i \left(h_t, a_{-i,t}, h_{i,t+1} \right) \right\}.$$

By being outside of the mechanism, the value of agent i is generated from the allocative decision of the remaining agents and naturally agent i neither influences their decision nor does she receive monetary payments. The interim participation constraint of agent i requires that for all h_t:

$$V_i(h_{t-1}, h_{i,t}) \geq O_i(h_{t-1}, h_{i,t}).$$

Again, we can strengthen the interim participation constraint to periodic ex post participation constraints for all h_t and ω^t :

$$\begin{aligned} & \omega_{i,t} \left(a^t \left(h_{t-1}, \omega^t \right) \right) + p_{i,t} \left(h_{t-1}, \omega^t \right) + \delta \mathbb{E} V_i \left(h_t, a_t, h_{i,t+1} \right) \backslash \\ \geq \; & \omega_{i,t} \left(a_{-i}^t \left(h_{t-1}, \omega_{-i}^t \right) \right) + \delta \mathbb{E} O_i \left(h_t, a_{-i,t}, h_{i,t+1} \right). \end{aligned}$$

The periodic ex post participation constraint requires that for all possible signal profiles of the remaining agents and induced allocations, agent i would prefer to stay in the mechanism rather than leave the mechanism. For the remainder of the text we shall refer to periodic ex post constraints simply as ex post constraints.

3 Scheduling: An Example

We begin the analysis with a class of scheduling problems. The scheduling model is kept deliberately simple to illustrate the insights and results which are then established for general environments in the subsequent sections.

We consider the problem of allocating time to use a central facility among competing agents. Each agent has a private valuation for the completion of a task which requires the use of the central facility. The facility has a capacity constraint and can only complete one task per period. The cost of delaying any task is given by the discount rate $\delta < 1$. The agents are competing for the right to use the facility at the earliest available time. The objective of the social planner is to sequence the tasks over time so as to maximize the sum of the discounted utilities.

We denote by $\omega_{i,t}(a^t)$ the private valuation for bidder $i \in \{1, ..., I\}$ in period t. The prior probability over valuation functions $\{\omega_{i,t}(\cdot)\}_{t=0}^{\infty}$ is given $\mathbb{P}_i$. An allocation policy in this setting is a sequence of choices $a_t \in \{0, 1, ..., I\}$, where a_t denotes the bidder chosen in period t. We allow for $a_t = 0$ and hence the possibility that no bidder is selected in t. Each agent has only one task to complete and the value $\omega_i \in \mathbb{R}_+$ of the task is constant over time and independent of the realization time (except for discouting). The utility function $\omega_{i,t}(\cdot)$ for bidder i from an allocation policy a^t is represented by:

$$\omega_{i,t}\left(a^t\right) = \begin{cases} \omega_i & \text{if} \quad a_t = i \text{ and } a_s \neq i \text{ for all } s < t, \\ 0 & \text{if} \quad \text{otherwise.} \end{cases} \tag{2}$$

For this scheduling model we find the marginal contribution of each agent and then derive the associated marginal contribution mechanism. We also show that a natural indirect mechanism, a dynamic bidding mechanism, will lead to the efficient scheduling of tasks over time with the same flow utilities. Finally we extend the scheduling model to allow each agent to have multiple tasks. In this slightly more general setting, the dynamic bidding mechanism fails to lead to an efficient allocation, but the marginal contribution mechanism continues to implement the efficient allocation.

Marginal Contribution We determine the marginal contribution of bidder i by comparing the value of the social program with and without i. We can assume without loss of generality (after relabelling) that the valuations ω_i of the agents are ordered with respect to their identity i:

$$\omega_1 \geq \cdots \geq \omega_I \geq 0. \tag{3}$$

With stationary valuations ω_i for all i, the optimal policy is clearly given by assigning in every period the alternative j with the highest remaining valuation, or

$$a_t^* = t + 1, \text{ for all } t < I.$$

The descending order of the valuations of the bidders allows us to identify each alternative i with the time period $i + 1$ in which it is employed along the efficient path and so:

$$W(h_0) = \sum_{t=1}^{I} \delta^{t-1} \omega_t. \tag{4}$$

Similarly, the efficient program in the absence of bidder i assigns the bidders in descending order, but necessarily skips bidder i in the assignment process. In consequence it assign all bidders after i one period earlier relative to the program with bidder i:

$$W_{-i}(h_0) = \sum_{t=1}^{i-1} \delta^{t-1} \omega_t + \sum_{t=i}^{I-1} \delta^{t-1} \omega_{t+1}. \tag{5}$$

By comparing the social program with and without i, (4) and (5), respectively, we find that the assignments for bidders $j < i$ remain unchanged after i is removed, but that each bidder $j > i$ is allocated the slot one period earlier than in the presence of i. The marginal contribution of i form the point of view of period 0 is:

$$M_i\left(h_0\right) = W\left(h_0\right) - W_{-i}\left(h_0\right) = \sum_{t=i}^{I} \delta^{t-1}\left(\omega_t - \omega_{t+1}\right),$$

and from the point of view of period h_{i-1} along the efficient path is

$$M_i\left(h_{i-1}\right) = W\left(h_{i-1}\right) - W_{-i}\left(h_{i-1}\right) = \sum_{t=i}^{I} \delta^{t-1}\left(\omega_t - \omega_{t+1}\right). \tag{6}$$

The social externality cost of agent i is now established in a straightforward manner. At time $t = i - 1$, i will complete her task and hence realize a gross value of ω_i. The immediate opportunity cost is given by the next highest valuation ω_{i+1}. But this alone would overstate the externality cost, because in the presence of i all less valuable tasks will now be realized one period later. In other words, the insertion of i into the program leads to the realization of a relatively more valuable task in all subsequent periods The externality cost of agent i is hence equal to the value of the next valuable task ω_{i+1} minus the improvement in future allocations due the delay of all tasks by one period:

$$p_{i,t}\left(h_t\right) = -\omega_{i+1} + \sum_{t=i+1}^{I} \delta^{t-i}\left(\omega_t - \omega_{t+1}\right). \tag{7}$$

Since by construction (see (3)), we have $\omega_t - \omega_{t+1} \geq 0$, it follows that the externality cost of agent i in the intertemporal framework is less than in the corresponding single allocation problem where it would be ω_{i+1}. Consequently, we can rewrite (7) to:

$$p_{i,t}\left(h_t\right) = -\left(1 - \delta\right) \sum_{t=i}^{I} \delta^{t-i}\omega_{t+1}, \tag{8}$$

which simply states that the externality cost of agent i is the cost of delay, namely $(1 - \delta)$ imposed on the remaining and less valuable tasks. With the monetary transfers given by (7), Theorem 1 will formally establish that the marginal contribution mechanism leads to thruthtelling with ex post incentive and ex post participation constraints.

In the present scheduling model, the relevant private information for all agents arrives in period $t = 0$ and by the stationarity assumption is not changing over time. It would therefore be possible to assign the tasks completely in $t = 0$ and also assess the appropriate transfers in $t = 0$. In this static version of the direct mechanism each bidder reports her value of the task and the allocation is determined in the order of the reported valuations. The static VCG mechanism then has a truthful dominant strategy equilibrium if the payments are set with reference to (8) as:

$$p_i = -(1-\delta)\sum_{t=i}^{I}\delta^t\omega_{t+1}, \tag{9}$$

which equals the payments in the dynamic directed mechanism appropriately discounted as the payments are now assessed at $t = 0$.

Dynamic Bidding Mechanism In this scheduling problem a number of bidders compete for a scare resource, namely timely access to the central facility. It is then natural to ask whether the efficient allocation can be realized through a bidding mechanism rather than a direct revelation mechanism. We find a dynamic version of the ascending price auction where the contemporaneous use of the facility is auctioned. As a given task is completed, the number of effective bidders decreases by one. We can then use a backwards induction algorithm to determine the values for the bidders starting from a final period in which only a single bidder is left without effective competition.

Consider then an ascending auction in which all tasks except that of bidder I have been completed. Along the efficient path, the final ascending auction will occur at time $t = I - 1$. Since all other bidders have vanished along the efficient path at this point, bidder I wins the final auction at a price equal to zero. By backwards induction, we consider the penultimate auction in which the only bidders left are $I-1$ and I. As agent I can anticipate to win the auction tomorrow even if she were to loose it today, she is willing to bid at most

$$b_I(\omega_I) = \omega_I - \delta(\omega_I - 0), \tag{10}$$

namely the net value gained by winning the auction today rather than tomorrow. Naturally, a similar argument applies to bidder $I-1$, by dropping out of the competition today bidder $I-1$ would get a net present discounted value of $\delta\omega_{I-1}$ and hence her maximal willingness to pay is given by

$$b_{I-1}(\omega_{I-1}) = \omega_{I-1} - \delta(\omega_{I-1} - 0).$$

Since $b_{I-1}(\omega_{I-1}) \geq b_I(\omega_I)$, given $\omega_{I-1} \geq \omega_I$, it follows that bidder $I-1$ wins the ascending price auction in $t = I-2$ and receives a net payoff:

$$\omega_{I-1} - (1-\delta)\,\omega_I.$$

We proceed inductively and find that the maximal bid of bidder $I-k$ in period $t = I-k-1$ is given by:

$$b_{I-k}(\omega_{I-k}) = \omega_{I-k} - \delta\left(\omega_{I-k} - b_{I-(k-1)}\left(\omega_{I-(k-1)}\right)\right) \tag{11}$$

In other words, bidder $I-k$ is willing to bid as much as to be indifferent between being selected today and being selected tomorrow, when she would be able to realize a net valuation of $\omega_{I-k} - b_{I-(k-1)}$, but only tomorrow, and so the net gain from being selected today rather than tomorrow is:

$$\omega_{I-k} - \delta\left(\omega_{I-k} - b_{I-(k-1)}\right) \tag{12}$$

The maximal bid of bidder $I-(k-1)$ generates the transfer price of bidder $I-k$ and by solving (11) recursively with the initial condition given by (10), we find that the price in the ascending auction equals the externality cost in the direct mechanism given by (8). In this class of scheduling problems, the efficient allocation can therefore be implemented by a bidding mechanism.[3]

Multiple Tasks We end this section with a minor modification of the scheduling model to allow for multiple tasks. For this purpose it will suffice to consider an example with two bidders. The first bidder has an infinite series of single period tasks, each delivering a value of ω_1. The second bidder has only a single task with a value ω_2. The utility function of bidder 1 is thus given by

$$\omega_{1,t}\left(a^t\right) = \begin{cases} \omega_i & \text{if} \quad a_t = 1 \text{ for all } t, \\ 0 & \text{if} \quad \text{otherwise.} \end{cases}$$

whereas the utility function of bidder 1 is as described earlier by (2).

The socially efficient allocation in this setting either has $a_t = 1$ for all t if $\omega_1 \geq \omega_2$ or $a_0 = 2,\ a_t = 1$ for all $t \geq 1$ if $\omega_1 < \omega_2$. For the remainder of this example, we will assume that $\omega_1 > \omega_2$. Under this assumption the efficient policy will never complete the task of bidder 2. The marginal contribution of each bidder is:

$$M_1\left(h_0\right) = \left(\omega_1 - \omega_2\right) + \frac{\delta}{1-\delta}\omega_1 \tag{13}$$

and

$$M_2\left(h_0\right) = 0.$$

Along any efficient path h_t, we have $M_i\left(h_0\right) = M_i\left(h_t\right)$ for all i and we compute the social externality cost of agent 1, $p_{1,t}$ for all t, by using (13):

$$p_{1,t} = -\left(1-\delta\right)\omega_2.$$

[3]The nature of the recursive bidding strategies bears some similarity to the construction of the bidding strategies for multiple advertising slots in the keyword auction of Edelman, Ostrovsky, and Schwartz (2007). In the auction for search keywords, the multiple slots are differentiated by their probability of receiving a hit and hence generating a value. In the scheduling model here, the multiple slots are differentiated by the time discount associated with different access times.

The externality cost is again the cost of delay imposed on the competing bidder, namely $(1-\delta)$ times the valuation of the competing bidder. This accurately represent the social externality cost of agent 1 in every period even though agent 2 will never receive access to the facility.

We contrast the efficient allocation and transfer with the allocation resulting in the dynamic ascending price auction. For this purpose, suppose that the equilibrium path generated by the dynamic bidding mechanism would be efficient. In this case bidder 2 would never be chosen and hence would receive a net payoff of 0 along the equilibrium path. But this means that bidder 2 would be willing to bid up to ω_2 in every period. In consequence the first bidder would receive a net payoff of $\omega_1 - \omega_2$ in every period and her discounted sum of payoff would then be:

$$\frac{1}{1-\delta}(\omega_1 - \omega_2) < M_1. \tag{14}$$

But more important than the failure of the marginal contribution is the fact that the equilibrium will not support the efficient assignment policy. To see this, notice that if bidder 1 looses to bidder 2 in any single period, then the task of bidder 2 is completed and bidder 2 will drop out of the auction in all future stages. Hence the continuation payoff for bidder 1 from dropping out in a given period and allowing bidder 2 to complete his task is given by:

$$\frac{\delta}{1-\delta}\omega_1. \tag{15}$$

If we compare the continuation payoffs (14) and (15) respectively, then we see that it is beneficial for bidder 1 to win the auction in all periods if and only if

$$\omega_1 \geq \frac{\omega_2}{1-\delta},$$

but the efficiency condition is simply $\omega_1 \geq \omega_2$. It follows that for a large range of valuations, the outcome in the ascending auction is inefficient and will assign the object to bidder 2 despite the inefficiency of this assignment. The reason for the inefficiency is easy to detect in this simple setting. The forward looking bidders consider only their individual net payoffs in future periods. The planner on the other hand is interested in the level of gross payoffs in the future periods. As a result, bidder 1 is strategically willing and able to depress the future value of bidder 2 by letting bidder 2 win today to increase the future difference in the valuations between the two bidders. But from the point of view of the planner, the differential gains for bidder 1 is immaterial and the assignment to bidder 2 represents an inefficiency. The rule of the ascending price auction, namely that the highest bidder wins, only internalizes the individual *equilibrium payoffs* but not the social payoffs.

This small extension to multiple tasks shows that the logic of the marginal contribution mechanism can account for subtle intertemporal changes in the payoffs. On the other hand, common bidding mechanisms may not resolve the dynamic allocation problem in an efficient manner. Indirectly, it suggests that suitable indirect mechanisms have yet to be devised for scheduling and other sequential allocation problems.

4 Marginal Contribution Mechanism

In this section we construct the marginal contribution mechanism for the general model described in Section 2. We show that it is the unique mechanism which guarantees the ex post incentive constraints, the ex post participation constraints and an efficient exit condition.

4.1 Characterization

In the static Vickrey auction, the price of the winning bidder is equal to the highest valuation among the loosing bidders. The highest value among the remaining bidders represents the social opportunity cost of assigning the object to the winning bidder.

The marginal contribution of agent i is her contribution to the social value. At the same time, it is the information rent that agent i can secure for herself if the planner wishes to implement the socially efficient allocation. In a dynamic setting if agent i can secure her marginal contribution in every continuation game of the mechanism, then she should be able to receive the *flow marginal contribution* $m_i(h_t)$ in every period. The flow marginal contribution accrues incrementally over time and is defined recursively:

$$M_i(h_t) = m_i(h_t) + \delta M_i(h_t, a_t^*).$$

As in the notations of the value functions, $W(\cdot)$ and $V_i(\cdot)$ above, $M_i(h_t, a_t)$ represents the marginal contribution of agent i in the *continuation problem* conditional on the history h_t and the allocation a_t today. The flow marginal contribution can be expressed more directly using the definition of the marginal contribution (1) as

$$m_i(h_t) = W(h_t) - W_{-i}(h_t) - \delta(W(h_t, a_t^*) - W_{-i}(h_t, a_t^*)). \tag{16}$$

We can replace the value functions $W(h_t)$ and $W_{-i}(h_t)$ by the corresponding flow payoffs and continuation payoffs to get the flow marginal contribution of agent i:

$$m_i(h_t) = \sum_j \omega_{j,t}\left(a_t^*, a^{t-1}\right) - \sum_{j \neq i} \omega_{j,t}\left(a_{-i,t}^*, a^{t-1}\right) + \delta\left(W_{-i}(h_t, a_t^*) - W_{-i}\left(h_t, a_{-i,t}^*\right)\right). \tag{17}$$

If the presence of i, leads the designer to adopt the allocation a_t^*, then this preempts the preferred allocation $a_{-i,t}^*$ for all agents but i. To the extent that a decision for a_t^* irrevocably changes the value (including continuation value) of the remaining agents, the difference in value represents the social externality cost of agent i in period t. It is natural to suggest that a monetary transfer by agent i such that the resulting flow net utility matches her flow marginal contribution will lead agent i to dynamically internalize her social externalities, or

$$p_{i,t}^*(h_t) \triangleq m_i(h_t) - \omega_{i,t}\left(a_t^*, a^{t-1}\right), \tag{18}$$

and inserting (17) into (18) we have the transfer payment of the dynamic marginal contribution mechanism:

$$p^*_{i,t}(h_t) = \sum_{j \neq i} \omega_{j,t}\left(a^*_t, a^{t-1}\right) + \delta W_{-i}(h_t, a^*_t) - \sum_{j \neq i} \omega_{j,t}\left(a^*_{-i,t}, a^{t-1}\right) - \delta W_{-i}\left(h_t, a^*_{-i,t}\right). \tag{19}$$

The monetary transfers based on the marginal contribution of each agent i can support the efficient allocation in the resulting dynamic direct mechanism. We observe that the transfer pricing (19) 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 VCG environment. The monetary transfers $p^*_{i,t}(h_t)$ are always non-positive as the policy $a^*_{-i,t}$ is by definition an optimal policy to maximize the social value of all agents exclusive of i. It follows that in every period t the sum of the monetary transfers across all agents generates a weak budget surplus. Thus the design of the transfers $p^*_{i,t}$ guarantees that the designer does not face a budget deficit in any single period.

Theorem 1 (Dynamic Marginal Contribution Mechanism)
The dynamic marginal contribution mechanism $\{a^*_t, p^*_t\}_{t=0}^{\infty}$ *is efficient and satisfies ex post incentive and ex post participation constraints for all* i *and all* h_t.

Proof. 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:

$$\begin{aligned} & \omega_{i,t}\left(a^*_t, a^{t-1}\right) - p^*_{i,t}\left(h_{t-1}, a_{t-1}, \omega^t_i, \omega^t_{-i}\right) + \delta M_i(h_t, a^*_t) \\ \geq \; & \omega_{i,t}\left(a_t, a^{t-1}\right) - p^*_{i,t}\left(h_{t-1}, a_{t-1}, r_{i,t}, \omega_{-i,t}\right) + \delta M_i(h_t, a_t), \end{aligned} \tag{20}$$

for all $r_{i,t} \in \mathcal{F}_{i,t}$ and all $\omega_{-i,t} \in \mathcal{F}_{-i,t}$, where a_t is the socially efficient allocation if the report $r_{i,t}$ would be the true information in period t, or $a_t = a^*_t\left(h_{t-1}, a_{t-1}, \omega^t_i, \omega^t_{-i}\right)$.

By construction of the transfer price $p^*_{i,t}(\cdot)$ in ( ), the lhs of (20) 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 to get

$$W(h_t) - W_{-i}(h_t) \geq \tag{21}$$
$$\omega_{i,t}\left(a_t, a^{t-1}\right) - p^*_{i,t}(h_{t-1}, a_{t-1}, r_{i,t}, \omega_{-i,t}) + \delta\left(W(h_t, a_t) - W_{-i}(h_t, a_t)\right).$$

By construction of the transfer price $p^*_{i,t}(\cdot)$, we can represent the price that agent i would have to pay if allocation a_t were to be chosen in terms of the marginal contribution if the reported signal $r_{i,t}$ were the true signal received by agent i. We can then insert the transfer price (19) associated with the history profile $(h_{t-1}, a_{t-1}, r_{i,t}, \omega_{-i,t})$ into (21) to obtain:

$$W(h_t) - W_{-i}(h_t) \geq$$
$$\omega_{i,t}\left(a_t, a^{t-1}\right) - \sum_{j\neq i}\omega_{j,t}\left(a^*_{-i,t}, a^{t-1}\right) - \delta W_{-i}\left(h_t, a^*_{-i,t}\right) + \sum_{j\neq i}\omega_{j,t}\left(a_t, a^{t-1}\right) + \delta W(h_t, a_t).$$

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:

$$W(h_t) - W_{-i}(h_t) \geq \sum_{j}\omega_{j,t}\left(a_t, a^{t-1}\right) + \delta W(h_t, a_t) - W_{-i}(h_t),$$

where the later is true by the social optimality of a^*_t at h_t. ■

Theorem 1 gives a characterization of the monetary transfer. In specific environments, as in the earlier scheduling problem or the licensing auction in the next section, we 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.

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 monetary transfer 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 where ex post refers to reports 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. Since the efficient allocation in $t+1$ depends on information reported in t, there is no reason to believe that truthful reporting remains an optimal strategy for an agent when other agents have misreported their information. It is possible, for example, that agents other than i report in period t information that results in a negative flow marginal contribution for i when the efficient allocation is calculated according to this report. If the reports are not truthful, there is no guarantee that i can recoup period t losses in future periods. Nevertheless, our argument shows that the weaker condition of ex post incentive compatibility can be satisfied.

4.2 Uniqueness

The marginal contribution mechanism specifies a unique monetary transfer in every period and after every history. This mechanism guarantees that the ex post incentive and ex post participation constraints are satisfied after every history h_t, but it is not the only mechanism to satisfy these constraints over time. In the intertemporal environment, each agent evaluates the monetary transfers to be paid in terms of the expected discounted transfers, but is indifferent (up to discounting) about the incidence of transfers over time. The natural consequence is a multiplicity of transfer schemes that support the same intertemporal incentives as the marginal contribution mechanism. In particular, the monetary transfers necessary to induce the efficient action in period t may always be due to transfers to be paid at a later period s, provided that the relevant transfers grow at the required rate of $1/\delta$ to maintain a constant net present value. Agent i may therefore be called to make a payment long after agent i ceased to be important for the mechanism in sense of influencing current or future allocative decisions.[4]

This temporal separation between allocative influence and monetary payments may be undesirable for may reasons. First, agent i could be tempted to leave the mechanisms and break her commitment *after* she ceases to have a pivotal role but *before* her payments come due. Second, if the arrival and departure of the agents were random, then an agent could falsely claim to depart to avoid future payments. Finally, the designer could wish to minimize communication cost by eliciting information and payments only from agents who are pivotal with positive probability. In the intertemporal environment it is then natural to require that if agent i ceases to influence current or future allocative decisions in period t, then she also ceases to have monetary obligations. Formally, for agent i let time τ_i be the first time such that the efficient social decision a_s will be unaffected by the absence of agent i for all possible future states of the world, or

$$\tau_i = \min \left\{ t \left| a_s^* \left(h_t, (a_t, \omega_{t+1}), ..., (a_{s-1}, \omega^s)\right) = a_{-i,s} \left(h_t, \left(a_t, \omega^{t+1}\right), ..., (a_{s-1}, \omega^s)\right), \forall s \geq t, \forall \omega^s \right. \right\}.$$

We now say that a mechanism satisfies the efficient exit condition if the end of economic influence coincides with the end of monetary payments.

Definition 1 (Efficient Exit)

A dynamic mechanism satisfies the efficient exit condition if for all i, h_t and τ_i :

$$p_{i,s}(h_s) = 0, \quad \textit{for all } s \geq \tau_i. \tag{22}$$

The efficient exit condition is sufficient to uniquely identify the marginal contribution mechanism among all dynamic mechanism which satisfies the ex post incentive and the ex post participation constraints.

[4] We would like to thank an anonymous referee to suggest to us a link between exit and uniqueness of the transfer rule.

Theorem 2 (Uniqueness)

If a dynamic direct mechanism is efficient, satisfies the ex post incentive constraints, the ex post participation constraints and the efficient exit condition, then it is the dynamic marginal contribution mechanism.

Proof. We fix an arbitrary efficient dynamic mechanism which satisfies the ex post incentive, ex post participation and efficient exit conditions with transfer payments $\{p_{i,t}(\cdot)\}_{t=0}^{\infty}$ for all i. We first establish that for the given mechanism and for every i, h_{t-1}, a_t and ω_{-i}^t, there exists some type ω_i^t such that the monetary transfer $p_{i,t}(h_{t-1}, a_{t-1}, \omega^t)$ for the efficient allocation a_t^* is equal to the transfer payment (19) under the marginal contribution mechanism. Consider a type ω_i^t of the form

$$\omega_i^t = (\omega_{i,t}(\cdot), \mathbf{0}, \mathbf{0}, \ldots). \tag{23}$$

In words, type ω_i^t of agent i has a valuation function $\omega_{i,t}(\cdot)$ today and a valuation of zero for all allocations beyond period t. By the efficient exit condition, it follows that $p_{i,s}(\cdot) = 0$ for all $s > t$. Given ω_{-i}^t, the optimal allocation in the absence of i is given by some $a_{-i,t}^*$. For an arbitrary allocation a_t, we can now always find a utility function $\omega_{i,t}(\cdot)$ with a sufficiently high valuation for a_t such that a_t is the socially efficient allocation today, or $a_t = a_t^*$ even though i will have zero valuations starting from tomorrow. In particular we consider

$$\omega_{i,t}(a_t') = \begin{cases} 0 & \text{if} \quad a_t' \neq a_t, \\ \omega_i & \text{if} \quad a_t' = a_t, \end{cases} \tag{24}$$

for some $\omega_i \in \mathbb{R}_+$. (We can always find a continuous approximation of $\omega_{i,t}(\cdot)$ to stay in the class of continuous utility functions.) Now if $\omega_i \in \mathbb{R}_+$ is sufficiently large so as to outweigh the social externality cost of imposing a_t as the efficient allocation, or

$$\omega_i > \sum_{j \neq i} \omega_{j,t}\left(a_{-i,t}^*, a^{t-1}\right) - \sum_{j \neq i} \omega_{j,t}\left(a_t, a^{t-1}\right) + \delta\left(W_{-i}\left(h_t, a_{-i,t}^*\right) - W_{-i}(h_t, a_t)\right),$$

then a_t is the efficient allocation in period t. By the efficient exit condition, the ex post incentive and participation constraints for type ω_i^t defined by (23) and (24) reduces to the static ex post incentive and ex post participation constraints. It now follows that the transfer payment $p_{i,t}(h_t)$ has to be exactly equal to (19). For, if $p_{i,t}(h_t)$ were smaller than $p_{i,t}^*(h_t)$ of (19), then there would be valuations ω_i above but close to the social externality cost such that agent i would find the transfer payment too large to report truthfully in an ex post equilibrium. Likewise if the monetary transfer to agent i would be above $p_{i,t}^*(h_t)$ of (19), then agent i would have an incentive to induce the allocation a_t even so it would not be the socially efficient decision.

Next we argue that for all i, h_t, and a_t, the monetary transfer is equal to or below (19). Suppose not, i.e. there exists an i and h_t such that $p_{i,t}$ is above the value $p_{i,t}^*(h_t)$ of (19). Then by the argument above, we can find a type of the form (23), who would want to claim $p_{i,t}$ even though a_t is not the socially efficient decision.

Finally, we argue that for all i and h_t the monetary transfer $p_{i,t}(h_t)$ cannot be below the value $p_{i,t}^*(h_t)$ of (19) either. We observe that we already showed that the monetary transfer $p_{i,t}(h_t)$ in any period will not exceed the value of $p_{i,t}^*(h_t)$. Thus if in any period t agent i receives less than indicated by (19), she will not able to recover her loss relative to the social externality cost (19) in any future period. But in the first argument we showed that i always has the possibility, i.e. for all h_t and ω_{-i}^t, to induce the efficient allocation a_t with a monetary transfer equal to (19) by reporting a type ω_i^t of the form (23). It follows that agent i will never receive less than $p_{i,t}^*(h_t)$. We thus have shown that the lower and upper bound of the monetary transfer under ex post incentive and ex post participation constraints are equal to $p_{i,t}^*(h_t)$ provided that the efficient exit condition holds. ■

The uniqueness results uses the richness of the set of current and future utility functions to uniquely identify the set of transfers which satisfy the efficient exit condition. The argument begins with the class of types ω_i^t which cease to be economic influence after period t and given by: $\omega_i^t = (\omega_{i,t}(\cdot), \mathbf{0}, \mathbf{0}, ...)$. For these types, the incentive and participation constraints are similar to the corresponding static constraints though the transfer remain forward looking in the sense that they incorporate information about future utilities of the other agents. We then show that for these types, the marginal contribution mechanism is the only efficient mechanism which satisfies the ex post incentive, ex post participation and efficient exit conditions. We can then show that in presence of the marginal contribution transfers $p_{i,t} = p^*_{i,t}$ for the above class of types ω_i^t, the flow transfers of *all* types then have to agree with the marginal contribution transfers. We establish this by first arguing that the flow transfers for any type $\omega_i^{t\prime}$ cannot be larger than $p^*_{i,t}$ or else some of the types $\omega_i^t = (\omega_{i,t}(\cdot), \mathbf{0}, \mathbf{0}, ...)$ would have an incentive to misrepresent. Finally with an upper bound on the transfers given by the marginal contribution mechanism, it follows that every type $\omega_i^{t\prime}$ has to receive the upper bound or else type $\omega_i^{t\prime}$ would have an incentive to misreport to receive a larger flow transfer without affecting the social decision.

5 Learning and Licensing

In this section, we show how our general model can be interpreted as one where the bidders learn gradually about their preferences for an object that is auctioned repeatedly over time. We use the insights from the general marginal contribution mechanism to deduce properties of the efficient allocation mechanism. A primary example of an economic setting that fits this model is the leasing of a resource or license over time.

In every period t, a single indivisible object can be allocated to a bidder $i \in \{1, ..., I\}$. The true valuation of bidder i is given by $\theta_i \in \Theta_i = [0, 1]$. The prior distribution of θ_i is given by $F_i(\theta_i)$ and the distributions are independent across bidders. In period 0, bidder i does not know the realization of θ_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 $G_i(s_i|\theta_i)$. In each subsequent period t, only the winning bidder in period $t-1$ receives additional information about her valuation θ_i in the form of an additional and conditionally independent signal $s_{i,t} \in S_i$ from the same conditional distribution $G_i(s_i|\theta_i)$. If bidder i does not win in period t, we assume that she gets no information, and we denote this by an uninformative signal $s_{i,t} = \emptyset$. Apart from the uninformative signals, $s_{i,t}$ is private information to bidder i.[5]

In terms of the notation of the general model, $\omega_{i,t}$ is the posterior expectation of θ_i conditional on the information revealed in previous periods:

$$\omega_{i,t}\left(a^t\right) = \begin{cases} \mathbb{E}\left[\theta_i \left| h_{i,t}\right.\right] & \text{if} \quad a_t = i, \\ 0 & \text{if} \quad \text{otherwise.} \end{cases}$$

The type ω_i^t of agent i is a sequence of posterior expectations of θ_i generated by F_i and $s_i^t = (s_{i,0}, ..., s_{i,t-1})$.

Social Efficiency The socially optimal assignment over time 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 can compute for every bidder i the Gittins index based exclusively on the information about bidder i. The index of bidder i after private history $h_{i,t}$ is the solution to the following optimal stopping problem:

$$\gamma_i\left(h_{i,t}\right) = \max_{\tau_i} \mathbb{E}\left\{ \frac{\sum_{k_i=0}^{\tau} \delta^s \omega_{i,t+k_i}\left(a^{t+k_i}\right)}{\sum_{k_i=0}^{\tau} \delta^{k_i}} \right\},$$

[5] We describe the arrival of new information as a Bayesian sampling process. The equilibrium characterization in Theorem 3 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.

where a^{t+k_i} denotes the path in which alternative i has been chosen k_i times following the allocation profile a^t and where the expectation is taken with respect to the signal realizations $s_{i,t+k}$. An important property of the index policy is that the index of alternative i can be computed independent of any information about the other alternatives. In particular, the index of bidder i remains constant if bidder i does not win the object. The socially efficient allocation policy $a^* = \{a_t^*\}_{t=0}^{\infty}$ is to choose in every period a bidder i if:

$$\gamma_i(h_{i,t}) \geq \gamma_j(h_{j,t}) \text{ for all } j.$$

Dynamic Direct Mechanism In the direct dynamic mechanism, we take the flow marginal contribution to be the net utility that each bidder should receive in each period t. We construct a transfer price such that under the efficient allocation, each bidder's net payoff coincides with her flow marginal contribution $m_i(h_t)$. We also show that this pricing rule makes truthtelling incentive compatible in the dynamic mechanism.

We consider first an efficient bidder i following a history h_t.and to match her net payoff to her flow marginal contribution, we must have:

$$m_i(h_t) = \omega_i(h_{i,t}) + p_i(h_t). \tag{25}$$

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

$$m_j(h_t) = p_j(h_t).$$

We expand the flow marginal contribution in (25) 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:

$$m_i(h_t) = \omega_i(h_{i,t}) - \omega_k(h_{k,t}) - \delta(W_{-i}(h_t, i) - W_{-i}(h_t, k)). \tag{26}$$

In (26), $W_{-i}(h_t, i)$ 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 but conditional on 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:

$$m_i(h_t) = \omega_i(h_{i,t}) - (1 - \delta) W_{-i}(h_t).$$

It follows that the transfer price should simply be given by:

$$p_i^*(h_t) = -(1 - \delta) W_{-i}(h_t), \tag{27}$$

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

A similar analysis, based on the flow marginal contribution (26) 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 loosing bidder is zero, and we have:

$$p_j^*(h_t) = m_j(h_t) = 0.$$

Our main result in this section collects this information on the transfers in the dynamic marginal contribution mechanism.

Theorem 3 (Dynamic Second Price Auction)
The socially efficient allocation rule a^* is ex post incentive compatible in the dynamic direct mechanism with the payment rule $\mathbf{p}^*$ where:

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

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. It should be noted, though, that the object is not necessarily assigned to the bidder with the highest current flow payoff.

With more than two bidders, the flow value of the social program without bidder i is different from the flow value of any remaining alternative. Since there are at least two bidders left after excluding i, the planner has the option to abandon any chosen alternative if its value happens to fall sufficiently much. This option value increases the social flow payoff and hence the transfer that the efficient bidder must pay. 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.

6 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-Clarke-Groves mechanism from a static environment to intertemporal 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. Our examples show that the most immediate generalizations of standard auction formats such as dynamic ascending price auction may fail to lead to efficient allocations in dynamic models. The direct mechanism calculated in this paper is straightforward from a theoretical point of view. Nevertheless, in practice the designer may wish to find equivalent bidding mechanisms in which reports are simply statements about the willingness to pay. The initial scheduling problem points to issue of defining and analyzing reasonable or simple auction mechanisms for dynamic allocation problems.

The dynamic mechanism considered here satisfies incentive compatibility and participation constraints with respect to the efficient allocation. It is natural to ask whether the approach here may yield insights into revenue maximizing problems in dynamic models. In order to make progress in this direction, a characterization of the set of implementable dynamic allocations would be necessary. In particular with intertemporal models the signal space of every agent inherently becomes multidimensional. 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 questions are left for future research.

References

Athey, S., and I. Segal (2007): "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*, 111, 23–48.

——— (2006): "Dynamic Price Competition," *Journal of Economic Theory*, 127, 232–263.

Cavallo, R., D. Parkes, and S. Singh (2006): "Optimal Coordinated Planning Amon Self-Interested Agents with Private State," in *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, Cambridge.

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.

Edelman, B., M. Ostrovsky, and M. Schwartz (2007): "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97, 242–259.

Freixas, X., R. Guesnerie, and J. Tirole (1985): "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52, 173–191.

Friedman, E., and D. Parkes (2003): "Pricing Wifi at Starbucks - Issues in Online Mechanism Design," in *4th ACM Conference on Electronic Commerce*, pp. 240–241.

Gittins, J. (1989): *Allocation Indices for Multi-Armed Bandits*. London, Wiley.

Green, J., and J. Laffont (1977): "Révélation Des Préférences Pour Les Biens Publics : Charactérisation Des Mécanismes Satisfaisants," *Cahiers du Séminaire d'Économétrie*, 13.

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.

Parkes, D., and S. Singh (2003): "An MDP-Based Approach to Online Mechanism Design," in *Proceedings of the 17th Annual Conference on Neural Information Processing System (NIPS '03)*.

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.