

THE DYNAMIC PIVOT MECHANISM

BY DIRK BERGEMANN AND JUUSO VÄLIMÄKI¹

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. The dynamic pivot mechanism here is equivalent to a modified second price auction.

KEYWORDS: Pivot mechanism, dynamic mechanism design, ex post equilibrium, marginal contribution, multiarmed bandit, Bayesian learning.

1. INTRODUCTION

IN THIS PAPER, we generalize the idea of the pivot mechanism (due to Green and Laffont (1977)) 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 agent 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 agent receives as her payment her marginal contribution to the social welfare in every conceivable continuation game.

¹We thank the editor and four 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, Xianwen Shi, and Tomasz Strzalecki for many informative conversations. The authors acknowledge financial support through National Science Foundation Grants CNS 0428422 and SES 0518929, and the Yrjö Jahnsson’s Foundation, respectively.

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 (apart from the commitment to the transfers promised for the current period). In contrast, in revenue-maximizing problems, it is well known that the optimal solution relies critically on the ability of the principal to commit to a contract, see Baron and Besanko (1984). Interestingly, Battaglini (2005) showed that in dynamic revenue-maximizing problems with stochastic types, the commitment problems are less severe than with constant types.

The dynamic pivot mechanism yields a positive monetary surplus for the planner in each period and, therefore, the planner does not need outside resources to achieve the efficient allocation. Finally, the dynamic pivot mechanism induces all agents to participate in the mechanism after all histories.

In the intertemporal environment there are many 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 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* (see Lavi and Nisan (2000)). We establish 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 allocation analogue to the static single-unit auction. The scheduling problem is kept deliberately simple and all the relevant private information arrives in the initial period. Subsequently, we use the dynamic pivot mechanism to derive the dynamic auction format for a model where bidders learn their valuations for a single object over time. In contrast to the scheduling problem where a static mechanism could still have implemented the efficient solution, a static mechanism now necessarily fails to support the efficient outcome as more information arrives over time. In turn, this requires a more complete understanding of the intertemporal trade-offs in the allocation process. By computing the dynamic marginal contributions, we can 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. Among the contributions

which focus on socially efficient allocation, Cavallo, Parkes, and Singh (2006) proposed a Markovian environment for general allocation problems and analyzed two different classes of sequential incentives schemes: (i) Groves-like payments and (ii) pivot-like payments. They established that Groves-like payments, which award every agent positive monetary transfers equal to the sum of the valuation of all other agents, guarantee interim incentive compatibility and ex post participation constraints after all histories. In contrast, pivot-like payments guarantee interim incentive compatibility and ex ante participation constraints. Athey and Segal (2007) considered a more general dynamic model in which the current payoffs are allowed to depend on the entire past history including past signals and past actions. In addition, they also allowed for hidden action as well as hidden information. The main focus of their analysis is on incentive compatible mechanisms that are budget balanced in every period of the game. Their mechanism, called balanced team mechanism, transfers the insight from the Arrow (1979) and D'Aspremont and Gerard-Varet (1979) mechanisms into a dynamic environment. In addition, Athey and Segal (2007) presented conditions in terms of ergodic distributions over types and patients agents such that insights from repeated games can be employed to guarantee interim participation constraints. In contrast, we emphasize voluntary participation without any assumptions about the discount factor or the ergodicity of the type distributions. We also define an efficient exit condition which allows us to single out the dynamic pivot mechanism in the class of efficient mechanisms.

The focus of the current paper is on the socially efficient allocation, but a number of recent papers have analyzed the design of dynamic revenue-maximizing mechanisms, beginning with the seminal contributions by Baron and Besanko (1984) and Courty and Li (2000), who considered optimal intertemporal pricing policies with private information in a setting with two periods. Battaglini (2005) considered the revenue-maximizing long-term contract of a monopolist in a model with an infinite time horizon when the valuation of the buyer changes in a Markovian fashion over time. In particular, Battaglini (2005) showed that the optimal continuation contracts for a current high type are efficient, as his payoff is determined by the allocations for the current low type (by incentive compatibility). The net payoffs of the types then have a property related to the marginal contribution here. But as Battaglini (2005) considered revenue-maximizing contracts, the lowest type served receives zero utility, and hence the notion of marginal contribution refers only to the additional utility generated by higher types, holding the allocation constant, rather than the entire incremental social value. Most recently, Pavan, Segal, and Toikka (2008) developed a general allocation model and derived the optimal dynamic revenue-maximizing mechanism. A common thread in these papers is a suitable generalization of the notion of virtual utility to dynamic environments.

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 quasilinear in the monetary transfer:

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

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 the state space Θ_i . The aggregate state is given by the vector $\theta_t = (\theta_{1,t}, \dots, \theta_{I,t})$ with $\Theta = \times_{i=1}^I \Theta_i$.

There is a common prior $F_i(\theta_{i,0})$ regarding the initial type $\theta_{i,0}$ of each agent i . 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 $t = 0$. The common prior $F_i(\theta_{i,0})$ and the stochastic kernels $F_i(\theta_{i,t+1}; \theta_{i,t}, a_t)$ are assumed to be independent across agents. At the beginning of each period t , each agent 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 agent i , $\theta_{i,t+1}$, does not depend on $\theta_{j,t}$ for $j \neq i$. The expected absolute value of the flow payoff is assumed to be bounded by some $K < \infty$ for every i, a, θ and allocation plan $a' : \Theta \rightarrow A$:

$$\int |v_i(a'(\theta'), \theta'_i)| dF(\theta'; a, \theta) < K.$$

The nature of the state space Θ depends on the application at hand. At this point, we stress that the formulation accommodates the possibility of random arrival or departure of the agents. The arrival or departure of agent i can be represented by an inactive state $\underline{\theta}_i$, where $v_i(a_t, \underline{\theta}_i) = 0$ for all $a_t \in A$ and a random time τ at which agent i privately observes her transition in or out of the inactive state.

Social Efficiency

All agents discount the future with a common discount factor δ , $0 < \delta < 1$. The socially efficient policy is obtained by maximizing the expected discounted

sum of valuations. Given the Markovian structure, the socially optimal program starting in period t at state θ_t can be written as

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

For notational ease, we omit the conditioning state in the expectation operator, when the conditioning event is obvious, as in the above, where $\mathbb{E}[\cdot] = \mathbb{E}_{\theta_t}[\cdot]$. 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 $\mathbf{a}^* = \{a_i^*\}_{i=0}^\infty$. The social externality cost of agent i is determined by the social value in the absence of agent i :

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

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

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

The marginal contribution of agent i is the change in the social value due to the addition of agent i .²

Mechanism and Equilibrium

We focus attention on direct mechanisms which truthfully implement the socially efficient policy \mathbf{a}^* . 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 or may not be truthful. The public history in period t is 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 we assume that the past 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*

²In symmetric information environments, we used the notion of marginal contribution to construct efficient equilibria in dynamic first price auctions; see Bergemann and Välimäki (2003, 2006).

mechanism is represented by a family of allocations and monetary transfers, $\{a_t^*, p_t\}_{t=0}^\infty: a_t^*: \Theta \rightarrow \Delta(A)$, and $p_t: H_t \times \Theta \rightarrow \mathbb{R}^I$. With the focus on efficient mechanisms, the allocation a_t^* depends only on the current (reported) state $r_t \in \Theta$, while the transfer p_t may depend on the entire public history.

A (pure) reporting strategy for agent i in period t is a mapping from the private history into the state space: $r_{i,t}: H_{i,t} \rightarrow \Theta_i$. For a given mechanism, the expected payoff of agent i from reporting strategy $\mathbf{r}_i = \{r_{i,t}\}_{t=0}^\infty$ given the strategies $\mathbf{r}_{-i} = \{r_{-i,t}\}_{t=0}^\infty$ is

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

Given the mechanism $\{a_t^*, p_t\}_{t=0}^\infty$ and the reporting strategies \mathbf{r}_{-i} , the optimal strategy of bidder i can be stated recursively:

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

The value function $V_i(h_{i,t})$ represents the continuation value of agent i given the current private history $h_{i,t}$. 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 state 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 *periodic ex post* participation constraints of each agent. After each history h_t , each agent i may opt out (permanently) from the mechanism. The value of the outside option is denoted $O_i(h_{i,t})$ and it is defined by the payoffs that agent i receives if the planner pursues the efficient policy \mathbf{a}_{-i}^* for the remaining agents. 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 say that a mechanism is *ex post incentive compatible and individually rational* if it satisfies the periodic ex post 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 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

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

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 the constant valuations over time for all i , the optimal policy is given by assigning in every period the alternative j with the highest remaining valuation. To simplify notation, we define the positive valuation $v_i \triangleq \theta_{i,0}$. We may assume without loss of generality (after relabelling) that the valuations v_i are ordered with respect to the index i : $v_1 \geq \dots \geq v_I \geq 0$. Due to the descending order of valuations, we identify each task i with the period $i + 1$ in which it is completed along the efficient path:

$$(2) \quad W(\theta_0) = \sum_{t=1}^I \delta^{t-1} v_t.$$

Similarly, the efficient program in the absence of task i assigns the tasks in ascending order, but necessarily skips task i in the assignment process:

$$(3) \quad W_{-i}(\theta_0) = \sum_{t=1}^{i-1} \delta^{t-1} v_t + \sum_{t=i}^{I-1} \delta^{t-1} v_{t+1}.$$

By comparing the social program with and without i , (2) and (3), respectively, we find that the assignments for agents $j < i$ remain unchanged after i is removed, but that each agent $j > i$ 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

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

The social externality cost of agent i is established in a straightforward manner. At time $t = i - 1$, agent i completes her task and realizes the value v_i . The immediate opportunity cost is the next highest valuation v_{i+1} . But this overstates the externality, because in the presence of i , all less valuable tasks are realized one period later. The externality cost of agent i is hence equal to the next valuable task v_{i+1} minus the improvement in future allocations due to the delay of all tasks by one period:

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

Since we have by construction $v_t - v_{t+1} \geq 0$, 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, the final expression states that the externality of agent i is the cost of delay imposed on the remaining and less valuable tasks.³

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:

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

³In the online Supplementary Material (Bergemann and Välimäki (2010)), we show that the socially efficient scheduling can be implemented through a bidding mechanism rather than the direct revelation mechanism used here. In a recent and related contribution, Said (2008) used the dynamic pivot mechanism and a payoff equivalence result to construct bidding strategies in a sequence of ascending auctions with entry and exit of the agents.

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

$$(5) \quad m_i(\theta_t) \triangleq W(\theta_t) - W_{-i}(\theta_t) - \delta \mathbb{E}[W(\theta_{t+1}) - W_{-i}(\theta_{t+1})].$$

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 , $a_t^* \triangleq a^*(\theta_t)$ and $a_{-i,t}^* \triangleq a_{-i}^*(\theta_{-i,t})$, may differ. The continuation value of the socially optimal program, conditional on current allocation a_t and state θ_t is

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

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 $m_i(\theta_t)$ is expressed as

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

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:

$$(6) \quad p_i^*(\theta_t) \triangleq v_i(a_t^*, \theta_{i,t}) - m_i(\theta_t).$$

We refer to $p_i^*(\theta_t)$ as the transfer of the dynamic pivot mechanism. The transfer $p_i^*(\theta_t)$ depends only on the current report θ_t and not on the entire public history h_t . We can express $p_i^*(\theta_t)$ in terms of the flow utilities and the social continuation values:

$$(7) \quad p_i^*(\theta_t) = \sum_{j \neq i} [v_j(a_{-i,t}^*, \theta_{j,t}) - v_j(a_t^*, \theta_{j,t})] + \delta[W_{-i}(\theta_{t+1}|a_{-i,t}^*, \theta_t) - W_{-i}(\theta_{t+1}|a_t^*, \theta_t)].$$

The transfer $p_i^*(\theta_t)$ for agent i depends on the report of agent i only through the determination of the social allocation which is a prominent feature of the static Vickrey–Clarke–Groves mechanisms. The monetary transfers $p_i^*(\theta_t)$ are always nonnegative, 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

$$(8) \quad \begin{aligned} v_i(a^*(\theta_t), \theta_{i,t}) - p_i^*(\theta_t) + \delta M_i(\theta_{t+1}|a_t^*, \theta_t) \\ \geq v_i(a^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - p_i^*(r_{i,t}, \theta_{-i,t}) + \delta M_i(\theta_{t+1}|a^*(r_{i,t}, \theta_{-i,t}), \theta_t) \end{aligned}$$

for all $r_{i,t} \in \Theta_i$ and all $\theta_{-i,t} \in \Theta_{-i}$, and we recall that we denote the socially efficient allocation at the true state profile θ_t by $a_t^* \triangleq a^*(\theta_t)$. By construction of p_i^* in (7), the left-hand side of (8) represents the marginal contribution of agent i . We can express the marginal contributions $M_i(\cdot)$ in terms of the different social values to get

$$(9) \quad \begin{aligned} W(\theta_t) - W_{-i}(\theta_t) \\ \geq v_i(a^*(r_{i,t}, \theta_{-i,t}), \theta_{i,t}) - p_i^*(r_{i,t}, \theta_{-i,t}) \\ + \delta(W(\theta_{t+1}|a^*(r_{i,t}, \theta_{-i,t}), \theta_t) - W_{-i}(\theta_{t+1}|a^*(r_{i,t}, \theta_{-i,t}), \theta_t)). \end{aligned}$$

We then insert the transfer price $p_i^*(r_{i,t}, \theta_{-i,t})$ (see (7)) into (9) to obtain

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

But now we reconstitute the entire inequality in terms of the respective social values:

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

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

The dynamic pivot mechanism specifies a unique monetary transfer after every history. It guarantees that the ex post incentive and ex post participation constraints are satisfied after every history. 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) over 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 who 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.⁴

This motivates the following efficient exit condition. Let state θ_{τ_i} in period τ_i be such that the probability that agent i affects the efficient social decision a_t^* in period t is equal to zero for all $t \geq \tau_i$, that is, $\Pr(\{\theta_t | a_t^*(\theta_t) \neq a_{-i,t}^*(\theta_t)\} | \theta_{\tau_i}) = 0$. In this case, agent i is irrelevant for the mechanism in period τ_i , and we say that the mechanism satisfies the efficient exit condition if agents neither make nor receive transfers in periods where they are irrelevant for the mechanism.

DEFINITION 1—Efficient Exit: A dynamic direct mechanism satisfies the efficient exit condition if for all $i, h_{\tau_i}, \theta_{\tau_i}$,

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

We 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.

ASSUMPTION 1—Diverse Preferences:

(i) For all i , there exists $\underline{\theta}_i \in \Theta_i$ such that for all a ,

$$v_i(a, \underline{\theta}_i) = 0 \quad \text{and} \quad F_i(\underline{\theta}_i; a, \underline{\theta}_i) = 1.$$

(ii) For all i, a , and $x \in \mathbb{R}_+$, there exists $\theta_i^{a,x} \in \Theta_i$ such that

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

and for all a_t ,

$$F_i(\theta_i^{a,x}; a_t, \theta_i^{a,x}) = 1.$$

⁴We would like to thank an anonymous referee for the suggestion to consider the link between exit and uniqueness of the transfer rule.

The diverse preference assumption assigns to each agent i a state, $\underline{\theta}_i$, which is an absorbing state and in which i gets no payoff from any allocation. In addition, each agent i has a state in which i has a positive valuation x for a specific current allocation a and no value for other current or any future allocations. The diverse preferences condition is similar to the rich domain conditions introduced in Green and Laffont (1977) 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 $\underline{\theta}_i$. 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 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*

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

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

$$(10) \quad x \geq W_{-i}(\theta_{-i,t}) - \sum_{j \neq i} v_j(a, \theta_{j,t}) - \delta W_{-i}(\theta_{-i,t+1} | a, \theta_{-i,t}),$$

then $p_i^*(\theta_i^{a,x}, \theta_{-i,t})$ is equal to the above social externality cost; otherwise it is 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 (10), must pay the same transfer. To see this, suppose that for some $x, y \in \mathbb{R}_+$ satisfying (10), we have $p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) < p_i(h_t, \theta_i^{a,y}, \theta_{-i,t})$. Now type $\theta_i^{a,y}$ has a strict incentive to misreport $r_{i,t} = \theta_i^{a,x}$, a contradiction. We therefore denote the transfer for all x and $\theta_i^{a,x}$ satisfying (10) by $p_i(h_t, a, \theta_{-i,t})$, and denote 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 incentive constraints of a type $\theta_i^{a,x}$ with a valuation x such that

$$(11) \quad p_i(h_t, a, \theta_{-i,t}) < x < p_i^*(a, \theta_{-i,t}).$$

If the inequality (11) 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

$$(12) \quad v_i(a^*(\theta_i^{a,x}, \theta_{-i,t}), \theta_i^{a,x}) - p_i(h_t, \theta_i^{a,x}, \theta_{-i,t}) \geq v_i(a, \theta_i^{a,x}) - p_i(h_t, a, \theta_{-i,t}).$$

Given that $\theta_i = \theta_i^{a,x}$, we rewrite (12) 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 (11), 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 positive subsidy violates the ex post incentive constraint of the absorbing type $\underline{\theta}_i$. By the efficient exit condition, type $\underline{\theta}_i$ should not receive any contemporaneous (or future) subsidies. But by misreporting her type to be $\theta_i^{a,x}$, type $\underline{\theta}_i$ would gain access to a positive subsidy without changing the social allocation. 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 . Q.E.D.

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 in the dynamic pivot mechanism.

LEMMA 2: *If $\{a_i^*, p_i\}_{i=0}^\infty$ 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 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 $p_i(h_\tau, \theta_\tau) < p_i^*(\theta_\tau)$. We 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 the incentive constraint of a type $\theta_i^{a^*,x}$ with $p_i(h_\tau, \theta_\tau) < x < p_i^*(\theta_\tau)$ at a misreport $\theta_{i,\tau}$:

$$(13) \quad v_i(a^*(\theta_i^{a^*,x}, \theta_{-i,\tau}), \theta_i^{a^*,x}) - p_i(h_\tau, \theta_i^{a^*,x}, \theta_{-i,\tau}) + \delta V_i(h_{i,\tau+1} | a^*(\theta_i^{a^*,x}, \theta_{-i,\tau}), (\theta_i^{a^*,x}, \theta_{-i,\tau})) \geq v_i(a^*(\theta_{i,\tau}, \theta_{-i,\tau}), \theta_i^{a^*,x}) - p_i(h_\tau, \theta_\tau) + \delta V_i(h_{i,\tau+1} | a^*(\theta_{i,\tau}, \theta_{-i,\tau}), (\theta_i^{a^*,x}, \theta_{-i,\tau})).$$

By hypothesis, we have $p_i(h_\tau, \theta_\tau) < x < p_i^*(\theta_\tau)$ and if $x < p_i^*(\theta_\tau)$, then we can infer from marginal contribution pricing that $a^*(\theta_i^{a^*,x}, \theta_{-i,\tau}) \neq a^*(\theta_{i,\tau}, \theta_{-i,\tau})$. But as the type $\theta_i^{a^*,x}$ has only a positive valuation for $a^*(\theta_{i,\tau}, \theta_{-i,\tau})$, it follows

that the left-hand side of (13) is equal to zero. However, the right-hand side is equal to $v_i(a^*(\theta_{i,\tau}, \theta_{-i,\tau}), \theta_i^{a^*,x}) - p_i(h_\tau, \theta_\tau) = x - p_i(h_\tau, \theta_\tau) > 0$, leading to a contradiction.

Suppose next that for some $\varepsilon > 0$, we have

$$(14) \quad M_i(\theta_0) - V_i(\theta_0) > \varepsilon.$$

By hypothesis of ex post incentive compatibility, we have for all reports $r_{i,0}$,

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

Given a_0^* , we can find, by the diverse preference condition, a type $\theta_i = \theta_i^{a_0^*,x}$ such that $a_0^* = a^*(\theta_i^{a_0^*,x}, \theta_{-i,0})$. Now by Lemma 1, there exists a report $r_{i,0}$ for agent i , namely $r_{i,0} = \theta_i^{a_0^*,x}$, such that a_0^* is induced at the price $p_i^*(\theta_0)$. After inserting $r_{i,0} = \theta_i^{a_0^*,x}$ into the above inequality and observing that $v_i(a^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0}) - p_i(h_0, r_{i,0}, \theta_{-i,0}) = m_i(\theta_0)$, we conclude that $M_i(\theta_0) - V_i(h_{i,1}|a_0^*(r_{i,0}, \theta_{-i,0}), \theta_{i,0}) > \varepsilon/\delta$.

Now we repeat the argument we started with (14) and find that there is a path of realizations $\theta_0, \dots, \theta_t$, 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. *Q.E.D.*

The above lemma can be viewed as a revenue equivalence result 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 involves a fixed utility for the lowest type. In the infinite-horizon model here, the diverse preference assumption gives us a natural candidate of a lowest type in terms of $\underline{\theta}_i$, 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 the diverse preference condition is satisfied and if $\{a_i^*, p_i\}_{i=0}^\infty$ 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 an agent 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 $p_i(h_\tau, \theta_\tau)$ violates the ex post incentive constraint of some type $\theta_i^{a,x}$. Consider now a type $\theta_i^{a,x}$ with a valuation x for the allocation a_i^* such that $x > p_i^*(\theta_\tau)$. Her ex post incentive constraints are given by

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

for all $r_{i,t} \in \Theta_i$. By the efficient exit condition, we have for all $r_{i,t}$,

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

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

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. *Q.E.D.*

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 an 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 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 . 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 multiarmed bandit problem and the optimal policy is characterized by an index policy (see Whittle (1982)). In particular, we can compute for every bidder i the 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 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) . 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 $\mathbf{a}^* = \{a_t^*\}_{t=0}^\infty$ is to choose in every period a bidder i if $\gamma_i(h_{i,t}) \geq \gamma_j(h_{j,t})$ 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 $m_i(\theta_i)$. 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 . To match her net payoff to her flow marginal contribution, we must have

$$(15) \quad m_i(\theta_t) = v_i(h_{i,t}) - p_i(\theta_t).$$

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 $m_i(\theta_t)$ by noting that i is the efficient assignment and that another bidder, say k , would be the efficient assignment in the absence of i :

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

The continuation value without i in $t + 1$, but conditional on having assigned the object to i in period t , is simply equal to the value conditional on θ_t , 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. The flow marginal contribution of the winning agent i is, therefore,

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

It follows that $p_i^*(\theta_t) = (1 - \delta)W_{-i}(\theta_t)$, which is the flow externality cost of assigning the object to agent i . A similar analysis leads to the conclusion that each losing bidder makes zero payments: $p_j^*(\theta_t) = -m_j(\theta_t) = 0$.

THEOREM 3—Dynamic Second Price Auction: *The socially efficient allocation rule \mathbf{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

- ARROW, K. (1979): "The Property Rights Doctrine and Demand Revelation Under Incomplete Information," in *Economics and Human Welfare: Essays in Honor of Tibor Scitovsky*, ed. by M. Boskin. New York: Academic Press, 23–39. [773]
- ATHEY, S., AND I. SEGAL (2007): "An Efficient Dynamic Mechanism," Discussion Paper, Harvard University and Stanford University. [773]
- BARON, D., AND D. BESANKO (1984): "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1, 267–302. [772,773]
- BATTAGLINI, M. (2005): "Long-Term Contracting With Markovian Consumers," *American Economic Review*, 95, 637–658. [772,773]
- BERGEMANN, D., AND J. VÄLIMÄKI (2003): "Dynamic Common Agency," *Journal of Economic Theory*, 111, 23–48. [775]
- (2006): "Dynamic Price Competition," *Journal of Economic Theory*, 127, 232–263. [775]
- (2010): "Supplement to 'The Dynamic Pivot Mechanism'," *Econometrica Supplemental Material*, 78, http://www.econometricsociety.org/ecta/Supmat/7260_extensions.pdf. [778]
- 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. [773]
- COURTY, P., AND H. LI (2000): "Sequential Screening," *Review of Economic Studies*, 67, 697–717. [773]
- D'ASPROMONT, C., AND L. GERARD-VARET (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25–45. [773]
- DOLAN, R. (1978): "Incentive Mechanisms for Priority Queuing Problems," *Bell Journal of Economics*, 9, 421–436. [777]
- GREEN, J., AND J. LAFFONT (1977): "Characterization of Satisfactory Mechanisms for the Revelation of the Preferences for Public Goods," *Econometrica*, 45, 427–438. [771,782]
- LAVI, R., AND N. NISAN (2000): "Competitive Analysis of Incentive Compatible Online Auctions," in *Proceedings of the 2nd Conference of Electronic Commerce*. New York: ACM Press, 233–241. [772]
- MOULIN, H. (1986): "Characterization of the Pivotal Mechanism," *Journal of Public Economics*, 31, 53–78. [782]
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2008): "Dynamic Mechanism Design: Revenue Equivalence, Profit Maximization and Information Disclosure," Discussion Paper, Northwestern University and Stanford University. [773]
- SAID, M. (2008): "Auctions With Dynamic Populations: Efficiency and Revenue Maximization," Discussion Paper, Yale University. [778]

WHITTLE, P. (1982): *Optimization Over Time*, Vol. 1. Chichester: Wiley. [786]

*Dept. of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT
06520, U.S.A.; dirk.bergemann@yale.edu*

and

*Dept. of Economics, Helsinki School of Economics, Arkadiankatu 7, 00100
Helsinki, Finland; juuso.valimaki@hse.fi.*

Manuscript received July, 2007; final revision received September, 2009.

SUPPLEMENT TO “THE DYNAMIC PIVOT MECHANISM”
(*Econometrica*, Vol. 78, No. 2, March 2010, 771–789)

BY DIRK BERGEMANN AND JUUSO VÄLIMÄKI

We show that the socially efficient solution to the scheduling problem in Section 3 of the paper can be realized through a bidding mechanism, specifically a dynamic version of the ascending price auction, rather than a direct revelation mechanism. We also give a slight modification of the example where the bidding mechanism is inefficient.

IN THE SCHEDULING PROBLEM in Section 3 of the main paper, a number of bidders compete for a scarce resource, namely early access to a central facility. We show here 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 backward 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 backward 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

$$(S1) \quad b_I(v_I) = v_I - \delta(v_I - 0),$$

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

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

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

$$(S2) \quad b_{I-k}(v_{I-k}) = v_{I-k} - \delta(v_{I-k} - b_{I-(k-1)}(v_{I-(k-1)})).$$

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 (S2) recursively with the initial condition given by (S1), 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.¹

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 assume that $v_1 > v_2$. Under this assumption, the efficient policy never completes the task of bidder 2. The marginal contributions of each bidder are

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

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 represents the social externality cost of agent 1 in every period even though agent 2 never receives access to the facility.

¹The 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.

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 is efficient. In this case, bidder 2 is never chosen and hence receives 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 receives a net payoff of $v_1 - v_2$ in every period and her discounted sum of payoff is then

$$(S3) \quad \frac{1}{1 - \delta}(v_1 - v_2) < M_1(\theta_0).$$

But more important than the failure of the marginal contribution is the fact that the equilibrium does 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 drops 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

$$(S4) \quad \frac{\delta}{1 - \delta}v_1.$$

If we compare the continuation payoffs (S3) and (S4), 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 assigns 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 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 are 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.

REFERENCE

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.

Dept. of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06520, U.S.A.; dirk.bergemann@yale.edu

and

Dept. of Economics, Helsinki School of Economics, Arkadiankatu 7, 00100, Helsinki, Finland; juuso.valimaki@hse.fi.

Manuscript received July, 2007; final revision received September, 2009.