Dynamic Learning-based Mechanism Design for Dependent Valued Exchange Economies

Swaprava Nath
Supervised by: Y. Narahari
Department of Computer Science and Automation
Indian Institute of Science, Bangalore 560 012, India
swaprava@gmail.com

ABSTRACT

Learning private information from multiple strategic agents poses challenge in many Internet applications. Sponsored search auctions, crowdsourcing, Amazon's mechanical turk, various online review forums are examples where we are interested in learning true values of the advertisers or true opinion of the reviewers. The common thread in these decision problems is that the optimal outcome depends on the *private* information of all the agents, while the decision of the outcome can be chosen only through reported information which may be manipulated by the strategic agents. The other important trait of these applications is their dynamic nature. The advertisers in an online auction or the users of mechanical turk arrive and depart, and when present, interact with the system repeatedly, giving the opportunity to *learn* their types. Dynamic mechanisms, which learn from the past interactions and make present decisions depending on the expected future evolution of the game, has been shown to improve performance over repeated versions of static mechanisms. In this paper, we will survey the past and current state-of-the-art dynamic mechanisms and analyze a new setting where the agents consist of buyers and sellers, known as *exchange economies*, and agents having value interdependency, which are relevant in applications illustrated through examples. We show that known results of dynamic mechanisms with independent value settings cannot guarantee certain desirable properties in this new significantly different setting. In the future work, we propose to analyze similar settings with dynamic types and population.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Economics, Game Theory, Mechanism Design

Keywords

Incentive Compatibility, Individual Rationality, Nash Equilibrium

1. INTRODUCTION

Dynamic mechanism design evolved primarily through auction theory. With the advent of the Internet, the sponsored search auction became one of the major resources of electronic commerce. Due to the dynamic nature of the interaction, the classical results, e.g., Vickrey [22], Myerson [13], dealing with static auction were needed to be extended to the settings where the agent-population or the types or both evolve over time. Though the literature on dynamic auction theory provides guarantees on truthfulness, voluntary participation and no deficit, it is restricted to the setting where agents are either all buyers or all sellers, with values independent across agents (this is also referred to as *private value* setting, [4]). In a typical auction, the auctioneer is not considered as a player, which is not suitable to model the crowdsourcing websites or dynamic project outsourcing problems, as the central organization is also critical for the strategic interaction. This framework, consisting of a population of both buyers and sellers, is known as *Exchange Economies* in the literature. We also note that the value of a crowdsourcing center would depend upon the quality of information provided by the crowd, inducing value dependency. In this paper, we will address the above interesting setting of exchange economy with dependent value structure, which has not been studied much in the literature on dynamic mechanism design. To illustrate, we discuss a simplified example of this general setting.

Motivating example: Consider a company, having received a contract for a project, wishes to select a set of employees to execute it. The company does not deterministically know the efficiencies of the employees. Company's revenue (value) increases for a faster completion of the project, and so, its goal is to allocate the project to the set of most efficient employees and make appropriate payments. A typical dynamics of this game is shown in Figure 1, where player 0 denotes the company and the players 1 through n are the employees. Let $N = \{0, 1, \dots, n\}$. In each round t, agent i observes her true types $\theta_{i,t}, i \in N$, analogous to the efficiencies for the employees and the project workload for the company. Since the true efficiency levels are private information of the employees, reported types ($\hat{\theta}_{i,t}$'s) may be different from the true ones. The decision problem for a mechanism designer is therefore to design a mechanism endowed with efficient allocation and payment rules to ensure truthfulness, voluntary participation etc. It is clear that the cost of the project depends on the completion time, which in turn depends on the efficiencies (types) of all the employees, and hence the costs to the employees (values) depend on the types of other agents, making the values dependent. The other point to note here is that the payment is made by

WWW 2011, March 28–April 1, 2011, Hyderabad, India.
ACM 978-1-4503-0637-9/11/03.

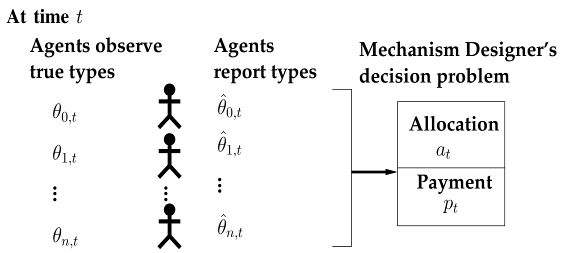

Figure 1: An illustration of the motivating example

the company to the employees, giving rise to an exchange economy.

The rest of the paper is organized as follows. In Section 2, we look into the recent literature on dynamic mechanisms. In Section 3, we present the challenges of designing mechanism in an exchange economy with dependent values and present some initial results which satisfy certain desirable properties under this setting. We conclude the paper in Section 4 with a roadmap of proposed future work.

Due to space constraints, we cannot provide formal definitions of *value*, *type*, *efficiency*, *quasi-linear payoff*, *ex-post incentive compatibility*, *within period ex-post individual rationality*, which can be found in literature on mechanism design ([15], [14], [4]) and will be clear in the context.

2. PREVIOUS WORK

In this section, we will take a look into the existing literature on *dynamic mechanism design*. The literature on static mechanism design is rich. However, they do not immediately extend to dynamic settings, since decisions in the dynamic setting often depend on the past behavior of agents which is absent in a static setting. The literature on dynamic mechanism design can be classified according to the nature of the dynamics. When the population of the agents is varying with time, e.g., they have a certain arrival and departure epochs in the system, the class is known as *dynamic population* mechanisms. While if the types of the agents are varying with time, the class is known as the *dynamic type* mechanisms. Figure 2 illustrates the different classes of dynamic mechanisms with a flowchart.

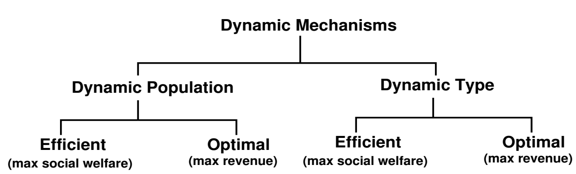

Figure 2: Different paradigms of dynamic mechanisms

2.1 Dynamic Population Mechanisms

Under this class of mechanisms, **efficient** mechanisms are those which maximize the total value (social welfare) of the agents. The sequential allocation problem in dynamic population setting has been studied by Parkes and Singh [18]. Each agent has an arrival time and a departure time, and has a utility function for decisions made while she is present, but the mechanism designer does not know any of this information. Their model considers *patient* buyers who are willing to maximize their long term payoff, and proposes *online VCG* mechanism as a solution. A sharp contrast to this problem with *impatient* buyers have been studied in Gershkov and Moldovanu [9], where the buyers wish to purchase an object immediately upon their arrival on the market.

Said [20] discusses the allocation of a sequence of indivisible goods to a dynamic population of buyers. The author presents an efficient dynamic *indirect* mechanism where patient buyers demand a single unit of a perishable homogeneous good.

The other class called **optimal** mechanisms maximize the revenue of an auction, following the taxonomy from Myerson [13]. The revenue-maximizing counterpart to the efficient online VCG mechanism [18] is discussed by Pai and Vohra [16] with patient buyers. The setting with impatient buyers has been analyzed by Gershkov and Moldovanu [8].

Said [20] examines the assignment of a sequence of perishable goods to a population of patient buyers, and develops and indirect mechanism.

2.2 Dynamic Type Mechanisms

The sub-classification into **efficient** and **optimal** mechanisms apply here as well. Let us discuss **efficient** mechanisms first. The infinite horizon dynamic type model has been analyzed in Bergemann and Välimäki [3]. The proposed efficient mechanism is called *dynamic pivot mechanism*, which is a generalization of the Clarke's pivotal mechanism in dynamic setting. Cavallo *et al.* [5] develop a mechanism similar to the dynamic pivot mechanism in a setting with agents whose type evolution follows a Markov process. In Cavallo *et al.* [6], the authors extend dynamic VCG to settings in which buyers are *periodically inaccessible* and are unable to make reports. Zoeter [23] applied dynamic VCG in sponsored search to make the auction cheating-proof. Sarma *et al.* [21] characterized mechanisms in multi-slot sponsored search auctions.

Athey and Segal [1] consider a similar setting to that of [3]. However, they are also interested in finding an efficient mechanism that is **budget balanced**. They essentially generalize the AGV mechanism [7] to dynamic setting.

Kakade *et al.* [11] characterize the **optimal** dynamic mechanism for Multi-armed Bandit problems, finding its applications in sponsored search auctions.

Pavan *et al.* [19] generalize the results of Myerson [13] to a general dynamic setting. In particular, they characterize incentive compatibility and revenue in multi-period settings with dynamic private information.

Literature summary: The literature on dynamic mechanism design addresses private value environment, and a large share of them address auction setting. We would like to emphasize that this paper asks questions about the dependent value setting in an exchange economy, which distinguishes itself from the above literature.

3. DEPENDENT-VALUED EXCHANGE ECONOMIES

In this section, we motivate the problem addressed and show the subtle differences from the existing literature and formally introduce the model, notation and assumptions.

3.1 Motivation for this work

Many social and network economics examples, discussed in Section 1, involve exchange economies with interdependent values. The dynamic pivot mechanism or dynamic VCG mechanism ([3], [6]), will not readily guarantee incentive compatibility in this setting. For interdependent valuations, if one restricts attention to single-stage *static* mechanisms, then incentive compatibility is inconsistent with making efficient decisions, even if one does not impose any budget balancing or individual rationality constraint (Jehiel and Moldovanu [10]). Since dynamic mechanism design is inherently harder than static ones, designing dynamic mechanisms under such setting is a distinct nontrivial problem.

3.2 The Model

We consider a simple exchange economy setting, with single buyer and multiple sellers trading over an infinite horizon. The agents' values at time t depend on the allocation in that round and the type profile of all the agents *present at that instant*. This has the implication that if an agent is removed from the system, she can no longer affect the valuations of other agents in that round. We consider a direct revelation mechanism with **quasi-linear payoffs**, where the buyer chooses a single (or a set of) seller(s) for procurement. Once selected, the agents *perfectly* observe their realized values. These observed values are functions of a random quantity called the *state of the world* (example: random disturbance for a project completion) and we will assume that it has a known *stationary* distribution. The observed types and values are *private information* of the agents.

3.3 Notation and Assumptions

The notation is summarized in Table 1. We summarize the assumptions as follows.

- After the agents are allocated, they can observe their realized *true values* perfectly.
- We consider a dynamic mechanism with types evolving in a Markov process with the transitions independent across agents, i.e., $F(\theta_{t+1}|a_t,\theta_t) = \prod_{i\in N} F_i(\theta_{i,t+1}|a_t,\theta_{i,t})$, $\forall\ t$, a_t is the action chosen at t.
- The value of the center at t depends only on the true types of the agents present at t.
- If an agent cannot increase her payoff by misreporting, then she sticks to truth-telling.
- The absolute value of the valuation function is *almost surely* bounded.
- The state of the world process has a stationary distribution.

3.4 Dynamics of the Game

Since interdependent valuations prevent a single stage mechanism to simultaneously guarantee efficiency and incentive compatibility ([10]), Mezzetti [12] decouples the allocation and payment decisions by proposing two-stage reporting mechanism for *static setting*. Similarly in *dynamic setting*, we assume that in each round, each agent observes her true value after allocation, which depends on the realization of the state of the world in addition to the true types of the agents. In the example of Section 1, the exact completion time of the project does not only depend on the efficiency of the selected employee, but also on a random disturbance (e.g., labor shortage, power cut etc.), and the cost is observed after the project is executed.

Let us illustrate the dynamics of the game with a single buyer (center) and n sellers (agents) as follows.

1. Agent i observes $\theta_{i,t}$, at time t, $i \in N$.
2. Agents are asked to report types. They report $\hat{\theta}_{i,t}$.
3. Mechanism decides the allocation of the job, $a_t \in A$ depending on $\hat{\theta}_t$.
4. The state of the world, ω, realizes.

$N = \{0, 1, \ldots, n\}$	Set of agents, 0 denotes the buyer.
Θ_i	Type set of agent $i \in N$.
$\Theta = \Theta_1 \times \cdots \times \Theta_n$	Set of type profiles.
A	Set of all possible allocations.
$\omega \in \Omega$	State of the world.
$\Pi \in \triangle(\Omega)$	Cumulative distribution of the state of the world.
$V_i : A \times \Theta \times \Omega \to \mathbb{R}$	Value function of agent $i \in N$. The value given an allocation depends on the types of the agents present at time t and on the state of the world.
$v_i : A \times \Theta \to \mathbb{R}$	Expected value function of agent $i \in N$, expected over the state of the world.
$t = 0, 1, 2, \ldots$	Discrete time index.
$\theta_{i,t}, \hat{\theta}_{i,t}$	True type, reported type of agent $i \in N$ at time t.
$\theta_t, \hat{\theta}_t$	True type profile, reported type profile at time t.
$\hat{V}_{i,t}$	Values reported by agent $i \in N$, at time t
$\hat{V}_t$	Reported value profile at t.
$a_t : \Theta \to A$	Allocation function acting on $\hat{\theta}_t$.
$F_i \in \triangle(\Theta_i)$	Prior of $\theta_{i,0}$.
$F \in \triangle(\Theta)$	Joint prior of θ_0.
$F_i(\theta_{i,t+1}\|a_t, \theta_{i,t})$	Stochastic kernel associated with agent i.
$F(\theta_{t+1}\|a_t, \theta_t)$	Stochastic kernel of the joint transition. We assume that the joint kernel is the product of the marginals.
δ	Common discount factor (infinite horizon).
$p_i^* : \Theta \times \mathbb{R} \to \mathbb{R}$	Monetary transfer to agent i at each stage given the reported types and values.
$W : \Theta \to \mathbb{R}$	Total social welfare function given a type profile θ_t.
$W_{-i} : \Theta_{-i} \to \mathbb{R}$	Social welfare function excluding agent i given a type profile $\theta_{-i,t}$.

Table 1: Notation.

5. Agents observe their true valuation $V_i(a_t, \theta_t, \omega)$, $i \in N$, which is a realization of a random variable.
6. All agents are asked to report their observed values. They report $\hat{V}_{i,t}$, $i \in N$.
7. Payments $p_t(\hat{\theta}_t, \hat{V}_t)$ are decided by the mechanism.

Figure 3 gives a graphical illustration of the above dynamics.

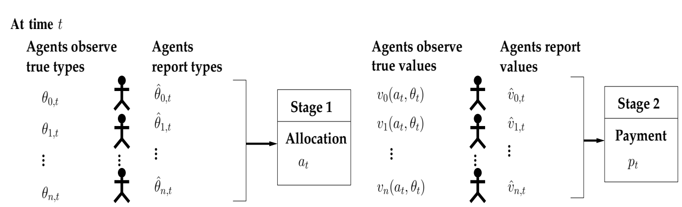

Figure 3: Graphical illustration of the proposed dynamic mechanism

3.5 The Generalized Dynamic Pivot Mechanism

Given the above dynamics of the game, the task of the mechanism designer is to design the allocation and payment rules. We note that the expected value of agent i (expected over the state of the world) is given by,

$$v_i(a_t, \theta_t) = \int_\Omega V_i(a_t, \theta_t, \omega) d\Pi(\omega) \tag{1}$$

where $V_i(a_t, \theta_t, \omega)$ is the *realized* value of the agent given *realized* state of the world ω. The objective of the social planner is to maximize social welfare given the type profile θ_t, which is defined as,

$$\begin{aligned} & W(\theta_t) \\ = \ & \max_{\pi_t} \mathbb{E}_{\pi_t, \theta_t} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i \in N} v_i(a_s, \theta_s) \right] \\ = \ & \max_{a_t} \mathbb{E}_{a_t, \theta_t} \left[\sum_{i \in N} v_i(a_t, \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a_t, \theta_t} W(\theta_{t+1}) \right] \end{aligned}$$

where $\pi_t = (a_t, a_{t+1}, \dots)$ is the sequence of actions starting from t. We assume that there exists a *stationary* policy which maximizes the social welfare, i.e., $\forall\ \theta_t \in \Theta_t,\ \exists\ a^*(\theta_t)$ such that,

$$a^*(\theta_t) \in \arg\max_{a_t} \mathbb{E}_{a_t, \theta_t} \left[\sum_{i \in N} v_i(a_t, \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a_t, \theta_t} W(\theta_{t+1}) \right] \tag{2}$$

We seek a socially efficient mechanism in this setting. In the following, we propose the generalized (two-stage) dynamic pivot mechanism (GDPM) which guarantees efficiency, incentive compatibility and individual rationality.

Mechanism 1 (GDPM). *Given the reported type profile $\hat{\theta}_t$, choose the agents $a^*(\hat{\theta}_t)$ according to Equation 2. Transfer to agent i after agents report $\hat{V}_t$,*

$$p_i^*(\hat{\theta}_t, \hat{V}_t) = \sum_{j \neq i} \hat{V}_{j,t} + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{t+1}) - W_{-i}(\hat{\theta}_t) \tag{3}$$

Proposition 1. *GDPM is efficient, within period ex-post incentive compatible, and within period ex-post individually rational.*

Proof: The allocation is efficient by choice. Let us focus on agent i. To prove ex-post incentive compatibility, we assume all agents except agent i report their true types. Hence, $\hat{\theta}_t = (\hat{\theta}_{i,t}, \theta_{-i,t})$. So, the discounted utility to agent i at t given the realized state of the world ω is,

$$\begin{aligned} & U_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t, \omega) \\ = \ & \underbrace{V_i(a^*(\hat{\theta}_t), \theta_t, \omega) + p_i^*(\hat{\theta}_t, \hat{V}_t)}_{\text{marginal flow utility}} + \\ & \underbrace{\delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - W_{-i}(\theta_{t+1})]}_{\text{discounted future marginal utility}} \\ = \ & V_i(a^*(\hat{\theta}_t), \theta_t, \omega) + \sum_{j \neq i} \hat{V}_{j,t} + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{t+1}) \\ & - W_{-i}(\hat{\theta}_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - W_{-i}(\theta_{t+1})] \end{aligned}$$

where the last equality comes from Equation 3. We notice that agent i's payoff does not depend on her value report $\hat{V}_{i,t}$. Hence, agent i has no incentive to misreport her observed valuation, and this applies to all agents. Therefore, by assumption, agents report their values truthfully, and we get,

$$\hat{V}_{i,t} = V_i(a^*(\hat{\theta}_t), \theta_t, \omega) \quad \forall\ i \in N \tag{4}$$

Hence,

$$\begin{aligned} & U_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t, \omega) \\ = \ & V_i(a^*(\hat{\theta}_t), \theta_t, \omega) + \sum_{j \neq i} V_j(a^*(\hat{\theta}_t), \theta_t, \omega) + \\ & \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{t+1}) - W_{-i}(\hat{\theta}_t) + \\ & \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - W_{-i}(\theta_{t+1})] \end{aligned} \tag{5}$$

Now, we note that,

$$\mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \hat{\theta}_t} W_{-i}(\theta_{t+1}) = \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{t+1}) \tag{6}$$

This is because when i is removed from the system (while computing $W_{-i}(\theta_{t+1})$), the values of all other agents do not depend on the type $\theta_{i,t+1}$ (this is by assumption, which is credible, since the revenue of a company cannot depend on the efficiency of an employee who is not present in the game). 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 $W_{-i}(\theta_{t+1})$. Similar arguments show that,

$$W_{-i}(\hat{\theta}_t) = W_{-i}(\theta_t) \tag{7}$$

Hence, Equation 5 reduces to,

$$\begin{aligned} & U_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t, \omega) \\ = \ & V_i(a^*(\hat{\theta}_t), \theta_t, \omega) + \sum_{j \neq i} V_j(a^*(\hat{\theta}_t), \theta_t, \omega) + \\ & \cancel{\delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W_{-i}(\theta_{t+1})} - W_{-i}(\theta_t) + \\ & \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} [W(\theta_{t+1}) - \cancel{W_{-i}(\theta_{t+1})}] \\ = \ & \sum_{i \in N} V_i(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_t) \end{aligned} \tag{8}$$

where we get the first equality combining Equations 6 and 7. The utility of agent i given by Equation 8 depends on a specific realization of the state of the world ω. It is clear that, this utility of agent i is indeed random. Hence, the correct quantity to consider would be the utility expected over the state of the world, i.e.,

$$\begin{aligned} & u_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t) \\ & = \int_\Omega U_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t, \omega) d\Pi(\omega) \end{aligned} \tag{9}$$

Hence, the expected discounted utility to agent i at t is, using Equations 8 and 9,

$$\begin{aligned} & u_i(a^*(\hat{\theta}_t), (p_i^*(\hat{\theta}_t, \hat{V}_t)), \theta_t) \\ = \ & \sum_{i \in N} \int_\Omega V_i(a^*(\hat{\theta}_t), \theta_t, \omega) d\Pi(\omega) + \\ & \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_t) \\ = \ & \sum_{i \in N} v_i(a^*(\hat{\theta}_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\hat{\theta}_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_t) \\ \leq \ & \sum_{i \in N} v_i(a^*(\theta_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\theta_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_t) \\ = \ & u_i(a^*(\theta_t), (p_i^*(\theta_t, V_t)), \theta_t) \end{aligned}$$

where $V_t = (V_i(a^*(\theta_t), \theta_t, \omega))_{i \in N}$. The second equality comes from Equation 1, and the next inequality comes by definition of $a^*(\theta_t)$, Equation 2. Hence ex-post IC. We notice that at this equilibrium, i.e., the ex-post Nash equilibrium, the utility of agent i is,

$$\begin{aligned} & u_i(a^*(\theta_t), (p_i^*(\theta_t, V_t)), \theta_t) \\ = & \sum_{i \in N} v_i(a^*(\theta_t), \theta_t) + \delta \mathbb{E}_{\theta_{t+1}|a^*(\theta_t), \theta_t} W(\theta_{t+1}) - W_{-i}(\theta_t) \\ = & W(\theta_t) - W_{-i}(\theta_t) \\ \geq & 0 \end{aligned}$$

Hence ex-post IR. ■

4. CONCLUSIONS AND FUTURE WORK

In this paper, we have motivated the need of studying dependent-valued exchange economies in dynamic setting. We have reviewed the literature on dynamic mechanism design for agents with independent value structure to show that this setting has not been addressed. In this work, we have performed a case study of this setting and proposed a mechanism, that serves to perform as truthful and individually rational. In general, the future work would be the following.

- Design of mechanisms which would additionally satisfy budget balance (i.e., the sum of the payments to the agents is non-positive, so that the mechanism does not run into deficit), *consistent payments* (i.e., buyers pay and sellers get paid in each round of this dynamic mechanism) under this setting.
- We would like to study the situation where the parameters of the decision problem is also learned over time.
- We would study revenue properties in order to design revenue optimizing mechanisms.
- We would be interested to study the dynamic population model with the dependent value structure.
- We would study trade-offs between the strong notions of truthfulness and the computational issues, some of which have been addressed in [2], [17].

We hope that some of these directions will turn out to be an exciting area of research and contribute to the community of Internet commerce.

5. ACKNOWLEDGMENTS

We would like to thank David C. Parkes for many useful discussions, and Pankaj Dayama for discussions and comments on the paper.

6. REFERENCES

[1] S. Athey and I. Segal. An efficient dynamic mechanism. Working Paper, Stanford University, Earlier version circulated in 2004, 2007.

[2] M. Babaioff, R. D. Kleinberg, and A. Slivkins. Truthful Mechanisms with Implicit Payment Computation. In *Proceedings of Electronic Commerce*, June 2010.

[3] D. Bergemann and J. Välimäki. The Dynamic Pivot Mechanism. *Econometrica*, 78:771–789, 2010.

[4] R. Cavallo. Efficiency and Redistribution in Dynamic Mechanism Design. In *Proceedings of Electronic Commerce*. ACM, July 8–12 2008.

[5] R. Cavallo, D. C. Parkes, and S. 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.

[6] R. Cavallo, D. C. Parkes, and S. Singh. Efficient Mechanisms with Dynamic Populations and Dynamic Types. Technical report, Harvard University, 2009.

[7] C. d'Aspremont and L. A. Gérard-Varet. Incentives and Incomplete Information. *Journal of Public Economics*, 11(1):25–45, 1979.

[8] A. Gershkov and B. Moldovanu. Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach. *American Economic Journal: Microeconomics*, 1(2):168–198, 2009.

[9] A. Gershkov and B. Moldovanu. Efficient Sequential Assignment with Incomplete Information. *Games and Economic Behavior*, 68(1):144–154, 2010.

[10] P. Jehiel and B. Moldovanu. Efficient Design with Interdependent Valuations. *Econometrica*, (69):1237–1259, 2001.

[11] S. Kakade, I. Lobel, and H. Nazerzadeh. An Optimal Dynamic Mechanism for Multi-Armed Bandit Processes. arXiv:1001.4598v2, 2010.

[12] C. Mezzetti. Mechanism Design with Interdependent Valuations: Efficiency. *Econometrica*, 2004.

[13] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(5873), 1981.

[14] Y. Narahari, D. Garg, R. Narayanam, and H. Prakash. *Game Theoretic Problems in Network Economics and Mechanism Design Solutions*. Springer-Verlag London Limited, 2009.

[15] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[16] M. Pai and R. Vohra. Optimal Dynamic Auctions. Unpublished Manuscript, Northwestern University.

[17] D. C. Parkes and A. D. Procaccia. Dynamic Social Choice: Foundations and Algorithms. Unpublished Manuscript, Harvard University, 2010.

[18] D. C. Parkes and S. Singh. An MDP-based Approach to Online Mechanism Design. In *Proceedings of the 17th Annual Conference on Neural Information Processing Systems*, 2003.

[19] A. Pavan, I. Segal, and J. Toikka. Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure. Unpublished manuscript, Northwestern University.

[20] M. Said. Auctions with Dynamic Populations: Efficiency and Revenue Maximization. In *Proceedings of AMMA*, 2009.

[21] A. D. Sarma, S. Gujar, and Y. Narahari. Multi-Armed Bandit Mechanisms for Multi-Slot Sponsored Search Auctions. Unpublished Manuscript, 2010, arxiv.org/pdf/1001.1414.

[22] W. Vickery. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 1961.

[23] O. Zoeter. On a form of advertiser cheating in sponsored search and a dynamic-VCG solution. In *World Wide Web*, April 2008.