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

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

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

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

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 $cy_t(h_t)$ 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

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

The maximand is strictly unimodal because its first order derivative is

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

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'). □

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

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

where $1_{\{y_i(h_i)>\varepsilon_i, i=1,2,\dots,t-1\}}$ 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

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

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

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

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

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

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. □