Optimal Long-Term Supply Contracts with Asymmetric Demand Information

Ilan Lobel and Wenqiang Xiao*

January 26, 2017

Abstract

We consider a manufacturer selling to a retailer with private demand information arising dynamically over an infinite time horizon. Under a backlogging model, we show that the manufacturer's optimal dynamic long-term contract takes a simple form: in the first period, based on her private demand forecast, the retailer selects a wholesale price and pays an associated upfront fee, and, from then on, the two parties stick to a simple wholesale price contract with the retailer's chosen price. Under a lost sales model, we show that the structure of the optimal long-term contract combines a menu of wholesale pricing contracts with an option that, if exercised by the retailer, reduces future wholesales prices in exchange for an immediate payment to the manufacturer.

1 Introduction

We study a supply chain model that is composed of a manufacturer (he) selling to a retailer (she) over an infinite time horizon. At the beginning of each period, the retailer observes a demand forecast containing valuable information about the random demand to be realized in the current period. Both the retailer's demand forecast and the actual realized demand at each period are privately observed by the retailer. We assume the retailer carries over unused inventory and that excess demand is either backlogged or lost. The retailer's inventory is also assumed to be private information. The manufacturer is allowed to offer the retailer a dynamic long-term contract of his choice. The contract is dynamic because the payment charged for a given order could potentially depend on the entire history of events. The contract is long-term because it specifies the terms of trade over the entire time horizon.

*Stern School of Business, New York University – {ilobel, wxiao}@stern.nyu.edu. Ilan Lobel thanks the National Science Foundation for its financial support (Grant CCF-1216004). We thank Hamid Nazerzadeh, the area and associate editors and two referees for their insights and suggestions.

The first main finding of our paper is that, in the case where the retailer backlogs excess demand, the manufacturer's optimal long-term contract takes a very simple form: the manufacturer offers a menu of wholesale prices and upfront payments associated with these prices; the retailer selects her preferred wholesale price, pays an associated upfront fee and, from then on, places orders with the manufacturer at the chosen wholesale price. Our second main finding is the structure of the optimal long-term contract for the case where excess demand is lost. In this case, the manufacturer also offers the retailer a menu of wholesale prices and associated upfront payments. However, these contracts have an additional feature. All of these contracts have an option embedded in them that allows the retailer to pay a fixed fee and, in return, have the wholesale price reduced from that point onwards to the marginal production cost.

Our results contrast with the results currently in the literature. In a static model where a manufacturer sells to a retailer who has private demand information, the manufacturer's optimal contract is a nonlinear quantity discount in which the marginal per-unit price is decreasing in the quantity purchased (see Burnetas et al. [2007]). Zhang et al. [2010] generalize the selling-to-newsvendor model to a multi-period setting in which the retailer's dynamically arising demand/inventory information is unknown to the manufacturer. They focus on a setting with dynamic short-term contracts, i.e., where contracting takes place in every period, and show that the optimal payment scheme is often of the form of a batch-order contract, where the seller makes a single take-it-or-leave-it offer to the retailer at each period.

Despite these theoretical predictions, what we mostly observe in practice are simple terms of trade whereby a manufacturer sells to a retailer at a per-unit price that remains unchanged over a long time horizon. In this paper, we aim to reconcile this paradox by showing that, in the case of backlogging, such simple trade practice is, in fact, the best choice of a manufacturer selling to a retailer over an infinite horizon when the retailer has private information about her demand.

The intuition as to why such a simple contract is optimal can be explained as follows. The manufacturer designs a contract with the goal of maximizing his own profits, which equals the total supply chain profits minus the retailer's information rent. Optimizing the manufacturer's profit imposes somewhat contradictory pressures on prices and the optimal way for the manufacturer to tradeoff revenue extraction and supply chain efficiency in a single-period problem is via a nonlinear pricing scheme. In a multi-period problem, the manufacturer has an additional flexibility: he can rearrange payments over time. The payments associated with supply chain efficiency necessarily have to take place over time in order to induce the retailer to order appropriate amounts. Any excess payments the manufacturer desires to impose for revenue extraction, however, can be freely moved in time. This gives rise to an optimal long-term contract with an upfront fee in the first period and wholesale pricing thereafter.

Our finding has several implications for the manufacturer. First, instead of using a nonlinear contract (as suggested by Burnetas et al. [2007] for static contracts), a batch-order contract (as proposed by Zhang et al. [2010] for dynamic settings where only short-term contracting is allowed) or, alternatively, a complex dynamic contract, it is in the manufacturer's best interest to sell to the retailer at a constant per-unit price with the commitment that this price does not change over time. Second, with demand information disadvantage, the manufacturer ought to let the retailer with superior demand information self-select the wholesale price at the beginning of the time horizon by paying a corresponding one-time upfront fee.

Long-term supply contracts are often used in practice, but are an understudied topic of research. A manufacturer-retailer supply contract is long-term if the duration of the contract is much longer than the interval between replenishment dates. For example, a supermarket might sign a multi-year contract with one of its suppliers that involves weekly deliveries to replenish inventory. With a long-term contract, the manufacturer and the retailer do not need to engage in a negotiation every time the retailer needs more units of inventory.

Long-term contracts of the form we suggest for the case of backlogging – menus of fixed fees plus wholesale prices – are widely used in retail. These contract formats are often called two-part tariffs when the manufacturer charges the retailer a fixed fee, and called slotting allowances when the manufacturer pays the retailer for the use of its shelf space (see Marx and Shaffer [2010]). There are many examples of companies that sell its products via menus of two-part tariff contracts. They include the warehouse club Costco Wholesale, which charges different prices to different customers depending on their membership level. They also include the online retailer Amazon.com, which offers free shipping to customers who agree to pay fixed annual fees to be Amazon Prime members. Another example is the rental car by-the-hour company Zipcar, which reduces its hourly rate to customers who agree to pay fixed monthly fees. There are now even academic papers studying how to optimize menus of two-part tariffs. In a recent contribution, Perakis and Thraves [2016] optimize the contracts offered by a firm whose main business is to provide satellite phone communication to maritime shipping companies. In that work, Perakis and Thraves use a data-driven approach to increase revenues by improving the menu of fixed fees and per-minute prices offered by the satellite phone company.

There are a few key assumptions that are needed to prove our results. Perhaps no assumption is more crucial than commitment power. In our model, we assume the manufacturer commits to a long-term contract and abides by it. We also assume there is no order lead time, an assumption that is innocuous in the backlogging case, but important in the lost sales model. We also make simple assumptions about demand and forecasting. We assume that demand in each period is independent and identically distributed and that the retailer has a private demand forecast only about the upcoming period's demand. Assuming a slightly more complex demand model is possible – our results would still hold if the demand forecasts were still independent but not identically distributed – but an arbitrary demand stochastic process, with a potentially arbitrary demand forecasting by the retailer, would make the contracting problem significantly more difficult to solve.

Our main results are presented in Section 4, with the proofs being offered in Appendix A. We also present the reader with an online appendix with additional results. In Appendix B, we extend our model to a finite horizon setting. We consider a problem where excess demand is backlogged in all periods, except the last, when unfulfilled orders are lost. We show that, under certain conditions, the optimal long-term contract takes the form of a sequence of wholesale prices and associated upfront fees, with different wholesale prices in force in different periods. In Appendix C, we consider the effect of allowing the retailer to break the contract at any time. We show that if the retail price is sufficiently high, the retailer will never want to terminate the contract early. However, if the retail price is too low, the manufacturer needs to add an additional feature to the contract in order to ensure dynamic participation in this case. This feature can take the form of an early termination fee, or take the form of periodic fixed payments to the retailer.

2 Literature

There is an extensive literature that explores how contracts should be designed to mediate interactions among self-interested firms with private information in a supply chain (see, for example, Ha [2001] on production cost information asymmetry, Nazerzadeh and Perakis [2011] on capacity constraint information asymmetry and Cachon and Lariviere [2001], Shin and Tunca [2010] and Taylor and Xiao [2010] on demand information asymmetry). All of these papers focus on static models with one-shot operational decisions. A finding that is relevant to our work is that in the single-period model where a manufacturer sells to a newsvendor retailer with private demand information, the manufacturer's optimal contract takes the form of a concave quantity discount where the marginal unit payment decreases in the order quantity (Burnetas et al. [2007]). A quantity discount contract can be interpreted as a menu of two-part tariffs, whereby the retailer is asked to choose a pair of wholesale price and upfront payment based on her private demand information.

Recently, several pioneering studies have explored multi-period contracting problems where private information arises over time and operational decisions need to be made dynamically based on available information. Zhang et al. [2010] studies a manufacturer selling to a retailer over multiple periods with asymmetric demand and inventory information. They focus on dynamic short-term contracts and show that the optimal contracts take the form of batch-order contracts under many circumstances. Compared with the multi-period model in Zhang et al. [2010], we allow the retailer to possess private demand forecast information at the beginning of every period, while maintaining the same assumption that the realized demand (and thus inventory) in every period is the retailer's private information. Since demand in our model is given by the sum of demand forecast and forecast error, the two papers' information structures are isomorphic. Crucially, we consider long-term contracts while they study short-term contracts.

In our work, we characterize the optimal dynamic long-term contract and show that it takes a simple form. We are not the first to advocate the efficiency of simple contracts such as the wholesale price contract in multi-period contracting problems. Ren et al. [2010] show that a wholesale price contract, coupled with a multi-period review strategy profile, is efficient in governing a long-term repeated interaction within a supply chain with demand information asymmetry.

To reconcile the apparent conflict between the theoretical suboptimality of wholesale pricing contracts and their prevalent use in practice, a stream of experimental research has been done in controlled laboratory settings that demonstrate the wholesale price contract is more efficient than what the theory predicts (Katok and Wu [2009], Kalkanci et al. [2011]); the wholesale price contract has also been shown to have desirable properties from a social welfare perspective (Cui et al. [2007]). In this regard, we identify a new appealing feature of linear wholesale pricing when used in multi-period interactions with information asymmetry: it eliminates the informed party's incentives of misreporting demand/inventory information earlier in the hope of gaining strategic advantage in future transactions.

Finally, our work is related to the recent literature on optimal dynamic mechanism design. We use a relaxation approach that was first pioneered by Ëso and Szentes [2007] and is by now standard in the literature (see Kakade et al. [2013], Pavan et al. [2014] or the survey by Bergemann and Said [2011]). The environment we consider satisfies the notions of independent shocks and regularity of Pavan et al. [2014]'s Theorem 1, which offers a dynamic version of the envelope theorem. We could thus use this theorem to remove payments from the manufacturer's optimization problem, à la Myerson [1981]. However, this theorem does not guarantee that the original problem and the relaxed problem generate equal profits. To prove this equivalency, we could use Proposition 1 from Pavan et al. [2014] if the environment were strongly monotone or Theorem 4.1 from Kakade et al. [2013] if the environment were additively or multiplicatively separable. However, our environment satisfies neither of these conditions. In particular, strong monotonicity requires the retailer's inventory to be unaffected by the manufacturer's actions, which clearly is not the case (see the definition of F-AUT in Pavan et al. [2014]). An important contribution of our paper is to use multi-period inventory theory to generate one such payment rule – wholesale pricing – that makes the ordering policy generated by the relaxation method dynamically incentive compatible, thus proving that the two problems generate equal profits. No general characterization is yet known of when the relaxation method produces solutions that can be made dynamically incentive compatible – see Section 6 in Kakade et al. [2013] for examples of dynamic mechanism design problems that are known not to be solvable by the relaxation method. In particular, a recent paper by Battaglini and Lamba [2015] shows that the relaxation approach fails for a large class of dynamic mechanism design problems. Fortunately, our supply chain contracting can indeed be solved by the relaxation method and the solution it generates is simple and practical. Our result does rely on the assumption that demand forecasts are drawn from a continuous distribution. A recent paper by Krähmer and Strausz [2015] shows that the relaxation method often fails when the first period signal is discrete.

3 The Two Models

In this section, we introduce two supply chain contracting models, the first being a backlogging model and the second being a lost sales one. The two models are identical except for the way in which excess demand is handled (backlogging vs. lost sales). In what follows, everything applies to both models unless specified otherwise.

We consider a manufacturer selling to a retailer, who subsequently sells to consumers at a given retail price p, over an infinite number of time periods, indexed by $t = 1, 2, \ldots$. In period t, the retailer has a demand forecast μ_t, which deviates from the true demand by a zero-mean forecast error ε_t. That is, the retailer's period t demand is $D_t = \mu_t + \varepsilon_t$. Both μ_t and ε_t are independent, identically distributed sequences of random variables and their respective cumulative probability distributions are $F(\cdot)$ and $G(\cdot)$, with densities $f(\cdot)$ and $g(\cdot)$, respectively.[1] Let $[\underline{\mu}, \overline{\mu}]$ and $[\underline{\varepsilon}, \overline{\varepsilon}]$ be the respective supports of μ_t and ε_t, and μ^* be the expected demand. To avoid negative demand, we assume $\mu_t + \varepsilon_t \geq 0$ almost surely, i.e., $\underline{\mu} + \underline{\varepsilon} \geq 0$.

The sequence of events is as follows. At the beginning of period t, the retailer observes her inventory level x_t and demand forecast μ_t for that period. Second, the retailer decides the order quantity q_t. We assume there is no order lead time.[2] The manufacturer incurs a production cost c for every unit produced. Third, the demand D_t is realized. If the demand D_t is higher than the retailer's available inventory $x_t + q_t$, then excess demand is backlogged with a per-unit penalty cost b under the backlogging model and is lost under the lost sales model. Otherwise, leftover inventory is carried over to the next period with a per-unit holding cost h. We assume that at the first period the retailer has not yet interacted with the manufacturer and, therefore, the initial inventory level

[1]The results in our paper still apply if the forecasts μ_t are independent, but not identically distributed. The same is not true if the forecasts errors are not independent or identically distributed.

[2]The assumption of no lead time is without loss of generality for the backlogging model. However, inventory control with positive lead times under lost sales is a difficult problem, where the structure of an optimal policy is not known (see Bijvank and Vis [2011] and Xin and Goldberg [2016]). Since base stock policies are not optimal under lost sales with positive lead times, our solution approach does not immediately generalize to this case.

x_1 is equal to 0.[3]

We assume that all demand and inventory information is private information of the retailer. That is, the manufacturer does not have access to any data on the retailer's inventory level $\{x_t\}$, the retailer's demand forecast $\{\mu_t\}$ or realized demand $\{D_t\}$. That is, the manufacturer only knows the quantities he has supplied the retailer $\{q_t\}$ and the distributions $F(\cdot)$ and $G(\cdot)$ of the retailer's demand forecast and demand forecasting error, respectively. The model is thus one of *dynamic asymmetric information*, where the retailer accumulates private inventory and demand information over time. The only piece of asymmetric information at period 1 is the demand forecast μ_1, but all other pieces of information that the retailer privately learns over time are also contract relevant. We make the following mild assumption on the distribution of demand forecasts and forecast errors: both $F(\cdot)$ and $G(\cdot)$ have an increasing failure rate, i.e., both $f(\cdot)/\overline{F}(\cdot)$ and $g(\cdot)/\overline{G}(\cdot)$ are increasing functions. The increasing failure rate assumption is satisfied by many common distributions such as the normal, uniform, and exponential distributions.

We study the manufacturer's problem of designing the long-term dynamic supply contract that maximizes his expected total discounted profits. The contracts we consider are long-term because they determine the terms of trade for all periods. The contracts are also dynamic because those terms of trade are allowed to evolve in any way of the manufacturer's choosing and are allowed to be contingent on the retailer's communication of newly observed information over time.

Myerson [1981] argued, within the context of static mechanism design, that any outcome that is implementable in equilibrium in an arbitrary mechanism can also be implemented in equilibrium via a direct mechanism, a result that is known as the revelation principle. In later work, the same author showed that the revelation principle also applies in dynamic settings as long as the mechanism designer has commitment power (see Myerson [1986]), a result that has become the starting point of the literature on dynamic mechanism design (see, for example, Pavan et al. [2014]). Since the manufacturer has commitment power in our model, it's sufficient to search over direct long-term dynamic contracts.[4]

In our setting, a direct mechanism is one where at each period t, the retailer is asked to report her demand forecast $\widehat{\mu}_t$ and the forecasting error $\widehat{\varepsilon}_t$ (the inventory level x_t can be deduced from these two pieces of information together with past orders, and thus need not to be reported). We represent the period t history of realized and forecasted demand by h_t and history of reports up

[3]Our approach would also extend to a setting with initial private inventory. In this case, the retailer's first period forecast would effectively be $\tilde{\mu}_1 = \mu_1 - x_1$. The distribution of the effective first period demand forecast would be different than $F(\cdot)$, but the methodology we propose would still apply.

[4]Since there is only one retailer in our model, and thus no issue of communication across agents over time, even the static version (Myerson [1981]) of the revelation principle applies. The manufacturer can commit to an allocation rule (how many units to order at each stage given the history of retailer reports) and a payment rule upfront, and the retailer will face a stochastic optimization problem after that.

to period t by $\widehat{h}_t$. That is, $h_1 = \mu_1$ and $h_t = \{\mu_1, \varepsilon_1, \mu_2, \varepsilon_2, ..., \mu_{t-1}, \varepsilon_{t-1}, \mu_t\}$ for $t \geq 2$. A direct long-term dynamic contract takes the form $\{(q_t(\widehat{h}_t), T_t(\widehat{h}_t))|_{t\in\mathbb{N}}\}$, where q_t represents the quantity of units delivered by the manufacturer to the retailer in period t and T_t represents the payment made by the retailer to the manufacturer in period t.

4 Optimal Long-Term Dynamic Contracts

In this section, we consider how to optimize long-term dynamic supply chain contracts. We first do so by considering the backlogging model, and we subsequently analyze the lost sales model.

4.1 The Backlogging Case

We begin by discussing the retailer's problem under a given contract $\{(q_t(\widehat{h}_t), T_t(\widehat{h}_t))|_{t\in\mathbb{N}}\}$. First, we describe the retailer's inventory dynamics under backlogging. In period t, the retailer's starting inventory can be determined by h_t and $\widehat{h}_{t-1}$. To see this, suppose the starting inventory in period $t-1$ is determined by h_{t-1} and $\widehat{h}_{t-2}$, and denote it by $x_{t-1}(h_{t-1}, \widehat{h}_{t-2})$. Clearly, the starting inventory in period t is equal to $x_{t-1}(h_{t-1}, \widehat{h}_{t-2})$ plus the order quantity $q_{t-1}(\widehat{h}_{t-1})$ less the demand $\mu_{t-1}+\varepsilon_{t-1}$, each of which can be determined by h_t and $\widehat{h}_{t-1}$. Hence, given h_t and $\widehat{h}_{t-1}$, the starting inventory in period t, denoted by $x_t(h_t, \widehat{h}_{t-1})$, is determined by the following recursive relation:

$$x_t(h_t, \widehat{h}_{t-1}) = x_{t-1}(h_{t-1}, \widehat{h}_{t-2}) + q_{t-1}(\widehat{h}_{t-1}) - \mu_{t-1} - \varepsilon_{t-1}, \tag{1}$$

with $x_1 = 0$. In period t, the order quantity is $q_t(\widehat{h}_t)$, resulting in a total after-order inventory of $x_t(h_t, \widehat{h}_{t-1}) + q_t(\widehat{h}_t)$. Excluding the units that are to be used to satisfy the demand forecast μ_t, the remaining inventory is $x_t(h_t, \widehat{h}_{t-1}) + q_t(\widehat{h}_t) - \mu_t$, which we call the *safety stock* since it will be used to hedge against the forecast error ε_t. Let $y_t(h_t, \widehat{h}_t) = x_t(h_t, \widehat{h}_{t-1}) + q_t(\widehat{h}_t) - \mu_t$. It follows from Eq. (1) that the safety stock in period t can be determined by the following recursive relation:

$$y_t(h_t, \widehat{h}_t) = y_{t-1}(h_{t-1}, \widehat{h}_{t-1}) + q_t(\widehat{h}_t) - \mu_t - \varepsilon_{t-1}, \tag{2}$$

with $y_0 = \varepsilon_0 = 0$.

Next we derive the recursive relation for the retailer's profit-to-go function. Consider any one period. Given the safety stock y and the forecast μ, the retailer's expected sales revenue is $p\mu$ because the forecast error ε has zero mean, and her expected holding cost is $\mathbf{E}_\varepsilon[h(y-\varepsilon)^+]$ and backlogging cost is $\mathbf{E}_\varepsilon[b(\varepsilon - y)^+]$, where $z^+ = \max\{0, z\}$. Consequently, the retailer's expected profit (excluding the payment to the manufacturer) in the period is $p\mu - L(y)$, where $L(y) = \mathbf{E}_\varepsilon[h(y-\varepsilon)^+ + b(\varepsilon - y)^+]$. At the beginning of period t, given h_t and $\widehat{h}_{t-1}$, if the retailer reports

$\widehat{\varepsilon}_{t-1}$ for the realized forecast error in the previous period and $\widehat{\mu}_t$ for the forecast in the current period, then her maximum expected total discounted profit-to-go (i.e., her profit in period t plus the total discounted profits from period $t+1$ onwards), denoted by $\Pi_t(h_t, \widehat{h}_t)$ with $\widehat{h}_t = \{\widehat{h}_{t-1}, \widehat{\varepsilon}_{t-1}, \widehat{\mu}_t\}$, satisfies the following Bellman equation:

$$\Pi_t(h_t, \widehat{h}_t) = p\mu_t - L(y_t(h_t, \widehat{h}_t)) - T_t(\widehat{h}_t) + \delta \mathbf{E}_{\varepsilon_t, \mu_{t+1}} \max_{\widehat{\varepsilon}_t, \widehat{\mu}_{t+1}} \{\Pi_{t+1}(h_{t+1}, \widehat{h}_{t+1})\}, \tag{3}$$

where $y_t(h_t, \widehat{h}_t)$ is determined by Eq. (2), $\delta \in (0,1)$ is the discount factor, $h_{t+1} = \{h_t, \varepsilon_t, \mu_{t+1}\}$ and $\widehat{h}_{t+1} = \{\widehat{h}_t, \widehat{\varepsilon}_t, \widehat{\mu}_{t+1}\}$.

We now turn to the manufacturer's problem. In period t, given h_t and that the retailer has been following the truth-telling strategy from period 1 to t (i.e., $\widehat{h}_t = h_t$), the manufacturer's profit is given by the discounted sum of the retailer's payments $T_t(h_t)$ minus the production cost $cq_t(h_t)$. Therefore, if the retailer is truthful, the manufacturer's expected total discounted profit is

$$\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} \left(T_t(h_t) - cq_t(h_t)\right)\right], \tag{4}$$

where the expectation is taken over h_∞.

To ensure truthful reporting by the retailer, the contract written by the manufacturer should satisfy dynamic incentive compatibility constraints. Specifically, for any given up-to-date information $h_t = \{h_{t-1}, \varepsilon_{t-1}, \mu_t\}$ at the beginning of period t, the retailer should be better off under truthful report h_t than under any other report $\widehat{h}_t$. Mathematically,

$$\Pi_t(h_t, h_t) \geq \Pi_t(h_t, \widehat{h}_t), \text{ for any } t,\ h_t \text{ and } \widehat{h}_t, \tag{IC}$$

where $\Pi_t(\cdot,\cdot)$ is determined by Eq. (3). The (IC) constraints ensure that it is in the retailer's best interest to truthfully reveal her private information in every period. Furthermore, the contract should satisfy individual rationality constraints. Specifically, when the contract is announced at the beginning of period 1, the retailer's information consists of only the forecast for the first period demand, i.e., $h_1 = \{\mu_1\}$. To ensure the retailer's acceptance of the contract, her expected total discounted profit $\Pi_1(\mu_1, \mu_1)$ under truth-telling should be no less than her reservation profit, which is normalized to zero without loss of generality. Mathematically,

$$\Pi_1(\mu_1, \mu_1) \geq 0, \text{for any } \mu_1. \tag{IR}$$

We consider a model with dynamic participation constraints in Appendix C.

To summarize, the manufacturer's problem of finding the optimal direct long-term dynamic mechanism can be formulated as follows:

$$\max_{\{(q_t(\cdot),T_t(\cdot))|_{t\in\mathbb{N}}\}} \mathbf{E}\left[\sum_{t=1}^{\infty}\delta^{t-1}\left(T_t(h_t)-cq_t(h_t)\right)\right] \tag{P}$$

$$\text{s.t.} \qquad \Pi_t(h_t,h_t)\geq \Pi_t(h_t,\widehat{h}_t), \text{ for any } t,\ h_t \text{ and } \widehat{h}_t, \tag{IC}$$

$$\Pi_1(\mu_1,\mu_1)\geq 0, \text{ for any } \mu_1. \tag{IR}$$

It is worthwhile to note that by considering long-term dynamic contracts, the space of allowable contracts is incredibly general. The fact that the contract is long-term by no means implies that it's a stationary one or that it has any other particular structure.

The above optimization problem might appear at first glance to be too complex to be tractable, but we will show that this problem is actually solvable and the optimal solution is actually quite simple. We now present our solution methodology for solving the manufacturer's contract optimization problem (P). The first step in our technique is to consider a relaxation of the original problem by imposing an assumption that the manufacturer can observe all of the retailer's demand forecasts and all of the realized errors, except the demand forecast μ_1 in period 1. We call this problem the *relaxed problem*. Because the manufacturer, as the contract designer, has more information that is contractible in the relaxed problem than in the original problem, the manufacturer should be no worse off under the relaxed problem, implying that solving the relaxed problem yields an upper bound on the maximum expected profit that the manufacturer can achieve under the original problem.

By the revelation principle, it is sufficient to look for direct long-term dynamic contracts whereby, in period 1, the retailer reports her demand forecast $\widehat{\mu}_1$, and the order quantity and payment in every period t is specified as a function of the retailer's report $\widehat{\mu}_1$ and the up-to-date demand information $h_t^{-1}=\{\mu_2,...,\mu_t,\varepsilon_1,...,\varepsilon_{t-1}\}$. A direct mechanism can be formally represented by the quantity-payment pair $(q_t(h_t^{-1},\widehat{\mu}_1),T_t(h_t^{-1},\widehat{\mu}_1))$ for each period $t\in\mathbb{N}$. That is, the only difference between a direct mechanism for the original problem and a direct mechanism for the relaxed problem is that, in the former, the order and payment functions take the reports of the demand forecasts and realizations as inputs, while in the latter the same functions take as inputs the initial forecast report $\widehat{\mu}_1$ together with actual forecasts and demand shocks h_t^{-1} for periods after $t=1$.

Consider any given direct mechanism $\{(q_t(h_t^{-1},\widehat{\mu}_1),T_t(h_t^{-1},\widehat{\mu}_1))|_{t\in\mathbb{N}}\}$ of the relaxed problem. We first discuss the retailer's problem. In period t, given h_t and the retailer's report $\widehat{\mu}_1$ submitted in period 1, all of the demand and orders before period t are known, and so is the starting inventory in period t. We denote this inventory by $x_t(h_t,\widehat{\mu}_1)$, which is determined by the following recursive relation:

$$x_t(h_t, \widehat{\mu}_1) = x_{t-1}(h_{t-1}, \widehat{\mu}_1) + q_{t-1}(h_{t-1}^{-1}, \widehat{\mu}_1) - \mu_{t-1} - \varepsilon_{t-1},$$

with $x_1 = 0$. Similarly, the safety stock in period t, denoted by $y_t(h_t, \widehat{\mu}_1)$, is determined by a recursive relation:

$$y_t(h_t, \widehat{\mu}_1) = y_{t-1}(h_{t-1}, \widehat{\mu}_1) + q_t(h_t^{-1}, \widehat{\mu}_1) - \mu_t - \varepsilon_{t-1}, \tag{5}$$

with $y_0 = \varepsilon_0 = 0$.

Because the retailer makes only a single decision in the relaxed problem, which is what to report after observing the true demand forecast μ_1 in period 1, it suffices to characterize her expected total discounted profit for any given demand forecast μ_1 and her report $\widehat{\mu}_1$, which we denote by $\Pi_1(\mu_1, \widehat{\mu}_1)$. In each period t, given h_t and $\widehat{\mu}_1$, the retailer's safety stock is $y_t(h_t, \widehat{\mu}_1)$, implying that she earns the one-period newsvendor profit $p\mu_t - L(y_t(h_t, \widehat{\mu}_1))$ less the payment $T_t(h_t^{-1}, \widehat{\mu}_1)$ to the manufacturer. Consequently, we have

$$\Pi_1(\mu_1, \widehat{\mu}_1) = \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}[p\mu_t - L(y_t(h_t, \widehat{\mu}_1)) - T_t(h_t^{-1}, \widehat{\mu}_1)] \;\middle|\; \mu_1\right]. \tag{6}$$

We now turn to the manufacturer's problem. Compared with the original problem (P), under the retailer's truth-telling strategy, the manufacturer's profit function remains unchanged, because $(\mu_1, h_t^{-1}) = h_t$. The (IR) constraints remain the same as well. However, the (IC) constraints are greatly simplified because only the (IC) constraints for period 1 matter now. The relaxed problem can be formally stated as follows:

$$\max_{\{(q_t(\cdot), T_t(\cdot))|_{t\in\mathbb{N}}\}} \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}\left(T_t(h_t) - cq_t(h_t)\right)\right] \tag{P'}$$

$$\text{s.t.} \qquad \Pi_1(\mu_1, \mu_1) \geq \Pi_1(\mu_1, \widehat{\mu}_1) \text{ for any } \mu_1, \text{and } \widehat{\mu}_1, \tag{IC'}$$

$$\Pi_1(\mu_1, \mu_1) \geq 0 \text{ for any } \mu_1. \tag{IR'}$$

Even though the (IC') and (IR') constraints in the relaxed problem are similar to those in a single period problem, the relaxed problem (P') is still a more complex problem than a single period one since it involves selecting a potentially elaborate order policy $\{q_t(h_t)\}$, where order quantities are allowed to depend not only the retailer's report in period 1 but also on the historical demand information h_t^{-1}.

We now provide an optimal mechanism for the relaxed problem (P'). We use the notation $\{(q_t^r(h_t), T_t^r(h_t))|_{t\in\mathbb{N}}\}$ to represent this mechanism, where the superscript r represents the relaxed problem. We call an ordering policy a base stock policy if it orders up to a given safety stock level, and does not order any units if the safety stock is already above the desired level.

Proposition 1. *Under the optimal mechanism of the relaxed problem* $\{(q_t^r(h_t), T_t^r(h_t))|_{t\in\mathbb{N}}\}$*, the retailer with demand forecast* μ_1 *in period 1 follows a base stock policy with constant safety stock level* $y^*(\mu_1)$ *in every period, where* $y^*(\mu_1)$ *satisfies*[5]

$$b - (b+h)G(y^*(\mu_1)) - (1-\delta)c - \frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y^*(\mu_1)) = 0.$$

Specifically, the order quantities are $q_1^r(\mu_1) = \mu_1 + y^*(\mu_1)$ *and* $q_t^r(h_t) = \mu_t + \varepsilon_{t-1}$ *for* $t \geq 2$*, with a single payment in the first period*

$$T_1^r(\mu_1) = \frac{p\mu^* - L(y^*(\mu_1))}{1-\delta} - \int_{\underline{\mu}}^{\mu_1} \left\{ p + \frac{-b + (b+h)G(y^*(\mu))}{1-\delta} \right\} d\mu$$

and no payments thereafter, i.e., $T_t^r(h_t) = 0$ *for all* h_t *and* $t \geq 2$.

The proposition above reveals that the ordering decision in each period depends on the historical demand information in a simple and intuitive way. Specifically, in period 1, the order quantity $q_1^r(\mu_1)$ intended for the type μ_1 retailer (who observed the demand forecast μ_1) pushes her inventory to $\mu_1 + y^*(\mu_1)$, with $y^*(\mu_1)$ effectively being the safety stock to cope with the uncertain forecast error ε_1; in any period $t \geq 2$, the order quantity $q_t^r(h_t)$ is simply equal to the forecast error ε_{t-1} realized in the previous period $t-1$ plus the current demand forecast μ_t, resulting in a constant safety stock $y^*(\mu_1)$ to satisfy the forecast error ε_t for the type μ_1 retailer under truth-telling.

To see the intuition for the result that the safety stock intended for each retailer type is kept at a constant level in every period, it is useful examine the role played by the safety stock. The literature on the problem of a manufacturer selling to a newsvendor retailer with private demand information has revealed that the tradeoff between improving system efficiency (improving the total pie of the supply chain) and limiting the retailer's information rent (shrinking the retailer's share of the total pie) determines the safety stock in a single period problem. In a multi-period setting where the only source of information asymmetry is the demand forecast in period 1, there are two reasons why the manufacturer faces the exact same tradeoff repeatedly in determining the safety stock in each period. First, since the problem is stationary, the impact of the safety stock on the system efficiency remains the same from period to period. Second, the impact of the safety stock on any type μ_1 retailer's information rents also remains the same from period to period. That is, since demand forecasts and shocks after the first period are observable to the manufacturer, the manufacturer can select the optimal resupply amount $\mu_t + \varepsilon_{t-1}$ without paying rents to the retailer

[5] This equation assumes that $y^*(\mu_1)$ is an interior solution of $[\underline{\varepsilon}, \overline{\varepsilon}]$. If the solution of the equation is below this interval, then $y^*(\mu_1) = \underline{\varepsilon}$. Similarly, if the solution is above the interval, then $y^*(\mu_1) = \overline{\varepsilon}$.

at any later period ($t \geq 2$). With this reorder, the problem at period $t+1$ becomes identical to the problem at period t. This way, any advantage that a high-type retailer has over a low-type retailer in period t, immediately carries over to any later period $t' > t$. Consequently, facing the same tradeoff between efficiency and rent in every period, the manufacturer's optimal strategy is to keep the safety stock intended for type μ_1 retailer at the constant level $y^*(\mu_1)$. Therefore, despite the fact that μ_1 is only a forecast of period 1 demand, the allocative distortion that the manufacturer generates to minimize the information rent associated with μ_1 is permanent and affects the safety stock in all periods.

We note that Proposition 1 assumes that the retailer chooses to sign a contract with the manufacturer. If the retail price p is too low, however, a retailer with a low first period forecast μ_1 might choose not to sign a contract with the manufacturer. If this were to be the case, there would be a minimal forecast $\mu' > \underline{\mu}$ such that contracts would only be signed with retailers with types above μ'. If this were to be the case, the structure of the optimal supply contract would still be a combination of wholesale prices and associated upfront fees.

Proposition 1 establishes that it is in the best interest of the manufacturer to induce the type μ_1 retailer to follow an ordering policy so that her safety stock is kept at $y^*(\mu_1)$ in every period. This goal can be achieved by the manufacturer in the relaxed problem because he can dictate the order quantity in period $t \geq 2$ to be equal to the sum of the realized forecast error in the previous period and the newly observed demand forecast in the current period, i.e., $q_t^r(h_t) = \mu_t + \varepsilon_{t-1}$ as such information is observable to him. Thus, there is no need for transfer payments after the first period in the relaxed problem and, therefore, we can construct a solution with $T_t^r(h_t) = 0$ for all h_t and $t \geq 2$.

However, back in the original problem, demand forecasts and realizations are no longer observable to the manufacturer, implying that the order quantities must be made contingent on the retailer's reports and triggering the need to satisfy incentive compatibility constraints in every period. The solution of the relaxed problem $\{(q_t^r(h_t), T_t^r(h_t))|_{t\in\mathbb{N}}\}$ obtained in Proposition 1, with its lack of payments after the first period, will fail to induce truth-telling in the original model.

We now construct a second optimal solution of the relaxed problem that does satisfy the incentive compatibility constraints of the original problem. Since the problem (P') is a relaxation of (P), any optimal solution of (P') that is also feasible in (P) is immediately an optimal solution of (P). The technique of solving a relaxed version of the mechanism design problem and then using that solution to create a feasible solution to the original problem has been used in the literature before (see Ëso and Szentes [2007], Kakade et al. [2013], Pavan et al. [2014]). However, there does not exist a universal technique to create a solution to the original problem from the solution of the relaxed problem. Nevertheless, we next show that the relaxation approach does work in our multi-period inventory model.

We now describe the key idea that leads to the tractability of (P). Our goal is to create a different contract than $\{(q_t^r(h_t), T_t^r(h_t))|_{t\in\mathbb{N}}\}$, but one with the same ordering policy $\{q_t^r(h_t)|_{t\in\mathbb{N}}\}$ and a different payment rule that properly incentivizes the retailer to follow this ordering policy. It is known from classical multi-period inventory theory that, to induce the retailer to follow the ordering policy that results in a constant safety stock level in every period, it is sufficient to charge the retailer a constant wholesale price for every unit of order quantity, and that there is a one-to-one mapping between the desired safety stock level and the wholesale price. These arguments lead to one of the main results of this paper, which is formally stated in the following proposition.

Proposition 2. *The optimal long-term dynamic contract of the backlogging model takes the following form: in period 1, the manufacturer offers a menu of contracts of the form* $\{w^*(\mu_1), T^*(\mu_1)\}$ *specifying wholesale prices* $w^*(\mu_1)$ *and a fixed upfront payment* $T^*(\mu_1)$ *associated with each wholesale price, where*

$$w^*(\mu_1) = [b - (b+h)G(y^*(\mu_1))]/(1-\delta)$$

and

$$T^*(\mu_1) = T_1^r(\mu_1) - w^*(\mu_1)(\mu_1 + y^*(\mu_1)) - \delta w^*(\mu_1)\mu^*/(1-\delta). \tag{7}$$

The retailer observing demand forecast μ_1 *selects the wholesale price* $w^*(\mu_1)$ *and pays an upfront fee* $T^*(\mu_1)$*. Forever after (including the first period), the supply chain operates under the wholesale contract* $w^*(\mu_1)$*.*

What makes the dynamic adverse selection problem difficult is the requirement of satisfying the dynamic incentive constraints which involve multiple pieces of unknown information. Proposition 2 suggests such complicated incentive requirement can be met by offering, from period 2 onwards, a simple and time-invariant payment scheme, which is linear in the order quantity, where the order quantity is simply the sum of the retailer's report of the realized forecast error in the previous period and of the demand forecast in the current period. Such an approach, due to its time-invariant property and linearity, greatly simplifies the complexity of the incentive requirement. Specifically, because of its time-invariant property, the retailer can no longer benefit from intentionally delaying orders in the hope for a more favorable purchase price, or ordering more than what is needed driven by the fear of higher purchase prices in the subsequent periods; similarly, the linearity property completely eliminates the retailer's incentives of upward or downward order manipulation that exist under nonlinear payment schemes. Incentives to strategically manipulate order quantities, which can be done by reporting a false realized or forecasted demand, are nonexistent when future payments are linear and time-invariant.

An implication from the result that a linear payment scheme is optimal from period 2 onwards is that the mechanism can be decomposed into two simple components: first, the manufacturer offers a family of wholesale price contracts for the retailer to choose from, each one with a different upfront payment. The retailer will select a wholesale price contract based on her period 1 forecast. The manufacturer will then supply the retailer at her chosen wholesale price. For every possible forecast μ_1 of the retailer, the manufacturer will add a wholesale price $w^*(\mu_1)$ to the menu of contracts. With this contract form, instead of directly reporting her initial forecast μ_1, the retailer simply selects the wholesale price $w^*(\mu_1)$ and, from then onwards, it will be in her best interest to maintain the safety stock associated with that wholesale price.

4.2 The Lost Sales Case

In this subsection, we consider contracts under a lost sales inventory model under which any unsatisfied demand in every period is lost. For consistency, we are reusing the variables from the backlogging model here, with $x_t(h_t, \widehat{h}_{t-1})$ representing period t's inventory, $y_t(h_t, \widehat{h}_t)$ representing period t's safety stock and so on. Though we use the same representation, these variables satisfy different dynamics in the backlogging and the lost sales models.

Specifically, given h_t and $\widehat{h}_t$, the inventory dynamics satisfy

$$x_t(h_t, \widehat{h}_{t-1}) = \Big(x_{t-1}(h_{t-1}, \widehat{h}_{t-2}) + q_{t-1}(\widehat{h}_{t-1}) - \mu_{t-1} - \varepsilon_{t-1}\Big)^+$$

and the safety stock now satisfies

$$y_t(h_t, \widehat{h}_t) = q_t(\widehat{h}_t) - \mu_t + x_t(h_t, \widehat{h}_{t-1}),$$

with $y_0 = \varepsilon_0 = 0$. The retailer's profit-to-go function satisfies the following recursive relation:

$$\Pi_t(h_t, \widehat{h}_t) = p\mu_t - L(y_t(h_t, \widehat{h}_t)) - T_t(\widehat{h}_t) + \delta \mathbf{E}_{\varepsilon_t, \mu_{t+1}} \max_{\widehat{\varepsilon}_t, \widehat{\mu}_{t+1}} \{\Pi_{t+1}(h_{t+1}, \widehat{h}_{t+1})\}, \tag{8}$$

where $L(y) = \mathbf{E}_\varepsilon[h(y-\varepsilon)^+ + p(\varepsilon - y)^+]$.

The fact that the equation above looks identical to Eq. (3) might lure us into thinking the optimal contract under lost sales is identical to the one under backlogging. However, this is not the case. The lost sales problem is more difficult to solve because the safety stock dynamics are linear under backlogging, but nonlinear under lost sales. We now describe the optimal long-term contract under lost sales. Let $y^*(\mu_1)$ be the solution to[6]

$$-L'(y^*(\mu_1)) - c(1 - \delta G(y^*(\mu_1))) - \left[\frac{\overline{F}(\mu_1)}{f(\mu_1)}(p+h) + \delta(\widehat{\pi}(0) - \pi(0))\right] g(y^*(\mu_1)) = 0,$$

where $\widehat{\pi}(\cdot)$ and $\pi(\cdot)$ are defined in the proof of Proposition 3.

Proposition 3. *The optimal long-term dynamic contract of the lost sales model takes the following form: in period 1, the manufacturer offers a menu of contracts of the form* $\{w^*(\mu_1), T^*(\mu_1), T_0^*(\mu_1)\}$ *specifying a wholesale price* $w^*(\mu_1)$*, a fixed upfront payment* $T^*(\mu_1)$ *and an option priced at* $T_0^*(\mu_1)$ *that can be exercised at the end of any period by the retailer, where*

$$w^*(\mu_1) = [p - (p+h)G(y^*(\mu_1))]/(1 - \delta G(y^*(\mu_1))),$$

$$\begin{aligned} T^*(\mu_1) = \frac{p\mu^* - L(y^*(\mu_1))}{1-\delta} &- \int_{\underline{\mu}}^{\mu_1} \left\{ p + \frac{-p + (p+h)G(y^*(\mu))}{1 - \delta G(y^*(\mu))} \right\} d\mu \\ &- (\mu_1 + y^*(\mu_1))w^*(\mu_1) - \delta w^*(\mu_1)\mu^*/(1-\delta), \end{aligned}$$

and $T_0^*(\mu_1)$ *is set to ensure the retailer observing demand forecast* μ_1 *would be indifferent between exercising the option and not when she stocks out. The retailer observing demand forecast* μ_1 *selects the wholesale price* $w^*(\mu_1)$ *and pays an upfront fee* $T^*(\mu_1)$*. The supply chain operates under the wholesale contract* $w^*(\mu_1)$ *until the first stockout event occurs, at which point the option is exercised and the retailer pays the manufacturer the fee* $T_0^*(\mu_1)$ *to lower the wholesale price to* c*.*

The proof of this result is based on a relaxation argument that is similar, yet more elaborate than the proof for the backlogging case.[7] When we relax the dynamic mechanism design problem, we are left with an inventory control problem that is similar to the one we found in the backlogging case, but with one crucial distinction. In the backlogging case, the first period's information creates a permanent information rent term. With lost sales, the information rent term is temporary. To be more precise, the information rent term persists until the first time the company stocks out. After the first stockout, the manufacturer needs to solve an undistorted (no information rent) problem.

The first stockout plays a significant role because period t's safety stock is affected linearly by the first period demand forecast if and only if the first stockout has not happened yet. After the

[6] As in Proposition 1, this equation assumes that $y^*(\mu_1)$ is an interior solution of $[\underline{\varepsilon}, \overline{\varepsilon}]$. If the solution of the equation is below this interval, then $y^*(\mu_1) = \underline{\varepsilon}$. Similarly, if the solution is above the interval, then $y^*(\mu_1) = \overline{\varepsilon}$.

[7] As was the case with Proposition 1, the solution presented here assumes that the retailer chooses to sign a contract with the manufacturer for any initial type. If the retail price p is too low, we would also need to set a cutoff on types such that retailers with low initial forecasts would not sign contracts with the manufacturer. The rest of the optimal contract structure would remain as shown in Proposition 3.

stockout occurs, the first period's demand forecast no longer affects the retailer's safety stock. This gives rise to a complex and uncommon inventory problem. The retailer's optimal ordering policy after the first stockout is obvious: maintain the safety stock level that maximizes the supply chain profits. However, before the first stock out, the solution is non-trivial. In particular, the retailer has an incentive to understock in order to cause a stockout and thus reduce future purchase prices (after a stockout occurs, the supply chain should operate efficiently, and thus the retailer should be charged only the production cost). In this proposition, we prove that despite this unusual feature of the inventory model, the optimal ordering policy before the first stockout is a base stock one, with a safety stock level below the one maintained in the equivalent backlogging problem.

Solving the relaxed problem is not sufficient, however. We also need to come up with payments that dynamically incentivize the retailer to use the optimal ordering policy. Because the safety stock level is increasing over time, a pair of wholesale prices —one pre-stockout and post-stockout— can create the right incentives. However, the manufacturer also needs to convince the retailer to truthfully report that a stockout happened. In order to achieve this, we add the option feature to the contract. The option is priced so that the retailer does not want to use it unless she has stocked out. Even when she stocks out, she is still only indifferent between exercising the option and not doing so. In equilibrium, she chooses to exercise the option the first time she stocks out. Once the option is exercised, the retailer makes a payment to the manufacturer and the new wholesale price becomes active. We note the option should be exercised after the retailer stocks out, but before the retailer obtains her new demand forecast. At that moment in-between time periods, all retailer types who stocked out are effectively equivalent. Perhaps the greatest difference between the backlogging and lost sales models is that the operation eventually becomes efficient under lost sales (after the first stockout), but the supply chain inefficiency never disappears under backlogging.

References

M. Battaglini and R. Lamba. Optimal Dynamic Contracting: the First-Order Approach and Beyond. *Working Paper, Princeton University*, 2015.

D. Bergemann and M. Said. Dynamic Auctions: A Survey. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.

M. Bijvank and I. Vis. Lost-sales Inventory Theory: A Review. *European Journal of Operational Research*, 215(1):1–13, 2011.

A. Burnetas, S. Gilbert, and C. Smith. Quantity Discounts in Single-period Supply Contracts with Asymmetric Demand Information. *IIE Transactions*, 29(5):465–479, 2007.

G. Cachon and M. Lariviere. Contracting to Assure Supply: How to Share Demand Forecasts in a Supply Chain. *Management Science*, 47(5):629–646, 2001.

T. Cui, J. Raju, and Z. Zhang. Fairness and Channel Coordination. *Management Science*, 53(8): 1303–1314, 2007.

P. Ëso and B. Szentes. Optimal Information Disclosure in Auctions and the Handicap Auction. *Review of Economic Studies*, 74(3):705–731, 2007.

A. Ha. Supplier-Buyer Contracting: Asymmetric Cost Information and Cutoff Level Policy for Buyer Participation. *Naval Research Logistics*, 48(1):41–64, 2001.

S. Kakade, I. Lobel, and H. Nazerzadeh. Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism. *Operations Research*, 61(4):837–854, 2013.

B. Kalkanci, K. Chen, and F. Erhun. Contract Complexity and Performance under Asymmetric Demand Information: An Experimental Evaluation. *Management Science*, 57(4):689–704, 2011.

S. Karlin. Optimal Inventory Policy for the Arrow-Harris-Marschak Dynamic Model. *Chapter in Studies in the mathematical theory of inventory and production (Eds. K. Arrow, S. Karlin, and H. Scarf). Stanford University Press*, pages 135–154, 1958.

E. Katok and D. Wu. Contracting in Supply Chains: A Laboratory Investigation. *Management Science*, 55(12):1953–1968, 2009.

D. Krähmer and R. Strausz. Ex Post Information Rents in Sequential Screening. *Games and Economic Behavior*, 90:257–273, 2015.

L. Marx and G. Shaffer. Slotting Allowances and Scarce Shelf Space. *J. Econ. Manage. Strat.*, 19 (3):575–603, 2010.

R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

R. Myerson. Multistage Games with Communications. *Econometrica*, 54(2):323–358, 1986.

H. Nazerzadeh and G. Perakis. Menu Pricing Competition with Private Capacity Constraints. *Proceedings of the ACM Conference on Electronic Commerce*, 2011.

A. Pavan, I. Segal, and J. Toikka. Dynamic Mechanism Design: A Myersonian Approach. *Econometrica*, 82(2):601–653, 2014.

G. Perakis and C. Thraves. A Two-part Tariff Pricing Approach for Services: A Data Driven Approach. *In preparation, MIT*, 2016.

Z. Ren, M. Cohen, T. Ho, and C. Terwiesch. Sharing Forecast Information in a Long-Term Supply Chain Relationship: The Role of Customer Review Strategy. *Operations Research*, 58(1):81–93, 2010.

H. Shin and T. Tunca. Do Firms Invest in Forecasting Efficiently? The Effect of Competition on Demand Forecast Investments and Supply Chain Coordination. *Operations Research*, 58(6): 1592–1610, 2010.

T. Taylor and W. Xiao. Does a Manufacturer Benefit from Selling to a Better-Forecasting Retailer? *Management Science*, 56(9):1584–1598, 2010.

L. Xin and D. Goldberg. Optimality Gap of Constant-order Policies Decays Exponentially in the Lead Time for Lost Sales Models. *Forthcoming in Operations Research*, 2016.

H. Zhang, M. Nagarajan, and G. Sosic. Optimal Dynamic Supplier Contracts under Asymmetric Inventory Information. *Operations Research*, 58(5):1380–1397, 2010.

Optimal Long-Term Supply Contracts with Asymmetric Demand Information

Appendix

Ilan Lobel Wenqiang Xiao
{ilobel, wxiao}@stern.nyu.edu
Stern School of Business, New York University

Appendix A: Proofs

Proof of Proposition 1. The proof is carried out in three steps. Step 1. For any given contract that satisfies (IC') and (IR'), we use the (IC') constraints together with the envelope theorem to rewrite the retailer's profit as a function of the order quantities specified in the contract. Step 2. The result from Step 1 allows us to express the manufacturer's profit as a function of the order quantities. We derive the optimal order quantities that maximize the manufacturer's objective without considering the constraints. We then obtain the corresponding payment scheme which, together with the unconstrained optimal order quantities, satisfies the first-order conditions of the (IC') constraints and the (IR') constraints. Step 3. Because we replace the (IC') constraints by their first-order necessary conditions in deriving the order quantity-payment contract in Step 2, such a contract yields an upper bound on the manufacturer's expected profit. It then suffices to verify that this contract satisfies the (IC') constraints. Our proof procedure is similar to the standard approach solving the single-period adverse selection, with a distinction that the retailer's order quantity in every period is allowed to depend not only on the retailer's report in period 1 but also the up-to-date realized demand information.

Step 1. Let $\Pi_1(\mu_1) = \Pi_1(\mu_1, \mu_1)$, which is the type μ_1 retailer's expected profit under truth-telling. It follows from the (IC') constraints and the envelope theorem that

$$\begin{aligned}\Pi_1'(\mu_1) &= \frac{\partial \Pi_1(\mu_1, \widehat{\mu}_1)}{\partial \mu_1}\Big|_{\widehat{\mu}_1=\mu_1} \\ &= p + \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}[-b + (b+h)G(y_t(h_t))] \;\middle|\; \mu_1\right],\end{aligned}$$

where the expectation is taken over h_∞^{-1} and the last equality is due to Eq. (6). Consequently, the type μ_1 retailer's profit is equal to

$$\Pi_1(\mu_1) = \int_{\underline{\mu}}^{\mu_1} \Pi_1'(\mu) d\mu + \Pi_1(\underline{\mu})$$
$$= \int_{\underline{\mu}}^{\mu_1} \left\{ p + \mathbf{E} \left[\sum_{t=1}^{+\infty} \delta^{t-1} [-b + (b+h)G(y_t(h_t))] \;\middle|\; \mu \right] \right\} d\mu + \Pi_1(\underline{\mu}). \tag{9}$$

Step 2. Using the above expression for $\Pi_1(\mu_1)$, we can rewrite the manufacturer's objective function as the total supply chain profits minus the retailer's profit, which can be rewritten as

$$\begin{aligned}
&\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} \left(T_t(h_t) - c q_t(h_t)\right)\right] \\
&= \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} [p\mu_t - L(y_t(h_t)) - c q_t(h_t)] - \Pi_1(\mu_1)\right] \\
&= \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} [(p-c)\mu_t - L(y_t(h_t)) - c(y_t(h_t) - y_{t-1}(h_{t-1}))] - \Pi_1(\mu_1)\right] \\
&= \frac{(p-c)\mu^*}{1-\delta} + \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} [-L(y_t(h_t)) - c(1-\delta) y_t(h_t)] - \Pi_1(\mu_1)\right],
\end{aligned}$$

where the second equality is obtained by replacing $q_t(h_t)$ according to Eq. (2) and the third equality is simply a rearrangement that ensures that each $c y_t(h_t)$ term appears only once in the summation, plus a replacement of μ_t by its expected value μ^*. Replacing $\Pi_1(\mu_1)$ with the right-hand side of Eq. (9) and using Myerson's change of order of integration, we obtain that the manufacturer's profit is equal to $\frac{(p-c)\mu^*}{1-\delta}$ plus

$$\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} \left[-L(y_t(h_t)) - c(1-\delta) y_t(h_t) - \frac{\overline{F}(\mu_1)}{f(\mu_1)} [-b + (b+h)G(y_t(h_t))]\right] - \frac{\overline{F}(\mu_1)}{f(\mu_1)} p\right] - \Pi_1(\underline{\mu}), \tag{10}$$

We now optimize the safety stock levels in the equation above pointwise, that is, we find the value of $y_t(h_t)$ for any t and h_t. The solution to this pointwise maximization is $y_t(h_t) = y^*(\mu_1)$ for any given h_t, where

$$y^*(\mu_1) = \arg\max_y \left\{ -L(y) - \frac{\overline{F}(\mu_1)}{f(\mu_1)} (b+h)G(y) - (1-\delta) c y \right\}.$$

The maximand is strictly unimodal because its first order derivative is

$$
\begin{aligned}
& b-(b+h)G(y)-\frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y)-(1-\delta)c \\
&= b\overline{G}(y)-hG(y)-\frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y)-(1-\delta)c \\
&= \overline{G}(y)\left[b-h\frac{G(y)}{\overline{G}(y)}-\frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)\frac{g(y)}{\overline{G}(y)}-(1-\delta)\frac{c}{\overline{G}(y)}\right], \qquad (11)
\end{aligned}
$$

which changes the sign only once because $\frac{g(y)}{\overline{G}(y)}$ increases in y. Hence, $y^*(\mu_1)$ satisfies the following first-order condition

$$
b-(b+h)G(y)-\frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y)-(1-\delta)c=0.
$$

Therefore, the safety stock level is kept at $y^*(\mu_1)$ in every period for any μ_1, implying that the corresponding ordering policy is $\{q_t^r(h_t)\}$ where $q_1^r(\mu_1)=\mu_1+y^*(\mu_1)$ and $q_t^r(h_t)=\mu_t+\varepsilon_{t-1}$ for $t\geq 2$. Such an ordering policy, while letting $\Pi_1(\underline{\mu})=0$, maximizes the manufacturer's objective function. The payment function can be determined by ensuring that the retailer's profit is $\Pi_1(\mu_1)$ with $\Pi_1(\underline{\mu})=0$, i.e.,

$$
\begin{aligned}
T_1^r(\mu_1) &= \frac{p\mu^*-L(y^*(\mu_1))}{1-\delta}-\Pi_1(\mu_1) \\
&= \frac{p\mu^*-L(y^*(\mu_1))}{1-\delta}-\int_{\underline{\mu}}^{\mu_1}\left\{p+\sum_{t=1}^{\infty}\delta^{t-1}[-b+(b+h)G(y^*(\mu))]\right\}d\mu \\
&= \frac{p\mu^*-L(y^*(\mu_1))}{1-\delta}-\int_{\underline{\mu}}^{\mu_1}\left\{p+\frac{-b+(b+h)G(y^*(\mu))}{1-\delta}\right\}d\mu
\end{aligned}
$$

and the payment is zero in every subsequent period.

Step 3. Note that $y^*(\mu_1)$ increases in μ_1, which follows from the first-order condition given in Eq. (11) and the assumption that $\frac{\overline{F}(\cdot)}{f(\cdot)}$ is a decreasing function and $\frac{g(\cdot)}{\overline{G}(\cdot)}$ is an increasing function. This implies that $q_1^r(\mu_1)$ increases in μ_1. This, together with the fact that $q_t^r(h_t)$ is independent of μ_1, is sufficient to show that the contract $\{(q_t^r(\cdot),T_t^r(\cdot))|_{t\in\mathbb{N}}\}$ described above satisfies the (IC') constraints, and hence solves (P'). $\square$

Proof of Proposition 2. Consider first any (IC) constraint for $t\geq 2$. Regardless of the history h_t and the history of reports $\widehat{h}_t$, the retailer faces a standard multi-period inventory problem with a linear ordering cost of $w^*(\widehat{\mu}_1)$. Since the current and future ordering cost will not be affected by any of her decisions, the retailer's optimal policy is simple: always keep a safety stock of $y^*(\widehat{\mu}_1)$. If the retailer has been truthful up to now, she will find herself with an inventory of $y^*(\mu_1)-\varepsilon_{t-1}$

and will have a demand forecast of μ_t. To bring the safety stock to $y^*(\mu_1)$, the retailer will order exactly $q_t^o(h_t) = \varepsilon_{t-1} + \mu_t$. Thus, the retailer will continue to be truthful in periods $t \geq 2$.

Now consider the (IC) constraints for $t = 1$. A type μ_1 retailer's expected total discounted profit by choosing the contract intended for type $\widehat{\mu}_1$ is

$$\begin{aligned}\Pi_1(\mu_1, \widehat{\mu}_1) &= \mathbf{E}\left[\max_{y_t(h_t)} \sum_{t=1}^{\infty} \delta^{t-1}[p\mu_t - L(y_t(h_t)) - w^*(\widehat{\mu}_1)(y_t(h_t) - y_{t-1}(h_{t-1}) + \mu_t)]\right] - T^*(\widehat{\mu}_1) \\ &= \mathbf{E}\left[\max_{y_t(h_t)} \sum_{t=1}^{\infty} \delta^{t-1}[(p - w^*(\widehat{\mu}_1))\mu_t - L(y_t(h_t)) - (1-\delta)w^*(\widehat{\mu}_1)y_t(h_t)]\right] - T^*(\widehat{\mu}_1),\end{aligned}$$

where $y_t(h_t)$ is the safety stock in period t given h_t. Note that the maximand is unimodal in $y_t(h_t)$. By pointwise optimization, the maximizer is at $y_t(h_t) = y^*(\widehat{\mu}_1)$ since it satisfies the first-order conditions from the definitions of $w^*(\widehat{\mu}_1)$ and $y^*(\widehat{\mu}_1)$. This, together with the definition of $T^*(\widehat{\mu}_1)$, implies that the type μ_1 retailer's expected total discounted profit under truth-telling is

$$\Pi_1(\mu_1) = \int_{\underline{\mu}}^{\mu_1} \left\{p + \left[\frac{-b + (b+h)G(y^*(\mu))}{1-\delta}\right]\right\} d\mu,$$

implying that

$$\Pi_1'(\mu_1) = p + \left[\frac{-b + (b+h)G(y^*(\mu_1))}{1-\delta}\right] = p - w^*(\mu_1),$$

where the last equality follows from the definition of $w^*(\mu_1)$. Note that $\Pi_1(\mu_1, \widehat{\mu}_1) = \Pi_1(\widehat{\mu}_1) + (p - w^*(\widehat{\mu}_1))(\mu_1 - \widehat{\mu}_1)$, implying that

$$\begin{aligned}\frac{\partial \Pi_1(\mu_1, \widehat{\mu}_1)}{\partial \widehat{\mu}_1} &= \Pi_1'(\widehat{\mu}_1) - (p - w^*(\widehat{\mu}_1)) - [w^*(\widehat{\mu}_1)]'(\mu_1 - \widehat{\mu}_1) \\ &= -[w^*(\widehat{\mu}_1)]'(\mu_1 - \widehat{\mu}_1)\end{aligned}$$

which is nonnegative for $\mu_1 \geq \widehat{\mu}_1$ and nonpositive for $\mu_1 \leq \widehat{\mu}_1$ because $[w^*(\widehat{\mu}_1)]' \leq 0$. Therefore, it is in the best interest of type μ_1 retailer to select the wholesale price $w^*(\mu_1)$, implying that $\{w^*(\mu_1), T^*(\mu_1)\}$ satisfies the (IC) constraint in period 1. Clearly, the retailer's ordering quantity decisions under $\{w^*(\mu_1), T^*(\mu_1)\}$ are the same as those under the optimal direct truth-telling mechanism that solves (P') (see Proposition 1), and so is the manufacturer's expected total discounted profit. Therefore, these two methods achieve the same performance in expectation for the manufacturer. □

Proof of Proposition 3. We first solve the relaxed problem, where only the demand forecast μ_1 in period 1 is unobservable to the manufacturer. The optimal objective value of the relaxed problem is an upper bound of the objective value of the original problem. We then construct a menu of contracts and show that the constructed menu satisfies all the constraints of the original problem and achieves the upper bound. Hence, the constructed menu solves the original problem.

1. Solution to the relaxed problem. In what follows we use the approach similar to that in the proof of Proposition 1 to solve the relaxed problem. Let $\Pi_1(\mu_1) = \Pi_1(\mu_1, \mu_1)$, which is the type μ_1 retailer's expected profit under truth-telling. It follows from the (IC') constraints and the envelope theorem that

$$\begin{aligned}\Pi_1'(\mu_1) &= \frac{\partial \Pi_1(\mu_1, \widehat{\mu}_1)}{\partial \mu_1}|_{\widehat{\mu}_1=\mu_1} \\ &= p + \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}[-p + (p+h)G(y_t(h_t))]1_{\{y_i(h_i)>\varepsilon_i, i=1,2,\ldots,t-1\}} \;\middle|\; \mu_1 \right],\end{aligned}$$

where $1_{\{y_i(h_i)>\varepsilon_i, i=1,2,\ldots,t-1\}}$ is equal to 1 if the condition inside the big brackets holds and zero otherwise, and the last equality is true due to Eq. (8). Consequently, the type μ_1 retailer's profit is equal to

$$\begin{aligned}\Pi_1(\mu_1) &= \int_{\underline{\mu}}^{\mu_1} \Pi_1'(\mu) d\mu + \Pi_1(\underline{\mu}) \\ &= \int_{\underline{\mu}}^{\mu_1} \left\{ p + \mathbf{E}\left[\sum_{t=1}^{+\infty} \delta^{t-1}[-p + (p+h)G(y_t(h_t))]1_{\{y_i(h_i)>\varepsilon_i, i=1,2,\ldots,t-1\}} \;\middle|\; \mu \right] \right\} d\mu + \Pi_1(\underline{\mu}).\end{aligned} \tag{12}$$

Using the above expression for $\Pi_1(\mu_1)$, we can rewrite the manufacturer's objective function as the total supply chain profits minus the retailer's profit, which can be rewritten as

$$\begin{aligned}&\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}[p\mu_t - L(y_t(h_t)) - cq_t(h_t)] - \Pi_1(\mu_1) \right] \\ &= \mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1}[p\mu_t - L(y_t(h_t)) - c(y_t(h_t) + \mu_t - x_t(h_{t-1})))] - \Pi_1(\mu_1) \right].\end{aligned}$$

Replacing $\Pi_1(\mu_1)$ with the right-hand side of Eq. (12) and using Myerson's change of order of integration, we obtain that the manufacturer's profit is equal to $(p-c)\mu^*/(1-\delta) - p\mathbf{E}\left[\frac{\overline{F}(\mu_1)}{f(\mu_1)}\right] - \Pi_1(\underline{\mu})$ plus

$$\mathbf{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} \left[-L(y_t(h_t)) - c(y_t(h_t) - x_t(h_{t-1})) - \frac{\overline{F}(\mu_1)}{f(\mu_1)}[-p + (p+h)G(y_t(h_t))]1_{\{y_i(h_i)>\varepsilon_i, i=1,2,\ldots,t-1\}} \right] \right], \tag{13}$$

We now optimize the safety stock levels $y_t(h_t)$ in the equation above by using a dynamic programming (DP) approach. The corresponding DP formulation is as follows. Take any period t. Let x be the starting inventory level at the beginning of period t. Let $\pi(x)$ be the profit-to-go function, given that the starting inventory level in every period prior to period t is always strictly positive. Let $\widehat{\pi}(x)$ be the profit-to-go function, assuming that a stockout has occurred prior to period t. It follows from Eq. (13) that

$$\pi(x) = \max_y \left\{ -L(y) - c(y-x) - \frac{\overline{F}(\mu_1)}{f(\mu_1)}[-p + (p+h)G(y)] + \delta \left[\int_{\underline{\varepsilon}}^{y} \pi(y-\varepsilon)g(\varepsilon)d\varepsilon + \widehat{\pi}(0)\overline{G}(y) \right] \right\}, \tag{14}$$

where the fact that the profit-to-go function for the subsequent period changes from $\pi(x)$ to $\widehat{\pi}(0)$ if a stockout occurs in the current period is captured by the last term, and

$$\widehat{\pi}(x) = \max_y \left\{ -L(y) - c(y-x) + \delta \mathbf{E}_\varepsilon \widehat{\pi}((y-\varepsilon)^+) \right\}, \tag{15}$$

where we ignore the nonnegative constraints for the order quantity (which are satisfied under the optimal policy).

Note that the DP problem in Eq. (15) is the same as the classical infinite-period inventory problem with lost sales (see, e.g., Karlin [1958]). Therefore, if a stockout has occurred, then the optimal policy is to keep the safety stock at a constant level $\widehat{y}^*$ from then on, where $\widehat{y}^*$ satisfies

$$-L'(\widehat{y}^*) - c(1 - \delta G(\widehat{y}^*)) = 0.$$

Following the same approach as that in Karlin [1958], we can solve the DP problem in Eq. (14) and show that if a stockout has never occurred before, then the optimal policy is to keep the safety stock at a constant level $y^*(\mu_1)$, where $y^*(\mu_1)$ satisfies

$$-L'(y^*(\mu_1)) - c(1 - \delta G(y^*(\mu_1))) - \left[\frac{\overline{F}(\mu_1)}{f(\mu_1)}(p+h) + \delta(\widehat{\pi}(0) - \pi(0)) \right] g(y^*(\mu_1)) = 0.$$

The corresponding ordering policy is $\{q_t^r(h_t)\}$ where $q_1^r(\mu_1) = \mu_1 + y^*(\mu_1)$, $q_t^r(h_t) = \mu_t + \varepsilon_{t-1}$ for $t \geq 2$ if a stockout has never occurred before t; $q_t^r(h_t) = \widehat{y}^* - x_t + \mu_t$ for $t \geq 2$ if a stockout has occurred before t.

To summarize, the manufacturer's optimal inventory policy is to keep a safety stock level at $y^*(\mu_1)$ in every period for any μ_1 when no stockout has ever happened before and to increase the safety stock level to $\widehat{y}^*$ after the first stockout event occurs. The payment function can be determined by ensuring that the retailer's profit is $\Pi_1(\mu_1)$ with $\Pi_1(\underline{\mu}) = 0$ and the payment is zero in every subsequent period. Note that $q_1^r(\mu_1)$ increases in μ_1, which is sufficient to ensure that the contract $\{(q_t^r(\cdot), T_t^r(\cdot))|_{t\in\mathbb{N}}\}$ described above satisfies the incentive compatibility constraint in the first period, and hence solves the relaxed problem.

2. Solution to the original problem. We construct the following menu:

$$w^*(\mu_1) = [p - (p+h)G(y^*(\mu_1))]/(1 - \delta G(y^*(\mu_1)))$$

and

$$T^*(\mu_1) = T_1^r(\mu_1) - (\mu_1 + y^*(\mu_1))w^*(\mu_1) - \delta w^*(\mu_1)\mu^*/(1-\delta).$$

Suppose the retailer has chosen the wholesale price $w^*(\mu_1)$ in the first period. Consider any period $t \geq 2$. Let x be the retailer's ending inventory in period t. The retailer has two options. One is to stick to the wholesale price $w^*(\mu_1)$, and the other is to pay the manufacturer the fixed fee $T_0^*(\mu_1)$ (to be specified below) to lower the wholesale price to the production cost c. Under the former, the retailer's expected profit-to-go function, denoted by $\Pi(x, \mu_1)$, is equal to $\Pi(0, \mu_1) + w^*(\mu_1)x$. This is because the retailer with the starting inventory x needs to order a positive quantity to bring the safety stock to $y^*(\mu_1)$ and, in comparison, the retailer with zero starting inventory needs to order x units more with the per unit order cost $w^*(\mu_1)$ to bring the safety stock to $y^*(\mu_1)$. Similarly, under the latter, the retailer's expected profit-to-go function, denoted by $\widehat{\Pi}(x, \mu_1)$, is equal to $\widehat{\Pi}(0, \mu_1) + cx$. We can set $T_0^*(\mu_1)$ so that $\Pi(0, \mu_1) = \widehat{\Pi}(0, \mu_1)$, implying that the retailer becomes indifferent between exercising the option and not when she stocks out. Because $w^*(\mu_1) > c$, we have $\Pi(x, \mu_1) > \widehat{\Pi}(x, \mu_1)$ for $x > 0$, implying that it is a best response for the retailer with an ending inventory x to stick to the wholesale price $w^*(\mu_1)$ if $x > 0$ and to exercise the option and pay the fee $T_0^*(\mu_1)$ if $x = 0$.

Now consider the (IC) constraints for $t = 1$. If the retailer with a true demand forecast μ_1 were to report a demand forecast $\widehat{\mu}_1$, then her best response from then onwards is to keep a safety stock of $y^*(\widehat{\mu}_1)$ until a stockout occurs, at which point she will exercise the stockout option. Therefore, the type μ_1 retailer's expected total discounted profit by choosing the contract intended for type $\widehat{\mu}_1$ is

$$\Pi_1(\mu_1, \widehat{\mu}_1) = \Pi_1(\widehat{\mu}_1) + (p - w^*(\widehat{\mu}_1))(\mu_1 - \widehat{\mu}_1),$$

implying that

$$\begin{aligned}
\frac{\partial \Pi_1(\mu_1, \widehat{\mu}_1)}{\partial \widehat{\mu}_1} &= \Pi_1'(\widehat{\mu}_1) - (p - w^*(\widehat{\mu}_1)) - [w^*(\widehat{\mu}_1)]'(\mu_1 - \widehat{\mu}_1) \\
&= p + \frac{-p + (p+h)G(y^*(\widehat{\mu}_1))}{1 - \delta G(y^*(\widehat{\mu}_1))} - (p - w^*(\widehat{\mu}_1)) - [w^*(\widehat{\mu}_1)]'(\mu_1 - \widehat{\mu}_1) \quad \text{(by Eq. (12))} \\
&= -[w^*(\widehat{\mu}_1)]'(\mu_1 - \widehat{\mu}_1) \text{ (by definition of } w^*(\cdot))
\end{aligned}$$

which is nonnegative for $\mu_1 \geq \widehat{\mu}_1$ and nonpositive for $\mu_1 \leq \widehat{\mu}_1$ because $[w^*(\widehat{\mu}_1)]' \leq 0$. Therefore, it is in the best interest of type μ_1 retailer to select the contract $w^*(\mu_1)$, implying that $\{w^*(\mu_1), T^*(\mu_1)\}$ satisfies the (IC) constraint in period 1. Clearly, the retailer's ordering quantity decisions under $\{w^*(\mu_1), T^*(\mu_1), T_0^*(\mu_1)\}$ are the same as those under the optimal direct truth-telling mechanism that solves the relaxed problem, and so is the manufacturer's expected total discounted profit. Therefore, the constructed menu solves the original problem. $\square$

Proof of Proposition 4. We follow the proof procedure used to prove Propositions 1 and 2, and focus on the modifications. Specifically, in Step 1,

$$\Pi_1'(\mu_1) = p + \mathbf{E}\left[\sum_{t=1}^{N} \delta^{t-1}[-b + (b+h)G(y_t(h_t))] + \ \delta^{N-1}(-p + (p-v)G(y_N(h_N)))] \middle| \ \mu_1\right].$$

In step 2, the manufacturer's objective function needs to be modified to the following

$$\begin{aligned}
&\mathbf{E}\left[\sum_{t=1}^{N-1} \delta^{t-1}[p\mu_t - L(y_t(h_t)) - cq_t(h_t)] + \ \delta^{N-1}[p\mu_N - \widetilde{L}(y_N(h_N)) - cq_N(h_N)] - \Pi_1(\mu_1)\right] \\
&= \mathbf{E}\left[\begin{array}{c} \sum_{t=1}^{N-1} \delta^{t-1}\left[-L(y_t(h_t)) - c(1-\delta)y_t(h_t) - \frac{\overline{F}(\mu_1)}{f(\mu_1)}[-b + (b+h)G(y_t(h_t))]\right] \\ + \ \delta^{N-1}\left[-\widetilde{L}(y_N(h_N)) - cy_N(h_N) - \frac{\overline{F}(\mu_1)}{f(\mu_1)}[-p + (p-v)G(y_N(h_N))]\right] - \frac{\overline{F}(\mu_1)}{f(\mu_1)}p \end{array}\right] \\
&\quad + \frac{(p-c)\mu^*}{1-\delta} - \Pi_1(\underline{\mu}).
\end{aligned}$$

We can use pointwise optimization to obtain the optimal safety stock: $y_t(h_t) = y^*(\mu_1)$ for any given h_t and $t = 1, 2, ..., N-1$, where $y^*(\mu_1)$ satisfies

$$b - (b+h)G(y^*(\mu_1)) - (1-\delta)c - \frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y^*(\mu_1)) = 0$$

and $y_N(h_N) = \widetilde{y}^*(\mu_1)$ for any given h_N, where $\widetilde{y}^*(\mu_1)$ satisfies

$$p - (p-v)G(\widetilde{y}^*(\mu_1)) - c - \frac{\overline{F}(\mu_1)}{f(\mu_1)}(p-v)g(\widetilde{y}^*(\mu_1)) = 0.$$

The nonnegative order quantity constraints are satisfied because $\widetilde{y}^*(\mu_1) > y^*(\mu_1)$ for any given μ_1 and the nonnegative demand assumption $\mu_t + \varepsilon_{t-1} \geq 0$. The base stock policy remains optimal, with the optimal safety stock level being $y^*(\mu_1)$ for the first $N-1$ periods and then increasing to $\widetilde{y}^*(\mu_1)$ in the last period. By following the remainder of the proofs of Propositions 1 and 2, we can establish the result that the optimal long-term dynamic contract takes the form of a menu of contracts $\{w_1^*(\mu_1), w_2^*(\mu_1), ..., w_N^*(\mu_1), T^*(\mu_1)\}$, where the wholesale prices $\{w_1^*(\mu_1), w_2^*(\mu_1), ..., w_N^*(\mu_1)\}$ are set to ensure that the retailer observing demand forecast μ_1 would optimally follow the base stock policy with the optimal safety stock level being $y^*(\mu_1)$ for the first $N-1$ periods and then increasing to $\widetilde{y}^*(\mu_1)$ in the last period N. $\square$

Appendix B: Finite Horizon

In this section, we consider a finite horizon model with time periods indexed by $1, 2, ..., N$, whereby the unfulfilled orders in periods 1 to $N-1$ can be backlogged at a per unit penalty cost b but the unfulfilled orders in period N are lost. The retailer's leftover inventory at the end of period N can be salvaged at per unit value v.

We start with formulating the retailer's problem under a given contract $\{(q_t(\widehat{h}_t), T_t(\widehat{h}_t))|_{t=1,2,...N}\}$. Similar to the backlogging model with infinite time horizon, for any given true information h_t and the retailer's report $\widehat{h}_t$, the retailer's safety stock in period t satisfies the following recursive relation:

$$y_t(h_t, \widehat{h}_t) = y_{t-1}(h_{t-1}, \widehat{h}_{t-1}) + q_t(\widehat{h}_t) - \mu_t - \varepsilon_{t-1},$$

with $y_0 = \varepsilon_0 = 0$ for $t = 1, 2, ..., N$. Similarly, in every period $t = 1, 2, ..., N-1$, the retailer's maximum expected total discounted profit-to-go, denoted by $\Pi_t(h_t, \widehat{h}_t)$ satisfies the following recursive relation:

$$\Pi_t(h_t, \widehat{h}_t) = p\mu_t - L(y_t(h_t, \widehat{h}_t)) - T_t(\widehat{h}_t) + \delta \mathbf{E}_{\varepsilon_t, \mu_{t+1}} \max_{\widehat{\varepsilon}_t, \widehat{\mu}_{t+1}} \{\Pi_{t+1}(h_{t+1}, \widehat{h}_{t+1})\},$$

where $L(y) = \mathbf{E}_\varepsilon[h(y-\varepsilon)^+ + b(\varepsilon - y)^+]$. Contrast emerges for the retailer's expected profit in period N, denoted by $\Pi_N(h_N, \widehat{h}_N)$. If the safety stock $y_N(h_N, \widehat{h}_N)$ is larger than the forecast error ε_N in period N, then the leftover inventory $(y_N(h_N, \widehat{h}_N) - \varepsilon_N)^+$ is salvaged at value v; otherwise the unfulfilled orders $(\varepsilon_N - y_N(h_N, \widehat{h}_N))^+$ are lost without generating any revenue. Therefore, the retailer's expected profit in period N is

$$\Pi_N(h_N, \widehat{h}_N) = p\mu_N - \widetilde{L}(y_N(h_N, \widehat{h}_N)) - T_N(\widehat{h}_N),$$

where $\widetilde{L}(y) = \mathbf{E}_\varepsilon[p(\varepsilon - y)^+ - v(y - \varepsilon)^+]$.

Now we turn to the manufacturer's problem. Our procedure to characterize the manufacturer's optimal dynamic long-term contracts is similar to that in the infinite-time horizon model with backlogging. We first consider the relaxed problem. By following the proof of Proposition 1, the relaxed problem can be simplified to the following maximization problem

$$\max_{\{y_t(h_t)|t=1,2,...,N\}} \mathbf{E} \left[\begin{array}{c} \sum_{t=1}^{N-1} \delta^{t-1} \left[-L(y_t(h_t)) - c(1-\delta) y_t(h_t) - \frac{\overline{F}(\mu_1)}{f(\mu_1)} [-b + (b+h) G(y_t(h_t))] \right] \\ +\delta^{N-1} \left[-\widetilde{L}(y_N(h_N)) - c y_N(h_N) - \frac{\overline{F}(\mu_1)}{f(\mu_1)} [-b + (b+h) G(y_N(h_N))] \right] \end{array} \right]$$

subject to the nonnegative order quantity constraints, i.e., $y_t(h_t) - y_{t-1}(h_{t-1}) + \mu_t + \varepsilon_{t-1} \geq 0$ for $t = 2, 3, ...N$. Let $y^*(\mu_1)$ and $\widetilde{y}^*(\mu_1)$ be the unconstrained maximizers corresponding to $t \leq N-1$ and $t = N$, respectively. Clearly, $y^*(\mu_1)$ satisfies

$$b - (b+h)G(y^*(\mu_1)) - (1-\delta)c - \frac{\overline{F}(\mu_1)}{f(\mu_1)}(b+h)g(y^*(\mu_1)) = 0$$

and $\widetilde{y}^*(\mu_1)$ satisfies

$$p - (p-v)G(\widetilde{y}^*(\mu_1)) - c - \frac{\overline{F}(\mu_1)}{f(\mu_1)}(p-v)g(\widetilde{y}^*(\mu_1)) = 0.$$

It is verifiable that $y^*(\mu_1) \leq \widetilde{y}^*(\mu_1)$ if and only if the ratio between cost of understocking and cost of overstocking in period N is no less than that in the previous periods, i.e.,

$$(p-c)/(c-v) \geq (b-(1-\delta)c)/(h+(1-\delta)c). \tag{16}$$

Note that the above inequality holds when one of the three parameters $\{p, h, v\}$ is sufficiently large or b is sufficiently small.

Two cases emerge. In the one case, the inequality in Eq. (16) holds. In this case, the nonnegative order quantity constraints are satisfied for any demand history, implying that in the relaxed problem, it is optimal for the manufacturer to induce the retailer to follow the base stock policy with a constant safety stock level $y^*(\mu_1)$ in the first $N-1$ periods and a higher safety stock level $\widetilde{y}^*(\mu_1)$ in the last period N. It then follows from the proof of Proposition 1 that for the original problem, it is optimal for the manufacturer to offer a menu of contracts, each consisting of wholesale prices for every period and fixed upfront payments, under which it is optimal for the retailer with demand forecast μ_1 to follow the base stock policy described in the solution to the relaxed problem. This leads to the following proposition.

Proposition 4. *If the inequality in Eq. (16) holds, then under the optimal mechanism of the relaxed problem, the retailer with demand forecast μ_1 follows a base stock policy with safety stock level $y^*(\mu_1)$ in period $t = 1, 2, ..., N-1$ and with safety stock level $\widetilde{y}^*(\mu_1)$ in period N. In this case, the optimal long-term dynamic contract of the original problem takes the form of a menu of contracts $\{w_1^*(\mu_1), w_2^*(\mu_1), ..., w_N^*(\mu_1), T^*(\mu_1)\}$ specifying wholesale prices for periods $t = 1, 2, ..., N$ and fixed upfront payments, respectively.*

In case the inequality in Eq. (16) does not hold, the nonnegative order quantity constraint in period N cannot be ignored, implying that the optimal safety stock in period N is $y_N(h_N) = \max(y_{N-1}(h_{N-1}) - \mu_N - \varepsilon_{N-1}, \widetilde{y}^*(\mu_1))$. This in turn implies that the safety stock $y_{N-1}(h_{N-1})$

influences the manufacturer's profits in periods $N-1$ and N. Although the manufacturer's profits in each period is a unimodal function of the safety stock, the sum of two unimodal functions may not be unimodal any more, implying that the manufacturer's profits as a function of $y_{N-1}(h_{N-1})$ may have multiple local maximizers. This undermines the optimality of the base stock policy in period $N-1$ for the relaxed problem, which in turn undermines the optimality of the menu of wholesale price contracts with upfront payments for the original problem.

From the above discussion, a key driver behind the optimality of wholesale price contracts with upfront payments is that the inventory ordering policy under the solution to the relaxed problem is of the base stock policy possibly with distinct safety stock levels in each period. If the base stock policy is optimal for the relaxed problem, then one can properly construct a menu of wholesale prices (possibly varying over time) together with upfront payments to induce the retailer with private demand and inventory information arising dynamically to follow the same base stock policy as that of the relaxed problem, thereby optimally solving the original problem. In contrast, if the base stock policy is no longer optimal for the relaxed problem, then the wholesale price contracts together with upfront payments would induce the retailer to follow a base stock policy in the original problem, which deviates from the optimal inventory policy under the relaxed problem, thereby undermining the optimality result of wholesale price contracts with upfront payments for the original problem. An implication of this insight is that the optimality result continues to hold if μ_t is drawn from a different distribution $F_t(\cdot)$ but all ε_t's are drawn from the same distribution $G(\cdot)$ or from a series of distributions $G_t(\cdot)$ that stochastically increase in t.

Appendix C: Early Termination

In our two basic models, we assume the contract requires an individual rationality (IR) constraint to hold in the first period in order to ensure that the retailer agrees to sign a long-term contract regardless of her initial demand forecast. We did not, however, assume the retailer has the option to end the contractual relationship at any time of her choice. In this section, we add this additional requirement to our contract design problem and study whether it changes the manufacturer's profit or the format of the optimal long-term contract.

To study this problem, we need to model what occurs with backlogged and leftover inventory if the relationship is terminated. We assume that if the contract is terminated at time t, unfulfilled orders are cancelled with payments returned to customers and unused inventory has salvage value v. Therefore, the retailer's profit when breaking the contract with inventory $x_t(h_t, h_{t-1})$ is equal to $p\min\{0, x_t(h_t, h_{t-1})\} + v\max\{0, x_t(h_t, h_{t-1})\}$. In order to guarantee dynamic participation, we add the following additional constraints to our backlogging model:

$$\Pi_t(h_t, h_t) \geq p \min\{0, x_t(h_t, h_{t-1})\} + v \max\{0, x_t(h_t, h_{t-1})\} \quad \text{for all } t \text{ and } h_t. \tag{17}$$

Only termination from a truthful history needs to be verified since the dynamic incentive compatibility constraints already ensure that reporting truthfully is weakly better for the retailer than reporting untruthfully.

The following proposition states that dynamic participation constraints are either immediately satisfied by the optimal contract, which occurs when the retail price is sufficiently high, or can be easily satisfied via a rearrangement of payments. That is, by moving some payments into the future, the manufacturer can ensure that the retailer does not want to quit the relationship at any point in time, regardless of the history. The manufacturer can also ensure dynamic participation is satisfied simply by charging an early termination fee equal or greater than the cost-to-go associated with the worst possible history. In equilibrium, the retailer would never choose to break the contract.

Proposition 5. *There exists a threshold $\underline{p}$ such that if the retail price $p \geq \underline{p}$, then the optimal contract from Proposition 2 satisfies dynamic participation constraints. Otherwise, the manufacturer should add to the optimal contract a fixed payment $Q \geq 0$ to the retailer in every period $t \geq 2$, and increase the upfront fee by $\delta Q/(1-\delta)$ in order to ensure dynamic participation. Alternatively, the manufacturer could add to the optimal contract an early termination fee equal to $Q/(1-\delta)$ to the contract in order to satisfy dynamic participation constraints.*

Proof. Using Eq. (3) , we can rewrite Eq. (17) as

$$\Pi_t(h_t, h_t) = p\mu_t - L(y_t(h_t, h_t)) - T_t(h_t) + \delta[\Pi(h_{t+1}, h_{t+1})] \geq vx_t(h_t, h_{t-1})^+ + px_t(h_t, h_{t-1})^-$$

for all t and h_t. Under the optimal long-term contract and under truthful behavior, the inventory, safety stock, payment function and value-to-go of the retailer will be respectively $x_t(h_t, h_{t-1}) = y^*(\mu_1) - \varepsilon_{t-1}$, $y_t(h_t, h_t) = y^*(\mu_1)$, $T_t(h_t) = w^*(\mu_1)(\mu_t + \varepsilon_{t-1})$ and

$$E[\Pi_t(h_{t+1}, h_{t+1})] = \frac{(p - w^*(\mu_1))\mu^* - L(y^*(\mu_1))}{1-\delta}.$$

Plugging these values in and moving the right-hand side to the left-hand side, we obtain the constraint

$$p\mu_t - w^*(\mu_1)(\mu_t + \varepsilon_{t-1}) + \frac{\delta(p - w^*(\mu_1))\mu^*}{1-\delta} - \frac{L(y^*(\mu_1))}{1-\delta} - v(y^*(\mu_1) + \varepsilon_{t-1})^+ + p(y^*(\mu_1) - \varepsilon_{t-1})^- \geq 0$$

for all $\mu_1, \mu_t, \varepsilon_{t-1}$. The left-hand side of the inequality above is continuous in μ_1, μ_t and ε_{t-1}, which are variables that belong to bounded sets. Furthermore, the left-hand side of the inequality is increasing in p. Therefore, if p is sufficiently high, the inequality above is satisfied for all μ_1, μ_t and ε_{t-1} completing the first part of the proof.

Now suppose the minimum of the left-hand side of the inequality over the space $\mu_1 \in [\underline{\mu}, \overline{\mu}]$, $\mu_t \in [\underline{\mu}, \overline{\mu}]$ and $\varepsilon_{t-1} \in [\underline{\varepsilon}, \overline{\varepsilon}]$ above is a negative quantity K. By continuity and boundedness of the decision space, K is finite. Then, we need to charge an early termination fee equal to or greater than $-K$ in order to ensure dynamic participation.

Alternatively, let $Q = -(1-\delta)K$. Consider a modified contract where a payment equal to Q is added by the manufacturer to the retailer in every period starting from $t = 2$. This will increase the retailer's value-to-go in every period after $t = 2$ by $\sum_{t=1}^{\infty} \delta^{t-1} Q = Q/(1-\delta) = -K$. With these additional payments, the retailer's dynamic participation constraints will be satisfied. The manufacturer's profit will not change if he adds a charge equal to $\delta Q/(1-\delta)$ to the upfront payment. The discounted sum of all payments will equal to 0. Therefore, this change does not affect the first period individual rationality constraint. $\square$