

Dynamic Mechanism Design

Rakesh V. Vohra*

Abstract

This paper provides a brief survey of developments in the study of mechanisms for dynamic settings. It interprets the literature as identifying conditions under which simple index policies are optimal.

1 Introduction

Mechanism design is a framework for thinking about what a given institution can achieve when the information necessary to make decisions is dispersed and privately held. The set of questions that have been investigated using mechanism design is large. For example, to achieve a given reduction in carbon emissions, is it more cost effective to rely on taxes or a cap and trade system? Is it better to sell an IPO via auction or the traditional book building approach? In each case the information needed to arrive at the correct answer is dispersed and privately held. In the emissions case, the relevant information is the actual cost of reducing emissions by each emitter of carbon. Ideally, one would want those with the lowest cost of reducing emissions to reduce their emissions first so as to achieve the targeted reduction in the lowest cost way. The most popular application of mechanism design is the examination of trading rules in static settings when the valuations of the buyers are private information.

In the last decade a concerted effort has been made to extend the boundaries of the subject to dynamic settings. Two kinds of applications motivate this effort. The first is to understand how a seller facing a stream of strategic buyers that arrive over time should adjust her prices. The second is when the set of buyers remains fixed, but their private information about valuations changes over time with experience.¹ The recent survey of Bergemann and Said (2010) provides a fair, balanced and eminently readable account of these developments. This review will be neither fair or balanced. Perhaps, not even readable! I depart from a conventional review that would summarize the literature around problems or models. Instead, I choose to emphasize one idea and argue that the literature can be understood as an attempt to enlarge the domain of application of that one idea. Fans of Isaiah Berlin might justly call this a hedgehog exercise.²

*Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University, Evanston IL 60208.

¹There are possibilities that are a mixture of the two which I ignore.

²Berlin offered a classification of thinkers inspired by a fragment of poetry due to Archilochus. Rendered in English, it reads: The fox knows many things, but the hedgehog knows one big thing.

What is the one idea? Index policies. In each period the agents with the highest index have the highest priority for access to resources. In the absence of restrictions on how the index is computed, this statement has no bite. Thus, we want the index to be easily computable and dependent only on ‘local’ information. The well known Gittins index (1979) is an example of the kind of policy I have in mind. It is easy to compute and simple to implement.

I begin with an overview of the classic static mechanism design problem of allocating a single object amongst a group of agents so as to maximize expected revenue (see Myerson (1981)). I will emphasize the index policy description of the optimal auction. Subsequently, I outline how much of the structure of that static set up can be ‘lifted’ into dynamic settings. Space limitations require that I assume the reader is familiar with the terminology of mechanism design.³

2 Static Allocation of a Single Unit

Here I will consider the allocation of a single object amongst a group of agents whose value for the object is private information. In this environment the type of a buyer is their value for the good. Let T be the type space and for convenience, take it to be discrete, i.e., $T = \{1, 2, \dots, m\}$. Denote the set of all n buyer profiles of types by T^n . Given a profile $\tau \in T^n$, we use $n_i(\tau)$ to denote the number of buyers in the profile with type i . Given the revelation principle we restrict attention to direct mechanisms. A direct mechanism consists of an allocation rule \mathbf{a} and a payment rule \mathbf{p} . The allocation rule \mathbf{a} specifies for each buyer the probability that it will receive the good as a function of the reported types of all buyers. Similarly, \mathbf{p} specifies the payment that each buyer must make as a function of the reported types of all buyers. I restrict attention to allocation and payment rules that are anonymous, i.e., independent of the names of buyers. This is without loss when each buyer’s type is an independent draw from the same distribution, F . Denote by f_i the probability a buyer has type i . The anonymity assumption means that we can focus on a single buyer, say, buyer 1.

Denote by τ a profile of types and let $\mathbf{a}_i(\tau)$ be the probability that agent 1 receives the good when his type is i and the reported profile types is τ . Let \mathcal{A}_i be the **interim allocation** probability that agent 1 with type i receives the good, i.e.,

$$\mathcal{A}_i = \sum_{\tau^{-1} \in T^{n-1}} \mathbf{a}_i[i, \tau^{-1}] \pi(\tau^{-1}). \quad (1)$$

Here $\pi(\tau^{-1})$ is the probability that the profile of types of the buyers other than buyer 1 are τ^{-1} . Similarly, we can define P_i to be the expected payment of agent 1 if she reports type i . The expected surplus that agent 1 of type i enjoys when she reports type j will be is $i\mathcal{A}_j - P_j$.

The usual conditions imposed on a mechanism (\mathbf{a}, \mathbf{p}) are individual rationality (IR) and Bayes-Nash incentive compatibility (BNIC). IR is

$$i\mathcal{A}_i - P_i \geq 0 \quad \forall i \in T. \quad (2)$$

³The reader unfamiliar with these matters would do well to consult the early chapters in Krishna (2002).

BNIC requires that

$$i\mathcal{A}_i - P_i \geq i\mathcal{A}_j - P_j \quad \forall j \neq i. \quad (3)$$

One can fold the IR constraint into BNIC by introducing a dummy type ‘0’ and setting $\mathcal{A}_0 = 0$. From now we will assume the type space to be $\{0, 1, \dots, m\}$.

The first key result characterizes allocation rules that are BNIC.

Theorem 1 *An allocation rule \mathbf{a} can be implemented by a BNIC mechanism (\mathbf{a}, \mathbf{p}) iff. the corresponding interim allocation probabilities are monotonic. That is, $r \leq s \iff \mathcal{A}_r \leq \mathcal{A}_s$.*

The proof is straightforward. Observe that (3) implies that $i\mathcal{A}_i - P_i \geq i\mathcal{A}_j - P_j$ as well as $j\mathcal{A}_j - P_j \geq j\mathcal{A}_i - P_i$. Adding these two inequalities together yields

$$(i - j)[\mathcal{A}_i - \mathcal{A}_j] \geq 0.$$

Theorem 1 means that the search for a BNIC mechanism can be restricted to choosing amongst allocation rule for which the corresponding interim allocation probabilities is monotone.

The second key result bounds the payments of a BNIC mechanism. Specifically,

$$P_i \leq \sum_{r=0}^i r(\mathcal{A}_r - \mathcal{A}_{r-1}) + K = i\mathcal{A}_i - \sum_{r=0}^{i-1} \mathcal{A}_r + K. \quad (4)$$

Here K is a constant that is usually determined by the IR constraint.⁴

2.1 Social Welfare Maximization

Here, I derive the incentive compatible allocation rule that maximizes the expected value of the allocation. Call a set of interim allocation probabilities, \mathcal{A}_i ’s feasible if there exists an allocation rule \mathbf{a} such that (1) holds for all i . The next result characterizes the set of feasible interim allocation probabilities.

Theorem 2 (Border’s Theorem, Border (1991, 2007)) *The interim allocation \mathcal{A}_i is feasible iff.*

$$n \sum_{i \in S} f_i \mathcal{A}_i \leq 1 - \left(\sum_{i \notin S} f_i \right)^n \quad \forall S \subseteq \{1, 2, \dots, m\}. \quad (5)$$

Three remarks are in order. First, Border’s theorem relies on the independence assumption on the distribution of types. Second, no use is made of the fact that types are one dimensional, i.e., Border’s Theorem holds even for multidimensional types. Third, the right hand side of (5), $1 - (\sum_{i \notin S} f_i)^n$, is submodular in S . Hence, the set of inequalities defined in (5) form a polymatroid.⁵

The problem of finding monotone interim allocation probabilities that maximize expected value becomes:

$$\begin{aligned} Z_0 = & \max n \sum_{i=1}^m i f_i \mathcal{A}_i \quad (OPT E) \\ \text{s.t.} & \quad n \sum_{i \in S} f_i \mathcal{A}_i \leq 1 - \left(\sum_{i \notin S} f_i \right)^n \quad \forall S \subseteq \{1, 2, \dots, m\} \\ & \quad 0 \leq \mathcal{A}_1 \leq \dots \leq \mathcal{A}_i \leq \dots \leq \mathcal{A}_m \end{aligned}$$

⁴When the type space is continuous, the inequality in (4) becomes an equality and the summation an integral.

⁵After a change of variables, $x_i = f_i \mathcal{A}_i$. For more on polymatroids, see Schrijver (2003).

Theorem 1 ensures that the optimal solution to (*OPT E*) is incentive compatible.

Ignoring the monotonicity constraints on the \mathcal{A}_i 's in program (*OPT E*) yields a polymatroid optimization problem. Such problems can be solved in greedy fashion, by picking the variable with the largest objective function coefficient first and making it as large as feasibly possible. At optimality, $\mathcal{A}_i = 0$ whenever $i = 0$. Therefore, $\sum_{i \geq r} f_i \mathcal{A}_i = \frac{1 - F(r-1)^n}{n}$ for all r . In particular at optimality $\mathcal{A}_m = \frac{1 - F(m-1)^n}{nf_m}$, and

$$\mathcal{A}_r = \frac{1 - F(r-1)^n}{nf_r} - \frac{1 - F(r)^n}{nf_r} = \frac{F(r)^n - F(r-1)^n}{nf_r}$$

for all $0 < r \leq m-1$. Straightforward algebra reveals that the \mathcal{A} 's chosen in this way are monotone. Therefore, the optimal solution to the relaxation of (*OPT E*) obtained by omitting the monotonicity constraints is feasible for (*OPT E*). Hence, it must be optimal for problem (*OPT E*).

The solution via the inequalities in Border's Theorem fixes the value of the \mathcal{A}_i 's but does not produce the actual allocation rule \mathbf{a} . Nevertheless, one can recover the corresponding \mathbf{a} 's quite easily. Specifically, treat the value of \mathcal{A}_i as a priority index for type i . In each profile, allocate the good to the agent with the highest priority. In case of a tie, allocate with equal probability amongst all agents with the same priority (see Manelli and Vincent (2010) or Vohra (2011)). Under monotonicity of the interim allocations, we know that higher types will have higher priorities. Thus, an equivalent rule is to allocate in each profile to the buyer with highest type. Notice this rule satisfies the stronger property of being ex-post efficient.

To summarize, the efficient mechanism in this instance is an example of an index policy where the index of an agent depends only on her type.

2.2 Revenue Maximization

I turn now to the problem of finding the incentive compatible mechanism that maximizes expected revenue. Since P_i is the expected payment received from agent 1 when she claims type i and the allocation rule is anonymous, the expected payment *per agent* will be $\sum_{i=0}^m f_i P_i$. Hence, the problem of finding the revenue maximizing mechanism can be formulated as:

$$Z = \max_{\{\mathbf{a}\}} n \sum_{i=0}^m f_i P_i \quad (\text{OPT1}) \quad (6)$$

$$\text{s.t. } P_i \leq i \mathcal{A}_i - \sum_{r=0}^{i-1} \mathcal{A}_r \quad \forall i \in T \quad (7)$$

$$0 = \mathcal{A}_0 \leq \mathcal{A}_1 \leq \dots \leq \mathcal{A}_i \leq \dots \leq \mathcal{A}_m \quad (8)$$

$$P_0 = 0 \quad (9)$$

$$\mathcal{A}_i = \sum_{\tau^{-1} \in T^{n-1}} \mathbf{a}_i[i, \tau^{-1}] \pi(\tau^{-1}) \quad (10)$$

$$\sum_{i=1}^m n_i(\tau) \mathbf{a}_i(\tau) \leq 1 \quad \forall \tau \in T^n \quad (11)$$

It is easy to see that at optimality, the constraints in (7) bind.⁶ Hence, $P_i = i\mathcal{A}_i - \sum_{r=0}^{i-1} \mathcal{A}_r$ for all $i \in T$. Therefore, after some tedious algebra,

$$\sum_{i=1}^m f_i P_i = \sum_{i=1}^m f_i \left\{ i - \frac{1 - F(i)}{f_i} \right\} \mathcal{A}_i.$$

The term $\left\{ i - \frac{1 - F(i)}{f_i} \right\}$ is called the **virtual value** of type i . It is also usual to assume $\frac{1 - F(i)}{f_i}$, is non-increasing in i . This is called the **monotone hazard condition** and many natural distributions satisfy it.

Since the expression $i - \frac{1 - F(i)}{f_i}$ will be ubiquitous, set $\nu(i) = i - \frac{1 - F(i)}{f_i}$. Under the monotone hazard condition, $\nu(i)$ is non-decreasing in i . Problem (OPT1) becomes:

$$\begin{aligned} Z_1 &= \max_{\{\mathbf{a}\}} n \sum_{i=1}^m f_i \mathcal{A}_i \nu(i) \quad (OPT2) \\ \text{s.t. } & 0 \leq \mathcal{A}_1 \leq \dots \leq \mathcal{A}_i \leq \dots \leq \mathcal{A}_m \\ & \mathcal{A}_i = \sum_{\tau^{-1} \in T^{n-1}} \mathbf{a}_i[i, \tau^{-1}] \pi(\tau^{-1}) \quad \forall i \\ & \sum_{i=1}^m n_i(\tau) \mathbf{a}_i(\tau) \leq 1 \quad \forall \tau \in T^n \end{aligned}$$

Using (5) we can reformulate the problem of finding the revenue maximizing mechanism for the sale of a single object as follows:

$$\begin{aligned} Z_1 &= \max n \sum_{i=1}^m f_i \mathcal{A}_i \nu(i) \quad (OPT3) \\ \text{s.t. } & n \sum_{i \in S} f_i \mathcal{A}_i \leq 1 - \left(\sum_{i \notin S} f_i \right)^n \quad \forall S \subseteq \{1, 2, \dots, m\} \\ & 0 \leq \mathcal{A}_1 \leq \dots \leq \mathcal{A}_i \leq \dots \leq \mathcal{A}_m \end{aligned}$$

Ignoring the monotonicity constraints on the \mathcal{A}_i 's in program (OPT3), we have a polymatroid optimization problem (recall the change of variables: $x_i = f_i \mathcal{A}_i$) that can be solved in greedy fashion, by picking the variable with the largest objective function coefficient first and making it as large as feasibly possible. At optimality, $\mathcal{A}_i = 0$ whenever $\nu(i) < 0$. Let i^* be the lowest type such that $\nu(i^*) \geq 0$. Assuming the monotone hazard condition, for all $r \geq i^*$ we have $\sum_{i \geq r} f_i \mathcal{A}_i = \frac{1 - F(r-1)^n}{n}$. In particular at optimality $\mathcal{A}_m = \frac{1 - F(m-1)^n}{nf_m}$, and

$$\mathcal{A}_r = \frac{1 - F(r-1)^n}{nf_r} - \frac{1 - F(r)^n}{nf_r} = \frac{F(r)^n - F(r-1)^n}{nf_r}$$

for all $i^* \leq r \leq m-1$. Straightforward algebra reveals that the \mathcal{A} 's chosen in this way are monotone. Therefore, the optimal solution to the relaxation of (OPT3) obtained by omitting the monotonicity constraints, is feasible for (OPT3). Hence, it must be optimal for problem (OPT3). Note also that the optimal solution to (OPT3) differs from that of (OPT1) only when types have negative virtual values.

⁶To maximize expected revenue and satisfy IR, set $K = 0$ in (4).

Theorem 3 *Under the monotone hazard condition the revenue maximizing mechanism for the sale of a single object allocates the object in each profile to the buyer with highest virtual value (provided it is non-negative).*

The idea mimics the case of efficient allocation. Treat the value of \mathcal{A}_i as a priority index for type i . Under the monotone hazard condition, higher types will have higher priorities. Thus, allocate in each profile to the buyer with highest virtual value. Once again we have an index policy where the index depends on the type of an agent as well as the distribution of types.

The derivation of the revenue maximizing mechanism in the static single good setting presented relies on what is known as the reduced form. Specifically, the use of the interim allocation probabilities, \mathcal{A}_i rather than the actual allocation rule \mathbf{a} . I have chosen the reduced form because it suggests how to generalize these results to a dynamic setting.

3 Dynamic Settings

There are two ways in which the single good setting can be made ‘dynamic’. The first is to keep the set of buyers fixed and have their types change over a time as a function of allocations selected in earlier periods. Such settings can be used to model the sale of experience goods, long term procurement relationships, indeed any situation where information that arrives over time changes the buyers valuations. Examples of papers of this variety are Athey and Segal (2007), Esö and Szentes (2007) and Akan *et al.* (2009). Call this the changing type case.

Modest reflection reveal the challenges the changing type case presents. Buyers must take into account not only own payoffs in the current period, but how the allocation in the current period will influence their type and allocation in a future period. Second, what should the appropriate notion of individual rationality be? Should one require individual rationality in each period or does it suffice for individual rationality to hold in the aggregate? A similar question applies in the case of incentive compatibility. By enforcing individual rationality, say, in each period one ensures the participation of the agent for the duration in which the mechanism is in force. Individual rationality in the aggregate might give an agent the incentive to depart in mid-stream. One can prevent early departure by requiring agents to post a bond that is forfeit if they depart early. The bond, however, may be so large as to be impractical.⁷

The second way to make the single good setting dynamic, is to expose the seller to a stream of buyers that arrive over time. Call this the changing buyer case. Here buyers types do not change but the mix of buyers changes over time. The central problem for the seller is the usual one of whether to make a sale today or wait until tomorrow. The difficulties arise when buyers can time their purchases. A buyer today may not be present tomorrow. Alternatively a buyer may choose to wait for an opportune moment to make a purchase.

In both cases a ‘simple’ index policy would be attractive. By simple, I mean that in any period the index of a buyer should depend only on his own type and the distribution of types of buyers

⁷One might argue that objections to a bond rely on the assumption that agents are budget constrained. If so, critics of the bond idea should impose the budget constraint in each period.

in that period. I view the literature as an attempt to find sufficient conditions for the existence of such simple index policies. I will provide an illustrative example of such an attempt for the changing type case as well as the changing buyer case.

4 Fixed Buyers with Changing Types: Bandit Setting

The canonical *infinite horizon* dynamic system with changing types is the multi armed bandit. One has a collection of projects (arms), N , each with a finite set of associated states.⁸ The set of states associated with project i are S^i . Project i engaged in state $j \in S^i$ generates a reward or value $v^i(j)$. When a project is engaged in a given period, the relevant reward is obtained and the state of the project changes according to a known Markov process. Projects not engaged in the current period remain in their current state.⁹ At most one project can be engaged in any period. The objective is to determine a policy for engaging projects in each period so as to maximize the expected discounted rewards obtained.¹⁰

When there is a single project, the question becomes a stopping problem: At what point should one stop engaging the project? The maximum expected discounted reward to be obtained from project i when it is initially in state j is called the Gittins index of the pair (i, j) . Call it $g^i(j)$. Note that the index does not depend on other projects. It is well known, that the multi-armed Bandit problem can be solved using Gittins indices. The optimal policy is to engage, in each time period, the project with the highest Gittins index.

An important observation for what is to come is that there is a reduced form representation of the set of feasible policies. Instead of specifying which project is to be engaged each time as a function of the history to date, one chooses x_j^i , the expected discounted number of times project i is engaged when it is in state j . Call a collection of x_j^i variables feasible if there is a policy for the bandit problem that yields those x_j^i 's. Bertsimas and Nino-Mora (1996) show that the set of feasible x_j^i 's is characterized by an extended polymatroid. If we denote the relevant extended polymatroid by E , then the bandit problem can be stated as:

$$V(N) = \max\left\{\sum_{i \in N} \sum_{j \in S^i} v^i(j)x_j^i : x \in E\right\} \quad (\text{MAB})$$

Many important properties of polymatroid optimization carry over to optimization over extended polymatroids (see Bhattacharya, Georgiadis and Tsoucas (1992)). In particular, there is a greedy algorithm for maximizing a linear function over an extended polymatroid. In the case of the Bandit problem, this greedy algorithm reduces to the Gittins index policy.

It is easy to graft a mechanism design problem onto the bandit setting. Agents play the role of projects. One can imagine engaging a project or agent as allotting them an object. The value agent i assigns to consuming the object in state $j \in S^i$, is $v^i(j)$. At the end of the period, the object is consumed and the state of the agent awarded the object changes. This change corresponds to

⁸For economy of exposition I ignore the case of an infinite number of states.

⁹This assumption is particularly important for tractability.

¹⁰For an example of a finite horizon setting see Athey and Segal (2007)

experience or learning. The private information of agent i in each period is their current state or equivalently $v^i(j)$. Further motivation for such a graft can be found in Bergemann and Välimäki (2010), Pavan *et al.* (2008) as well as Kakade *et al.* (2010).

4.1 Social Welfare Maximization

We examine the easiest case first: identify an incentive compatible, efficient mechanism for the bandit setting.¹¹ This case is considered in full generality in Bergemann and Välimäki (2010) (see also Carvalho, Parkes and Singh (2006)). Here I examine a special case.

Suppose an auctioneer who in each period must allocate an indivisible, non-storable object amongst a set N of bidders. Let δ be the common discount rate of the bidders, s_t^i the state of agent i at time t and $g^i(s_t^i)$ the Gittins index of the state s_t^i . For each i and t we order the states by increasing Gittins index. The value that agent i enjoys from consuming the object in time period t in state s_t^i is $v^i(s_t^i)$. We assume a one to one correspondence between values and states. At each period t the efficient allocation rule chooses the agent i for which $g^i(s_t^i)$ is largest.

Consider the reduced form of the Bandit problem, i.e., problem (MAB). This is a static problem and the natural thing to do is run the Vickrey-Clarke-Groves (VCG) mechanism. Each agent is asked to report their vector of v^i 's. Note that since the bandit process has not 'begun', the agents are only reporting what their values in each of the possible states they may find themselves in during the course of the bandit process. Next, solve (MAB) with respect to the reported v^i 's. Let x^* be the optimal solution of (MAB). Subsequently, charge payments according to the (VCG) rule. This mechanism is both efficient and dominant strategy incentive compatible. It achieves incentive compatibility by giving each agent a surplus equal to her marginal product. The marginal product of agent i is $V(N) - V(N \setminus i)$.

Given x^* , we can back out the optimal policy. Now, execute the policy. There's the rub. To implement x^* we must rely on each agent truthfully reporting their current state at each time period not just their values in each possible state. Second, applying VCG to (MAB) tells us only what the *total* expected discounted payments should be not when they should be charged. If payments are extracted at the end, an agent might find it profitable to exit the mechanism at some intermediate stage. This would happen if its realized surplus exceeded the expected surplus promised under VCG.

Thus, the challenge is to determine per period payments needed to ensure three things:

1. Agents have an incentive to participate in each period.
2. Agents have an incentive to truthfully report their state in each period.
3. The expected discounted payments equal the expected payments under VCG.

We meet the challenge by observing that in each period we have a problem similar to the static problem of allocating a single good. Thus, an expression similar to (4) should hold.

Conditioned on the history, let $x_t^i(j)$ denote the interim probability of selecting agent i at time t if she reports the state j . Similarly, denote by $p_t^i(j)$ her expected payment. Let $V_{t+1}(S|E_t)$

¹¹The analysis described here was developed in collaboration with Sven de Vries.

denote the value of the efficient allocation from time $t + 1$ onwards involving only the agents in S conditional on event E_t occurring at time period t . Similarly $MP_{t+1}^i(E_t)$ is the marginal product of agent i from time $t + 1$ onwards conditional on E_t .

Let $\{E_t^i, r\}$ be the event that agent i wins at time t when she reports state r . Let $\{F_t^i, r\}$ be the complement of this event. Then, agent i 's payoff assuming truthful reporting is

$$\begin{aligned} & [v^i(s_t^i) + \delta MP_{t+1}^i(\{E_t^i, s_t^i\})]x_t^i(s_t^i) + [\delta MP_{t+1}^i(\{F_t^i, s_t^i\})](1 - x_t^i(s_t^i)) - p_t^i(s_t^i) \\ &= \{v^i(s_t^i) + \delta MP_{t+1}^i(\{E_t^i, s_t^i\}) - \delta MP_{t+1}^i(\{F_t^i, s_t^i\})\}x_t^i(s_t^i) - p_t^i(s_t^i) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) \end{aligned}$$

It will be useful to rewrite the expression in braces:

$$v^i(s_t^i) + \delta MP_{t+1}^i(\{E_t^i, s_t^i\}) - \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) \quad (12)$$

Observe, first that

$$V_t(N \setminus i | \{F_t^i, s_t^i\}) = V_{t+1}(N \setminus i | \{E_t^i, s_t^i\}).$$

Furthermore, if k is the agent that wins at time t if agent i does not, and w_t^k is the state of agent k , it follows that

$$V_t(N \setminus i | \{F_t^i, s_t^i\}) - v^k(w_t^k) = \delta V_{t+1}(N \setminus i | \{F_t^i, s_t^i\}).$$

Therefore expression (12) is

$$v^i(s_t^i) - v^k(w_t^k) + (1 - \delta)V_t(N \setminus i | \{F_t^i, s_t^i\}) + \delta[V_{t+1}(N | \{E_t^i, s_t^i\}) - V_{t+1}(N | \{F_t^i, s_t^i\})].$$

For economy of notation, set

$$h^i(s_t^i) = v^i(s_t^i) - v^k(w_t^k) + (1 - \delta)V_t(N \setminus i) + \delta[V_{t+1}(N | \{E_t^i, s_t^i\}) - V_{t+1}(N | \{F_t^i, s_t^i\})].$$

If agent i reports she is in state r instead, her expected payoff, *assuming* truthful reporting in all subsequent periods, will be:

$$\begin{aligned} & [v^i(s_t^i) + \delta MP_{t+1}^i(\{E_t^i, s_t^i\})]x_t^i(r) + [\delta MP_{t+1}^i(\{F_t^i, s_t^i\})](1 - x_t^i(r)) - p_t^i(r) \\ &= \{v^i(s_t^i) + \delta MP_{t+1}^i(\{E_t^i, s_t^i\}) - \delta MP_{t+1}^i(\{F_t^i, s_t^i\})\}x_t^i(r) - p_t^i(r) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) \\ &= h^i(s_t^i)x_t^i(r) - p_t^i(r) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}). \end{aligned}$$

The notion of incentive compatibility imposed is bayesian incentive compatibility and captured in the following:

$$h^i(s_t^i)x_t^i(s_t^i) - p_t^i(s_t^i) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) \geq h^i(s_t^i)x_t^i(r) - p_t^i(r) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) \quad \forall r \in S^i. \quad (13)$$

The first term on the left hand side of (13), $h^i(s_t^i)x_t^i(s_t^i) - p_t^i(s_t^i)$, is the surplus in the current period from reporting truthfully. The second term, $\delta MP_{t+1}^i(\{F_t^i, s_t^i\})$, on the left hand side, is *promised* future surplus assuming truthful reporting in the next period. The two terms on the right hand side can be similarly interpreted assuming the agent lies about her type. We want (13) to hold for all histories.

Interchanging the roles of s_t^i and r in (13) yields:

$$h^i(r)x_t^i(r) - p_t^i(r) + \delta MP_{t+1}^i(\{F_t^i, r\}) \geq h^i(r)x_t^i(s_t^i) - p_t^i(s_t^i) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}). \quad (14)$$

Adding them together yields:

$$[h^i(s_t^i) - h^i(r)][x_t^i(s_t^i) - x_t^i(r)] \geq 0. \quad (15)$$

Suppose that $h^i(r)$ is increasing in r , i.e., as $g^i(r)$ increases, $h^i(r)$ increases. Then, (15) implies that x_t^i is increasing in the state. In words, the higher the Gittins index, the higher the probability of being engaged.

To show that $h^i(r)$ is increasing in r , consider the part of $h^i(s_t^i)$ that depends on s_t^i :

$$v^i(s_t^i) + \delta[V_{t+1}(N|E_t^i(s_t^i)) - V_{t+1}(N|\{F_t^i, s_t^i\})].$$

Then

$$v^i(s_t^i) + \delta[V_{t+1}(N|E_t^i(s_t^i)) - V_{t+1}(N|\{F_t^i, s_t^i\})] = V_t(N) - [v^k(w_t^k) + \delta V_{t+1}(N|\{F_t^i, s_t^i\})] + v^k(w_t^k).$$

Now, the term $V_t(N) - [v^k(w_t^k) + \delta V_{t+1}(N|\{F_t^i, s_t^i\})]$ is the change in welfare when we switch from i to k at time t . This difference should increase the larger the Gittins index of agent i at time t . In fact it does by the interchange lemma in Gittins' original proof. Thus, $h^i(r)$ is increasing in r .

Let \hat{r} be the lowest state in S^i . Since h^i is monotone, it follows by standard arguments (similar to that used to derive (4)) that

$$p_t^i(s_t^i) = h^i(s_t^i)x_t^i(s_t^i) - \sum_{q=\hat{r}}^{h^i(s_t^i)} [h^i(q) - h^i(q+1)]x_t^i(q) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) + K.$$

Here K is a constant determined by the IR constraint.

As long as q has a Gittins index that is at least as large the Gittins index of agent k , $x_t^i(q) = 1$. Let r_t^i be the smallest state for agent i for which $g^i(r_t^i)$ is equal to the Gittins index of agent k . Then

$$p_t^i(s_t^i) = h^i(r_t^i) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) + K.$$

Observe that the efficient allocation rule will be indifferent between awarding the object to i in state r_t^i and to agent k . Thus,

$$v^i(r_t^i) + \delta V_{t+1}(N|\{E_t^i, r_t^i\}) = v^k(w_t^k) + \delta V_{t+1}(N|\{E_t^k, w_t^k\}) = v^k(w_t^k) + \delta V_{t+1}(N|\{F_t^i, r_t^i\}). \quad (16)$$

Rearranging (16) we conclude that

$$\delta V_{t+1}(N|\{E_t^i, r_t^i\}) - \delta V_{t+1}(N|\{F_t^i, r_t^i\}) = v^k(w_t^k) - v^i(r_t^i).$$

Hence,

$$h^i(r_t^i) = v^i(r_t^i) - v^k(w_t^k) + (1 - \delta)V_t(N \setminus i) + \delta[V_{t+1}(N|\{E_t^i, r_t^i\}) - V_{t+1}(N|\{F_t^i, r_t^i\})] = (1 - \delta)V_t(N \setminus i).$$

Thus,

$$p_t^i(s_t^i) = (1 - \delta)V_t(N \setminus i) + \delta MP_{t+1}^i(\{F_t^i, s_t^i\}) + K.$$

4.2 Revenue Maximization

Turn now to revenue maximization. As before, I consider the reduced form, i.e. (MAB) and check to see if we can decompose the payments across periods. Unfortunately, the reduced form of the problem is an instance of revenue maximization when private information is *multi-dimensional*. This is because an agent's private information is $\{v_j^i\}_{j \in S^i}$.

Since revenue maximization with multidimensional private information is hard even in the non-dynamic case, we will have to make an assumption to reduce the dimensionality of private information. Specifically, each agent i has a one dimensional type θ^i drawn independently from a common distribution F (with density f) over $\{1, 2, \dots, m\}$. Agent i 's value from consuming the object when in state $j \in S^i$ is $v^i(j|\theta^i)$. So as to mimic the analysis of section 4.1 suppose that

$$v^i(j|\theta^i) = \theta^i + r_j^i.$$

We can interpret this to mean that agent i initially has a value θ^i for the object. After consuming the object when she is in state j , she updates her value to $\theta^i + r_j^i$. Such a restriction is considered in Pavan *et al.* (2008) as well as Kakade *et al.* (2010).¹²

We now have a bandit system where the state of project i is of the form (θ^i, j) . When a project i is in state (θ^i, j) it cannot transition to a state (σ^i, k) for $k \in S^i$ and $\sigma^i \neq \theta^i$. A project i in state (θ^i, j) generates a reward of $\theta^i + r_j^i$. Let E denote the extended polymatroid associated with this bandit system.

Now, follow the path taken in section 4.1. Let $x_j^i(\theta^i)$ be the expected discounted number of times agent i is awarded the object when she is in state (θ^i, j) . In what follows assume for the moment that the second component of the state, r_j^i is known to the seller. Equivalently, agent i will truthfully report r_j^i when she is in state (θ^i, r_j^i) . Subsequently we relax this assumption.

Let $p^i(\theta^i)$ denote the total expected payment of agent i when she reports type θ^i . Let $X^i(\theta^i)$ denote the total discounted number of times agent i receives the object when she reports type θ^i . Observe that $X^i(\theta^i) = \sum_{j \in S^i} x_j^i(\theta^i)$. Bayesian incentive compatibility implies that

$$\theta^i X^i(\theta^i) + \sum_{j \in S^i} r_j^i x_j^i(\theta^i) - p^i(\theta^i) \geq \theta^i X^i(\sigma^i) + \sum_{j \in S^i} r_j^i x_j^i(\sigma^i) - p^i(\sigma^i) \quad \forall \sigma^i \neq \theta^i. \quad (17)$$

Interchanging the roles of θ^i and σ^i yields:

$$\sigma^i X^i(\sigma^i) + \sum_{j \in S^i} r_j^i x_j^i(\sigma^i) - p^i(\sigma^i) \geq \sigma^i X^i(\theta^i) + \sum_{j \in S^i} r_j^i x_j^i(\theta^i) - p^i(\theta^i) \quad \forall \theta^i \neq \sigma^i. \quad (18)$$

Adding (17) to (18) produces

$$(X^i(\theta^i) - X^i(\sigma^i))(\theta^i - \sigma^i) \geq 0. \quad (19)$$

The rest of the argument now follows the pattern of the single object case. We conclude that bayesian incentive compatibility is equivalent to monotonicity of $X^i(\cdot)$. Similarly we deduce that (4) applies, i.e.,

$$p^i(\theta^i) \leq \theta^i X^i(\theta^i) - \sum_{\sigma < \theta^i} X^i(\sigma) + \sum_{j \in S^i} r_j^i x_j^i(\theta^i).$$

¹²Both papers consider environments more general than this.

The problem of finding the revenue maximizing mechanism can be formulated as:

$$\max \sum_{i \in N} \sum_{\theta^i} f(\theta^i) p^i(\theta^i) \quad (20)$$

$$\text{s.t. } X^i(\theta^i) \leq X^i(\sigma^i) \quad \forall i, \forall \theta^i \leq \sigma^i \quad (21)$$

$$p^i(\theta^i) \leq \theta^i X^i(\theta^i) - \sum_{\sigma < \theta^i} X^i(\sigma) + \sum_{j \in S^i} r_j^i x_j^i(\theta^i) \quad \forall i, \theta^i \quad (22)$$

$$X^i(\theta^i) = \sum_{j \in S^i} x_j^i(\theta^i) \quad \forall i, \theta^i \quad (23)$$

$$\{x_j^i(\theta^i)\}_{i \in N, \theta^i} \in E \quad (24)$$

Note the similarity to the standard single object auction. By analogy we conclude that the above program can be rewritten as

$$\max \sum_{i \in N} \sum_{\theta^i} f(\theta^i) \left[\theta^i - \frac{1 - F(\theta^i)}{f(\theta^i)} \right] X^i(\theta^i) + \sum_{i \in N} \sum_{\theta^i} f(\theta^i) \sum_{j \in S^i} r_j^i x_j^i(\theta^i) \quad (25)$$

$$\text{s.t. } X^i(\theta^i) \leq X^i(\sigma^i) \quad \forall i, \forall \theta^i \leq \sigma^i \quad (26)$$

$$X^i(\theta^i) = \sum_{j \in S^i} x_j^i(\theta^i) \quad \forall i, \theta^i \quad (27)$$

$$\{x_j^i(\theta^i)\}_{i \in N, \theta^i} \in E \quad (28)$$

Eliminating the $X^i(\theta^i)$ variables and dropping the monotonicity constraint yields the following relaxation:

$$\max \sum_{i \in N} \sum_{j \in S^i} \sum_{\theta^i} f(\theta^i) \left[\theta^i - \frac{1 - F(\theta^i)}{f(\theta^i)} \right] + r_j^i x_j^i(\theta^i) \quad (MAB_\theta) \quad (29)$$

$$\text{s.t. } \{x_j^i(\theta^i)\}_{i \in N, \theta^i} \in E \quad (30)$$

The expression $\theta^i - \frac{1 - F(\theta^i)}{f(\theta^i)}$ should be familiar from the static case. Notice there is no adjustment to the r_j^i . This is because we assumed that the agent would truthfully report r_j^i .

Next, invoke the monotone hazard rate condition on F to ensure that the solution to the relaxation satisfies the omitted monotonicity constraints. Observe that the Gittins index of the state (θ^i, r_j^i) is at least as large as the Gittins index of the state (σ^i, r_j^i) whenever $\theta^i > \sigma^i$. Under the monotone hazard rate condition

$$\theta^i - \frac{1 - F(\theta^i)}{f(\theta^i)} > \sigma^i - \frac{1 - F(\sigma^i)}{f(\sigma^i)} \iff \theta^i > \sigma^i.$$

This suggests modifying each state (θ, r_j^i) to $(\theta - \frac{1 - F(\theta)}{f(\theta)}, r_j^i)$. Hence, the Gittins index of $(\theta^i - \frac{1 - F(\theta^i)}{f(\theta^i)}, r_j^i)$ (virtual Gittins index) is at least as large as the Gittins index of $(\sigma^i - \frac{1 - F(\sigma^i)}{f(\sigma^i)}, r_j^i)$ whenever $\theta^i > \sigma^i$. Therefore, $X^i(\theta^i) > X^i(\sigma^i)$ i.e., the omitted monotonicity constraints are satisfied.

Let $\{\bar{x}_j^i\}_{\{i \in N, \theta^i\}} \in E$ be an optimal solution(MAB_θ). As before we can back out the optimal policy. It will be to award the good to the agent with the highest virtual Gittins index in each period. The problem is that we assumed that each agent i can conceal θ^i , but not r_j^i . What happens when each agent i can conceal r_j^i ? We can use same approach from Section 4.1 to get around this. As before the challenge is to decompose the payments across periods so as to give each agent an incentive to reveal r_j^i as well as to participate. The argument mimics the one given in the case of finding the efficient allocation. There is, however, an important wrinkle. In the very first period, agent i 's private information is $\theta^i + r_j^i$, not simply r_j^i . Thus, in the first period the private information is two dimensional! There are two ways to deal with this. The first is to assume that the set of types θ^i is sufficiently rich so that any misreport $\sigma^i + r_k^i$ can be mimicked by a report $\mu^i + r_j^i$ (see Pavan *et al.* (2008)). Thus, the agent gains nothing by lying about both θ^i and r_j^i . It suffices to misreport θ^i only. The second, is to use a second round of reporting, see Kakade *et al.* (2010).

The optimal mechanism derived here can be summarized as follows. Execute the efficient mechanism assuming a type in each period equal to the virtual ' θ ' plus reported shocks, r_j^i . In effect, the agent receives rents commensurate with the fact that θ is private information but is charged payments equal to their expected future value. This idea originates in Esö and Szentes (2007) (see also Baron and Besanko (1984)).

5 Changing Buyer

The problem of optimally selling a finite number of indivisible goods to buyers arriving over time has been studied for at least half a century. Early papers ignored the incentive issues. See for example Karlin (1962), Derman *et al.* (1972) and Albright (1974). Subsequent papers, for example, Stokey (1979) and Bulow (1982), incorporated strategic buyers. In the non-strategic cases, the optimal policy allocates a good to an agent if its index exceeds a threshold corresponding to the opportunity cost of the good. Indeed, one way to view the literature with strategic buyers is as an attempt to show such index polices continue to be optimal in the presence of private information.

For selfish reasons, I offer Pai and Vohra (2010) as typical of this line of work. It generalizes some of the earlier models and sits exactly on the border between the tractable and intractable. Pai and Vohra (2010) assumes a risk neutral seller seeking to sell C identical and indivisible units over T discrete time periods, indexed by t . Buyers with unit demand arrive over time, and have private valuations for a unit of the good. In particular, buyer i has a valuation v_i for one unit of the good, excess units are worthless. Furthermore, each buyer has an arrival time t_i and a deadline \bar{t}_i . This means buyer i cannot participate in the auction before time t_i or after time \bar{t}_i . Goods assigned to agent i outside the interval $[t_i, \bar{t}_i]$ have no value to him. As will be seen below, it is the ability for the buyer to be strategic about when they depart the 'system' that causes difficulties.

No buyer observes how many units remain, how many buyers are present or what messages are sent. The 3-tuple (v_i, t_i, \bar{t}_i) is buyer i 's type and assumed to be the private information of buyer i . Thus buyer i is in a position to claim a valuation smaller or larger than v_i , an arrival later than t_i and a deadline earlier than \bar{t}_i . Therefore, there is a finite set of buyers' types $\mathbb{T} \subset V \times T^2$. A buyer

with a cumulative probability a of getting allotted good between periods t and \bar{t} enjoys a monetary value of:

$$v(a|(v, t, \bar{t})) = va. \quad (31)$$

In each period t , N_t risk neutral potential buyers arrive. N_t is a non-negative, discrete valued random variable, distributed according to a known probability mass function $g(\cdot)$, with finite support $\{0, 1, \dots, \bar{N}\}$. There is a common prior (a probability distribution with p.d.f. f) on the space of types \mathbb{T} . For any entry time t and exit time $\bar{t} \geq t$, let $f_{(t, \bar{t})}(\cdot)$ and $F_{(t, \bar{t})}(\cdot)$ denote the pdf and cdf respectively of the distribution of a buyer's valuation conditional on his entry at time t and exit at \bar{t} .

Agents arriving at time t receive i.i.d. draws from the posterior of this distribution (i.e. given time of arrival is t). Formally, the probability that an agent arriving at time t has type (v, t, \bar{t}) is

$$f_t(v, \bar{t}) \equiv \mathbb{P}\{\text{type} = (v, t, \bar{t}) | \text{arrival time} = t\} = \frac{f(v, t, \bar{t})}{\sum_{v' \in V, \bar{t}' \geq t} f(v', t, \bar{t}')}. \quad (32)$$

Thus, the valuation and exit time of a buyer arriving at time t is independent of both the number and type of other buyers.

Consider a partial order \succeq on the space \mathbb{T} , defined as

$$(v', t', \bar{t}') \succeq (v, t, \bar{t}) \equiv (v' \geq v) \wedge (t' \leq t) \wedge (\bar{t}' \geq \bar{t}). \quad (33)$$

In other words, a type is said to be partially ordered above another according to \succeq if it has a (weakly) higher valuation, arrives earlier and has a later deadline. One can extend the standard increasing hazard rate condition to respect this partial order, i.e. have a higher hazard rate.

Definition 1 (Monotone Hazard Rate) *A distribution with pdf f on the space of types \mathbb{T} satisfies the monotone hazard rate condition if*

$$(v', t', \bar{t}') \succeq (v, t, \bar{t}) \Rightarrow \frac{f_{(t', \bar{t}')} (v')}{1 - F_{(t', \bar{t}')} (v')} \geq \frac{f_{(t, \bar{t})} (v)}{1 - F_{(t, \bar{t})} (v)}, \quad (34)$$

The corresponding notion of a virtual value is $v - \frac{f_{(t, \bar{t})} (v)}{1 - F_{(t, \bar{t})} (v)}$.

The monotone hazard rate condition implies two restrictions. First, fixing entry and exit times, the distribution of valuations of types with that entry and exit time has an increasing hazard rate. Second: fix (t, \bar{t}) and (t', \bar{t}') such that $[t, \bar{t}] \subseteq [t', \bar{t}']$. The condition implies that:¹³

$$\mathbb{E}[v|(t, \bar{t})] \geq \mathbb{E}[v|(t', \bar{t}')].$$

In the airline motivation we may imagine that impatient buyers are business travelers, and patient buyers are 'leisure' travelers. The hazard rate condition means that the former are expected to have higher valuations than the latter. Alternately, fixing an arrival time, high valuation travelers tend to be more impatient than low valuation travelers. This is orthogonal to the standard assumption in

¹³It actually has a stronger implication, i.e. that $F_{(t, \bar{t})}$ first order stochastically dominates $F_{(t', \bar{t}')}$.

this literature that leisure travelers arrive early, while business travelers arrive late— neither implies nor contradicts the other.

The dynamic nature of the problem, and the fact that possible misreports by a type are restricted by assumption mean the revelation principle cannot be applied directly. Standard proofs of the revelation principle are for settings where any type can send *any* feasible message: by contrast, in our setting the periods in which an agent is present, and can hence send messages, depends on her type.¹⁴ One result in Pai and Vohra (2010) shows that it suffices to consider mechanisms in which buyers never receive an object before their deadline. This allows one to employ a suitable version of the revelation principle to restrict attention to direct revelation mechanisms.

In this setting, Pai and Vohra (2010) ask when the revenue maximizing mechanism can be described by a ‘simple’ index rule? The index rule is easy to describe. To each type, assign an index equal to the virtual value. Give the good to a buyer only if his index (virtual value) exceeds the expected index that extra unit would have generated in the future. Under the hazard rate condition, Pai and Vohra (2010) show the index rule to be optimal.

When such a ‘simple’ index policy is not incentive compatible, a dynamic analogue of Myerson’s ironing procedure is needed. This means the allocation rule cannot be written in a recursive fashion, or alternately, the index of a buyer depends on the allocation rule in the future. As a result the optimal allocation rule is both hard to compute, and hard to describe. To summarize

1. If the deadlines, \bar{t}_i , are common knowledge, then, the distribution of types satisfying an appropriate monotone hazard rate condition is necessary and sufficient for the index rule to characterize the optimal allocation rule.
2. If we drop the assumption of common knowledge of the deadlines, then the appropriate monotone hazard rate condition is *not sufficient*. A form of ironing will be required in spite of the maintained assumption of a monotone hazard rate. Equivalently, the algorithm for computing the optimal allocation rule in this case cannot be recursive.

What causes the difficulty in the second case? Incentive compatibility requires that the allocation rule should be exit monotone, i.e., a later exit from the system increases probability of allotment (all other things being equal). However, the monotone hazard rate condition by itself does not force exit monotonicity in the allocation rule in the same way that it forces the allocation rule to be monotone in valuation. To see why, consider a seller with 1 good for sale, and a buyer arriving in the first period, with a last period deadline and an intermediate valuation. With some probability, a high value buyer will arrive after period 1, get the good, and leave our original buyer with no allocation. If the original buyer had a period 1 deadline instead, then he might even get the good outright if his virtual value is high enough. This suggests that a sufficient condition for allocations to be monotonic in exit time is that types with lower exit times *ceteris paribus* have a substantially lower hazard rate.¹⁵

¹⁴Green and Laffont (1986) identify a sufficient condition for the revelation principle to hold in settings where the set of messages an agent can send is a function of her type. However, they consider a static setting.

¹⁵The difficulty remains even if privately known exit times are replaced with privately known discount rates.

The question of *how* to ‘iron’ in dynamic settings does not admit an easy answer. Recent work by Mierendorff (2009) considers a simpler model (2 periods, 2 buyers, simpler type space). He completely characterizes the optimal mechanism in this setting, but the methods appear technical, and hard to generalize. There is some computational work on implementing ironing in a related dynamic setting (see Parkes *et al.* (2007)). Neither paper will satisfy those interested in a complete characterization, but they do highlight the difficulties of ironing in a dynamic setting.

A sequence of papers beginning with Vulcano *et al.* (2002) can be seen as special cases or variations of Pai and Vohra (2010). In each case, the revenue maximizing mechanism is an index rule of the type described above: allocate in each period to the agent with largest virtual value provided it exceeds a certain threshold. They do so by limiting the dimensionality of the buyers’ type. In Vulcano *et al.* (2002) each buyer is present for at most one period (i.e., $\bar{t}_i = t_i$ for all i). The only private information is the buyer’s valuation for a single unit. Dizdar *et al.* (2010) extends this to allow buyers to demand multiple units. Both marginal valuations and maximum amount demanded are private information but the misreports a buyer can make about the quantity demanded is limited. Gallien (2006) considers a monopolist selling a finite number of identical items over an infinite horizon to time sensitive buyers with unit demands and private value and time of entry, whose inter-arrival times have an increasing failure rate. Time sensitivity of buyers is modeled using a discount factor common to all buyers that is known to the seller.¹⁶ Gershkov and Moldovanu (2009) consider a similar setting, except that the monopolist now sells heterogenous goods which are commonly ranked by all buyers, and each buyer’s private information is one dimensional. Further, their model considers a finite horizon.¹⁷

6 Non-Bayesian Approaches

My summary has focused on Bayesian approaches because I am interested in the conditions that lead to a simple index policy being optimal. There is also a non-Bayesian perspective. One can interpret this literature as identifying conditions under which simple index policies yield outcomes with good worst case behavior. In the changing buyer case, Lavi and Nisan (2004), consider the same model as Vulcano *et al.* (2002) but in a setting where no prior distribution over types is assumed. They propose a index policy and perform a worst case analysis of the revenue achieved. Ng, Parkes and Seltzer (2003) investigate a closely related model where the seller has C perishable units of a good in each period. They exhibit a dominant strategy mechanism, and perform a worst case analysis of the revenue achieved. Hajiaghayi, Kleinberg, Mahdian and Parkes (2005) consider the same model but achieve a better competitive ratio. A variant of the changing buyer case that has its roots in the Secretary problem has been examined in Babaioff, Immorlica, and Kleinberg (2007) as well as Buchbinder, Singh and Jain (2010) . Finally, for the changing type case there is Nazarzadeh, Saberi and Vohra (2008) as well as Babaioff, Slivkin and Sharma (2009).

¹⁶See Board (2008) and Board and Skrzypacz (2010) for related papers.

¹⁷Gershkov and Moldovanu (2008) derive the efficient mechanism in the same setting.

Acknowledgements

I thank two anonymous referees, Mallesh Pai and Ahmad Peivandi for many useful comments.

References

- Akan, M., B. Ata and J. Dana. "Revenue Management by Sequential Screening," manuscript, 2009.
- Albright, S. C. "Optimal sequential assignments with random arrival times," **Management Science**, 21, 60-70, 1974.
- Athey, S. and I. Segal. "An efficient dynamic mechanism," manuscript, 2007.
- Babaioff, M., N. Immorlica, and R. Kleinberg. "Matroids, Secretary Problems, and Online Mechanisms," **Proceedings 18th ACM-SIAM Symposium on Discrete Algorithms**, 2007.
- Babaioff, M., A. Slivkins and Y. Sharma. "Characterizing truthful multi-armed bandit mechanisms," **Proceedings 10th ACM Conference on Electronic Commerce**, 2009.
- Baron, D. and D. Besanko. "Regulation and Information in a Continuing Relationship," **Information Economics and Policy** 1, 267-302, 1984.
- Bergemann, D. and J. Välimäki. "The Dynamic Pivot Mechanism," **Econometrica**, Volume 78, Issue 2, 771-789, 2010.
- Bergemann, D. and M. Said. "Dynamic Auctions: A Survey," Wiley Encyclopedia of Operations Research and Management Science, 2011.
- Bertsimas, D. and J. Nino-Mora. "Conservation Laws, Extended Polymatroids and Multiarmed Bandit Problems; A Polyhedral Approach to Indexable Systems," **Mathematics of Operations Research**, Vol. 21, No. 2, 257-306, 1996.
- Bhattacharya, P., L. Georgiadis and P. Tsoucas. "Extended Polymatroids: Properties and Optimization," **Proceedings of the Conference on Integer Programming and Combinatorial Optimization**, Carnegie Mellon University, Mathematical Programming Society, 1992.
- Board, S. "Durable-Goods Monopoly with Varying Demand," **Review of Economic Studies** Vol. 75, No. 2, 391-413, 2008.
- Board, S. and A. Skrzypacz. "Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-sales," manuscript, 2010.

- Border, K. "Implementation of Reduced Form Auctions: A Geometric Approach," **Econometrica**, Vol. 59, No. 4, 1175-1187, 1991.
- Border, K. "Reduced Form Auctions Revisited," **Economic Theory**, 31, 167-181, 2007.
- Buchbinder, N., K. Jain and M. Singh. "Incentives in Online Auctions via Linear Programs," manuscript, 2010.
- Bulow, J. "Durable-Goods Monopolists," **Journal of Political Economy**, Vol. 90, 2, 314-332, 1982.
- Cavallo, R., D. Parkes, and S. Singh. "Optimal Coordinated Planning Among Self-Interested Agents with Private State," **Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence**, Cambridge, 2006.
- Derman, C. G. J. Lieberman and S. M. Ross. "A sequential stochastic assignment problem," **Management Science**, 18, 349-355, 1972.
- Dizdar, D., A. Gershkov and B. Moldovanu. "Revenue Maximization in the Dynamic Knapsack Problem," **Theoretical Economics**, 2010.
- Esö, P., and B. Szentes. "Optimal Information Disclosure in Auctions," **Review of Economic Studies**, 74, 705-731, 2007.
- Gallien, J. "Dynamic Mechanism Design for Online Commerce," **Operations Research**, Vol. 54, No. 2, 291-310, 2006.
- Gershkov, A. and B. Moldovanu. "Dynamic revenue maximization with heterogeneous objects: A mechanism design approach," **American Economic Journal: Microeconomics**, Vol. 1, No. 2, 168-198, 2009.
- Gershkov, A. and B. Moldovanu. "Efficient Sequential Assignment with Incomplete Information," manuscript, 2008.
- Green, J.R. and J. J. Laffont. "Partially verifiable information and mechanism design," **Review of Economic Studies**, 447-456, 1986.
- Gittins, J. C.. "Bandit Processes and Dynamic Allocation Indices," **Journal of the Royal Statistical Society**, Series B, Vol. 41, No. 2, 148177, 1979.
- Hajiaghayi, M., R. Kleinberg, M. Mahdian and D. Parkes. "Online Auctions with Re-usable

- Goods,” **Proceedings 6th ACM Conference on Electronic Commerce**, 165-175, 2005.
- Kakade, S., I. Lobel and H. Nazerzadeh. “An Optimal Dynamic Mechanism for Multi-Armed Bandit Processes,” manuscript, 2010.
- Karlin, S. “Stochastic models and optimal policy for selling an asset,” *Studies in Applied Probability and Management Science*, K. J. Arrow, S. Karlin and H. Scarf, Eds., Stanford University Press, 148-158, 1962.
- Krishna , V. . *Auction Theory*, Academic Press, San Diego, 2002.
- Lavi, R. and N. Nisan. “Competitive Analysis of Incentive Compatible On-Line Auctions,” **Theoretical Computer Science**, 310(1),159-180, 2004.
- Manelli, A and D. Vincent. “Bayesian and Dominant Strategy Implementation in the Independent Private Values Model, ” **Econometrica**, forthcoming, 2010
- Mierendorff, K. “Optimal dynamic mechanism design with deadlines,” manuscript, University of Bonn, 2009.
- Myerson, R. “Optimal Auction Design,” **Mathematics of Operations Research**, Vol. 6, No. 6, 58-73, 1981.
- Nazerzadeh, H., A. Saberi, and R. V. Vohra. “Dynamic Cost-Per-Action Mechanisms and Applications to Online Advertising,” manuscript, 2008.
- Ng, C., D. Parkes and M. Seltzer. “Virtual Worlds: Fast and Strategyproof Auctions for Dynamic Resource Allocation,” **Fourth ACM Conference on Electronic Commerce**, 2003.
- Parkes, D. and Q. Duong. “An Ironing-Based Approach to Adaptive Online Mechanism Design in Single-Valued Domains,” **Proc. 22nd National Conference on Artificial Intelligence**, 2007.
- Pai, M. and R. V. Vohra. “Optimal Dynamic Auctions,” manuscript, 2010.
- Pavan, A., I. Segal, and J. Toikka. “Dynamic Mechanism Design: Revenue Equivalence, Profit Maximization and Information Disclosure,” manuscript, 2008.
- Schrijver, A. *Combinatorial Optimization*, Springer-Verlag, 2003.
- Stokey, N. “Intertemporal Price Discrimination,” **Quarterly Journal of Economics**, vol 93, No. 3, 355-371, 1979.

Vohra, R. V.. Mechanism Design: A Linear Programming Approach, Cambridge University Press, 2011.

Vulcano, V., G. van Ryzin and C. Maglaras. "Optimal Dynamic Auctions for Revenue Management," **Management Science**, Vol. 48, No. 11, 1388-1407, 2002.