Working Paper

EFFICIENT DYNAMIC MECHANISMS IN INTERDEPENDENT VALUATION ENVIRONMENTS

Heng Liu*

In this paper, we provide a generic existence result for efficient mechanisms in dynamic interdependent valuation environments. We use the intertemporal correlation of different agents' private information to construct a sequence of payment schedules under which truth-telling consists of an ex post equilibrium after all histories. While previous results in static mechanism design with interdependent valuations are mostly negative, we show that intertemporal consideration can offer an easy and intuitive solution to efficient implementation.

1. INTRODUCTION

This paper studies the problem of implementing the socially efficient allocation in dynamic environments with private information and interdependent valuations. When agents' payoff functions are quasi-linear, the social objective is to maximize the expected discounted sum of the individual payoffs. Because agents' payoff-relevant private information may evolve over time and allocation decisions are made in each period, the main issue is to design a sequence of transfer payments for each agent so that she is willing to reveal her information truthfully at every stage.

In dynamic private valuation environments, Bergemann and Välimäki (2010) construct an efficient dynamic mechanism, which generalizes the well-known Vickrey-Clarke-Groves (VCG) mechanism, in the sense that each agent reports truthfully and receives her flow marginal contribution in each period. A key feature of the dynamic pivot mechanism in Bergemann and Välimäki (2010) is that in each period each agent is made a residual claimant under the payment schedule so that her incentive is aligned with the social objective. This is achieved by letting the agent pay the flow externality she imposes on other agents. When an agent's utility depends only on her own private information, this externality depends on her report only through the realized allocation. Therefore, truthful report in each period consists of an equilibrium.

Preliminary and incomplete.

Acknowledgment: I am grateful to Hari Govindan and Paulo Barelli for their guidance and encouragement. All errors are my own.

*Department of Economics, University of Rochester, heng.liu@rochester.edu

Although the private value assumption is realistic in some interesting situations, may fail in many important settings where each agent's utility depends not only on her own information but also on other agents' private information. In these interdependent valuation environments, the VCG mechanism and its dynamic extension fail to provide incentive for agents to reveal their private information. The problem with these mechanisms is that, in order to make each agent a residual claimant, the payment rule has to depend directly on all agents' reports, which generates incentives for an agent to manipulate her report. Since interdependent valuation or informational externality is common in many practical instances, it is important to understand whether there exists efficient mechanisms which are the interdependent-value counterparts of the VCG mechanism and the dynamic pivot mechanism.

In static models where agents only report their private information once, results from previous work (e.g. Dasgupta and Maskin (2000), Jehiel and Moldovanu (2001), Bergemann and Välimäki (2002), Jehiel, et. al. (2006)) indicate that it is in general impossible to ex post implement the efficient allocation when agents' preferences are interdependent and private information is multidimensional. Moreover, even if private information is one dimensional, efficient mechanisms can be constructed only under certain restrictions on the utility functions.[1] Therefore, interdependent valuations create a fundamental difficulty for efficient mechanism design problems, at least in static models.

In dynamic models, interdependent valuations may arise even if each agent's utility function depends only on her own information. For example, when agents are uncertain about the distributions of their future private information and all current private information can be used for inference, an agent's expected continuation utility and hence her total payoff may depend directly on other agents' information. This suggests that a slight modification of the dynamic private valuation environment would undermine the validity of the dynamic pivot mechanism. Furthermore, given that there would be a lot more incentive constraints because of intertemporal considerations in dynamic models, truthful implementation of the efficient allocation may seem rather challenging. Do the

[1]In auctions with interdependent valuation, under an appropriate single-crossing condition, Crémer and McLean (1985) first constructed a payment rule that generalized the Vickrey auction and implemented the efficient outcome. This payment rule, which was later generalized to more abstract mechanism design setting by Jehiel and Moldovanu (2001) and Bergemann and Välimäki (2002), is often referred to as the generalized VCG mechanism.

impossibility results in static setup persist in dynamic models? Under what conditions do efficient dynamic mechanisms exist? How restrictive are those conditions?

In this paper, we consider an intertemporal mechanism design setting with interdependent valuations where each agent's private information changes over time and allocation decisions are made in each period. We provide sufficient conditions under which there exist incentive compatible payment rules that truthfully implement the dynamic efficient allocations in dynamic interdependent valuation environments. Surprisingly, the conditions for the existence of efficient mechanisms in dynamic models are much less restrictive than those in static models. In contrast to the generic impossibility results in static models, our results suggest that efficient allocations are implementable under an identification condition on the evolution of private information, even if each agent's private information is multidimensional in most periods. The identification condition, which essentially requires different agents' private information to be correlated intertemporally, is satisfied for an open set of parameters for the transition functions. Therefore, the possibility of long-term interactions provides a solution to the efficient mechanism design problem with interdependent valuation.

The intuition of our efficient dynamic mechanism is simple. Suppose that an agent's private signal today is correlated with other agents' signals tomorrow. Then, given that the dynamic mechanism could elicit all agents' true signals tomorrow, we can amend the original payment schedule tomorrow with a carefully chosen function of these signals so that, in expectation, the agent's total payoff today is equal to the social welfare less a term that is independent of her own reports. Therefore, just like in the VCG mechanism and dynamic pivot mechanism in private value models, the agent, as a residual claimant, is willing to report her private information truthfully today.

While previous work on dynamic mechanism design often assumes that each agent's private information evolves independently in order to avoid the possibility of full surplus extraction mechanism à la Crémer and McLean (1988), the efficient mechanism constructed in this paper suggests that there exists a different way to utilize information correlation in dynamic environments, which would not appear in static settings.

When the identification condition is violated, exploiting the intertemporal correlation of signals may not be possible for some utility functions. In this case, we restrict the analysis to one dimensional private information and construct a sequence of payment schedules which is the dynamic counterpart of the payment rule in the generalized VCG mechanism (e.g. Jehiel and Moldovanu (2001), Bergemann and Välimäki (2002)). Furthermore, the sequence of payments reduces to the dynamic pivot mechanism in Bergemann and Välimäki (2010) when there is no interdependence of preferences. In this one dimensional setting, we also identify sufficient conditions, which are generalizations of the single-crossing condition in static models, for these payment rules to be incentive compatible.

The rest of the paper is organized as follows. In Section 2 we review the literature and compare the present paper with related work. In Section 3 we describe the model and define dynamic efficiency and mechanisms. In Section 4, we first use a simple example to illustrate the main ideas of our mechanisms, then we present the main results in greater generality. Extensions are considered in Section 5. Section 6 concludes. All formal proofs are relegated to the Appendix.

2. RELATED LITERATURE

This paper is related to the following strands of literature.

Efficient mechanisms with interdependent valuations. In static quasilinear environments with interdepedent valuation, previous results have pointed that efficient mechanism desgin is generically impossible. Dasgupta and Maskin (2000), and Jehiel and Moldovanu (2001) prove that efficient allocations cannot be implemented by any Bayesian incentive compatible mechanisms, if private signals are multidimensional and statistically independent. Jehiel, et. al. (2006) further show that if a more robust notion of implementation–ex post implementation–is required, then only constant allocation rules are ex post incentive compatible in generic frameworks with multidimensional signals. Given these negative results, dynamic implementaion of efficiency seems rather hopeless. Yet our main result suggests that efficient mechanism design is possible in a large class of dynamic environments. One crucial feature of the dynamic mechanism design framework is that private information may evolve over time. Agents' private signals may be correlated over time even they are conditionally indepedently distributed in each period. We show precisely how to use the intertemporal correlation to construct payment rules that implement the dynamic efficient allocation, without imposing much structure on the signal space.

Conceptually, our mechanism is related to the two-stage VCG mechanism in Mezzetti (2004). Mezzetti provides one way to get over the impossibility result and implement the efficient allocation, under the assumption that all agents can observe their realized utilities and transfers can be made based on the reported utilities. However, the assumptions in Mezzetti (2004) are not very realistic in many situations. Another drawback with Mezzetti's mechanism is that agents are indifferent among all messages when they are asked to report their realized utilities, since each agent's transfer payment only depends on other agents' reported utilities. If reporting the realized utilities is costly for agents, then no matter how small the cost is, agents would want to walk away from the mechanism in this stage. Our efficient dynamic mechanism indicates that, instead of assuming realized payoffs are observable, it is sufficient to have some intertemporal correlation in different agents' private signals. Since agents maximize their continuation payoffs in dynamic environments, delayed payments are constructed so that an agent's expected continuation valuation is equal to the social surplus in each period. In our mechanism, truth-telling consists of a perfect equilibrium and no agent is ever indifferent among all messages at any time.

On the other hand, a few postive results on the existence of ex post incentive compatible efficient mechanisms are established in a very small class of environments. In interdependent-value auctions with one dimensional private information, Crémer and McLean (1985) construct the generalized VCG mechanism and show that it is ex post incentive compatible under an additional single-crossing condition on the utility functions.[2] Jehiel and Moldovanu (2001) and Bergemann and Välimäki (2002) later extend the generalized VCG mechanism to abstract mechanism design settings, again assuming one dimensional signals and certain single-crossing conditions. This paper also extends these postive findings to dynamic environments with one dimensional signals. With certain single-crossing assumptions, we construct a sequence of transfer payments that mirrors the generalized VCG payments and implements the dynamic efficient allocation.

Dynamic mechanism design. Recent papers by Bergemann and Välimäki (2010) and Athey and Segal (2012) have constructed efficient dynamic mechanisms in private valuation environments. They focus on finding the dynamic counterparts of the classic VCG mechanism and the expected externality mechanism in static mechanism design.

[2]See also Dasgupta and Maskin (2000).

The main departure of our model from previous research is that we study a general model with interdependent valuations. As the existing results on efficient design with private valuations are not applicable in our setup, we propose conceptually a quite different way solving the design problem. Our mechanism also applies to private-value models under the same assumptions.

There is also a large body of literature on revenue maximizing dynamic mechanisms. Early studies (Baron and Besanko (1984), Courty and Li (2000), Battaglini (2005) ,Eso and Szentes (2007)) consider several special models and characterize the optimal mechanism under regularity assumptions. A recent work (Pavan, Segal and Toikka (2012)) extends the first order approach for incentive compatibility in static models to dynamic settings with evolving private information and provides general sufficient conditions for incentive compatibility and revenue equivalence.[3] However, none of these papers considers the possibility of correlated information.

Full surplus extraction. The idea of exploiting correlated information in our dynamic mechanism has some similarity with the results on full surplus extraction (e.g. Crémer and McLean (1985), Crémer and McLean (1988), McAfee and Reny (1992)). The idea of full surplus extraction is to construct lottery payments as entrance fees so that if agents' private signals are correlated, the mechanism designer's expected revenue from the mechanism is equal to the entire social surplus. The focus of this paper is efficient design. We consider another type of correlated information and use lotteries or contingent payments to make each agent a residual claimant, and hence provides incentive for truth-telling in each period. Our construction is based on a different generalization of the convex independence condition in Crémer and McLean (1988).

3. MODEL

3.1. *The Environment*

We consider a dynamic interdependent valuation environment with N ($N \geq 2$) agents. Time is discrete, indexed by $t \in \{1, 2, \dots, T\}$, where $T < \infty$.[4] In each period t, each agent $i \in \{1, 2, \dots, N\}$ privately observes a state variable $\theta_t^i \in \Theta_t^i$, where Θ_t^i is a compact subset

[3]See also the recent papers of Battaglini and Lamba (2912), Boleslavsky and Said (2013), Deb and Said (2012), Pai and Vohra (2012) and Skrzypacz and Toikka (2012).

[4]We assume T is finite in the current draft. Results for the infinite horizon case will be incorporated in the later version of the paper.

of a metric space endowed with the Borel σ-algebra. The state space in period t is $\Theta_t = \prod_{i=1}^N \Theta_t^i$ with a generic element $\theta_t = (\theta_t^1, \ldots, \theta_t^N)$. For each i and t, denote the private information held by agents other than i in period t by $\theta_t^{-i} = (\theta_t^1, \ldots, \theta_t^{i-1}, \theta_t^{i+1}, \ldots, \theta_t^N) \in \prod_{j \neq i} \Theta_t^j$.

In each period t, the flow utility of agent i is determined by the current state profile θ_t, the current allocation $a_t \in A$ and the current monetary transfer $p_t^i \in \mathbb{R}$, where A is a finite set of social alternatives. The flow utility of each agent is assumed to be quasilinear in monetary transfers and agents have a common discount factor $\delta \in (0,1)$. Given sequences of states $\{\theta_t\}_{t=1}^T$, allocations $\{a_t\}_{t=1}^T$ and monetary transfers $\{p_t^1, \ldots, p_t^N\}_{t=1}^T$, the total payoff of each agent i is

$$\sum_{t=1}^{T} \delta^{t-1} \left[u^i(a_t, \theta_t) - p_t^i \right],$$

where the functions $u^i : A \times \Theta_t \to \mathbb{R}_+$ are bounded and measurable. Let $M \in \mathbb{R}_+$ denote a bound for all u_i's, i.e., for every i, t, a_t and θ_t,

$$u^i(a_t, \theta_t) < M.$$

The private information of each agent evolves over time and is modeled as a Markov process. Specifically, in the initial period, the state profile θ_1 follows a prior probability measure $\mu_1 \in \Delta(\Theta_1)$. In any period $t > 1$, the distribution of current state profile θ_t is determined by the realized state profile θ_{t-1} and the allocation decision a_{t-1} in the previous period, which is represented by a transition probability $\mu_t : A_{t-1} \times \Theta_{t-1} \to \Delta(\Theta_t)$. The utility functions u_i, the prior μ_1 and the transition probabilities μ_t are assumed to be common knowledge.

Notice that in contrast to previous work which often assumes that the prior distribution and the transitions are independent across agents, here we specify a rather general Markov process for the evolution of states, which allows correlation of private information. While the existence of efficient mechanisms does not depend on whether correlation is allowed or not in private valuation environment as shown by Athey and Segal (2012), it will be clear in Section 4 how and what type of correlation is going to make a difference for efficient implementation with interdependent valuations.

3.2. *Efficiency and Mechanism*

In our setup, a socially efficient allocation rule is a sequence of measurable mappings[5] $\{a_t^* : \Theta_t \to A\}_{t=1}^T$ which solves the following social program

$$\max_{\{a_t\}_{t=1}^T} \mathbb{E}\left[\sum_{t=1}^T \delta^{t-1} \sum_{i=1}^N u^i(a_t, \theta_t)\right].$$

Since the flow utility depends only on current state profile which is assumed to be Markovian, the social program can also be written in a recursive form. Specifically, for each $t \in \{1, 2, \ldots, T\}$

$$W_t(\theta_t) = \max_{a_t \in A} \sum_{i=1}^N u^i(a_t, \theta_t) + \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) | a_t, \theta_t\right], \tag{1}$$

where $W_t(\theta_t)$ is the social surplus starting from period t with realized state profile θ_t and $W_{T+1} \equiv 0$. By the principle of optimality, a_t^* is also a solution to this recursive problem.

By the revelation principle (Myerson (1986)), we focus on truthful equilibria of direct mechanisms which implement the socially efficient allocations $\{a_t^*\}_{t=1}^T$. In a direct mechanism, in each period t, each agent i is asked to make a public report $r_t^i \in \Theta_t^i$ of her current private state θ_t^i. Then a public allocation decision a_t and a transfer payment p_t^i for each agent i are made as functions of the current report profile $r_t = (r_t^i)_{i=1}^N$ and the period-t public history h_t. The period-t public history contains all reports and allocations up to period $t-1$, i.e.,

$$h_t = (r_1, a_1, r_2, a_2, \ldots, r_{t-1}, a_{t-1}).$$

Let H_t denote the set of possible period-t public histories. Formally, an *efficient direct revelation mechanism* $\Gamma = \{\Theta_t, a_t^*, p_t\}_{t=1}^T$ consists of (i) a message space Θ_t in each period t; (ii) a sequence of allocation rules $a_t^* : \Theta_t \to A$, and (iii) a sequence of monetary transfers $p_t : H_t \times \Theta_t \to \mathbb{R}^N$.

The period-t private history h_t^i of each agent i contains the period-t public history and the sequence of her realized private states until period t, i.e.,

$$h_t^i = (r_1, a_1, \theta_1^i, r_2, a_2, \theta_2^i, \ldots, r_{t-1}, a_{t-1}, \theta_{t-1}^i, \theta_t^i).$$

[5] We assume that the program always has a solution so the efficient rule is always well-defined. Note that the solution to the program may not be unique due to the possibility of indifference. If this case, we simply choose a measurable selection of the solution correspondence.

Let H_t^i denote the set of agent i's possible period-t private histories. With a slight abuse of notation, a behavioral strategy for agent i in period t is a mapping $r_t^i : H_t^i \to \Theta_t^i$ assigning a report to each of her period-t private history. A strategy for agent i is *truthful* if it always reports agent i's private state θ_t^i truthfully in each period t, regardless of her private history.

Given a mechanism $\Gamma = \{\Theta_t, a_t^*, p_t\}_{t=1}^T$ and a strategy profile $r = \{(r_t^i)_{i=1}^N\}_{t=1}^T$, agent i's expected discounted payoff is

$$\mathbb{E} \sum_{t=1}^{T} \delta^{t-1} \left[u^i(a_t^*(r_t), \theta_t) - p_t^i(h_t, r_t) \right].$$

The equilibrium concept we adopt is periodic ex post equilibrium defined by Bergemann and Välimäki (2010) and Athey and Segal (2012). We say that the mechanism is *periodic ex post incentive compatible*, or equivalently, the truthful strategy profile is a *periodic ex post equilibrium* if for each agent and in each period, truth-telling is always a best response regardless of the private history and the current state of other agents given that other agents adopt truthful strategies. Formally, let $V_t^i(h_t^i)$ be agent i's continuation payoff given period-t private history, given that other agents always report truthfully. That is,

$$V_t^i(h_t^i) = \max_{r_t^i \in \Theta_t^i} \mathbb{E} \left[u^i(a_t^*(r_t^i, \theta_t^{-i}), \theta_t) - p_t^i(h_t, r_t^i, \theta_t^{-i}) + \delta V_{t+1}^i(h_{t+1}^i) \right], \tag{2}$$

with $V_{T+1}^i \equiv 0$. The efficient mechanism is *periodic ex post incentive compatible* if for each i, t and h_t^i,

$$\theta_t^i \in \arg\max_{r_t^i \in \Theta_t^i} u^i(a_t^*(r_t^i, \theta_t^{-i}), \theta_t) - p_t^i(h_t, r_t^i, \theta_t^{-i}) + \delta \mathbb{E} \left[V_{t+1}^i(h_{t+1}^i) | a_t^*(r_t^i, \theta_t^{-i}), \theta_t \right],$$

for each $\theta_t \in \Theta_t$.

As suggested by Bergemann and Välimäki (2010), ex post incentive compatibility notions need to be qualified within each period in a dynamic environment, as an agent may wish to change her report in some previous round based on the new information she has received in later periods. Since dominant strategy equilibrium is too strong with interdependent valuations, periodic ex post incentive compatibility is probably the most robust solution concept we could achieved in our setup.

With some exgenously given outside options for each agent after each private history, we can define the periodic ex post individual rationality condition in a similar way as in

Bergemann and Välimäki (2010). The mechanism is *periodic ex post individually rational* if in the truthful equilibrium each agent's payoff is no less than her outside option after any history.

Finally, budget balancedness notions can also be defined naturally. The mechanism is *ex ante budget balanced* if

$$\mathbb{E}\left[\sum_{t=1}^{T}\delta^{t-1}\sum_{i=1}^{N}p_t^i\right]\geq 0.$$

The mechanism is *budget balanced* if for each t,

$$\sum_{i=1}^{N}p_t^i\equiv 0.$$

In situations where the mechanism designer can use outside financing, a balanced budget requires the expected present value of all transfer payments from agents to be nonnegative. Without outside financing, a balanced budget means that in each period agents' transfer payments sum to zero.

4. EFFICIENT MECHANISM DESIGN

4.1. *A Two-Period Example*

Consider the following two-round repeated auction example. Two bidders, i and j, compete for a non-storable good in each of the two periods. Suppose there is no discounting. In each period $t\in\{1,2\}$, bidder i (j) observes a private signal θ_t^i (θ_t^j). The allocation in period t is $a_t\in\{i,j\}$ where $a_t=i$ means that agent i obtains the good and $a_t=j$ otherwise. The flow payoffs are $\mathbf{1}_{\{a_t=i\}}(\theta_t^i+\gamma\theta_t^j)-p_t^i$ for bidder i and $\mathbf{1}_{\{a_t=j\}}(\theta_t^j+\gamma\theta_t^i)-p_t^j$ for bidder j, where $\mathbf{1}_{\{\cdot\}}$ is an indicator function and $\gamma<1$ is a constant. Signals in the first period $\theta_1=(\theta_1^i,\theta_1^j)$ are drawn from the joint distribution $F_1(\theta_1)=F_1^i(\theta_1^i)F_1^j(\theta_1^j)$. Given θ_1 and a_1, signals in the second period $\theta_2=(\theta_2^i,\theta_2^j)$ are drawn from the transition kernel $F_2(\theta_2|a_1,\theta_1)=F_2^i(\theta_2^i|a_1,\theta_1)F_2^j(\theta_2^j|a_1,\theta_1)$.

An efficient allocation (a_1^*,a_2^*) maximizes the social surplus. In particular, $a_2^*(\theta_2)=i$ if $\theta_2^i>\theta_2^j$, and the efficient social surplus in period 2 is

$$W_2(\theta_2)=\mathbf{1}_{\{a_2^*=i\}}(\theta_2^i+\gamma\theta_2^j)+\mathbf{1}_{\{a_2^*=j\}}(\theta_2^j+\gamma\theta_2^i).$$

In period 1, $a_1^*(\theta_1) = i$ if

$$(\theta_1^i + \gamma\theta_1^j) + \mathbb{E}[W_2(\theta_2)|i, \theta_1] > (\theta_1^j + \gamma\theta_1^i) + \mathbb{E}[W_2(\theta_2)|j, \theta_1].$$

Given $\gamma < 1$, in the second period, the generalized VCG mechanism

$$p_2^i = 0, p_2^j = (1+\gamma)\theta_2^i \quad \text{if } \theta_2^i < \theta_2^j; \quad p_2^i = (1+\gamma)\theta_2^j, p_2^j = 0, \quad \text{if } \theta_2^i \geq \theta_2^j$$

is ex post incentive compatible in period 2. Under this payment rule, the flow payoffs of bidders i and j in the second period are $V_2^i = \mathbf{1}_{\{\theta_2^i \geq \theta_2^j\}}(\theta_2^i - \theta_2^j)$ and $V_2^j = \mathbf{1}_{\{\theta_2^i \leq \theta_2^j\}}(\theta_2^j - \theta_2^i)$, respectively.

Suppose the payment rule in period 2 is the generalized VCG mechanism momentarily and consider the incentive problem in the first period. Assume bidder j always reports truthfully and fix $\theta_1^j = 0.5$. Assume that the transition kernel $F_2(\theta_2|a_1, \theta_1)$ is specified as in Table I. Then the efficient allocation in the first period satisfies $a_1^* = j$ if $\theta_1^i = 0.4$,

TABLE I

TRANSITION KERNEL $F_2(\theta_2|a_1, \theta_1)$

(θ_1^i, a_1)		θ_2^i	θ_2^j
$(0.4, i)$	$\to$	uniform$[1, 2]$	uniform$[0, 1]$
$(0.4, j)$	$\to$	uniform$[0, 1]$	uniform$[1, 2]$
$(0.6, i)$	$\to$	uniform$[0, 1]$	uniform$[1, 2]$
$(0.6, j)$	$\to$	uniform$[1, 2]$	uniform$[0, 1]$

and $a_1^* = i$ if $\theta_1^i = 0.6$. Now consider the first period payment rule for agent i. We need to find $p_1^i(0.4)$ and $p_1^i(0.6)$ so that the following two incentive constraints are satisfied.

$$0.6 + 0.5\gamma - p_1^i(0.6) + 0 \geq 0 - p_1^i(0.4) + 1$$
$$0 - p_1^i(0.4) + 0 \geq 0.4 + 0.5\gamma - p_1^i(0.6) + 1$$

However, summing up the two inequalities gives $0.2 \geq 2$, which is impossible. The reason for this impossibility is that given the generalized VCG payment in period 2, bidder i has a strong incentive to misreport in period 1 in order to gain advantage in period 2. Notice that an implicit assumption in the above mechanism is that the generalized VCG payment rule in the second period $p_2 : \Theta_2 \to \mathbb{R}^2$ is a function of current reports only and is independent of the reports and allocation in the first period. While focusing on this kind of history independent payment rules for efficient implementation is without loss in private valuation models as illustrated in Bergemann and Välimäki (2010) and Athey and Segal (2012), it may be too restrictive in the interdependent valuation setup.

In fact, agent i's signal in the first period is statistically correlated with agent j's signal in the second period according to the transition kernel in Table I. We can indeed exploit this correlation by adding another history-dependence payments to the generalized VCG payments in period 2 to obtain incentive compatibility in period 1. Consider the following period-2 transfer payment schedule $\bar{p}_2^i : \Theta_2 \times A_1$ for bidder i

$$\bar{p}_2^i(\theta_2; a_1) = p_2^i(\theta_2) + \tilde{p}_2^i(\theta_2^j; a_1),$$

where p_2^i is the generalized VCG payment defined above and $\tilde{p}_2^i$ satisfies

$$\tilde{p}_2^i(a_1^* = i) = -1.5 - 0.5\gamma, \ \text{ if } \theta_2^j \in [1,2]; \quad \tilde{p}_2^i(a_1^* = i) = -0.5 - 0.5\gamma, \ \text{ if } \theta_2^j \in [0,1];$$
$$\tilde{p}_2^i(a_1^* = j) = -2 - 0.9\gamma, \ \text{ if } \theta_2^j \in [1,2]; \quad \tilde{p}_2^i(a_1^* = j) = -1 - 1.1\gamma, \ \text{ if } \theta_2^j \in [0,1].$$

Since the generalized VCG payment p_2 is incentive compatible in period 2 and $\tilde{p}_2^i$ is independent of bidder i's report, period-2 ex post incentive compatibility still holds under $\bar{p}_2^i$. Set $p_1^i \equiv 0$, then bidder i's incentive constraints in period 1 are:

$$0.6 + 0.5\gamma + 1.5 + 0.5\gamma \geq 0.5 + 0.6\gamma + 1.5 + 0.5\gamma,$$
$$0.5 + 0.4\gamma + 1.5 + 0.5\gamma \geq 0.4 + 0.5\gamma + 1.5 + 0.5\gamma,$$

which hold obviously given the assumption $\gamma < 1$.

The intuition for this modified mechanism is the following. By exploring the intertemporal correlation between θ_1^i and θ_2^j conditional on any a_1, the payment rule $\bar{p}_2^i$ simply makes each agent a claimant for the entire social surplus in the first period with out affecting her incentive for truthful report in the second period. Given bidder j adheres to truthful strategies, it is optimal for bidder i to be truthful so as to maximize the social surplus and hence her own payoff.

Finally, we note that bidders' signals are assumed to be conditionally independent in this example and the only correlation in signals comes from the dynamic aspect of the model. Without any correlation, it is impossible to construct such a "VCG-style" mechanism whenever preferences are interdependent.

4.2. *Main Results*

In this section, we first consider the existence of periodic ex post incentive compatible efficient dynamic mechanisms under general state transition dynamics. Theorem 4.1 shows that under a generic intertemporal correlation condition (Assumption 1) and a uniform lower bound condition (Assumption 2) on the transition probabilities, and some restrictions on utility functions and state spaces in the last period (Assumption 3),[6] such a dynamic mechanism always exists. If the correlation condition is violated, the construction of the dynamic mechanism in the proof of Theorem 4.1 may not work for some utility functions. In this case, we restrict our attention to a one-dimensional setting (Assumption 4) and construct a payment schedule which extends the generalized VCG mechanism to dynamic settings. Theorem 4.4 shows that our dynamic generalized VCG mechanism achieves periodic ex post incentive compatibility under several restrictions on utilities and transitions.

First consider the following assumptions on the relation of different agents' private information over time.

Assumption 1 (Identification condition) For any t, i, a_t, θ_t, there exists no measurable function $h : \Theta_t^i \to \mathbb{R}$ such that for any measurable subset $\tilde{\Theta}_{t+1}^{-i} \subset \Theta_{t+1}^{-i}$,

$$\int_{\Theta_t^i} \mu_{t+1}^{-i}(\tilde{\Theta}_{t+1}^{-i}|a_t, \theta_t^i, \theta_t^{-i}) h(\theta_t^i) \mu_t^i(d\theta_t^i|a_{t-1}, \theta_{t-1}) = 0,$$

where μ_t^i and μ_t^{-i} are the marginals of μ_t on Θ_t^i and Θ_t^{-i}, respectively.

Assumption 2 (Uniform lower bound) There exists $\varepsilon > 0$ such that for each $i, t, a_{t-1}, a_t, \theta_{t-1}$ and θ_t, we have

$$\mu_{t+1}^{-i}(\tilde{\Theta}_{t+1}^{-i}|a_t, \theta_t^i, \theta_t^{-i}) \geq \varepsilon \int_{\Theta_t^i} \mu_{t+1}^{-i}(\tilde{\Theta}_{t+1}^{-i}|a_t, \theta_t^i, \theta_t^{-i}) \mu_t^i(d\theta_t^i|a_{t-1}, \theta_{t-1}),$$

for each measurable subset $\tilde{\Theta}_{t+1}^{-i} \subset \Theta_{t+1}^{-i}$.[7]

Assumption 1 says that each of the transition probability exhibits intertemporal correlation among different agents' signals. In particular, conditional on each a_t and θ_t^{-i}, agent i's current private state θ_t^i is correlated with other agents' private states θ_{t+1}^{-i} in the next period. Independent evolution of private information across agents is ruled out

[6] For the infinite horizon case, we do not need to impose any restriction on utilities or state spaces.

[7] For $t = 1$, we use $\mu_1^i(a_0, \theta_0) \equiv \mu_1^i$ to simplify notations.

by this assumption. Assumption 2 is a technical assumption, which essentially requires that conditional on any a_t and θ_t^{-i}, the transition from θ_t^i to θ_{t+1}^{-i} admits a density that is bounded from below.

To motivate the information correlation in Assumption 1, suppose that there is a underlying state of nature ω_t with possible values in a compact set Ω in each period t. In addition, ω_t follows a hidden Markov process which evolves over time and is not observed by any agent. In each period t, the relationship between the state of nature ω_t and agents' private information θ_t is described by a joint distribution ξ_t over $\Omega \times \Theta_t$. If each agent's private information θ_t^i provides useful information about ω_t, i.e., the conditional $\xi_t(\omega_t|\theta_t^i)$ varies with θ_t^i, then as long as ω_t is not independently distributed, θ_t^i is correlated with θ_{t+1}^{-i} even conditional on θ_t^{-i} and a_t.

If for each i and each t, Θ_t^i is a subset of a finite-dimensional Euclidean space and if each transition probability $\mu_t(a_{t-1}, \theta_{t-1})$ have a density representation $f_t(\theta_t|a_{t-1}, \theta_{t-1})$ which is strictly positive, then Assumptions 1 and 2 are equivalent to the following requirements on the transition densities.

Assumption 1' (Identification condition with densities) For any t, i, a_t, θ_t, there exists no measurable function $h : \Theta_t^i \to \mathbb{R}$ such that for almost all θ_{t+1}^{-i},

$$\int_{\Theta_t^i} f_{t+1}^{-i}(\theta_{t+1}^{-i}|a_t, \theta_t^i, \theta_t^{-i}) h(\theta_t^i) d\theta_t^i = 0.$$

Where $f_{t+1}^{-i}(\cdot|a_t, \theta_t)$ is the marginal density of $f_{t+1}(\cdot|a_t, \theta_t)$ on Θ_{t+1}^{-i}.

Assumption 2' (Uniform lower bound on densities) There exists $\varepsilon > 0$ such that for each $t > 1, a_{t-1}, \theta_{t-1}$ and θ_t, the transition density $f_t(\theta_t|a_{a-1}, \theta_{t-1})$ satisfies

$$f_t(\theta_t|a_{a-1}, \theta_{t-1}) \geq \varepsilon,$$

If each state space Θ_t^i is finite, then fix any pair (a_t, θ_t^{-i}), each $f_{t+1}^{-i}(\theta_{t+1}^{-i}|a_t, \theta_t^i, \theta_t^{-i})$ is an element of a $|\Theta_{t+1}^{-i}| \times |\Theta_t^i|$ stochastic transition matrix. In this case, Assumption 1' is similar to the convex independence condition in Crémer and McLean (1988), which requires the transition matrix to have full column rank. Assumption 2' requires each probability distribution to have full support on. Unlike Crémer and McLean (1988), we focus on ex post efficient mechanisms in this paper and leave the analysis of dynamic surplus extraction for later research.

Next we impose some conditions on the last-period primitives to guarantee ex post implementability in the last period. Note that in our setup, the allocation problem in period T is a static one. Thus, we simply adopt a set of sufficient conditions from the existing research (Bergemann and Välimäki (2002) in particular) on static mechanism design.

Let the finite allocation space be $A = \{a^1, \dots, a^K\}$. Fix any i and θ_T^{-i}. For each $k = 1, \dots, K$ define the set $\Theta_T^{i,k} \subset \Theta_T^i$ as

$$\Theta_T^{i,k} = \left\{ \theta_T^i \in \Theta_T^i \,\middle|\, \sum_i u^i(a^k, \theta_T) \geq \sum_i u^i(a^l, \theta_T), \forall a^l \neq a^k \right\}.$$

We say the collections of sets $\{\Theta_T^{i,k}\}_{k=1}^K$ satisfies *monotonicity* if for each k, $\theta_T^i, \tilde{\theta}_T^i \in \Theta_T^{i,k}$ implies that for each $\lambda \in [0,1]$, $\lambda\theta_T^i + (1-\lambda)\tilde{\theta}_T^i \in \Theta_T^{i,k}$. That is, $\{\Theta_T^{i,k}\}_{k=1}^K$ is monotone if each $\Theta_T^{i,k}$ is connected. Under monotonicity, there exists an efficient allocation a_T^* in period T such that after relabeling the social alternatives, Θ_T^i can be partitioned into successive intervals $\{S_T^{i,1}, \dots, S_T^{i,K}\}$ and each a^k is chosen if and only if $\theta_T^i \in S_T^{i,k}$. Then there is a linear order $\prec_T^i$ (which also depends on θ_T^{-i}) on A: $a^1 \prec_T^i \dots \prec_T^i a^K$.

Assumption 3 (Implementation conditions in period T**)** For each i and θ_T, monotonicity holds. Moreover, for each i and θ_T, the utility function $u^i(a^k, \theta_T^i, \theta_T^{-i})$ is supermodular in (a^k, θ_T^i).

The sufficient conditions for ex post efficient implementation in static models are restrictive given the impossibility results in Dasgupta and Maskin (2000) and Jehiel and Moldovanu (2001). In particular, period-T signals are one dimensional and utility functions satisfy a single-crossing condition under Assumption 3. If we relax ex post incentive compatibility to Bayesian incentive compatibility in period T, then there exist less restrictive conditions when period-T signals are correlated.

We also emphasize that no assumption is made on the private signals from period 1 to $T-1$. Any θ_t^i $(t < T)$ can be a vector in a multidimensional or even infinite dimensional space. We can think of a situation where agents trade an new asset with each other in multiple periods. Initially, each agent's private information may be multidimensional since there is much uncertainty about many aspects of the asset. As agents trade over time, they gradually learn about the asset. And in the last period, each agent's signal is simply a real number which represents her estimation of the asset value.

Now we state and prove the main result in this paper which generalizes the idea of the example in Section 4.1.

THEOREM 4.1 *Under Assumptions 1,2 and 3, there exists an efficient dynamic mechanism which is periodic ex post incentive compatible.*

PROOF: See Appendix, Section 7.2. □

REMARK 4.2 The identification condition on the joint transition probability (Assumption 1) is generically satisfied even if we assume that signals are independently distributed conditional on the underlying state of nature in each period. Accordingly, efficient dynamic mechanisms can be constructed in a very large class of dynamic settings provided that ex post implementability is achievable in the last period.[8] Therefore, contrary to the generic impossibility results in static environments, an implication of Theorem 4.1 is that efficient dynamic mechanisms "always" exist even preferences are interdependent. Repeated interactions in a sense help us achieve efficiency rather than hampering it.

REMARK 4.3 From a theoretical point of view, the dynamic mechanism constructed in the proof of Theorem 4.1 is closely related to the classical VCG mechanism and the dynamic pivot mechanism. As we explained in the introduction, with interdependent valuations the idea of making agents residual claimants is not applicable if they interact only once. When agents interact for multiple periods, our dynamic mechanism uses (one-period) delayed payments to link different agents' information over time so that each agent's expected continuation payoff in each period is equal to the continuation social surplus plus a term that is independent of her own report, when all agents adopt the truthful strategy.

Suppose that Assumption 1 is violated, then in a restricted environment, we have the following result which extends the generalized VCG mechanism in static one dimensional models to dynamic settings. We say that a payment rule $\{p_t\}_{t=1}^T$ or a mechanism $(\{a_t^*, p_t\}_{t=1}^T)$ is *history-independent* if for each t and θ_t, and for any two public histories

[8] A later version of the paper will show that Assumption 2 is not needed for Theorem 4.1 to hold in the infinite horizon case.

h_t and h'_t,

$$p_t(h_t, \theta_t) = p_t(h'_t, \theta_t).$$

That is, a history-independent payment p_t depends only on the reported profile $r_t \in \Theta_t$ in period t. Under a history-independent mechanism, agent i's period-t continuation payoff (2) depends only on her private state θ_t^i, i.e.,

$$V_t^i(\theta_t^i) = \max_{r_t^i \in \Theta_t^i} \mathbb{E}\left[u^i(a_t^*(r_t^i, \theta_t^{-i}), \theta_t) - p_t^i(r_t^i, \theta_t^{-i}) + \delta V_{t+1}^i(\theta_{t+1}^i)\right].$$

In this case, we also define $V_t^i(a_t, \theta_t)$ as

$$V_t^i(a_t, \theta_t) = u^i(a_t, \theta_t) + \delta \mathbb{E}\left[V_{t+1}^i(\theta_{t+1}^i) | a_t, \theta_t\right].$$

Assumption 4 (One dimensional private states) For each i and each t, $\Theta_t^i = [0, 1]$.

Under Assumption 4, we can define monotonicity and related notions similar to those considered in Assumption 3. For any i, t and θ_t^{-i}. Define the set $\Theta_t^{i,k} \subset \Theta_t^i$ as

$$\Theta_t^{i,k} = \left\{ \theta_t^i \in \Theta_t^i \;\middle|\; \begin{array}{l} u^i(a^k, \theta_t) + \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) | a^k, \theta_t\right] \\ \quad \geq \sum_i u^i(a^l, \theta_t) + \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) | a^l, \theta_t\right], \end{array} \forall a^l \neq a^k \right\}.$$

The collections of sets $\{\Theta_t^{i,k}\}_{k=1}^K$ satisfies *monotonicity* if for each k, $\theta_t^i, \tilde{\theta}_t^i \in \Theta_t^{i,k}$ implies that for each $\lambda \in [0, 1]$, $\lambda \theta_t^i + (1 - \lambda) \tilde{\theta}_t^i \in \Theta_t^{i,k}$. Under monotonicity, there exists an efficient allocation a_t^* in period t such that after relabeling the social alternatives, Θ_t^i can be partitioned into successive intervals $\{S_t^{i,1}, \ldots, S_t^{i,K}\}$ and each a^k is chosen if and only if $\theta_t^i \in S_t^{i,k}$. Then for each i, t and θ_t^{-i}, there is a linear order $\prec_t^i$ (which also depends on θ_t^{-i}) on A:

$$a^1 \prec_t^i \ldots \prec_t^i a^K.$$

Assumption 5 (Independent transitions) For $t = 1$, $\mu_1 = \prod_{i=1}^N \mu_1^i$, where for each i, $\mu_1^i \in \Delta(\Theta_1^i)$. For each $t > 1$, $\mu_t(a_{t-1}, \theta_{t-1}) = \prod_{i=1}^N \mu_t^i(a_{t-1}, \theta_{t-1}^i)$, where for each i, $\mu_t^i : A \times \Theta_{t-1}^i \to \Delta(\Theta_t^i)$ is a transition probability.

Suppose $a_t^*(\theta_t) = a^k$, then consider the following history-independent transfer payment

$$p_t^{i*}(\theta_t) = \sum_{\kappa=1}^{k} \sum_{j\neq i} \left[u^j(a^{\kappa-1}, x_t^i(\kappa, \theta_t^{-i}), \theta_t^{-i}) - u^j(a^{\kappa}, x_t^i(\kappa, \theta_t^{-i}), \theta_t^{-i})\right] \\ + \sum_{\kappa=1}^{k} \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) - V_{t+1}^i(\theta_{t+1}) | a^{\kappa-1}, x_t^i(\kappa, \theta_t^{-i}), \theta_t^{-i}\right] \\ - \sum_{\kappa=1}^{k} \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) - V_{t+1}^i(\theta_{t+1}) | a^{\kappa}, x_t^i(\kappa, \theta_t^{-i}), \theta_t^{-i}\right], \tag{3}$$

where $x_t^i(\kappa, \theta_t^{-i}) = \inf\{\theta_t^i : a_t^*(\theta_t^i, \theta_t^{-i}) = a^\kappa\}$. Note that $p_t^{i*}(\theta_t)$ does not depend directly on θ_t^i under Assumption 5.

Finally, we define a function that will be used in the next result. Recall that $W_t(\theta_t)$ is the continuation social surplus given period-t state profile θ_t. For each a_t and θ_t, define $W_t(a_t, \theta_t)$ as

$$W_t(a_t, \theta_t) = \sum_{i=1}^{N} u^i(a_t, \theta_t) + \delta \mathbb{E}\left[W_{t+1}(\theta_{t+1}) | a_t, \theta_t\right].$$

In Theorem 4.4, we show that the payment rule constructed in (3) is periodic ex post incentive compatible under some restrictions. Therefore, it is an appropriate extension of the generalized VCG mechanism.

THEOREM 4.4 *Suppose that Assumptions 4 and 5 hold. There exists a periodic ex post incentive compatible mechanism (a_t^*, p_t) with history-independent payment rules if for each t, i and θ_t^{-i}, there exists an order on the set of allocations A such that*

1. *$W_t(a_t, \theta_t^i, \theta_t^{-i})$ is single-crossing in (a_t, θ_t^i),*
2. *$V_t^i(a_t, \theta_t^i, \theta_t^{-i})$ has increasing difference in (a_t, θ_t^i)*

PROOF: See Appendix, Section 7.3. □

REMARK 4.5 The payment rule (3) is also a generalization of the dynamic pivot mechanism proposed by Bergemann and Välimäki (2010). To see this, suppose that each u^i does not depend on θ_t^{-i} and that private information is statistically independent across agent, then the payment rule in equation (3) can be written as

$$\begin{aligned} p_t^{i*}(\theta_t) = &\sum_{j\neq i}\left[u^j(a_t^*(\underline{\theta}^i,\theta^{-i}),\theta_t^{-i}) - u^j(a_t^*(\theta_t),\theta_t^{-i})\right] \\ &+ \delta\mathbb{E}\left[W^{-i}(\theta_{t+1})|a_t^*(\underline{\theta}^i,\theta^{-i}),\theta_t\right] \\ &- \delta\mathbb{E}\left[W^{-i}(\theta_{t+1})|a_t^*(\theta_t),\theta_t\right], \end{aligned}$$

where

$$W^{-i}(\theta_t) = W(\theta_t) - V^i(\theta_t) = \max_{\{a_s\}_{s\geq t}} \mathbb{E}\left[\sum_{s\geq t}\delta^{s-t}\left(u^i(a_s,\underline{\theta}^i) + \sum_{j\neq i}u^j(a_s,\theta_s^j)\right)\right].$$

Therefore each agent i's payment $p_t^{i^*}$ in every period t is the flow externality cost she imposes on other agents.

Remark 4.6 We have explored on the existence of history-independent efficient mechanisms in Theorem 4.4. History-independent mechanisms are appealing in practice as they do not depend on the entire history of allocations and reports. So the mechanism designer does not need to keep track of what happened in the past when deciding current payments. In addition, we can show that focusing on history-independent rules are without loss of generality under Assumption 5. Specifically, if there is a periodic ex post incentive compatible mechanism $\{a_t^*, p_t\}_{t=1}^T$ with history-dependent payments, then there exists a sequence of history-independent payments $\{\bar{p}_t\}_{t=1}^T$ such that the resulting mechanism $\{a_t^*, \bar{p}_t\}_{t=1}^T$ is also periodic ex post incentive compatible.

5. EXTENSIONS

Many questions remain to be explored. Here we discuss a few extensions that will be formalized in later versions of the paper.

First, if the time horizon is infinite, there exists efficient dynamic mechanisms that are ex post incentive compatible under Assumptions 1 and 2. The construction in the proof of Theorem 4.1 can be extended to the infinite horizon case without imposing Assumption 3. The idea is again to use the intertemporal correlation of signals with one-period delayed transfers such that an agent's transfer payment tomorrow in today's expectation is equal to the sum of other agents' flow utility today. Using the one-shot deviation principle, we can show that truth-telling strategy is still an periodic ex post equilibrium in the infinite horizon case.

Second, our efficient dynamic mechanism always runs a budget deficit, since the designer needs to make subsidies in each period in order to make each agent a claimant of the social surplus. Therefore, it is important to understand whether we can construct efficient dynamic mechanisms that are either budget balanced or ex ante budget balanced, depending on different instances. Assuming signals are conditionally independent across agents in each period, we would like to explore whether or how the balanced team mechanism in Athey and Segal (2012) can be adapted to our setup. It is even more interesting to find necessary and sufficient conditions for our mechanism to be periodic ex post individually rational and (ex ante) budget balanced.

Third, we would like to incorporate arrivals and departures of agents into consideration. With interdependent valuations, allowing new agents to join or existing agents to leave will alter the amount of information available to the mechanism designer and all current participating agents. One way to handle the departure of an agent is to integrate out her information from the moment of her departure.

Fourth, since we consider direct revelation mechanisms in this paper, it is of interest to consider whether our efficient dynamic mechanism can be implemented in a more practical manner. This question is particularly relevant in dynamic auctions.

Finally, while our focus is dynamic efficiency here, it would be interesting to consider problems related to revenue maximization. For example, given that efficiency is achievable, can the mechanism designer fully extract the surplus with some carefully designed lotteries as in Crémer and McLean (1985)? This question is not as straightforward as previous ones. One easy case is to assume private information is correlated aross agents in each period and there is no intertemporal correlation. Then applying the Crémer-McLean mechanism repeatedly would give full surplus extraction in each period if we consider interim incentive compatibility in each period. Suppose instead that private signals are conditionally independently distributed in each period but are intertemporal correlated as in our setup, can we get full surplus extraction with a periodic ex post incentive compatible mechanism? More generally, it is interesting to characterize conditions for full surplus extraction in dynamic environments.

6. CONCLUSION

The main result in this paper is that efficient mechanisms generically exist in dynamic interdependent valuation environments. This is in striking contrast to the impossibility results in static settings. We construct an efficient dynamic mechanism based on the intertemporal correlation of different agents' private information, a feature that does not exist in static models. An important implication of the result is that information from future interactions can be used to influence agents' behavior in the current interaction.

7. APPENDIX

7.1. *A Useful Lemma*

We first prove a lemma (Lemma 7.1) which generalizes the characterization of convex independence in Crémer and McLean (1988) to infinite type spaces.

Let $(S, \mathcal{S})$ and $(T, \mathcal{T})$ be compact metric spaces endowed with the corresponding Borel σ-algebras. Let λ be a regular probability measure on $\mathcal{S} \otimes \mathcal{T}$ with marginals λ_s on $\mathcal{S}$ and λ_t on $\mathcal{T}$. Let $\mu : T \times \mathcal{S} \to [0, 1]$ denote a version of the transition probability From T to S. Then for any $S' \in \mathcal{S}$ and $T' \in \mathcal{T}$, we have $\lambda(S' \times T') = \int_{T'} \int_{S'} \mu(t, ds) d\lambda_t$. Similarly, let $\nu : S \times \mathcal{S} \to [0, 1]$ be a version of the transition probability From S to T.

LEMMA 7.1 *Assume that there exists $\varepsilon > 0$ such that for each $s \in S$ and each $T' \in \mathcal{T}$, we have $\nu(s, T') \geq \varepsilon \lambda_t(T')$. For any bounded and measurable function $\pi : T \to \mathbb{R}_+$, there exists a bounded measurable function $p : S \to \mathbb{R}_+$ such that for λ_t-a.e. $t \in T$,*

$$\pi(t) = \int_S p(s) \mu(t, ds)$$

if and only if there does not exist a non-zero signed measure η on $(T, \mathcal{T})$ that is absolutely continuous with respect to λ_t such that

$$\int_T \int_{S'} \mu(t, ds) \eta(dt) = 0,$$

for any $S' \in \mathcal{S}$.

PROOF: For the only if part, suppose by contradiction that there exists a non-zero signed measure η that is absolutely continuous with respect to λ_t such that

$$\int_T \int_{S'} \mu(t, ds) \eta(dt) = 0,$$

for any $S' \in \mathcal{S}$. By absolute continuity, there exists a measurable function $h : T \to \mathbb{R}$ such that $h = d\eta/d\lambda_t$ for λ_t-a.e. $t \in T$. Using Fubini's Theorem, we have

$$0 = \int_T \int_{S'} \mu(t, ds) h(t) d\lambda_t = \int_{S'} \int_T h(t) \nu(s, dt) d\lambda_s.$$

for any $S' \in \mathcal{S}$. Therefore, for λ_s-a.e. $s \in S$, we have

$$\int_T h(t) \nu(s, dt) = 0.$$

Since for any bounded and measurable function $\pi : T \to \mathbb{R}_+$, there exists a bounded and measurable function $p : S \to \mathbb{R}_+$ such that

$$\pi(t) = \int_S p(s) \mu(t, ds)$$

for λ_t-a.e. $t \in T$, by Fubini's Theorem, we have

$$\begin{aligned}
\int_T \pi(t) \eta(dt) &= \int_T \pi(t) h(t) d\lambda_t \\
&= \int_T \int_S p(s) \mu(t, ds) h(t) d\lambda_t \\
&= \int_S p(s) \int_T h(t) \nu(s, dt) d\lambda_s \\
&= \int_S p(s) \cdot 0 d\lambda_s = 0.
\end{aligned}$$

Since π is an arbitrary measurable function, it follows that $\eta(t) = 0$ for λ_t-a.e. $t \in T$, which is a contradiction.

Next consider the if part. The proof is in two steps. We first prove that the result holds for any simple function $\pi(t)$. Then we use an approximation argument to show that the result holds for any measurable function $\pi(t)$.

Step 1. Let $\pi(t)$ be a nonnegative simple function, i.e., $\pi(t) = \sum_{i=1}^n \pi_i \mathbf{1}_{T_i}$, where $(\pi_1, \ldots, \pi_n) \in \mathbb{R}_+^n$ and $\{T_1, \ldots, T_n\}$ is a finite partition of T where for each i, $\lambda_t(T_i) > 0$. For each $i = 1, \ldots, n$, define a mapping $\tau_i : S \to [0, 1]$ by

$$\tau_i(s) \triangleq \frac{\int_{T_i} \nu(s, dt)}{\lambda_t(T_i)} \geq \varepsilon$$

for each $s \in S$. Consider the set $W \subset \mathbb{R}_+^n$ defined as

$$W = \left\{ w \in \mathbb{R}_+^n : \begin{array}{l} \exists\ p : S \to \mathbb{R}_+ \text{ bounded \& measurable s.t.} \\ w = \left(\int_S p(s) \tau_1(s) d\lambda_s, \ldots, \int_S p(s) \tau_n(s) d\lambda_s \right) \end{array} \right\}.$$

Note that W is closed and convex.

If $(\pi_1, \ldots, \pi_n) \in W$, that is, there exists a measurable function $p : S \to \mathbb{R}_+$ such that for each $i = 1, \ldots, n$, $\pi_i = \int_S p(s)\tau_i(s)d\lambda_s$,then for λ_t-a.e. $t \in T$,

$$\pi(t) \; = \; \int_S p(s)\mu(t, ds).$$

Moreover, since $\tau_i(s) \geq \varepsilon$ for each i and s, it follows that for λ_s-a.e. $s \in S$,

$$p(s) \leq \frac{1}{\varepsilon} \cdot \max\{\pi_1, \ldots, \pi_n\}. \tag{4}$$

On the other hand, suppose that $(\pi_1, \ldots, \pi_n) \notin W$, then by the Separating Hyperplane Theorem, there exists a non-zero vector $(\eta_1, \ldots, \eta_n) \in \mathbb{R}^n$ such that for any measurable function $p : S \to \mathbb{R}_+$, we have

$$\eta_1 \cdot \int_S p(s)\tau_1(s)d\lambda_s + \cdots + \eta_n \cdot \int_S p(s)\tau_n(s)d\lambda_s = 0.$$

By the definition of τ_i's, we have

$$\begin{aligned}
0 &= \int_S p(s) \sum_{i=1}^n \eta_i \tau_i(s) d\lambda_s \\
&= \int_S p(s) \sum_{i=1}^n \eta_i \int_{T_i} \nu(s, dt) d\lambda_s \\
&= \int_S p(s) \int_T \left(\sum_{i=1}^n \eta_i \mathbf{1}_{T_i} \right) \nu(s, dt) d\lambda_s \\
&= \int_T \int_S p(s)\mu(t, ds) \left(\sum_{i=1}^n \eta_i \mathbf{1}_{T_i} \right) d\lambda_t \\
&\triangleq \int_T \int_S p(s)\mu(t, ds)\eta(dt),
\end{aligned}$$

where in the last equality the signed measure η on $(T, \mathcal{T})$ is absolutely continuous with respect to λ_t. Since $p : S \to \mathbb{R}_+$ is arbitrary, it follows that for any $S' \in \mathcal{S}$,

$$\int_T \int_{S'} \mu(t, ds)\eta(dt) = 0,$$

which contradicts the hypothesis that there does not exist such a signed measure η.

Step 2. Let $\pi : T \to \mathbb{R}_+$ be any bounded and measurable function. Let $K < \infty$ be a bound on π, i.e., $\|\pi\| = \sup_{t \in T} |\pi(t)| < K$. Then there exists a sequence of simple functions $\{\pi_n(t)\}$ such that $\|\pi_n\| \leq K$ and for each t,

$$\lim_{n\to\infty} \pi_n(t) = \pi(t).$$

By the Bounded Convergence Theorem, for any $T' \in \mathcal{T}$,

$$\lim_{n\to\infty} \int_{T'} \pi_n(t) d\lambda_t = \int_{T'} \pi(t) d\lambda_t.$$

By Step 1, for each n, there exists a bounded measurable function $p_n : S \to \mathbb{R}_+$ such that for λ_t-a.e. $t \in T$,

$$\pi_n(t) = \int_S p_n(s)\mu(t, ds).$$

Since each simple function π_n satisfies $\|\pi_n\| < K$, by the inequality (4) in step 1, we have

$$p_n(s) \leq \frac{K}{\varepsilon}, \tag{5}$$

for λ_s-a.e. $s \in S$. Therefore, for each n we have

$$\int_{T'} \pi_n(t) d\lambda_t = \int_{T'} \int_S p_n(s)\mu(t, ds) d\lambda_t = \int_S p_n(s) \int_{T'} \nu(s, dt) d\lambda_s = \int_S p_n(s)\nu(s, T') d\lambda_s.$$

Note that $p_n \in L^\infty(S, \mathcal{S}, \lambda_s)$ and for each $T' \in \mathcal{T}$, $\nu(\cdot, T') \in L^1(S, \mathcal{S}, \lambda_s)$. From inequality (5), it follows that $\|p_n\| = \operatorname{ess\,sup}_{s\in S} |p_n(s)| \leq K/\varepsilon$ for every n. Thus by Alaoglu's Theorem,[9] the sequence $\{p_n\}$ lies in a weak-* compact subset of L^∞. By passing to a subsequence, it follows that there exists $p_\infty \in L^\infty(S, \mathcal{S}, \lambda_s)$ with $\|p_\infty\| \leq K/\varepsilon$ such that for any $T' \in \mathcal{T}$

$$\lim_{n\to\infty} \int_S p_n(s)\nu(s, T') d\lambda_s = \int_S p_\infty(s)\nu(s, T') d\lambda_s.$$

Note that for every $t \in T$,

$$\int_S p_\infty(s)\mu(t, ds) \leq \int_S \frac{K}{\varepsilon}\mu(t, ds) = \frac{K}{\varepsilon}. \tag{6}$$

[9]See Theorem 6.21 in Aliprantis and Border (2006), page 235.

Applying Fubini's Theorem again gives

$$\int_S p_\infty(s)\nu(s,T')d\lambda_s = \int_S p_\infty(s)\int_{T'}\nu(s,dt)d\lambda_s = \int_{T'}\int_S p_\infty(s)\mu(t,ds)\lambda_t.$$

Therefore, we have

$$\int_{T'}\pi(t)d\lambda_t = \int_{T'}\int_S p_\infty(s)\mu(t,ds)\lambda_t,$$

for any $T' \in \mathcal{T}$. Hence,

$$\pi(t) = \int_S p_\infty(s)\mu(t,ds)$$

for λ_t-a.e. $t \in T$.

□

7.2. *Proof of Theorem 4.1*

The proof consists of two lemmas. The first lemma (Lemma 7.2) shows that if the transition probability satisfies the condition in Assumption 1, then any bounded measurable function can be represented as a linear combination of the transition probability.

LEMMA 7.2 *For each i and each $t < T$, consider any bounded and measurable function $\pi^i : A \times \Theta_t^i \times \Theta_t^{-i} \to \mathbb{R}_+$. Then under Assumption 1 and Assumption 2, there exists a bounded and measurable function $\tilde{p}^i_{t+1} : \Theta_{t+1}^{-i} \times A \times \Theta_t^{-i} \to \mathbb{R}_+$ such that*

$$\pi^i(a_t,\theta_t^i,\theta_t^{-i}) = \int_{\Theta_{t+1}} \tilde{p}^i_{t+1}(\theta_{t+1}^{-i}, a_t, \theta_t^{-i})d\mu_{t+1}(a_t,\theta_t),$$

for every a_t, θ_t^{-i} and $\mu_t^i(a_{t-1},\theta_{t-1})$-a.e. $\theta_t^i \in \Theta_t^i$, where $\mu_t^i(a_{t-1},\theta_{t-1})$ is the marginal of $\mu_t(a_{t-1},\theta_{t-1})$ on Θ_t^i.

PROOF: Given any $i, t, a_{t-1}, \theta_{t-1}$, and any bounded measurable function $\pi^i(a_t,\theta_t^i,\theta_t^{-i})$, note that for any pair (a_t,θ_t^{-i}), $\pi^i(a_t,\theta_t^i,\theta_t^{-i})$ is measurable in θ_t^i. By Assumption 1, there is no measurable function $h_t^i(\theta_t)$ such that

$$\int_{\Theta_t^i} \mu_{t+1}^{-i}(\tilde{\Theta}_{t+1}^{-i}; a_t, \theta_t^i, \theta_t^{-i})h(\theta_t^i)\mu_t^i(d\theta_t^i; a_{t-1},\theta_{t-1}) = 0,$$

for any measurable subset $\tilde{\Theta}^{-i}_{t+1} \subset \Theta^{-i}_{t+1}$. Then Assumption 2 allows us to invoke Lemma 7.1, which implies that for each fixed pair (a_t, θ^{-i}_t) there exists a bounded function $q^i_{t+1}(\theta^{-i}_{t+1}; a_t, \theta^{-i}_t)$ which is measurable in θ^{-i}_{t+1} such that

$$\pi^i(a_t, \theta^i_t, \theta^{-i}_t) = \int_{\Theta^{-i}_{t+1}} q^i_{t+1}(\theta^{-i}_{t+1}; a_t, \theta^{-i}_t) d\mu^{-i}_{t+1}(a_t, \theta_t)$$

for $\mu^i_t(a_{t-1}, \theta_{t-1})$-a.e. $\theta^i_t \in \Theta^i_t$. Since $q^i_{t+1}(\theta^{-i}_{t+1}; a_t, \theta^{-i}_t)$ is independent of θ^i_{t+1}, we have

$$\pi^i(a_t, \theta^i_t, \theta^{-i}_t) = \int_{\Theta_{t+1}} q^i_{t+1}(\theta^{-i}_{t+1}; a_t, \theta^{-i}_t) d\mu_{t+1}(a_t, \theta_t)$$

for $\mu^i_t(a_{t-1}, \theta_{t-1})$-a.e. θ^i_t.

Applying the measurable selection theorem in Mertens (2003), there exists a bounded measurable function $\tilde{p}^i_{t+1}(\theta^{-i}_{t+1}, a_t, \theta^{-i}_t)$ on $\Theta^{-i}_{t+1} \times A \times \Theta^{-i}_t$ such that

$$\pi^i(a_t, \theta^i_t, \theta^{-i}_t) = \int_{\Theta_{t+1}} \tilde{p}^i_{t+1}(\theta^{-i}_{t+1}, a_t, \theta^{-i}_t) d\mu_{t+1}(a_t, \theta_t),$$

for for every a_t, θ^{-i}_t and $\mu^i_t(a_{t-1}, \theta_{t-1})$-a.e. θ^i_t.

□

The next lemma (lemma 7.3) establishes the existence result in Theorem 4.1.

LEMMA 7.3 *Under Assumptions 1, 2 and 3, there exists a sequence of payment rules $\bar{p}_t : H_t \times \Theta_t \to \mathbb{R}^N$ such that the efficient dynamic mechanism $\{a^*_t, \bar{p}_t\}^T_{t=1}$ is periodic ex post incentive compatible.*

PROOF: The proof is by backward induction. Let $W_t(\theta_t)$ denote the expected period-t continuation social surplus given state profile θ_t, i.e.,

$$W_t(\theta_t) = \mathbb{E}\left[\sum_{i=1}^N u^i(a^*_t(\theta_t), \theta_t)\right].$$

First consider the problem in period T. By Assumption 3, there exists an ex post incentive compatible payment $p_T : \Theta_T \to \mathbb{R}^N$ that implements the efficient allocation a^*_T. Given (a^*_T, p^*_T), the payoff V^i_T for each agent i in the truth-telling equilibrium is given by

$$V^i_T(\theta_T) = u^i(a^*_T(\theta_T), \theta_T) - p^i_T(\theta_T),$$

for each θ_T.

Next consider agent i's incentive problem in period $T-1$ with an arbitrary public history $h_{T-1} = (r_1, a_1, r_2, a_2, \ldots, r_{t-1}, a_{t-1})$. Suppose that agents other than i always report truthfully. For each pair (a_{T-1}, θ_{T-1}), define

$$\pi^i_{T-1}(a_{T-1}, \theta_{T-1}) = \sum_{j \neq i} u^j(a_{T-1}, \theta_{T-1}) + \delta \mathbb{E}\left[W(\theta_T) - V^i_T(\theta_T) | a_{T-1}, \theta_{T-1}\right].$$

Since the resulting function π^i_{T-1} is nonnegative, bounded and measurable, by Lemma 7.2 there exists a bounded measurable function $\tilde{p}^i_T(\theta^{-i}_T, a_{T-1}, \theta^{-i}_{T-1})$ such that for every a_{T-1}, θ^{-i}_{T-1} and $\mu^i_{T-1}(a_{T-2}, r_{T-2})$-a.e. θ^i_{T-1},

$$\pi^i_{T-1}(a_{T-1}, \theta_{T-1}) = \delta \int_{\Theta_T} \tilde{p}^i_T(\theta^{-i}_T, a_{T-1}, \theta^{-i}_{T-1}) d\mu_T(a_{T-1}, \theta_{T-1}).$$

Define a new period-T payment rule $\bar{p}^i_T : \Theta^{-i}_{T-1} \times A_{T-1} \times \Theta_T \to \mathbb{R}$ for agent i as

$$\bar{p}^i_T(\theta^{-i}_{T-1}, a_{T-1}, \theta_T) = p^i_T(\theta_T) - \tilde{p}^i_T(\theta^{-i}_T, a_{T-1}, \theta^{-i}_{T-1}).$$

Note that $\tilde{p}^i_T$ is independent of θ^i_T, so agent i still finds it optimal to report truthfully in period T under this new payment $\bar{p}^i_T$. Suppose agent i reports r^i_{T-1} in period $T-1$, then for any realized state profile θ_{T-1}, her expected continuation payoff from $T-1$ on is $\mu^i_{T-1}(a_{T-2}, r_{T-2})$-a.e. equal to

$$u^i(a^*_{T-1}(r^i_{T-1}, \theta^{-i}_{T-1}), \theta_{T-1}) + \delta \mathbb{E}\left[V^i(\theta_T) | a^*_{T-1}(r^i_{T-1}, \theta^{-i}_{T-1}), \theta_{T-1}\right] + \pi^i_{T-1}(a^*_{T-1}(r^i_{T-1}, \theta^{-i}_{T-1}), \theta_{T-1})$$

$$= \sum_{i=1}^N u^i(a^*_{T-1}(r^i_{T-1}, \theta^{-i}_{T-1}), \theta_{T-1}) + \delta \mathbb{E}\left[W_T(\theta_T) | a^*_{T-1}(r^i_{T-1}, \theta^{-i}_{T-1}), \theta_{T-1}\right].$$

Denote the exceptional set by $\hat{\Theta}^i_{T-1}$. Note that $\mu^i_{T-1}(\hat{\Theta}^i_{T-1} | a_{T-2}, r_{T-2}) = 0$.

By definition, the allocation rule $a^*_{T-1} : \Theta_{T-1} \to A$ maximizes the social surplus from period $T-1$ onward. Given that other agents always report truthfully, it follows that for $\mu^i_{T-1}(a_{T-2}, r^i_{T-2}, \theta^{-i}_{T-2})$-a.e. realized signal θ^i_{T-1}, it is optimal for agent i to report $r^i_{T-1} = \theta^i_{T-1}$.

Note that we actually need incentive compatibility to hold for $\mu^i_{T-1}(a_{T-2}, \theta^i_{T-2}, \theta^{-i}_{T-2})$-a.e. θ^i_{T-1}. If $\mu^i_{T-1}(\hat{\Theta}^i_{T-1} | a_{T-2}, \theta_{T-2}) = 0$, then the payment rule is also ex post incentive compatible under the true transitions. On the other hand, if $\mu^i_{T-1}(\hat{\Theta}^i_{T-1} | a_{T-2}, \theta_{T-2}) > 0$, then we have to consider the following possibility. Since in period $T-1$, there may exist a state $\theta^i_{T-1} \in \hat{\Theta}^i_{T-1}$ that occurs with positive probability under the true transition $\mu^i_{T-1}(a_{T-2}, \theta^i_{T-2}, \theta^{-i}_{T-2})$ so that agent i wants to manipulate her report. Therefore, agent i may in turn want to misreport in period $T-2$. To deal with this type of incentive of lying, we modify the payments by adding punishments in period $T-1$ for types in $\hat{\Theta}^i_{T-1}$ as follows.

Since each function u^i is bounded by some $M < \infty$,[10] we can define a function $\hat{p}^i_{T-1} : \Theta^i_{T-1} \to \mathbb{R}$ as

$$(7) \qquad \hat{p}^i_{T-1}(\theta^i_{T-1}) = \begin{cases} \frac{MN}{1-\delta} & \text{if } \theta^i_{T-1} \in \hat{\Theta}^i_{T-1} \\ 0 & \text{otherwise.} \end{cases}$$

That is, all types in $\hat{\Theta}^i_{T-1}$ receive a large punishment under $\hat{p}^i_{T-1}$. This implies that there is also no benefit for agent i from manipulating mechanism designer's perceived transition probabilities μ^i_{T-1}. Note that on the equilibrium path, $\hat{\Theta}^i_{T-1}$ has measure zero, so the punishment does not affect the truth-telling incentive for almost every state. Therefore, given $\hat{p}^i_{T-1}$ and $\bar{p}^i_T$, it is always optimal for agent i to report truthfully in period $T-1$. Also note that for $\mu_{T-1}(a_{T-2}, \theta_{T-2})$-a.e. state profile θ_{T-1}, agent i's continuation payoff V^i_{T-1} in the truth-telling equilibrium is

$$V^i_{T-1}(\theta_{T-1}) = W_{T-1}(\theta_{T-1}).$$

Now for any $t < T$, suppose that there exist payment schedules $\hat{p}^i_t$ and $\{\bar{p}^i_{s+1}\}^{T-1}_{s=t}$ for every agent i such that truth-telling consists of a periodic ex post equilibrium at any period $s = t, \ldots, T$ and each agent i's continuation payoff in the truth-telling equilibrium is $V^i_t(\theta_t) = W_t(\theta_t)$ for $\mu_{t-1}(a_{t-2}, \theta_{t-2})$-a.e. θ_t.

We would like to construct a transfer payment $\bar{p}^i_t : \Theta^{-i}_{t-1} \times A_{t-1} \times \Theta_t \to \mathbb{R}$ for each agent i as

$$\bar{p}^i_t(\theta^{-i}_{t-1}, a_{t-1}, \theta_t) = \hat{p}^i_t(\theta_t) - \tilde{p}^i_t(\theta^{-i}_t, a_{t-1}, \theta^{-i}_{t-1}),$$

where $\tilde{p}^i_t$ satisfies

$$\sum_{j \neq i} u^j(a_{t-1}, \theta_{t-1}) = \delta \int_{\Theta_t} \tilde{p}^i_t(\theta^{-i}_t, a_{t-1}, \theta^{-i}_{t-1}) d\mu_t(a_{t-1}, \theta_{t-1}),$$

[10] Together with finite time horizon, or discounting in the infinite horizon case.

for every a_{t-1}, θ_{t-1}^{-i} and $\mu_{t-1}^i(a_{t-2},\theta_{t-2})$-a.e. θ_{t-1}^i. The existence of such a bounded and measurable function $\tilde{p}_t^i$ again follows from Lemma 7.2. Since $\bar{p}_t^i$ is independent of θ_t^i, incentive constraints for truth-telling in periods $s=t,\dots,T$ still hold.

For $\mu_{t-1}^i(a_{t-2},\theta_{t-2})$-a.e. realized state profile θ_{t-1}, suppose agent i reports r_{t-1}^i, then her expected continuation payoff from $t-1$ on is

$$\sum_{i=1}^N u^i(a_{t-1}^*(r_{t-1}^i,\theta_{t-1}^{-i}),\theta_{t-1}) + \delta\mathbb{E}\left[W_t(\theta_t)|a_{t-1}^*(r_{t-1}^i,\theta_{t-1}^{-i}),\theta_{t-1}\right].$$

Denote the exceptional set by $\hat{\Theta}_{t-1}^i$ and $\mu_{t-1}^i(\hat{\Theta}_{t-1}^i;a_{t-2},\theta_{t-2})=0$ and define $\hat{p}_{t-1}^i:\Theta_{t-1}^i\to\mathbb{R}$ by

$$\hat{p}_{t-1}^i(\theta_{t-1}^i)=\begin{cases}\frac{MN}{1-\delta} & \text{if } \theta_{t-1}^i\in\hat{\Theta}_{t-1}^i\\ 0 & \text{otherwise.}\end{cases}\tag{8}$$

By the definition of a_{t-1}^* and similar arguments along those following equation (7), for each agent i, any report $r_{t-1}^i\in\Theta_{t-1}^i$ in period $t-1$ other than θ_{t-1}^i is suboptimal under $\hat{p}_{t-1}$ and $\{\bar{p}_s\}_{s=t}^T$.

Finally, note that in period $t-1$, agent i's continuation payoff in the truth-telling equilibrium is

$$V_{t-1}^i(\theta_{t-1})=W_{t-1}(\theta_{t-1}),$$

for $\mu_{t-1}(a_{t-2},\theta_{t-2})$-a.e. state profile θ_{t-1}.

Inducting on t backwards, we have a sequence of payments $\{\bar{p}_t\}_{t=1}^T$, where $\bar{p}_1^i=\hat{p}_1^i$ for each i. By the one-shot deviation principle, truth-telling consists of a periodic ex post equilibrium under the efficient dynamic mechanism $\{a_t^*,\bar{p}_t\}_{t=1}^T$.

□

7.3. *Proof of Theorem 4.4*

Using the one-shot deviation principle, the proof is by backward induction on t. For each t, the argument follows the same lines as the proof of Proposition 3 in Bergemann and Välimäki (2002)[11] with the payment rule defined in (3).

[11] Pages 1029–1030.

REFERENCES

ALIPRANTIS C. and K. BORDER (2006). *Infinite dimensional analysis: A hitchhiker's guide, 3rd edition*, Springer, Berlin.

ATHEY S. and I. SEGAL (2012). An efficient dynamic mechanism. *Working paper.*

BARON, D. and D. BESANKO (1984). Regulation and information in a continuing relationship. *Information Economics and Policy* **1**(3), 267–302.

BATTAGLINI, M. (2005). Long-term contracting with Markovian consumers. *American Economic Review* **95**(3), 637–658.

BATTAGLINI, M and R. LAMBA (2012). Optimal dynamic contracting. *Working paper.*

BERGEMANN, D. and M. SAID (2011). Dynamic auctions: A survey. In J.J. Cochran, with L.A. Cox, P. Keskinocak, J.P. Kharoufeh, and J.C. Smith, eds., *Wiley Encyclopedia of Operations Research and Management Science*, Volume 2, Wiley, 2011, 1511-1522.

BERGEMANN, D. and J. VÄLIMÄKI (2002). Information acquisition and efficient mechanism design. *Econometrica* **70**(3), 1007–1033.

BERGEMANN, D. and J. VÄLIMÄKI (2010). The dynamic pivot mechanism. *Econometrica* **78**(2), 771–789.

BOARD, S. (2007). Selling options. *Journal of Economic Theory* **136**(1), 324–340.

BOARD, S. and A. SKRZYPACZ (2010). Revenue management with forward looking buyers. *Working paper.*

BOLESLAVSKY, R. and M. SAID (2013). Progressive screening: Long-term contracting with a privately known stochastic process. *The Review of Economic Studies* **80**(1), 1–34.

CAVALLO, R., D. PARKES and S. SINGH (2009). Efficient mechanisms with dynamic populations and dynamic types. *Working paper.*

CHEN, Y. and S. XIONG (2012). Genericity and robustness of full surplus extraction. *Econometrica.* forthcoming.

COURTY, P. and H. LI (2000). Sequential screening. *The Review of Economic Studies* **67**, 697–717.

CRÉMER, J. and R. MCLEAN (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* **53**(2), 345–361.

CRÉMER, J. and R. MCLEAN (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* **56**(6), 1247–1257.

DASGUPTA, P. and E. MASKIN (2000). Efficient auctions. *The Quarterly Journal of Economics* **115**(2), 341–388.

DEB, R. and M. SAID (2012). Dynamic screening with limited commitment. *Working paper.*

DILME, F. and F. LI (2012). Revenue management without commitment: Dynamic pricing and periodic fire sales. *Working paper.*

ELY, J., D. GARRETT and T. HINNOSAAR (2012). Overbooking. *Working paper.*

ESŐ, P. and B. SZENTES (2007). Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies* **74**(3), 705–731.

FIESELER, K., T. KITTSTEINER and B. MOLDOVANU (2003). Partnerships, lemons, and efficient trade. *Journal of Economic Theory* **113**(2), 223–234.

GARRETT, D. (2012). Incoming demand with private uncertainty. *Working paper*.

GERSHKOV, A. and B. MOLDOVANU (2009a). Learning about the future and dynamic efficiency. *American Economic Review* **99**(4), 1576–1588.

GERSHKOV, A. and B. MOLDOVANU (2009b). Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics* **1**(2), 168–198.

GERSHKOV, A. and B. MOLDOVANU (2010). Efficient sequential assignment with incomplete information. *Games and Economic Behavior* **68**(1), 144–154.

HEIFETZ, A. and Z. NEEMAN (2006). On the (im)possibility of full surplus extraction in mechanism design. *Econometrica* **74**(1), 213–233.

JEHIEL, P. and B. MOLDOVANU (2001). Efficient design with interdependent valuations. *Econometrica* **69**(5), 1237–1259.

JEHIEL, P., M. MEYER-TER-VEHN, B. MOLDOVANU and W. ZAME (2006). The limits of ex post implementation. *Econometrica* **74**(3), 585–610.

KRÄHMER, D. and R. STRAUSZ (2012). The benefits of sequential screening. *Working paper*.

KRISHNA, V. and M. PERRY (2000). Efficient mechanism design. *Working paper*.

MCAFEE, P. and P. RENY (1992). Correlated information and mechanism design. *Econometrica* **60**(2), 395–421.

MERTENS, J.-F. (2003). A measurable "measurable choice" theorem. *Stochastic Games and Applications*, NATO Science Series **570**, 107–130.

MEZZETTI, C. (2004). Mechanism design with interdependent valuations: Efficieny. *Econometrica* **72**(5), 1617–1626.

MEZZETTI, C. (2007). Mechanism design with interdependent valuations: Surplus extraction. *Economic Theory* **31**(3), 473–488.

MIERENDORFF, K. (2011a). Optimal dynamic mechanism design with deadlines. *Working paper*.

MIERENDORFF, K. (2011b). The dynamic Vickrey auction. *Working paper*.

MILGROM, P. and R. WEBER (1982). A theory of auctions and competitive bidding. *Econometrica* **50**(5), 1089–1122.

MYERSON, R. (1986). Multistage games with communication. *Econometrica* **54**(2), 323–358.

MYERSON, R. and M. SATTERTHWAITE (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory* **29**(2), 256–281.

PAI, M. and R. VOHRA (2012). Optimal dynamic actions and simple index rules. *Mathematics of Operations Research*. forthcoming.

PARKES, D. and S. SINGH (2003). An MDP-based approach to online mechanism design. *Proc. 17th Annual Conf. on Neural Information Processing Systems*, 2003.

PAVAN, A., I. SEGAL and J. TOIKKA (2012). Dynamic mechanism design. *Working paper*.

PERRY, M. and P. RENY (2002). An efficient auction. *Econometrica* **70**(3), 1199–1212.

PERRY, M. and P. RENY (2005). An efficient multi-unit ascending auction. *The Review of Economic*

Studies **72**(2), 567–592.

Rahman, D. (2012). Surplus extraction on arbitrary type spaces. *Theoretical Economics*. forthcoming.

Said, M. (2012). Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory* **147**(6), 2419–2438.

Segal, I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review* **93**(3), 509–529.

Skrzypacz, A. and J. Toikka (2012). Mechanisms for repeated bargaining. *Working paper*.

Wilson, R. (1987). Game theoretic analysis of trading processes. In *Advances in Economic Theory*, ed. by T. Bewley, Cambridge University Press.