Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

A Mechanism Design Approach to Vendor Managed Inventory

Bharadwaj Kadiyala, Özalp Özer, Alain Bensoussan

To cite this article:
Bharadwaj Kadiyala, Özalp Özer, Alain Bensoussan (2019) A Mechanism Design Approach to Vendor Managed Inventory. Management Science

Published online in Articles in Advance 22 Oct 2019

. https://doi.org/10.1287/mnsc.2019.3297

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–25
ISSN 0025-1909 (print), ISSN 1526-5501 (online)
http://pubsonline.informs.org/journal/mnsc

A Mechanism Design Approach to Vendor Managed Inventory

Bharadwaj Kadiyala,[a] **Özalp Özer,**[b] **Alain Bensoussan**[b]

[a] School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong;
[b] Naveen Jindal School of Management, The University of Texas at Dallas, Richardson, Texas 75080

Contact: bkadiyala@ust.hk, https://orcid.org/0000-0003-0603-1054 (BK); oozer@utdallas.edu,
https://orcid.org/0000-0001-5552-0968 (ÖÖ); alain.bensoussan@utdallas.edu (AB)

Received: October 27, 2015
Revised: February 4, 2018; August 30, 2018
Accepted: November 25, 2018
Published Online in Articles in Advance: October 22, 2019

https://doi.org/10.1287/mnsc.2019.3297

Abstract. This paper studies an inventory management problem faced by an upstream supplier that is in a collaborative agreement, such as vendor-managed inventory (VMI), with a retailer. A VMI partnership provides the supplier an opportunity to manage inventory for the supply chain in exchange for point-of-sales (POS)- and inventory-level information from the retailer. However, retailers typically possess superior local market information and as has been the case in recent years, are able to capture and analyze customer purchasing behavior beyond the traditional POS data. Such analyses provide the retailer access to market signals that are otherwise hard to capture using POS information. We show and quantify the implication of the financial obligations of each party in VMI that renders communication of such important market signals as noncredible. To help institute a sound VMI collaboration, we propose *learn and screen*—a dynamic inventory mechanism—for the supplier to effectively manage inventory and information in the supply chain. The proposed mechanism combines the ability of the supplier to learn about market conditions from POS data (over multiple selling periods) and dynamically determine when to screen the retailer and acquire his private demand information. Inventory decisions in the proposed mechanism serve a strategic purpose in addition to their classic role of satisfying customer demand. We show that our proposed dynamic mechanism significantly improves the supplier's expected profit and increases the efficiency of the overall supply chain operations under a VMI agreement. In addition, we determine the market conditions in which a strategic approach to VMI results in significant profit improvements for both firms, particularly when the retailer has high market power (i.e., when the supplier highly depends on the retailer) and when the supplier has relatively less knowledge about the end customer/market compared with the retailer.

History: Accepted by Gad Allon, operations management.
Funding: B. Kadiyala and Ö. Özer acknowledge financial support from the National Science Foundation [Grant 1644935]. A. Bensoussan acknowledges financial support from the National Science Foundation [Grant 1612880], the Research Grants Council of Hong Kong [Grants 500113 and 11303316], and the City University of Hong Kong.
Supplemental Material: The e-companion is available at https://doi.org/10.1287/mnsc.2019.3297.

Keywords: value of information • vendor-managed inventory • Bayesian inventory management • incentives of information sharing • Dynamic programming • mechanism design with endogenous learning • censored demand

1. Introduction

Practitioners and scholars have long shown that centralized inventory management together with information sharing, such as in vendor-managed inventory (VMI), allows supply chains to be efficient and responsive to customer needs (see, for example, Aviv 2007 and Simchi-Levi et al. 2008 for an extensive review of this literature). However, recent empirical and anecdotal evidence also suggests that VMI-type agreements have proven difficult to maintain over multiple planning horizons, resulting in companies terminating such agreements (e.g., Kouvelis et al. 2006, Brinkhoff et al. 2015). One frequently cited reason for such failed relationships has been the decline of trust (in terms of both credibility and capability) among firms implementing VMI. For such dynamic settings, this paper first determines key reasons for why VMI relationships fail and then provides a mechanism to reinforce and better manage an ongoing VMI agreement that evolves over time.

Aforementioned breakdown among collaborating firms often manifests itself in the following fashion. Consider a supplier (vendor) and a retailer who operate in a VMI agreement. Under this agreement, the supplier (she) takes the sole responsibility, including financial and operational control, of inventory in the supply chain.[1] The retailer (he) takes the responsibility of store-level execution to satisfy end customer demand as much as possible. The retailer uses information technology, such as electronic data interchange (EDI), to share customer sales information through point-of-sales (POS) data and inventory levels with the supplier at the end of each selling period. The POS data help the supplier adapt to dynamic market conditions and improve her inventory replenishment process over time. However, the retailer also obtains new demand information regarding customers' taste and demand because of his proximity and close relationship (e.g., through loyalty cards) with his customers. Such information is not immediately available in POS data. Nevertheless, the supplier could improve her forecasts and hence, her inventory decisions by learning about the retailer's private demand information using POS data. However, relying on the learning process poses an important challenge for the supplier; such a learning process may take too long to be effective, and also, POS data often do not contain information, such as unobserved lost sales. The leftover inventory is the supplier's (and not the retailer's) liability. Hence, the retailer is always better off depicting a positive outlook of the market and reporting higher levels of demand (as well as lost sales if they were observable) to induce the supplier to ship sufficient inventory during all selling periods. Knowing this incentive, the supplier may disregard such information even when the retailer accurately reports his information. The resulting relationship, therefore, often boils down to only exchanging POS data and the supplier learning about the market through POS data. However, such learning takes time, during which *unobserved* lost sales result in the supplier carrying even less inventory, leading to additional lost sales for the retailer. Such lack of coordination diminishes the retailer's patience for having to relinquish control of his inventory, leading to a "lose-lose" outcome for both firms and resulting in termination of VMI agreements.[2]

Retailers have been forming collaborative relationships with their suppliers in an effort to make their operations lean while maintaining high customer service. For example, Walmart and Procter & Gamble, Tesco and Nestle, and Northern Foods and Sainsbury's are involved in VMI practices (see Lee et al. 1997, Watson 2005, The Grocer 2009). Centralizing inventory control by moving it up the supply chain closer to the source mitigates the well-known bullwhip effect (Lee et al. 1997). The flow of POS data upstream also improves demand forecasts, leading to fewer stockouts at the retail store. Technology companies, such as Dell and Apple, have managed to avoid selling through resellers by vertically integrating with the downstream. They, however, also practice VMI with their upstream suppliers (e.g., Katariya et al. 2014).

The classical VMI agreements, however, do not provide the supplier and the retailer with a means to credibly share demand information beyond the POS data.[3] Note that, after a VMI agreement is signed, the retailer can still obtain private demand information, because he has access to local market conditions (such as competitor's store openings/closures), in-store promotions, and his customers' information (e.g., obtained through loyalty/reward membership programs). All such information could be beneficial for the supplier in making effective replenishment decisions. However, the retailer may not be able to credibly share this information (because of aforementioned incentive conflict). The following excerpt provides a case in point.

> Sainsbury's discovered that a cereal brand called Grape-Nuts was worth stocking—despite weak sales—because the shoppers who bought it were extremely loyal to Sainsbury's and often big spenders. (Ferguson 2013)

Sainsbury's in this case may not want to share low-demand information with the supplier of Grape-Nuts if both firms are in a VMI-type relationship.

In such situations, the supplier may decide to dynamically learn about customer demand through the POS data and disregard retailer's private demand information. In the meantime, the supplier continues to maintain higher/lower than necessary inventory levels. Consequences of improperly stocked shelves could be dire. Recently, Rosenblum (2014) reported that Walmart, a pioneer in VMI, lost around $3 billion owing to out-of-stock items. Spartan Stores ended its VMI program a year after its inception, citing the supplier's inability to take into account promotional events at the retail store (Mathews 1995). In fact, after taking off in early 1990s, VMI practices faced tough opposition in industry resulting from frictions between the supply chain members participating in it. However, the recent boom in information technology (IT) infrastructure developments and widespread use of data analytics in business will likely lead to an increase in the desire to have VMI-type collaborations. For such partnerships to have sustained success and for VMI to deliver on its promise, reexamining and aligning incentives of firms under such agreements are imperative.

The above observations motivate us to study the following questions. How should the supplier dynamically and optimally manage centralized inventory over multiple selling periods when lost sales are unobserved and the retailer has private demand information? Can the supplier use her inventory decisions to gain long-term leverage with the retailer? What mechanism can the supplier use in an ongoing VMI agreement to credibly elicit demand information from the retailer while effectively managing inventory over a planning horizon? To address these questions, we propose a learn-and-screen approach that effectively combines dynamic inventory management with mechanism design.

In the *learn-and-screen* approach, the supplier and the retailer sign a VMI agreement and agree on a wholesale price. The wholesale price could be the outcome of a bargaining process that may reflect each firm's market power and the retailer's informational advantage. The supplier then owns and replenishes inventory at the retailer's site, and the retailer pays the supplier when a product is sold (i.e., the supplier manages inventory periodically under unobserved lost sales). During the rest of the agreement horizon, the supplier can continue to improve her belief about the market demand (including her belief about the retailer's private demand information) by *learning* through POS data. Alternatively, at an optimal time during the VMI relationship, the supplier can offer a menu of *screening* contracts to the retailer. The menu of contracts offers the retailer base stock levels to choose from for the supply chain in exchange for a lump sum (or equivalently, per period) payment (to or from the supplier). If the retailer agrees and picks a base stock level, the supplier continues to manage the inventory under unobserved lost sales following this chosen base stock level. If the retailer rejects and does not select any base stock level, the prevailing supply chain ensues. In other words, the supplier continues to make replenishment decisions and reverts to the status quo VMI agreement with a base stock policy and level that capture the supplier's updated information learned up to that point.

We remark that our proposed approach also endogenously accounts for both firms' market power (other than the wholesale price). On one hand, the supplier can mitigate the retailer's informational advantage by learning from POS data before offering the screening contracts. The supplier can optimally determine the time to offer the menu of (screening) contracts as a function of her belief about the market condition and the on-hand inventory level. On the other hand, the retailer's outside option (i.e., the status quo VMI agreement) also improves as the supplier obtains information through POS data. The retailer's outside option (i.e., when any such screening contract is rejected) is no longer just exogenous (e.g., a constant) to the system, but instead, it depends on the retailer's ongoing agreement with the supplier and how this relationship evolves. Hence, the learning process and optimal screening mechanism are coupled (i.e., affect each other) and evolve over time as the supplier improves her belief using the POS data.

The rest of the paper formalizes this process and shows that our proposed mechanism outperforms the status quo VMI agreements as well as a simple static screening mechanism. We also characterize how the optimal screening mechanism evolves over time and how learning effects the structure of this mechanism/contract. We quantify how unobserved lost sales impact the learning process as well as the screening contracts. In addition, we compare the performance of the learn-and-screen approach with the learn (only) and screen (immediately) approaches. We find that both the supplier and the retailer are better off from using the learn-and-screen approach (creating a win-win outcome). This approach also lowers the on-hand inventory levels maintained by the supplier compared with the learn (only) approach. Yet, the supplier makes more profit with the lower inventory levels by adopting the learn-and-screen approach. In this sense, learning and screening have a synergistic effect on each other.

2. Literature Review

Many operations management scholars have explored and documented the benefits that accrue from the practice of information sharing in supply chains (e.g., Cachon and Fisher 2000, Lee et al. 2000, Aviv 2001, Gallego and Özer 2001, Ren et al. 2010, Ha et al. 2011, Dong et al. 2014, Shang et al. 2016). The value of information sharing, in particular, of demand forecasts within the supply chain has been shown to play an important role in determining success of collaborative partnerships, such as VMI and collaborative planning, forecasting, and replenishment to name a few (e.g., Aviv 2002, 2007; Chen and Lee 2009; Brinkhoff et al. 2015). Aviv (2007) shows that supply chain characteristics, such as the retailer's ability to observe superior market signals and the supplier's agility in production, contribute to a win-win situation in a collaborative forecasting partnership. We note that these partnerships improve visibility of POS information upstream (which is verifiable) and/or centralizing replenishment processes. Improved visibility of demand in turn helps the supplier resolve some of demand variability over the planning horizon, albeit rather slowly. The question of whether valuable information that is private and unverifiable, such as the retailer's subjective assessment of demand, can be credibly shared is a natural extension to this line of investigation.

Researchers have also provided several contractual remedies to alleviate the credibility issue that may arise when self-interested firms report demand forecasts. Cachon and Lariviere (2001) consider capacity decisions under demand information asymmetry. They provide some properties of an incentive mechanism to enable credible information sharing and show the existence of separating equilibria in a signaling game. Özer and Wei (2006) model the forecast sharing game using both a screening model and a signaling model in a unified context which enables them to design and compare several contracts and choose the most effective contract for different market conditions. To do so, they first show that, used alone, wholesale price contract is a reason for distorted forecast. They then design several contracts such as the capacity reservation contracts (with quantity discounts) and advance purchase contracts that enable credible information sharing. Several scholars have since extended this line of research and consider various important and relevant supply chain settings in strategic single-period interactions with many practical implications (e.g., Burnetas et al. 2007, Li and Zhang 2008, Babich et al. 2012, Gümüş 2014, Li et al. 2014, Nazerzadeh and Perakis 2016 and Shamir and Shin 2016 and references therein). We extend this stream of literature by considering a *multiperiod* inventory model in which the supplier improves her demand forecasts over time by incorporating historical POS data.

More recently, scholars started to consider solutions for the aforementioned incentive problem in dynamic supply chain settings. Most papers in this line of work primarily consider what is often referred to as dynamic contracts. Dynamic refers to the possibility of the contract terms being history dependent, in which the underlying state variables evolve in a stochastic fashion. Thus, the contract terms specify outcomes for any possible evolution of the state variables. The principal could offer and commit to long-term contracts at the beginning of the planning horizon (Lutze and Özer 2008, Zhang and Zenios 2008, Lobel and Xiao 2017). Alternately, when the principal lacks the commitment power, short-term contracts can be redesigned in each period based on the updated information (Zhang et al. 2010). In this stream of literature, the proposed contract is often exogenous to the overall problem. In other words, the principal's or the agent's actions both before and after the menu of contracts is offered do not affect the structure of the menu. The menu of contracts also does not depend on the information revelation before the time that the menu of contracts is offered. In addition, most papers in this literature assume that, if the agent was to reject the contract, the remaining problem for each firm is exogenous to the mechanism design. Such a modeling assumption means that, if the agent rejects the offer, then the principal and the agent part ways (i.e., do not engage further) and that their follow-up actions (including the rejection of the contract) have no consequence to the menu of contracts offered at the beginning. As a result, the design of the menu of contracts does not endogenously depend on what each party does during the planning horizon. Nevertheless, this literature contributes to our understanding of these problems in dynamic settings. It also provide actionable policies (contract mechanisms) that firms can follow to effectively manage supply chains in corresponding dynamic settings.

A few papers, however, represent a significant departure from the aforementioned group of papers and bring the literature one step closer to the more general dynamic supply chain settings. Specifically, the proposed solutions to the mechanism design problems depend on the principal's and/or agent's actions both before and/or after the mechanism is executed as well as any information revelation during the planning horizon. Oh and Özer (2013) propose a general framework to model multiple evolutions of forecasts generated by multiple firms. Using this framework, they introduce the Martingale Model of Asymmetric Forecast Evolutions and propose a dynamic mechanism for a supplier to elicit a retailer's information credibly before making an irreversible capacity decision. The offered dynamic mechanism is endogenous to the system in that it depends on the evolution of asymmetric forecasts and the agent's actions after the menu of contract is offered as well as the ongoing relationship, even when the agent rejects the proposed contract. Feng et al. (2015) model a dynamic bargaining game between a buyer (with private demand information) and a seller that ensues before a one-time demand realization. The negotiation continues until an agreement on quantity and payment for the trade of a product is reached. In the process, the contract offers are updated by each party based on outcomes of the previous negotiation stages. Our paper also falls in this group in that we design and solve a dynamic mechanism problem in which the proposed mechanism/contract depends on (1) the retailer's (i.e., agent's) private demand information, (2) the supplier's inventory replenishment decisions (i.e., principal's actions) both before and after a mechanism is offered, and (3) the retailer's profit, which is a function of supplier's inventory policy even when the retailer rejects the mechanism. In addition, the optimal mechanism (i.e., the menu of contracts) also evolves depending on what the supplier learns through the retailer's POS data over time.

The dynamic nature of demand information asymmetry in our setting arises from the fact that the supplier updates her demand forecasts using the periodic POS data. The statistical evolution of demand forecasts has been modeled in literature using various approaches, such as time series (e.g., Aviv 2001, 2002, 2007), Martingale models of forecast evolution (e.g., Oh and Özer 2013), and Bayesian inference (e.g., Scarf 1959, 1960; Azoury 1985; Lovejoy 1990). We adopt the Bayesian approach. Because of demand censoring, forecast evolution in our problem resembles that of the unobserved lost sales Bayesian inventory problem (Lariviere and Porteus 1999, Chen and Plambeck 2008, Chen 2010, Bisi et al. 2011). A noteworthy aspect of our Bayesian forecast evolution model is that inventory decisions made by the supplier determine the extent of censoring of demand data in each sales period. Thus, the evolution of forecasts is *endogenized* through the supplier's inventory decisions. Our paper adds a new dimension to the classical Bayesian inventory problem. In particular, we show how this classical inventory problem impacts the contract/mechanism design problem and vice versa. We also show how to manage inventory under lost sales together with the strategic issue of credible demand information sharing in a decentralized supply chain.

3. The Model

Consider a supplier and a retailer who participate in a VMI agreement. The supplier is responsible for periodically producing and maintaining on-hand inventory over a planning horizon. At the beginning of a selling period n, the supplier decides how many units to produce and deliver at unit cost c to the retailer before consumer demand is realized. The retailer then satisfies demand to the extent possible from on-hand inventory. Unmet demand is *lost*, and neither the supplier nor the retailer observe lost sales. For every unit sold to customers, the retailer earns r and pays a wholesale price w to the supplier. The supplier, who is liable for the leftover inventory, incurs a unit holding cost h on the leftover inventory carried over to the next period. Appendix A provides a glossary of notation for an easy reference.

Demand in each period is independent and identically distributed with a cumulative distribution function (cdf) $G(z)$ and corresponding probability density function $g(z)$ for $z \geq 0$. Both the supplier and the retailer are uncertain about demand before its realization in each selling period. When the VMI agreement is negotiated and signed, both firms may possess the same demand information. In this case, the supplier faces the classical lost sales inventory control problem for which base stock policy is optimal (Karlin and Scarf 1958). Hence, maintaining VMI agreement over a long horizon would seem simple and possible, because the supplier only needs to maintain a stationary base stock level (that both firms can agree on at the time when the VMI agreement is signed) to replenish the retailer's inventory (in a stationary environment). However, over time, market conditions often change: for example, because of store promotions, new store opening/closure, or even changes in consumer tastes. Often, the retailer is in a better position to obtain additional demand information given his proximity to customers. Without loss of generality, we denote the period in which the retailer receives the additional information as the start of the planning horizon (i.e., $n = 1$). The retailer's demand information could contain useful market signals, which provide the retailer with an improved estimate of the average market size. Essentially, using this information, the retailer is able to accurately estimate a parameter ξ of the demand distribution. Larger ξ represents larger average demand (i.e., $\xi_1 \leq \xi_2$ implies that $G(z|\xi_1) \geq G(z|\xi_2)$ for all $z \geq 0$). Therefore, the retailer's demand information is composed of complete knowledge of the underlying demand distribution. The supplier, however, consolidates her prior demand information at the beginning of the planning horizon in the form of a belief $\pi_1 \in \mathscr{A}$, representing the pdf over $\Theta := [\underline{\xi}, \overline{\xi})$, the set of values that ξ takes. The set $\mathscr{A} := \{\pi : \Theta \to \mathbb{R}^+ | \int_\Theta \pi(\xi) d\xi = 1\}$ denotes the collection of all pdfs defined on the set Θ. The supplier can use, for example, sales information from previous selling periods and market research to develop the initial prior belief.

We define a probability space $(\Omega, \mathscr{L}, \mathbb{P})$ hosting the random variable ξ, the demand process D_n, and the sales observation process Z_n for $n \geq 1$. At the beginning of a selling period n, the supplier raises the on-hand inventory level from x_n to y_n.[4] Demand for that period D_n is then realized. The supplier only observes the POS information: that is, $Z_n := \min\{y_n, D_n\}$. At the beginning of period $n+1$, using historical sales information and the sales observation in period n, the supplier updates her belief about ξ using the Bayes' rule as follows:

$$\pi_{n+1}(\xi) = \mathbf{1}_{\{z_n = y_n\}} \cdot \overbrace{\frac{\overline{G}(y_n|\xi)\pi_n(\xi)}{\int_{\underline{\xi}}^{\overline{\xi}} \overline{G}(y_n|\eta)\pi_n(\eta)d\eta}}^{\pi^c_{n+1}(\xi|y_n)} + \mathbf{1}_{\{z_n < y_n\}} \cdot \overbrace{\frac{g(z_n|\xi)\pi_n(\xi)}{\int_{\underline{\xi}}^{\overline{\xi}} g(z_n|\eta)\pi_n(\eta)d\eta}}^{\pi^e_{n+1}(\xi)}, \quad n \geq 1, \tag{1}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function, π_1 is the supplier's initial belief, z_n is a realization of the random variable Z_n, and $\bar{G}(\cdot) = 1 - G(\cdot)$. The first term $\pi^c_{n+1}(\xi|y_n)$ in Equation (1) is the supplier's posterior belief when the demand realization in period n is greater than the on-hand inventory level (i.e., demand information is censored). The second term $\pi^e_{n+1}(\xi)$ is the supplier's posterior belief when the supplier observes the exact demand realization. We differentiate the notation of π^e_{n+1} from π^c_{n+1} to emphasize the dependence of the posterior distribution on the on-hand inventory level y_n, when the demand realization is *censored*.

The supplier's total expected profit under this VMI agreement is

$$\sum_{n=1}^{\infty} \alpha^{n-1} \mathbb{E}_{\xi, D(\xi)}\left[w \min\{y_n, D_n\} - c(y_n - x_n) - h(y_n - D_n)^+\right]$$
$$= \sum_{n=1}^{\infty} \alpha^{n-1} \mathbb{E}_{\xi, D(\xi)}\left[c x_n + (w-c) y_n - (w+h)\int_0^{y_n} Q_n(z) dz\right], \text{ where} \tag{2}$$

$$Q_n(z) := \int_0^z q_n(u)\, du \text{ and } q_n(z) := \int_\Theta g(z|\xi)\pi_n(\xi) d\xi \tag{3}$$

denote the posterior predictive distribution and the density of demand in period n, respectively. The expectation in the supplier's profit function is with respect to random demand and the unknown market signal ξ, which is the retailer's private information. Establishing a VMI-type partnership requires firms to set up costly IT infrastructure (e.g., to share POS data) and also reorganize workforce (e.g., to monitor inventory and manage the relationship). Thus, the decision to enter into a VMI agreement is typically a long-term one. Hence, we adopt an infinite horizon formulation to model this relationship.[5]

The retailer's total expected profit, given his demand information ξ, is

$$\sum_{n=1}^{\infty} \alpha^{n-1} \mathbb{E}_{D(\xi)}[(r-w)\min\{y_n, D_n\}] = (r-w)\sum_{n=1}^{\infty} \alpha^{n-1} \mathbb{E}_{D(\xi)}\left[y_n - \int_0^{y_n} G(z|\xi)\,\mathrm{d}z\right],$$

which is increasing in on-hand inventory level y_n. Therefore, the retailer has incentive to report optimistic demand to induce the supplier to allocate high on-hand inventory in each period.

3.1. The Learn-and-Screen Approach

We propose a learn-and-screen approach to help the supplier with the joint management of inventory, demand estimation, and incentive conflicts, all of which collectively arise in a VMI framework. The sequence of events in the learn-and-screen approach is as follows. At the beginning of each period $n \geq 1$, the supplier decides between learning about the retailer's private demand information from the POS data or offering a menu of screening contracts to the retailer. If the supplier opts for the former, she raises the on-hand inventory level up to y_n and updates (via Bayes' rule) her belief over ξ using the sales z_n observed in that period. The problem proceeds to the next period. Otherwise, the supplier offers a menu of contracts, $\{S(\xi|x_n, \pi_n), P(\xi|x_n, \pi_n)\}_{\xi\in\Theta}$. We note that these contracts are a function of the on-hand inventory level x_n and the updated posterior belief π_n. The retailer decides whether to accept a contract from this menu. If he accepts and chooses the contract $S(\hat{\xi}), P(\hat{\xi})$, the supplier procures inventory for the remaining planning horizon following the base stock level $S(\hat{\xi})$. That is, she produces enough to bring the on-hand inventory level up to $S(\hat{\xi})$ in each period. The retailer pays $P(\hat{\xi})$ to the supplier at the beginning of every period. Equivalently, the retailer pays a one-time lump sum after accepting one of the contracts. The retailer then satisfies the realized demand to the extent possible and makes a profit of $(r-w)$ for every unit sold; the supplier updates inventory, and this repeats next period. If the retailer rejects and does not choose any contract from this menu, the prevailing supply chain relationship ensues. The supplier continues to make replenishment decisions and does not offer another menu of contracts. Both firms make profit from the sales realized in each period. Essentially, rejection of the menu of contracts reverts the supply chain partnership back to the status quo VMI agreement with the base stock policy and level that captures the supplier's updated information learned up to that point.

We highlight three important benefits of the proposed learn-and-screen approach. (i) The ongoing VMI agreement between the firms is unaffected regardless of whether the retailer accepts one of the base stock levels in the menu. The ownership of inventory continues to remain with the supplier. We will show that the menu of contracts act as an instrument to facilitate credible communication of private demand information. In the event that the retailer rejects the menu of contracts, the ongoing VMI agreement ensues. (ii) The form of the contract is optimal (i.e., best among all possible forms), because the supplier faces the classical periodic review inventory control problem with lost sales after demand information is (and can be) credibly shared. For such an inventory problem, Karlin and Scarf (1958) have shown the optimality of base stock policy, thus justifying the contract terms. (iii) Monitoring the contract terms, after they are accepted, requires minimal effort. The supplier collects a one-time payment from the retailer, and the retailer periodically monitors the inventory level maintained by the supplier. Current VMI frameworks, such as PeopleSoft Enterprise Inventory and Fulfillment Management by Oracle, already implements this feature. For example, an automated message is delivered to the retailer as soon as inventory on his shelf is replenished (see Oracle 2009, p. 1040). Therefore, implementing the learn-and-screen approach does not alter the financial transactions between the supplier and the retailer after the period in which the contracts are offered (even if the retailer rejects them). As in the ongoing VMI agreement, in all of the subsequent periods, the retailer pays a unit wholesale price w for each sold in that period.

3.1.1. The Contract Design Problem. Suppose that the supplier offers the menu of contracts, $\{S(\cdot), P(\cdot)\}$, in period $n \geq 1$. Suppose also that a retailer with private demand information ξ (i.e., type ξ retailer for short) chooses the contract $S(\tilde{\xi}), P(\tilde{\xi})$ from this menu; the supplier then delivers inventory following the base stock level $S(\tilde{\xi})$. The supplier offers base stock levels that are at least as much as the on-hand inventory level x_n.[6] Thus, $y_j = S(\tilde{\xi})$, $j \geq n$, and the inventory evolves as follows:

$$x_{j+1} = [S(\tilde{\xi}) - D_j]^+, \quad j \geq n.$$

Given type ξ retailer's choice of the contract from the menu, the expected profits of the supplier, the retailer, and the total supply chain over the remaining horizon are given by, respectively,

$$\begin{aligned}
&\Pi^s(x, S(\tilde{\xi}), P(\tilde{\xi})) \\
&= \mathbb{E}_{D(\xi)}\left[\sum_{i=n}^{\infty} \alpha^{i-n}\left(w \min\{S(\tilde{\xi}), D_i\} - c(S(\tilde{\xi}) - x_i)\right.\right. \\
&\qquad \left.\left. - h(S(\tilde{\xi}) - D_i)^+ + P(\tilde{\xi})\right) \Big| x_n = x\right], \\
&\Pi^r(S(\tilde{\xi}), P(\tilde{\xi}), \xi) \\
&= \mathbb{E}_{D(\xi)}\left[\sum_{i=n}^{\infty} \alpha^{i-n}\left((r-w)\min\{S(\tilde{\xi}), D_i\} - P(\tilde{\xi})\right)\right], \text{ and} \\
&\Pi^{tot}(x, S(\tilde{\xi}), \xi) \\
&= \mathbb{E}_{D(\xi)}\left[\sum_{i=n}^{\infty} \alpha^{i-n}\left(r \min\{S(\tilde{\xi}), D_i\} - c(S(\tilde{\xi}) - x_i)\right.\right. \\
&\qquad \left.\left. - h(S(\tilde{\xi}) - D_i)^+\right) \Big| x_n = x\right].
\end{aligned}$$

Constant $\alpha \in [0, 1)$ is the discount factor. Each menu, $\{S(\cdot), P(\cdot)\}$, determines a Bayesian game in which the retailer chooses, in equilibrium, a contract that maximizes his total expected profit over the remaining planning horizon[7]:

$$\Pi^r(S(\xi), P(\xi), \xi) = \max_{\eta} \Pi^r(S(\eta), P(\eta), \xi), \quad \forall\, \xi \in \Theta. \tag{IC}$$

To ensure the retailer's participation, the supplier guarantees at least as much profit to the retailer as he would make (in expectation) by rejecting the offered menu of contracts. The value of the retailer's outside option is given by

$$\begin{aligned}
&\Pi^r(S(\xi), P(\xi), \xi) \\
&\geq \overbrace{(r-w)\sum_{i=n}^{\infty} \alpha^{i-n}\mathbb{E}_{D(\xi)}[\min\{y_i^o, D_i(\xi)\} | x_n = x]}^{\Pi^r_{\min}(x, \mathbf{y}^o, \xi)}, \\
&\qquad\qquad \forall \xi \in \Theta, \quad \text{(PC)}
\end{aligned}$$

where $\mathbf{y}^o := (y_n^o, y_{n+1}^o, \ldots)$ denotes the order up to levels used by the supplier under the status quo VMI agreement if the retailer was to reject the menu of contracts. The value of the outside option for the retailer depends on the inventory level in the period that the contracts are offered, the order up to levels maintained by the supplier over the remaining horizon, and his type (via the type-dependent demand distribution).

The supplier's incentive problem can be summarized as follows:

$$\begin{aligned}
&\tilde{\Pi}^{sr}(x, \pi) := \max_{S(\cdot), P(\cdot)} \mathbb{E}_{\xi}[\Pi^s(x_n, S(\xi), P(\xi) | x_n = x, \pi_n = \pi]; \\
&\quad \text{subject to } S(\cdot) \geq x, \text{ (IC) , and (PC).}
\end{aligned} \tag{4}$$

We remark that the above problem is drastically different from the classical static mechanism design problems, the solutions of which generally follow the work of Mirrlees (1971). We highlight these differences when we analyze the problem in Section 4.

3.1.2. The Bayesian Inventory Control and Optimal Stopping Problem.

Let $\tilde{V}(x, \pi)$ denote the supplier's maximum profit using the learn-and-screen approach starting with the initial on-hand inventory x and belief over ξ given by π. For all $(x, \pi) \in \mathbb{R}^+ \times \mathcal{A}$, the value function is given by

$$\begin{aligned}
\tilde{V}(x, \pi) := \sup_{(\mathbf{y}, \tau) \in \mathcal{M}} \mathbb{E}&\left[\sum_{n=1}^{\tau-1} \alpha^{n-1}\left(cx_n + (w-c)y_n\right.\right. \\
&\left. - (w+h)\int_0^{y_n} Q_n(z)\mathrm{d}z\right) \\
&\left. + \alpha^{\tau-1}\tilde{\Pi}^{sr}(x_\tau, \pi_\tau) \Bigg| \begin{matrix} x_1 = x, \\ \pi_1 = \pi \end{matrix}\right],
\end{aligned} \tag{5}$$

where $\mathbf{y} := (y_1, y_2, \ldots, y_{\tau-1})$, $\tau \in \{1, 2, \ldots\} \cup \{+\infty\}$ is a stopping time of the filtration generated by the sales process, and $\mathcal{M}$ denotes the set of all admissible policies.[8] Admissible policies essentially mean that the supplier's decisions in period n can only be based on the information gathered from the sales observations up to period $n-1$. In each period, the supplier optimally decides whether to continue learning about demand through POS information or stop and offer a menu of contracts to the retailer so as to credibly obtain the retailer's private demand information. It follows from the principle of optimality that the value function associated with the learn-and-screen approach solves the following (functional) dynamic programming (DP) equation:

$$\tilde{V}(x, \pi) = \max\{\tilde{\Pi}^{lr}(x, \pi), \tilde{\Pi}^{sr}(x, \pi)\}, \text{where} \tag{6}$$

$$\begin{aligned}
\tilde{\Pi}^{lr}(x, \pi) &:= cx + \max_{y \geq x} \tilde{L}(y, \pi) \\
&:= cx + \max_{y \geq x}\left\{(w-c)y - (w+h)\int_0^y Q(z)\,\mathrm{d}z\right. \\
&\qquad + \alpha(1 - Q(y))\tilde{V}(0, \pi^c(\cdot | y)) \\
&\qquad \left. + \alpha\int_0^y q(z)\tilde{V}(y-z, \pi^e)\,\mathrm{d}z\right\}.
\end{aligned}$$

The first term in Equation (6) represents the supplier's maximum profit if she decides to learn and update her belief about the retailer's private demand information, ξ, using the POS sales data. The second term is the supplier's maximum profit if she decides to offer the menu of contracts and screen the retailer's private demand information. The first three terms of $\tilde{L}(y,\pi)$ represent the myopic profit from raising the on-hand inventory level to y in the current period, and the last two terms correspond to future profit stream. The DP in Equation (6) can be simplified (see Appendix B for details) following the transformation $V(x,\pi) := \tilde{V}(x,\pi) - cx$ and $\Pi^{sr}(x,\pi) := \tilde{\Pi}^{sr}(x,\pi) - cx$. The resulting DP equation is

$$V(x,\pi) = \max\{\Pi^{lr}(x,\pi),\ \Pi^{sr}(x,\pi)\}, \text{ where} \tag{7}$$

$$\begin{aligned}\Pi^{lr}(x,\pi) &:= \max_{y\ge x} L(y,\pi)\\ &:= \max_{y\ge x}\left\{(w-c)y-(w+h-\alpha c)\int_0^y Q(z)\,\mathrm{d}z\right.\\ &\quad + \alpha(1-Q(y))V(0,\pi^c(\cdot|y))\\ &\quad \left. + \alpha\int_0^y q(z)V(y-z,\pi^e)\,\mathrm{d}z\right\}.\end{aligned} \tag{8}$$

The existence of the value function $V(x,\pi)$ and an optimal control $(\mathbf{y}^*,\tau^*)\in\mathcal{M}$ is established in Section EC.2 of the e-companion. Furthermore, we also propose a successive approximation scheme to compute the value function and the optimal policy.

3.2. Relation to the Classical Bayesian Inventory Problem

Consider the following scenario in which the supplier decides to never offer the screening contracts to the retailer: that is, $\tau = +\infty$ (the learn-only approach). In this case, the supplier faces a Bayesian inventory management problem with unobserved lost sales. Even this simpler inventory management problem is difficult to solve for two reasons. First, the objective function, $\tilde{L}(y,\pi)$, is not concave in y. In particular, future beliefs are affected by inventory decisions through $\pi^c(\cdot|y)$ in Equation (1). For an example, we refer the reader to theorem 2(i) in Bisi et al. (2011). Hence, a simple policy, such as a state-dependent base stock policy, does not need to be optimal. Second, to update the belief over ξ at the end of period n, one needs to keep track of π_1 (the supplier's belief in period $n=1$) and the entire history of sales $z_1,\ldots,z_{n-1}$. Hence, the state space of the DP grows with time, and one quickly runs into the curse of dimensionality.

To overcome these analytical and computational challenges, researchers have focused on the newsvendor class of distributions (see Lariviere and Porteus 1999, Bisi et al. 2011). These demand distributions possess several desirable statistical properties when the on-hand inventory level censors demand observations (as is the case in the classical newsvendor problem and hence, the name newsvendor class of distributions).

Definition 1. A cdf $F(z|\xi)$, $z\ge 0$, belongs to the newsvendor family (denoted henceforth as $\mathcal{N}$) if it can be expressed as $1-e^{-\frac{t(z)}{\xi}}$, where ξ is the parameter and $t(z)$ is a nonnegative increasing function.

Given n sales realizations $z_1,\ldots,z_n$, of which m are uncensored demand observations, the two-dimensional sufficient statistic for the newsvendor likelihood and the unknown parameter ξ are $(m;\ \sum_{i=1}^n t(z_i))$. That is, all information contained in the sample $z_1,\ldots,z_n$ regarding the unknown parameter ξ can be summarized by these two numbers. Thus, the state space of the DP can be reduced to three variables. Braden and Freimer (1991) were the first to identify the newsvendor family of distributions. Furthermore, they present distributions, including Weibull, that belong to the newsvendor family. We remark that the learn-and-screen approach adds a new dimension to the classical Bayesian inventory problem by incorporating demand information asymmetry in decentralized supply chains.

In summary, the Bayesian learning approach may be tractable when the demand distribution is from the newsvendor family. In this case, the supplier can optimally determine a state-dependent order up to policy, which can vary from one period to the other. Implementing the learning approach over a short horizon (a few selling periods) is practical under a VMI setting. However, learning over a long horizon can be challenging and may also prove controversial, because the retailer can face low inventory levels in some periods after observing high inventory levels. To circumvent this potential issue, the supplier can stop the learning process after a short period and use her demand information to manage the inventory. In that case, the supplier essentially faces the classical lost sales inventory management problem, for which the base stock policy is optimal (Karlin and Scarf 1958). The optimal base stock level based only on her updated belief (at the beginning of period n) about demand is given by

$$S^o(\pi_n) := Q_n^{-1}\left(\frac{w-c}{w+h-\alpha c}\right),\ n\ge 1, \tag{9}$$

where $Q_n(z)$ is defined in Equation (3). Henceforth, we will use $S^o(\pi)$ and S^o_π interchangeably, where π denotes the supplier's belief in the period when the menu of contracts is rejected.

4. Analysis

We solve for the supplier's optimal strategy using backward induction. In Section 4.1, we provide tractable sufficient conditions for the incentive and participation constraints to hold. These conditions allow us to characterize the optimal menu of contracts in closed form. In Section 4.2, we study the impact of Bayesian (demand) learning on the evolution of the optimal contracts. In Section 4.3, we investigate the structure of the supplier's optimal timing of the menu of contracts in an ongoing VMI agreement.

4.1. Optimal Menu of Contracts

The supplier needs to consider three important issues in designing a menu of contracts when she implements the learn-and-screen approach. First, the supplier has to offer information rent in addition to what the retailer obtains from the status quo VMI agreement (so as to encourage the retailer to accept a new contract and hence, the revised VMI agreement). The retailer's profit from the status quo VMI agreement depends on not only his private demand information but also, the supplier's inventory policy during the remaining planning horizon, which is a base stock policy with a fixed base stock level S^o_π. Note that the supplier's belief π in the period when the contracts are rejected is affected by the supplier's inventory decisions before that period. Furthermore, these inventory decisions are impacted by the supplier's option to screen the retailer in a future period. In this sense, the belief π and hence, S^o_π are endogenous to the contracts offered by the supplier. In other words, the retailer's "outside" option (after rejecting the menu of contracts) is no longer exogenous (as in classical adverse selection problems) but instead, endogenous to the contract design problem. Second, using the POS data, the supplier can improve her belief about the retailer's private demand information before offering a menu of contracts. Visibility of POS data upstream in a VMI agreement facilitates this dynamic learning. This learning aspect also adds a new dimension to the adverse selection problems seen in operations management/economics literature. Third, the supplier's inventory decisions before offering a menu of contracts affect the learning process. The on-hand inventory censors the demand realization in an unobserved lost sales environment (which is typical in most retail stores). Therefore, a larger on-hand inventory reduces the amount of lost sales, enabling POS data to better capture the true demand distribution and resulting in a better demand learning process for the supplier at the expense of higher holding costs.

To design the menu of contracts, we first characterize the retailer's outside option (i.e., his reservation profit) if the menu of contracts is rejected. Following the retailer's rejection of the menu of contracts, the supplier continues to manage inventory with unobserved lost sales using a base stock policy with base stock level S^o_π, defined in Equation (9), for the rest of the planning horizon. After the retailer is given an option to credibly reveal his private information, the supplier no longer puts forth effort and resources to continue to update belief about the retailer's private information. The retailer knows this fact. In other words, if the retailer rejects the menu of contracts, the supplier simply follows the status quo VMI agreement with the optimal base stock level S^o_π from then on. Base stock policy and this base stock level are optimal for the supplier, because the supplier faces the classical lost sales inventory management problem for the remaining planning horizon (Karlin and Scarf 1958). Hence, type ξ retailer's reservation profit for the remaining planning horizon satisfies the following functional equation

$$\begin{aligned}\Pi^r_{\min}(x, S, \xi) &:= (r-w)\int_0^y \overline{G}(z|\xi)\mathrm{d}z + \alpha\Pi^r_{\min}(0, S, \xi)\overline{G}(y|\xi)\\ &\quad + \alpha\int_0^y \Pi^r_{\min}(y-z, S, \xi)g(z|\xi)\mathrm{d}z, \qquad (10)\end{aligned}$$

where $y = \max\{x, S\}$ and $S = S^o_\pi$.[9] Next, we establish upper and lower bounds for the retailer's reservation profit.

Theorem 1. *For any* $\xi \in \Theta$,

$$\begin{aligned}\underline{\Pi}^r_{\min}(S^o_\pi, \xi) \le \Pi^r_{\min}(x, S^o_\pi, \xi) \le \overline{\Pi}^r_{\min}(x, S^o_\pi, \xi)\\ := (r-w)(x - S^o_\pi)^+ + \underline{\Pi}^r_{\min}(S^o_\pi, \xi),\end{aligned}$$

where $\underline{\Pi}^r_{\min}(S^o_\pi, \xi) := \frac{(r-w)\int_0^{S^o_\pi}\overline{G}(z|\xi)\mathrm{d}z}{1-\alpha}$.

We remark that the upper bound is essentially equal to the retailer's reservation profit if the on-hand inventory is lower than the base stock level offered when the menu of contracts is rejected. In other words, for those cases, the upper bound gives us the retailer's true reservation profit. To design the menu of contracts, we use this upper bound (i.e., use a conservative estimate) for the retailer's outside option. We replace the participation constraint (PC) constraint in the optimization problem of Equation (4) with this upper bound. The solution to this optimization problem with a tighter constraint also satisfies the incentive compatibility (IC) constraint in the original problem. Hence, it is a feasible solution for the original optimization problem. In Section 5, we also solve the optimization problem by replacing the retailer's reservation profit with the lower bound.[10] These two solutions define an upper and lower bound on the supplier's optimal profit (obtained as a solution to the original problem in Equation (4) in which the retailer's reservation profit is $\Pi^r_{\min}(x, S^o_\pi, \xi)$). Hence, these two solutions provide the optimality gap and show how well the menu of contracts (designed by using the upper bound on the retailer's reservation profit) performs.

Lemma 1. *An incentive-compatible menu of contracts $\{S(\cdot), P(\cdot)\}$ satisfies the revised (PC) if $S(\xi) \geq S^o_\pi$, $\forall \xi$ and $\Pi^r(S(\underline{\xi}), P(\underline{\xi}), \underline{\xi}) = \overline{\Pi}^r_{\min}(x, S^0_\pi, \underline{\xi})$.*

Recall that, if the retailer rejects the menu of contracts, the supplier continues to implement the base stock policy used in the status quo VMI agreement but uses her updated belief and sets the base stock level to S^o_π. Lemma 1 shows that, to encourage the retailer to accept the revised VMI, the supplier needs to offer a base stock level that is at least as much as what the retailer would get in the status quo VMI.

The (IC) constraint provides first-order conditions on the feasible menu of contracts. Using Lemma 1, (IC) can be restated as follows.

Lemma 2. *A menu of contracts $\{S(\cdot), P(\cdot)\}$ satisfies the (IC) constraint if and only if the menu also satisfies the following.*

(i) *The expected profit of the retailer for the remaining periods is given by*

$$\Pi^r(\xi) := \Pi^r(S(\xi), P(\xi), \xi) = \overline{\Pi}^r_{\min}(x, S^o_\pi, \underline{\xi}) - \underbrace{\frac{(r-w)}{1-\alpha} \int_{\underline{\xi}}^{\xi} \int_0^{S(\eta)} \frac{\delta}{\delta \eta} G(z|\eta) \mathrm{d}z \mathrm{d}\eta}_{\textit{Information Rent}}, \quad \forall\, \xi \text{ and } n. \tag{11}$$

(ii) *The base stock level $S(\xi)$ is increasing.*[11]

The conditions in Lemmas 1 and 2 can be used to reformulate the supplier's dynamic contract design problem and obtain her resulting optimal profit.

Lemma 3. *The following optimization problem determines the optimal menu of contracts:*

$$\tilde{\Pi}^{sr}(x, \pi) = cx - \overline{\Pi}^r_{\min}(x, S^o_\pi, \underline{\xi}) + \frac{1}{1-\alpha} \max_{\left\{\substack{\frac{\mathrm{d}S(\xi)}{\mathrm{d}\xi} \geq 0, \\ S(\underline{\xi}) \geq \max\{x, S^o_\pi\}}\right\}} \int_\Theta \pi(\xi) H(S(\xi), \xi|\pi) \mathrm{d}\xi, \tag{12}$$

where
$$H(S, \xi|\pi) := (r-c)S - (r+h-\alpha c)\int_0^S G(z|\xi)\mathrm{d}z + \frac{(r-w)}{\lambda(\xi)} \int_0^S \frac{\partial}{\partial \xi} G(z|\xi) \mathrm{d}z \tag{13}$$

and $\lambda(\xi) := \frac{\pi(\xi)}{\int_\xi^{\bar{\xi}} \pi(\eta)\mathrm{d}\eta}$, $n \geq 1$ is the failure rate function corresponding to the pdf, π.

Solving the above problem determines the optimal menu of contracts to be offered to the retailer. In Lemmas 2 and 3, we use the classical approach (Mirrlees 1971) to express (IC) as a differential equation that governs the marginal information rent offered to the retailer. However, the similarity with the classical mechanism design solution approach ends here. First, note that the supplier's (principal's) belief π is dynamically updated using the Bayes' rule (see Equation (1)). Thus, the information structure in our problem deviates from the classical principal agent problem in which the principal's belief remains static. Second, the supplier's dynamic *virtual surplus*, $H(S, \xi|\pi)$ (defined in Equation (13)), depends on the prior information π_1, historical POS data $z_1, z_2, \ldots$, and inventory decisions $y_1, y_2, \ldots$ through the updated failure rate function $\lambda(\cdot)$. The function $\lambda(\xi)$ measures the supplier's belief that the parameter of demand distribution is ξ given that the parameter is at least ξ. Larger ξ translates to having greater average demand. Hence, λ captures the dynamic learning aspect of the learn-and-screen approach.

Third, consider the maximization problem in Equation (12) without the monotonicity constraint $\frac{\mathrm{d}S(\xi)}{\mathrm{d}\xi} \geq 0$. The standard solution approach involves showing that the optimal base stock menu for this relaxed problem is increasing in the retailer's type, thus establishing its optimality for the constrained optimization problem. However, unlike in the classical mechanism design problems, the function $H(S, \xi|\pi)$, which results from a multiperiod lost sales inventory problem (after accounting for the retailer's information rent), is not concave in S for every type ξ and a given prior π. Thus, we look for weaker structural properties, such as unimodality of $H(\cdot, \xi|\pi)$, that would ensure that first-order conditions are not only necessary but, also sufficient for existence and uniqueness of the optimal solution. To this end, the next theorem characterizes a family of demand distributions for which $H(S, \xi|\pi)$ is unimodal in S and determine its maximizer using the first-order conditions. Lemma C.2 in Appendix C provides conditions to verify unimodality.

Theorem 2. *Suppose that the demand distribution, $G(z|\xi)$, is from the exponential family (see section 3.4 in* Berger and Casella 2002*) (i.e., $g(z|\xi) = k(z)l(\xi)e^{-s(\xi)t(z)}$, $z \geq 0$, and functions $k(\cdot), l(\cdot), t(\cdot)$ and $s(\cdot)$ are differentiable, where $k(\cdot), t(\cdot)$ are defined over $\mathbb{R}^+$ and $l(\cdot), s(\cdot)$ are defined over Θ). If $l(\cdot), s(\cdot)$ are decreasing and $t(\cdot)$ is increasing, then the following statements hold.*

1. *The family of demand distributions $\{G(z|\xi)\}_{\xi \in \Theta}$ is stochastically increasing.*
2. *$H(\cdot, \xi|\pi)$ is unimodal for all $\xi \in \Theta$.*
3. *$S^*(\xi|x, \pi) := \max\{x, S^o_\pi, \hat{S}(\xi|\pi)\}$ is the maximizer of $H(S, \xi|\pi)$ over $S(\underline{\xi}) \geq \max\{x, S^o_\pi\}$, where $\hat{S}(\xi)$ solves the following first-order condition:*

$$(r-c) - (r+h-\alpha c)G(S|\xi) + \frac{(r-w)}{\lambda(\xi)} \frac{\partial}{\partial \xi} G(S|\xi) = 0. \tag{14}$$

Normal, gamma, and Weibull are some distributions that satisfy the sufficient conditions (all with unknown *scale* parameter) in Theorem 2. Unimodality of $H(\cdot, \xi|\pi)$ ensures a simple characterization of the menu of base stock levels. Such a characterization emphasizes the practical value of offering the menu of contracts to the retailer in a VMI relationship. After the retailer accepts a contract, the inventory policy for the remaining horizon is a simple base stock policy. We note that the characterization in Theorem 2 also include the newsvendor family of distributions,[12] which have been widely used in unobserved lost sales Bayesian inventory literature (recall Section 3.2). In Theorem 3, we provide sufficient conditions that guarantee monotonicity of the base stock levels S^* and characterize the optimal menu of contracts for the incentive problem in Equation (4) with a more stringent reservation profit given by the upper bound $\overline{\Pi}^r_{\min}(x, S^o_\pi, \xi)$.

Theorem 3. *Suppose that* $\pi \in \mathcal{A}$ *has increasing failure rate (IFR) property,* $\{G(\cdot|\xi)\}_\xi \subset \mathcal{N}$, *and* $\frac{r-c}{r+h-\alpha c} \leq 1 - e^{-2}$. *Then, the menu of base stock levels* $S^*(\xi)$ *is increasing. Therefore,* $\{S^*(\cdot), P^*(\cdot)\}$ *is an optimal menu of contracts, where*

$$P^*(\xi) := (r-w)\left(\int_0^{S^*(\xi)} \overline{G}(z|\xi)\mathrm{d}z + \int_{\underline{\xi}}^{\xi}\int_0^{S^*(\eta)} \frac{\partial}{\partial \eta} G(z|\eta)\mathrm{d}z\mathrm{d}\eta\right) - (1-\alpha)\overline{\Pi}^r_{\min}(x, S^o_\pi, \underline{\xi}). \tag{15}$$

The evolution of $\lambda_n(\cdot)$ via Bayesian updating determines the dynamic aspect of the optimal contracts. Note that $\lambda_{n+1}(\xi) > \lambda_n(\xi)$ means that, compared with the previous period n, the supplier is more confident in period $n+1$ that the underlying market signal is ξ given that it is at least ξ. Given this updated belief, it can be verified from Equation (14) that, in period $n+1$, the supplier offers a higher base stock level than what she offered in period n. The evolution of failure rate function depends on the supplier's historical inventory decisions and POS information. This dependence of contract structure on demand information, which also depends on how the supplier manages inventory, ties together the supplier's learning and screening problems. In Section 4.2, we further investigate the impact of inventory decisions under learn-and-screen approach on the supply chain operations.

The first condition in Theorem 3 (that π has IFR property) is a standard assumption in static, exogenous information mechanism design problems (see Tirole 2002, p. 156). The mechanism that we consider has a dynamic and *endogenous* information structure with learning.[13] The subtle issue here is the preservation of IFR property under Bayesian learning. The inverse gamma distribution is a conjugate prior for the newsvendor likelihood, and this motivates the second condition. The inverse gamma distribution has IFR property for most of its parameter space, and its conjugacy ensures that the IFR property is preserved. The last sufficient condition on the cost parameters is satisfied, for example, in industries with low to medium profit margins.[14] This assumption might seem restrictive at first. However, lower margins are typically a characteristic of the fast-moving consumer goods in the retail industry (see Taylor and Mauer 2013). We also remark that these conditions are sufficient but are not *necessary*.

Note from Equation (15) that the payment $P^*(\cdot)$ is increasing with the corresponding base stock level. Thus, there is a one-to-one correspondence between the payment and the corresponding base stock level. Hence, the supplier can construct the optimal payment schedule $P^*(S^*)$ equivalent to the menu of contracts $\{S^*(\cdot), P^*(\cdot)\}$. The supplier can then offer this payment schedule to the retailer. Type ξ retailer chooses order up to level $S^*(\xi)$ and pays $P^*(S)$. In other words, the retailer does not need to explicitly communicate his private information using the payment schedule. Instead, the retailer simply communicates the inventory level that he finds suitable given his market demand and makes the corresponding payment (or receives a payment if $P(\xi)$ is negative). In VMI relationships, retailers are generally expected to communicate changing market conditions to their suppliers and in doing so, help suppliers make informed inventory decisions. Such engagements are considered to be a good practice for VMI relationships.[15] However, such communications remain informal and hence, are prone to manipulation or misinterpretation (i.e., they may be perceived as manipulation). The proposed learn-and-screen mechanism ensures that the retailer stands monetarily accountable for providing his input.

4.2. Impact of Bayesian Learning on the Proposed Menu of Contracts

Here, we investigate how the supplier's observation of sales in a period (i.e., learning from POS data) affects the structure of the incentives (i.e., the menu of base stock levels and corresponding payments) offered to the retailer in the following period. Note from Part 3 of Theorem 2 that the proposed menu of base stock levels in any period n is the maximum of the on-hand inventory x_n, the base stock level $S^o(\pi_n)$ that the supplier follows if the retailer were to reject the menu of contracts (a factor related to (PC)), and the maximizer $\hat{S}(\xi|\pi_n)$ of $H(\cdot, \xi|\pi_n)$ (the function that determines information rent arising from incentive compatibility). Given an on-hand inventory level, the remaining two factors evolve over time as a result of the learning process summarized in the supplier's belief π_n about the retailer's private demand information. Learning from sales observations (i.e., POS data) can have a contrasting effect on these two factors (i.e., the retailer's alternative option when he rejects the contracts and the incentive compatibility).[16] Note also that the proposed menu of base stock level S^* is increasing with $\hat{S}$. Similarly, the corresponding payment P^* is also monotone with respect to S^*. Hence, to isolate the effect of learning on primarily the incentive structure (and to simplify our discussions), we focus on how learning affects the menu of contracts $\{\hat{S}(\xi|\pi_n), \hat{P}(\xi|\pi_n)\}$, where $\hat{P}(\cdot)$ is defined as in Equation (15) but with S^* replaced by $\hat{S}$.

4.2.1. Evolution of the Menu of Base Stock Levels. We proceed by investigating the impact of sales observations on the supplier's belief process and in turn, how the updated belief effects the base stock levels $\hat{S}(\cdot)$. The following theorem shows that, by Bayesian updating, the supplier's posterior belief is stochastically ordered in the sense of failure (hazard) rate with respect to the prior belief.[17] The direction of ordering depends on whether the demand realization in a period is censored.

Theorem 4. *The following statements hold for any $n \geq 1$.*
1. *If the sales observation in period n is censored (i.e., $z_n = y_n$), then $\lambda_n(\xi) \geq \lambda_{n+1}(\xi|y_n)$ and $\hat{S}(\xi|\pi_n) \geq \hat{S}(\xi|\pi_{n+1})$.*
2. *If the sales observation in period n is uncensored (i.e., $z_n < y_n$), then*
   (i) *for types $\xi \geq t(z_n)$, $\lambda_n(\xi) \leq \lambda_{n+1}(\xi)$ and $\hat{S}(\xi|\pi_n) \leq \hat{S}(\xi|\pi_{n+1})$.*[18]
   (ii) *for types $\xi < t(z_n)$, such that $g(z_n|\xi) < g(z_n|\overline{\xi})$, we have $\lambda_n(\xi) > \lambda_{n+1}(\xi)$ and $\hat{S}(\xi|\pi_n) > \hat{S}(\xi|\pi_{n+1})$.*

Theorem 4 highlights and quantifies how the two factors, namely the occurrence of a stockout event and the magnitude of the sales, determine the evolution of the incentive structure offered to the retailer. A censored demand observation (i.e., POS data with the knowledge that on-hand inventory at the retailer dropped to zero during the sales period) always lowers the base stock level $S(\cdot)$ offered to the retailer. Thus, censored information negatively impacts the retailer by reducing the incentive (i.e., the base stock level) offered to him (who always prefers a high base stock level). This result identifies a new drawback of censored demand information (in addition to the well-known drawback on customers' perception of the retail store). In contrast, an uncensored demand observation (i.e., POS data with no stockout) translates into a reduced incentive only for a retailer with low demand information. How much the offered base stock level changes as a result of learning depends on the magnitude of the observed POS data (through the learning process specified in Equation (1)). Next, we explore the intuition behind these results.

A censored demand observation suggests to the supplier that the average market size must be larger than what was expected in the previous period. An interesting consequence of this ordering is that the menu of base stock levels offered in the following period is smaller. Intuitively, one would expect the supplier to offer higher base stock levels in the following period due to the learning dynamic, that is, account for her updated belief in higher demand. However, careful analysis and thought reveal that the result is in fact the opposite. The optimal base stock levels are driven by both the learning dynamic and the incentive that needs to be offered to the retailer to facilitate credible communication. The increased confidence in a larger market size (following a censored demand observation) implies that the retailer makes greater expected profit than the previous period. As a result, the supplier lowers the menu of base stock levels (and hence, the incentive) offered to the retailer, while still facilitating credible communication (Part 1 of Theorem 4).

An uncensored demand observation suggests to the supplier that the average market size may be smaller or larger than what was expected in the previous period. The direction of this ordering depends on the magnitude of the sales observation. Given that the average market size is at least ξ, a small demand observation increases the supplier's confidence that the average market size is ξ in the following period (refer to the definition of λ_n in Lemma 3). This intuition explains the failure rate ordering in Part 2(i) of Theorem 4. One may expect, as a result, the menu of base stock levels to become smaller. However, this is not the case. The supplier offers higher base stock levels in the following period (greater incentive) to ensure that the retailers with more optimistic demand information (i.e., retailer types $\xi > t(z_n)$) are able to credibly share their demand information. Increasing the menu of base stock levels $\hat{S}(\cdot)$ in the following period provides sufficient incentive to deter the retailer with larger market size ($\xi > t(z_n)$) from choosing a base stock level meant for a smaller market. The situation, however, gets interesting as the magnitude of the demand observation increases. As the magnitude increases, the supplier becomes more confident that the underlying average market size is large. In the event of a large uncensored demand observation, the supplier mimics her actions following a censored demand observation, i.e., resorts to lowering the menu of base stock levels (Part 2(ii) of Theorem 4).

4.2.2. Symmetric Demand Information. Next, we study the symmetric information setting to understand how learning impacts the efficiency of screening contracts in the learn-and-screen approach. In the symmetric setting, the supplier and the retailer have identical information about demand. That is, market signal ξ is common knowledge, whereas demand realization in a period is still unknown to both players at the beginning of the period. In this case, the supplier faces the classical lost sales inventory control problem (Karlin and Scarf 1958). The optimal base stock level for this problem is $G^{-1}(\frac{w-c}{w+h-\alpha c}|\xi)$, which depends on the underlying market conditions ξ. In this case, the wholesale price contract, commonly used in VMI agreements, prohibits the supplier from coordinating the supply chain. In Theorem 5, we characterize linear base stock level contracts that coordinate this decentralized supply chain and examine its properties.

Theorem 5. *Under symmetric demand information, the linear base stock level contract with base stock level $S^{fb}(\xi) := G^{-1}(\frac{r-c}{r+h-\alpha c}|\xi)$ and marginal price $p^{fb} := \frac{(r-w)(h+c(1-\alpha))}{r+h-\alpha c}$ coordinates the channel. That is, this contract maximizes the total supply chain profit.*

In this case, the supplier does not need to learn from POS data, because she has the same demand information as the retailer. The theorem shows that the supplier can design a simple supply chain coordinating contract, which offers the base stock level S^{fb} to the retailer in exchange for a payment of $p^{fb}S^{fb}$ from the retailer in each period. The supplier implements this contract by announcing the marginal price (constant) p^{fb} rather than offering a menu of contracts. In response, the retailer would choose the base stock level S^{fb} to maximize his profit.

The symmetric information setting may also materialize as a special case in the learn-and-screen approach if the supplier's belief π_n converges to the Dirac delta function δ_ξ with the probability mass concentrated at ξ. That is, the supplier believes with probability one that the average market size is ξ.[19] For this case, Equation (14) is still valid, and it provides exactly one solution $\hat{S}(\xi|\delta_\xi)$, which is equal to the supply chain coordinating base stock level $S^{fb}(\xi)$. This equivalence can be verified by noting that $\lambda(\xi) = \infty$. Thus, over time, if the supplier's belief about market conditions converges to a singular distribution (Dirac delta), the base stock level in the learn-and-screen approach equals the coordinating base stock level. Note, however, that the convergence of the belief process and hence, the base stock levels in the learn-and-screen approach are a function of the sales observation process (see Theorem 4). In the following section, we further explore the impact of learning on the payment function in the learn-and-screen approach. We use the linear price contract $p^{fb}S$ as a benchmark.

4.2.3. The Evolution of $\hat{P}(\xi)$.

Here, we show how Bayesian learning about the retailer's private demand information affects payment $\hat{P}(\xi)$ in the learn-and-screen approach. Note that monotonicity of $\hat{P}(\cdot)$ and $\hat{S}(\cdot)$ implies that the menu of contracts can be implemented as a payment schedule $\hat{P}(S)$.

Theorem 6. *The following statements hold for any $n \geq 1$.*

1. *The marginal price paid by the retailer is always greater than the first best price: that is, $\frac{d\hat{P}_n}{d\hat{S}_n} \geq p^{fb} = \frac{(r-w)(h+c(1-\alpha))}{r+h-\alpha c} > 0$ for all $\xi \in \Theta$.*[20]
2. *If the demand realization in period n is censored, then $\frac{d\hat{P}_{n+1}}{d\hat{S}_{n+1}} \geq \frac{d\hat{P}_n}{d\hat{S}_n}$.*
3. *If the demand realization in period n is uncensored, then $\frac{d\hat{P}_{n+1}}{d\hat{S}_{n+1}} \leq \frac{d\hat{P}_n}{d\hat{S}_n}$ for all $\xi \geq t(z_n)$.*

Part 1 shows that the retailer pays a higher marginal price than in the symmetric demand information setting to convince the supplier about the credibility of his demand information. This monetary commitment on the part of the retailer assures credibility of the information that he shares. Parts 2 and 3 illustrate the impact of learning from sales observation on the marginal price. Censored demand observation negatively impacts the retailer (and the supply chain). It results in the retailer choosing a lower base stock level (Part 1 of Theorem 4) and paying more for the contract in the following period. However, an uncensored small demand observation benefits the retailer and the supply chain. In particular, the retailer pays a lower marginal price and reserves a more suitable base stock level given his market conditions. Lower marginal price provides sufficient incentives to deter the retailer with larger market from choosing a low base stock level. These observations corroborate the discussion of Theorem 4. Figure 1 illustrates the dynamics of $\hat{P}(S)$ driven by observed POS data. We remark that the payment function illustrated in Figure 1 is concave, which implies a quantity discount scheme (although we do not formally prove this result to hold for the proposed contract). The concavity of a screening contract was first formally shown in part 5 of theorem 1 in Özer and Wei (2006). However, this concavity does not necessarily hold in a dynamic setting. For example, Oh and Özer (2013) show that the optimal contract is neither concave nor convex in a dynamic environment that they study.

4.3. Timing of Contracts in VMI

Suppose that the supplier starts a period with on-hand inventory level x and belief π. The supplier has to decide between offering the optimal menu of contracts or continuing to learn via Bayesian updating. By delaying to offer the menu of contracts, the supplier

Figure 1. Dynamics of Optimal Contracts

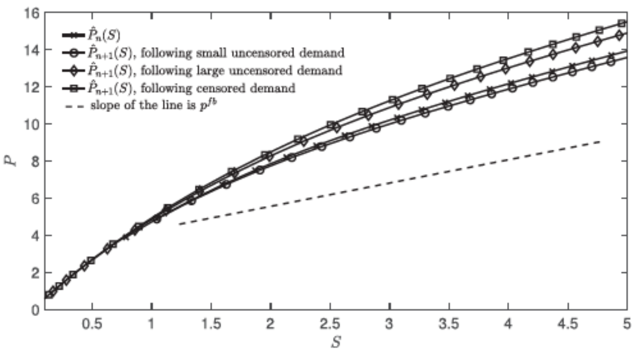

Notes. In this figure, cost and demand parameters are as follows: $r = 12, w = 6, c = 3, h = 2, \alpha = 0.8$, and $\Theta = [2, 12]$. Demand follows exponential(ξ) distribution, and its prior is inverse gamma (IG) distribution with shape and scale parameters (9,30), respectively. In the small uncensored case, $z_n = 1$, corresponding to updated IG distribution with parameters (10,31). In the large uncensored case and censored cases, $z_n = 10$, corresponding to updated parameters (10,40) and (9,40), respectively. Section EC.1.2 of the e-companion provides the details on belief updating.

gives herself a chance to improve her knowledge about market conditions, contingent on her prior inventory decisions. If, after many selling periods, the supplier is able to accurately estimate market conditions (i.e., there is no longer information asymmetry), she can extract all of the information rent by screening the retailer (see Theorem 5 for the coordinating contract). Thus, delaying has potential benefits in terms of greater rent extraction. However, until the contracts are offered, the supplier makes inventory decisions based on her limited knowledge of demand, potentially resulting in lost sales over several periods.

The aforementioned tradeoff drives the supplier's inventory decisions and the timing of screening contracts. The following theorem provides a partial characterization of the optimal time to offer the menu of contracts.

Theorem 7. *The following statements hold for any* $(x, \pi) \in \mathbb{R}^+ \times \mathcal{A}$.

1. $\Pi^{sr}(x, \pi)$ *is decreasing in* x.
2. *If* $x > \max\left\{G^{-1}\left(\frac{r-w(1-\alpha)-\alpha c}{r-w(1-\alpha)+\alpha^2(h-c)} \,|\overline{\xi}\right),\ S^o(\pi),\ \hat{S}(\overline{\xi}|\pi)\right\}$,

then it is optimal for the supplier to continue learning using sales data.

On-hand inventory level serves as an indicator for the supplier to determine whether it is better to postpone offering contracts. With a high on-hand inventory level, the supplier incurs holding cost for excess inventory regardless of screening the retailer. The supplier could, therefore, benefit from learning while inventory level is high and eventually, offer the menu of contracts when she is better informed about market conditions. Next, we illustrate and quantify the value of the learn-and-screen approach.

5. Value of Learn-and-Screen Approach

Here, we quantify the value of three approaches that the supplier could use to manage inventory and information in an ongoing VMI framework. (i) In the *learn* approach ($\tilde{V}^{lr}$), the supplier statistically improves her demand forecasts using POS data until the end of the planning horizon. Note that $\tilde{V}^{lr}$ can be computed using the recursion in Equation (6) by setting $\tilde{\Pi}^{sr} \equiv -\infty$. (ii) In the *screen* approach ($\tilde{V}^{sr}$), the supplier offers the optimal menu of contracts to the retailer at the beginning of the planning horizon. $\tilde{V}^{sr}$ can be computed by setting $\tilde{\Pi}^{lr} \equiv -\infty$ in Equation (6). However, in this approach, the supplier undermines her ability to learn more about demand over time. (iii) In the learn-and-screen approach ($\tilde{V}$), the supplier dynamically evaluates on-hand inventory level and her belief about market conditions to determine the timing of the contracts. The learn-and-screen approach also helps us quantify how much value screening adds to the learning approach (i.e., the value of screening $\left(\frac{\tilde{V}-\tilde{V}^{lr}}{\tilde{V}^{lr}} \times 100\right)$) and how much value dynamically learning adds to the screening approach (i.e., the value of learning $\left(\frac{\tilde{V}-\tilde{V}^{sr}}{\tilde{V}^{sr}} \times 100\right)$). In addition, we report the centralized supply chain profit ($\tilde{V}^{cs}$) as a benchmark. In the centralized setting, there is a single decision maker for the supply chain. This decision maker faces the classical lost sales inventory problem with complete demand information. The value function of centralized supply chain is computed using $\tilde{V}^{cs}_{\xi}(x, \xi) := cx + V^{cs}_{\xi}(x, \xi)$, which is defined in Equation (C.6). We report (with a slight abuse of notation) $\tilde{V}^{cs}(x, \pi) = \mathbb{E}_{\xi}[\tilde{V}^{cs}(x, \xi)]$.

To evaluate performance of these three inventory management approaches, we consider a two-point prior and exponential demand distribution under a range of market conditions. The average demand is either high ($\overline{\xi}$) or low ($\underline{\xi}$). The supplier's initial prior p denotes her belief that average demand is high in the ongoing season. The two-point prior simplifies the state space of the dynamic program in Equation (7) to two variables—inventory level (x) and probability of high type demand (p). We direct the reader to Section EC.1.1 of the e-companion for more details.

We compute $\tilde{V}$ and $\tilde{V}^{sr}$ using a stricter version of the participation constraint, $\Pi^r(\xi) \geq \overline{\Pi}^r_{\min}(x, S^o_{\pi}, \xi)$, which implies (PC). The value function obtained using $\overline{\Pi}^r_{\min}$, denoted by $\tilde{V}^{LB}$, is a lower bound to the value function with the (PC). We also report an upper bound $\tilde{V}^{UB}$ to the value function using the lower bound $\underline{\Pi}^r_{\min}(S^o_{\pi}, \xi)$ for the retailer's reservation profit. The optimality gap, which is defined as the percentage difference between the upper and the lower bounds, $\frac{\tilde{V}^{UB}-\tilde{V}^{LB}}{\tilde{V}^{LB}} \times 100$, gives us a measure of the performance of our proposed solution (with $\overline{\Pi}^r_{\min}$) to the supplier's problem.

Prevailing market conditions are quantified using two measures: double marginalization (DM) and degree of information asymmetry (DIA) in the supply chain. Double marginalization is the fraction of the total profit margin captured by the retailer (i.e., $\frac{r-w}{r-c}$). Higher double marginalization (with r, c fixed[21]) implies that the retailer has sufficient market power to capture a larger share of the total margin $r-c$ in an ongoing VMI agreement. The degree of information asymmetry, $\rho := \frac{\mathrm{Var}(D)}{\mathbb{E}_\xi[\mathrm{Var}(D|\xi)]}$ (which can be computed by the supplier), measures the variance in demand seen by the supplier relative to the expected variance in demand seen by the retailer given her current prior. Because $\mathrm{Var}(D) = \mathrm{Var}_\xi(\mathbb{E}[D|\xi]) + \mathbb{E}_1[\mathrm{Var}(D|\xi)]$,

$$\rho := \frac{\mathrm{Var}(D)}{\mathbb{E}_\xi[\mathrm{Var}(D|\xi)]} = 1 + \frac{\mathrm{Var}_\xi(\mathbb{E}[D|\xi])}{\mathbb{E}_\xi[\mathrm{Var}(D|\xi)]} = 1 + \frac{\mathrm{Var}_\xi(\xi)}{\mathbb{E}_\xi[\xi^2]}$$
$$= 1 + \frac{\mathbb{E}_\xi[\xi^2] - \mathbb{E}^2_\xi[\xi]}{\mathbb{E}_\xi[\xi^2]} = 2 - \frac{\mathbb{E}^2_\xi[\xi]}{\mathbb{E}_\xi[\xi^2]},$$

where the third equality follows from $D|\xi \sim \exp(\xi)$. From above, it follows that $\rho \in [1, 2]$. Note that $\rho = 1$ denotes the symmetric information setting and that $\rho = 2$ corresponds to the highest degree of information asymmetry in the supply chain. The degree of information asymmetry evolves endogenously in the VMI agreement, whereas double marginalization is exogenously determined ahead of the season.

We fix parameters $r = 12, c = 3, h = 2, \alpha = 0.8, \underline{\xi} = 2$ and vary other parameters as follows: $w = \{10, 6\}$, denoting low (L) and high (H) double marginalization, and $\overline{\xi} = \{6, 12\}$, representing small and large demand variability. We consider two priors, each representing L/H degree of information asymmetry in the supply chain.[22] The initial on-hand inventory level x is assigned integer values from zero to seven, and the supplier's profit, averaged across different x, is reported in Table 1.

We first consider the high double-marginalization condition as in innovative product supply chains in which the retailer (such as Apple) exercises greater market power. In such a supply chain, adopting either the learning approach or the screening approach is inefficient compared with the learn-and-screen approach. In the learning approach, the supplier makes a smaller margin on each unit of the product sold (relative to the retailer) and hence, is likely to maintain lower on-hand inventory levels than the coordinating level. Lower inventory levels further reduce the scope to observe and learn from uncensored demand data. In the screening approach, the supplier has to offer steep information rent and reservation profit (because of the retailer's market power) to learn the private demand information—especially when the degree of information asymmetry is high. In such market conditions, the supplier benefits from lowering the degree of information asymmetry by learning through POS data before offering the screening contracts to the retailer.

Next, we consider the low double-marginalization condition observed, for example, in functional product supply chains as in the fast-moving consumer good retail industry (such as the Sainsbury's example mentioned in Section 1). In such markets, the supplier exercises greater market power, and the retailer is typically squeezed for profit (e.g., Fisher 1997, Taylor and Mauer 2013). Despite low reservation profit of the retailer, the supplier may benefit from postponing offering the contracts to lower the information rent offered to the retailer. When the demand information asymmetry is high, the supplier gains up to 12.21% by delaying the offer of the contracts. Compared with the learning approach, the supplier gains up to 22.65% by strategically screening the retailer. The value of learn-and-screen approach is lowest when the prevailing market conditions are characterized by low double marginalization and low degree of information asymmetry.

Comparing $\tilde{V}^{LB}$ and $\tilde{V}^{cs}$ in Table 1, we note that the supplier makes up to 68.7% of centralized supply chain profit using the learn-and-screen approach (while ensuring that the retailer makes at least his reservation profit). Also, for the same parameter instances, we compare $\tilde{V}^{UB}$ and $\tilde{V}^{LB}$ and find that the

Table 1. Value of Learn-and-Screen Approach

$\overline{\xi}$	DM	DIA	$\tilde{V}^{lr}$	$\tilde{V}^{sr}$	$\tilde{V}^{LB}$	$\tilde{V}^{UB}$	$\tilde{V}^{cs}$	Value of screening (%)	Value of learning (%)	Optimality gap (%)
6	L	L	99.65	101.92	102.12	102.12	157.19	2.50	0.17	**8e-4**
6	L	H	55.18	54.90	60.47	60.59	94.55	9.97	12.21	0.20
6	H	L	34.31	52.44	61.46	61.87	157.19	84.32	29.35	0.67
6	H	H	19.56	20.48	33.69	35.44	94.55	77.99	172.02	5.13
12	L	L	200.66	201.88	202.62	202.63	305.59	0.98	0.37	**7e-3**
12	L	H	70.09	83.31	85.38	85.87	124.14	22.65	2.91	0.57
12	H	L	37.34	124.12	124.12	124.12	305.59	99.65	0.00	0.00
12	H	H	19.56	46.96	55.85	59.23	124.14	149.50	53.14	6.31

average optimality gap is 1.61%. This finding suggests that the supplier does not lose much by using $\overline{\Pi}^r_{\min}(x, S^o_\pi, \xi)$ instead of $\Pi^r_{\min}(x, S^o_\pi, \xi)$ as the retailer's reservation profit. The optimality gap is largest when both the degree of double marginalization and the degree of information asymmetry are high.

5.1. Impact of Screening Contracts on Inventory Decisions

Next, we numerically investigate the effect of screening the retailer on the supplier's inventory decisions before offering the menu of contracts. As a benchmark, we consider the supplier's inventory decisions without the option to screen the retailer (i.e., the learning approach). Figure 2 illustrates the base stock levels in the learn approach versus the learn-and-screen approach. We first note that the optimal inventory policy is a state-dependent base stock policy for both approaches. The supplier's updated belief p about ξ determines the base stock levels. The base stock levels in the learn-and-screen approach are lower than when the supplier does not have the option to offer screening contracts to the retailer. In essence, having the option to strategically screen the retailer for private demand information lowers the value of maintaining higher inventory levels to learn about demand through Bayesian updates. Yet, despite the lower inventory levels compared with the learning approach, the supplier and the retailer make more profit using the learn-and-screen approach. We remark that these observations are robust to varying other parameter $(r, c, h, \underline{\xi})$ values.

5.2. Timing of Contracts

The four panels in Figure 3 illustrate the optimal contract offering region under different market conditions. The supplier offers the optimal menu of contracts if her on-hand inventory level x (x axis) and her belief about demand p (y axis) fall within the shaded region. A larger region indicates that there is a greater likelihood that contracts are offered.

For a given belief (resp., on-hand inventory) level, the optimal stopping time has a state-dependent threshold (resp., control band) structure in the supplier's on-hand inventory level (resp., belief level). At the beginning of each period, for a given belief level, if the supplier's on-hand inventory is lower than a threshold, then the supplier should optimally offer the menu of contracts. Likewise, for a given on-hand inventory level, if the supplier's belief is within a control band, then the supplier should optimally offer the menu of contracts. Generally, we observe that the supplier is more likely to offer the menu of contracts when she is more confident that the underlying demand is high in view of the greater sales associated with the larger demand.

Comparing the upper panels with the lower panels in Figure 3 highlights the impact of demand variability on the timing of contracts. As the demand variability increases, the contract offering region enlarges. This suggests that learning about underlying demand through Bayesian updates can be slower, and the supplier is better off screening the retailer. Comparing the left panels with the right panels in Figure 3 illustrates the impact of double

Figure 2. Impact of the Learn-and-Screen Approach on Inventory Decisions

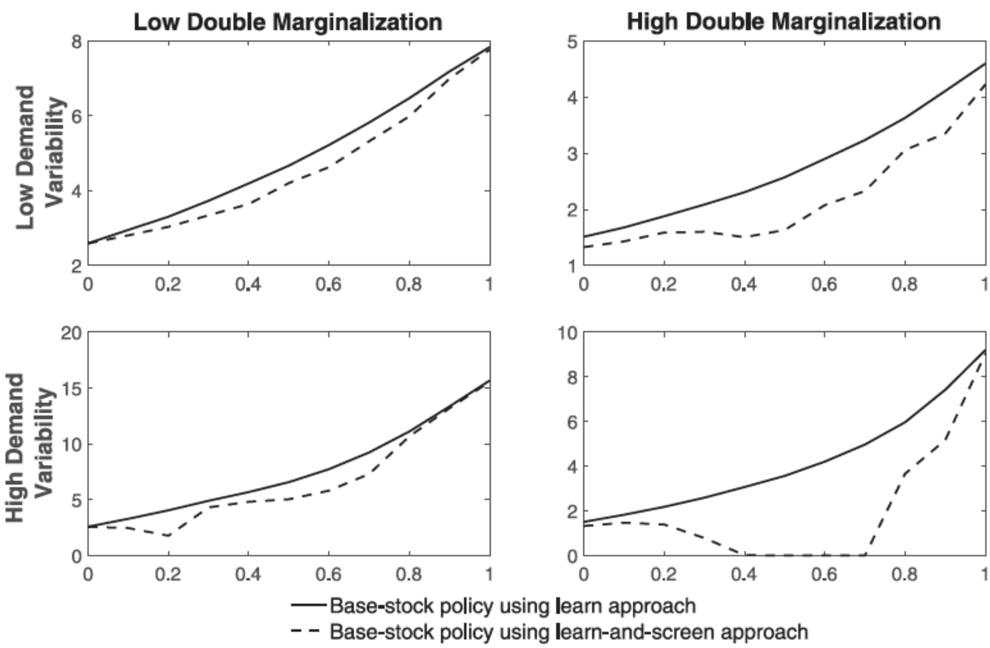

Figure 3. Optimal Timing of Screening Contracts in VMI

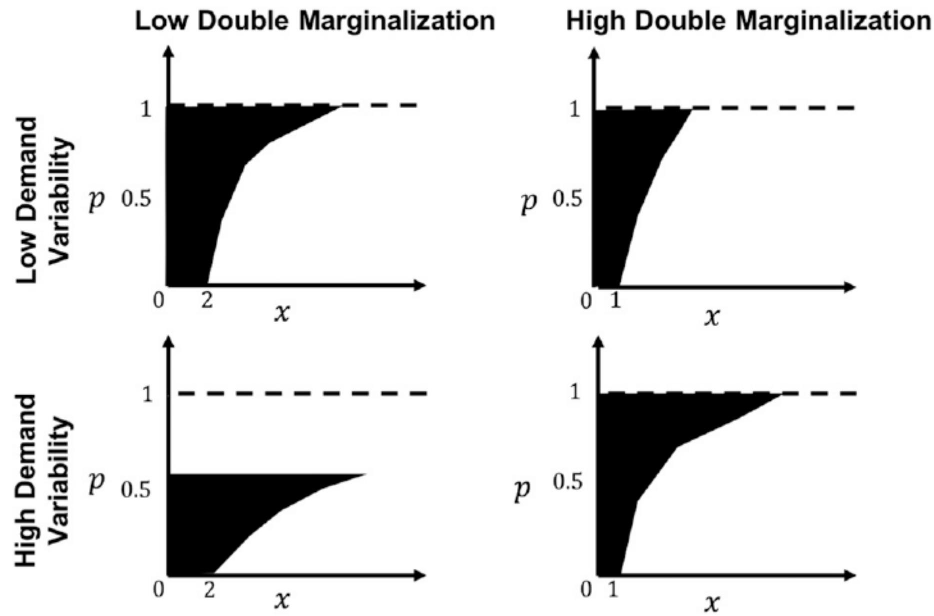

marginalization (or the retailer's market power) on the timing of the contracts. With higher double marginalization (retailer has greater market power), the contract offering region shrinks. All else equal, the retailer is guaranteed higher profit (compared with the low double-marginalization condition) if he accepts one of the contracts. Thereby, the supplier compensates for the increase in the retailer's market power by waiting longer before offering the contracts.

5.3. Dynamics of Information Rent

The two-point prior considered in this section also enables us to illustrate in a transparent way the impact of the supplier's updated belief on the information rent offered to the retailer. Specifically, we illustrate how much the supplier gains by improving her belief before offering contracts. We refer the reader to Section EC.1.1 of the e-companion for the analytical derivations supporting our discussion in this section. The supplier offers two contracts $(\underline{S}^*, \underline{P}^*)$ and $(\overline{S}^*, \overline{P}^*)$ if she decides to screen. The retailer's expected profit (see Equation (11) and Section EC.1.1 of the e-companion) in the two-point prior case simplifies to

$$\mathbb{E}_\xi[\Pi^r(\xi)] = \underbrace{\frac{(r-w)p}{1-\alpha}\int_0^{\underline{S}^*}\left(G(z|\underline{\xi}) - G(z|\overline{\xi})\right)dz}_{\text{Information Rent}} + \overline{\Pi}^r_{\min}(x, S^o_p, \underline{\xi}),$$

where $\underline{S}^* = \max\{x, S^o_p, \hat{\underline{S}}\}$ and $e^{-\frac{\hat{\underline{S}}}{\underline{\xi}}} - e^{-\frac{\hat{\underline{S}}}{\overline{\xi}}}\frac{p(r-w)}{p(r-w)+(1-p)(r+h-\alpha c)} = \frac{(1-p)(h+c(1-\alpha))}{p(r-w)+(1-p)(r+h-\alpha c)}$.

By delaying offering the contracts, the supplier potentially improves her knowledge about market conditions: that is, p tends toward zero or one (ρ tends to one). Figure 4 shows that the expected information rent, which is the difference between $\mathbb{E}_\xi[\Pi^r(\xi)] - \overline{\Pi}^r_{\min}(x, S^o_p, \underline{\xi})$, first increases and then decreases with p. Thus, by improving her belief, the supplier reduces the information rent offered to the retailer.

Figure 4. Retailer's Expected Profit, Expected Information Rent, and Reservation Profit as a Function of the Supplier's Belief When $w = 6$; $\overline{\xi} = 12$ with $x = 0$

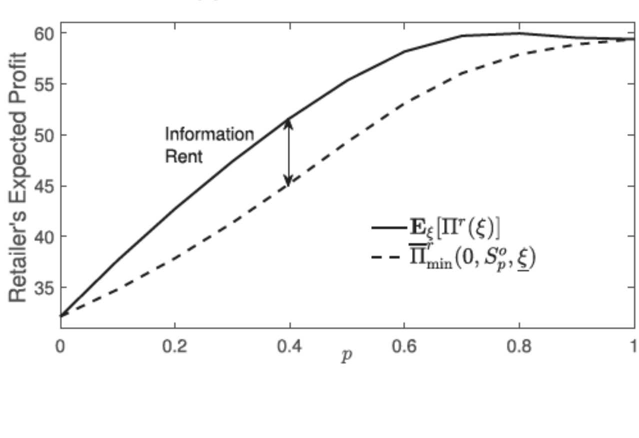

6. Conclusion

From the supply chain point of view, there are two compelling reasons to adopt VMI. First, there is a single centralized inventory manager for the supply chain who is closer to the upstream. Second, the inventory manager has access to the periodic point-of-sale information. Both of these reasons significantly combat the well-documented bullwhip effect in supply chains. However, there remain unresolved issues, such as credibly sharing demand information beyond POS

Figure 5. Comparison of Inventory Management Approaches in VMI

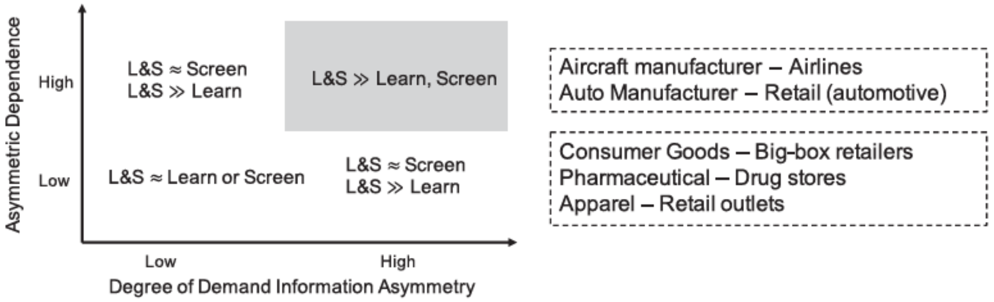

Notes. In the figure, Learn, Screen, and L&S denote the value of implementing learn, screen, and learn-and-screen approaches, respectively. In the learn approach, the supplier relies on the POS data to learn about the demand and make inventory decisions. In the screen approach, the supplier offers the proposed contracts in the first possible period of the planning horizon. L&S ≈ Screen means that the learn-and screen approach is approximately equivalent to screening the retailer early on. The supplier's asymmetric dependence in the examples listed is determined as the ratio of the retailer's average (industry-wide) gross margin to the supplier's average (industry-wide) gross margin of publicly traded firms.

data, between the supply chain members in an ongoing and dynamic VMI agreement. As a result, important demand information that is observed locally at the retail store could be ignored by the VMI manager at the supplier's site when making inventory decisions. This lack of credible communication between VMI members has led to tensions and eventually, falling out of VMI agreements. We propose and characterize a dynamic learn-and-screen approach, which addresses this key issue and provides a channel for credible communication in an ongoing VMI agreement. We note that, to implement the learn-and-screen approach, the VMI manager requires minimal additional monitoring effort to oversee the approach.

In the proposed learn-and-screen approach, the supplier learns about market conditions via POS data and then, offers a menu of screening contracts to the retailer. The retailer communicates his private demand information by choosing the base stock level from the menu of contracts that is most appropriate given his private information about the market conditions. In exchange for maintaining inventory at the mutually agreed on level, a one-time fee (or equivalently, a per period payment) is exchanged between the two firms. The proposed mechanism takes into account the retailer's market power, which evolves endogenously to the mechanism. Thus, offering the menu of contracts ensures that both the supplier and the retailer are better off from the status quo VMI agreement.

Our analysis provides structural insights about the impact of learning from POS data on the design of the contracts. In particular, two aspects of the learning process—occurrence of a stockout event and the magnitude of the sales data—determine the evolution of the supplier's belief process and hence, the incentives offered to the retailer to share his demand information credibly. Thus, the learn-and-screen mechanism allows us to explore the dynamic interplay between inventory decisions and evolution of incentives—highlighting the strategic aspect of inventory management.

The value of implementing the learn-and-screen inventory management approach in an ongoing VMI is summarized in Figure 5. The upstream inventory manager should consider the supplier's *asymmetric dependence* and the degree of demand information asymmetry to assess the value of various inventory management approaches within a VMI agreement. The supplier's asymmetric dependence is defined as the difference between the supplier's dependence on the retailer and the retailer's dependence on the supplier in the mutual partnership (Brinkhoff et al. 2015). High asymmetric dependence implies that the supplier is highly dependent on the retailer, resulting in the retailer capturing a larger share of the total profit margin.

For supply chains with low asymmetric dependence, the supplier benefits from updating the ongoing VMI operations by *immediately* eliciting the retailer's private demand information (via the proposed contracts). Examples of such a scenario include VMI agreements between big box retailers and consumer goods manufacturers or between pharmaceutical companies and drug stores for commonly prescribed drugs. In these supply chains, when a new product or a drug is being introduced into the market or when the retailer plans to run store promotions, the supplier can credibly elicit the retailer's expectation of the upcoming demand in exchange for a payment and in turn, make appropriate inventory decisions.

For supply chains with high asymmetric dependence (e.g., airline manufacturers and commercial airlines) the supplier is better off updating her ongoing VMI operations by eliciting the retailer's private demand information via the (dynamic) learn-and-screen approach. The timing of contracts in the learn-and-screen approach depends on the degree of information asymmetry in the supply chain. When the degree of information asymmetry is low, offering the contracts immediately alleviates the effect of double marginalization. The low degree of information asymmetry implies that the supplier offers lower information rent to update the VMI agreement. As the degree of information asymmetry grows, the supplier benefits significantly by learning from POS data and dynamically deciding the timing of the contracts.

Acknowledgments

The authors thank Gad Allon, the anonymous associate editor, and two anonymous reviewers for their constructive and helpful suggestions. The discussions during the authors' presentations at the following conferences and colloquia were also beneficial: the 2015 INFORMS Annual Meeting in Philadelphia, the 2015 MSOM Annual Conference in Toronto, the 2015 POMS Annual Conference in Washington, DC, North Carolina State University, Nanyang Technological University, the University of Utah.

The first two authors' names for this article are listed in alphabetical order.

Appendix A. Glossary of Notation

Cost parameters	Description	Demand parameters	Description
r	Unit retail price	ξ	The retailer's private demand information
w	Unit wholesale price	$\Theta := [\underline{\xi}, \overline{\xi}]$	Possible values that ξ takes
c	Unit cost of production	$\pi_n(\cdot)$	The supplier's belief (pdf) at beginning of each period $n \geq 1$
		$\mathcal{A}$	The set of all pdfs on Θ
h	Unit holding cost per period	$g(\cdot\|\xi), G(\cdot\|\xi)$	pdf and cdf of demand, respectively
α	Discount factor	$q_n(\cdot), Q_n(\cdot)$	Predictive demand density and distribution at the beginning of each period $n \geq 1$, respectively

Profit functions	Description
$\Pi^s(S, P)$	The supplier's expected profit if the retailer chooses contract $\{S, P\}$ from the menu
$\Pi^r(S, P, \xi)$	Type ξ retailer's expected profit if he chooses contract $\{S, P\}$ from the menu
$\Pi^{tot}(S)$	Expected total supply chain profit if the retailer chooses base stock level S from the menu
$\Pi^r_{\min}(x, y, \xi)$	Type ξ retailer's expected profit if he rejects the menu of contracts
$\Pi^r_{\min}(x, S, \xi)$	Type ξ retailer's expected profit if he rejects the menu of contracts and the supplier follows the base stock policy with the base stock level S
$\overline{\Pi}^r_{\min}(x, S, \xi)$	Upper bound on type ξ retailer's reservation profit $\Pi^r_{\min}(x, S, \xi)$
$\underline{\Pi}^r_{\min}(S, \xi)$	Lower bound on type ξ retailer's reservation profit $\Pi^r_{\min}(x, S, \xi)$
$\tilde{V}(x, \pi)$	The supplier's maximum expected profit obtained from using the learn-and-screen approach starting the planning horizon with x, π
$\tilde{V}^{lr}(x, \pi)$	The supplier's maximum expected profit obtained from using the learn approach with initial conditions x, π
$\tilde{V}^{sr}(x, \pi)$	The supplier's maximum expected profit obtained from offering the menu of screening contracts at the beginning of the planning horizon

Decision variables	Description
τ	Stopping time with respect to the filtration of the observed sales
$\mathbf{y} = (y_1, y_2, \ldots, y_{\tau-1})$	Order up to levels used by the supplier in the status quo VMI agreement before offering the contracts
$\mathbf{y}^o = (y^o_\tau, y^o_{\tau+1}, \ldots,)$	Order up to levels used by the supplier in the status quo VMI agreement if the menu of contracts is rejected
$\mathcal{M}$	Denotes the set of all admissible $(\mathbf{y}, \tau)$
$\{S^*(\cdot), P^*(\cdot)\}$	Optimal menu of contracts offered
$S^o(\pi)$	Base stock level used by the supplier if the menu of contracts is rejected
$S^{fb}(\cdot)$	Optimal base stock levels under symmetric demand information using the linear coordinating contract
p^{fb}	Coordinating price under symmetric information

Appendix B. Simplification of the DP

Using the definitions of transformations V and Π^{sr}, we can simplify $\tilde{L}(y,\pi)$ as follows: $\tilde{L}(y,\pi):=(w-c)y-(w+h)\cdot\int_0^y Q(z)dz+\alpha\overline{Q}(z)\tilde{V}(0,\pi^c(\cdot|y))+\alpha\int_0^y q(z)\tilde{V}(y-z,\pi^e)dz=(w-c)\cdot y-(w+h)\int_0^y Q(z)dz+\alpha\cdot\overline{Q}(z)V(0,\pi^c(\cdot|y))+\alpha\int_0^y q(z)[V(y-z,\pi^e)+c(y-z)]dz=(w-c)y-(w+h)\int_0^y Q(z)+\alpha\overline{Q}(z)V(0,\pi^c(\cdot|y))+\alpha\int_0^y q(z)V(y-z,\pi^e)dz+\alpha c\int_0^y q(z)(y-z)dz=(w-c)y-(w+h-\alpha c)\int_0^y Q(z)dz+\alpha\bar{Q}(z)V(0,\pi^c(\cdot|y))+\alpha\int_0^y q(z)V(y-z,\pi^e)dz =: L(y,\pi)$.

Appendix C. Proofs

Proof of Theorem 1. First, we establish the lower bound: $\underline{\Pi}^r_{\min}(x,S,\xi)=(r-w)\mathbb{E}[\sum_{i=1}^{\infty}\alpha^{i-1}\min\{y_i^o,D_i(\xi)\}|x_1=x]=(r-w)\cdot\mathbb{E}[\sum_{i=1}^{\infty}\alpha^{i-1}\min\{\max\{x_i,S\},D_i(\xi)\}|x_1=x]\ge(r-w)\mathbb{E}[\sum_{i=1}^{\infty}\alpha^{i-1}\cdot\min\{S,D_i(\xi)\}|x_1=x]=\frac{\int_0^S\overline{G}(z|\xi)dz}{1-\alpha}$. We establish the upper bound by first providing a successive approximation scheme that monotonically (increasingly) converges to $\Pi^r_{\min}(x,S,\xi)$. Next, we show that $\overline{\Pi}^r_{\min}$ is an upper bound to each function in the approximation scheme and to the limiting function—establishing the necessary result. Let $\mathcal{B}:=\{w:\mathbb{R}^+\to\mathbb{R}^+|\sup|w(x)|<\infty\}$. For every $\xi\in\Theta$ and base stock level S, $\Pi^r_{\min}(x,S,\xi)\in\mathcal{B}$, which follows from the above lower bound and the upper bound given by $\Pi^r_{\min}(x,S,\xi)=(r-w)\sum_{i=1}^{\infty}\alpha^{i-1}\cdot\mathbb{E}_{D(\xi)}[\min\{y_i^o,D_i(\xi)\}|x_1=x]\le(r-w)\sum_{i=1}^{\infty}\alpha^{i-1}\mathbb{E}_{D(\xi)}[D_i(\xi)]=\frac{(r-w)}{1-\alpha}\mathbb{E}[D(\xi)]<\infty$.

Next, we define an operator $\mathbf{T}^r_{S,\xi}$ acting on functional space $\mathcal{B}$. For a fixed $\xi\in\Theta$ and S, $(\mathbf{T}^r_{S,\xi}w)(x):=(r-w)\cdot\int_0^y\overline{G}(z|\xi)dz+\alpha w(0)\overline{G}(y|\xi)+\alpha\int_0^y w(y-z)g(z|\xi)dz$, where $y=\max\{x,S\}$.

Lemma C.1. *The operator $\mathbf{T}^r_{S,\xi}$ acting on metric space $\mathcal{B}$ satisfies the following properties for any $\xi\in\Theta$ and $S\in\mathbb{R}^+$.*

(1) *For $w\in\mathcal{B}$, $\mathbf{T}^r_{S,\xi}w\in\mathcal{B}$.*

(2) *For $w_1,w_2\in\mathcal{B}$ with $w_1\le w_2$, we have $\mathbf{T}^r_{S,\xi}w_1\le\mathbf{T}^r_{S,\xi}w_2$.*

(3) *$\mathbf{T}^r_{S,\xi}$ is a contraction mapping on $(\mathcal{B},\|\cdot\|_\infty)$, where $\|w\|_\infty:=\sup_{x\in\mathbb{R}^+}|w(x)|$.*

Proof of Lemma C.1. Parts 1 and 2 can be directly verified from definition of $\mathbf{T}^r_{S,\xi}$. Let $w_1,w_2\in\mathcal{B}$; then, for any ξ,S, $\mathbf{T}^r_{S,\xi}w_1-\mathbf{T}^r_{S,\xi}w_2=\alpha\overline{G}(y|\xi)(w_1(0)-w_2(0))+\alpha\int_0^y(w_1(y-z)-w_2\cdot(y-z))g(z|\xi)dz\le\alpha\overline{G}(y|\xi)\|w_1-w_2\|_\infty+\alpha\int_0^y\|w_1-w_2\|_\infty g(z|\xi)dz\le\alpha\overline{G}(y|\xi)\|w_1-w_2\|_\infty+\alpha G(y|\xi)\|w_1-w_2\|_\infty\le\alpha\|w_1-w_2\|_\infty$. Because $\alpha<1$, $\mathbf{T}^r_{S,\xi}$ is a contraction mapping. □

Note that $(\mathcal{B},\|\cdot\|_\infty)$ is a complete metric (Banach) space. It follows from the Banach Fixed Point Theorem that there exists a *unique* fixed point $w^*_{S,\xi}\in\mathcal{B}$ satisfying $w^*_{S,\xi}=\mathbf{T}^r_{S,\xi}w^*_{S,\xi}$. In fact, $w^*_{S,\xi}=\Pi^r_{\min}(x,S,\xi)$. Furthermore, the successive approximation scheme

$$w^{k+1}_{S,\xi}:=\begin{cases}0, & \text{if } k=0\\ \mathbf{T}^r_{S,\xi}w^k_{S,\xi}, & \text{if } k\ge1;\end{cases}$$

monotonically (increasingly) converges to $w^*_{S,\xi}$.

In this scheme, note that $w^1_{S,\xi}\le w^2_{S,\xi}$. From Part 2 of Lemma C.1, it follows that $w^k_{S,\xi}\uparrow w^*_{S,\xi}\equiv\Pi^r_{\min}(\cdot,S,\xi)$. We prove that $\overline{\Pi}^r_{\min}$ is an upper using induction arguments. For a fixed S,ξ and $k=2$, we have $w^2_{S,\xi}(x)=(r-w)\cdot\int_0^{\max\{x,S\}}\overline{G}(z|\xi)dz=(r-w)\left[\int_0^S\overline{G}(z|\xi)dz+\int_S^{\max\{x,S\}}\overline{G}(z|\xi)dz\right]\le(r-w)\left[\int_0^S\overline{G}(z|\xi)dz+[x-S]^+\right]\le\overline{\Pi}^r_{\min}(x,S,\xi)$.

Suppose that $w^k_{S,\xi}\le\overline{\Pi}^r_{\min}(\cdot,S,\xi)$. We prove that the inequality holds for $k+1$: $w^{k+1}_{S,\xi}(x)=(\mathbf{T}^r_{S,\xi}w^k_{S,\xi})(x)=(r-w)\cdot\int_0^y\overline{G}(z|\xi)dz+\alpha w^k_{S,\xi}(0)\overline{G}(y|\xi)+\alpha\int_0^y w^k_{S,\xi}(y-z)g(z|\xi)dz$ (note that $y=\max\{x,S\}$) $\le(r-w)\int_0^y\overline{G}(z|\xi)dz+\frac{\alpha(r-w)}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz\cdot\overline{G}(y|\xi)+\alpha\int_0^y(r-w)\left\{(y-z-S)^++\frac{1}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz\right\}g(z|\xi)dz=(r-w)\int_0^y\overline{G}(z|\xi)dz+\frac{\alpha(r-w)}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz\cdot\overline{G}(y|\xi)+\alpha(r-w)\int_0^y(y-z-S)^+g(z|\xi)dz+\frac{(r-w)\alpha}{1-\alpha}\int_0^S\overline{G}(z|\xi)dzG(y|\xi)=(r-w)\{\int_0^y\overline{G}(z|\xi)dz+\alpha\cdot\int_0^{y-S}(y-S-z)g(z|\xi)dz\}+\frac{\alpha(r-w)}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz=(r-w)\left\{\int_0^y\overline{G}(z|\xi)dz+\alpha(y-S)G(y-S|\xi)-\alpha\left[(y-S)G(y-S|\xi)-\int_0^{y-S}G(z|\xi)dz\right]\right\}+\frac{\alpha(r-w)}{1-\alpha}\cdot\int_0^S\overline{G}(z|\xi)dz=(r-w)\left\{\int_0^S\overline{G}(z|\xi)dz+\int_S^y\overline{G}(z|\xi)dz+\alpha\int_0^{y-S}G(z|\xi)dz\right\}+\frac{\alpha(r-w)}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz=(r-w)\left\{\int_S^y\overline{G}(z|\xi)dz+\alpha\int_0^{y-S}G(z|\xi)dz\right\}+\frac{(r-w)}{1-\alpha}\cdot\int_0^S\overline{G}(z|\xi)dz=(r-w)\left\{(\max\{x,S\}-S)-\int_S^y G(z|\xi)dz+\alpha\int_0^{y-S}G(z|\xi)\cdot dz\right\}+\frac{(r-w)}{1-\alpha}\int_0^S\overline{G}(z|\xi)dz=\overline{\Pi}^r_{\min}(x,S,\xi)-(r-w)(\int_S^y G(z|\xi)dz-\alpha\int_0^{y-S}G(z|\xi)dz)=\overline{\Pi}^r_{\min}(x,S,\xi)-M(x)$, where $M(x):=\begin{cases}0, & \text{if } x\le S;\\ \int_S^x G(z|\xi)dz-\alpha\int_0^{x-S}G(z|\xi)dz, & \text{if } x>S.\end{cases}$ Note that $M(x)>0$ for all $x>S$, because $\frac{d}{dx}M(x)=G(x|\xi)-\alpha G(x-S|\xi)\ge G(x|\xi)-G(x-S|\xi)\ge0$. Thus, $w^k_{S,\xi}(x)\le\overline{\Pi}^r_{\min}(x,S,\xi)$, $\forall k$. It follows then, for every x,S,ξ, that the limit point $w^*_{S,\xi}(x)\le\overline{\Pi}^r_{\min}(x,S,\xi)$, which concludes the proof. □

Proof of Lemma 1. It follows from the definition of $\overline{\Pi}^r_{\min}$ in Theorem 1 that $\frac{\partial}{\partial\xi}\overline{\Pi}^r_{\min}(x,S^o_\pi,\xi)=-\frac{(r-w)}{1-\alpha}\int_0^{S^o_\pi}\frac{\partial}{\partial\xi}G(z|\xi)\,dz$. Next, we show that, if the conditions stated in the lemma are satisfied, then the profit obtained by accepting a contract grows faster than the upper bound on the outside option. For any contract S,P, the type ξ retailer's profit is given by $\Pi^r(S,P,\xi)=\dfrac{1}{1-\alpha}\{\mathbb{E}[(r-w)\min\{S,D_i\}]-P\}=\dfrac{1}{1-\alpha}\Big\{(r-w)\cdot\left[S-\int_0^S G(z|\xi)dz\right]-P\Big\}$. The (IC) constraint and the envelope theorem together imply that

$$\begin{aligned}\frac{\partial}{\partial\xi}\Pi^r(S(\xi),P(\xi),\xi)&=\frac{\partial}{\partial\xi}\Pi^r(S(\eta),P(\eta),\xi)|_{\eta=\xi}\\&=-\frac{(r-w)}{1-\alpha}\int_0^{S(\xi)}\frac{\partial}{\partial\xi}G(z|\xi)\,.\end{aligned}\tag{C.1}$$

If $S(\xi)\ge S^o(\pi)$, then, $\frac{\partial}{\partial\xi}\Pi^r(S(\xi),P(\xi),\xi)\ge\frac{\partial}{\partial\xi}\overline{\Pi}^r_{\min}(x,S^o_\pi,\xi)$. In addition, if $\Pi^r(S(\underline{\xi}),P(\underline{\xi}),\underline{\xi})=\overline{\Pi}^r_{\min}(x,S^o_\pi,\underline{\xi})$, then $\Pi^r(S(\xi),P(\xi),\xi)\ge\overline{\Pi}^r_{\min}(x,S^o_\pi,\xi)\ge\Pi^r_{\min}(x,S^o_\pi,\xi)$, $\forall\xi$, where the first inequality follows, because Π^r grows faster than $\overline{\Pi}^r_{\min}$, and the second inequality follows from Theorem 1. □

Proof of Lemma 2. We first prove that (IC) implies Lemma 2, (i) and (ii). It follows from Equation (C.1) and the participation constraint (which is binding for the lowest type) that $\Pi^r(\xi)=\Pi^r(\underline{\xi})-\frac{(r-w)}{1-\alpha}\int_{\underline{\xi}}^{\xi}\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\,dz\,d\eta=\overline{\Pi}^r_{\min}(x,S^o_\pi,\underline{\xi})-\frac{(r-w)}{1-\alpha}\int_{\underline{\xi}}^{\xi}\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\,dz\,d\eta$. To prove (ii), note that $\frac{d^2\Pi^r(S,P,\xi)}{dS^2}=-\frac{(r-w)}{1-\alpha}g(S|\xi)<0$ and that $\frac{d^2\Pi^r(S,P,\xi)}{dSd\xi}=-\frac{(r-w)}{1-\alpha}\frac{\partial G(S|\xi)}{d\xi}>0$. Then, $S'(\cdot)>0$. Otherwise, for $\xi_1>\xi_2$, $0=\frac{\partial\Pi^r(S,P,\xi_1)}{\partial S}|_{S=S(\xi_1)}>\frac{\partial\Pi^r(S,P,\xi_1)}{\partial S}|_{S=S(\xi_2)}>\frac{\partial\Pi^r(S,P,\xi_2)}{\partial S}|_{S=S(\xi_2)}$. The first equality follows from (IC) and the two inequalities because of the signs of the second derivatives. The last inequality contradicts (IC), because $S(\xi_2)$ is the maximizer for type ξ_2.

To prove (i) and (ii), imply (IC): $\Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\xi) = \int_{\underline{\xi}}^{\xi} \frac{\mathrm{d}}{\mathrm{d}x}\Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),x)\,\mathrm{d}x + \Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\underline{\xi}) = \int_{\underline{\xi}}^{\tilde{\xi}} \frac{\mathrm{d}}{\mathrm{d}x}\Pi^r(S(\tilde{\xi}),P(\tilde{\xi}), x)\,\mathrm{d}x + \int_{\tilde{\xi}}^{\xi} \frac{\mathrm{d}}{\mathrm{d}x}\Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),x)\,\mathrm{d}x + \Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\underline{\xi}) = \Pi^r(S(\tilde{\xi}), P(\tilde{\xi}),\tilde{\xi}) - \Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\underline{\xi}) + \int_{\tilde{\xi}}^{\xi} \frac{\mathrm{d}}{\mathrm{d}x}\Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),x)\,\mathrm{d}x + \Pi^r(S(\tilde{\xi}), P(\tilde{\xi}),\underline{\xi}) = \Pi^r(S(\tilde{\xi}), P(\tilde{\xi}), \tilde{\xi}) - \frac{(r-w)}{1-\alpha}\int_{\tilde{\xi}}^{\xi}\int_0^{S(\tilde{\xi})}\frac{\delta}{\delta x}G(z|x)\mathrm{d}z\,\mathrm{d}x = \Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\tilde{\xi}) - \frac{(r-w)}{1-\alpha}\Big(\int_{\tilde{\xi}}^{\xi}\Big[\int_0^{S(\tilde{\xi})}\frac{\delta}{\delta x}G(z|x)\mathrm{d}z - \int_0^{S(x)}\frac{\delta}{\delta x}G(z|x)\cdot \mathrm{d}z\Big]\mathrm{d}x - \int_{\tilde{\xi}}^{\xi}\int_0^{S(x)}\frac{\delta}{\delta x}G(z|x)\mathrm{d}z\mathrm{d}x\Big) = \Pi^r(S(\tilde{\xi}),P(\tilde{\xi}),\tilde{\xi}) - \frac{(r-w)}{1-\alpha}\int_{\tilde{\xi}}^{\xi}\Big[\int_0^{S(\tilde{\xi})}\cdot \frac{\delta}{\delta x}G(z|x)\mathrm{d}z - \int_0^{S(x)}\frac{\delta}{\delta x}G(z|x)\mathrm{d}z\Big]\mathrm{d}x + \Pi^r(S(\xi), P(\xi),\xi) - \Pi^r(S(\tilde{\xi}), P(\tilde{\xi}),\tilde{\xi}) = \Pi^r(S(\xi),P(\xi),\xi) - \frac{(r-w)}{1-\alpha}\int_{\tilde{\xi}}^{\xi}\int_{S(x)}^{S(\tilde{\xi})}\frac{\delta}{\delta x}G(z|x)\mathrm{d}z\mathrm{d}x$. The last equality implies (IC). When $\xi > \tilde{\xi}$, $S(x) > S(\tilde{\xi})$ for all x in $(\tilde{\xi},\xi)$. Hence, the value of the second integral is positive; thus, it is not optimal for the retailer of type ξ to choose $S(\tilde{\xi}),P(\tilde{\xi})$. By a similar argument, one can rule out the case $\xi < \tilde{\xi}$. □

Proof of Lemma 3. If the type ξ retailer accepts the contract $S(\xi), P(\xi)$, then $y_j = S(\xi)$, for all $j \geq 1$. The total surplus generated in the supply chain is (we denote $S(\xi) = S$) $\sum_{j=1}^{\infty}\alpha^{j-1}\cdot [r\min\{S, D_j\} - c(S - x_j) - h(S - D_j)^+] = cx + \sum_{j=1}^{\infty}\alpha^{j-1}[(r-c)S - (r+h)[S-D_j]^+] + c\sum_{j=n}^{\infty}\alpha^j[S-D_j]^+ = cx + \sum_{j=1}^{\infty}\alpha^{j-1}[(r-c)S-(r+h-\alpha c)[S-D_j]^+]$. The total expected profit of the supply chain is

$$
\begin{aligned}
&\Pi^{tot}(x, S(\xi), \xi) \\
&= cx + \sum_{j=1}^{\infty}\alpha^{j-1}\mathbb{E}_{D(\xi)}[(r-c)S(\xi) - (r+h-\alpha c)[S(\xi) - D_j]^+] \\
&= cx + \frac{1}{1-\alpha}\left[(r-c)S(\xi) - (r+h-\alpha c)\int_0^{S(\xi)} G(z|\xi)\mathrm{d}z\right].
\end{aligned} \tag{C.2}
$$

When the retailer's type is ξ, the supplier's profit for any incentive-compatible menu $\{S(\cdot), P(\cdot)\}$ is the difference in profits, $\Pi^{tot}(x, S(\xi), \xi) - \Pi^r(\xi)$. Using Equations (11) and (C.2), the supplier's expected profit from screening the retailer is given by $\mathbb{E}_\xi[\Pi^{tot}(x, S(\xi), \xi) - \Pi^r(\xi)] = cx - \overline{\Pi}^r_{\min}(x, S^o_\pi, \underline{\xi}) + \frac{1}{1-\alpha}\mathbb{E}_\xi\Big[(r-c)S(\xi) - (r+h-\alpha c)\int_0^{S(\xi)}G(x|\xi)\,\mathrm{d}x + (r-w) \int_{\underline{\xi}}^{\xi}\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\mathrm{d}z\,\mathrm{d}\eta\Big]$. The following simplification is needed for the last term: $\int_{\underline{\xi}}^{\overline{\xi}}\Big(\overbrace{\int_{\underline{\xi}}^{\xi}\mathrm{d}\eta\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\mathrm{d}z}^{u}\Big)\overbrace{\pi(\xi)\mathrm{d}\xi}^{\mathrm{d}v} = \int_{\underline{\xi}}^{\xi}\pi(\eta)\mathrm{d}\eta\int_{\underline{\xi}}^{\xi}\mathrm{d}\eta\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\mathrm{d}z\Big|_{\underline{\xi}}^{\overline{\xi}} - \int_{\underline{\xi}}^{\overline{\xi}}\int_{\underline{\xi}}^{\xi}\pi(\eta)\mathrm{d}\eta\int_0^{S(\xi)}\frac{\partial}{\partial\xi}G(z|\xi)\cdot \mathrm{d}z\mathrm{d}\xi = \int_{\underline{\xi}}^{\overline{\xi}}\mathrm{d}\eta\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)\mathrm{d}z\mathrm{d}\eta - \int_{\underline{\xi}}^{\overline{\xi}}\int_{\underline{\xi}}^{\xi}\pi(\eta)\mathrm{d}\eta\int_0^{S(\xi)}\frac{\partial}{\partial\xi}G(z|\xi)\cdot \mathrm{d}z\mathrm{d}\xi = \int_{\underline{\xi}}^{\overline{\xi}}\int_{\xi}^{\overline{\xi}}\pi(\eta)\,\mathrm{d}\eta\int_0^{S(\xi)}\frac{\partial}{\partial\xi}G(z|\xi)\mathrm{d}z\mathrm{d}\xi$. □

Lemma C.2. *Let $f : [a, b] \to \mathbb{R}$ be a twice differentiable function over (a, b). If the following two conditions hold: (i) $\exists\hat{x} \in (a, b)$ such that $f_{xx} < 0$ for all $x < \hat{x}$, $f_{xx} > 0$ for all $x > \hat{x}$, and $f_{xx}(\hat{x}) = 0$ and (ii) $f_x(a+) > 0$ and $f_x(b-) < 0$, the following statements are true:*

(1) f_x has a unique root x^* in (a, b) such that $x^* < \hat{x}$.

(2) The function f is unimodal with peak at $x = x^*$.

Proof of Lemma C.2.

(1) Because $f_x(a) > 0$, $f_x(b) < 0$, and f_x is continuous, there exists at least one root of f_x in (a, b). Suppose that η_1, η_2 are two roots, such that $\eta_1 < \eta_2$ without loss of generality. Then, $f_x(\eta_1) = f_x(\eta_2) = 0$. By the Mean Value Theorem, there exists $u \in (\eta_1, \eta_2)$ such that $f_{xx}(u) = 0$. There are three possibilities: (i) $u > \hat{x}$, (ii) $u < \hat{x}$, or (iii) $u = \hat{x}$. The first two possibilities result in a contradiction, because f_{xx} has a unique root. Hence, $u = \hat{x}$. This implies that $\eta_1 < \hat{x} < \eta_2$. The function $f_{xx} > 0$ for $x > \hat{x}$. Thus, $f_x(y) > f_x(\eta_2) = 0$ for all $y > \eta_2$. Taking limit as $y \to b$, we get $f_x(b-) \geq 0$, which contradicts assumption (ii) in the lemma. Thus, there exists a unique root for f_x denoted by x^*. Suppose that $x^* \geq \hat{x}$. Using assumption (ii), we get $f_x(y) > f_x(x^*) = 0$ for all $y > x^*$. Taking limit on $y \to b$, we get a contradiction. Hence, $x^* < \hat{x}$.

(2) Because $f_x(a+) > 0$ by assumption (ii), it follows from part (i) of the proof that $f_x > 0$ for all $x < x^*$, $f_x < 0$ for $x > x^*$, and $f_x(x^*) = 0$. This clearly implies that the function f is a unimodal function with its maximum attained at $x^* < \hat{x}$. □

Proof of Theorem 2.

(1) Differentiating $G(z|\xi)$ wrt ξ yields $\frac{\partial G(z|\xi)}{\partial\xi} = \frac{\partial}{\partial\xi}\int_0^z g(y|\xi)\mathrm{d}y = \int_0^z\frac{\partial}{\partial\xi}g(y|\xi)\mathrm{d}y$. The result hence is determined by the behavior of $\frac{\partial}{\partial\xi}g(y|\xi)$ as a function of y. $\frac{\partial}{\partial\xi}g(y|\xi)$ changes sign at least once, because any two density functions cross each other (and the total area under any pdf is one). If the sign change happens exactly once and is from negative to positive, then first-order stochastic dominance follows (see Lemma C.3 below). For a pdf from the exponential family, $\frac{\partial}{\partial\xi}g(y|\xi) = k(y)e^{-t(y)s(\xi)}\cdot [l'(\xi) - l(\xi)t(y)s'(\xi)]$. Because $l', s' < 0$ and $t' > 0$ (by assumption), it follows that $[l'(\xi) - l(\xi)t(y)s'(\xi)]$ is initially negative, becomes 0 at $t(y^*(\xi)) := \frac{l'(\xi)}{l(\xi)s'(\xi)} > 0$, and remains positive thereafter. Note that, because t is increasing and $t(0) = 0$, we have $y^*(\xi) > 0$. Lemma C.3 guarantees that $\int_0^z\frac{\partial}{\partial\xi}g(y|\xi)\mathrm{d}y \leq 0$ for all $\xi \in \Theta$ and $z \in \mathbb{R}^+$, which in turn, implies first-order stochastic dominance of the family of demand distributions considered in Theorem 2.

Lemma C.3. *Let $f(\cdot)$ be a continuous function on $\mathbb{R}^+$ such that $f(x) < 0$, for all $x < \hat{x}$, $f(\hat{x}) = 0$, and $f(x) > 0$, for all $x > \hat{x}$. Define $b(z) := \int_0^z f(y)\mathrm{d}y$. If $\lim_{z\to\infty} b(z) = 0$, then $b(z) \leq 0$ for all $z \in \mathbb{R}^+$.*

Proof of Lemma C.3. Suppose not (i.e., $\exists\tilde{z} \in \mathbb{R}^+$) such that $b(\tilde{z}) > 0$. Then, $\tilde{z} > \hat{x}$; otherwise, $b(\tilde{z}) \leq 0$, because $f(\cdot)$ is negative until $\hat{x}$. This implies that $b(\tilde{z}) \leq b(\tilde{z}) + \int_{\tilde{z}}^{\infty} f(x)\,\mathrm{d}x = \int_0^{\infty} f(x)\,\mathrm{d}x = b(\infty) = 0$. The first inequality follows, because $f(x)$ is strictly positive for $x > \tilde{z} > \hat{x}$, and the last equality follows by the assumption in the lemma. Thus, we have arrived at a contradiction. Therefore, $b(z) \leq 0$ for all $z \in \mathbb{R}^+$. □

Note that the assumptions in the above lemma are satisfied by $\frac{\partial}{\partial\xi}g(x|\xi)$, because $\int_0^{\infty}\frac{\partial}{\partial\xi}g(x|\xi)\mathrm{d}x = 0$.

2. The first derivative $\frac{\partial}{\partial S}H(S,\xi)=(r-c)-(r+h-\alpha c)\cdot G(S|\xi)+\frac{r-w}{\lambda(\xi)}\frac{\partial}{\partial \xi}G(S|\xi)$. Note that $\lim_{z\to\infty}\frac{\partial G(z|\xi)}{\partial \xi}=\lim_{z\to\infty}\frac{\partial}{\partial \xi}\cdot\int_0^z g(y|\xi)\,dy=\lim_{z\to\infty}\int_0^z\frac{\partial}{\partial \xi}g(y|\xi)\,dy=\int_0^\infty\frac{\partial}{\partial \xi}g(y|\xi)\,dy=0$. The last equality follows, because $\int_0^\infty g(z|\xi)\,dz=1, \forall\,\xi\in\Theta\Rightarrow\frac{\partial}{\partial \xi}\int_0^\infty g(z|\xi)\,dz=0\Rightarrow\int_0^\infty\frac{\partial}{\partial \xi}g(z|\xi)\,dz=0,\ \forall\,\xi\in\Theta$. Thus, $\frac{\partial}{\partial S}H(0,\xi)=r-c>0$ and $\frac{\partial}{\partial S}H(\infty,\xi)=-c(1-\alpha)-h<0$. Because $H(\cdot,\xi)$ is continuous, there exists at least one critical point where it equals zero. The second derivative $\frac{\partial^2}{\partial S^2}H(S,\xi)=-(r+h-\alpha c)\cdot\left\{g(S|\xi)-\frac{\varpi}{\lambda(\xi)}\frac{\partial g(S|\xi)}{\partial \xi}\right\}$, where $\varpi:=\frac{r-w}{r+h-\alpha c}<1$. We evaluate the term inside the brackets on the right hand side for a density from the exponential family as $\left\{g(S|\xi)-\dfrac{\varpi}{\lambda(\xi)}\dfrac{\partial g(S|\xi)}{\partial \xi}\right\}=k(S)e^{-t(S)s(\xi)}[l(\xi)-\frac{\varpi}{\lambda(\xi)}(l'(\xi)-l(\xi)t(S)s'(\xi))]$, where l', s' denote derivatives. Then, $\frac{\partial^2 H(S,\xi)}{\partial S^2}>0\iff t(S)>\frac{-\lambda_n(\xi)}{\varpi s'(\xi)}+\frac{l'(\xi)}{l(\xi)s(\xi)}\iff S>\bar{S}(\xi)$, where

$$\bar{S}(\xi):=t^{-1}\left(\frac{-\lambda(\xi)}{\varpi s'(\xi)}+\frac{l'(\xi)}{l(\xi)s(\xi)}\right). \tag{C.3}$$

The first inequality follows, because $s'<0$ and $l'<0$ by assumption. Because $t(\cdot)$ is increasing, t^{-1} exists and is monotone. Note that $\bar{S}(\xi)>0$ for all $\xi\in\Theta$. All of the conditions in Lemma C.2 are satisfied, and Part 2 follows.

3. It follows from unimodality of $H(\cdot,\xi)$ that $\frac{\partial H(\hat{S}(\xi),\xi)}{\partial S}=0$ characterizes the unique maximizer of $H(\cdot,\xi|\pi)$ and $\max\{x,S^o(\pi),\hat{S}(\xi)\}$ is the maximizer of the constrained problem. □

Proof of Theorem 3. Because $\hat{S}(\xi)$ satisfies $\frac{\partial}{\partial S}H(\hat{S}(\xi),\xi)=0$, for all $\xi\in\Theta$, we have

$$\frac{d}{d\xi}\frac{\partial}{\partial S}H(\hat{S}(\xi),\xi)=\frac{\partial^2}{\partial S^2}H(\hat{S}(\xi),\xi)\frac{d\hat{S}(\xi)}{d\xi}+\frac{\partial^2}{\partial S\partial\xi}H(\hat{S}(\xi),\xi)=0. \tag{C.4}$$

We first establish $\hat{S}(\xi)\le\bar{S}(\xi)$ (defined in Equation (C.3)). Because $s(\xi)=l(\xi)=\frac{1}{\xi}$ and $k(z)=1$ in the case of newsvendor family of distributions, the definition of $\bar{S}(\xi)$ simplifies to $\bar{S}(\xi)=t^{-1}\left(\frac{\xi^2\lambda(\xi)}{\varpi}+\xi\right)$. Thus,

$$\begin{aligned}\frac{\partial H(\bar{S}(\xi),\xi)}{\partial S}&=(r-c)-(r+h-\alpha c)\left(1-e^{-\left(\frac{\xi\lambda(\xi)}{\varpi}+1\right)}\right)\\&\quad-\frac{r-w}{\lambda(\xi)}e^{-\left(\frac{\xi\lambda(\xi)}{\varpi}+1\right)}\left(\frac{1}{\xi}+\frac{\lambda(\xi)}{\varpi}\right)\\&=-h-c(1-\alpha)\ -\frac{r-w}{\lambda(\xi)\xi}e^{-\left(\frac{\xi\lambda(\xi)}{\varpi}+1\right)}<0.\end{aligned}$$

It then follows from unimodality of $H(\cdot,\xi)$ that $\hat{S}(\xi)<\bar{S}(\xi)$. From Equation (C.3), it also follows that $\frac{\partial^2 H(\hat{S}(\xi),\xi)}{\partial S^2}<0$ for all ξ. Finally,

$$\begin{aligned}\frac{\partial^2 H(\hat{S}(\xi),\xi)}{\partial S\partial\xi}&=-(r+h-\alpha c)\frac{\partial G(\hat{S}|\xi)}{\partial\xi}\\&\quad+(r-w)\frac{d}{d\xi}\left(\frac{\int_{\underline{\xi}}^{\bar{\xi}}\pi(\eta)d\eta}{\pi(\xi)}\right)\frac{\partial G(\hat{S}|\xi)}{\partial\xi}\\&\quad+(r-w)\frac{\int_{\underline{\xi}}^{\bar{\xi}}\pi(\eta)d\eta}{\pi(\xi)}\frac{\partial^2}{\partial\xi^2}G(\hat{S}|\xi).\end{aligned}$$

If $G(\cdot|\xi)$ is from the newsvendor family, $\frac{\partial^2 G(\hat{S}|\xi)}{\partial\xi^2}=t(\hat{S})\cdot\xi^{-4}e^{\frac{-t(\hat{S})}{\xi}}(2\xi-t(\hat{S}))$.

Note that $2\xi-t(\hat{S}(\xi))\ge0\iff\hat{S}(\xi)\le t^{-1}(2\xi)\iff G(\hat{S}(\xi)|\xi)\le G(t^{-1}(2\xi)|\xi)$ for all $\xi\in\Theta$. From Equation (14), it follows that $G(\hat{S}(\xi)|\xi)\le\frac{(r-c)}{r+h-\alpha c}$, because $\frac{\partial G(\hat{S}(\xi)|\xi)}{\partial\xi}\le0$. Thus, if $G(t^{-1}(2\xi)|\xi)=1-e^{-2}\ge\frac{(r-c)}{r+h-\alpha c}$, then $2\xi-t(\hat{S}(\xi))>0$. Therefore, the condition on cost parameters along with the first-order stochastic dominance of $G(S|\xi)$ implies that $\frac{\partial^2}{\partial S\partial\xi}H(\hat{S}(\xi),\xi)$ is positive. This along with (C.4) implies that $\hat{S}(\xi)$ is increasing. It follows that $S^*(\xi)=\max\{x,S^o(\pi),\hat{S}(\xi)\}$ is increasing as well. For any feasible menu of base stock levels $S(\xi)$ offered to the retailer, the corresponding payments can be determined from (11):

$$\begin{aligned}P(\xi)=(r-w)&\left(\int_0^{S(\xi)}\bar{G}(z|\xi)\,dz+\int_{\underline{\xi}}^{\xi}\int_0^{S(\eta)}\frac{\partial}{\partial\eta}G(z|\eta)dzd\eta\right)\\&-(1-\alpha)\overline{\Pi}^r_{\min}(x,S^o_\pi,\underline{\xi}).\end{aligned} \tag{C.5}$$

□

Proof of Theorem 4. If, in period n, the supplier decides to wait an additional period before offering the menu of contracts, she raises the inventory level in period n to y_n. Two scenarios are possible in period n.

(1) If the sales observations in period n are censored (i.e., $y_n=z_n$), we have

$$\begin{aligned}\lambda_{n+1}(\xi)&=\frac{\bar{G}(y_n|\xi)\cdot\pi_n(\xi)}{\int_{\underline{\xi}}^{\bar{\xi}}\bar{G}(y_n|\eta)\cdot\pi_n(\eta)\,d\eta}=\frac{\pi_n(\xi)}{\int_{\xi}^{\bar{\xi}}\frac{\bar{G}(y_n|\eta)}{\bar{G}(y_n|\xi)}\cdot\pi_n(\eta)\,d\eta}\\&\le\frac{\pi_n(\xi)}{\int_{\xi}^{\bar{\xi}}\pi_n(\eta)\,d\eta}=\lambda_n(\xi).\end{aligned}$$

The first equality follows from the definition of $\lambda_{n+1}(\cdot)$. The inequality follows from first-order stochastic dominance of demand distributions, $\bar{G}(y_n|\eta)\ge\bar{G}(y_n|\xi)$, for all $\eta\in[\xi,\bar{\xi}]$. This implies the following $\frac{\partial H_{n+1}(S,\xi)}{\partial S}\le\frac{\partial H_n(S,\xi)}{\partial S}$ for all $S\in\mathbb{R}^+$ and $\xi\in\Theta$. $H_{n+1}(\cdot,\xi)$ is unimodal by Theorem 2.[23] Thus, $\frac{\partial H_{n+1}(\hat{S}_n(\xi),\xi)}{\partial S}\le\frac{\partial H_n(\hat{S}_n(\xi),\xi)}{\partial S}=0$. This implies that $\hat{S}_{n+1}(\xi)\le\hat{S}_n(\xi)$.

(2) If the sales observation in period n is uncensored (i.e., $z_n<y_n$), we have

$$\lambda_{n+1}(\xi)=\frac{g(z_n|\xi)\cdot\pi_n(\xi)}{\int_{\underline{\xi}}^{\bar{\xi}}g(z_n|\eta)\cdot\pi_n(\eta)\,d\eta}=\frac{\pi_n(\xi)}{\int_{\xi}^{\bar{\xi}}\frac{g(z_n|\eta)}{g(z_n|\xi)}\cdot\pi_n(\eta)\,d\eta}.$$

Therefore, the ratio of densities $\frac{g(z_n|\eta)}{g(z_n|\xi)}$ determines the relation between $\lambda_{n+1}(\xi)$ and $\lambda_n(\xi)$. For the newsvendor family of demand distributions, $\frac{\partial g(z_n|\xi)}{\partial\xi}=\frac{t'(z_n)}{\xi^3e^{\frac{t(z_n)}{\xi}}}(t(z_n)-\xi)$. Because

$$t'(\cdot)>0,\ \frac{\partial g(z_n|\xi)}{\partial\xi}\ \begin{cases}<0, & \text{if } \xi>t(z_n);\\ \ge0, & \text{if } \xi\le t(z_n).\end{cases}$$

(i) For types $\xi\ge t(z_n)$, it follows that $g(z_n|\eta)<g(z_n|\xi)$ for all $\eta>\xi$. Hence,

$$\lambda_{n+1}(\xi)=\frac{\pi_n(\xi)}{\int_{\xi}^{\bar{\xi}}\frac{g(z_n|\eta)}{g(z_n|\xi)}\cdot\pi_n(\eta)\,d\eta}\ge\frac{\pi_n(\xi)}{\int_{\xi}^{\bar{\xi}}\pi_n(\eta)\,d\eta}=\lambda_n(\xi).$$

Therefore, $\frac{\partial H_{n+1}(S,\xi)}{\partial S} \geq \frac{\partial H_n(S,\xi)}{\partial S}$, and unimodality of $H_{n+1}(\cdot,\xi)$ and $H_n(\cdot,\xi)$ implies that $\hat{S}_{n+1}(\xi) \geq \hat{S}_n(\xi)$ for all types $\xi \geq t(z_n)$.

(ii) For types $\xi < t(z_n)$ such that $g(z_n|\xi) \leq g(z_n|\overline{\xi})$, note that $g(z_n|\eta) > g(z_n|\xi)$ for all $\eta > \xi$, because $g(z_n|\eta)$ is increasing up to $t(z_n)$ and decreasing thereafter. Therefore, $\lambda_{n+1}(\xi) = \frac{\pi_n(\xi)}{\int_\xi^{\overline{\xi}} \frac{g(z_n|\eta)}{g(z_n|\xi)} \pi_n(\eta)\,d\eta} \leq \frac{\pi_n(\xi)}{\int_\xi^{\overline{\xi}} \pi_n(\eta)\,d\eta} = \lambda_n(\xi)$. By arguments similar to those used in earlier part, we $\hat{S}_{n+1}(\xi) < \hat{S}_n(\xi)$. □

Proof of Theorem 5. In the symmetric setting, $\xi \in \Theta$ is common knowledge. Under a linear price contract, the retailer pays pS for base stock level S. Subsequently, he pays a fixed wholesale price of w per unit to satisfy demand to the extent possible. The retailer's profit if he chooses base stock level S is $\frac{1}{1-\alpha}\left((r-w-p)S - (r-w)\int_0^S G(z|\xi)\,dz\right)$. The retailer chooses the inventory level that maximizes his profit: $\frac{\partial}{\partial S}\frac{1}{1-\alpha}(r-w-p-(r-w)G(S|\xi)) = 0$. To coordinate the channel, the inventory level must be determined using the critical fractile, $S^{fb}(\xi) = G^{-1}(\frac{r-c}{r+h-\alpha c}|\xi)$. To ensure that the retailer chooses the coordinating inventory level $S^{fb}(\xi)$, the supplier sets $G(S|\xi) = G(S^{fb}(\xi)|\xi) = \frac{r-c}{r+h-\alpha c}$. The resulting coordinating marginal price is $p = p^{fb} := \frac{(r-w)(h+c(1-\alpha))}{r+h-\alpha c}$. The total profit in the coordinated supply chain can be calculated using the following recursion:

$$
\begin{aligned}
V^{cs}(x,\xi) &= (r-c) - (r+h-\alpha c)\int_0^{\hat{y}(\xi)} G(z|\xi)dz \\
&\quad + \alpha(1-G(\hat{y}_n(\xi)|\xi)V^{cs}(0,\xi) \\
&\quad + \alpha\int_0^{\hat{y}(\xi)} g(z|\xi)V^{cs}(\hat{y}(\xi)-z,\xi)dz, \qquad \text{(C.6)}
\end{aligned}
$$

where $\hat{y}(\xi) := \max\{S^{fb}(\xi), x\}$. □

Proof of Theorem 6.

(1) Differentiating payment function $P(\xi)$ wrt ξ in Equation (C.5) gives $\frac{dP(\xi)}{d\xi} = (r-w)(S'(\xi) - G(S(\xi)|\xi)S'(\xi) - \int_0^{S(\xi)} \frac{\partial G(z|\xi)}{\partial \xi}dz + \int_0^{S(\xi)} \frac{\partial}{\partial \xi}G(z|\xi)dz) = (r-w)S'(\xi)(1-G(S(\xi)|\xi)) \geq 0$. It follows from above that $\frac{dP}{dS} = \frac{\frac{dP(\xi)}{d\xi}}{\frac{dS(\xi)}{d\xi}} = (r-w)(1-G(S|\xi))$. Then, $\frac{d\hat{P}}{d\hat{S}} = (r-w)(1-G(\hat{S}(\xi)|\xi)) \geq (r-w)(1-G(S^{fb}(\xi)|\xi)) = (r-w)(1-\frac{r-c}{r+h-\alpha c})$. The inequality follows, because the supplier never maintains inventory level higher than the first best level $S^{fb}(\xi)$ for any demand type ξ.

(2) The marginal price that the retailer of type ξ pays is $\frac{d\hat{P}_n}{d\hat{S}_n} = (r-w)(1-G(\hat{S}_n(\xi)|\xi)) \leq (r-w)(1-G(\hat{S}_{n+1}(\xi)|\xi)) = \frac{d\hat{P}_{n+1}}{d\hat{S}_{n+1}}$, using Part 1 of Theorem 4.

(3) For all $\xi > t(z_n)$, $\hat{S}_{n+1}(\xi) \geq \hat{S}_n(\underline{\xi})$ from Part 2(i) of Theorem 4. Therefore, $\frac{d\hat{P}_n}{d\hat{S}_n} = (r-w)(1-G(\hat{S}_n(\xi)|\xi)) \geq (r-w)\cdot(1-G(\hat{S}_{n+1}(\xi)|\xi)) = \frac{d\hat{P}_{n+1}}{d\hat{S}_{n+1}}$. □

Proof of Theorem 7.

(1) Differentiating $\Pi^{sr}(x,\pi)$ with respect to x gives $\frac{d}{dx}\Pi^{sr}(x,\pi) = \frac{1}{1-\alpha}\int_\Theta \pi(\xi)\frac{d}{dx}H(S^*(\xi),\xi)d\xi$. For any

$$
\xi \in \Theta, \frac{d}{dx}H(\max\{x, S^o_\pi, \hat{S}(\xi)\}, \xi|\pi)
\begin{cases} = 0, & \text{if } x \leq \max\{S^o_\pi, \hat{S}(\xi)\}; \\ \leq 0, & \text{if } x > \max\{S^o_\pi, \hat{S}(\xi)\}, \end{cases}
$$

because $H(\cdot,\xi|\pi)$ is a unimodal function and $\hat{S}(\xi)$ is the unique maximizer of $H(\cdot,\xi|\pi)$ and characterized by the first-order condition.

(2) To establish the result, we first construct a lower bound for $L(y,\pi)$. Then, the lower bound is used to find an upper bound on the difference $\Pi^{sr}(x,\pi) - \max_{y\geq x} L(y,\pi)$. Last, we find the range of values of x for which the difference is always less than zero, thus establishing the result. Note that $L(y,\pi) = (w-c)y - (w+h-\alpha c)\int_0^y Q(z)dz + \alpha\overline{Q}(y)V(0,\pi^c(\cdot|y)) + \alpha\int_0^y q(z)\cdot V(y-z,\pi^e)dz$, where $Q(z) = \int_\Theta \pi(\xi)G(z|\xi)d\xi$ is the predictive demand distribution. To find a lower bound on $V(z,\pi)$, consider a policy in which the supplier never produces and never offers contracts to the retailer (i.e., $\tau = +\infty$). Then, $V(z,\pi) \geq \mathbb{E}[w\min\{z,D_1\} - h(z-D_1)^+ + \alpha(w\min\{x_2,D_2\} - h(x_2 - D_2)^+) + \ldots] \geq -hz$. Hence, $L(y,\pi) \geq (w-c)y - (w+h-\alpha c)\cdot\int_0^y Q(z)dz - h\alpha\int_0^y q(z)(y-z)dz = (w-c)y - (w+h(1+\alpha)-\alpha c)\cdot\int_0^y Q(z)dz =: \overline{L}(y,\pi)$. Note that $\underline{L}(y,\pi)$ is concave in y and that it is maximized at $\underline{y}(\pi) := Q^{-1}\left(\frac{w-c}{w+h(1+\alpha)-\alpha c}\right) < \hat{S}(\overline{\xi})$. For $x \geq \max\{S^o(\pi), \hat{S}(\overline{\xi})\} \geq \max\{S^o(\pi), \hat{S}(\xi)\}$, it follows that $\Pi^{sr}(x,\pi) - \max_{y\geq x} L(y,\pi) \leq -\overline{\Pi}^r_{\min}(x, S^o_\pi, \underline{\xi}) + \frac{1}{1-\alpha}\int_\Theta \pi(\xi)H(x,\xi|\pi)d\xi - \max_{y\geq x}\underline{L}(y,\pi) \leq \frac{1}{1-\alpha}\int_\Theta \pi(\xi)H(x,\xi|\pi)d\xi - \underline{L}(x,\pi) \leq \frac{1}{1-\alpha}\int_\Theta \pi(\xi)\cdot\left\{(r-c)x - (r+h-\alpha c)\int_0^x G(z|\xi)dz + \frac{(r-w)}{\lambda(\xi)}\int_0^x \frac{\partial}{\partial\xi}G(z|\xi)\,dz\right\}d\xi - \underline{L}(x,\pi) \leq \frac{1}{1-\alpha}\int_\Theta \pi(\xi)\left\{(r-c)x - (r+h-\alpha c)\int_0^x G(z|\xi)dz\right\}d\xi - (w-c)x - (w+h(1+\alpha)-\alpha c)\int_0^x Q(z)dz \leq \frac{1}{1-\alpha}\int_\Theta \pi(\xi)\{(r-c-(1-\alpha)\cdot(w-c))x - (r+h-\alpha c-(1-\alpha)(w+h(1+\alpha)-\alpha c))\int_0^x G(z|\xi)dz\}\,d\xi \leq \frac{(r-c-(1-\alpha)(w-c))}{1-\alpha}\int_\Theta \pi(\xi)\left\{x - \theta\int_0^x G(z|\xi)dz\right\}d\xi \leq \frac{(r-c-(1-\alpha)(w-c))}{1-\alpha}\left\{x - \theta\int_0^x G(z|\overline{\xi})dz\right\}$, where $\theta := \frac{r+h-\alpha c-(1-\alpha)(w+h(1+\alpha)-\alpha c)}{r-c-(1-\alpha)(w-c)} = \frac{r-w(1-\alpha)+\alpha^2(h-c)}{r-w(1-\alpha)-\alpha c} > 1$. The function inside the integral on the last line is concave in x, and it is negative when $x > U$ for all ξ. U is implicitly defined as the nonzero solution to $U = \theta\int_0^U G(z|\overline{\xi})d\xi$. Thus, the last integral is nonpositive for all $x > K$, where $K := \max\{U, S^o(\pi), \hat{S}(\overline{\xi})\}$. □

Endnotes

[1] For example, Walmart only owns products briefly as they pass through the checkout scanner (see p. 156 in Simchi-Levi et al. 2008).

[2] In the classical case, Hammond (2006) describes the stern opposition to VMI practice from Barilla's distributors. The article points out the difficulty that Barilla had in incorporating promotional data, which are separate from the usual EDI information, into their forecasting process. Giorgio Maggiali, then Director of Logistics at Barilla, noted, "We're grappling with how to treat these promotions in our

operations planning processes, including forecasting, manufacturing, and logistics." Ineffectively managing inventory lead to disappointment of some distributors (e.g., the "disgruntled DO" mentioned on p. 3 of Hammond 2006) over VMI implementation and eventually, falling out of the relationship.

[3] Henceforth, we consider a VMI setting in which inventory is managed in a consignment fashion.

[4] We restrict attention to inventory policies such that y_n can be determined based on the information available to the supplier at the start of period n. That is, for any $u \in \mathbb{R}^+$, the event $\{y_n > u\} \in \mathfrak{F}^{Z,y}_{n-1}$, where $\{\mathfrak{F}^{Z,y}_i\}_{i\geq 1}$ denotes the natural filtration associated with the sales observation process.

[5] We remark that our main results continue to hold with the finite horizon formulation. Therefore, these results are relevant in managing short- or long-term supply chain relationships. Hence, the proposed approach and results can be used for supply chains distributing innovative products (e.g., consumer electronics) or functional products (e.g., grocery products). Using infinite horizon formulation simplifies notation and mitigates the end-of-horizon effects.

[6] Offering a contract term with a base stock level S lower than the retailer's current on-hand inventory level x_n is often impractical (although technically possible) in a VMI setting. Note that reducing the base stock level indicates reduction of inventory; hence, the retailer faces low inventory. Recall also that more inventory always benefits the retailer, whereas it only costs the supplier. As a result, reducing base stock level to a point lower than the current on-hand inventory can prove controversial. This constraint also provides analytical tractability, but as we shall note later, it may not be binding in the optimal solution.

[7] The *revelation principle* (Myerson 1981) ensures that, without loss of generality, the supplier can restrict herself to those menus of contracts in which the retailer truthfully communicates his private assessment of demand in Bayesian equilibrium. Thus, we narrow our search for the optimal menu to those menus of contracts for which $\tilde{\xi} = \xi$ is the best response function for the retailer of type ξ.

[8] To be precise, $\mathcal{M} := \{(\mathbf{y}, \tau):\ y_i \in \mathbb{R}^+, y_i \geq x_i$ with $\{y_i > u\} \in \mathfrak{F}^{Z,y}_{i-1}$, $\forall u \in \mathbb{R}, 1 \leq i \leq \tau - 1$, and $\{\tau = n\} \in \mathfrak{F}^{Z,y}_{n-1}$, for all $n \geq 1\}$.

[9] With a slight abuse of notation, we let $\Pi^r_{\min}(x, S, \xi)$ denote the case in which the supplier uses a base stock policy with a fixed base stock level S in each period after the menu of contracts is rejected. In this case, the vector $\mathbf{y}^o = (y^o_n, y^o_{n+1}, \ldots)$ defined in (PC) simplifies as $y^o_n = \max\{x_n, S\}$.

[10] All of our analytical results also hold when we replace the retailer's reservation profit with the lower bound.

[11] We use increasing/decreasing in the weak sense throughout the paper.

[12] Setting $s(\xi) = l(\xi) = \xi^{-1}$ and $h(z) = k'(z)$, we obtain the newsvendor family.

[13] The supplier's action $(y_n - x_n)$ influences the information signal (i.e., periodic sales z_n) that she observes, and together, they govern the evolution of her belief process over time.

[14] The condition is met if the total profit margin in the supply chain, $\frac{r-c}{r}$, is less than 87%.

[15] See http://www.vendormanagedinventory.com/pitfalls.php.

[16] For example, for a two-point prior distribution, censored demand observation (suggesting larger market demand) increases the base stock level $S^o(\pi_n)$ but lowers $\hat{S}(\cdot|\pi_n)$ in the following period (as shown in Part 1 of Theorem 4). The evolution of $S^o(\pi_n)$ is purely driven by the learning dynamic, whereas $\hat{S}(\cdot|\pi_n)$ also accounts for evolution of the information rent.

[17] Let G, H be two continuously differentiable distribution functions, and let μ_1, μ_2 denote their failure rate functions, respectively. Then, $G \leq_{fr} H$ if and only if $\mu_1(x) \geq \mu_2(x)$ for all $x \geq 0$.

[18] Recall that $t(\cdot)$ is defined in Theorem 2. As an example, $t(z) = z$ if $G(x|\xi)$ has an exponential distribution.

[19] The Bayesian update in (1) does not alter the belief after this point.

[20] We use $\hat{S}_n, \hat{P}_n$ to denote $\hat{S}(\xi|\pi_n), \hat{P}(\xi|\pi_n)$, respectively.

[21] Note that the wholesale price w is generally negotiated by the two firms in the VMI agreement and hence, may better indicate relative market power of each firm.

[22] We set the prior $p = 0.2$ to capture the high degree of information asymmetry (H). This value of p results in large ρ for both $\overline{\xi} = 6$ and 12. We set the prior $p = 0.8$ and 0.9 to capture the low degree of information asymmetry for $\overline{\xi} = 6$ and 12, respectively.

[23] We use $H_n(\cdot, \xi)$ to denote $H(\cdot, \xi|\pi_n)$.

References

Aviv Y (2001) The effect of collaborative forecasting on supply chain performance. *Management Sci.* 47(10):1326–1343.

Aviv Y (2002) Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing Service Oper. Management* 4(1):55–74.

Aviv Y (2007) On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Sci.* 53(5):777–794.

Azoury KS (1985) Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* 31(9):1150–1160.

Babich V, Li H, Ritchken P, Wang Y (2012) Contracting with asymmetric demand information in supply chains. *Eur. J. Oper. Res.* 217(2):333–341.

Berger RL, Casella G (2002) *Statistical Inference*, 2nd ed. (Duxbury Press, Pacific Grove, CA).

Bisi A, Dada M, Tokdar S (2011) A censored-data multiperiod inventory problem with newsvendor demand distributions. *Manufacturing Service Oper. Management* 13(4):525–533.

Braden DJ, Freimer M (1991) Informational dynamics of censored observations. *Management Sci.* 37(11):1390–1404.

Brinkhoff A, Özer Ö, Sargut G (2015) All you need is trust? An examination of inter-organizational supply chain projects. *Production Oper. Management* 24(2):181–200.

Burnetas A, Gilbert SM, Smith CE (2007) Quantity discounts in single-period supply contracts with asymmetric demand information. *IIE Trans.* 39(5):465–479.

Cachon GP, Fisher M (2000) Supply chain inventory management and the value of shared information. *Management Sci.* 46(8): 1032–1048.

Cachon GP, Lariviere MA (2001) Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* 47(5):629–646.

Chen L (2010) Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales. *Oper. Res.* 58(2):396–413.

Chen L, Lee HL (2009) Information sharing and order variability control under a generalized demand model. *Management Sci.* 55(5):781–797.

Chen L, Plambeck EL (2008) Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing Service Oper. Management* 10(2):236–256.

Dong Y, Dresner M, Yao Y (2014) Beyond information sharing: An empirical analysis of vendor-managed inventory. *Production Oper. Management* 23(5):817–828.

Feng Q, Lai G, Lu LX (2015) Dynamic bargaining in a supply chain with asymmetric demand information. *Management Sci.* 61(2): 301–315.

Ferguson D (2013) How supermarkets get your data—and what they do with it. *The Guardian* (June 8), http://www.theguardian.com/money/2013/jun/08/supermarkets-get-your-data.

Fisher ML (1997) What is the right supply chain for your product? *Harvard Bus. Rev.* 75(March/April):105–116.

Gallego G, Özer Ö (2001) Optimal use of demand information in supply chain management. Song J, Yano D, eds. *Supply Chain Structures* (Kluwer Academic Publishers, Dordrecht, Netherlands), 119–160.

Gümüş M (2014) With or without forecast sharing: Competition and credibility under information asymmetry. *Production Oper. Management* 23(10):1732–1747.

Ha AY, Tong S, Zhang H (2011) Sharing demand information in competing supply chains with production diseconomies. *Management Sci.* 57(3):566–581.

Hammond JH (2006) Barilla SpA (D). Harvard Business School Case 9-695-066, Harvard University, Cambridge, MA.

Karlin S, Scarf H (1958) Inventory models of the Arrow-Harris-Marschak type with time lag. Arrow K, Karlin S, Scarf H, eds. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press, Stanford, CA), 155–178.

Katariya AP, Çetinkaya S, Tekin E (2014) Cyclic consumption and replenishment decisions for vendor-managed inventory of multisourced parts in Dell's supply chain. *Interfaces* 44(3):300–316.

Kouvelis P, Chambers C, Wang H (2006) Supply chain management research and production and operations management: Review, trends, and opportunities. *Production Oper. Management* 15(3): 449–469.

Lariviere MA, Porteus EL (1999) Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* 45(3):346–363.

Lee HL, Padmanabhan V, Whang S (1997) Information distortion in a supply chain: The Bullwhip effect. *Management Sci.* 43(4): 546–558.

Lee HL, So KC, Tang CS (2000) The value of information sharing in a two-level supply chain. *Management Sci.* 46(5):626–643.

Li L, Zhang H (2008) Confidentiality and information sharing in supply chain coordination. *Management Sci.* 54(8):1467–1481.

Li T, Tong S, Zhang H (2014) Transparency of information acquisition in a supply chain. *Manufacturing Service Oper. Management* 16(3):412–424.

Lobel I, Xiao W (2017) Technical note—Optimal long-term supply contracts with asymmetric demand information. *Oper. Res.* 65(5): 1275–1284.

Lovejoy WS (1990) Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* 36(6):724–738.

Lutze H, Özer Ö (2008) Promised lead-time contracts under asymmetric information. *Oper. Res.* 56(4):898–915.

Mathews R (1995) *Spartan Pulls the Plug on VMI* (Progressive Grocer, Deerfield, IL).

Mirrlees JA (1971) An exploration in the theory of optimum income taxation. *Rev. Econom. Stud.* 38(2):175–208.

Myerson RB (1981) Optimal auction design. *Math. Oper. Res.* 6(1): 58–73.

Nazerzadeh H, Perakis G (2016) Technical note—Nonlinear pricing competition with private capacity information. *Oper. Res.* 64(2):329–340.

Oh S, Özer Ö (2013) Mechanism design for capacity planning under dynamic evolutions of asymmetric demand forecasts. *Management Sci.* 59(4):987–1007.

Oracle (2009) Peoplesoft enterprise inventory 9.1 peoplebook. Technical report, Oracle Corporation, Pleasanton, CA.

Özer Ö, Wei W (2006) Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Sci.* 52(8):1238–1257.

Ren ZJ, Cohen MA, Ho TH, Terwiesch C (2010) Information sharing in a long-term supply chain relationship: The role of customer review strategy. *Oper. Res.* 58(1):81–93.

Rosenblum P (2014) Walmart's out of stock problem: Only half the story? *Forbes* (April 15), http://www.forbes.com/sites/paularosenblum/2014/04/15/walmarts-out-of-stock-problem-only-half-the-story/.

Scarf H (1959) Bayes solutions of the statistical inventory problem. *Ann. Math. Statist.* 30(2):490–508.

Scarf HE (1960) Some remarks on Bayes solutions to the inventory problem. *Naval Res. Logist. Quart.* 7(4):591–596.

Shamir N, Shin H (2016) Public forecast information sharing in a market with competing supply chains. *Management Sci.* 62(10): 2994–3022.

Shang W, Ha AY, Tong S (2016) Information sharing in a supply chain with a common retailer. *Management Sci.* 62(1):245–263.

Simchi-Levi D, Kaminsky P, Simchi-Levi E (2008) *Designing & Managing the Supply Chain*, 3rd ed. (McGraw-Hill/Irwin, New York).

Taylor R, Mauer G (2013) Margin unlocked: Integrated margin management to deliver breakthrough performance inconsumer products. Report, Ernst & Young, London, UK.

The Grocer (2009) Northern in talks over Sainsbury's VMI deal. Accessed March 10, 2019, http://www.thegrocer.co.uk/companies/northern-in-talks-over-sainsburys-vmi-deal/199333.article.

Tirole J (2002) *The Theory of Industrial Organization* (MIT Press, Cambridge, MA).

Watson E (2005) Nestlé switches to vendor managed inventory with Tesco. Accessed March 10, 2019, http://www.foodmanufacture.co.uk/Supply-Chain/Nestle-switches-to-vendor-managed-inventory-with-Tesco.

Zhang H, Zenios S (2008) A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. *Oper. Res.* 56(3):681–696.

Zhang H, Nagarajan M, Sošić G (2010) Dynamic supplier contracts under asymmetric inventory information. *Oper. Res.* 58(5):1380–1397.