Dynamic Contracting under Positive Commitment

Ilan Lobel,[1] **Renato Paes Leme**[2]

[1]NYU Stern, [2]Google Research

ilobel@stern.nyu.edu, renatoppl@google.com

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, and characterize the perfect Bayesian equilibrium of this dynamic game. We prove the equilibrium is efficient and that the seller's revenue is a function of the buyer's ex ante utility under a no commitment model. In particular, all goods are sold in advance to the buyer at what we call the *positive commitment price*.

Introduction

Motivation. 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 valuation 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? Should the seller sell the products in advance? If so, for how much?

Solving this problem is crucial to the design of certain marketplaces such as the large and fast growing 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 bidding, where individual impressions are sold milliseconds before they are delivered. Online ads are also sold via contracts, where a seller agrees in advance to deliver a specific number of impressions to a buyer over a given time horizon for a pre-specified price. In our work, we aim to understand if there is a fundamental reason behind this kind of fragmented market design, and if so, what should be the relationship between prices in the real-time market and the pre-sale contract market.

Selling with or without commitment power. Consider the seller's problem. The seller 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 prices, where prices are chosen according to Myerson (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 seller 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 real-world 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, Lobel, and Nazerzadeh 2013) or (Pavan, Segal, and Toikka 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. This offer is made to the buyer before he learns his valuation for any of the items. In equilibrium, the buyer accepts the offer, leading to a 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 pa-

pers have proposed ways to incorporate bundling ideas into market design. This includes work on dynamic bundling by (Ashlagi, Daskalakis, and Haghpanah 2016), (Balseiro, Mirrokni, and Paes Leme 2018), (Mirrokni et al. 2016a) and (Mirrokni et al. 2016b) as well as work on preferred deals by (Mirrokni and Nazerzadeh 2017).

The case for partial commitment power. 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.

Our proposed framework. 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 positive commitment price is the highest price the seller can charge the buyer for an impression in the contract channel if both players take into account the fact that impressions that are not pre-sold via contracts are eventually sold in the real-time exchange.

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 per-item revenue are $3/8$.

In equilibrium, the buyer purchases all the items in the stream of products at the positive commitment price. 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. 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 how the surplus gets split between the seller and the buyer.

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

Skreta (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. Skreta (Skreta 2015) is a follow-up paper that considers a multi-buyer problem and finds the sequence of optimal auctions for selling a single item.

Skreta's result builds on the seminal idea of the Coase conjecture ((Coase 1972), (Gul, Sonnenschein, and Wilson 1986)). Coase conjectured that a seller with a single durable good and no commitment power would have to sell at or near marginal cost, as the seller would in effect compete with a more informed future version of herself. In our work, a similar effect appears where the seller competes with a future version of herself. However, this dynamic competition between the seller and her future self are quite different in our model and in Coase's. In our model, in contrast to Coase's, there is no persistent private information as our goods are non-durable. Therefore, the future version of the seller does not possess any information the current seller does not. The mechanism that Coase elucidated is thus very different from the one that takes place in our model. What happens in our setting is that the future seller is in a weaker bargaining position vis-a-vis the buyer than the present seller. That is, the seller must compromise today in order to avoid having to negotiate with the buyer after he learns his private value. There is also a body of literature on contract renegotiation that builds on the Coase conjecture framework. The same distinction to our work applies. In (Hart and Tirole 1988), the seller is constrained by the fact that her future self will be more knowledgeable about the buyer than her present self. Meanwhile, our model has no persistent private information and the tension emerges from the buyer expecting to gain private information in the future.

The formulation of our problem as an optimization program resembles a Markov Decision Process (MDP) and therefore our analysis shares similarities with the papers on dynamic mechanism design for MDPs, e.g. (Parkes, Yanovsky, and Singh 2005; Mierendorff 2016; Parkes and Singh 2004). Those papers, however, all assume commitment power from the seller, even though decisions are online in the sense that they use only present and past data to make decisions.

Deb and Said (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 the allocation rule in period 1, but she cannot commit to the mechanism being offered in period 2. As in our work, buyers can choose to wait rather than contract immediately upon arrival. 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.

Implications for contracting under symmetric information. A classic result in contract theory is that first-best can be achieved in contracting problems without information asymmetry via a "sell the firm" contract ((Hölmstrom 1979)). That is, the seller's optimal strategy is to sell her entire output in advance to the buyer, in exchange for an upfront payment equal to the expected value of the output. In the supply chain literature, this kind of contract is called a coordinating contract ((Cachon and Lariviere 2005)). Our paper points out that this classic result relies on the buyer not expecting to gain an informational advantage in the future. If the buyer expects to gain private information as time elapses, he could use this fact to gain some bargaining power. In this case, the seller would not be able to charge more than the positive commitment price of the output.

Full version. A complete version of our paper containing the missing proofs as well as a longer discussion of our results and extensions is available online at (Lobel and Paes Leme 2018).

The Model

We consider a model with one seller and one buyer who interact over $T+1$ 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. We also allow interaction in period $t = 0$, before any item arrives. 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 $f(\cdot)$, and 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)), an assumption which is satisfied by many of the commonly used distributions, including the normal, the lognormal, the uniform and the exponential distributions. We denote by p^* the seller's optimal monopoly 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. 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 in-

---

[1]In the full version we note that we can extend our results to irregular distributions and even to multi-item problems by borrowing techniques from the literature. We opt however to present our results assuming regularity in order to simplify the notation and presentation.

clude discounting in our model, though discounting would not significantly alter our results.

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 optimal (static) monopoly price p^* in each period. In this case, the seller's expected per-item revenue is $p^* \cdot (1 - F(p^*))$. Under full commitment, the seller's optimal strategy would be to offer at period $t = 0$ a take-it-or-leave-it offer at price $E[\sum_{t=1}^{T} v_t] = \mu \cdot T$ — see, for example, (Kakade, Lobel, and Nazerzadeh 2013) or (Pavan, Segal, and Toikka 2014). In equilibrium, the buyer would accept the offer, leading to a per-item revenue of μ. The seller's revenue is significantly greater under full commitment than under no commitment, as she is able to extract the entire welfare of the system under full commitment.

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 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 items for the price of $\mu \cdot T$. The leave-it part of the offer is the crucial bit that would be hard to enforce. Suppose the buyer declines 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 future after the first offer is rejected.

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 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 any of the standard versions of the revelation principle (see (Myerson 1981; 1986)) to simplify our contracting 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 $X_0 = \{1, 2, ..., t\}$ 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" message, that we represent by $a_\emptyset \in A_t$. By formally including a "no trade" message $a_\emptyset$ in our model, we can analyze what happens after the buyer declines an offer from the seller.

The allocation rule $y_t : A_t \to 2^{\{t, t+1, ..., T\}}$ determines which goods 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 payments 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:

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

The buyer and the seller's per-item utilities are equal to U^b/T and U^s/T, respectively.

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 $\sigma^s = \{\sigma_t^s\}_{t=0,...,T}$ 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 $\sigma^b = \{\sigma_t^b\}_{t=0,...,T}$ 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., $\sigma_t^b :$

$\mathcal{X}_t \times \mathcal{M}_t \times \mathbb{R}^+ \to A_t$.[2] We refer to a pair $\sigma = (\sigma^s, \sigma^b)$ as a strategy profile.

Definition 1. *A strategy profile σ constitutes a 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 now introduce some additional notation that we will use to clarify the notion of a PBE in our game. Given a mechanism M_t chosen by the seller at time t, we define $\tilde{U}_t^{b,\sigma}(X_t, M_t, v, a)$ as the value-to-go of the buyer with valuation $v_t = v$ given state X_t, assuming he plays action $a \in A_t$ at time t and both agents play according to σ afterwards:

$$\tilde{U}_t^{b,\sigma}(X_t, M_t, v_t, a) = v_t \cdot |y_t(a) \cap \{t\}| + \mu \cdot |y_t(a) \setminus \{t\}| - r_t(a) + V_{t+1}^{b,\sigma}(X_{t+1}), \quad (2)$$

where X_{t+1} is defined in Eq. (1) and $V_t^{b,\sigma}(X_t)$ is the buyer's equilibrium value-to-go at time t given state X_t, before knowing his valuation v_t or the seller's choice of mechanism M_t:

$$V_t^{b,\sigma}(X_t) = E\left[\tilde{U}_t^{b,\sigma}(X_t, \sigma_t^s(X_t), v_t, \sigma_t^b(X_t, \sigma_t^s(X_t), v_t))\right]. \quad (3)$$

For $t > T$, $V_t^{b,\sigma}(\cdot) = 0$. Eq. (2) states that the buyer's value-to-go $\tilde{U}_t^{b,\sigma}(X_t, M_t, v_t, a)$ is given by the addition of four terms: (i) the utility from the allocation of item t at period t, (ii) the expected utility from the allocation of items $t' > t$, (iii) the negative of the payment to the seller at time t, and (iv) the expected value-to-go from future interactions after time t.

Given a mechanism M_t and state X_t, the buyer has no incentive to deviate from strategy σ^b at time t if

$$\tilde{U}_t^{b,\sigma}\left(X_t, M_t, v, \sigma_t^b(X_t, M_t, v)\right) \geq \tilde{U}_t^{b,\sigma}\left(X_t, M_t, v, a\right)$$
$$\text{for all } v \in \mathbb{R}^+, a \in A_t. \quad (BIC_t)$$

Similarly, we now introduce the seller's value-to-go at time t given state X_t under equilibrium σ, $V_t^{s,\sigma}(X_t)$. The seller has no incentive to deviate in period t if his value-to-go is maximized by his choice of mechanism:

$$V_t^{s,\sigma}(X_t) = \sup_{M_t \in \mathcal{M}_t} E\left[r_t(\sigma_t^b(X_t, M_t, v_t)) + V_{t+1}^{s,\sigma}(X_{t+1})\right] \quad (4)$$
$$\text{s.t. } (BIC_t).$$

As for the buyer, $V_t^{s,\sigma}(\cdot) = 0$ if $t > T$. The feasibility constraints of the mechanism are implicitly included above through the constraint $M_t \in \mathcal{M}_t$. If Eq. (4) is satisfied for all $t = 0, ..., T$ and $X_t \in \mathcal{X}_t$, then σ is a PBE.

---

[2]The buyer's strategy in period $t = 0$ should not include a valuation v_0 since there is no item available at time zero. However, in order to make the notation of the paper consistent over the different periods, we add a fake valuation v_0 for the non-existent period zero item. Our results would not change if we formalized period zero separately without adding v_0.

Dynamic Contracting

We now use dynamic mechanism design to study the PBE of the problem defined in the prior section. The letter Z will denote the buyer's ex ante expected utility under the optimal price in a single-shot problem, i.e.,

$$Z = \int_{p^*}^{\infty} (x - p^*) dF(x). \quad (5)$$

A quantity that emerges as particularly important is the positive commitment price.

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

We now state the main result of our paper.

Theorem 1. *Under any PBE, the seller's utility per-item is $\hat{p}$ and the buyer's per-item utility is Z. Furthermore, for any $t = 1, ..., T$, item t is allocated to the buyer before period T.*

In the remainder of this section, we will provide the analysis that proves the theorem above. Our first step is to obtain a "revelation principle"-type result for our problem.

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 us 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, the actions being removed are never used in any subgame: we can remove unused actions from the action set via backwards induction, starting from period T and going

back all the way to period 0. Therefore, the new mechanism will generate the exact same value-to-go for both the seller and the buyer.

In sum, without loss of generality, we can consider only "direct mechanisms" where the buyer either reports his type or declines to trade, i.e., $A_t = \mathbb{R}^+ \cup \{a_\emptyset\}$ for all $t = 0, ..., T$. We use "direct mechanism" in quotes because of the addition of an extra message in the game beyond the classical reporting of types, the message $a_\emptyset$.

Simplifying the formulation. We now use the revelation principle argument above to simplify our problem formulation. 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 $q_{t,r}^\sigma(X_t, v)$ represent whether item r is allocated at period t to the buyer given state X_t and $v_t = v$ under the equilibrium σ. The allocation $q_{t,r}^\sigma(X_t, v)$ will be subject to the natural feasibility constraints (F_t) that only items still available in period t can be allocated:

$$q_{t,r}^\sigma(X_t, v) \in \{0, 1\} \quad \text{for all } r \in X_t, v \in \mathbb{R}^+;$$
$$q_{t,r}^\sigma(X_t, v) = 0 \quad \text{for all } r \notin X_t, v \in \mathbb{R}^+. \quad (F_t)$$

Similarly, for any t, we let $p_t^\sigma(X_t, v)$ 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 the state transition function $\hat{X}_{t+1}^\sigma(X_t, v)$ as the period $t+1$ state given the period t state X_t and the buyer's value $v_t = v$, under the equilibrium σ.

Using the revelation principle argument above, we can simplify the notation for the buyer's value-to-go. We let $U_t^{b,\sigma}(X_t, M_t, v, z)$ be the value-to-go of the buyer at time t given state X_t and mechanism M_t assuming his true value is $v_t = v$ but that he acts as if his true value were z — that is, plays action $\sigma_t^b(X_t, M_t, z)$ — assuming that all future actions of both the seller and the buyer will conform to the equilibrium σ. That is,

$$U_t^{b,\sigma}(X_t, M_t, v, z) = v \cdot q_{t,t}^\sigma(X_t, z) + \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, z)$$
$$- p_t^\sigma(X_t, z) + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, z)\right), \quad (6)$$

where the value-to-go $V_t^{b,\sigma}(X_t)$ can be simplified from Eq. (3) to

$$V_t^{b,\sigma}(X_t) = E_{v_t}[U_t^{b,\sigma}(X_t, \sigma_t^s(X_t), v_t, v_t)].$$

We can therefore simplify the original incentive constraints and replace them with simpler constraints, which we call (IC_t) and (IR_t) respectively:

$$U_t^{b,\sigma}(X_t, M_t, v, v) \geq U_t^{b,\sigma}(X_t, M_t, v, z) \quad \text{for all v, z} \in \mathbb{R}^+$$
$$U_t^{b,\sigma}(X_t, M_t, v, v) \geq V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) \quad \text{for all v} \in \mathbb{R}^+,$$

where incentive compatibility (IC_t) captures deviations to other messages that are used by other types and individual rationality (IR_t) 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 period t problem can thus be simplified from Eq. (4) to:

$$V_t^{s,\sigma}(X_t) = \sup E_{v_t}\left[p_t^\sigma(X_t, v_t) + V_{t+1}^{s,\sigma}(\hat{X}_{t+1}^\sigma(X_t, v_t))\right] \quad (7)$$
$$s.t. \quad (F_t), (IC_t), (IR_t).$$

Analyzing the seller's problem. We now proceed to analyze the seller's period t problem, as given in Eq. (7). 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. (6)), we get:

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

The (IC_t) constraints from problem (7) can therefore be rewritten as

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

The equation above is a very standard incentive compatibility constraint. We can therefore use the standard technique from Myerson (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:

$$U_t^{b,\sigma}(X_t, M_t, v, v) = U_t^{b,\sigma}(X_t, M_t, 0, 0) + \int_0^v q_{t,t}^\sigma(X_t, x)dx. \quad (8)$$

Applying the standard Myersonian transformation, we can eliminate the (IC_t) and (IR_t) constraints obtaining the following simplified problem (the details are omitted due to space constraints and the fact that the transformation is standard):

$$V_t^{s,\sigma}(X_t) = \sup E_{v_t}\left[v_t \cdot q_{t,t}^\sigma(X_t, v_t) - \quad (9)\right.$$
$$\int_0^{v_t} q_{t,t}^\sigma(X_t, x)dx -$$
$$V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) + \mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v_t) +$$
$$\left. V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t)\right)\right]$$
$$s.t. \quad (F_t), \quad q_{t,t}^\sigma(X_t, \cdot) \text{ is a nondecreasing function.}$$

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,

$$V_t^{s,\sigma}(X_t) = W_t^{P,\sigma}(X_t) + W_t^{F,\sigma}(X_t) - V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) \quad (10)$$

where

$$W_t^{P,\sigma}(X_t) = \sup_{q_{t,t}^\sigma(\cdot)} E_{v_t}\left[v_t \cdot q_{t,t}^\sigma(X_t, v_t) \right. \tag{11}$$

$$\left. - \int_0^{v_t} q_{t,t}^\sigma(X_t, x)dx \right]$$

$$s.t. \quad q_{t,t}^\sigma(X_t, \cdot) \text{ is a nondecreasing function,}$$

$$q_{t,t}^\sigma(X_t, v) \in \{0,1\} \quad \text{if } t \in X_t, v \in \mathbb{R}^+,$$

$$q_{t,t}^\sigma(X_t, v) = 0 \quad \text{if } t \notin X_t, v \in \mathbb{R}^+.$$

and

$$W_t^{F,\sigma}(X_t) = \sup E_{v_t}\left[\mu \sum_{r \in X_t \setminus \{t\}} q_{t,r}^\sigma(X_t, v_t) \right.$$

$$\left. + V_{t+1}^{b,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t) \right) + V_{t+1}^{s,\sigma}\left(\hat{X}_{t+1}^\sigma(X_t, v_t) \right) \right]$$

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

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

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,

$$W_t^{P,\sigma}(X_t) = p^* \cdot (1 - F(p^*)) \cdot I\{t \in X_t\}, \tag{12}$$

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 $\mu \cdot |X_t \setminus \{t\}|$, yielding the following bound:

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

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

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,

$$W_t^{F,\sigma}(X_t) = \mu \cdot |X_t \setminus \{t\}|. \tag{13}$$

We have now determined that there exist a 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. (10), (12) and (13), we have established that for all PBE, the seller's value-to-go satisfies for all t and all $X_t \in \mathcal{X}_t$:

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

$$- V_{t+1}^{b,\sigma}(X_t \setminus \{t\}). \tag{14}$$

Computing the value-to-go functions. We now use Eq. (14) to recursively determine the buyer's and the seller's value-to-go functions.

Recall that Z denotes the buyer's ex ante expected utility under the optimal price in a single-shot game (see Eq. (5)). We now prove by backward induction that $V_t^{b,\sigma}(X_t) = Z \cdot |X_t|$ for any t, X_t and any PBE σ. Let's consider 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 utility under Myerson: $V_T^{b,\sigma}(X_T) = Z$. 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 $V_{t+1}^{b,\sigma}(X_t \setminus \{t\}) = Z \cdot |X_t \setminus \{t\}|$. Plugging this into Eq. (8), we obtain

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

$$- Z \cdot |X_t \setminus \{t\}|. \tag{15}$$

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

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

$$+ \mu \cdot |X_t \setminus \{t\}|. \tag{16}$$

Subtracting Eq. (15) from Eq. (16), we obtain

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

completing our proof.

Plugging the buyer's value-to-go function into the seller's value-to-go function from Eq. (14), we obtain

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

Therefore, the buyer's starting value-to-go is $V_0^{b,\sigma}(X_1) = Z \cdot T$ and the seller's starting value-to-go is $V_0^{s,\sigma}(X_1) = (\mu - Z) \cdot T$, proving Theorem 1.

The Equilibrium Outcome

Market design. The prior section establishes the key properties of any PBE of our dynamic game. Under any PBE, all items are allocated to the buyer at the positive commitment price $\hat{p}$ in advance of the items becoming available, and the allocation is socially efficient.

Our analysis actually specifies that the seller should have two different sales channels: one for present items and a different one for future items. The present one, if available, is always offered to the buyer at price p^*, while future items are offered at the lower price $\hat{p}$. This decomposition into two different sales channels is matched by what we observe in practice in the online advertising business, with the market being divided into an ahead-of-time contract channel and a real-time bidding channel.

Multiplicity of equilibria and upfront payments. Our game has a multiplicity of PBE. There exists a PBE, which we will label the upfront PBE, where all goods $t =$

$1, 2, ..., T$ are allocated at period zero. There exists another equilibrium where the $t + 1$ good is always allocated at time t. 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 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.

The Equilibrium Outcome

Market design. The prior section establishes the key properties of any PBE of our dynamic game. Under any PBE, all items are allocated to the buyer at the positive commitment price $\hat{p}$ in advance of the items becoming available, and the allocation is socially efficient.

Our analysis actually specifies that the seller should have two different sales channels: one for present items and a different one for future items. The present one, if available, is always offered to the buyer at price p^*, while future items are offered at the lower price $\hat{p}$. This decomposition into two different sales channels is matched by what we observe in practice in the online advertising business, with the market being divided into an ahead-of-time contract channel and a real-time bidding channel.

Multiplicity of equilibria and upfront payments. Our game has a multiplicity of PBE. There exists a PBE, which we will label the upfront PBE, where all goods $t = 1, 2, ..., T$ are allocated at period zero. There exists another equilibrium where the $t + 1$ good is always allocated at time t. 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 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.

References

Ashlagi, I.; Daskalakis, C.; and Haghpanah, N. 2016. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of EC Conference*.

Balseiro, S.; Mirrokni, V.; and Paes Leme, R. 2018. Dynamic mechanisms with martingale utilities. Management Science (forthcoming).

Cachon, G. P., and Lariviere, M. A. 2005. Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science* 51(1):30–44.

Coase, R. H. 1972. Durability and monopoly. *The Journal of Law and Economics* 15(1):143–149.

Deb, R., and Said, M. 2015. Dynamic screening with limited commitment. *Journal of Economic Theory* 159:891–928.

Gul, F.; Sonnenschein, H.; and Wilson, R. 1986. Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory* 39(1):155–190.

Hart, O. D., and Tirole, J. 1988. Contract renegotiation and Coasian dynamics. *The Review of Economic Studies* 55(4):509–540.

Hölmstrom, B. 1979. Moral hazard and observability. *The Bell journal of economics* 74–91.

Kakade, S. M.; Lobel, I.; and Nazerzadeh, H. 2013. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research* 61(4):837–854.

Lobel, I., and Paes Leme, R. 2018. Dynamic contracting under positive commitment (full version). In *Online Appendix in* https://ssrn.com/abstract=2916284.

Mierendorff, K. 2016. Optimal dynamic mechanism design with deadlines. *Journal of Economic Theory* 161:190–222.

Mirrokni, V., and Nazerzadeh, H. 2017. Deals or no deals: Contract design for online advertising. In *Proceedings of WWW Conference*.

Mirrokni, V.; Leme, R. P.; Tang, P.; and Zuo, S. 2016a. Optimal dynamic mechanisms with ex-post IR via bank accounts. Ad Auctions Workshop.

Mirrokni, V.; Paes Leme, R.; Tang, P.; and Zuo, S. 2016b. Non-clairvoyant dynamic mechanism design. Working paper, Google Research.

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

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

Parkes, D. C., and Singh, S. P. 2004. An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, 791–798.

Parkes, D. C.; Yanovsky, D.; and Singh, S. P. 2005. Approximately efficient online mechanism design. In *Advances in Neural Information Processing Systems*, 1049–1056.

Pavan, A.; Segal, I.; and Toikka, J. 2014. Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2):601–653.

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

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