

IRONING IN DYNAMIC REVENUE MANAGEMENT: POSTED PRICES & BIASED AUCTIONS

RAHUL DEB[†] AND MALLESH M. PAI[‡]

SEPTEMBER 7, 2012

ABSTRACT: We consider the design of the revenue maximizing mechanism for a seller with a fixed capacity of C units selling over T periods to buyers who arrive over time. The buyers have single unit demand and multi-dimensional private information— both their value for the object and the deadline by which they must make a purchase are unknown to the seller. This contrasts with previous work where buyers have single dimensional private information— deadlines are publicly observed and only values are private. Here, the optimal mechanism can be computed by running a dynamic stochastic knapsack algorithm. However, these mechanisms are only optimal with private deadlines when the calculated allocation rule is monotone— buyers with higher values and later deadlines should be allocated with higher probability. Such monotonicity only arises in very special cases.

By contrast, in the classic static environment of Myerson [5] monotonicity is only violated for ‘irregular’ value distributions. Myerson characterizes the optimal mechanism by a procedure he calls ‘ironing.’ We characterize the optimal mechanism in our general dynamic environment by providing the dynamic counterpart of ironing. We show that only a subset of the monotonicity constraints can bind in a solution of the seller’s dynamic programming problem. The optimal mechanism can be characterized by ‘relaxing’ these constraints with their appropriate dual multiplier. Further, the optimal mechanism can be implemented by a series of posted prices followed by a ‘biased’ auction in the final period where buyers have the auction biased in their favor depending on their arrival time. Our theoretical characterization complements the existing computational approaches for ironing in these settings (e.g. Parkes et al. [8]).

[†]DEPARTMENT OF ECONOMICS, UNIVERSITY OF TORONTO, RAHUL.DEB@UTORONTO.CA

[‡]DEPARTMENT OF ECONOMICS, UNIVERSITY OF PENNSYLVANIA, MALLESH@ECON.UPENN.EDU

1. INTRODUCTION

Sellers often have fixed deadlines by which they must sell to buyers who arrive and leave the market at different instants. Consider, for example, an airline selling tickets for an upcoming flight or a hotel renting rooms for a future date. In both cases, there is a fixed capacity and a date beyond which this capacity has no value. In such an environment, sellers aiming to maximize profits must solve a dynamic problem: they could sell to buyers currently in the market and ‘lock in’ some revenue, or keep the capacity for buyers who may arrive in the future by which time the currently present buyers may have left. Variants of this problem have been studied in the operations research literature (where it is known as the ‘single leg, multi-period revenue management problem’) and, more recently, in the dynamic mechanism design literature.

Here, we focus on studying the seller optimal mechanism when buyers have multidimensional private information. Specifically, we allow buyers to privately know both how much they value the good and how long they plan to stay in the market. Consider the airline ticket example in which some travelers may be willing to delay their purchase for a potentially better deal. For others, the flight may be a part of a larger trip and hence may choose to change plans and not make a purchase if tickets are currently too expensive. The seller’s mechanism must therefore take this into account. This realistic additional aspect of private information considerably complicates the seller’s dynamic revenue management problem and is avoided by the majority of the literature.

IRONING. To contextualize our results, we first describe a procedure of mechanism design called ‘ironing’. The celebrated paper of Myerson [5] studies revenue maximizing mechanisms for a seller of a single unit of a good to buyers who privately know their values. In this paper, there are two key insights. The first, the Revelation Principle, reduces the problem to one in which each agent truthfully reveals his type to the seller (as opposed to a bid) and is given incentives to do so. The second insight is the notion of ‘virtual value’- this, roughly speaking, is the buyer’s value minus the cost of the incentives to get him to reveal his type. Myerson’s analysis shows that revenue maximization is the same as maximizing efficiency over virtual values *subject to* a monotonicity constraint. This constraint requires that higher value types must win the good more often in expectation. Myerson therefore considers two cases. The first, which he calls the ‘regular case’ involves an assumption on the distribution of values such that the monotonicity constraint is satisfied by the solution to the problem of maximizing virtual efficiency.¹ In this case, a knapsack algorithm where buyers’ weights are given by their virtual values can compute the optimal allocation rule.

The other case is more involved. Myerson defines an ‘ironing’ procedure which modifies the virtual values to restore monotonicity.² The ironed virtual values can then be used to compute the optimal allocation rule via the same knapsack algorithm.

¹Formally, the condition is that the virtual values, $v - \frac{1-F(v)}{f(v)}$, be non-decreasing in v , where f is the density and F is the CDF of the distribution.

²At a technical level, it adds to the virtual value of a buyer the value of the dual variable corresponding to the monotonicity constraint.

DYNAMIC MECHANISM DESIGN AND DYNAMIC IRONING. Various results in dynamic mechanism design follow the same path in a dynamic setting.³ The approach of most papers in this field can be described thus. First, an appeal to a generalization of the revelation principle results in an analogous reduction as the static case. Analysis of the resulting dynamic problem reveals a notion of virtual value for the setting at hand. It can therefore be shown that revenue maximization boils down to maximizing virtual efficiency subject to an appropriate ‘monotonicity’ constraint. It is here that the duality ends. There are normally no conditions one can impose on the distribution of buyer values to ensure that this monotonicity constraint is satisfied. Further, unlike the static Myerson setting, there is no counterpart to the ironing procedure if monotonicity is violated.

With no guidance from economic theory, progress has been made in computational approaches in interesting dynamic settings— see for example Constantin and Parkes [3], Parkes [7] and Parkes and Duong [8]. These papers restore monotonicity via computational methods (see the discussion in Section 4). While they can perform well in practice they are provably suboptimal. Our theoretical results complement this approach and provide a benchmark for them to be measured against. Moreover, our methodology provides an avenue for the design of computational approaches which can yield higher revenue (again, see discussion in Section 4).

OVERVIEW OF MODEL AND RESULTS. We consider a seller who has a finite capacity C to sell over T periods to buyers who arrive over time. Specifically, one buyer arrives in each period, and buyers have a privately known value for a unit and, in addition, privately know whether they are *impatient* (only present in that period) or *patient* (present till the final period). The model, arrival processes, and value distributions are all common knowledge among buyers and the seller. Further, at the time of arrival, a buyer knows all past reports made to and allocation decisions made by the mechanism. It is in this setting that we derive the mechanism that maximizes the seller’s expected revenue, with a Bayes-Nash solution concept.

In a simpler model where buyer’s patience is publicly known (we will refer to this as the public information case), the seller’s optimal mechanism is easier to characterize. She offers a posted price in each period to impatient buyers arriving in that period, while all patient buyers compete in a standard second price auction with a reserve price. Of course, this may not remain incentive compatible when, in addition, the buyers’ patience levels are private, and therefore ironing is required. Ironing in environments with multi dimensional private information typically leads to very complex contracts often involving lotteries. Surprisingly, our main result shows that the optimal mechanism in our environment still takes a surprisingly simple form and is therefore practical. There are two distortions over the mechanism we describe above. First, we show that (under mild regularity conditions), the seller still offers posted prices to impatient buyers. However, these posted prices are weakly higher than the public patience case. This makes it less attractive for a patient buyer to misreport herself as impatient. Second, patient buyers are still invited to bid for an auction that is conducted in the final period. However, this auction is biased and need not award the good to the highest bidder— benefits are provided to early arriving patient buyers for reporting their patience truthfully by distorting the auction in their favor.

³A comprehensive survey of the field of dynamic mechanism design is beyond the scope of this paper. We refer the interested reader to excellent surveys by Bergemann and Said [1] and Vohra [9].

In the public deadline benchmark the optimal mechanism is a dynamic stochastic knapsack algorithm. It allocates a unit if the ‘virtual value’ of the buyer exceeds the expected incremental revenue from 1 unit in the future (computed by dynamic programming). In the private deadline case, for each period t , and each history of what has transpired up to $t - 1$, \bar{h}_{t-1} , the mechanism assigns a penalty/bonus $\eta(\bar{h}_{t-1})$. The same dynamic knapsack mechanism is run, however impatient buyers’ virtual values in that period are penalized by $\eta(\bar{h}_{t-1})$, while patient buyers in that period get a bonus of $\eta(\bar{h}_{t-1})$. This η is the dual variable for the incentive compatibility constraint corresponding to a patient buyer arriving in period t with the highest possible value misreporting as impatient.

RELATED LITERATURE. This paper is closely related to three main papers. The first is the paper of Pai & Vohra [6]. This paper considers a generalization of the model above— multiple buyers may arrive in each period, and a buyer may have any private deadline between the current period and the final period. However in this more general setting, they are only able to provide a limited characterization of the optimal mechanism. Specifically, and in contrast with this work, they focus on the case where the additional private information on departures plays no role in the optimal mechanism. The second paper is a recent work by Mierendorff [4]. This paper characterizes the optimal mechanism in the same setting as ours, but restricted to 2 periods. Our model is thus a special case of [6] and a generalization of [4]. Like [6] and unlike [4], we consider a model with discrete types. This allows for less technical proofs than those of the continuum-of-types-case. Additionally, the discrete type setting is more amenable to algorithmic analysis.

Finally, Board & Skrzypacz [2] consider a continuous time setting where buyers with private value, but commonly known identical discount rates, arrive over time. They show the optimal mechanism for the seller can be implemented by a sequence of posted prices over time, coupled with an auction at the end. By contrast, we show that when buyers have additional private information on their level of patience (the counterpart of discounting in their model), the optimal mechanism involves distorting the auction in favor of more patient buyers.

2. MODEL & NOTATION

There is a seller with C units of an indivisible good to be sold by time T . Time is discrete.

BUYER TYPE SPACE: Buyers arrive and leave over time. A buyer’s type is a 3-tuple with the following components:

- (1) $v \in V$ is the valuation of the agent for one unit of the good. V , the set of all possible valuations is a finite set of positive numbers. While V can be arbitrarily fine, for economy of notation, we consider $V = \{1, 2, \dots, \bar{v}\}$ (which is without loss of generality). Buyers are risk neutral.⁴
- (2) $t \leq T$ is the time that the buyer learns of his demand for unit, we will also call it his entry/ arrival time.

⁴This implies that a buyer with value v who receives the good with probability A making a payment P gets utility $Av - P$.

- (3) The third component x , is either i , denoting that the buyer is impatient (present only for the entry period) or p , denoting that he is patient and willing to wait till the last period. An impatient buyer only values a good he receives in the period of his arrival, while a patient buyer's value for the good does not depend on when he receives it.

The set of possible types who arrive in period t is denoted by the set $\mathbb{T}_t = V \times \{t\} \times \{i, p\}$. We abuse notation slightly and define $v(r_t)$, and $x(r_t)$ as the value and patience level corresponding to a type $r_t \in \mathbb{T}_t$. We can partition the set \mathbb{T}_t into patient types \mathbb{T}_t^p and impatient types \mathbb{T}_t^i .

DISTRIBUTION OVER TYPES: We assume that in every period t one buyer arrives. It is easy to accommodate probabilistic arrivals in our model, we do not do so here to reduce notation. Further, thinking of the discrete time periods as an approximation to 'real' time, it seems reasonable to assume that no two buyers arrive at the same discrete interval if T is large.

A period t arriving buyer is patient/impatient with probability ρ_t^p, ρ_t^i where $\rho_t^p + \rho_t^i = 1$. A patient period t buyer has a private value drawn from a distribution $F_t^p(\cdot)$ and with density $f_t^p(\cdot)$. An impatient buyer has corresponding distribution $F_t^i(\cdot)$ with density $f_t^i(\cdot)$. In what follows, we maintain the following regularity conditions on these distributions:

DEFINITION 1 (Regularity Conditions). *We assume the following regularity conditions on the distribution of buyer's types:*

- (1) *Monotone Hazard rate: Patient buyers arriving in any period have a distribution of values with increasing hazard rate, i.e.*

$$\frac{f_t^p(v)}{1 - F_t^p(v)} \text{ is increasing in } v, \forall t.$$

- (2) *Decreasing Density: Patient buyers arriving in any period have a distribution of values with decreasing density, i.e. $f_t^p(v)$.*
- (3) *Concave revenue function: Impatient buyers arriving in any period have a distribution of values such that $v(1 - F_t^i(v))$ is concave in v (formally, has decreasing differences).*

Condition 1 is a standard assumption in the mechanism design literature and it is satisfied by almost all commonly used distributions. Condition 2 is employed for technical reasons, and has been seen before in several multidimensional settings. Finally, Condition 3 is simply the discrete analog of requiring the monopoly profit from impatient buyers to be concave in the price. This is a ubiquitous assumption in the pricing literature. The latter condition is the only one required for our results. If either of the former two conditions is not satisfied, the solution can still be characterized by additional (static) ironing in the auction for the final period.

THE MECHANISM. We invoke the revelation principle which allows us to restrict attention to direct revelation mechanisms. Put differently, this implies that it is without loss of generality for the seller to ask each arriving buyer to reveal her type. In each period, all past activity is made known to the arriving buyer. In other words, at the time of arrival a buyer knows his own type, past arrivals, and past actions by previous buyers and the mechanism. At each time period, the mechanism decides who will be allocated the good and the payments made by each agent as a function of all previous reports. We now define the mechanism formally.

We begin by defining the relevant histories for the mechanism. We use the variable $c_{\tau,t}$ to denote whether a buyer who arrived at period $\tau \leq t$ has received a good by period t . We set $c_{\tau,t} = 1$ if such an allocation did occur and the if it did not we set $c_{\tau,t} = 0$. Each arriving buyer is asked to report her type $r_t \in \mathbb{T}_t$ (of course, the mechanism needs to provide incentives for the buyer to report her type truthfully).

There are two classes of relevant histories for our mechanism. History h_t is all reports up to and including t , and all allocation decisions up to period $t - 1$. History \bar{h}_t is the history of all reports and allocations up to and including period t . Formally:

$$\begin{aligned}\bar{h}_t &= ((c_{1,t}, \dots, c_{t,t}), (r_1, \dots, r_t)), \\ h_t &= (\bar{h}_{t-1}, r_t).\end{aligned}$$

In a slight abuse of notation, we use $c_{\tau,t}(\bar{h}_t)$ to denote whether in history \bar{h}_t a buyer arriving in τ is allocated at period t or before. Similarly, we use $r_\tau(h_t)$ to denote the announced type of the buyer who arrived in period $\tau \leq t$ according to history h_t . We will let H_t denote the set of all feasible 'beginning of period t ' histories, and \bar{H}_t denote all feasible 'end of period t ' histories.

The seller announces a mechanism before the first period. A mechanism is a set of functions

$$a_{t',t} : H_t \rightarrow [0, 1], \quad q_{t',t} : H_t \rightarrow \mathbb{R}.$$

In words, $a_{t',t}(h_t)$ is the probability with which the buyer who arrived in period $t' \leq t$ and reported $r_{t'}(h_t)$ expects to get the good in period t when history h_t has realized. Similarly $q_{t',t}(h_t)$ is the amount that this buyer has to pay in period t . Since buyers are risk neutral, they only care about their *expected* (interim) allocation probabilities and their *expected* payments, where expectations are taken over the truthful reports of all other players. We denote these by

$$\begin{aligned}A_t(\bar{h}_{t-1}, r_t) &= \sum_{\tau=t}^T \mathbb{E}_{h_\tau} [a_{t,\tau}(h_\tau) | (\bar{h}_{t-1}, r_t)], \\ P_t(\bar{h}_{t-1}, r_t) &= \sum_{\tau=t}^T \mathbb{E}_{h_\tau} [q_{t,\tau}(h_\tau) | (\bar{h}_{t-1}, r_t)].\end{aligned}$$

Notice that the above expectations are conditioned on the current history (\bar{h}_{t-1}, r_t) . Thus the expected utility of a buyer arriving in period t with value v_t who makes a report r_t is $v_t A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t)$. Importantly for the seller, the choice of a 's also determines the distribution of future histories that realize, by changing how many units remain to be allotted in subsequent periods.

SELLER'S PROBLEM. The seller's problem is to maximize expected revenue subject to a few constraints. The first few are relatively straightforward. Individual rationality (IR) requires that the expected profit from participating is non-negative:

$$v(r_t) A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t) \geq 0 \quad \forall r_t \in \mathbb{T}_t, \forall \bar{h}_{t-1}. \quad (\text{IR})$$

The next feasibility constraint simply requires that, at all histories, the seller can never assign more goods than his remaining inventory:

$$\sum_{\tau=1}^t a_{\tau,t}(\bar{h}_{t-1}, r_t) \leq C - \sum_{\tau=1}^{t-1} c_{\tau,t-1}(\bar{h}_{t-1}) \quad \forall r_t \in \mathbb{T}_t, \forall \bar{h}_{t-1}. \quad (\text{Capacity})$$

We impose some standard restrictions on allocations– allocation probabilities must lie between 0 and 1 at all histories, and no buyer is allotted more than one unit.

$$0 \leq a_{\tau,t}(h_t) \leq 1 \quad \forall h_t, \forall \tau \leq t. \quad (\text{Feas1})$$

$$c_{\tau,t-1}(\bar{h}_{t-1}) = 1 \implies a_{\tau,t}(\bar{h}_{t-1}, r_\tau) = 0 \quad \forall \bar{h}_{t-1}, \forall \tau \leq t, \forall r_\tau \quad (\text{Feas2})$$

Finally, the incentive compatibility (IC) conditions, i.e. that a buyer is incentivized to report his true type. We break these into two parts– IC for impatient buyers and IC for patient buyers. This is useful because impatient buyers cannot report themselves as patient. In this sense, their space of reports is identical to bidders in the literature on (static) auction design:

$$v(r_t)A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t) \geq v(r_t)A_t(\bar{h}_{t-1}, r'_t) - P_t(\bar{h}_{t-1}, r'_t) \quad \forall \bar{h}_{t-1}, \forall r_t, r'_t \in \mathbb{T}_t^i, \quad (\text{IC-I})$$

$$v(r_t)A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t) \geq v(r_t)A_t(\bar{h}_{t-1}, r'_t) - P_t(\bar{h}_{t-1}, r'_t) \quad \forall \bar{h}_{t-1}, \forall r_t \in \mathbb{T}_t^p, r'_t \in \mathbb{T}_t. \quad (\text{IC-P})$$

We can now write down the problem the seller must solve:

$$\begin{aligned} & \max_{a,P} \sum_t \mathbb{E}_{h_t} P_t(h_t) && (\text{Expected Revenue}) \\ & \text{subject to } && (\text{IR}), (\text{Capacity}), (\text{Feas1}), (\text{Feas2}), (\text{IC-I}), (\text{IC-P}). \end{aligned}$$

The main difficulty in solving the above problem analytically is in determining which of the above constraints bind. In most single dimensional mechanism design problems, the approach is as follows. The objective function is first maximized by focusing only on local downward IC constraints, that is, the only constraints that are imposed is that the agent should not want to report her type as one lower. It is then shown that the solution to this relaxed problem also satisfies global IC. We too will employ this approach. However, agents in our problem have a second dimension of private information– their patience. This results in a large number of IC constraints. In the next section, we will show that due to the structure of our dynamic problem, ‘most’ of the new constraints can be relaxed. The only IC constraints corresponding to patience that need to be considered are those of a patient buyer of the *highest possible value* (\bar{v}) reporting himself as impatient.

3. THE RELAXED APPROACH AND MAIN RESULT

In this section, we present our main result. As mentioned above, the optimal mechanism is hard to solve for as there are a large number of constraints. We thus approach the problem in three steps. We first argue that the space over which optimization is being done can be reduced by eliminating the prices and only focusing on allocation probabilities. This technique is standard in mechanism design. Second, we then define a relaxed problem with only a subset of the constraints in the original. This is the main challenge as classical static mechanism design does not inform on which constraints are important in dynamic problems with multidimensional private information. We then derive the solution to the relaxed problem. Since our actual problem contains more constraints than the relaxed problem, if the solution to the latter is feasible in the original problem then it must be the solution we are seeking. We now provide a sketch of this argument, details are in the appendix.

3.1. Solving the Seller's Problem

We begin by showing that the space over which optimization is being conducted can be reduced. First, the following lemma argues that it is without loss of generality to restrict attention to allocation rules which only allot to patient buyers at the last period T . This lemma was shown by [6].

LEMMA 1 (Allot at Exit). *It is without loss of generality to only consider allocation rules where patient buyers are only considered for allotment in the final period (impatient buyers must, by definition, be considered for allotment only in their arrival period).*

Loosely speaking, the intuition for this lemma is the following. Any other feasible allocation rule can be converted to an allot-at-exit rule by simply moving the allocation for patient buyers to the last period. The seller has no incentive to allocate to patient buyers earlier- by waiting till the final period T , the seller has the most possible information and therefore can make better allocation decisions, while patient buyers are indifferent between receiving the good earlier or later. Therefore the seller weakly prefers allot-at-exit rules. This lemma simplifies the problem since it implies that $a_{t',t}(h_t) = 0$, whenever $t' \leq t < T$ and $x(r_{t'}(h_t)) = p$.

Buyers in our model can simultaneously misreport both values and patience. It is easy to see therefore that our IC constraints, (IC-I) and (IC-P) can equivalently be written as two constraints. The first considers misreports of value by a buyer (while correctly reporting patience), while the second considers patient buyers misreporting as impatient (while correctly reporting their value). It is easy to see that if these two constraints are satisfied, incentive compatibility with respect to 'double-misreports,' i.e. a buyer misreporting his value and his patience is satisfied for free:

$$v(r_t)A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t) \geq v(r_t)A_t(\bar{h}_{t-1}, r'_t) - P_t(\bar{h}_{t-1}, r'_t) \quad r_t, r'_t \in \mathbb{T}_t^x, x \in \{i, p\} \quad (\text{IC-Value})$$

$$v(r_t)A_t(\bar{h}_{t-1}, r_t) - P_t(\bar{h}_{t-1}, r_t) \geq v(r_t)A_t(\bar{h}_{t-1}, r'_t) - P_t(\bar{h}_{t-1}, r'_t) \quad r_t \in \mathbb{T}_t^p, r'_t = (v(r_t), t, i) \quad (\text{IC-Patience})$$

Since (IC-Value) is a standard incentive compatibility condition from single dimensional mechanism design, we can use it to eliminate payment variables $P_t(\cdot, \cdot)$ from the seller's problem. The following lemma expresses this formally.

LEMMA 2 (Revenue Equivalence). *For any interim allocation rule A , there exists P satisfying (IC-Value) if and only if A is monotone in value, i.e.*

$$\forall t, \bar{h}_{t-1}, v, x \in \{i, p\} : A_t(\bar{h}_{t-1}, (v, t, x)) \leq A_t(\bar{h}_{t-1}, (v+1, t, x)). \quad (\text{Value Monotonicity})$$

Further, if A satisfies (Value Monotonicity), the maximum possible P is:

$$\forall t, \bar{h}_{t-1}, v, x \in \{i, p\} : P_t(\bar{h}_{t-1}, (v, t, x)) = vA(\bar{h}_{t-1}, (v, t, x)) - \sum_{v=1}^{v-1} A(\bar{h}_{t-1}, (v, t, x)). \quad (\text{Payments})$$

Lemma 1 implies that the interim allocations are simply $A_t(\bar{h}_{t-1}, (v, t, i)) = a_{t,t}(\bar{h}_{t-1}, (v, t, i))$ and $A_t(\bar{h}_{t-1}, (v, t, p)) = \mathbb{E}_{h_T}[a_{t,T}(h_T) | (\bar{h}_{t-1}, (v, t, p))]$, and allows us to eliminate a whole set of a 's. Lemma 2 then allows us to eliminate all the P 's from the objective function and the constraints. In

particular, note that when P is given by (Payments), (IR) is satisfied. Further, (IC-Value) is replaced by (Value Monotonicity), while (IC-Patience) can be rewritten as:

$$\sum_{v=1}^{v-1} A_t(\bar{h}_{t-1}, (v, t, p)) \geq \sum_{v=1}^{v-1} A_t(\bar{h}_{t-1}, (v, t, i)) \quad \forall t, \bar{h}_{t-1}, v. \quad (\text{IC-Patience2})$$

Finally, substituting P into the objective function, and changing the order of summation, we have the analog of the celebrated virtual value formula:

$$\begin{aligned} \sum_t \mathbb{E}_{h_t} P_t(h_t) &= \sum_t \mathbb{E}_{\bar{h}_{t-1}} \mathbb{E}_{r_t} P_t(\bar{h}_{t-1}, r_t) = \sum_t \mathbb{E}_{\bar{h}_{t-1}} \mathbb{E}_{r_t} \left[v(r_t) A_t(\bar{h}_{t-1}, r_t) - \sum_{v=1}^{v(r_t)-1} A_t(\bar{h}_{t-1}, (v, t, x(r_t))) \right] \\ &= \sum_t \mathbb{E}_{\bar{h}_{t-1}} \mathbb{E}_{r_t} \left[\varphi_t(r_t) A_t(\bar{h}_{t-1}, r_t) \right]. \quad (\text{Expected Revenue2}) \end{aligned}$$

Here,

$$\varphi_t(r_t) = \begin{cases} v - \frac{1-F_i^i(v)}{f_i^i(v)} & \text{if } r_t = (v, t, i), \\ v - \frac{1-F_i^p(v)}{f_i^p(v)} & \text{if } r_t = (v, t, p). \end{cases} \quad (\text{Virtual Value})$$

Therefore the seller's problem can be rewritten as:

$$\begin{aligned} \max_a \sum_t \mathbb{E}_{\bar{h}_{t-1}} \mathbb{E}_{r_t} \left[\varphi_t(r_t) A_t(\bar{h}_{t-1}, r_t) \right], \\ \text{s.t. } (\text{Value Monotonicity}), (\text{IC - Patience2}), (\text{Capacity}), (\text{Feas1}), (\text{Feas2}). \end{aligned}$$

The 'problematic' constraints are (IC-Patience2). Without further guidance, we do not know which of these constraints bind in the optimal solution- and there is 1 constraint corresponding to each value v after each history. To further understand the structure of the optimal contract, the following lemma is helpful.

LEMMA 3. Suppose for some time t and history \bar{h}_{t-1} , there exists some v^* such that

$$A_t(\bar{h}_{t-1}, (v, t, i)) = \begin{cases} 0 & \text{if } v < v^*, \\ 1 & \text{if } v > v^*. \end{cases} \quad (\text{Zero-One A})$$

Then (IC-Patience2) is satisfied at \bar{h}_{t-1} for \bar{v} if and only if satisfied for all $v \leq \bar{v}$. Therefore it is sufficient to consider the seller's problem with (IC-Patience2) for value \bar{v} at each history \bar{h}_{t-1} .

PROOF. The if direction is trivial, we only need to show the only if direction. To see this, suppose that for some $v' \leq \bar{v}$, (IC-Patience2) is violated or:

$$\sum_{v=1}^{v'-1} A_t(\bar{h}_{t-1}, (v, t, p)) < \sum_{v=1}^{v'-1} A_t(\bar{h}_{t-1}, (v, t, i)),$$

However from the former inequality, $A_t(\bar{h}_{t-1}, (v' - 1, t, i)) > 0$. By the assumption of the lemma, therefore, $A_t(\bar{h}_{t-1}, (v, t, i)) = 1$ for all $v \geq v'$. By definition, $A_t(\bar{h}_{t-1}, (v, t, p)) \leq 1$. Therefore:

$$\sum_{v=1}^{\bar{v}-1} A_t(\bar{h}_{t-1}, (v, t, p)) \not\geq \sum_{v=1}^{\bar{v}-1} A_t(\bar{h}_{t-1}, (v, t, i)). \quad \square$$

Intuitively the assumption of Lemma 3 is satisfied if impatient buyers face posted prices- in this case all values below the posted price get an allocation probability of 0, while all values above the posted price get the good for sure 1. In this case, the Lemma shows that the only patience misreport we need be concerned with is the highest possible value patient buyer misreporting as impatient. Roughly speaking, the reason is that all A 's are probabilities- therefore they must lie between 0 and 1. If for some other value $v' < \bar{v}$, (IC-Patience2) is violated, it will also be violated for \bar{v} since $A(\bar{h}_{t-1}, (v, t, i)) = 1 \geq A(\bar{h}_{t-1}, (v, t, p))$ for any $v > v'$.

This simple insight allows us to greatly reduce the number of constraints considered- we only need to (IC-Patience2) constraints corresponding to each history \bar{h}_{t-1} and the highest possible value \bar{v} .

Our final step simply involves taking a relaxation of each of these constraints to the objective function. We will only give a heuristic sketch here, the formal proof is left in the appendix (i.e. verifying that a relaxation approach works for this setting etc.). Since there will be one constraint for each history \bar{h}_{t-1} , let us denote the corresponding dual variable by $\eta(\bar{h}_{t-1})$. Collecting terms, the problem can be written as:

$$\begin{aligned} \max_a \sum_t \mathbb{E}_{\bar{h}_{t-1}} \mathbb{E}_{r_t} \left[\varphi_t^M(r_t, \bar{h}_{t-1}) A_t(\bar{h}_{t-1}, r_t) \right], & \quad (\text{Relaxed OPT}) \\ \text{s.t. } & \quad (\text{Value Monotonicity}), (\text{Capacity}), (\text{Feas1}), (\text{Feas2}). \end{aligned}$$

Here,

$$\varphi_t^M(r_t, \bar{h}_{t-1}) = \begin{cases} v - \frac{1-F_t^i(v)}{f_t^i(v)} - \frac{\eta(\bar{h}_{t-1})}{f_t^i(v)} & \text{if } r_t = (v, t, i), \\ v - \frac{1-F_t^i(v)}{f_t^p(v)} + \frac{\eta(\bar{h}_{t-1})}{f_t^p(v)} & \text{if } r_t = (v, t, p). \end{cases} \quad (\text{Modified VV})$$

Note that in the absence of (Value Monotonicity), the solution to problem (Relaxed OPT) can be computed by a dynamic stochastic knapsack algorithm where the 'weight' of a buyer reporting r_t at history \bar{h}_{t-1} is given by $\varphi_t^M(r_t, \bar{h}_{t-1})$. In the appendix, we show that the maintained regularity assumption on the distribution, Definition 1, implies that the output of the dynamic stochastic knapsack algorithm satisfies both (Zero-One A) and (Value Monotonicity).

To roughly see why, note that at any stage before the final period, the algorithm trades off the current weight $v - \frac{1-F_t^i(v)}{f_t^i(v)} - \frac{\eta(\bar{h}_{t-1})}{f_t^i(v)}$ with the expected incremental value from one extra unit in the future, $\Delta_f(h_t)$. Condition (3) of Definition 1 implies that $v - \frac{1-F_t^i(v)}{f_t^i(v)} - \frac{\eta(\bar{h}_{t-1})}{f_t^i(v)}$ is increasing in v , implying (Value Monotonicity) and (Zero-One A) for impatient buyers. In the final period, the algorithm will assign the remaining goods to the highest non-negative φ^M s among patient buyers. It follows that conditions (1) and (2) of Definition 1 implies (Value Monotonicity) for patient buyers.

3.2. Main Result

We can summarize the argument above into our main result, which characterizes the form of the optimal mechanism. Importantly, the optimal mechanism can be "implemented" in a simple

and transparent way. This implies that the mechanism is straightforward enough to allow unsophisticated buyers to purchase optimally. This is critical for utilizing a mechanism in practice as it has been shown that even reasonably sophisticated bidders do not understand the rules of complex mechanisms (such as the Vickrey auction).

PROPOSITION 1. *The seller’s optimal mechanism can be described as follows— as buyers arrive, they are asked to declare whether they are patient or impatient. An impatient buyer is offered a posted price (a take it or leave it offer at price). Patient buyers are invited to bid in an auction that will clear in the final period. The auction need not be a ‘standard’ auction, i.e. the highest bidder need not win the good. Buyers who bid in different periods will be treated differently, even if their distributions are the same (i.e. $f_t^p = f_t^p$).*

4. DISCUSSION & CONCLUDING REMARKS

As we mentioned earlier, prior to this work, the main approaches to ironing in a dynamic setting were computational. These approaches identified that the ‘problematic’ constraints were (IC-Patience2). However, in the absence of a counterpart to Lemma 3, this approach instead imposes the *stronger* constraint: $A_t(\bar{h}_{t-1}, (v, t, p)) \geq A_t(\bar{h}_{t-1}, (v, t, i))$ for each v . This (stronger) constraint is easier to work with. When considering allotting the good to buyer (v, t, i) at time t after history \bar{h}_{t-1} , the seller ‘looks ahead’ in his dynamic program to whether he would have allotted the good to the buyer, had the buyer instead been of type (v, t, p) (an exact look-ahead is often too taxing computationally, so Monte Carlo simulations are used). If not, the algorithm allots to neither and therefore ensures that the stronger monotonicity constraint is satisfied.

By contrast, our mechanism is provably optimal *if* the correct η ’s are known. Indeed once the η ’s are known, a dynamic stochastic knapsack with φ^M as the weights will output the optimal mechanism. Efficient algorithms to find $\eta(\cdot)$ are beyond the scope of this paper. In simple settings, such as the case of 2 time periods, there is a single η which can be found by binary search up to a desired tolerance (note that $0 \leq \eta \leq 1$). However, for even the case of 3 time periods, naïve approaches are computationally forbidding. Given the importance and relevance of the problem at hand, we hope that this will spur further work characterizing how to efficiently find (or approximate) the optimal η ’s, and leave it as an open question.

EXTENSIONS. The main assumption in this paper is on buyers’ patience— each buyer is either completely impatient or completely patient. Generalizing this to allow a buyer in period t to have any deadline $t \leq \bar{t} \leq T$ is non-trivial. If the buyer has a deadline of \bar{t} , $t < \bar{t} < T$, then from his perspective in period t , he faces a random price in \bar{t} . This is because the price will depend on the other arrivals between t and \bar{t} . We therefore lose the structure provided by Lemma 3, which allows us to relax most of the (IC-Patience) constraints. This in turn makes characterization of the optimal mechanism difficult. We leave this as an open question.

REFERENCES

- [1] D. Bergemann and M. Said. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.
- [2] Simon Board and Andy Skrzypacz. Revenue management with forward looking buyers. Technical report, 2010, Stanford University.
- [3] F. Constantin and D.C. Parkes. Self-correcting sampling-based dynamic multi-unit auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 89–98. ACM, 2009.
- [4] Konrad Mierendorff. Optimal dynamic auctions with deadlines. Technical report, Bonn University, 2010.
- [5] R.B. Myerson. Optimal auction design. *Mathematics of Operations Research*, pages 58–73, 1981.
- [6] Mallesh M Pai and Rakesh Vohra. Optimal dynamic auctions and simple index rules. Technical report, Northwestern University, 2010.
- [7] D. Parkes. When analysis fails: Heuristic mechanism design via self-correcting procedures. *SOFSEM 2009: Theory and Practice of Computer Science*, pages 62–66, 2009.
- [8] D.C. Parkes and Q. Duong. An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proceedings of the National Conference on Artificial Intelligence*, volume 22, page 94, 2007.
- [9] R.V. Vohra. Dynamic mechanism design. *Surveys in Operations Research and Management Science*, 17(1):60–68, 2012.