Dynamic Mechanism Design under Positive Commitment

Ilan Lobel and Renato Paes Leme

February 13, 2017

Abstract

We consider a firm that sells products that arrive over time to a buyer. We study this problem under a notion we call *positive commitment*, where the seller is allowed to make binding positive promises to the buyer about items arriving in the future, but is not allowed to commit not to make further offers to the buyer in the future. We model this problem as a dynamic game where the seller chooses a mechanism at each period subject to a sequential rationality constraint. We characterize the perfect Bayesian equilibrium of this dynamic game. We prove the equilibrium is long-term efficient and that the seller's revenue is a function of the buyer's ex ante utility under a no commitment model.

1 Introduction

Consider a setting with a seller (she) and a buyer (he) interacting over time. In each period, the seller obtains one new product she can sell to the buyer. The products are heterogeneous and the buyer learns his value for a given product just before the transaction takes place. What kind of marketplace should the seller set up to sell her products to the buyer over time?

Solving this problem is crucial to the design of certain marketplaces such as the market for online advertisement. Online ads arrive in the form of a stream of impressions (or user views) that can be sold to advertisers. Impressions are highly differentiated products whose value to advertisers depend on many pieces of information, from the webpage being viewed to the user's IP address and cookie data. Online ads are currently sold via highly fragmented ecosystem. One part of this ecosystem consists of real-time auctions, where individual impressions are sold milliseconds before they are delivered. Online ads are also sold via contracts, where a seller agrees to deliver a specific number of impressions to a buyer over a given time horizon for a pre-specified price.

Consider the seller's problem. She would like to design a revenue-maximizing marketplace for selling her products to the buyer. If she is familiar with the classical mechanism design literature, her first instinct might be to sell the stream of products via a sequence of optimal auctions, where each individual auction is designed à la Myerson (1981). This market design would closely mimic the real-time auctions that are currently in use for selling online ad impressions. For the case of a single buyer, which is the focus of our paper, the optimal auction reduces to a posted pricing scheme. Therefore, a seller well-versed in the classical mechanism design literature would probably choose to simply post items for sale as they arrive and offer them to the buyer at the revenue-maximizing price. If, for example, the buyer valuations are drawn independently from a uniform distribution over the $[0, 1]$ interval, the buyer would post each item at a price of $1/2$ and would sell each item with probability of 50%, leading to an average per-item revenue of $1/4$. In fact, if the seller has no commitment power, this is indeed the best she could do. No commitment power means the seller is not able to make any binding promises about future allocations and prices.

However, the strategy outlined above is too pessimistic and, therefore, leaves a significant amount of revenue on the table. In most realistic situations, the seller has at least some commitment power which she can use to extract more revenues. Let's now assume the seller is well-versed in the recent literature on dynamic mechanism design. In most of this literature (see, for example, Kakade et al. (2013) or Pavan et al. (2014)), the seller is assumed to have full commitment power. With full commitment power, the seller's optimal strategy involves offering the entire future stream of products to the buyer at a take-it-or-leave-it price that is equal to the buyer's expected value for the entire stream of impressions. In our example where the buyer valuations are drawn uniformly from $[0, 1]$, the seller can offer the buyer the entire set of products arriving after the first period at a price of $1/2$ per item. In equilibrium, the buyer accepts the offer, leading to a long-term average per-item revenue of $1/2$, which is double the revenue from the environment without commitment. Motivated by this large potential boost in revenues, several recent papers have proposed ways to incorporate bundling ideas into market design. This includes work on dynamic bundling by Ashlagi et al. (2016), Balseiro et al. (2016), Mirrokni et al. (2016a) and Mirrokni et al. (2016b) as well as work on preferred deals by Mirrokni and Nazerzadeh (2017).

While the no commitment assumption is too pessimistic, the full commitment assumption is excessively optimistic. The full commitment model assumes the seller is able to extract the full surplus via a take-it-or-leave-it offer. This equilibrium outcome crucially depends on what happens if the buyer declines the offer. In the revenue-maximizing contract, the seller is committed not to trade with the buyer after her initial offer is rejected. However, it would be difficult to enforce that commitment. After having her first offer rejected, it would be in the seller's interest to make a new offer to the buyer. To be more precise, after having her offer rejected in the first period, the seller would find herself in the second period in an almost identical position to her initial position: she will have a stream of products for sale and a potentially willing buyer. She will be very tempted to make a new offer at that point, and the buyer could take that into account when making his initial accept/reject decision.

The story in the paragraph above supports adding some kind of subgame perfection refinement to our model. With subgame perfection, both the seller and the buyer will take their future actions into account when negotiating. However, merely adding a subgame perfection requirement is equivalent to assuming no commitment. Hence, almost all papers in the literature either assume no or full commitment.

We argue that not all promises are equally credible. There are some promises the seller could make that would be easy to enforce and there are other promises that would not be enforceable. In particular, a promise to deliver a unit of inventory in a given period in the future should be easy to enforce in a court of law. Therefore, a realistic model of commitment should allow sellers to offer future goods to buyers. Such contracts are routinely signed and enforced in practice. At the same time, a promise not to trade in the future with a particular buyer is not particularly credible and unlikely to be enforceable through the justice system. There are of course ways the seller could, in theory, enforce such a contract. For example, the seller could sign a contract with a charitable organization promising the donate the items to them if the buyer does not accept the first offer. However, such contracts are mostly theoretical curiosities and not typically used in practice. After all, under such a contract, a minor negotiating mistake could cause the seller to donate her entire inventory.

In this paper, we propose a positive commitment model that is able to capture this distinction that is quite important in practice, but is rarely captured in the academic literature. Under positive commitment, the seller is able to make credible promises regarding future allocations, but is not able to make credible promises regarding future mechanisms that she will make available. Under this framework, the seller must choose a new mechanism to offer the buyer at each point in time, and this sequence of mechanisms must satisfy a sequential rationality condition. At the same time, the seller's sequence of mechanisms must satisfy a dynamic feasibility rule. That is, the seller cannot resell a unit of inventory that she has previously committed to allocating to the buyer.

Our main finding is that the seller should price the current item and future items at different prices. The current item should be sold at the standard optimal monopoly price. Meanwhile, future items should be sold at a different price that we denote the *positive commitment price*. The positive commitment price is strictly less than the buyer's expected value, which would be the per-item price under full commitment. This solution mimics the real-world phenomenon of having two sales channels for selling online ads, one for real-time auctions and the other involving contracts that guarantee future allocations.

The idea behind the proof is as follows: suppose we arrive at the last period and that period's product has not yet been sold. The seller's optimal strategy then is to offer the product to the buyer at the optimal monopoly price. Therefore, the buyer knows in advance that he has an option to buy that product at the optimal monopoly price after learning his valuation for it. Thus, he shouldn't accept any offer from the seller for that product that guarantees him less utility than his ex ante expected utility under the optimal price. The seller must therefore make an offer to the buyer that is cognizant of his endogenous outside option, which consists of waiting until the item is offered at the monopoly price. In our running example of valuations being uniform in $[0, 1]$, the agent's ex ante expected utility under static pricing is $1/8$. Therefore, the seller must discount her advance offer from the expected item value of $1/2$ by $1/8$. Thus, both the positive commitment price and the seller's long-term average per-item revenue are $3/8$.

In equilibrium, the buyer purchases all the items in the stream of products at the positive commitment price (except for the first period item). This equilibrium has several appealing properties when compared to the no commitment and full commitment equilibria. In contrast with the no commitment solution, the allocation of items is welfare maximizing. That is, as in the full commitment case, there is no allocation distortion due to information rents (except for the first period item). At the same time, the seller does not extract 100% of the surplus, which is a hard-to-believe equilibrium outcome of the full commitment model. Even though there is no information-rent-related distortion, the ex-ante expected value of the information rent plays a significant role in determining in our model how the surplus gets split between the seller and the buyer.

Closely related papers. Our positive commitment framework is inspired by the limited commitment models of Skreta (2006, 2015) and Deb and Said (2015).

Skreta (2006) considers the interaction between a seller and a buyer over a finite horizon. The seller owns one item. The seller chooses a mechanism in period 1. If the item is still unallocated at the end of the period, she chooses a new mechanism in the next period. As in our paper, the seller cannot commit not to trade with the buyer in the future. Skreta starts her paper by arguing that the revelation principle does not hold in an environment without commitment. Regardless, she shows that she can still optimize over mechanisms with arbitrarily large message spaces. Her main result is that the seller should use a sequence of posted prices. We do not use Skreta's technique in our paper, but we do imitate her in allowing for arbitrarily large message spaces and then reconstructing a result similar to the revelation principle from scratch.

Her paper is more complex than ours in some dimensions, but simpler in others. She allows for persistent private information, something we rule out by assuming the buyer valuations over time are independently distributed. At the same time, her model includes only one good and does not allow for future trading, as our model does. Skreta (2015) is a follow-up paper that considers a multi-buyer problem and finds the sequence of optimal auctions. It is still considers only a single-item problem.

Deb and Said (2015) consider a dynamic screening problem with limited commitment. Buyers arrive over two periods and consumption occurs in the second period. There are no capacity constraints. As in our work, the seller can commit to future allocations in period 1, but she cannot commit to the mechanism being offered in period 2. Their main finding is that the equilibrium first-period contract is non-monotone, with the seller contracting with low and high types immediately, but deferring contracting with intermediate types until the second period.

2 The Model

We consider a model with one seller and one buyer who interact over T time periods. At each period $t = 1, 2, ..., T$, a new indivisible good becomes available to the seller, which we refer to as item t. The seller's outside value for all items is normalized to zero. The buyer's value for item t is given by v_t. We assume that the valuations v_t are drawn i.i.d. from a commonly known distribution $F(\cdot)$ with density $f(\cdot)$. We denote the expected value of

v_1 by μ. To simplify the presentation,[1] we assume the valuation distribution satisfies the standard regularity assumption that the virtual value function $\phi(v) = v - (1 - F(v))/f(v)$ is a non-decreasing function (see Myerson (1981)). We denote by p^* the seller's optimal static reserve price, i.e., $\phi(p^*) = 0$, which we assume to be unique to avoid multiplicity of equilibria. We assume the buyer learns his valuation for item t at the beginning of period t. After period t, item t has no value to either the seller or the buyer. Both the seller and the buyer desire to maximize the sum of their utilities over time, or equivalently, their average per-item utilities. The seller's utility is simply the expected sum of payments she obtains from the buyer over the T periods, and the buyer's utility is the expected value of the goods allocated to him minus the payments he makes to the seller. We do not include discounting in our model, though discounting would not significantly alter our results.

The setting described above is motivated by a problem in online advertising. In this context, the seller is a publishing platform and the buyer is an advertiser. Each item is a particular impression that the platform might sell to the advertiser. At period t, impression t arrives with the specific characteristics of the user who will potentially see an ad, which could include the IP address and cookie data. The characteristics of the user are only learned at period t and, therefore, the advertiser only learns v_t at the beginning of period t. After period t, the user is no longer on the publisher's website and, therefore, the impression is gone.

No commitment and full commitment. There several different contracting environments we could consider, which vary in the level of assumed seller commitment power. The best understood environments are the no commitment and the full commitment ones. Without commitment power, the seller's best strategy is to use the static reserve price p^* in each period. In this case, the seller's expected per-item revenue is $p^*(1 - F(p^*))$. Under full commitment, the seller's optimal strategy would be to offer the period 1 item at price p^* and all future items at a take-it-or-leave-it offer at price of $E[\sum_{t=2}^{T} v_t] = \mu(T-1)$ — see, for example, Kakade et al. (2013) or Pavan et al. (2014). In equilibrium, the buyer would accept the offer, leading to a long-term average revenue of μ. The seller's revenue is significantly greater under full commitment than under no commitment, as she is able to extract almost the entire welfare of the system under full commitment. For example, consider the case where the buyer's valuations are drawn from a uniform distribution over the interval $[0, 1]$. In this case, the static reserve price is $p^* = 1/2$ and the seller earns an average revenue of $1/4$ under no commitment. Under full commitment, the seller is able to extract almost the full surplus, i.e., her long-term average revenue is $1/2$. For this particular distribution of valuations, commitment power allows the seller to double her revenues.

The seller's optimal auction design and her ability to extract revenue thus crucially depend on her commitment power. How much commitment power does a seller truly have? In this paper, we argue that the no commitment power assumption is too pessimistic and the full commitment power assumption is too optimistic. The no commitment power model

[1] We note in the last section of the paper that we can extend the results to any distribution via the standard ironing technique. We opt however to present the result for regular distributions not to unnecessarily complicate the notation.

assumes that sellers cannot make binding promises to deliver units of inventory in the future. This is obviously an unrealistic assumption. Essentially all supply chain contracts used in practice involve promises of future delivery of goods, and such contracts are routinely enforced through the court system. In contrast, the full commitment power assumption involves the seller making binding promises that would be hard to enforce in a court of law. In particular, the full commitment power solution requires the seller to make a take-it-or-leave-it offer for $T-1$ items for the price of $\mu(T-1)$. The leave-it part of the offer is the crucial bit that would be hard to enforce. Suppose the buyer says "no, thanks" after receiving the offer. The optimal contract then requires the seller not to try to sell those units anymore. That is, the seller would to need to sign some sort of enforceable contract that would prohibit herself from trading with the buyer in the case where the buyer is not interested in purchasing an entire stream of items upfront at the seller's chosen price.

Positive commitment. We now propose a commitment framework that we believe is more realistic than the excessively pessimistic no commitment model and the excessively optimistic full commitment model. Under positive commitment, the seller is able to commit future units of inventory, but that the seller is not able to commit to what other offers she might make to the buyer in the future.

Without full commitment power, we cannot rely on the revelation principle to simplify our mechanism design problem. Instead, we must allow for general messaging spaces, which could potentially be larger than the space of valuations. We must then search for perfect Bayesian equilibria (PBE) of the resulting game.

Let X_t represent the set of unsold units of inventory at beginning of time t, with $X_1 = \{1, 2, ..., t\}$ and $X_t \subseteq \{t, t+1, ..., T\}$ for all $t = 1, ..., T$. We will refer to X_t as the state of the system at period t and to $\mathcal{X}_t$ as the set of possible states at time t. At each period $t = 1, ..., T$, the seller will choose a mechanism to offer to the buyer at time t. The set of available mechanisms at time t is denoted by $\mathcal{M}_t$. A potential mechanism that is offered at time t is a triplet $M_t = (A_t, y_t, r_t)$, where A_t represents a messaging space, y_t represents an allocation rule and r_t represents a payment rule. Since we are not relying on the revelation principle, the messaging set A_t could be an arbitrary set. The only restriction we impose is that it must contain a "no trade" option, that we represent by $a_\emptyset \in A_t$.

The allocation rule $y_t : A_t \to 2^{\{t,t+1,...,T\}}$ determines the goods that will be allocated to the buyer as a function of his chosen message $a_t \in A_t$. To be feasible, an allocation rule can only allocate to the buyer goods that the seller still owns, i.e., $y_t(a_t) \subseteq X_t$ for all $a_t \in A_t$. The state transition dynamics are given by

$$X_{t+1} = X_t \setminus (\{t\} \cup y_t(a_t)) \tag{1}$$

since item t becomes obsolete at the end of period t and all units allocated to the buyer at time t are removed from the system. The function $r_t : A_t \to \mathbb{R}$ determines the buyer's payment to the seller at time t. If the buyer chooses the "no trade" message $a_\emptyset \in A_t$ in period t, he will not be charged a payment and no items will be allocated to him in that period, i.e., $r_t(a_\emptyset) = 0$ and $y_t(a_\emptyset) = \emptyset$.

The buyer's purchase set is the union of all items allocated by the seller to him, i.e.,

$P = \cup_{t=1}^{T} y_t(a_t)$. The buyer's utility is given by the value he earns from the items allocated to him minus his payment to the seller:

$$U^b = \sum_{t \in P} v_t - \sum_{t=1}^{T} r_t(a_t).$$

The seller's utility is given by the sum of payments she receives from the buyer:

$$U^s = \sum_{t=1}^{T} r_t(a_t).$$

We consider the set of perfect Bayesian equilibria (PBE) of this game. For most games, defining a PBE requires defining strategies and beliefs. Because of our assumptions that valuations are i.i.d. and that the buyer only learns his valuation for item t at time t, there is no persistent information asymmetry in our model. Therefore, we can define a PBE in our model solely in terms of the strategies σ^s and σ^b for the seller and the buyer. A strategy of the seller $\sigma^s = \{\sigma_t^s\}_{t=1,\dots,T}$ determines which mechanism the seller should offer at each period t given the state of the system, i.e., $\sigma_t^s : \mathcal{X}_t \to \mathcal{M}_t$. A buyer's strategy $\sigma^b = \{\sigma_t^b\}_{t=1,\dots,T}$ determines the action of the buyer given at period t given the state of the system $X_t \in \mathcal{X}_t$, the seller's chosen mechanism $M_t \in \mathcal{M}_t$ and his valuation $v_t \in \mathbb{R}^+$, i.e., $\sigma_t^s : \mathcal{X}_t \times \mathcal{M}_t \times \mathbb{R}^+ \to A_t$. We refer to a pair $\sigma = (\sigma^s, \sigma^b)$ as a strategy profile.

Definition 1. *A strategy profile σ constitutes a perfect Bayesian equilibrium (PBE) if it satisfies sequential rationality. That is, σ^s maximizes the seller's expected utility starting from any period t and state X_t assuming the buyer will play according to σ^b. At the same time, σ^b maximizes the buyer's expected utility starting from any period t and triplet (X_t, M_t, v_t) assuming the seller will play according to σ^s.*

We note that the full commitment solution, which involves a take-it-or-leave-it offer at time $t = 1$ for items 2 through T, is not a PBE since the seller will be tempted to offer a new mechanism to the buyer at time $t = 2$ if her offer at time $t = 1$ is rejected (we represent rejection by the buyer choosing message $a_\emptyset$).

3 Dynamic Mechanism Design

We now use dynamic mechanism design to study the PBE of the problem defined in the prior section. We use the notation $\hat{U}^b = \lim_{T\to\infty} U^b/T$ and $\hat{U}^s = \lim_{T\to\infty} U^s/T$ to represent the buyer's and the seller's respective long-term average utilities. The notation Z^b will denote the buyer's ex ante expected utility under the optimal price in a single-shot problem ($T = 1$), i.e.,

$$Z^b = \int_{p^*}^{\infty} (x - p^*) dF(x). \tag{2}$$

We now state the main result of our paper.

Theorem 1. *Under any PBE, the seller's long-term average utility is $\hat{U}^s = \mu - Z^b$ and the buyer's long-term average utility is $\hat{U}^b = Z^b$.*

In the remainder of this section, we will provide the analysis that proves the theorem above. Because we do not have full commitment, we cannot rely on any of the standard versions of the revelation principle (see Myerson (1981, 1986)). Instead, we need to recreate the revelation principle from scratch for our setting. We now define a series of terms that we will use in our construction. Consider any PBE σ. For any pair of periods $t \leq r$, state $X_t \in \mathcal{X}_t$ and buyer's value $v \in \mathbb{R}^+$, we let $q^\sigma_{t,r}(X_t, v)$ represent whether item r is allocated at period t to the buyer given the state X_t and $v_t = v$ under the equilibrium σ. Similarly, for any t, we let $p^\sigma_t(X_t, v)$ represent the period t payment to seller given state X_t and buyer's value $v_t = v$ under the equilibrium σ. Based on Eq. (1), we define $\hat{X}^\sigma_{t+1}(X_t, v)$ as the period $t+1$ state given period t state X_t and buyer's value $v_t = v$ under the equilibrium σ.

We also define $U^{b,\sigma}_t(X_t, v, z)$ as the value-to-go of the buyer at time t given state X_t assuming his true value is $v_t = v$ but he acts as if his true value was z — that is, plays action $\sigma^b_t(X_t, M_t, z)$ — assuming that all future actions of both the seller and the buyer will conform to the equilibrium σ. We further define $V^{b,\sigma}_t(X_t)$ to be the buyer's true period t expected value-to-go before he learns his period t value v_t, under equilibrium σ. Thus,

$$V^{b,\sigma}_t(X_t) = E_{v_t}[U^{b,\sigma}_t(X_t, v_t, v_t)].$$

The buyer's value-to-go function must satisfy the following equation:

$$U^{b,\sigma}_t(X_t, v, z) = v \cdot q^\sigma_{t,t}(X_t, z) + \mu \sum_{r \in X_t \setminus \{t\}} q^\sigma_{t,r}(X_t, z) - p^\sigma_t(X_t, z) + V^{b,\sigma}_{t+1}\left(\hat{X}^\sigma_{t+1}(X_t, z)\right), \quad (3)$$

where the first term on the right-hand side captures the value of the period t item, the second term captures the expected value of any future items allocated at time t, the third term captures the period t payment and the final term captures the utility obtained by the buyer from trading in the future.

Modeling the seller's problem. We can determine the period t value-to-go function for the seller via an optimization formulation where the seller optimizes over current and future allocations, as well as payments, in order to maximize present plus future payments. The seller's value-to-go $V^{s,\sigma}_t(X_t)$ at period t and state X_t under equilibrium σ is thus given by

$$\begin{aligned} V^{s,\sigma}_t(X_t) = \sup_{q^\sigma_{t,t}(\cdot),\ldots,q^\sigma_{t,T}(\cdot),p^\sigma_t(\cdot)} \; & E_{v_t}\left[p^\sigma_t(X_t, v_t) + V^{s,\sigma}_{t+1}(\hat{X}^\sigma_{t+1}(X_t, v_t))\right] \\ s.t. \quad & \text{buyer does not want to deviate from } \sigma^b_t \text{ for any } v_t \text{ at period } t \\ & q^\sigma_{t,r}(X_t, v) \in \{0,1\} \qquad \forall r \in X_t, v \in \mathbb{R}^+ \\ & q^\sigma_{t,r}(X_t, v) = 0 \qquad\quad\; \forall r \notin X_t, v \in \mathbb{R}^+ \end{aligned}$$

The revelation principle. The seller's problem involves the feasibility constraints — only items still available at state X_t can be allocated to the buyer — and buyer incentive constraints. Let's now try to simplify the buyer's incentive constraints. In direct revelation mechanisms, we only need to reason about deviations where buyers with value v act as if they had a different value. Since here buyers are allowed to choose any message from A_t, a buyer could deviate by choosing a message that is not chosen by any other type. In the spirit of the revelation principle, we argue that simply removing from the seller's mechanism messages that are not optimal by any buyer type doesn't alter the equilibrium and allows us to associate messages with buyer types. This excludes the "no trade" message $a_\emptyset$, which we must keep regardless of whether it's used by any type since the seller does not have the power to remove the message.

Let A_t be the set of messages the seller chooses in the solution above at time t under state X_t. Let $\tilde{A}_t$ be a subset of A_t where only messages that are used by the buyer in equilibrium are kept, plus the "no trade" message $a_\emptyset$. That is,

$$\tilde{A}_t = \{a \in A : \exists v \in \mathbb{R}^+ \text{ s.t. } \sigma_t^b(X_t, M_t, v) = a\} \cup \{a_\emptyset\}.$$

Let's create a new mechanism $\tilde{M}_t$ which is a restricted version of M_t where only messages in $\tilde{A}_t$ are allowed. We now argue that removing the actions $A_t \setminus \tilde{A}_t$ from the mechanism does not change either the seller's or the buyer's value-to-go functions. The seller clearly does not lose or gain from having these unused actions available to the buyer at time t. The argument for the buyer is slightly more subtle. In principle, the buyer could gain from having unused actions in a dynamic game because they serves as threats. However, for this to be true, the buyer must use the action in some subgame, perhaps one that is never reached in equilibrium. However, we are discussing actions that are never used in any subgame. We can therefore remove unused actions from the action set via backwards induction, starting from period T and going back all the way to period 1. Therefore, an independence of irrelevant alternative holds and the new mechanism will generate the exact same value-to-go for both the seller and the buyer. We can therefore simplify the original incentive constraints and replace them with simpler constraints:

$$\begin{aligned} &U_t^{b,\sigma}(X_t, v, v) \geq U_t^{b,\sigma}(X_t, v, z) && \text{for all v, z} \in \mathbb{R}^+ \quad \text{(IC)}; \\ &U_t^{b,\sigma}(X_t, v, v) \geq V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) && \text{for all v} \in \mathbb{R}^+ \quad \text{(IR)}, \end{aligned}$$

where incentive compatibility (IC) captures deviations to other messages that are used by other types and individual rationality (IR) captures deviations to the message $a_\emptyset$. If the buyer chooses $a_\emptyset$, there will be no transfers in period t and the state of the system in period $t+1$ will be $X_t \setminus \{t\}$. The seller's value-to-go function can thus be simplified to

$$\begin{aligned} V_t^{s,\sigma}(X_t) = \sup_{q_{t,t}^\sigma(\cdot),\ldots,q_{t,T}^\sigma(\cdot),p_t^\sigma(\cdot)} \; & E_{v_t}\left[p_t^\sigma(X_t, v_t) + V_{t+1}^{s,\sigma}(\hat{X}_{t+1}^\sigma(X_t, v_t))\right] && (4) \\ s.t. \quad & U_t^{b,\sigma}(X_t, v, v) \geq U_t^{b,\sigma}(X_t, v, z) && \forall v, z \in \mathbb{R}^+ \\ & U_t^{b,\sigma}(X_t, v, v) \geq V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) && \forall v \in \mathbb{R}^+ \\ & q_{t,r}^\sigma(X_t, v) \in \{0, 1\} && \forall r \in X_t, v \in \mathbb{R}^+ \\ & q_{t,r}^\sigma(X_t, v) = 0 && \forall r \notin X_t, v \in \mathbb{R}^+ \end{aligned}$$

Analyzing the seller's problem. We now proceed to analyze the seller's value-to-go problem, as given in (4).

In order to bring the structure of the problem closer to a standard mechanism design formulation, we create the notion of a fake payment $\tilde{p}_t^\sigma(X_t, v)$. The fake payment $\tilde{p}_t^\sigma(X_t, v)$ will include not only the true payment $p_t^\sigma(X_t, v)$, but also all other terms in the buyer utility that do not depend on the true value v_t. That is,

$$\tilde{p}_t^\sigma(X_t, z) = p_t^\sigma(X_t, z) - \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, z) - V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, z)\right).$$

Replacing these right-hand side payments into the buyer's value-to-go (Eq. (3)), we get:

$$U_t^{b,\sigma}(X_t, v, z) = v \cdot q_{t,t}^\sigma(X_t, z) - \tilde{p}_t^\sigma(X_t, z).$$

The IC constraints from problem (4) can therefore be rewritten as

$$v \cdot q_{t,t}^\sigma(X_t, v) - \tilde{p}_t^\sigma(X_t, v) \geq v \cdot q_{t,t}^\sigma(X_t, z) - \tilde{p}_t^\sigma(X_t, z) \qquad \forall v, z \in \mathbb{R}^+.$$

The equation above is a very standard incentive compatibility constraint. We can therefore use the standard technique from Myerson (1981) and replace it with a monotonicity constraint on the allocation $q_{t,t}^\sigma(X_t, \cdot)$ and obtain the standard characterization of the buyer's utility as a function of the integral of the allocation:

$$U_t^{b,\sigma}(X_t, v, v) = U_t^{b,\sigma}(X_t, 0, 0) + \int_0^v q_{t,t}^\sigma(X_t, x) dx. \tag{5}$$

Combining the equation above with Eq. (3), we obtain:

$$\begin{aligned} U_t^{b,\sigma}(X_t, 0, 0) + \int_0^v q_{t,t}^\sigma(X_t, x) dx = v \cdot q_{t,t}^\sigma(X_t, v) + \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v) \\ - p_t^\sigma(X_t, v) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v)\right). \end{aligned} \tag{6}$$

We now simplify problem (4) by replacing the IC constraints by a monotonicity constraint on $q_{t,t}^\sigma(X_t, \cdot)$, using Eq. (5) to simplify the IR constraints and using Eq. (6) to remove the payment $p_t^\sigma(X_t, v_t)$ from the objective function.

$$
\begin{aligned}
&V_t^{s,\sigma}(X_t) = \\
&\sup_{q_{t,t}^\sigma(\cdot),\ldots,q_{t,T}^\sigma(\cdot)} E_{v_t}\Bigg[v_t \cdot q_{t,t}^\sigma(X_t, v_t) - \int_0^{v_t} q_{t,t}^\sigma(X_t, x)dx - U_t^{b,\sigma}(X_t, 0, 0) \\
&\qquad\qquad + \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v_t) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right)\Bigg] \\
&\qquad s.t. \quad q_{t,t}^\sigma(X_t, \cdot) \text{ is a nondecreasing function} \\
&\qquad\qquad U_t^{b,\sigma}(X_t, 0, 0) + \int_0^v q_{t,t}^\sigma(X_t, x)dx \geq V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) \qquad \forall v \in \mathbb{R}^+ \\
&\qquad\qquad q_{t,r}^\sigma(X_t, v) \in \{0, 1\} \qquad \forall r \in X_t, v \in \mathbb{R}^+ \\
&\qquad\qquad q_{t,r}^\sigma(X_t, v) = 0 \qquad \forall r \notin X_t, v \in \mathbb{R}^+
\end{aligned}
$$

Tthe tightest IR constraint is the one for $v = 0$ since the allocation is always a non-negative quantity. Therefore, $U_t^{b,\sigma}(X_t, 0, 0) = V_{t+1}^{b,\sigma}(X_t \setminus \{t\})$, leading to the following re-casting of problem (4):

$$
\begin{aligned}
&V_t^{s,\sigma}(X_t) = \\
&\sup_{q_{t,t}^\sigma(\cdot),\ldots,q_{t,T}^\sigma(\cdot)} E_{v_t}\Bigg[v_t \cdot q_{t,t}^\sigma(X_t, v_t) - \int_0^{v_t} q_{t,t}^\sigma(X_t, x)dx - V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) \\
&\qquad\qquad + \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v_t) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right)\Bigg] \\
&\qquad s.t. \quad q_{t,t}^\sigma(X_t, \cdot) \text{ is a nondecreasing function} \\
&\qquad\qquad q_{t,r}^\sigma(X_t, v) \in \{0, 1\} \qquad \forall r \in X_t, v \in \mathbb{R}^+ \\
&\qquad\qquad q_{t,r}^\sigma(X_t, v) = 0 \qquad \forall r \notin X_t, v \in \mathbb{R}^+
\end{aligned}
\tag{7}
$$

The problem above can be decomposed into two independent optimization problems, a present problem $W_t^{P,\sigma}(X_t)$ and a future problem $W_t^{F,\sigma}(X_t)$, minus a constant term $V_{t+1}^{b,\sigma}(X_t \setminus \{t\})$. This occurs because the first two terms in the objective depend only on the present item, the third term is a constant that does not depend on any allocation decisions and the last three terms depend only on future items. Furthermore, there are no constraints connecting the inter-period allocation decisions. In particular,

$$
V_t^{s,\sigma}(X_t) = W_t^{P,\sigma}(X_t) + W_t^{F,\sigma}(X_t) - V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) \tag{8}
$$

where

$$W_t^{P,\sigma}(X_t) = \sup_{q_{t,t}^\sigma(\cdot)} E_{v_t}\left[v_t \cdot q_{t,t}^\sigma(X_t, v_t) - \int_0^{v_t} q_{t,t}^\sigma(X_t, x)dx\right]$$

$$\begin{aligned} s.t. \quad & q_{t,t}^\sigma(X_t, \cdot) \text{ is a nondecreasing function} & \\ & q_{t,t}^\sigma(X_t, v) \in \{0,1\} & \text{if } t \in X_t, v \in \mathbb{R}^+ \\ & q_{t,t}^\sigma(X_t, v) = 0 & \text{if } t \notin X_t, v \in \mathbb{R}^+ \end{aligned}$$

and

$$W_t^{F,\sigma}(X_t) =$$

$$\sup_{q_{t,t+1}^\sigma(\cdot),\ldots,q_{t,T}^\sigma(\cdot)} E_{v_t}\left[\mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v_t) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right)\right]$$

$$\begin{aligned} s.t. \quad & q_{t,r}^\sigma(X_t, v) \in \{0,1\} & \forall r \in X_t \setminus \{t\}, v \in \mathbb{R}^+ \\ & q_{t,r}^\sigma(X_t, v) = 0 & \forall r \in X_t \setminus \{t\}, v \in \mathbb{R}^+. \end{aligned}$$

We can solve both of these optimization problems. We first consider the present problem. When $t \notin X_t$, the only feasible answer is $q_{t,t}^\sigma(X_t, v) = 0$ for all v. When $t \in X_t$, we can consider the relaxation where $q_{t,t}^\sigma(X_t, v) \in [0,1]$ for all v. This problem is identical to the single-shot Myerson (1981) problem with a single buyer. Since we assumed virtual values are non-decreasing, the optimal answer is to allocate item t to the buyer if, and only if, $v_t \geq p^*$. Note that this solution is integral, and thus feasible for the present problem $W_t^{P,\sigma}(X_t)$. Therefore,

$$W_t^{P,\sigma}(X_t) = p^*(1 - F(p^*))I\{t \in X_t\}, \tag{9}$$

where $I\{\cdot\}$ denotes an indicator function.

We now turn to the future problem $W_t^{F,\sigma}(X_t)$, which we can solve by using an almost trivial bound. The objective function of $W_t^{F,\sigma}(X_t)$ depends only on the value of the allocation of future items. Therefore, it cannot exceed the expected buyer's value of all future items $\mu|X_t \setminus \{t\}|$, yielding the following bound:

$$\mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v)\right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v)\right) \leq \mu\ |X_t \setminus \{t\}|.$$

This bound is achievable by assigning all items to the buyer right away: $q_{t,r}^\sigma(X_t, v) = 1$ for all v and all $r \in X_t \setminus \{t\}$. Thus,

$$W_t^{F,\sigma}(X_t) = \mu|X_t \setminus \{t\}|. \tag{10}$$

We have now determined that there exist PBE where the present item is allocated if $v_t \geq p^*$ and all future items are allocated to the buyer. This PBE is not unique, as we discuss later. However, by combining Eqs. (8), (9) and (10), we have established that for all PBE, the seller's value-to-go satisfies for all t and all $X_t \in \mathcal{X}_t$:

$$V_t^{s,\sigma}(X_t) = p^*(1 - F(p^*))I\{t \in X_t\} + \mu|X_t \setminus \{t\}| - V_{t+1}^{b,\sigma}(X_t \setminus \{t\}). \tag{11}$$

Computing the value-to-go functions. We now use Eq. (11) to recursively determine the buyer's and the seller's value-to-go functions.

Recall that Z^b denotes the buyer's ex ante expected utility under the optimal price in a single-shot game (see Eq. (2)). We now prove by backward induction that $V_t^{b,\sigma}(X_t) = Z^b|X_t|$ for any t, X_t and any PBE σ. Let's consider first period T. If $X_T = \emptyset$, then $V_T^{b,\sigma}(X_T) = 0$. If $X_T = \{T\}$, then we are faced with a standard Myerson problem in period T. The buyer's expected value-to-go is therefore his expected his expected utility under Myerson: $V_T^{b,\sigma}(X_T) = Z^b$. The base case is thus proved.

We now assume the inductive hypothesis is true for period $t+1$ and will prove it implies the hypothesis is also true for period t. The inductive hypothesis says that $V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) = Z^b|X_t \setminus \{t\}|$. Plugging this into Eq. (5), we obtain

$$V_t^{s,\sigma}(X_t) = p^*(1 - F(p^*))I\{t \in X_t\} + \mu|X_t \setminus \{t\}| - Z^b|X_t \setminus \{t\}|. \tag{12}$$

We also know that any PBE must assign the present item according to Myerson and all future items are allocated to the buyer. Therefore, the sum of the seller's and the buyer's value-to-go functions must satisfy

$$V_t^{s,\sigma}(X_t) + V_t^{b,\sigma}(X_t) = (p^*(1 - F(p^*) + Z^b) \cdot I\{t \in X_t\} + \mu|X_t \setminus \{t\}|. \tag{13}$$

Subtracting Eq. (12) from Eq. (13), we obtain

$$V_t^{b,\sigma}(X_t) = Z^b \cdot I\{t \in X_t\} + Z^b|X_t \setminus \{t\}| = Z^b|X_t|,$$

completing our proof.

Plugging the buyer's value-to-go function into the seller's value-to-go function from Eq. (11), we obtain that

$$V_t^{s,\sigma}(X_t) = p^*(1 - F(p^*))I\{t \in X_t\} + (\mu - Z^b)|X_t \setminus \{t\}|.$$

Therefore, the buyer's starting value-to-go is therefore $V_1^{b,\sigma}(X_1) = Z^b T$ and the seller's starting value-to-go is $V_1^{s,\sigma}(X_1) = p^*(1 - F(p^*)) + (\mu - Z^b)(T - 1)$. Dividing both starting value-to-go functions by T and taking the limit as T tends to infinity, we obtain Theorem 1.

4 The Equilibrium Outcome

The prior section establishes the key properties of any PBE of our dynamic game. A quantity that emerges as particularly important is what we call the positive commitment price.

Definition 2. *We call $\hat{p} = \mu - Z^b$ the positive commitment price.*

Under any PBE, all items except for the period 1 good are allocated to the buyer at a price of $\hat{p}$. The period 1 item is offered to the buyer at the static optimal price of p^*. With the exception of the period 1 good, the outcome is socially efficient.

Multiplicity of equilibria. Our game has a multiplicity of PBE. There exists a PBE, which we will label the upfront PBE, where all goods $t = 2, 3, ..., T$ are allocated at period 1. There exists another equilibrium where the $t + 1$ good is always allocated at time t (except for the period 1 good). We call this second equilibrium the day-before PBE.

These equilibria are not identical. In the upfront PBE, a large payment is made upfront in return for $T - 1$ items. In the day-before PBE, the seller only charges the buyer for one good in advance. It, therefore, does not rely on large payments. Regardless, all PBE lead to the same ex post allocation and the same sum of payments from the buyer to the seller.

Recovering the no-trade inefficiency. We now discuss the main benefit of using the positive commitment solution over the no commitment solution. For illustrative purposes, we do so while considering our running example of uniform valuations between 0 and 1. Figure 1 shows the different surplus regions under the monopoly price $p^* = 1/2$. The seller's surplus is $p^*(1 - F(p^*)) = 1/4$. The buyer's surplus is $Z^b = 1/8$. There is also a no-trade inefficiency which deducts 1/8 of total surplus from the expected buyer value of $\mu = 1/2$.

Under a PBE of the positive commitment game, the outcome associated with items $t > 1$ is efficient. The seller is able to recover the no-trade inefficiency and add it to her surplus, now earning $\hat{p} = \mu - Z^b = 1/2 - 1/8 = 3/8$. The buyer obtains the same surplus as in the static game: $Z^b = 1/8$ per item.

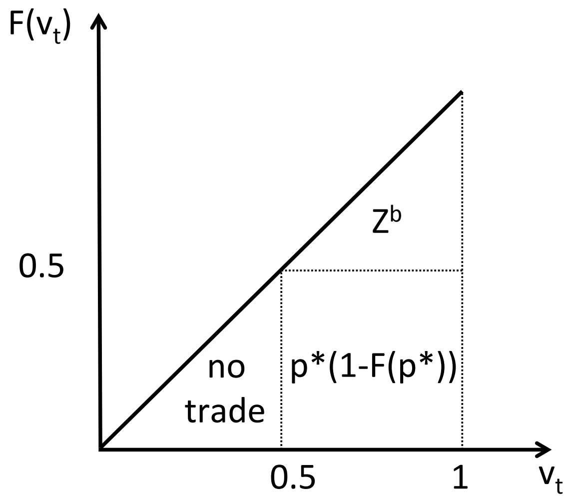

Figure 1: The surplus regions under the optimal static price $p^* = 1/2$ assuming uniform buyer valuations.

The long-term buyer surplus under positive commitment is the same as under no commitment. The long-term welfare is the same as under full commitment, since positive commitment allows us to recover the no-trade inefficiency.

Closer to no commitment or to full commitment? One natural question is whether the long-term average revenue under positive commitment price $\hat{p}$ is "closer" to the no commitment revenue of $p^*(1-F(p^*))$ or to the full commitment price μ. That is, in instances where $p^*(1-F(p^*))$ and μ are very different from each other, which one does $\hat{p}$ approximate?

The equal revenue distribution would seem like a suitable problem for studying this question since full commitment is extremely valuable under that distribution. However, the equal revenue distribution does not satisfy the unique optimal price p^* assumption. Therefore, we consider two variants of the equal revenue distribution.

Let L be a large positive constant. In scenario 1, the buyer's valuations are drawn from the distribution with density $f_1(x) = \frac{L}{L-1} \cdot \frac{1}{x^2}$ with $x \in [1, L]$. In scenario 2, the buyer's valuations are drawn from a distribution with density $f_2(x) = \frac{L+1}{L} \cdot \frac{1}{(x+1)^2}$ with $x \in [0, L]$.

In both scenarios, the buyer's surplus under no commitment $p^*(1 - F(p^*))$ is approximately 1 and the the buyer's surplus under full commitment μ is approximately $\log L$. That is, from a static mechanism design point-of-view, these two distributions lead to fairly similar revenue outcomes.

However, from a positive commitment point-of-view, these two distributions lead to very different outcomes. Under the first scenario, we have $\hat{p} = 1$. Under the second scenario, we have $\hat{p} \approx \frac{1}{2} \log L$. That is, under scenario 1, the positive commitment price generates the same revenue as the no commitment solution and, under scenario 2, the positive commitment price generates revenues that are close to the full commitment solution.

The difference is due to the different optimal static prices under the two scenarios. In scenario 1, we have $p^* = 1$, while $p^* \approx \sqrt{L}$ in scenario 2. In scenario 1, the static pricing solution is fully efficient due to the low monopoly price. In scenario 2, the solution is very inefficient due to the high monopoly price. The positive commitment price helps the seller by eliminating the "no trade" inefficiency. Therefore, positive commitment power is very useful to the seller in the inefficient scenario 2, but useless to the seller in the already efficient scenario 1.

5 Extensions

There are many potential avenues for extending the results in this paper, some of them straightforward and others less so.

Irregular distributions and the multi-item case. When introducing the model, we assumed the buyer's value distribution was regular. That is, we assumed the virtual values were non-decreasing. We did so in order to simplify our presentation by having the optimal single-shot mechanism to be a price p^*. This is not a critical assumption of the paper. Using ironing, we can straightforwardly extend our model to irregular distributions.

In fact, we could allow for much more complicated single-period models. Suppose the seller had multiple distinct items to sell in each period. In this, the period 1 items would be sold according to a much more complex auction than simply setting a price p^* (see Cai et al. (2013)). Future allocations would still be sold in advance at a price equal to their expected values minus the buyer's ex ante expected utility under a single-shot mechanism.

Our model could be easily adapted to handle complex valuations in individual periods as long as the valuation of the buyer is additive across periods. An interesting directions, that would require the development of new techniques, is to study more complex cross-period valuations, such as submodular and subadditive valuations.

The multi-buyer case. Our model assumes a single buyer and an important extension is to analyze the multi-buyer case. With multiple buyers, our definition of positive commitment needs to be refined. In a setting with many buyers, the seller would not want to pre-allocate items to individual buyer. The seller would want to wait until valuations are realized to decide which buyer gets a particular item. One possible relaxation of our positive commitment framework is to allow the seller to commit to minimum allocation probabilities in advance. A key challenge here would be to make sure the seller does not cherry pick low valuation items for a buyer who signed an allocation probability contract. Another possibility is to allow the seller to promise the buyers properties of the auction, such as which reserves prices each buyer will face.

We conjecture that our single-agent results can be naturally extended to a multi-buyer case. In this case, the seller would offer reserve price reduction contracts to buyers. The seller would then sell the items using an auction with lazy personalized prices (see Dhangwatnotai et al. (2015) or Paes Leme et al. (2016) for a definition of lazy personalized prices). Buyers without a contract would be faced with Myersonian reserve prices. Buyers under contract would be assigned a reserve price of zero. In equilibrium, future items would be allocated according to the VCG auction.

Persistent information. Our model also assumed the buyer's valuations are drawn i.i.d. from some distribution $F(\cdot)$. The identically distributed assumption is not key to our results and can easily be relaxed. The independence assumption, however, would be much more difficult to relax as it would lead to persistent private information. We would need to develop new tools to solve our model in the presence of persistent private information.

References

Ashlagi, Itai, Constantinos Daskalakis, Nima Haghpanah. 2016. Sequential mechanisms with ex-post participation guarantees. *Proceedings of EC Conference*.

Balseiro, Santiago, Vahab Mirrokni, Renato Paes Leme. 2016. Dynamic mechanisms with martingale utilities. Working paper, Duke University.

Cai, Yang, Constantinos Daskalakis, S Matthew Weinberg. 2013. Understanding incentives: Mechanism design becomes algorithm design. *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 618–627.

Deb, Rahul, Maher Said. 2015. Dynamic screening with limited commitment. *Journal of Economic Theory* **159** 891–928.

Dhangwatnotai, Peerapong, Tim Roughgarden, Qiqi Yan. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* **91** 318–333.

Kakade, Sham M, Ilan Lobel, Hamid Nazerzadeh. 2013. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research* **61**(4) 837–854.

Mirrokni, Vahab, Renato Paes Leme, Pingzhong Tang, Song Zuo. 2016a. Optimal dynamic mechanisms with ex-post ir via bank accounts. Working paper, Google Research.

Mirrokni, Vahab, Hamid Nazerzadeh. 2017. Deals or no deals: Contract design for online advertising. *Proceedings of WWW Conference*.

Mirrokni, Vahab, Renato Paes Leme, Pingzhong Tang, Song Zuo. 2016b. Oblivious dynamic mechanism design. Working paper, Google Research.

Myerson, Roger B. 1981. Optimal auction design. *Mathematics of Operations Research* **6**(1) 58–73.

Myerson, Roger B. 1986. Multistage games with communication. *Econometrica: Journal of the Econometric Society* 323–358.

Paes Leme, Renato, Martin Pál, Sergei Vassilvitskii. 2016. A field guide to personalized reserve prices. *Proceedings of WWW Conference*.

Pavan, Alessandro, Ilya Segal, Juuso Toikka. 2014. Dynamic mechanism design: A myersonian approach. *Econometrica* **82**(2) 601–653.

Skreta, Vasiliki. 2006. Sequentially optimal mechanisms. *The Review of Economic Studies* **73**(4) 1085–1111.

Skreta, Vasiliki. 2015. Optimal auction design under non-commitment. *Journal of Economic Theory* **159** 854–890.