

DYNAMIC MECHANISM DESIGN WITH INTERDEPENDENT VALUATIONS

Swaprava Nath¹, Onno Zoeter², Y. Narahari¹, and Christopher R. Dance²

¹Indian Institute of Science, Bangalore

²Xerox Research Centre Europe, Meylan, France

May 27, 2013

Abstract

We consider an infinite horizon dynamic mechanism design problem with interdependent valuations. In this setting the types of the agents are assumed to be evolving according to a first order Markov process and the types are independent across agents in every round and across rounds. However, the valuations of the agents are functions of the types of all the agents, which makes the problem fall into an interdependent value model. Designing mechanisms in this setting is non-trivial in view of an impossibility result which says that for interdependent valuations, any efficient and ex-post incentive compatible mechanism must be a constant mechanism, even if the mechanism is static. Mezzetti circumvents this problem by splitting the decisions of allocation and payment into two stages. However, Mezzetti's result is limited to a static setting and moreover in the second stage of that mechanism, agents are weakly indifferent about reporting their valuations truthfully. This paper provides a first attempt at designing a dynamic mechanism which is *strict* ex-post incentive compatible and efficient in interdependent value setting with Markovian type evolution. In a restricted domain, which appears often in real-world scenarios, we show that our mechanism is ex-post individually rational as well.

1 INTRODUCTION

Organizations often face the problem of executing a task for which they do not have enough resources or expertise. It may also be difficult, both logistically and economically, to acquire those resources. For example, in the area of healthcare, it has been observed that there are very few occupational health professionals and doctors and nurses in all specialities at the

hospitals in the UK (Nicholson, 2004). With the advances in computing and communication technologies, a natural solution to this problem is to outsource the tasks to experts outside the organization. Hiring experts beyond an organization was already in practice. However, with the advent of the Internet, this practice has extended even beyond the international boundaries, e.g., some U.S. hospitals are outsourcing the tasks of reading and analyzing scan reports to companies in Bangalore, India (Associated-Press, 2004). Gupta et al. (2008) give a detailed description of how the healthcare industry uses the outsourcing tool.

The organizations where the tasks are outsourced (let us call them vendors) have quite varied efficiency levels. For tasks like healthcare, it is extremely important to hire the right set of experts. If the efficiency levels of the vendors and the medical task difficulties of the hospitals are observable by a central management (controller), and these levels vary over time according to a Markov process, the problem of selecting the right set of experts reduces to a Markov Decision Problem (MDP), which has been well studied in the literature (Bertsekas, 1995; Puterman, 2005). Let us call the efficiency levels and task difficulties together as *types* of the tasks and resources.

However, the types are usually observed privately by the vendors and hospitals (agents), who are rational and intelligent. The efficiencies of the vendors are private information of the vendors (depending on what sort of doctors they hire, or machinery they use), and they might misreport this information in order to win the contract and to increase their net returns. At the same time the difficulty of the medical task is private to the hospital, and is unknown to the experts. A strategic hospital, therefore, can misreport this information to the hired experts as well. Hence, the asymmetry of information at different agents' end transforms the problem from a *completely* or *partially* observable MDP into a *dynamic game* among the agents.

Motivated by examples of this kind, in this paper, we analyze them using a formal mechanism design framework. We consider only cases where the solution of the problem involves monetary transfer, which makes the payoffs quasi-linear. The reporting strategy of the agents and the decision problem of the controller is dynamic since we assume that the types of the tasks and resources are varying with time. In addition, the above problem has two characteristics, namely, *interdependent values*: the task execution generates values to the task owners that depend on the efficiencies of the assigned resources, and *exchange economy*: a trade environment where both buyers (task owners) and sellers (resources) are present. In this paper, the theme of modeling and analysis would be centered around the settings of task outsourcing to strategic experts. We aim to have a socially efficient mechanism, and at the same time, that would demand truthfulness and voluntary participation of the agents in this setting.

Valuations	STATIC	DYNAMIC
Independent	VCG Mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973)	Dynamic Pivot Mechanism (Bergemann and Välimäki, 2010; Athey and Segal, 2007; Cavallo et al., 2006)
Interdependent	Generalized VCG (Mezzetti, 2004)	Mechanism MATRIX (this paper)

Table 1: The different paradigms of mechanism design problems with their solutions.

1.1 PRIOR WORK

The above properties have been investigated separately in literature on *dynamic mechanism design*. Bergemann and Välimäki (2010) have proposed an efficient mechanism called the *dynamic pivot mechanism*, which is a generalization of the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) in a dynamic setting, which also serves to be truthful and efficient. Athey and Segal (2007) consider a similar setting with an aim to find an efficient mechanism that is budget balanced. Cavallo et al. (2006) develop a mechanism similar to the dynamic pivot mechanism in a setting with agents whose type evolution follows a Markov process. In a later work, Cavallo et al. (2009) consider *periodically inaccessible* agents and dynamic private information jointly. Even though these mechanisms work for an exchange economy, they have the underlying assumption of *private values*, i.e., the reward experienced by an agent is a function of the allocation and her own private observations. Mezzetti (2004, 2007), on the other hand, explored the other facet, namely, interdependent values, but in a static setting, and proposed a truthful mechanism. The mechanism proposed in these papers use a two-stage mechanism, since it is impossible to design a single-stage mechanism satisfying both truthfulness and efficiency even for a static setting (Jehiel and Moldovanu, 2001).

1.2 CONTRIBUTIONS

In this paper, we propose a dynamic mechanism named **MDP-based Allocation and TRansfer in Interdependent-valued eXchange economies** (abbreviated **MATRIX**), which is designed to address the class of *interdependent values*. It extends the results of Mezzetti (2004) to a dynamic setting, and serves as an efficient, truthful mechanism. Under a certain realistic domain restriction, agents receive non-negative payoffs by participating in it. The key feature that distinguishes our model and results from that of the existing dynamic mechanism literature is that we address the interdependent values and dynamically varying types (in an exchange economy) jointly. In Table 1, we have summarized the different paradigms of the mechanism design problem, and their corresponding solutions in the literature.

Our main contributions in this paper can be summarized as follows.

- We propose a dynamic mechanism **MATRIX**, that is *efficient*, and *truthful* for the participants in an *interdependent-valued exchange economy* (Theorem 1).
 - ▶ This extends the classic mechanism proposed by Mezzetti (2004) to a dynamic setting.
 - ▶ It solves the issue of weak indifference by the agents in the second stage of the classic mechanism.
- Under a restricted domain, we show that this mechanism is also *individually rational* for the agents (Theorem 2).
- We discuss why known mechanisms like a fixed payment mechanism, VCG, or the dynamic pivot mechanism (Bergemann and Välimäki, 2010) do not satisfy all the properties that **MATRIX** satisfies (beginning of Section 3 and Section 3.2).

We discuss that **MATRIX** comes at a computational cost which is the same as that of its independent value counterpart (Section 3.4). This paper provides a first attempt of designing a dynamic mechanism which is *strict* ex-post incentive compatible and efficient in interdependent value setting with Markovian type evolution. In a restricted domain, which appears often in real-world scenarios, we show that our mechanism is ex-post individually rational as well.

The rest of the paper is organized as follows. We introduce the formal model in Section 2, and present the main results in Section 3. In Section 4, we discuss what kind of results would be interesting to prove in order to characterize the space of interdependent valued dynamic mechanisms. We conclude the paper in Section 5 with some potential future works.

2 BACKGROUND AND MODEL

Let the set of agents be given by $N = \{1, \dots, n\}$, who interact with each other for a countably infinite time horizon indexed by time steps $t = 0, 1, 2, \dots$. The time-dependent type of each agent is denoted by $\theta_{i,t} \in \Theta_i$ for $i \in N$. We will use the shorthands $\theta_t \equiv (\theta_{1,t}, \dots, \theta_{n,t}) \equiv (\theta_{i,t}, \theta_{-i,t})$, where $\theta_{-i,t}$ denotes the type vector of all agents excluding agent i . We will refer to θ_t as the type profile, $\theta_t \in \Theta \equiv \times_{i \in N} \Theta_i$.

STATIONARY MARKOV TYPE TRANSITIONS, SMTT The combined type θ_t follows a first order Markov process which is governed by the transition probability function $F(\theta_{t+1}|a_t, \theta_t)$, which is independent across agents, defined formally below.

DEFINITION 1 (Stationary Markov Type Transitions, SMTT) We call the type transitions to follow stationary Markov type transitions if the joint distribution F of the types of the agents $\theta_t \equiv (\theta_{1,t}, \dots, \theta_{n,t})$, and the marginals F_i 's exhibit the following for all t .

$$F(\theta_{t+1}|a_t, \theta_t, \theta_{t-1}, \dots, \theta_0) = F(\theta_{t+1}|a_t, \theta_t), \text{ and}$$

$$F(\theta_{t+1}|a_t, \theta_t) = \prod_{i \in N} F_i(\theta_{i,t+1}|a_t, \theta_{i,t}). \quad (1)$$

We will assume the types to follow SMTT throughout this paper.

The allocation set is denoted by A . In each round t , the mechanism designer chooses an allocation a_t from this set and decides a payment $p_{i,t}$ to agent i . The allocation leads to a valuation to agent i , $v_i : A \times \Theta \rightarrow \mathbb{R}$. This is in contrast to the classical independent valuations (also called *private values*) case where valuations are assumed to depend only on i 's own type; $v_i : A \times \Theta_i \rightarrow \mathbb{R}$. However, we assume for all i , $|v_i(a, \theta)| < \infty$, for all a and θ .

In the later part of the paper, we will restrict our attention to a restricted space of allocations and valuations as discussed below.

SUBSET ALLOCATION, SA Let us motivate this restriction with the medical task assignment example given in the previous section. The organizations outsource tasks to experts for a payment, where the expert may have different and often time-varying capabilities of executing the task. The task owners come with a specific task difficulty (type of the task owner), which is usually privately known to them, while the workers' capabilities (types of the workers) are their private information. A central planner's job in this setting is to efficiently assign the tasks to the workers. Clearly, in this setting, the set of possible allocations is the set of the subsets of agents, i.e., $A = 2^N$. Note that, for a finite set of players, the allocation set is always finite. So, we can formally define this setting as follows.

DEFINITION 2 (Subset Allocation, SA) When the set of allocations is the set of all subsets of the agent set, i.e., $A = 2^N$, we call the domain a subset allocation domain.

PEER INFLUENCED VALUATIONS, PIV Even though the valuation of agent i is affected by not only her private type but also by the types of others, it is often the case that the valuation is affected by the types (e.g. the efficiencies of the workers in a joint project) of only the *selected* agents. The valuation therefore is a function of the types of the allocated agents and not the whole type vector, $v_i : A \times \Theta_A \rightarrow \mathbb{R}$. We also assume that the value of a non-selected agent is zero. Formally, we define this setting as *peer influenced valuations* (PIV).

DEFINITION 3 (Peer Influenced Valuations, PIV) This is a special case of interdependent valuations in the SA domain, where the valuation of agent i is a function of the types

of other selected agents, $v_i : A \times \Theta_A \rightarrow \mathbb{R}$. In particular, the value function is given by,

$$v_i(a, \theta) = \begin{cases} v_i(a, \theta_a) & \text{if } i \in a \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

EFFICIENT ALLOCATION, EFF The mechanism designer aims to maximize the sum of the valuations of task owners and workers, summed over an infinite horizon, geometrically discounted with factor $\delta \in (0, 1)$. The discount factor accounts for the fact that a future payoff is less valued by an agent than a current stage payoff. We assume δ to be common knowledge. If the designer would have perfect information about the θ_t 's, his objective would be to find a *policy* π_t , which is a sequence of allocation functions from time t , that yields the following for all t and for all type profiles θ_t ,

$$\pi_t \in \operatorname{argmax}_{\gamma} \mathbb{E}_{\gamma, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i \in N} v_i(a_s(\theta_s), \theta_s) \right], \quad (3)$$

where $\gamma = (a_t(\cdot), a_{t+1}(\cdot), \dots)$ is any arbitrary sequence of allocation functions. Here we use $\mathbb{E}_{\gamma, \theta_t}[\cdot] = \mathbb{E}[\cdot | \theta_t; \gamma]$ for brevity of notation. We point to the fact that the allocation policy γ is not a random variable in this expectation computation. The policy is a functional that specifies what action to take in each time instant for a given type profile. Different policies will lead to different sequences of allocation functions over the infinite horizon, and the efficient allocation is the one that maximizes the expected discounted sum of the valuations of all the agents.

In general, the allocation policy π_t depends on the time instant t . However, for the special kind of stochastic behavior of the type vectors, namely SMTT, and due to the infinite horizon discounted utility, this policy becomes stationary, i.e., independent of t . We will denote such a stationary policy by $\pi = (a(\cdot), a(\cdot), \dots)$. Thus, the efficient allocation under SMTT reduces to solving for the optimal action in the following stationary Markov Decision Problem (MDP).

$$\begin{aligned} W(\theta_t) &= \max_{\pi} \mathbb{E}_{\pi, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N} v_j(a(\theta_s), \theta_s) \right] \\ &= \max_{a \in A} \mathbb{E}_{a, \theta_t} \left[\sum_{j \in N} v_j(a, \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a, \theta_t} W(\theta_{t+1}) \right]. \end{aligned} \quad (4)$$

Here, with a slight abuse of notation, we have used a to denote the actual action taken in t rather than the allocation function. The second equality comes from a standard recursive argument for stationary infinite horizon MDPs. We refer an interested reader to standard text (Puterman, 2005, e.g.) for this reduction. Above we have used the following shorthand, $\mathbb{E}_{\theta_{t+1} | a, \theta_t}[\cdot] = \sum_{\theta_{t+1}} p(\theta_{t+1} | \theta_t; a)[\cdot]$. We will refer to W as the social welfare. The efficient allocation under SMTT is defined as follows.

DEFINITION 4 (Efficient Allocation, EFF) An allocation policy $a(\cdot)$ is **efficient** under SMTT if for all type profiles θ_t ,

$$a(\theta_t) \in \operatorname{argmax}_{a \in A} \mathbb{E}_{a, \theta_t} \left[\sum_{j \in N} v_j(a, \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a, \theta_t} W(\theta_{t+1}) \right]. \quad (5)$$

2.1 CHALLENGES IN MECHANISM DESIGN WITH INTERDEPENDENT VALUATIONS

At time t

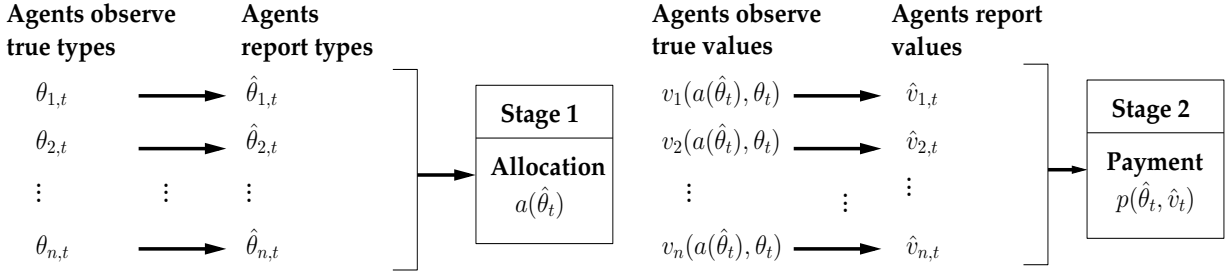


Figure 1: Graphical illustration of a candidate dynamic mechanism in an interdependent value setting.

The value interdependency among the agents poses a challenge for designing mechanisms. Even in a static setting, if the allocation and payment are decided simultaneously under the interdependent valuation setting, efficiency and incentive compatibility together can only be satisfied by a constant mechanism (Jehiel and Moldovanu, 2001). This strong negative result compels us to split the decisions of allocation and payment in two separate stages. We would mimic the two-stage mechanism of Mezzetti (2004) for each time instant of the dynamic setting (see Figure 1). The designer decides the allocation $a(\hat{\theta}_t)$ after the agents report their types $\hat{\theta}_t$ in first stage. After allocation, the agents observe their valuations $v_i(a(\hat{\theta}_t), \theta_t)$'s, and report $\hat{v}_{i,t}$'s to the designer. The payment decision is made after this second round of reporting. Our definition of incentive compatibility is accordingly modified for a two stage mechanism.

Due to SMTT and the infinite horizon of the MDP, we will focus only on stationary mechanisms, that give a stationary allocation and payment to the agents in each round of the dynamic game. Let us denote a typical two-stage dynamic mechanism by $M = \langle a, p \rangle$. The function $a : \Theta \rightarrow A$ yields an allocation for a reported type profile $\hat{\theta}_t$ in round t . Depending on the reported types in the first stage, the mechanism designer decides the allocation $a(\hat{\theta}_t)$, and due to which agent i experiences a valuation of $v_i(a(\hat{\theta}_t), \theta_t)$ in round t . Let us suppose that in the second stage, the reported value vector is given by \hat{v}_t . The payment function p

is a vector where $p_i(\hat{\theta}_t, \hat{v}_t)$ is the payment received by agent i at instant t . Combining the value and payment in each round we can write the expected discounted utility of agent i in the *quasi-linear* setting, denoted by $u_i^M(\hat{\theta}_t, \hat{v}_t|\theta_t)$, when the true type vector is θ_t and the reported type and value vectors are $\hat{\theta}_t$ and \hat{v}_t respectively. This utility has two parts: (a) the current round utility, and (b) expectation over the future round utilities. The expectation over the future rounds is taken on the true types. Thus the effect of manipulation is limited only to the current round in this utility expression. This is enough to consider due to the *single deviation principle* of [Blackwell \(1965\)](#).

$$\begin{aligned}
& u_i^M(\hat{\theta}_t, \hat{v}_t|\theta_t) \\
&= \underbrace{v_i(a(\hat{\theta}_t), \theta_t) + p_i(\hat{\theta}_t, \hat{v}_t)}_{\text{current round utility}} + \underbrace{\mathbb{E}_{\pi, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a(\theta_s), \theta_s) + p_i(\theta_s, v_s)) \right]}_{\text{expected discounted future utility}} \quad (6)
\end{aligned}$$

Here π denotes the stationary policy of actions, $(a(\cdot), a(\cdot), \dots)$. For the SMTT, the type evolution is dependent on only the current type profile and action. To avoid confusion, we will use π , $a(\hat{\theta}_t)$, or $a(\theta_s)$, $s \geq t+1$, according to the context.

Equipped with this notation, we can now define *incentive compatibility*.

DEFINITION 5 (w.p. EPIC) *A mechanism $M = \langle a, p \rangle$ is within period Ex-post Incentive Compatible (w.p. EPIC) if for all agents $i \in N$, for all possible true types θ_t , for all reported types $\hat{\theta}_{i,t}$, for all reported values $\hat{v}_{i,t}$, and for all t ,*

$$\begin{aligned}
& u_i^M(\theta_t, (v_i(a(\theta_t), \theta_t), v_{-i}(a(\theta_t), \theta_t))|\theta_t) \\
& \geq u_i^M((\hat{\theta}_{i,t}, \theta_{-i,t}), (\hat{v}_{i,t}, v_{-i}(a(\hat{\theta}_{i,t}, \theta_{-i,t}), \theta_t))|\theta_t)
\end{aligned}$$

That is, reporting the types and valuations in the two stages truthfully is an ex-post Nash equilibrium. In this context, *individual rationality* is defined as follows.

DEFINITION 6 (w.p. EPIR) *A mechanism $M = \langle a, p \rangle$ is within period Ex-post Individually Rational (w.p. EPIR) if for all agents $i \in N$, for all possible true types θ_t and for all t ,*

$$u_i^M(\theta_t, (v_i(a(\theta_t), \theta_t), v_{-i}(a(\theta_t), \theta_t))|\theta_t) \geq 0.$$

That is, reporting the types and valuations in the two stages truthfully yields non-negative expected utility.

3 THE MATRIX MECHANISM

In the setting mentioned above, our goal is to design a mechanism which is efficient (Def. 4), w.p. EPIC (Def. 5), and w.p. EPIR (Def. 6). Before we present the mechanism, let us discuss why it is non-trivial to design such a mechanism in this setting. We start with discussing a couple of naïve attempts to decide the allocation and the payment.

Fixed payment mechanism: A candidate mechanism that is often applied in organizations is a fixed payment mechanism. The allocation is done using the performance history of the agents. That is, select the agent(s) who has(have) been proved to be the most capable of doing the task in the past. But one can immediately notice that this fixed payment mechanism would not be efficient since the capabilities (types) of the agents vary over time. To make the correct decision on the allocation, it is important to know the realized type of the agent. Since this is private, the history will give only a (possibly incorrect) estimate of the type. Hence, this is not efficient.

Repeated static VCG mechanism: We know that the VCG mechanism is truthful in dominant strategies for static settings. The mechanism works on an efficient allocation for a single stage game, and pays each agent her marginal contribution. One can think of applying the static VCG in each stage of the dynamic setting. However, now, social welfare is no longer the sum of the values at the current stage, rather is the expected discounted sum of the values over the horizon of the dynamic game. Hence the allocation given by the static VCG mechanism would not be efficient in a dynamic setting.

Discussion The above two candidate mechanisms are designed to be truthful in each round, however, they fail to be efficient. It suggests that, to achieve efficient allocation in a dynamic setting, one needs to consider the expected future evolution of the types of the agents, which would reflect in the allocation and payment decisions. The value interdependency among the agents plays a crucial role here. The reason the above approach does not work is not an accident. We have already mentioned that even in a static interdependent value setting, if the allocation and payment are decided simultaneously, one cannot guarantee efficiency and incentive compatibility together (Jehiel and Moldovanu, 2001). One way out is to split the decision of allocation and payment in two stages (Mezzetti, 2004).

Following this observation, we propose MDP-based Allocation and TRansfer in Interdependent-valued eXchange economies (MATRIX), which we prove to satisfy EFF and w.p. EPIC for general interdependent valuations, and w.p. EPIR under the restricted setting of SA and PIV.

Given the above dynamics of the game as illustrated in Figure 1, the agents report their types in the first stage, and then the allocation is decided. In the second stage, they report their experienced values and the payment is decided. The task of the mechanism designer, therefore, is to design the allocation and payment rules $\langle a, p \rangle$.

We have already defined the social welfare given by Eq. (4). Let us also define the maximum social welfare excluding agent i to be $W_{-i}(\theta_{-i,t})$, which is the same as Eq. (4) except now the sum of the valuations and the allocations are over all agents $j \neq i$.

$$W_{-i}(\theta_{-i,t}) = \max_{a_{-i} \in A_{-i}} \mathbb{E}_{a_{-i}, \theta_{-i,t}} \left[\sum_{j \in N \setminus \{i\}} v_j(a_{-i}, \theta_{-i,t}) + \delta \mathbb{E}_{\theta_{-i,t+1} | a_{-i}, \theta_{-i,t}} W_{-i}(\theta_{-i,t+1}) \right] \quad (7)$$

Notice that, when i is absent, the following two notations are equivalent: $\mathbb{E}_{\theta_{t+1} | a_{-i,t}, \theta_t}[\cdot] = \mathbb{E}_{\theta_{-i,t+1} | a_{-i,t}, \theta_{-i,t}}[\cdot]$, since the type of i will be unchanged when she is not in the game. However, we adopt the former for the ease of notation.

Using the definitions above and in the previous section, now we formally present MATRIX.

MECHANISM 1 (MATRIX) *Given the reported type profile $\hat{\theta}_t$ in stage 1, choose the agents $a^*(\hat{\theta}_t)$ as follows.*

$$a^*(\hat{\theta}_t) \in \operatorname{argmax}_a \mathbb{E}_{a, \hat{\theta}_t} \left[\sum_{j \in N} v_j(a, \hat{\theta}_t) + \delta \mathbb{E}_{\theta_{t+1} | a, \hat{\theta}_t} W(\theta_{t+1}), \right] \quad (8)$$

and transfer to agent i after agents report \hat{v}_t in stage 2, a payment of,

$$p_i^*(\hat{\theta}_t, \hat{v}_t) = \left(\sum_{j \neq i} \hat{v}_{j,t} \right) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{-i,t+1}) - W_{-i}(\hat{\theta}_{-i,t}) - \left(\hat{v}_{i,t} - v_i(a^*(\hat{\theta}_t), \hat{\theta}_t) \right)^2. \quad (9)$$

The last quadratic term in the above equation is agent i 's penalty of not being consistent with the first stage report. The intuition of charging a penalty is to make sure that agent i be consistent with her reported type $\hat{\theta}_{i,t}$ in the first stage and her value report $\hat{v}_{i,t}$ in the second stage, given that others are reporting their types and values truthfully. We will argue that when all agents other than agent i reports their types and values truthfully in

Algorithm 1 MATRIX

for all time instants t **do**

Stage 1:

for agents $i = 0, 1, \dots, n$ **do**

 agent i observes $\theta_{i,t}$;

 agent i reports $\hat{\theta}_{i,t}$;

end for

 compute allocation $a^*(\hat{\theta}_t)$ according to Eq. 8;

Stage 2:

for agents $i = 0, 1, \dots, n$ **do**

 agent i observes $v_i(a^*(\hat{\theta}_t), \theta_t)$;

 agent i reports $\hat{v}_{i,t}$;

end for

 compute payment to agent i , $p_i^*(\hat{\theta}_t, \hat{v}_t)$, Eq. 9;

 types evolve $\theta_t \rightarrow \theta_{t+1}$ according to SMTT;

end for

those stages, it is the best response for agent i to do so as well. This term distinguishes our mechanism from that given by Mezzetti (2004), where the agents are weakly indifferent between reporting true and false values in the second round.

It is worth mentioning that we have used this quadratic term for the ease of exposition. However, it is easy to show that any non-negative function $g(x, \ell)$ having the property that $g(x, \ell) = 0 \Leftrightarrow x = \ell$ would still satisfy the claims made in this paper.

3.1 EFFICIENCY AND INCENTIVE COMPATIBILITY

We summarize the dynamics of MATRIX using an algorithmic flowchart in Algorithm 1. The following theorem shows that MATRIX satisfies two desirable properties in the unrestricted setting.

THEOREM 1 *Under SMTT, MATRIX is EFF and w.p. EPIC. In addition, the second stage of MATRIX is strictly EPIC.*

Note that the above theorem does not put any restriction on the allocation space and the valuation functions. MATRIX is a two stage mechanism, and we need to ensure that truth-telling is a best response in both these stages when other agents also do the same. Also, the strict EPIC in the second stage of this mechanism improves upon the mechanism given by Mezzetti (2004). Let us prove the above theorem.

Proof: Clearly, given true reported types, the allocation of MATRIX is efficient by Definition 4. Hence, we need to show only that MATRIX is w.p. EPIC.

To show that MATRIX is w.p. EPIC, let us assume that the true type profile at time t is θ_t , and all agents $j \neq i$ report their true types and values in each round $s = t, t + 1, \dots$ etc. Only agent i reports $\hat{\theta}_{i,t}$ and $\hat{v}_{i,t}$ in the two stages. Therefore, $\hat{\theta}_t = (\hat{\theta}_{i,t}, \theta_{-i,t})$ and $\hat{v}_{j,t} = v_j(a^*(\hat{\theta}_t), \theta_t)$, for all $j \neq i$. Using the *single deviation principle* (Blackwell, 1965), we conclude that it is enough to consider only a single shot deviation from the true report of the type. Hence, without loss of generality, let us assume that agent i deviates only in round t of this game.

Let us write down the discounted utility to agent i at time t .

$$\begin{aligned}
& u_i^{\text{MATRIX}}((\hat{\theta}_{i,t}, \theta_{-i,t}), (\hat{v}_{i,t}, v_{-i}(a(\hat{\theta}_{i,t}, \theta_{-i,t}), \theta_t)) | \theta_t) \\
&= \underbrace{v_i(a^*(\hat{\theta}_t), \theta_t) + p_i^*(\hat{\theta}_t, \hat{v}_t)}_{\text{current round utility}} + \underbrace{\mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right]}_{\text{expected discounted future utility}} \\
&= v_i(a^*(\hat{\theta}_t), \theta_t) + \sum_{j \neq i} \hat{v}_{j,t} + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{-i,t+1}) - W_{-i}(\hat{\theta}_{-i,t}) \\
&\quad - \left(\hat{v}_{i,t} - v_i(a^*(\hat{\theta}_t), \hat{\theta}_t) \right)^2 + \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right]
\end{aligned}$$

We use the shorthand π^* to denote the allocation policy under **MATRIX**. This gives rise to the allocations $a(\cdot)$ in each round given the type profiles (either reported or true). The first equality is from Eq. (6). The second equality comes by substituting the expression of payment from Eq. (9).

Now, from the previous discussion on the $\hat{v}_{j,t}$'s and $\hat{\theta}_{j,t}$'s, $j \neq i$, we get,

$$\begin{aligned}
& u_i^{\text{MATRIX}}((\hat{\theta}_{i,t}, \theta_{-i,t}), (\hat{v}_{i,t}, v_{-i}(a(\hat{\theta}_{i,t}, \theta_{-i,t}), \theta_t)) | \theta_t) \\
&= v_i(a^*(\hat{\theta}_t), \theta_t) + \sum_{j \neq i} v_j(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{-i,t+1}) - W_{-i}(\theta_{-i,t}) \\
&\quad - \left(\hat{v}_{i,t} - v_i(a^*(\hat{\theta}_t), \hat{\theta}_t) \right)^2 + \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right] \\
&\leq v_i(a^*(\hat{\theta}_t), \theta_t) + \sum_{j \neq i} v_j(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{-i,t+1}) - W_{-i}(\theta_{-i,t}) \\
&\quad + \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right] \tag{10}
\end{aligned}$$

The equality comes because of the assumption that all agents $j \neq i$ report their types and values truthfully. The inequality is because we are ignoring a non-positive term. Now, let us consider the last term of the above equation.

$$\begin{aligned}
& \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right] \\
&= \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) \right. \\
&\quad \left. + \sum_{j \neq i} v_j(a^*(\theta_s), \theta_s) + \delta \mathbb{E}_{\theta_{s+1} | a^*(\theta_s), \theta_s} W_{-i}(\theta_{-i, s+1}) - W_{-i}(\theta_{-i, s})) \right] \\
&= \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \left(\sum_{j \in N} v_j(a^*(\theta_s), \theta_s) + \delta \mathbb{E}_{\theta_{s+1} | a^*(\theta_s), \theta_s} W_{-i}(\theta_{-i, s+1}) \right. \right. \\
&\quad \left. \left. - W_{-i}(\theta_{-i, s}) \right) \right]
\end{aligned}$$

The first equality comes from Eq. (9). We can now rearrange the expectation for the first term above using the Markov property of θ_t that gives, $\mathbb{E}_{\pi^*, \theta_t}[\cdot] = \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t}[\mathbb{E}_{\pi^*, \theta_{t+1}}[\cdot]]$. Therefore,

$$\begin{aligned}
& \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (v_i(a^*(\theta_s), \theta_s) + p_i^*(\theta_s, v_s)) \right] \\
&= \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t} \left[\mathbb{E}_{\pi^*, \theta_{t+1}} \left(\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N} v_j(a^*(\theta_s), \theta_s) \right) \right] \\
&\quad + \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (\delta \mathbb{E}_{\theta_{s+1} | a^*(\theta_s), \theta_s} W_{-i}(\theta_{-i, s+1}) - W_{-i}(\theta_{-i, s})) \right] \\
&= \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t} (\delta W(\theta_{t+1})) \\
&\quad + \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (\delta \mathbb{E}_{\theta_{s+1} | a^*(\theta_s), \theta_s} W_{-i}(\theta_{-i, s+1}) - W_{-i}(\theta_{-i, s})) \right] \tag{11}
\end{aligned}$$

The last equality comes from the definition of $W(\theta_{t+1})$. Let us now focus on the last term of the above equation.

$$\begin{aligned}
& \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} (\delta \mathbb{E}_{\theta_{s+1}|a^*(\theta_s), \theta_s} W_{-i}(\theta_{-i, s+1}) - W_{-i}(\theta_{-i, s})) \right] \\
&= \frac{\delta^2 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2})}{\delta^2 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2})} - \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{-i, t+1}) \\
&\quad + \frac{\delta^3 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+3})}{\delta^3 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+3})} - \frac{\delta^2 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2})}{\delta^2 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2})} \\
&\quad + \dots - \frac{\delta^3 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+3})}{\delta^3 \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+3})} \\
&\quad + \dots - \dots \\
&= -\delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{-i, t+1}) \tag{12}
\end{aligned}$$

Let us show the reduction from the first term on the LHS to the first term on the RHS above. The reduction of the other terms comes from similar exercises which is straightforward and not shown here.

$$\begin{aligned}
& \mathbb{E}_{\pi^*, \theta_t} \mathbb{E}_{\theta_{t+2}|a^*(\theta_{t+1}), \theta_{t+1}} W_{-i}(\theta_{-i, t+2}) \\
&= \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [\mathbb{E}_{\theta_{t+2}|a^*(\theta_{t+1}), \theta_{t+1}} W_{-i}(\theta_{-i, t+2})] \\
&= \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [\mathbb{E}_{\theta_{t+2}|a^*(\theta_{t+1}), \theta_{t+1}, \theta_t} W_{-i}(\theta_{-i, t+2})] \\
&= \mathbb{E}_{\theta_{t+2}|\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2}) \\
&= \mathbb{E}_{\pi^*, \theta_t} W_{-i}(\theta_{-i, t+2})
\end{aligned}$$

The first equality above comes from the fact that the function inside bracket is only a function of θ_{t+1} , and the second equality is due to the Markov property.

Hence, combining Equations 10, 11, and 12, we get,

$$\begin{aligned}
& u_i^{\text{MATRIX}}((\hat{\theta}_{i,t}, \theta_{-i,t}), (\hat{v}_{i,t}, v_{-i}(a(\hat{\theta}_{i,t}, \theta_{-i,t}), \theta_t)) | \theta_t) \\
&\leq v_i(a^*(\hat{\theta}_t), \theta_t) + \sum_{j \neq i} v_j(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{-i, t+1}) \\
&\quad - W_{-i}(\theta_{-i, t}) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - W_{-i}(\theta_{-i, t+1})] \tag{13}
\end{aligned}$$

We also note that,

$$\mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{-i, t+1}) = \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{-i, t+1}) \tag{14}$$

This is because when i is removed from the system (while computing $W_{-i}(\theta_{-i, t+1})$), the values of none of the other agents will depend on the type $\theta_{i, t+1}$. And due to the independence of type transitions, i 's reported type $\hat{\theta}_{i, t}$ can only influence $\theta_{i, t+1}$. Hence, the reported value of agent i at t , i.e., $\hat{\theta}_{i, t}$ cannot affect $W_{-i}(\theta_{-i, t+1})$.

Hence, Equation 13 can be rewritten and we can show the following inequality.

$$\begin{aligned}
& u_i^{\text{MATRIX}}((\hat{\theta}_{i,t}, \theta_{-i,t}), (\hat{v}_{i,t}, v_{-i}(a(\hat{\theta}_{i,t}, \theta_{-i,t}), \theta_t)) | \theta_t) \\
& \leq v_i(a^*(\hat{\theta}_t), \theta_t) + \sum_{j \neq i} v_j(a^*(\hat{\theta}_t), \theta_t) + \cancel{\delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{-i,t+1})} \\
& \quad - W_{-i}(\theta_{-i,t}) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - \cancel{W_{-i}(\theta_{-i,t+1})}] \text{ (from Eq. 14)} \\
& = \sum_{j \in N} v_j(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\hat{\theta}_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_{-i,t}) \\
& \leq \sum_{j \in N} v_j(a^*(\theta_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\theta_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_{-i,t}) \\
& \quad \text{(by definition of } a^*(\theta_t), \text{ Eq. 8)} \\
& = u_i^{\text{MATRIX}}(\theta_t, (v_i(a(\theta_t), \theta_t), v_{-i}(a(\theta_t), \theta_t)) | \theta_t). \tag{15}
\end{aligned}$$

This shows that utility of agent i is maximized when $\hat{\theta}_{i,t} = \theta_{i,t}$ and $\hat{v}_{i,t} = v_i(a^*(\theta_t), \theta_t)$. This proves that MATRIX is within period ex-post incentive compatible.

We now argue that the second stage is strictly EPIC for an agent i . This happens because of the quadratic penalty term $(\hat{v}_{i,t} - v_i(a^*(\hat{\theta}_t), \hat{\theta}_t))^2$ in the payment p_i^* (Eq. (9)). Notice that if all the agents except i report the types and values truthfully, and agent i also reports her type truthfully in the first stage, then the penalty term will always penalize her if $\hat{v}_{i,t}$ is different from $v_i(a^*(\theta_t), \theta_t)$, which is her true valuation. Hence, the best response of agent i would be to report the true values in the second stage, which makes MATRIX strictly EPIC in this stage. \blacksquare

3.2 WHY A DYNAMIC PIVOT MECHANISM WOULD NOT WORK IN THIS SETTING

It is interesting to note that, if we tried to use the dynamic pivot mechanism (DPM), (Bergemann and Välimäki, 2010), unmodified in this setting, the true type profile θ_t in the first summation of Eq. (13) would have been replaced by $\hat{\theta}_t$, since this comes from the payment term (Eq. (9)). The proof for the DPM relies on the *private value* assumption (see the beginning of Section 2 for a definition) such that, when reasoning about the valuations for the other agents $j \neq i$, we have $v_j(a^*((\hat{\theta}_{i,t}, \theta_{-i,t})), (\hat{\theta}_{i,t}, \theta_{-i,t})) = v_j(a^*(\hat{\theta}_t), \theta_{j,t})$, with which the EPIC claim of DPM can be shown. But in the interdependent value setting, we cannot do such a substitution, and hence the proof of EPIC in DPM does not work. We have to invoke the second stage of value reporting in order to satisfy the EPIC.

3.3 EX-POST INDIVIDUAL RATIONALITY FOR A RESTRICTED DOMAIN

In this section, we consider subset allocation (SA) and the values to be peer influenced (PIV). Note that now the valuation of agent i is given by $v_i(a, \theta_a)$, and the maximum social welfare would be given by,

$$\begin{aligned} W(\theta_t) &= \max_{\pi} \mathbb{E}_{\pi, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N} v_j(a(\theta_s), \theta_{a(\theta_s)}) \right] \\ &= \max_{a \in A} \mathbb{E}_{a, \theta_t} \left[\sum_{j \in N} v_j(a, \theta_a) + \delta \mathbb{E}_{\theta_{t+1} | a, \theta_t} W(\theta_{t+1}) \right] \end{aligned} \quad (16)$$

Similarly the maximum social welfare excluding agent i is given by,

$$\begin{aligned} W_{-i}(\theta_{-i, t}) &= \max_{a_{-i}} \mathbb{E}_{a_{-i}, \theta_t} \left[\sum_{j \in N \setminus \{i\}} v_j(a_{-i}, \theta_{a_{-i}}) + \delta \mathbb{E}_{\theta_{t+1} | a_{-i}, \theta_t} W_{-i}(\theta_{-i, t+1}) \right] \end{aligned} \quad (17)$$

Both these definitions are the same as Definitions 4 and 7, but redefined in this restricted domain. We now state the following theorem on individual rationality.

THEOREM 2 (Individual Rationality) *When the allocations are chosen from class SA, values are in PIV, and types evolve in SMTT, MATRIX is w.p. EPIR.*

Proof: We observe that the allocation set is the set of subsets of N , the player set. Therefore, the set of allocations excluding agent i , denoted by $A_{-i} = 2^{N \setminus \{i\}}$, is already contained in the set of allocations including i , denoted by $A = 2^N$. Formally, this means $a_{-i} \in A_{-i} \subseteq A \ni a$. Therefore, the policies $\pi_{-i} \in A_{-i}^{\infty} \subseteq A^{\infty} \ni \pi$. From Equation 4, we can write the optimal social welfare in terms of the optimal policy π^* as follows.

$$\begin{aligned} W(\theta_t) &= \sum_{j \in N} v_j(a^*(\theta_t), \theta_{a^*(\theta_t)}) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\theta_t), \theta_t} W(\theta_{t+1}) \\ &= \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N} v_j(a^*(\theta_s), \theta_{a^*(\theta_s)}) \right] \end{aligned} \quad (18)$$

Hence in the ex-post Nash equilibrium, the utility of agent i is given by,

$$\begin{aligned} u_i^{\text{MATRIX}}(\theta_t, (v_i(a(\theta_t), \theta_t), v_{-i}(a(\theta_t), \theta_t)) | \theta_t) &= \sum_{j \in N} v_j(a^*(\theta_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1} | a^*(\theta_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_{-i, t}) \\ &= W(\theta_t) - W_{-i}(\theta_{-i, t}). \end{aligned}$$

The first equality comes from the last equality in Equation 15. The last expression in the equation above can be written as,

$$\begin{aligned}
& W(\theta_t) - W_{-i}(\theta_{-i,t}) \\
&= \mathbb{E}_{\pi^*, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N} v_j(a^*(\theta_s), \theta_{a^*(\theta_s)}) \right] \\
&\quad - \mathbb{E}_{\pi_{-i}^*, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N \setminus \{i\}} v_j(a_{-i}^*(\theta_{-i,s}), \theta_{a_{-i}^*(\theta_{-i,s})}) \right] \\
&\geq \mathbb{E}_{\pi_{-i}^*, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N \setminus \{i\}} v_j(a_{-i}^*(\theta_{-i,s}), \theta_{a_{-i}^*(\theta_{-i,s})}) \right] \\
&\quad - \mathbb{E}_{\pi_{-i}^*, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{j \in N \setminus \{i\}} v_j(a_{-i}^*(\theta_{-i,s}), \theta_{a_{-i}^*(\theta_{-i,s})}) \right] \\
&= 0
\end{aligned} \tag{19}$$

The inequality holds since while choosing the optimal policy including agent i , i.e., π^* , one has the option of choosing π_{-i}^* as well, as we are in the SA domain, and the fact that the valuations of the unallocated agents are zero, a consequence of the PIV domain. If this inequality was not true, then there would exist some $\pi_{-i}^* \in A_{-i}^\infty$ which would have achieved a social welfare more than the maximum, which is a contradiction. This proves that MATRIX is within period ex-post individually rational. ■

3.4 COMPLEXITY OF COMPUTING THE ALLOCATION AND PAYMENT

The non-strategic version of the resource to task assignment problem was that of solving an MDP, whose complexity was polynomial in the size of state-space (Ye, 2005). Interestingly, for the proposed mechanism, the allocation and payment decisions are also solutions of MDPs (Equations 8, 9). Hence the proposed mechanism MATRIX has polynomial time complexity in the number of agents and size of the state-space, which is the same as that of the dynamic pivot mechanism (Bergemann and Välimäki, 2010).

In order to get a feel for the theory behind the mechanism MATRIX, let us illustrate the mechanism through an example in the following section.

4 DISCUSSION OF THE CHARACTERIZATION OF DYNAMIC MECHANISMS

Let us consolidate our findings in this section. We have discussed the interesting and challenging domain of mechanism design with dynamically varying types and interdependent valuations. There is very little work where dynamic types and interdependent values have been addressed together. Hence, there is very little known on the limits of achievable properties in this domain. We have provided one mechanism, namely **MATRIX**, that is w.p. EPIC, strict in the second stage, and under a restricted domain, even w.p. EPIR. However, we do not know what mechanism characterizes those properties in this domain. For example, a question that may arise is “Is this the only dynamic mechanism that satisfies strict w.p. EPIC in an interdependent value setting?”. For the static setting with independent values we have the Green-Laffont characterization result that answers this question for efficiency and DSIC. However, such a characterization result is absent for interdependent valuations for both static and dynamic mechanisms. Developing such a full characterization would be worthwhile.

5 CONCLUSIONS AND FUTURE WORK

This paper provides a first attempt of designing a dynamic mechanism that is *strict* ex-post incentive compatible and efficient in an interdependent value setting with Markovian type evolution. In a restricted domain, which appears often in real-world scenarios, we show that our mechanism is ex-post individually rational as well. This mechanism, **MATRIX**, extends the mechanism proposed by Mezzetti (2004) to a dynamic setting and connects it to the mechanism proposed by Bergemann and Välimäki (2010).

However, we do not know if this is the maximal set of properties that can be satisfied by any mechanism in dynamic setting. The full set of mechanisms that satisfy the properties studied in this paper is also not characterized. Both these questions form interesting directions for future research.

Acknowledgments

The authors like to thank Ruggiero Cavallo and David C. Parkes for many useful comments on the paper. The first author is supported by Tata Consultancy Services (TCS) Doctoral Fellowship. This work is part of a collaborative project between Xerox Research and Indian Institute of Science on incentive compatible learning.

REFERENCES

Associated-Press. Some U.S. hospitals outsourcing work, June 2004. URL <http://www.msnbc.msn.com/id/6621014>.

- Susan Athey and Ilya Segal. An efficient dynamic mechanism. Working Paper, Stanford University, Earlier version circulated in 2004, 2007.
- Dirk Bergemann and Juuso Välimäki. The Dynamic Pivot Mechanism. *Econometrica*, 78: 771–789, 2010.
- Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume I. Athena Scientific, 1995.
- David Blackwell. Discounted dynamic programming. *The Annals of Mathematical Statistics*, pages 226–235, 1965.
- Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI 2006)*, pages 55–62, 2006.
- Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Efficient Mechanisms with Dynamic Populations and Dynamic Types. Technical report, Harvard University, 2009.
- Edward Clarke. Multipart Pricing of Public Goods. *Public Choice*, (8):19–33, 1971.
- Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–31, July 1973.
- Amar Gupta, Raj K Goyal, Keith A Joiner, and Sanjay Saini. Outsourcing in the healthcare industry: Information technology, intellectual property, and allied aspects. *Information Resources Management Journal (IRMJ)*, 21(1):1–26, 2008.
- Philippe Jehiel and Benny Moldovanu. Efficient Design with Interdependent Valuations. *Econometrica*, (69):1237–1259, 2001.
- Claudio Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, 72(5):1617–1626, 2004.
- Claudio Mezzetti. Mechanism Design with Interdependent Valuations: Surplus Extraction. *Economic Theory*, 31:473–488, 2007.
- P. J. Nicholson. Occupational health services in the UK - challenges and opportunities. *Occupational Medicine*, 54:147–152, 2004. doi: 10.1093/occmed/kqg125.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Interscience, 2005.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- Yinyu Ye. A New Complexity Result on Solving the Markov Decision Problem. *Math. Oper. Res.*, 30:733–749, August 2005.