

Optimal Dynamic Mechanism Design via a Virtual VCG Mechanism

SHAM M. KAKADE

Department of Statistics, Wharton School, University of Pennsylvania.

and

ILAN LOBEL

Stern School of Business, New York University.

and

HAMID NAZERZADEH

Microsoft Research, New England Lab.

We consider the problem of designing revenue maximizing mechanisms for dynamic settings. We propose a simple optimal mechanism for rather general “separable” settings. The proposed mechanism is implemented via an affine transformation of a dynamic VCG mechanism.

Categories and Subject Descriptors: K.4.4 [Computers and Society]: Electronic Commerce

General Terms: Economics, Theory, Algorithms

Additional Key Words and Phrases: Mechanism Design, Dynamics, Optimal Mechanisms, Efficient Mechanisms, VCG, Electronic Marketplaces, Internet Advertising

1. INTRODUCTION

Extracting revenue from strategic agents in *dynamic* environments — in which agents’ private information changes over time — is a central issue in many markets such as Internet advertising. While the theory of optimal mechanism design has enjoyed much success in static settings, the celebrated mechanism proposed by Myerson [1981] does not easily translate into an optimal mechanism for dynamic settings.

One compelling application is that of ad auctions where publishers sell the space on their webpages to the advertisers. In Internet advertising (e.g., sponsored search), typically, an advertiser places an ad in order to: first, draw a user to visit the advertiser’s website (via a click on the displayed ad), and then, subsequently, have the user perform a desired transaction (e.g., purchasing a product). The expected value that an advertiser obtains from a displayed ad depends on *both* the “click-through rate” (the probability that a user clicks on the ad, sending the user to the advertiser’s website) and the “conversion rate” (the probability that the user who visits the website performs a desired transaction). This is a dynamic environment in which both advertisers and the publisher (the mechanism) learn and update their estimates of these rates over time. A central question that arises in this setting is that of designing mechanisms for such *sequence of auctions* and determining how much revenue such mechanisms obtain.

Contact: skakade@wharton.upenn.edu, ilobel@stern.nyu.edu, hamidnz@microsoft.com

Analogously to static mechanism design, two natural objectives in the dynamic setting are maximizing the long term social welfare of all agents (*efficiency*) and maximizing the long term revenue of a seller (*optimality*). With regards to maximizing the future social welfare in a dynamic setting, there are elegant extensions of the efficient (VCG) mechanism that are applicable to quite general dynamic settings by [Bergemann and Välimäki 2010; Parkes and Singh 2003; Athey and Segal 2007]. With regards to *optimal* dynamic mechanisms the state of affairs is more murky. While there are detailed characterizations of necessary conditions that incentive compatible dynamic mechanisms must satisfy, there are only a few rather restricted special cases for which optimal mechanisms are characterized (see [Éso and Szentes 2007; Pavan et al. 2008; Deb 2008]).

In this work, we provide a certain structural characterization of a “separable” environment in which we characterize a simple, optimal dynamic mechanism. Our characterization is rather rich; it permits general stochastic processes and is applicable to certain natural formalizations of the aforementioned sponsored search setting (which will appear in a forthcoming paper [Kakade et al. 2011]). Our construction draws a rather close connection to efficient mechanism design, where our optimal mechanism builds upon a dynamic VCG mechanism.

2. DYNAMIC AUCTIONS AND SEPARABLE VALUES

The setting is a discrete-time, discounted infinite-horizon model that consists of one seller and n agents. At each time t , the seller decides upon an action (or an allocation) a_t at a cost of $c(a_t)$.

At every period, each agent i receives a private signal $s_{i,t}$. The signal $s_{i,t}$ that agent i receives at time t is maybe correlated to her previous signals $s_{i,1}, \dots, s_{i,t-1}$ and the past actions of the seller a_1, \dots, a_{t-1} . We assume it is independent of the past and current signals of the other agents (when conditioned on the seller’s actions a_1, \dots, a_{t-1}). At each period t , agent i has value $v_{i,t}$, which is a function of all her private information and the all of the seller’s actions.

A *direct* dynamic mechanism is one where each agent is asked to bid their private information every round – at each period, each agent reports all signals ($s_{i,1}, \dots, s_{i,t}$) she has received up to time t , not just her current signal $s_{i,t}$. The reporting strategies of the agent and the mechanism \mathcal{M} determine a stochastic process, which evolves as follows:

- (1) Each agent i receives her private signal $s_{i,t}$.
- (2) Each agent i provides a report of her private information ($s_{i,1}, \dots, s_{i,t}$) – this mechanism requires the agent to re-report at every period all her previous signals.
- (3) As a function of the previous reports, the previous actions, and the current reports, the mechanism determines the action a_t , and the payments $p_{i,t}$ for each agent i .

The (instantaneous) utility of agent i at time t is the value she obtains minus the payment she makes. The long term utility of the agent is the discounted sum of these instantaneous utilities. An *optimal mechanism* is one which both satisfies certain incentive constraints and maximizes the sellers’ expected discounted revenue

(or profit) among all mechanism which satisfies these incentive constraints. The incentive constraints that a feasible mechanism must satisfy are:

- (Incentive Compatibility) At the first period, truthful reporting of one’s type is a best response to all other agents being truthful. That is, truthfulness maximizes the expected utility of an agent when other agents are truthful.
- (Individual Rationality) All agents choose to participate, i.e., their expected discounted utility is non-negative.

We only require feasible mechanisms to satisfy the rather weak notions of both incentive compatibility and individual rationality above (they are only with respect to the future utility at $t = 1$). However, we might hope that an optimal mechanism can be implemented with stronger notions of incentive compatibility and individual rationality. In fact, we do provide much stronger guarantees with respect to our optimal mechanism, similar to those provided by Bergemann and Välimäki [2010] for efficient dynamic mechanisms. See [Kakade et al. 2011] for more details.

The setting we consider makes two structural restrictions on the environment. The first is a statistical one, where for each agent i the first signal $s_{i,1}$ is assumed to be independent of all future signals — this assumption can be made without loss of generality. The second is a functional one, where we assume that the value function is either multiplicatively or additively separable. A multiplicatively (additively) separable value function $v_{i,t}$ is given by a product (sum) of two functions: one that depends only on the first signal, another that depends on all actions and all signals besides the first one.

The sponsored search example can be described as multiplicatively separable environment if the first signal $s_{i,1}$ of an agent contains information about the marginal value to the advertiser of a transaction (a sale, for example), while the other signals $s_{i,2}, s_{i,3}, \dots$ capture the information the advertiser obtains over time about the click-through and conversion rates.

3. OPTIMAL MECHANISMS

We first describe some noteworthy properties of the optimal mechanism using a special case with a single agent.

3.1 Single Agent Case

For multiplicatively separable values, there exists an exceedingly simple mechanism that is optimal for selling a sequence of items to a single agent. Suppose the production cost of each item is equal to c — the actions of the mechanism correspond to either allocating the item to the agent at cost c or not allocating it at no cost. The mechanism selects a decreasing curve $M(\cdot)$ and asks the agent to quote a price $p > c$. The agent is then immediately charged a fee of $M(p)$. Afterwards, the mechanism offers to sell every single item to the agent at the take-it-or-leave-it price of p . That is, irrespective of changes in the agent’s valuation or the history of purchases, all items are offered at a constant price. This is perhaps surprising as one might expect that if, for example, the agent’s valuation follows a stochastic process with an upwards drift, then the mechanism would try to capture those rising values by using an increasing sequence of prices.

This simple optimal mechanism has the structure of a subscription service: an agent who expects a priori that she will have high valuations for the items on sale will pay a high membership fee, but be offered all products at a constant low price, while conversely an agent who expects a priori that she will have low valuations will pay a low membership fee, but will be faced with all items being offered at a constant high price.

3.2 Multiple Agents Case

For the case with multiple agents, the optimal mechanism is more complex than the optimal one for a single agent, but it still retains the flavor of a subscription service. Essentially, it is composed of two phases: on Phase 1, which occurs at $t = 1$, all agents bid for their “importance” in the social welfare function (they effectively bid on the affine transformation that will be used on their value, when it is placed in an affinely transformed social welfare function). On Phase 2, which occurs at all subsequent periods $t = 2, 3, \dots$, we implement a dynamic (transformed) social welfare maximizing mechanism (using the transformations that were determined in Phase 1).

The mechanism combines two main ideas. First, using the virtual value formulation of [Myerson 1981], it determines each agent’s “importance” in an affine transformation of the social welfare. Then, implements such affine transformation of social welfare using the dynamic pivot mechanism [Bergemann and Välimäki 2010]. We therefore call this optimal dynamic mechanism the Virtual Pivot Mechanism.

REFERENCES

- ATHEY, S. AND SEGAL, I. 2007. An efficient dynamic mechanism. *Working paper*.
- BERGEMANN, D. AND VÄLIMÄKI, J. 2010. The dynamic pivot mechanism. *Econometrica* 78, 771–789.
- DEB, R. 2008. Optimal contracting of new experience goods. *Working paper*.
- ËSO, P. AND SZENTES, B. 2007. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies* 74, 3, 705–731.
- KAKADE, S. M., LOBEL, I., AND NAZERZADEH, H. 2011. Optimal dynamic mechanism design and the virtual pivot mechanism. *Working paper*.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research*.
- PARKES, D. C. AND SINGH, S. P. 2003. An mdp-based approach for online mechanism design. In *Proceedings of the 17th Conference on Neural Information Processing Systems*.
- PAVAN, A., SEGAL, I., AND TOIKKA, J. 2008. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. *Working paper*.