

THE DYNAMIC PIVOT MECHANISM

By

Dirk Bergemann and Juuso Välimäki

August 2008

COWLES FOUNDATION DISCUSSION PAPER NO. 1672



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

The Dynamic Pivot Mechanism*

Dirk Bergemann[†]

Juuso Välimäki[‡]

First Version: September 2006

Current Version: August 2008

Abstract

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

We develop the dynamic pivot mechanism in detail for a repeated auction of a single object in which each bidder learns over time her true valuation of the object. We show that the dynamic pivot mechanism is equivalent to a modified second price auction.

JEL CLASSIFICATION: C72, C73, D43, D83.

KEYWORDS: Pivot Mechanism, Dynamic Mechanism Design, Ex-Post Equilibrium, Marginal Contribution, Multi-Armed Bandit, Bayesian Learning.

*We thank the editor, Eddie Dekel, and three anonymous referees for many helpful comments. The current paper is a major revision and supersedes “Dynamic Vickrey-Clarke-Groves Mechanisms” (2007). We are grateful to Larry Ausubel, Jerry Green, Paul Healy, John Ledyard, Benny Moldovanu, Michael Ostrovsky, David Parkes, Alessandro Pavan, Ilya Segal and Xianwen Shi for many informative conversations. The authors gratefully acknowledge 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, Duke University, LSE, Northwestern University, Ohio State University, Princeton University, University of Iowa, University of Madrid, University of Maryland and UCL 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

1 Introduction

In this paper, we generalize the idea of the pivot mechanism (due to Green and Laffont (1977b)) 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 the 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 signals and allocations. We assume throughout that the information 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 agent 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 game.

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 pivot mechanism is time-consistent and the social choice function can be implemented by a sequential mechanism without any ex-ante commitment by the designer.¹ Furthermore, the mechanism yields a net surplus in each period, and therefore the mechanism designer does not need outside resources to achieve the efficient allocation. Since marginal contributions

¹In revenue-maximizing problems, the “ratchet effect” leads to very distinct solutions for mechanisms with and without intertemporal commitment ability, see Baron and Besanko (1984) and Freixas, Guesnerie, and Tirole (1985).

are positive by definition, the dynamic pivot mechanism induces all productive agents to participate in the mechanism after all histories.

In the intertemporal environment there is a multiplicity in transfer schemes that support the same incentives as the pivot mechanism. In particular, the monetary transfers necessary to induce the efficient action in period t may become due at some later period s provided that the net present value of the transfers remains constant. We say that a mechanism supports efficient exit if an agent who ceases to affect the current and future allocations also ceases to pay and receive transfers. This condition is similar to the requirement often made in the scheduling literature that the mechanism be an *online mechanism*.² Our main characterization result shows that in an environment with diverse preferences, the dynamic pivot mechanism is the only efficient mechanism that satisfies ex-post incentive compatibility, ex-post participation and efficient exit conditions.

The basic idea of the dynamic pivot 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 to the static single-unit auction. It is easy to see that standard static mechanisms fail to produce efficient outcomes in the dynamic context. Hence a more complete understanding of the intertemporal trade-offs in the allocation process is needed. In section 5, we use the dynamic pivot mechanism to derive the optimal dynamic auction format for a model where bidders learn their valuations for a single object over time. We use the construct of the dynamic marginal contribution to derive 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 (2007b) consider a similar model to ours. Their focus is on mechanisms that are budget balanced in every period of the game. The same repeated game strategies are employed by Athey and Segal (2007a) with a focus on repeated bilateral trade. In contrast, we emphasize voluntary participation, in particular the efficient exit condition, as one of the key ingredients of our mechanism.

²The term online mechanism was coined by Lavi and Nisan (2000).

This allows us to single out the dynamic pivot mechanism in the class of efficient mechanisms. 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. They present necessary and sufficient conditions when an affine but *report-contingent* combination of dynamic allocation 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 independent 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 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 pivot mechanism in the general environment. Section 5 analyzes the implications of the general model for a licensing auction with learning.

2 Model

Uncertainty We consider an environment with private and independent values in a discrete-time, infinite-horizon model. The flow utility of agent $i \in \{1, 2, \dots, I\}$ in period $t \in \mathbb{N}$ is determined by the current allocation $a_t \in A$, the current monetary transfer $p_{i,t} \in \mathbb{R}$ and a state variable $\theta_{i,t} \in \Theta_i$. The von Neumann Morgenstern utility function u_i of agent i is assumed to be quasi-linear in the monetary transfer:

$$u_i(a_t, p_{i,t}, \theta_{i,t}) \triangleq v_i(a_t, \theta_{i,t}) - p_{i,t}.$$

We assume that $v_i(a_t, \theta_{i,t})$ is nonnegative for all i , a_t and $\theta_{i,t}$. The current allocation $a_t \in A$ is an element of a finite set A of possible allocations. The state of the world $\theta_{i,t}$ for agent i is a general Markov process on a state space Θ_i . The aggregate state of the system is given by the vector

$$\theta_t = (\theta_{1,t}, \dots, \theta_{I,t}) \in \times_{i=1}^I \Theta_i \triangleq \Theta.$$

The current state $\theta_{i,t}$ and the current action a_t define a probability distribution for next period state variables $\theta_{i,t+1}$ on Θ_i . We assume that this distribution can be represented by a stochastic kernel $F_i(\theta_{i,t+1}; \theta_{i,t}, a_t)$.

The utility functions $u_i(\cdot)$ and the probability transition functions $F_i(\cdot; a_t, \theta_{i,t})$ are common knowledge at $t = 0$. There is also a common prior $F_i(\theta_{i,0})$ regarding the initial type of each player i , and the common prior is independent across agents. At the beginning of each period t , each player i observes $\theta_{i,t}$ privately. At the end of each period, an action $a_t \in A$ is chosen and payoffs for period t are realized. The asymmetric information is therefore generated by the private observation of $\theta_{i,t}$ in each period t . We observe that by the independence of the priors and the stochastic kernels across i , the information of player, $\theta_{i,t+1}$ does not depend on $\theta_{j,t}$ for $j \neq i$. The expected flow payoff of every agent is assumed to be bounded for every allocation plan $a' : \Theta \rightarrow A$:

$$\int u_i(a'(\theta'), \theta'_i) dF(\theta'; a, \theta) < K,$$

for some $K < \infty$ for all i , all a and all θ .

The nature of the state space Θ will depend on the application at hand. At this point, we should stress that the formulation allows us to accommodate the possibility of random arrival or random departure of new agents. It is, for example, quite natural to model the arrival or the departure of a player i through an inactive state θ_i^0 , where $v_i(a_t, \theta_i^0) = 0$ for all $a_t \in A$ and a random time τ at which agent i privately observes her transitions in and out of the inactive state. We discuss the role of the richness of the signal space and the associated valuation profiles in more detail in the context of Theorem 2 which presents a uniqueness result for efficient mechanisms.

Social Efficiency All agents discount the future with a common discount factor δ , $0 < \delta < 1$. 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 the Markovian structure of the stochastic process, the socially optimal program starting in period t at state θ_t can be written simply as

$$W(\theta_t) \triangleq \max_{\{a_s\}_{s=t}^{\infty}} \mathbb{E} \left\{ \sum_{s=t}^{\infty} \delta^{s-t} \sum_{i=1}^I v_i(a_s, \theta_{i,s}) \right\}.$$

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

$$W(\theta_t) = \max_{a_t} \mathbb{E} \left\{ \sum_{i=1}^I v_i(a_t, \theta_{i,t}) + \delta \mathbb{E} W(\theta_{t+1}) \right\}.$$

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:

$$W_{-i}(\theta_t) \triangleq \max_{\{a_s\}_{s=t}^{\infty}} \mathbb{E} \left\{ \sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \neq i} v_j(a_s, \theta_{j,s}) \right\}.$$

The efficient policy when agent i is excluded is denoted by $\mathbf{a}_{-i}^* = \{a_{-i,t}^*\}_{t=0}^{\infty}$. The *marginal contribution* $M_i(\theta_t)$ of agent i at signal θ_t is defined by:

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

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

Mechanism and Equilibrium A dynamic direct mechanism asks every agent i to report her state $\theta_{i,t}$ in every period t . The report $r_{i,t} \in \Theta_i$ may be truthful or not depending on the incentives provided in the mechanism. The public history in period t is then a sequence of reports and allocations until period $t-1$, or $h_t = (r_0, a_0, r_1, a_1, \dots, r_{t-1}, a_{t-1})$,

where each $r_s = (r_{1,s}, \dots, r_{I,s})$ 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 we assume that the past and current reports of each agent are observable to all the agents. The private history of agent i in period t consists of the public history and the sequence of private observations until period t , or $h_{i,t} = (\theta_{i,0}, r_0, a_0, \theta_{i,1}, r_1, a_1, \dots, \theta_{i,t-1}, r_{t-1}, a_{t-1}, \theta_{i,t})$. The set of possible private histories in period t is denoted by $H_{i,t}$. An (*efficient*) *dynamic direct mechanism* is then represented by a family of allocations and monetary transfers, $\{a_t^*, p_t\}_{t=0}^\infty$, specifically a sequence of allocations:

$$a_t^* : \Theta \rightarrow \Delta(A),$$

and a sequence of monetary transfers:

$$p_t : H_t \times \Theta \rightarrow \mathbb{R}^I,$$

such that the decisions in period t respond to the reported information of all agents up to and including period t . With the focus on efficient mechanisms, the allocation a_t^* depends only on the current state θ_t . In contrast, the determination of the transfer may depend on the entire history of reports and actions.

In a dynamic direct mechanism, a (pure) reporting strategy for agent i in period t is a mapping from the private history into the signal space:

$$r_{i,t} : H_{i,t} \rightarrow \Theta_i.$$

For a given mechanism, the expected payoff for agent i from reporting strategy $r_i = \{r_{i,t}\}_{t=0}^\infty$ given that the others agents are reporting $r_{-i} = \{r_{-i,t}\}_{t=0}^\infty$ is given by

$$\mathbb{E} \sum_{t=0}^{\infty} \delta^t [v_i(a_t^*(r_t), \theta_{i,t}) - p_i(h_t, r_t)].$$

The allocations $a_t^*(r_t)$ are determined by the current reports r_t . Given the mechanism and the reporting strategies r_{-i} , the optimal reporting strategy of bidder i solves a sequential optimization problem which can be phrased recursively in terms of value functions, or

$$V_i(h_{i,t}) = \max_{r_{i,t} \in \Theta_i} \mathbb{E} \{v_i(a_t^*(r_{i,t}, r_{-i,t}), \theta_{i,t}) - p_i(h_t, r_{i,t}, r_{-i,t}) + \delta V_i(h_{i,t+1})\}.$$

The value function $V_i(h_{i,t+1})$ represents the continuation value of agent i given the current private history $h_{i,t}$, the current reports r_t , the current allocation a_t and tomorrow's private signal $\theta_{i,t+1}$ as $h_{i,t+1} = (h_{i,t}, r_t, a_t, \theta_{i,t+1})$. We say that a dynamic direct mechanism is *interim incentive compatible*, if for every agent and every history, truthtelling is a best response given that all other agents report truthfully. We say that the dynamic direct mechanism is *periodic ex-post incentive compatible* if truthtelling is a best response regardless of the history and the current signal realization of the other agents.

In the dynamic context, the notion of ex-post incentive compatibility is 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 arises 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 define the interim participation constraints of each agent. After each history h_t , each agent i may opt out (permanently) from the mechanism, and receive the outside option value $O_i(h_{i,t})$. We use payoffs generated by the efficient policy \mathbf{a}_{-i}^* for the remaining agents to calculate $O_i(h_{i,t})$ for the rest of the paper.³ The periodic participation constraint requires that each agent's equilibrium payoff after each history weakly exceeds $O_i(h_{i,t})$. For the remainder of the text we shall say that a mechanism is *ex-post incentive compatible and individually rational* if it satisfies the periodic incentive and participation constraints.

3 Scheduling: An Example

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

³The allocation decision a may in itself indicate which players are active in the game. As a result, the payoff to the remaining players will reflect the exit decisions through allocation decisions.

of the social planner is to sequence the tasks over time so as to maximize the sum of the discounted utilities. In an early contribution, Dolan (1978) developed a static mechanism to implement a class of related scheduling problems with private information.

An allocation policy in this setting is a sequence of choices $a_t \in \{0, 1, \dots, I\}$, 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 $\theta_{i,0} \in \mathbb{R}_+$ of the task is constant over time and independent of the realization time (except for discounting). The transition function is then given by:

$$\begin{aligned} \theta_{i,t+1} &= 0, & \text{if } a_t^* &= i, \\ \theta_{i,t+1} &= \theta_{i,t} & \text{if } a_t^* &\neq i. \end{aligned}$$

The utility function $v_i(a_t, \theta_{i,t})$ for bidder i from the efficient allocation policy a^* is given by:

$$v_i(a_t, \theta_{i,t}) = \begin{cases} \theta_{i,t} & \text{if } a_t = i, \\ 0 & \text{if otherwise.} \end{cases}$$

For this scheduling model, we find the marginal contribution of each agent and derive the associated dynamic pivot mechanism. We determine the marginal contribution of bidder i by comparing the value of the social program with and without i . With constant valuations $v_i(\cdot)$ over time for all i , the optimal policy is clearly given by assigning in every period the alternative j with the highest remaining valuation. To simplify notation, we let

$$v_i \triangleq v_i(i, \theta_{i,t}),$$

and by convention define $v_t \triangleq 0$ for all $t > I$. We may assume without loss of generality (after relabelling) that the valuations v_i of the agents are ordered with respect to their identity i :

$$v_1 \geq \dots \geq v_I \geq 0. \tag{2}$$

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:

$$W(\theta_0) = \sum_{t=1}^I \delta^{t-1} v_t. \tag{3}$$

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:

$$W_{-i}(\theta_0) = \sum_{t=1}^{i-1} \delta^{t-1} v_t + \sum_{t=i}^{I-1} \delta^{t-1} v_{t+1}. \quad (4)$$

By comparing the social program with and without i , (3) and (4), 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 from the point of view of period 0 is:

$$M_i(\theta_0) = W(\theta_0) - W_{-i}(\theta_0) = \sum_{t=i}^I \delta^{t-1} (v_t - v_{t+1}).$$

Correspondingly, the present value of marginal contribution of i at the time $t = i - 1$ at which she realizes her task is

$$M_i(\theta_{i-1}) = W(\theta_{i-1}) - W_{-i}(\theta_{i-1}) = \sum_{t=i}^I \delta^{t-i} (v_t - v_{t+1}).$$

The social externality cost of agent i is now established in a straightforward manner. At time $t = i - 1$, bidder i will complete her task and hence realize a gross value of v_i . The immediate opportunity cost is given by the next highest valuation v_{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 v_{i+1} minus the improvement in future allocations due to the delay of all tasks by one period:

$$p_i(\theta_t) = v_{i+1} - \sum_{t=i+1}^I \delta^{t-i} (v_t - v_{t+1}). \quad (5)$$

Since by construction (see (2)), we have $v_t - v_{t+1} \geq 0$, 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 v_{i+1} . Consequently, we can rewrite (5) to:

$$p_i(\theta_t) = (1 - \delta) \sum_{t=i}^I \delta^{t-i} v_{t+1},$$

which simply states that the externality cost of agent i is the cost of delay imposed on the remaining and less valuable tasks. With the monetary transfers given by (5), Theorem 1 will formally establish that the dynamic pivot mechanism leads to truth-telling with ex-post incentive and ex-post participation constraints.

We next show that 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 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 lose it today, she is willing to bid at most

$$b_I(v_I) = v_I - \delta(v_I - 0), \quad (6)$$

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 δv_{I-1} and hence her maximal willingness to pay is given by

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

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

$$v_{I-1} - (1 - \delta)v_I.$$

We proceed inductively and find that the maximal bid of bidder $I - k$ in period $t =$

$I - k - 1$ is given by:

$$b_{I-k}(v_{I-k}) = v_{I-k} - \delta (v_{I-k} - b_{I-(k-1)}(v_{I-(k-1)})) \quad (7)$$

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 $v_{I-k} - b_{I-(k-1)}$, but only tomorrow, and so the net gain from being selected today rather than tomorrow is:

$$v_{I-k} - \delta (v_{I-k} - b_{I-(k-1)})$$

The maximal bid of bidder $I - (k - 1)$ generates the transfer price of bidder $I - k$ and by solving (7) recursively with the initial condition given by (6), we find that the price in the ascending auction equals the externality cost in the direct mechanism. In this class of scheduling problems, the efficient allocation can therefore be implemented by a bidding mechanism.⁴

We end this section with a minor modification of the scheduling model to allow for multiple tasks. For this purpose it is sufficient to consider an example with two bidders. The first bidder has an infinite series of single-period tasks, each delivering a value of v_1 . The second bidder has only a single task with a value v_2 . The utility function of bidder 1 is thus given by

$$v_1(a_t, \theta_{1,t}) = \begin{cases} v_1 & \text{if } a_t = 1 \text{ for all } t, \\ 0 & \text{if otherwise.} \end{cases}$$

whereas the utility function of bidder 2 is as described earlier.

The socially efficient allocation in this setting either has $a_t = 1$ for all t if $v_1 \geq v_2$ or $a_0 = 2$, $a_t = 1$ for all $t \geq 1$ if $v_1 < v_2$. For the remainder of this example, we will assume that $v_1 > v_2$. Under this assumption the efficient policy will never complete the task of

⁴The 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.

bidder 2. The marginal contribution of each bidder is:

$$M_1(\theta_0) = (v_1 - v_2) + \frac{\delta}{1 - \delta}v_1,$$

and

$$M_2(\theta_0) = 0.$$

Along any efficient allocation path, we have $M_i(\theta_0) = M_i(\theta_t)$ for all i and the social externality cost of agent 1, $p_1^*(\theta_t)$ for all t , is $p_1^*(\theta_t) = -(1 - \delta)v_2$. 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 v_2 in every period. In consequence the first bidder would receive a net payoff of $v_1 - v_2$ in every period and her discounted sum of payoff would then be:

$$\frac{1}{1 - \delta}(v_1 - v_2) < M_1(\theta_0). \quad (8)$$

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 loses 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:

$$\frac{\delta}{1 - \delta}v_1. \quad (9)$$

If we compare the continuation payoffs (8) and (9) respectively, then we see that it is beneficial for bidder 1 to win the auction in all periods if and only if

$$v_1 \geq \frac{v_2}{1 - \delta},$$

but the efficiency condition is simply $v_1 \geq v_2$. 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 The Dynamic Pivot Mechanism

We now construct the dynamic pivot mechanism for the general model described in Section 2. The marginal contribution of agent i is her contribution to the social value. In the dynamic pivot mechanism, the marginal contribution will also be 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* $m_i(\theta_t)$ in every period. The flow marginal contribution accrues incrementally over time and is defined recursively:

$$M_i(\theta_t) = m_i(\theta_t) + \delta \mathbb{E} M_i(\theta_{t+1}).$$

The flow marginal contribution can be expressed directly in terms of the social value functions, using the definition of the marginal contribution given in (1), as:

$$m_i(\theta_t) = W(\theta_t) - W_{-i}(\theta_t) - \delta \mathbb{E}[W(\theta_{t+1}) - W_{-i}(\theta_{t+1})]. \quad (10)$$

We can further replace the value functions $W(\theta_t)$ and $W_{-i}(\theta_t)$ by the corresponding flow payoffs and continuation payoffs to express the flow marginal contribution of agent i in terms of flow and continuation payoffs. The continuation payoffs of the social programs with and without i , respectively, may be governed by different transition probabilities as the respective social decisions in period t , $a_t^* \triangleq a_t^*(\theta_t)$ and $a_{-i,t}^* \triangleq a_{-i,t}^*(\theta_{-i,t})$, may differ. We denote the expected continuation value of the socially optimal programs, conditional on the current state and the current action, by:

$$W(\theta_{t+1} | a_t, \theta_t) \triangleq \mathbb{E}_{F(\theta_{t+1}; a_t, \theta_t)} W(\theta_{t+1}),$$

where the transition from state θ_t to state θ_{t+1} is controlled by the allocation a_t . For notational ease we omit the expectations operator \mathbb{E} from the conditional expectation. We adopt the same notation for the marginal contributions $M_i(\cdot)$ and the individual value functions $V_i(\cdot)$. The flow marginal contribution $m_i(\theta_t)$ can be expressed as:

$$m_i(\theta_t) = \sum_{j=1}^I v_j(a_t^*, \theta_{j,t}) - \sum_{j \neq i} v_j(a_{-i,t}^*, \theta_{j,t}) + \delta [W_{-i}(\theta_{t+1} | a_t^*, \theta_t) - W_{-i}(\theta_{t+1} | a_{-i,t}^*, \theta_t)]. \quad (11)$$

A monetary transfer $p_i^*(\theta_t)$ such that the resulting flow net utility matches the flow marginal contribution leads agent i to internalize her social externalities:

$$p_i^*(\theta_t) \triangleq v_i(a_t^*, \theta_{i,t}) - m_i(\theta_t). \quad (12)$$

We refer to $p_i^*(\theta_t)$ as the transfer of the dynamic pivot mechanism. We observe that the transfer $p_i^*(\theta_t)$ depends only on the current report θ_t and does not depend on the past public history h_t . Inserting (11) into (12) we can express the transfer payment of the dynamic pivot mechanism in terms of the flow utilities and the continuation social values:

$$p_i^*(\theta_t) = \sum_{j \neq i} [v_j(a_{-i,t}^*, \theta_{j,t}) - v_j(a_t^*, \theta_{j,t})] + \delta [W_{-i}(\theta_{t+1} | a_{-i,t}^*, \theta_t) - W_{-i}(\theta_{t+1} | a_t^*, \theta_t)]. \quad (13)$$

The transfer price (13) for agent i depends on the report of agent i only through the determination of the social allocation which is a prominent feature already in the static Vickrey-Clarke-Groves mechanisms. The monetary transfers $p_i^*(\theta_t)$ are always non-negative 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.

Theorem 1 (Dynamic Pivot Mechanism)

The dynamic pivot mechanism $\{a_t^, p_t^*\}_{t=0}^\infty$ is ex-post incentive compatible and individually rational.*

Proof. By the unimprovability principle, it suffices to prove that if agent i receives as her continuation value her marginal contribution, then truthtelling is incentive compatible for agent i in period t , or:

$$\begin{aligned} v_i(a_t^*(\theta_t), \theta_{i,t}) - p_i^*(\theta_t) + \delta M_i(\theta_{t+1} | a_t^*, \theta_t) &\geq \\ v_i(a_t^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - p_i^*(r_{i,t}, \theta_{-i,t}) + \delta M_i(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t), \end{aligned} \tag{14}$$

for all $r_{i,t} \in \Theta_i$ and all $\theta_{-i,t} \in \Theta_{-i}$, where $a_t^* = a^*(\theta_{i,t}, \theta_{-i,t})$ is the socially efficient allocation if the report is $r_{i,t} = \theta_{i,t}$. By construction of the transfer price p_i^* in (13), the lhs of (14) represents the marginal contribution of agent i . We can express $M_i(\theta_{t+1} | a_t^*, \theta_t)$ and $M_i(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t)$, respectively, in terms of the values of the different social programs to get:

$$\begin{aligned} W(\theta_t) - W_{-i}(\theta_t) &\geq v_i(a_t^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - p_i^*(r_{i,t}, \theta_{-i,t}) \\ &+ \delta (W(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t) - W_{-i}(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t)). \end{aligned} \tag{15}$$

By construction of p_i^* , we can represent the transfer that agent i would pay if allocation $a^*(r_{i,t}, \theta_{-i,t})$ were 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 (13) into (15)

to obtain:

$$\begin{aligned}
W(\theta_t) - W_{-i}(\theta_t) &\geq v_i(a_t^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - \sum_{j \neq i} v_j(a_{-i,t}^*, \theta_{j,t}) - \delta W_{-i}(\theta_{t+1} | a_{-i,t}^*, \theta_t) \\
&+ \sum_{j \neq i} v_j(a_t^*(r_{i,t}, \theta_{-i,t}), \theta_{j,t}) + \delta W(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t).
\end{aligned}$$

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:

$$W(\theta_t) - W_{-i}(\theta_t) \geq \sum_{j=1}^I v_j(a_t^*(r_{i,t}, \theta_{-i,t}), \theta_{j,t}) + \delta W(\theta_{t+1} | a_t^*(r_{i,t}, \theta_{-i,t}), \theta_t) - W_{-i}(\theta_t).$$

The above inequality holds for all $r_{i,t}$ by the social optimality of $a_t^*(\theta_t)$ in state θ_t . ■

The dynamic pivot 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 . 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 the transfers over time. This temporal separation between allocative decisions and monetary decisions may be undesirable for many reasons. First, if the agents and the principal do not have the ability to commit to *future* transfer payments, then delays in payments become problematic. In consequence, an agent which is not pivotal should not receive or make a payment. Second, if it is costly (in a lexicographic sense) to maintain accounts of future monetary commitments, then the principal wants to close down (as early as possible) the accounts of those agents who are no longer pivotal.⁵

This motivates the following efficient exit condition. Let state θ_{τ_i} in period τ_i be a state such that the probability that agent i affects the efficient social decision a_t^* for all $t \geq \tau_i$ is equal to zero:

$$\Pr(a_t^*(\theta_t) \neq a_{-i,t}^*(\theta_t), \forall t \geq \tau_i | \theta_{\tau_i}) = 0.$$

We now say that a mechanism satisfies the efficient exit condition if for every agent i the end of her allocative influence coincides with the end of her monetary transfers .

⁵We would like to thank an anonymous referee for the suggestion to consider the link between exit and uniqueness of the transfer rule.

Definition 1 (Efficient Exit)

A dynamic direct mechanism satisfies the efficient exit condition if for all $i, \tau_i, \theta_{\tau_i}$:

$$p_i(\theta_{\tau_i}) = 0.$$

We now establish the uniqueness of the dynamic pivot mechanism in an environment with diverse preferences and the efficient exit condition. The assumption of diverse preferences allows for rich preferences over the current allocations and indifference over future allocations. We maintain this assumption for the remainder of the paper.

Assumption 1 (Diverse Preferences)

1. For all i , there exists $\theta_i^0 \in \Theta_i$ such that for all a , $v_i(a, \theta_i^0) = 0$ and $F_i(\theta_i^0; a, \theta_i^0) = 1$.
2. For all i , and for all a and all $x \in \mathbb{R}_+$, there exists $\theta_i^{a,x} \in \Theta_i$ such that

$$v_i(a_t, \theta_i^{a,x}) = \begin{cases} x & \text{if } a_t = a, \\ 0 & \text{if } a_t \neq a, \end{cases},$$

$$\text{and for all } a_t, F_i(\theta_i^0; a_t, \theta_i^{a,x}) = 1.$$

The first part of the diverse preference assumption assigns to each agent a state in which she gets no payoff from any allocation, and that this state is an absorbing state. The second part requires that each agent have a state in which she has a positive valuation x for a specific current allocation a and no value for other current or future allocations. Assuming diverse preferences is similar to imposing the rich domain conditions introduced in Green and Laffont (1977a) and Moulin (1986) to establish the uniqueness of the Groves and the Pivot mechanism in a static environment. Relative to their conditions, we augment the diverse (flow) preferences with the certain transition into the absorbing state θ_i^0 . With this transition we ensure that the diverse flow preferences continue to matter in the intertemporal environment.

The assumption of diverse preference in conjunction with the efficient exit condition guarantees that in every dynamic direct mechanism there are *some* types, specifically the

types of the form $\theta_i^{a,x}$, that receive exactly the flow transfers they would have received in the dynamic pivot mechanism.

Lemma 1

If $\{a_t^*, p_t\}_{t=0}^\infty$ is ex-post incentive compatible and individually rational, and satisfies the efficient exit condition, then:

$$p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) = p_i^*(\theta_i^{a,x}, \theta_{-i,t}), \text{ for all } i, a, x, \theta_t, h_t.$$

Proof. Consider any arbitrary history h_t and type realization $(\theta_i^{a,x}, \theta_{-i,t})$ in period t . The ex-post incentive constraints of type $\theta_i^{a,x}$ at type profile $(\theta_i^{a,x}, \theta_{-i,t})$ are:

$$\begin{aligned} v_i(a^*(\theta_i^{a,x}, \theta_{-i,t}), \theta_i^{a,x}) - p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) + \delta V_i(h_{i,t+1} | a^*(\theta_i^{a,x}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})) \geq \\ v_i(a^*(r_{i,t}, \theta_{-i,t}), \theta_i^{a,x}) - p_i(h_t, r_{i,t}, \theta_{-i,t}) + \delta V_i(h_{i,t+1} | a^*(r_{i,t}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})), \end{aligned}$$

for all $r_{i,t}$. Given $(\theta_i^{a,x}, \theta_{-i,t})$, the continuation payoff for i along the equilibrium path satisfies $V_i(h_{i,t+1} | a^*(\theta_i^{a,x}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})) = 0$ by the efficient exit condition.

In the dynamic pivot mechanism, if the valuation x of type $\theta_i^{a,x}$ for allocation a exceeds the social externality cost, or

$$\begin{aligned} x \geq \sum_{j \neq i} [v_j(a_{-i}^*(\theta_{-i,t}), \theta_{j,t}) - v_j(a, \theta_{j,t})] \\ + \delta W_{-i}(\theta_{-i,t+1} | a_{-i}^*(\theta_{-i,t}), \theta_{-i,t}) - \delta W_{-i}(\theta_{-i,t+1} | a, \theta_{-i,t}), \end{aligned} \tag{16}$$

then the transfer price $p_i^*(\theta_i^{a,x}, \theta_{-i,t})$ would be:

$$\begin{aligned} p_i^*(\theta_i^{a,x}, \theta_{-i,t}) = \sum_{j \neq i} [v_j(a_{-i}^*(\theta_{-i,t}), \theta_{j,t}) - v_j(a, \theta_{j,t})] \\ + \delta W_{-i}(\theta_{-i,t+1} | a_{-i}^*(\theta_{-i,t}), \theta_{-i,t}) - \delta W_{-i}(\theta_{-i,t+1} | a, \theta_{-i,t}) \end{aligned}$$

otherwise it would be equal to zero.

We now argue by contradiction. By the ex-post incentive compatibility constraints, all types $\theta_i^{a,x}$ of agent i where x satisfies the inequality (16) must pay the same transfer. To see this, suppose that for $x, y \in \mathbb{R}_+$ satisfying (16)

$$p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) < p_i(h_t, \theta_i^{a,y}, \theta_{-i,t}).$$

Now type $\theta_i^{a,y}$ has a strict incentive to mis-report $r_{i,t} = \theta_i^{a,x}$, a contradiction. We therefore denote the constant transfer for all $\theta_i^{a,x}$ and x satisfying (16) by $p_i(h_t, a, \theta_{-i,t})$ and the corresponding dynamic pivot transfer by $p_i^*(a, \theta_{-i,t})$.

Suppose next that $p_i(h_t, a, \theta_{-i,t}) > p_i^*(a, \theta_{-i,t})$. This implies that the ex-post participation constraint for some x with $p_i(h_t, a, \theta_{-i,t}) > x > p_i^*(a, \theta_{-i,t})$ is violated, contradicting the hypothesis of the lemma. Suppose to the contrary that $p_i(h_t, a, \theta_{-i,t}) < p_i^*(a, \theta_{-i,t})$, and consider the ex-post incentive constraints of a type $\theta_i^{a,x}$ with a valuation x such that

$$p_i(h_t, a, \theta_{-i,t}) < x < p_i^*(a, \theta_{-i,t}). \quad (17)$$

If the inequality (17) is satisfied then it follows that $a^*(\theta_i^{a,x}, \theta_{-i,t}) = a_{-i}^*(\theta_{-i,t})$, and in particular that $a^*(\theta_i^{a,x}, \theta_{-i,t}) \neq a$. If the ex-post incentive constraint of type $\theta_i^{a,x}$ were satisfied, then we would have

$$v_i(a^*(\theta_i^{a,x}, \theta_{-i,t}), \theta_i^{a,x}) - p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) \geq v_i(a, \theta_i^{a,x}) - p_i(h_t, a, \theta_{-i,t}). \quad (18)$$

Given that $\theta_i = \theta_i^{a,x}$, we can thus rewrite (18) as:

$$0 - p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) \geq x - p_i(h_t, a, \theta_{-i,t}).$$

But given (17), this implies that $p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) < 0$. In other words, type $\theta_i^{a,x}$ receives a strictly positive subsidy even though her report is not pivotal for the social allocation as $a^*(\theta_i^{a,x}, \theta_{-i,t}) = a_{-i}^*(\theta_{-i,t})$. Now, a negative transfer (i.e. a positive subsidy) necessarily violates the ex-post incentive constraint of the absorbing type θ_i^0 . By the efficient exit condition, type θ_i^0 should not receive any contemporaneous (or future) subsidies. But by mis-reporting her type to be $\theta_i^{a,x}$, type θ_i^0 would gain access to a positive subsidy without changing the social allocation, which would leave her with a strictly positive net utility. It thus follows that $p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) = p_i^*(\theta_i^{a,x}, \theta_{-i,t})$ for all a and all x . ■

Given that the transfers of the dynamic pivot mechanism are part of every dynamic direct mechanism with diverse preferences, we next establish that every type $\theta_{i,0}$ in $t = 0$ has to receive the same ex-ante expected utility as she would in the dynamic pivot mechanism.

Lemma 2

If $\{a_t^*, p_t\}_{t=0}^\infty$ is ex-post incentive compatible and individually rational, and satisfies the efficient exit condition, then for all i and all θ_0 : $V_i(\theta_0) = M_i(\theta_0)$.

Proof. The argument is by contradiction. Consider an i such that $V_i(\theta_0) \neq M_i(\theta_0)$. Suppose first that $V_i(\theta_0) > M_i(\theta_0)$. Then there is a history h_τ and a state θ_τ such that $p_i(h_\tau, \theta_\tau) < p_i^*(\theta_\tau)$. The socially efficient allocation in state θ_τ is denoted $a_\tau^* \triangleq a^*(\theta_\tau)$. We now show that such a transfer $p_i(h_\tau, \theta_\tau)$ leads to a violation of the ex-post incentive constraint for some type $\theta_i^{a,x} \in \Theta_i$. Specifically consider a type $\theta_i^{a_\tau^*, x}$ such that $p_i(h_\tau, \theta_\tau) < x < p_i^*(\theta_\tau)$. The ex-post incentive constraints of type $\theta_i^{a_\tau^*, x}$ imply that

$$\begin{aligned} 0 &= v_i\left(a^*\left(\theta_i^{a_\tau^*, x}, \theta_{-i, \tau}\right), \theta_i^{a_\tau^*, x}\right) - p_i\left(h_\tau, \theta_i^{a_\tau^*, x}, \theta_{-i, \tau}\right) + \delta V_i\left(h_{i, \tau+1} \mid a^*\left(\theta_i^{a_\tau^*, x}, \theta_{-i, \tau}\right), \theta_\tau\right) \\ &\geq x - p_i(h_\tau, \theta_\tau) + \delta V_i\left(h_{i, \tau+1} \mid a^*(r_{i, \tau}, \theta_{-i, \tau}), \theta_i^{a_\tau^*, x}, \theta_{-i, \tau}\right) > 0, \end{aligned}$$

leading to a contradiction.

Suppose next that

$$M_i(\theta_0) - V_i(\theta_0) > \varepsilon, \tag{19}$$

for some $\varepsilon > 0$. By hypothesis of ex-post incentive condition we have for all $r_{i,0}$:

$$M_i(\theta_0) - [v_i(a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0}) - p_i(h_0, r_{i,0}, \theta_{-i,0}) + \delta V_i(h_{i,1} \mid a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0})] > \varepsilon. \tag{20}$$

But by Lemma 1, we know that there exists a report $r_{i,0} = \theta_i^{a^*(\theta_0), x}$ for agent i such that $a^*(\theta_0)$ is induced at the price $p_i^*(\theta_0)$ associated with the dynamic pivot mechanism. After inserting $r_{i,0} = \theta_i^{a^*(\theta_0), x}$ into (20) and observing that

$$v_i(a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0}) - p_i(h_0, r_{i,0}, \theta_{-i,0}) = m_i(\theta_0),$$

we are lead to conclude that

$$\delta (M_i(\theta_1) - V_i(h_{i,1} \mid a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0})) > \varepsilon,$$

or

$$M_i(\theta_1) - V_i(h_{i,1} \mid a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0}) > \frac{\varepsilon}{\delta}.$$

But now we can repeat the argument we started with (19) and find that there is a path of realizations of $\theta_0, \dots, \theta_t$, such that the difference between the marginal contribution and the value function of agent i grows without bound. But the marginal contribution of agent i is finite given that the expected flow utility of agent i is bounded by some $K > 0$, and thus eventually the ex-post participation constraint of the agent is violated, and we obtain the desired contradiction. ■

The above lemma can be viewed as a revenue equivalence results of all (efficient) dynamic direct mechanisms. As we are analyzing a dynamic allocation problem with an infinite horizon, we cannot appeal to the revenue equivalence results established for static mechanisms. In particular, the statement of the standard revenue equivalence results involve a fixed utility for the lowest type. In the infinite horizon model here, the diverse preference assumption give us a natural candidate of a lowest type in terms of θ_i^0 , and the efficient exit condition determines her utility. The remaining task is to argue that among all intertemporal transfers with the same expected discounted value, only the time profile of the dynamic pivot mechanism satisfies the relevant conditions. Alternative payments streams could either require an agent to pay earlier or later relative to the dynamic pivot transfers. If the payments were to occur later, payments would have to be lower in an earlier period by the above revenue equivalence result. This would open the possibility for a “short-lived” type $\theta_i^{a,x}$ to induce action a at a price below the dynamic pivot transfer and hence violate incentive compatibility. The reverse argument applies if the payments were to occur earlier relative to the dynamic pivot transfer, for example if the agent were to be asked to post a bond at the beginning of the mechanism.

Theorem 2 (Uniqueness)

If $\{a_t^, p_t\}_{t=0}^\infty$ is ex-post incentive compatible and individually rational, and satisfies the efficient exit condition, then it is the dynamic pivot mechanism.*

Proof. The proof is by contradiction. Suppose not, then by Lemma 2 there exists a player i , a history h_τ and an associated state $\theta_{i,\tau}$ such that $p_i(h_\tau, \theta_\tau) \neq p_i^*(\theta_\tau)$. Suppose first that $p_i(h_\tau, \theta_\tau) < p_i^*(\theta_\tau)$. We show that the current monetary transfer $p_i(h_\tau, \theta_\tau)$ leads

to the violation of the ex-post incentive constraint of some type $\theta_i^{a,x}$. The socially efficient allocation at the true profile θ_τ is given by $a_\tau^* = a^*(\theta_\tau)$. Consider now a type $\theta_i^{a_\tau^*,x}$ with a valuation x for the allocation a_τ^* such that $x > p_i^*(\theta_\tau)$. Her ex-post incentive constraint is given by

$$\begin{aligned} v_i(a^*(\theta_i^{a,x}, \theta_{-i,t}), \theta_i^{a,x}) - p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) + \delta V_i(h_{i,t+1} | a^*(r_{i,t}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})) \geq \\ v_i(a^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - p_i(h_t, r_{i,t}, \theta_{-i,t}) + \delta V_i(h_{i,t+1} | a^*(r_{i,t}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})). \end{aligned}$$

By the efficient exit condition, we have

$$V_i(h_{i,t+1} | a^*(r_{i,t}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})) = V_i(h_{i,t+1} | a^*(r_{i,t}, \theta_{-i,t}), (\theta_i^{a,x}, \theta_{-i,t})) = 0.$$

By Lemma 1, we also have that $p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) = p_i^*(\theta_i^{a,x}, \theta_{-i,t}) = p_i^*(\theta_\tau)$. Consider then the report of $r_{i,\tau} = \theta_{i,\tau}$ by type $\theta_i^{a,x}$. The ex-post incentive constraints now reads: $x - p_i^*(\theta_\tau) \geq x - p_i(h_\tau, \theta_\tau)$, which leads to a contradiction as by hypothesis we had $p_i(h_\tau, \theta_\tau) < p_i^*(\theta_\tau)$.

Suppose next that $p_i(h_\tau, \theta_\tau) > p_i^*(\theta_\tau)$. Now by Lemma 2, it follows that the ex-ante expected payoff is equal to the value of the marginal contribution of agent i in period 0. It therefore follows from $p_i(h_\tau, \theta_\tau) > p_i^*(\theta_\tau)$ that there also exists another time τ' and state $\theta_{\tau'}$ such that $p_i(h_\tau, \theta_\tau) < p_i^*(\theta_\tau)$. By repeating the argument in the first part of the proof, we obtain a contradiction. ■

We should reiterate that in the definition of the ex-post incentive and participation conditions, we required that a candidate mechanism satisfies these conditions after all possible histories of past reports. It is in the spirit of the ex-post constraints that these constraints hold for all possible states rather than strictly positive probability events. In the context of establishing the uniqueness of the mechanism it allows us to use the diverse preference condition without making additional assumption about the transition probability from a given state $\theta_{i,t}$ into a specific state $\theta_i^{a,x}$. We merely require the existence of these types in the establishing the above result.

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 pivot 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 $i \in \{1, \dots, I\}$, and the allocation decision $a_t \in \{1, 2, \dots, I\}$ simply determines which bidder gets the object in period t . In order to describe the uncertainty explicitly, we assume that the true valuation of bidder i is given by $\omega_i \in \Omega_i = [0, 1]$. Information in the model represents therefore the bidder's prior and posterior beliefs on ω_i . In period 0, bidder i does not know the realization of ω_i , but she has a prior distribution $\theta_{i,0}(\omega_i)$ on Ω_i . The prior and posterior distributions on Ω_i are assumed to be independent across bidders. In each subsequent period t , only the winning bidder in period $t - 1$ receives additional information leading to an updated posterior distribution $\theta_{i,t}$ on Ω_i according to Bayes' rule. If bidder i does not win in period t , we assume that she gets no information, and consequently the posterior is equal to the prior. In the dynamic direct mechanism, the bidders simply report their posteriors at each stage.

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:

$$\gamma_i(h_{i,t}) = \max_{\tau_i} \mathbb{E} \left\{ \frac{\sum_{l=0}^{\tau_i} \delta^l v_i(a_{t+l})}{\sum_{l=0}^{\tau_i} \delta^l} \right\},$$

where a_{t+l} is the path in which alternative i is chosen l times following a given past allocation (a_0, \dots, a_t) , and where the expectation is taken with respect to the realized posteriors $\theta_{i,t+l}$. 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:

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

In the dynamic direct mechanism, we construct a transfer price such that under the efficient allocation, each bidder's net payoff coincides with her flow marginal contribution $m_i(\theta_t)$. We consider first the payment of the bidder i who has the highest index in state θ_t and who should therefore receive the object in period t . In order to match her net payoff to her flow marginal contribution, we must have:

$$m_i(\theta_t) = v_i(h_{i,t}) - p_i(\theta_t). \quad (21)$$

The remaining bidders, $j \neq i$, should not receive the object in period t and their transfer price must offset the flow marginal contribution: $m_j(\theta_t) = -p_j(\theta_t)$. We expand the flow marginal contribution in (21) 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 :

$$m_i(\theta_t) = v_i(h_{i,t}) - v_k(h_{k,t}) - \delta (W_{-i}(\theta_{t+1} | i, \theta_t) - W_{-i}(\theta_{t+1} | k, \theta_t)).$$

We notice that with private values, the continuation value of the social program without i but conditional on the object being assigned to agent i in period t is simply equal to the value of the program conditional on θ_t alone, or

$$W_{-i}(\theta_{t+1} | i, \theta_t) = W_{-i}(\theta_t).$$

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:

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

It follows that the transfer price should simply be given by $p_i^*(\theta_t) = (1 - \delta) W_{-i}(\theta_t)$, which is the flow social opportunity cost of assigning the object today to agent i . A similar analysis

leads to the conclusion that the losing bidders makes zero payments: $p_j^*(\theta_t) = -m_j(\theta_t) = 0$. Our main result in this section collects this information on the transfers in the dynamic pivot 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:*

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

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 property 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. 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 all periods in which the ownership of the object does not change, the transfer price stays constant as well, even though the value of the object to the winning bidder may change.

References

- ATHEY, S., AND I. SEGAL (2007a): “Designing Efficient Mechanisms for Dynamic Bilateral Trading Games,” *American Economic Review Papers and Proceedings*, 97, 131–136.
- (2007b): “An Efficient Dynamic Mechanism,” Harvard University and Stanford University.
- BAPNA, A., AND T. WEBER (2005): “Efficient Dynamic Allocation with Uncertain Valuations,” Stanford University.
- BARON, D., AND D. BESANKO (1984): “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 267–302.
- 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 Among Self-Interested Agents with Private State,” in *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, Cambridge.
- DOLAN, R. (1978): “Incentive Mechanisms for Priority Queuing Problems,” *Bell Journal of Economics*, 9, 421–436.
- 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.
- GITTINS, J. (1989): *Allocation Indices for Multi-Armed Bandits*. London, Wiley.

- GREEN, J., AND J. LAFFONT (1977a): “Characterization of Satisfactory Mechanisms for the Revelation of the Preferences for Public Goods,” *Econometrica*, 45, 427–438.
- (1977b): “R v lation Des Pr f rences Pour Les Biens Publics: Caract risation Des M canismes Satisfaisants,” *Cahiers du S minaire d’ conom trie*, 13.
- LAVI, R., AND N. NISAN (2000): “Competitive Analysis of Incentive Compatible Online Auctions,” in *Proceedings of 2nd Conference of Electronic Commerce*, pp. 233–241, New York. ACM Press.
- MOULIN, H. (1986): “Characterization of the Pivotal Mechanism,” *Journal of Public Economics*, 31, 53–78.
- WHITTLE, P. (1982): *Optimization Over Time*, vol. 1. Wiley, Chichester.