

INSEAD

The Business School
for the World®

Faculty & Research Working Paper

Synchronizing Global Supply Chains:
Advance Purchase Discounts

Wenjie TANG
Karan GIROTRA
2010/07/DS/TOM

Synchronizing Global Supply Chains: Advance Purchase Discounts

by
Wenjie Tang*

and
Karan Girotra**

February 2010

* PhD Candidate in Decision Sciences at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France, wenjie.tang@insead.edu

** Assistant Professor of Technology and Operations Management at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France, karan.girotra@insead.edu

A working paper in the INSEAD Working Paper Series is intended as a means whereby a faculty researcher's thoughts and findings may be communicated to interested readers. The paper should be considered preliminary in nature and may require revision.

Printed at INSEAD, Fontainebleau, France. Kindly do not reproduce or circulate without permission.

SYNCHRONIZING GLOBAL SUPPLY CHAINS: ADVANCE PURCHASE DISCOUNTS

WENJIE TANG AND KARAN GIROTRA

ABSTRACT. We study the economics of sharing demand information between a dual sourcing firm and its network of retailers. Our analysis demonstrates that employing Advance Purchase Discount (APD) contracts in these supply chains removes significant impediments to information sharing. Essentially, these contracts synchronize the timeline of the dual sourcing firm’s decisions with the actions of its network of retailers. This enables *accurate, timely, and self-enforcing* information sharing, which reduces the demand–supply mismatches, and improves the profitability of each of the agents in the supply chain. We provide prescriptions on the appropriate design of contracts that enable this Pareto-improving information sharing. Next, we extend this analysis to incorporate realistic constraints on the dual sourcing firm’s limited knowledge of its retailers’ operational costs and information quality. We characterize “certainty-equivalent” values of the unknown retailer parameters, which facilitate analogous prescriptions for the design of APD contracts in these realistic settings. It is interesting that the unobservability of retailer parameters leads to an asymmetric and “degree-of-unobservability dependent” departure from the full-observability design of the APD contract. If the uncertainty in the unobserved parameter is small, then the optimal discount is *higher* compared to the case of full observability; conversely, when the uncertainty is large, the optimal discount is *lower*. It is significant that, despite the departure in the design of the APD contract, this analysis reiterates that the benefits of APDs persist even under this practical constraint. Finally, we report on the application of our prescription to a U.S.-based apparel wholesaler. Using data derived from the operations of this firm, we estimate that APD-enabled improved information sharing can increase net profits by about 17%.

February, 2010

1. INTRODUCTION

Over the last two decades, improved communication technologies, differences in input costs between geographies, and geographical skill specialization have led firms to employ increasing levels of global sourcing: sourcing products from locations far away from the market for the product (Hausman et al. 2005). The long distances between the sourcing locations and the market lead to long delivery lead times and necessitate that firms make sourcing quantity decisions well before actual demand for

the product is realized. For firms sourcing short-life cycle products with uncertain demand, such as fashion wear, this implies that the sourcing quantity decisions must be made with limited demand-relevant information. This leads to significant demand–supply mismatches and is a major drain on firm profitability in these industries. Strategies to gather early, accurate demand information can reduce these demand supply mismatches and have the potential to be a key competitive advantage in this era of global sourcing (Donohue 2000, Taylor and Xiao 2009).

In addition to a firm’s best information about demand, its downstream partners in the supply chain, who are closer to the customers, often have some additional demand-relevant information. Gathering this information in a timely fashion and using it to place more accurate orders at long–lead time suppliers presents itself as an opportunity to reduce demand–supply mismatches. However, traditional methods of acquiring this information—such as customer surveys, retailer shows, industry conventions, buyer events, and so forth—lead to poor accuracy and inappropriate timing of the information. Specifically, the information shared by the partners is liable to be manipulated by the communicating party to favor itself. Further more, the information cannot always be gathered at a time of the firm’s choosing. It may be gathered too soon, in which case it will not reflect the latest trends from the marketplace; or it may be gathered too late, in which case it cannot be used to place orders at long–lead time suppliers.

This study proposes a contracting mechanism: an Advance Purchase Discount (APD) contract that enhances the effectiveness of a firm’s global sourcing strategy by incentivizing *accurate information communication* from downstream tiers *at a time chosen* by the firm. Under such a scheme, in addition to placing orders during the sales season, the firm offers its downstream retailers an opportunity to place an additional order, in advance, at a discounted price. These early orders are timed to coincide with the latest order time that ensures delivery from the long–lead time global source. This advances some of the retailer’s demand and transfers some of the mismatch risk to the retailers (Cachon 2004). More interesting and hitherto unidentified is the impact of this contract on the information-sharing issues highlighted previously. Acting in their own best interests, the downstream retailers use the best information available to them in deciding the early order quantities. From observing these orders, the upstream firm can infer the information available to the retailers and can make its global sourcing decisions using the demand information inferred. Essentially, the early order opportunity synchronizes the timeline of orders placed by retailers with the timeline of decisions that the global sourcing firm must make. This alignment drives the timely communication

of accurate demand information by the downstream retailers to the firm. In the process, it also redistributes risks in the supply chain in an efficient fashion. We show that, when taken together, the timely communication of information and the redistribution of risk improve the supply chain's performance as a whole and can also improve the profits of both the retailers and the firm. We demonstrate the robustness of this result by considering two different product designs: one where products are customized for each market and one where the same generic product is sold in each market.

Next, we examine the application of these contracts in a realistic setting. The contractual remedy just described for inducing accurate and timely information sharing relies on designing incentives that make information sharing preferable to any potential gains that the downstream retailers may accrue by either manipulating or not sharing the information. As with most other contractual remedies, designing these incentives requires knowledge about the retailers' benefits from manipulating or not sharing the information. These benefits depend on the operational costs and quality of the information available to the retailers. It is theoretically appealing to assume public knowledge of these retailer side parameters, but this is unrealistic in most practical situations, and such assumptions often limit the actionable prescriptions from the contractual remedies proposed. To address such concerns, we extend our proposal for employing APD contracts to this more practical setting where the firm may have limited knowledge about its retailers' operations. In this setup we characterize the optimal strategy in terms of a "certainty-equivalent" value of the unobserved parameter, a value that the firm can use as if it knew the unobserved parameter for sure. It is interesting that this value leads to an asymmetric and "degree-of-unobservability dependent" departure from the full-observability design of the Advance Purchase Discount. When the uncertainty in the unobserved parameter is small, the optimal discount is *higher* than in the case of full observability; when the uncertainty is large, the firm should set a *lower* discount. It is significant that this analysis reiterates that the benefits of APDs—improving agents' profits and the supply chain's profits as a whole—persist even under this practical constraint.

To validate the applicability of our analytical setup, to estimate the benefits accruing from our prescriptions, and to develop practical guidelines for application, we close this study by reporting on the implementation of our APD scheme at a New Jersey-based fashion-wear wholesaler. The subject of our study, Costume Gallery Inc. (www.costumegallery.net) is one of the largest wholesalers of dance costumes in the United States. The demand for dance costumes is highly seasonal, and almost

all of the annual demand is concentrated near the end of the school year in April. This demand is highly influenced by cultural trends (this season, Michael Jackson costumes are exceptionally popular). Costume Gallery sources extensively from Asia and has some local production, but it sells the dance costumes through a network of U.S.-based dance schools (the retailers). These dance schools typically buy costumes for dance productions they are organizing for their students. Consequently, they have access to significant advance knowledge about the popular styles for the season and the overall popularity of dancing as an activity. The supply chain structure, the sourcing economics, and the informational environment in this setup are completely in line with our analytical model. Further, the proposed APD contract was found to be simple and faced little resistance in implementation. In particular, existing arrangements for managing accounts payables at most of the affected firms were already employing similar terms for the purposes of cash flow management. Our proposal was thus perceived to not depart significantly from the existing business culture. Using data derived from Costume Gallery's operations, we estimate that implementing our proposed scheme would enable this firm to increase its net profits by about 17%. We close by reporting on some practical issues encountered during the implementation.

The theoretical analysis demonstrating the effectiveness of our proposed contracting mechanism, the case-based validation of its assumptions, and the enumeration of the accrued benefits reported in this study make three important contributions to the theory and the practice of information sharing in supply chains.

- (1) We demonstrate that advance purchase discounts can be used in a supply chain to synchronize the timeline of different tiers of the supply chain; this induces *accurate, timely, and self-enforcing information sharing*, which can make each of the agents in a supply chain better off. We identify key drivers of these benefits and, based on this analysis, we develop prescriptions for how a firm can go about designing and utilizing such a scheme to improve its profitability.
- (2) We extend our analysis to a realistic setting in which the firm has only limited knowledge about downstream retailers' operations. We illustrate analytically that all the benefits highlighted previously persist in this more realistic setup, and we develop novel insights for the design of APD contracts in these settings. We then report on the implementation of our prescriptions at a real organization. This provides case-based evidence of the applicability of our analysis and an estimate of the potential benefits derived from our prescription. Taken

together, the model extension and the study on implementation demonstrate the validity of our analysis, improve the applicability of our prescriptions, and reinforce the utility of advance purchase discounts in increasing firm profitability.

- (3) Our illustration of advance purchase discounts as a Pareto-improving mechanism for information sharing brings together and enhances the largely parallel literatures on demand information sharing and advance purchase discounts. Like this work, the existing demand information sharing literature has proposed contracts to achieve improved information sharing (cf. Cachon and Lariviere 2001, Ozer and Wei 2006). We enhance this literature by considering an unexplored but widely prevalent dual sourcing arborescent supply chain structure, by incorporating novel practical constraints on supply chain visibility and contract acceptance, and by proposing a simple, optimized contract. The Advance Purchase Discounts literature has studied these contracts as instruments for aligning incentives and sharing risks (Cachon 2004, Donohue 2000), and for incentivizing early consumer purchases (McCardle et al. 2004, Tang et al. 2004). We enhance this literature by identifying a hitherto unrecognized benefit of these discounts: credible, self-enforcing information sharing between different tiers of a supply chain.

The rest of the paper is organized as follows. Section 2 reviews the parallel literatures on demand information sharing and advance purchase discounts. Section 3 provides analytical results and insights demonstrating the utility of advance purchase discounts. Section 4 illustrates the persistence of these benefits even when there is limited knowledge about the downstream supply chain. Finally, Section 5 reports on applying to a real organization the results from Section 3 and 4. We conclude in Section 6.

2. RELATED LITERATURE

This study brings together and enhances the largely parallel literatures on information sharing and advance purchase discounts. We review our contributions to each of the two literatures in the following two subsections.

2.1. Literature on Information Sharing. The information sharing literature typically studies a supply chain in which a downstream agent has access to superior demand information that could be used by an upstream agent to make better decisions regarding production, inventory, or pricing.

A key concern in this literature is the value of information sharing. Specifically, major findings are that the value of information sharing is affected by coordination and information structures (Anand and Mendelson 1997, Aviv 2001, 2002), production flexibility (Fisher et al. 1994, Fisher and Raman 1996, Cachon and Fisher 2000, Fisher et al. 2001, Milner and Kouvelis 2002), demand characteristics (Lee et al. 2000, Milner and Kouvelis 2005), and competition and product characteristics (Ha et al. 2008, Ha and Tong 2008, Li 2002). Our analytical model follows this literature and models a similar asymmetric informational environment. Lack of information sharing is also a prime driver of the bullwhip effect (cf. Chen and Lee [2009] for a comprehensive treatment of this effect).

More recent work has identified impediments to realizing this potential value from information sharing. While information sharing may create value for one agent in the supply chain, it might not be beneficial to every other agent in the supply chain (Taylor 2006, Miyaoka and Hausman 2008). Further, information leakage might occur among competing agents (Anand and Goyal 2009). Finally, the downstream agent may have the incentive to manipulate the information he obtains for his own benefits, therefore impairing the truthful information sharing premise implied by the value of information sharing literature.

Later developments in the information sharing literature provide remedies to these impediments. A popular remedy is to design appropriate incentive and coordinating contracts so that, on one hand, it's in the best interest of the downstream agents to truthfully share their private information and on the other hand, the truthful information makes the supply chain better off. Our study continues in this tradition but achieves this remedy in our dual sourcing context by employing a simple, intuitively appealing, and novel contract—the Advance Purchase Discount Contract. Previous work has identified different contract types that achieve the same goal in other contexts; these include return and rebate policies (Arya and Mittendorf 2004, Taylor and Xiao 2009), contracts with commitment and options (Cachon and Lariviere 2001, Ozer and Wei 2006), and wholesale price contract under confidentiality (Li and Zhang 2008) and with customer review strategies (Ren et al. 2008).

Our remedy to enable information sharing distinguishes itself from the previously proposed ones in terms of the applicable context, its robustness to realistic constraints, in the simplicity and intuitiveness of the proposed contract, and in terms of *optimally* incentivizing self-enforcing information sharing.

Our applicable context is unique in terms of the supply chain structure and the nature of the supply chain relationships. Our contract is designed to enable information sharing between a dual

sourcing firm and its network of retailers. We explicitly model and consider the unique economics of this widely prevalent supply chain structure. Specifically, this implies that unlike the stylized economic models in Arya and Mittendorf [2004] and Li and Zhang [2008] which abstract from the internal physical flows in the supply chain, our operational analysis models individual agents in the supply chain such that they are allowed to carry physical inventories and face the risk of over- and under-supply. Further, unlike the single-source capacity reservation model of Ozer and Wei [2006], our study applies to products where shortage of production capacity is not a significant concern. Finally, in contrast to Cachon and Lariviere [2001], in our context, it is the upstream agent who needs the information offers a contract to infer private information from the downstream agent.

The supply chain relationships considered in our model depart from those studied in Ren et al. [2008], where the relationships are characterized by long-term commitments to enable inter-temporal trade-offs which significantly alter the enforcement incentives. Our proposed remedy applies equally even in the absence of long-term commitments.

In terms of the specifics of our contract design, unlike Ozer and Wei [2006] who also consider a time sensitive discount contract in a restricted capacity context, we endogenize a previously exogenous contract parameter—the discount in the APD contract, to *optimally* incentivize *self-enforcing* information sharing and create just the right incentives for each participating agent.

Our work also departs from the existing information sharing literature in incorporating two new practical constraints: First, our extended model explicitly considers a major practical constraint that has limited the implementation of most contracts—the contract designing agent’s limited knowledge about the cost structure and quality of information available to the opposing agent. All existing literature assumes full knowledge of these parameters, which in an information asymmetry context, is conflicting with the philosophical purpose of the contract itself. Our extended model considers this limited visibility and provides actionable prescriptions and appropriate contract designs for these contexts too. To the best of our knowledge, we are the first to consider such lack of visibility in the information sharing literature. Secondly, the contract type proposed in our model is simple and intuitive—getting a discount for timely actions and is very often used by businesses in the context of account payables; therefore it is easier to introduce it in within the constraints of the existing business culture.

Finally, we also report on the practical application of our contract in a real organization. To the best of our knowledge, this is the first paper that reports on a practical application of an information

sharing contract and provides some data-based estimates of the benefits of employing an information sharing contract.

To summarize, our work departs from the extensive information sharing literature by considering a novel but widely prevalent dual sourcing arborescent supply chain structure and a broad spectrum of supply chain relationships, by incorporating practical constraints on supply chain visibility and contract acceptance, and by proposing an optimized simple contract.

2.2. Advance Purchase Discount Literature. Our work is also related to the Advance Purchase Discounts literature. Advance purchase discounts were first studied as direct offers from a retailer to consumers in the market: Tang et al. [2004] consider an “advance booking discount” program that entices consumers to commit to their orders at a discounted price prior to the selling season, which shifts some of the demand earlier and enables retailers to better manage inventory and capacity. McCardle et al. [2004] further develop these consumer advance purchase discount contracts to a setup that includes retailer competition. Xie and Shugan [2001], Dana [1998], Gundepudi et al. [2001], and Akan et al. [2009] also consider consumer advance purchase discount contracts and establish their utility as a price discrimination mechanism. Li and Zhang [2009] study the interaction of the above two benefits of an APD contract: advance demand information and price discrimination. More relevant to our study is the literature on APDs offered in a supply chain. While the APD contract offered to agents in the supply chain is similar to consumer APD contracts, agents in a supply chain are typically assumed to have different utility and incentive structures from those of consumers, which leads to significantly different analysis and implications. In particular, new issues of risk distribution and incentive coordination that are not relevant in consumer APDs play an important role in these settings. The pioneering works by Donohue [2000], Cachon [2004], and Ozer et al. [2007] were the first to identify these effects. However, neither Donohue [2000] nor Cachon [2004] considers the information asymmetry between different tiers of the supply chain. Ozer et al. [2007] does consider information asymmetry, but the private information possessed by the downstream agent is not shared within the supply chain. More recently, Bernstein et al. [2009] apply these contracts to a supply chain with multiple locations and transshipment, while Cho and Tang [2009] study market price risks along with the usual demand–supply risks.

Finally, our work also extends traditional information-sharing and contracting models to capture an additional type of practically relevant asymmetric information—namely, asymmetric information

that arises from limited knowledge about the costs and information precision of the supply chain partners. This asymmetry has been previously studied in isolation with other information-sharing issues (cf. Kayis et al. [2007]) and our work integrates these two approaches. There is a growing body of literature that studies supply chain management with limited information about the environment (cf. Perakis and Roels [2008]). Our extension with asymmetric information continues in this tradition.

3. BENEFITS OF ADVANCE PURCHASE DISCOUNTS: CUSTOMIZED PRODUCTS

In this section, we develop a stylized model in order to illustrate the benefits of employing Advance Purchase Discount (APD) contracts, to understand the key drivers of these benefits, and to identify some challenges in employing these contracts.

Our model builds on the basic setup described in the celebrated case on Sport Obermeyer (Fisher et al. 1994, Hammond and Raman 1994, Fisher and Raman 1996, Parlar and Weng 1997, Donohue 2000), where the procurement of a short-life cycle product must be allocated between a “speculative” production source with long lead times and a “reactive” production source with short lead times. The reactive production source can make use of the latest demand information to meet any shortfalls from the speculative source. Our model replicates the speculative-reactive sourcing choice of a firm like Sport Obermeyer, yet we add three interesting features to this much-used model as follows. (i) In addition to modeling the *upstream* sourcing choices of the firm, we also explicitly model the actions and strategic behavior of the firm’s *downstream* supply chain—that is, we model the network of retailers through which the dual sourcing firm sells its products. (ii) We explicitly model the information environment in which the firm and its retailers make these choices, including allowances for any private information that may be available to the retailers and the consequent information asymmetry. (iii) We analyze and compare alternate contracts between a firm similar to Sport Obermeyer and its downstream retail network.

3.1. Supply Chain Setup and Information Environment. Consider a firm that sells a short-life cycle product through a network of N geographically dispersed retailers (Figure 1). Products with seasonal demand (e.g., fashion wear, toys, seasonal sports goods), rapidly evolving consumer electronics, and holiday decorations all fit this setup. In addition, perishable products that fit the classic newsvendor setup also fit. The retailers source the product from the firm with near

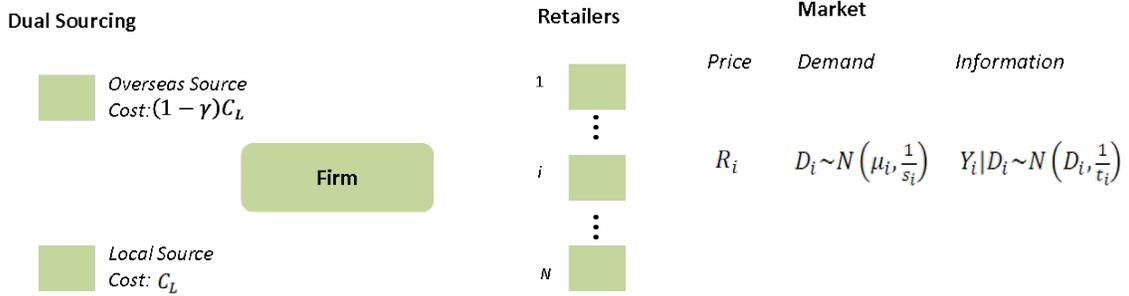


FIGURE 1. Setup: The Supply Chain

instantaneous lead times and sell it in the market for a unit price, R_i . We denote the selling season demand from each of the multiple retailers as D_i , where $i \in \{1, 2, \dots, N\}$.

In this section, we build a model where this product must be customized for each of the different retailers or territories. This setup applies to products sold as private-label, store-branded products, or products that must be customized for different regulatory standards in different territories. We further assume that it is not efficient to employ strategies involving delayed differentiation or postponement; in other words, the product design may not be modular or strategic, and marketing considerations limit the effectiveness of postponement strategies (Anand and Girotra 2007). Consequently, the product must be procured and inventoried separately for each of the different retailers/territories.

In Section 7 we present an alternate model where the same generic product is sold through all retailers under identical contract terms, thus enabling further efficiencies from joint production and common inventories. The analysis and the dynamics of the alternate model are somewhat different than the model of this section, but each of the key results presented here holds also in the case with generic products.

Demand in each of the nonoverlapping markets is independent and distributed normally with mean μ_i and variance $\sigma_i^2 = 1/\rho_i$. Each retailer has access to some intelligence about the local market conditions, private information that we model as a signal of demand, Y_i . Following Li and Zhang [2008], we assume that the signal is distributed conditionally normal, or that $Y_i | D_i$ is distributed normally with mean D_i and finite variance $1/t_i$. The signal is thus an unbiased estimate of the true demand, where t_i can be interpreted as the precision or quality of the retailer's local market knowledge. Further, we assume that all agents utilize information signals in a Bayesian fashion (see Winkler 1981). Consequently, retailers, having access to this private information signal can

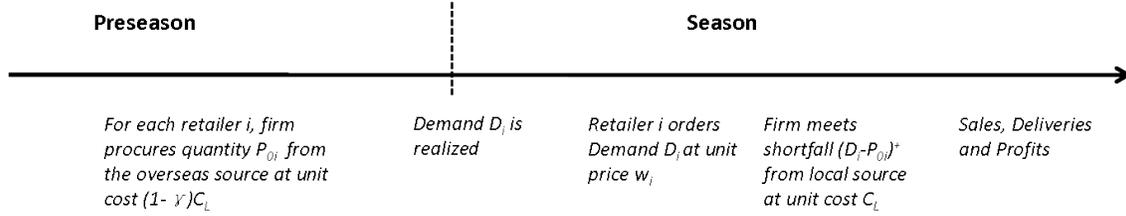


FIGURE 2. Wholesale Price Contracts: Sequence of Events

update the common prior on the demand distribution with the private information signal to obtain the private posterior distribution of demand, $D_i|Y_i$, which is distributed normally with mean $\mu'_i = (\rho_i\mu_i + t_iY_i) / (\rho_i + t_i)$ and variance $\sigma_i'^2 = 1 / (\rho_i + t_i)$.

The firm can produce or source these products using multiple different sources, each with different delivery lead time and associated costs. To simplify our analysis, we will consider two extreme alternate sources for this product: the product can be sourced at unit cost C_L from a *local* source with nearly instantaneous lead times. As an alternative, the product can be sourced at unit cost $(1 - \gamma)C_L$, where $\gamma \in (0, 1)$, from an *overseas* source with a significantly longer lead time. In particular, orders to the overseas source must be placed well ahead of the selling season of the product, which necessitates a *make-to-stock* inventory strategy. In contrast, a *make-to-order* strategy may be employed with respect to the local source. All of the described parameters, except the private information signal Y_i , are assumed to be common knowledge to all agents.¹

3.2. Wholesale Price Contracts: Sequence of Events and Performance. We first examine the conventional supply chain practice of using wholesale price contracts; this will enable us to develop a benchmark for comparison with an alternate contracting form, the Advance Purchase Discount (APD) scheme to be described in Section 3.3. The sequence of events is illustrated in Figure 2. First, well ahead of the selling season, for each retailer i , the firm places an order of P_{0i} customized units from the *overseas source* at unit cost $(1 - \gamma)C_L$. During the selling season, after demand is known, the retailers order the realized demand from the firm at unit price w_i .² Next,

¹In line with Taylor and Xiao [2008] and with most existing work on supply chain contracting, we assume that the firm has full knowledge of retailer-side operational parameters. In particular, the retailer’s signal quality t_i and the retailer’s ordering costs are all discerned from past experience. To the extent that this past experience is sufficient, the assumption that the retailer’s signal quality is common knowledge is appropriate. However, our own experience with applying the results of this study did not show this to be the case. Nevertheless, we first build the basic intuitions under this assumption. Then in Section 4, we relax this assumption to capture a more realistic setting in which the firm is not fully aware of all relevant retailer-side parameters and can only speculate on their distribution.

²Setting the wholesale price is often influenced and constrained by various aspects of the firm’s market strategy, competitive landscape, etc. Within the operations perspective of our paper, we assume that the wholesale price is

the firm examines its inventories; if its³ inventory is not sufficient to meet demand from any of the retailers, the firm procures additional units for that retailer from the local source at unit cost C_L . Finally, each retailer's demand is fully met by supply from the overseas sources and, in some cases, extra production from the local source. Any items left over from the initial orders at the overseas source are worthless to the firm.

In this setup, the retailers place their orders after market demand is realized and thus need not maintain any inventories or face any inventory risk. Each retailer's expected profits are calculated as her margin multiplied by the demand mean. Given the customization necessary for each product, the firm needs to maintain separate inventories for each of the retailers. Each of these setups is individually a classic newsvendor setup with recourse to a reactive capacity source, the local source. Consequently, the firm's overseas order quantities are $P_{0i} = \mu_i + \sigma_i z_\gamma$ for all i , and the consequent expected profits are $\sum_{i=1}^N [(w_i - (1 - \gamma) C_L) \mu_i - C_L \phi(z_\gamma) \sigma_i]$, where $z_\gamma = \Phi^{-1}(\gamma)$ and where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal pdf and cdf respectively.

3.3. The Advance Purchase Discount Scheme. We now propose an alternate contracting form between the firm and its network of retailers, the Advance Purchase Discount (APD) scheme. The setup of the market, the retailers, the firm, the informational environment, and the economics of sourcing from the local and global source are all identical to the wholesale price contract described previously. However, unlike the wholesale price contract, the firm now offers each retailer i an opportunity to place an order in advance of the selling season at a discounted price of $(1 - \delta_i)w_i$, where $\delta_i \in [0, 1]$. The retailer may choose to purchase a certain quantity Q_i during the pre-season at this discounted price. The retailer's tangible administrative cost incurred as a result of placing this additional order is captured by K_i , where $0 \leq K_i \leq w_i \mu_i$.⁴ This administrative cost may include financing costs, order costs, retailer investments in forecasting, and all other tangible fixed costs associated with an additional pre-season order. During the selling season and after the demand uncertainty is resolved, each retailer can order additional units at a price w_i . From the firm's perspective, aside from the APD scheme offered, the operations remain identical to the benchmark wholesale price contract case.

exogenously given, and instead we focus on the examination of novel operational practices that improve the Supply Chain's performance for each given value of wholesale price.

³Throughout the paper, we use "it" in reference to the firm, and "she" in reference to any retailer.

⁴In order to avoid trivial solutions where the administrative cost is extremely prohibitive, we assume that the administrative cost is bounded—in particular, it is lower than the total full-price order-cost for the retailer.

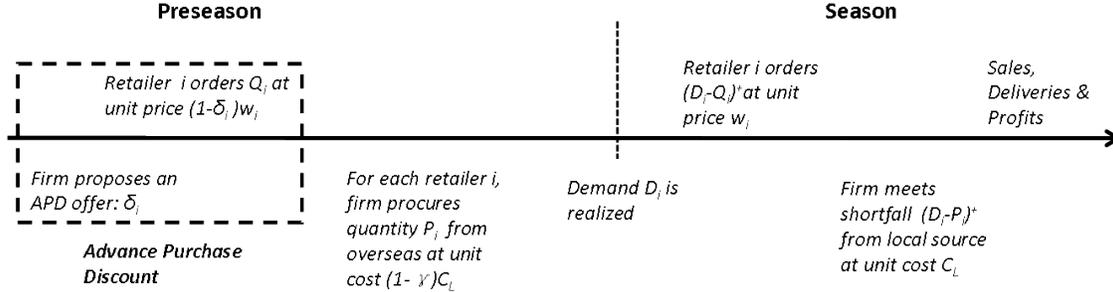


FIGURE 3. The Advance Purchase Discount Scheme: Sequence of Events

The sequence of events now proceeds as in Figure 3. As before, the retailers have access to their private demand signals. In advance of the selling season, the firm proposes the APD contract with discount δ_i . Next, the retailer decides if she wants to participate in the APD scheme. If she decides to participate in the scheme, the retailer updates the common prior on the demand distribution with the private signal to obtain the posterior distribution of demand, $D_i|Y_i$ and orders Q_i units at price $(1 - \delta_i) w_i$. For each retailer i , the firm then places its orders P_i from the overseas source at a unit cost of $(1 - \gamma) C_L$. After demand D_i is realized, the retailers place additional orders for $(D_i - Q_i)^+$ units at price w_i . Next, the firm places orders $(D_i - P_i)^+$ from the local source when the overseas orders cannot satisfy the demand from retailer i . Finally, all deliveries and payments are made and profits are realized.⁵ This setup differs from the previously described benchmark wholesale price contract setup in only one way: the Advance Purchase Discount, which gives retailers the ability to place preseason orders at a discounted price.

Similar schemes have previously been proposed in two different contexts in the existing literature. In contrast to our setup, where the scheme is offered to different retailers within a supply chain, the APD scheme is offered to a market’s atomistic heterogeneous customers in Dana [1998], Gundepudi et al. [2001], and Tang et al. [2004]. These customers have different incentive structures than our retailers, and our analysis will therefore illustrate some novel dynamics. Cachon [2004] and Donohue [2000] also introduce a similar scheme. As in our setup, their scheme is offered by an upstream firm to its retailers, but in the absence of demand information asymmetry, leading to a different analysis.

⁵Observe that in this setup the retailers cannot return “over-ordered” units to the firm. We make this assumption for two reasons. First, return policies are typically adopted to mitigate retail risk by incentivizing retailers to order the supply chain optimal quantities. However, given the APD contract and our two order modes, the supply chain is always coordinated regardless of retailers’ order quantities. Therefore, a return policy does not help coordinate the supply chain. The second reason is that, owing to the nature of customized products, the returned units are useless to the firm; hence, the firm also prefers to avoid such a policy. We relax this assumption in Section 7 to examine the case of generic products, where it may be in the firm’s interest to allow the return of extra units from retailers’ early orders when those units could be used to supply other retailers.

The focus of these two papers is on supply chain coordination. The two-order mode of our setup—retailers can also order after demand is realized to compensate for insufficient early orders—captures the demand characteristics of seasonal products and also eliminates concerns about a loss of supply chain efficiency. That is, in employing either a wholesale price contract or an APD contract, there is full coordination in the firm-retailer supply chain. This allows us to focus on the information-sharing issues central to our study and to abstract away from the previously examined supply chain coordination issues. We next provide an analysis of the equilibrium strategies of the retailers and the firm under this Advance Purchase Discount scheme.

3.4. Strategies under Advance Purchase Discounts. The setup described so far is a multi-stage game of imperfect information, where the private information Y_i of each retailer defines the type space. We describe the Bayesian subgame perfect strategies in this game.

The strategies and equilibrium actions for decisions made after the realization (by the firm and its retailers) of demand follow directly from the setup in Section 3.3. Essentially, each retailer orders the required quantity, if any, from the firm to meet all available demand. Similarly, the firm meets its market-specific shortfall by procuring from the local source. The preseason stage of the game is more interesting. In this stage, the firm must decide on the specific APD parameters to offer and the retailers must choose their advance purchase quantities, if any, on the basis of these parameters. The following lemma characterizes the retailers' strategy in response to the firm's offer.

Lemma. Retailers' Strategy Profile: *Each retailer orders an advance purchase quantity Q_i and earns expected profits $\mathbb{E}\pi_i^R$, where*

$$(3.1) \quad \begin{aligned} Q_i &= \mu'_i + \sigma'_i z_{\delta_i}, & \mathbb{E}\pi_i^R &= (R_i - (1 - \delta_i) w_i) \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i - K_i & \text{if } \delta_i \geq \underline{\delta}_i, \\ Q_i &= 0, & \mathbb{E}\pi_i^R &= (R_i - w_i) \mu_i & \text{otherwise,} \end{aligned}$$

$$(3.2) \quad \underline{\delta}_i = \left\{ \delta_i \mid \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i = K_i \right\},$$

and $z_{\delta_i} = \Phi^{-1}(\delta_i)$.

Proof. All proofs are provided in the Appendix. □

The retailer's strategy profile stipulates her order in the preseason for any given APD contract offered by the firm. The retailer's strategy can be viewed as consisting of two decisions: first, whether the

retailer should participate in the APD scheme, thereby incurring an additional administrative cost; and second, how much should the retailer order if she participates.

We demonstrate that there exists a $\underline{\delta}_i$ such that, at this discount, retailer i is indifferent between participating in the APD scheme or not—that is, the expected profits from participating in the APD scheme equal those from not participating. Moreover, for any discount greater than $\underline{\delta}_i$, retailer i is willing to participate but otherwise is not. We call this threshold level $\underline{\delta}_i$, the retailer’s minimum acceptable discount (MAD).

More specifically, this threshold level is characterized by the discount at which the benefits from the APD scheme are equal to the costs (Equation 3.2). The benefits of the APD scheme to the retailer arise from sourcing some fraction of the product at a discounted price, which is captured by the term $\delta_i w_i \mu_i$. There are two sources of costs from participating in the APD scheme. First, the retailer now places orders in advance of the demand realization and consequently bears some demand–supply mismatch costs; these are captured by the term $w_i \phi(z_{\delta_i}) \sigma'_i$. Second, the retailer now incurs the fixed administrative cost K_i of an additional preseason order.

The retailer compares the benefits and the costs of participating in the discount scheme. If the preseason discount offered is deeper than the threshold discount, then the savings due to sourcing at the discounted price are more than the mismatch and administrative costs and so the retailer chooses to participate in the APD scheme. In contrast with supply chain APD schemes previously considered in literature (e.g., Cachon 2004), we endogenize and explicitly model the retailer’s decision to participate in the APD scheme. This leads to the novel dynamics just described.

Once the retailer decides to participate in the APD scheme, she obtains the private information and faces a newsvendor-like choice, with consequences from ordering too much (penalty of leftover inventories) or from ordering too little (penalty of sourcing later at full price rather than the discounted price). Note that if the retailer decides to participate in the APD scheme, she chooses her order quantity on the basis of all the information available to her. In particular, she uses the private demand information signal to update the common knowledge prior demand and uses the consequent posterior estimates on demand. This leads to an expression for Q_i in Equation (3.1), that is a function of the posterior estimates.

Recall that μ'_i , the posterior estimate of the mean of the demand distribution, is a function of Y_i , we remark that the retailer’s order quantity Q_i is a fully invertible function of the private signal

Y_i . Thus, by observing the retailer's order, *the firm can fully discern the retailer's private signal*. This and previous observations suggest that, if a firm so desires, it can offer a discount substantial enough to induce a retailer's participation in the APD scheme and thereby perfectly infer her private information!

Next, we examine the firm's strategy before outlining the equilibrium outcome.

Lemma. *Firm's Strategy Profile:* *The firm offers each retailer a discount δ_i and places an overseas order P_i for each product,*

$$\begin{aligned} \delta_i &= \underline{\delta}_i, & P_i &= \mu'_i + \sigma'_i z_\gamma && \text{if } K_i \leq \bar{K}_i, \\ \delta_i &= 0, & P_i &= \mu_i + \sigma_i z_\gamma && \text{otherwise,} \end{aligned}$$

where $\bar{K}_i = C_L \phi(z_\gamma)(\sigma_i - \sigma'_i)$. The firm's expected profits $\mathbb{E}\pi^F$ are $\sum_{i=1, \dots, N} \mathbb{E}\pi_i^F$, where

$$\mathbb{E}\pi_i^F = \begin{cases} ((1 - \delta_i) w_i - (1 - \gamma) C_L) \mu_i + (w_i \phi(z_{\delta_i}) - C_L \phi(z_\gamma)) \sigma'_i & \text{if } K_i \leq \bar{K}_i, \\ (w_i - (1 - \gamma) C_L) \mu_i - C_L \phi(z_\gamma) \sigma_i & \text{otherwise.} \end{cases}$$

The firm's preseason strategy consists of two key decisions: what discount δ_i to offer, and how many units P_i to order from the overseas source. These decisions must be made while bearing in mind the retailer's reaction to the contract offered. The discount to offer determines the retailer's decision to participate or not. It is thus driven by two factors: the benefits and the costs of inducing the retailer to participate.

Once the retailer participates in the APD scheme, the firm can infer her private information. This allows the firm to update its demand forecast; in particular, the demand forecast now has a lower variance. Hence the firm can place its overseas orders on the basis of this updated, lower-variance demand forecast, which reduces the firm's demand-supply mismatch costs. Specifically, those costs are reduced by $C_L \phi(z_\gamma)(\sigma_i - \sigma'_i)$. The first part of this expression, $C_L \phi(z_\gamma)$, captures the cost differential between the overseas and local sources, which drives the mismatch costs. The next part, $(\sigma_i - \sigma'_i)$, captures the reduction in variance from obtaining the retailer's private information. In addition to reducing the effective mismatch costs, the firm also transfers an additional amount of these costs to the retailer. In particular, since the retailer now carries inventories, she also bears some of the oversupply risk. These two components constitute the firm's benefits from retailer participation.

From the firm's point of view, there are also costs associated with incentivizing the retailer to participate. The firm must offer a discount deep enough that the retailer's benefits from the discount exceed the sum of her mismatch costs and her administrative cost K_i .

Considering both the benefits and costs of inducing retailer participation, we can see that if the administrative cost for the retailer is small, then the firm's cost to incentivize her participation is smaller and so the firm finds that the potential benefits from offering APD outweigh the costs. Specifically, we show that there exists a threshold administrative cost \bar{K}_i such that, if the administrative cost is lower than this threshold, the firm finds it profitable to incentivize the retailer to participate. A firm that finds this profitable will set the discount just deep enough to make the retailer indifferent between participating and not. As discussed in the previous lemma, this threshold discount is the retailer's MAD, $\underline{\delta}_i$.

The decision regarding the overseas order quantity follows in a relatively straightforward fashion. The firm is now facing a newsvendor situation, albeit with demand forecasts updated on the basis of information discerned from observing the retailer's order.

3.5. Benefits of Advance Purchase Discounts. In this section we examine the effects of deploying an Advance Purchase Discount scheme. The next theorem characterizes conditions and a class of APD contracts that make the retailers, the firm, and the supply chain as a whole better off.

Theorem 1. *For all i such that $K_i < \bar{K}_i$, there exists an advance purchase discount $\delta_i \in (\underline{\delta}_i, \bar{\delta}_i)$ such that retailer i , the firm, and the supply chain are all strictly better off. The benefits are given as*

$$\begin{aligned}\Delta \mathbb{E}\pi_i^R &= \mathbf{1}_{K_i < \bar{K}_i} \left(\delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i - K_i \right) > 0, \\ \Delta \mathbb{E}\pi^F &= \sum_{i=1}^N \mathbf{1}_{K_i < \bar{K}_i} \left(-\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma'_i + \bar{K}_i \right) > 0, \\ \Delta \mathbb{E}\pi_i^S &= \sum_{i=1}^N \mathbf{1}_{K_i < \bar{K}_i} \left(\bar{K}_i - K_i \right) > 0,\end{aligned}$$

where $\bar{\delta}_i = \left\{ \delta_i \mid \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i = \bar{K}_i \right\}$ and where $\underline{\delta}_i$ and \bar{K}_i are defined in the preceding lemmas.

This theorem demonstrates the existence of a set of feasible discounts that make the retailers and the firm strictly better off. Essentially, these benefits accrue owing to a combination of the following four effects, each arising out of offering the Advance Purchase Discount.

- (1) **Effect Admin:** To avail the APD, the retailers must process an additional preseason order, which leads to an *additional administrative cost* of K_i .
- (2) **Effect Discount:** By availing the APD, the *average transfer price* between the firm and the retailers is reduced; this leads to higher expected profits for the retailers but to lower ones for the firm. This effect is captured by the term $\delta_i w_i \mu_i$.
- (3) **Effect Information:** In availing the APD, the *retailers signal to the firm their private information* about demand. Hence, the firm can utilize this information to update its demand forecasts and reduce their variance. Specifically, for demand from each retailer i , the firm's forecast variance is reduced to $\sigma_i'^2$ from the prior demand variance σ_i^2 . This allows the firm to better match its overseas order with the market demand and thereby reduce its demand–supply mismatch costs, as captured by the term $C_L \phi(z_\gamma)(\sigma_i - \sigma_i')$. An even stronger effect is that σ_i' is lower for retailers with better private information. Effect Information is novel in our setting and captures the role of APD in promoting better forecast sharing in supply chains. The literature on consumer-centric APD schemes includes a conceptually similar effect (cf. Tang et al. 2004), but it does not have the interesting interaction with our next effect on demand–supply mismatch risks.
- (4) **Effect Risk:** In availing the APD, the retailers end up bearing some new demand–supply mismatch risk. In the absence of an APD, the retailers did not stock any inventory. After taking advantage of the APD, retailers generally have leftover inventories, and this reduces their expected profits. On the other hand, by the same mechanism the firm ends up transferring some of this oversupply risk to the retailers, reducing its oversupply risk by the same amount. This is captured by the term $w_i \phi(z_{\delta_i}) \sigma_i'$. A similar risk transfer effect is captured by the model in Cachon [2004], but its lack of information asymmetry means that there is no interplay with Effect Information.

The interplay of these effects drives the net advantage (or disadvantage) of an APD scheme for each of the relevant agents. Let's first consider the retailers. Effects Admin and Risk reduce the expected profits of the retailers, and Effect Discount increases their expected profits. The expression

for $\Delta \mathbb{E}\Pi_i^R$ captures the trade-off between these effects. For the range of discounts $\delta_i \in (\underline{\delta}_i, \bar{\delta}_i)$, the net effect is that the retailers are better off by participating in the APD scheme.

Next consider the firm: Effect Discount reduces its average sales price and its expected profits; whereas Effect Information and Risk increase its expected profits. For each retailer that the firm faces with $K_i < \bar{K}_i$, there exists a discount at which the net effect for the firm is positive. Note that the firm cannot actually accrue these benefits unless the retailer has nonnegative benefits of participating in the scheme; although the firm proposes the scheme, it is the retailers who actually participate in it and generate the benefits.

Finally, consider the supply chain as a whole. Effects Discount and Risk are internal effects: they involve transfers between the firm and retailer, and do not affect the supply chain as a whole. The impact of APDs on the supply chain is composed of (i) the benefits due to better transfer of information from the tier (the retailers) that has access to it to the tier (the firm) that makes a crucial supply chain-relevant decision and (ii) the loss due to the additional administrative costs. By definition of \bar{K}_i , the benefit of the information sharing is higher than the administrative cost when $K_i < \bar{K}_i$.

Theorem 1 shows that, if the administrative cost of an additional order is not excessive, then the supply chain as a whole gains by better information sharing. This highlights the previously overlooked advantage of advance purchase discount as a *self-enforcing information-sharing mechanism*. The total benefits accruing to the supply chain by virtue of this superior information sharing can now be divided between the firm and the retailers to make everyone better off. In particular, the extent of the discount, δ_i , determines what fraction of the total gain accrues to the firm and what fraction to the retailer. If $\delta_i = \bar{\delta}_i$ then all benefits accrue to the retailer; on the other hand, if $\delta_i = \underline{\delta}_i$ then all benefits accrue to the firm. In our setup, the firm moves first and proposes a take-it-or-leave-it APD contract, so it enjoys all the bargaining power. Not surprisingly, the firm uses this power to ensure that it retains all the benefits from the APD; the firm does this by setting the discount to be just substantial enough that it ensures the retailers' participation but does not share any benefits beyond that. The next section illustrates this intuition.

3.6. Equilibrium Outcome. The equilibrium outcome follows from the strategy profiles for the retailers and the firm, which were described in Section 3.4. If $K_i \leq \bar{K}_i$, then it is in the firm's interest to offer the Advance Purchase Discount and it is in the retailer's interest to purchase some

quantities in advance. This purchase serves as a signal of the information available to the retailer. On the other hand, if $K_i > \bar{K}_i$, then the firm does not find it profitable to offer the APD and the usual wholesale price contract ensues. The following theorem formalizes this equilibrium outcome.

Theorem 2. *Equilibrium Outcome*

- (1) **Actions:** *The tuple of actions $\{\delta_i^*, Q_i^*, P_i^*\}$ characterizes a perfect Bayesian equilibria of the setup just described. For all i ,*

$$\begin{aligned} \delta_i^* &= \underline{\delta}_i, & Q_i^* &= \mu'_i + \sigma'_i z_{\delta_i}, & P_i^* &= \mu'_i + \sigma'_i z_{\gamma} & \text{if } K_i \leq \bar{K}_i, \\ \delta_i^* &= 0, & Q_i^* &= 0, & P_i^* &= \mu_i + \sigma_i z_{\gamma} & \text{otherwise.} \end{aligned}$$

- (2) **Profits:** *The equilibrium expected profits for retailer i are $\mathbb{E}\Pi_i^R$, and the equilibrium expected profits for the firm, $\mathbb{E}\Pi^F$ are $\sum_{i=1}^N \mathbb{E}\Pi_i^F$, where*

$$\begin{aligned} \mathbb{E}\Pi_i^R &= (R_i - w_i) \mu_i, & \mathbb{E}\Pi_i^F &= (w_i - \gamma C_L) \mu_i - C_L \phi(z_{\gamma}) \sigma'_i - K_i & \text{if } K_i \leq \bar{K}_i, \\ \mathbb{E}\Pi_i^R &= (R_i - w_i) \mu_i, & \mathbb{E}\Pi_i^F &= (w_i - \gamma C_L) \mu_i - C_L \phi(z_{\gamma}) \sigma_i & \text{otherwise,} \end{aligned}$$

and where z_{γ} , z_{δ} , $\underline{\delta}_i$, and \bar{K}_i are as defined in the previous lemmas.

The setup of this model effectively includes a take-it-or-leave-it APD offer, which gives the firm first-mover advantage and all the bargaining power. Consequently, in equilibrium, the retailer obtains the same expected profits as not participating in the APD scheme while the firm appropriates all the supply chain benefits, as demonstrated in Theorem 1. Yet Theorem 1 also suggests that, with alternate bargaining and negotiation mechanisms, the equilibrium contract may be more egalitarian in that it could strictly improve every agent's expected profits.

In Section 4 we will develop a more realistic model of this setup, where the firm will still have the same first-mover advantage but would no longer have full knowledge about each retailer's operations. We will see that such a setup also leads to some restoration of balance; indeed, both the retailers and the firm can have strictly positive benefits from participating in the APD scheme.

In the model just described, the firm has an additional advantage that may not extend to other situations: because the products are customized for each retailer/market, the firm can set retailer- and market-specific discounts. Our analysis in Section 7 examines such a setup, again with reduced discrimination power for the firm. That is, the firm can offer only one discount to all retailers, so the retailers derive strictly positive benefits from the supply chain.

Our model can also be simply modified to capture the case where the administrative cost actually accrues to the retailer. Under this interpretation, the cost K_i becomes the retailer's required participation incentive. Corbett et al. [2004] consider such a participation incentive in their model. The qualitative results described previously will continue to hold, and retailers will derive positive benefits from participating in the APD.

4. RETAILERS' COSTS/INFORMATION PRECISIONS ARE UNOBSERVABLE

In the model described so far, the firm's decision to offer the advance purchase discount scheme (Theorem 2) and the extent of the discount to be offered (Theorem 1) both depend crucially on the retailer's administration cost K_i and the quality of the retailer's private information t_i , which translates into the posterior standard deviation of demand, σ'_i . The administrative cost and the quality of the retailer's private information jointly influence the retailer's minimum acceptable discount (MAD) δ_i . In the previously described equilibrium, the firm offers this minimum acceptable level to the retailer, and consequently all the benefits of the APD scheme end up accruing to the firm. Our assumption in the preceding analysis was that these two parameters are common knowledge—in other words, the firm was fully aware of the retailer's operations and thus knew both the administrative cost and the quality of the private information that the retailer could provide. In practice, these variables are rarely known with certainty by upstream firms. Administrative cost depends on the retailer's internal mechanisms, organizational inertia, and other costs that are incurred at the retailer end. Thus, it is unreasonable to suppose that the upstream firm knows them with certainty. Similarly, the quality of the retailer's information is inherently a function of the kind of market research conducted by the retailer and the efficacy of her information systems, and as such it is not visible to upstream firms in the normal course of business (Taylor and Xiao 2009).

In this section we extend the analysis of previous sections to a more realistic setup, where the firm does not accurately know the retailer's minimum acceptable discount (MAD). This setup presents us with a few key questions: (i) How does the firm implement an efficient advance purchase discount scheme even while being handicapped by the lack of knowledge about the retailers? (ii) Do the benefits of Advance Purchase Discounts that are derived from information sharing persist in such a setup?

Note that unobservability of both the information quality and the administrative cost translate equivalently into an uncertainty about the minimum acceptable discount (Equation 3.2). In the

interest of brevity, in this section the analysis assumes that the firm does not know the administrative cost and that alone leads to the uncertainty about the MAD. The intuition follows similar lines if it were the uncertain information quality that alone drove the uncertainty about the MAD.

4.1. Retailers' and Firm's Strategies. The retailers' profit function and strategy profile do not change from the setup described in Section 3. The firm, however, faces a new challenge in computing the discount to offer. Essentially, it faces a trade-off between the likely results of a deeper discount and a more modest one: a modest discount would yield the firm higher benefits from the APD scheme, though it would reduce the likelihood of retailer participation; whereas a deep discount would increase the chances of retailer participation but at the cost of reduced benefits to the firm. The following example describes the essence of this trade-off and the corresponding strategy for the firm in a simple setting where the retailer's administrative cost could be one of two possible values.

Example: Retailer's administrative cost takes one of two values. Consider a setting where the unobservable administrative cost of retailer i can take one of two values: K_{i1} with probability b or K_{i2} with probability $1 - b$, where $0 \leq K_{i2} \leq K_{i1} \leq \bar{K}_i$.⁶ By Equation (3.2), retailer i 's corresponding MAD becomes either $\underline{\delta}_{i1}$ or $\underline{\delta}_{i2}$, where $0 \leq \underline{\delta}_{i2} \leq \underline{\delta}_{i1} \leq 1$. The following lemma describes the firm's strategy.

Lemma. *The firm chooses the discount level to be $\underline{\delta}_{i1}$ if $b(\bar{K}_i - K_{i1}) \geq (1 - b)(K_{i1} - K_{i2})$ and $\underline{\delta}_{i2}$ otherwise. Recall that, as before, $\bar{K}_i = C_L \phi(z_\gamma)(\sigma_i - \sigma'_i)$.*

Setting a discount δ_i is thus a trade-off between the potential loss of benefits due to an "excessive" discount and the risk of losing the retailer's participation and thereby all the APD benefits. If the firm sets the discount at $\underline{\delta}_{i1}$ then there is no risk of losing the retailer's participation and the APD benefits, but there is a loss of benefits due to an excessive discount. This arises when the firm sets the discount at $\underline{\delta}_{i1}$ when $\underline{\delta}_{i2}$ would have been sufficient (when the retailer's true cost is K_{i2}). The extent of this loss is $K_{i1} - K_{i2}$, and the probability of this contingency is $1 - b$. This consequence of setting the discount at $\bar{\delta}_{i1}$ is captured by the RHS of the critical inequality in the lemma just stated.

⁶If $K_{i1} > \bar{K}_i$ and $K_{i2} \leq \bar{K}_i$, then the firm's optimal strategy would be to simply set the discount at $\underline{\delta}_{i2}$; and if $K_{i1} \geq K_{i2} > \bar{K}_i$, the firm does not offer the APD. Here, we display results for the most interesting non-trivial case where $0 \leq K_{i2} \leq K_{i1} \leq \bar{K}_i$.

On the other hand, if the firm sets the discount at $\bar{\delta}_{i2}$, then it never offers an excessive discount but may lose its share of the benefits from the APD scheme. This happens when the true required discount is $\bar{\delta}_{i1}$ (i.e., when $K_i = K_{i1}$), which happens with probability b . In such a case, the firm loses its share of the benefits from the APD; as before, these benefits are equal to $\bar{K}_i - K_{i1}$. The total expected consequence of this loss is captured by the LHS of the critical inequality in our lemma. The appropriate discount is determined by comparing the expected excessive discount loss from setting the discount at $\bar{\delta}_{i1}$ and the expected loss of APD benefits from setting the discount at $\bar{\delta}_{i2}$.

We now extend the intuition from this example to a case where retailer i 's administrative cost K_i follows a distribution with a general pdf $g_i(\cdot)$ and cdf $G_i(\cdot)$, and positive support in the interval $[A_i, B_i]$. Further more, we assume that $G_i(\cdot)$ has a decreasing reversed hazard rate (DRHR).⁷ The firm's strategy is captured by the following theorem.

Theorem 3. Firm's Strategy Profile: *The firm offers each retailer a discount $\hat{\delta}_i$:*

$$\begin{aligned} \hat{\delta}_i &= \hat{\delta}_i(\hat{k}_i) && \text{if } \hat{k}_i \leq \bar{K}_i, \\ \hat{\delta}_i &= 0 && \text{otherwise.} \end{aligned}$$

Here

$$\hat{\delta}_i(\hat{k}_i) = \left\{ \delta_i \left| \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' = \hat{k}_i \right. \right\},$$

$$(4.1) \quad \hat{k}_i = \mathbf{1}_{A_i \leq k_{i0} \leq B_i} k_{i0} + \mathbf{1}_{k_{i0} < A_i} A_i + \mathbf{1}_{k_{i0} > B_i} B_i,$$

and

$$(4.2) \quad k_{i0} = \left\{ k_i \left| \frac{(\bar{K}_i - k_i) g_i(k_i)}{G_i(k_i)} = 1 \right. \right\}.$$

Theorem 3 illustrates that, when it cannot observe the retailer's administrative cost, the firm still treats the situation *as if* the retailer's administrative cost is known and takes a certain value.

Essentially, the firm uses a ‘‘certainty equivalent’’ value of the unobservable administrative cost and

⁷DRHR is equivalent to log-concavity for a cdf (Chandra and Roy 2001); the assumption is not at all restrictive because it includes most common distributions. Distributions with increasing failure rate—such as the Weibull, gamma, normal, uniform, and log-normal distribution—are all found to be DRHR distributions (Block et al. 1998).

then proceeds to set the discount in exactly the same fashion as in Section 3.⁸ The optimal value for this “certainty equivalent” is defined in Equation (4.1) and (4.2). Not surprisingly, it is closely related to the firm’s beliefs about the distribution of K_i ’s value (i.e., $g_i(\cdot)$, $G_i(\cdot)$, and the support $[A_i, B_i]$). Moreover, as discussed in Section 3, if the optimal “certainty equivalent” is no more than the maximum benefits the firm could receive from employing the APD scheme, \bar{K}_i , then the firm offers the APD to the retailer with the corresponding optimal discount $\hat{\delta}_i(\hat{k}_i)$; otherwise, the firm does not offer the APD.

The foregoing analysis characterizes the firm’s optimal strategy when the administrative cost is unobservable. Next we compare this optimal strategy with the strategy computed in the preceding section, where the administrative cost was common knowledge.

Recall that, when K_i is observable to the firm, the optimal discount is influenced only by the magnitude of K_i and that the higher K_i , the higher the optimal discount. This is no longer true when the firm cannot observe K_i , as illustrated by the following example.

Example: Effect of the Degree of Uncertainty about Retailer’s Cost. Consider three different retailers $i = 1, 2$, and 3 , where the administrative cost K_i is either l_i or h_i with equal probability. According to Theorem 3, the optimal discounts to offer are as computed in Table 1.

Retailer	l_i	h_i	$\hat{\delta}_i (\times 10^{-3})$
1	0.1	0.2	11
2	0.1	0.3	16
3	0.1	2	5

TABLE 1. Discounts for three retailers with different administrative cost distributions. The results are calculated assuming $C_L = 1.8$, $w_i = 1.9$, $\gamma = 0.5$, $\mu_i = 10$, $\sigma_i = 3$, and $\sigma'_i = 0.1$.

The administrative costs for both retailer 2 and retailer 3 are, on average, higher than that of retailer 1, and they are also larger in a first-order stochastic dominance sense. The logic of Section 3 would suggest that the discount offered to either retailer 2 or 3 would be higher than the discount offered to retailer 1. This logic holds as far as retailer 2 is concerned; however, the optimal discount offered to retailer 3 is actually smaller than that offered to either retailer 1 or 2! This example illustrates that neither the mean of the unknown administrative cost nor first-order stochastic dominance fully determines the order of the optimal discount. The uncertainty about the estimate of the

⁸A conceptually similar treatment of a different context is presented in Tang et al. [2009].

unobservable administrative cost is the missing explanatory variable. We next examine its impact on the firm's optimal strategy.

Write retailer i 's administrative cost as $K_i = u_i + s_i \epsilon_i$ with support $[A_i, B_i]$, where ϵ_i is a random variable with DRHR, standard deviation one, and cdf $G_{i0}(\epsilon_i)$ over support $[(A_i - u_i)/s_i, (B_i - u_i)/s_i]$. Then the distribution of K_i has cdf $G_i(k_i|u_i, s_i) = G_{i0}((k_i - u_i)/s_i)$ and a decreasing reversed hazard rate. The standard deviation of the distribution $G_i(k_i|u_i, s_i)$ equals s_i , and increasing (decreasing) s_i corresponds to stretching (contracting) around u_i . In that sense, s_i serves as a measure of the firm's uncertainty about retailer i 's administrative cost K_i . Observe that when $s_i = 0$, the distribution of K_i is degenerate with the mass point at u_i , and then the optimal discount to offer equals $\hat{\delta}_i(u_i)$. The next theorem shows that if s_i is small, the optimal discount is higher than $\hat{\delta}_i(u_i)$, and if s_i is large, the optimal discount is lower than $\hat{\delta}_i(u_i)$.

Theorem 4. *Denote by $\hat{k}_i(u_i, s_i)$ the solution to Equation (4.2) for a distribution with cdf $G_i(k_i|u_i, s_i) = G_{i0}((k_i - u_i)/s_i)$. The optimal discount is $\hat{\delta}_i(\hat{k}_i(u_i, s_i))$. If $s_i \leq s_{i0}$ then $\hat{\delta}_i(\hat{k}_i(u_i, s_i)) \geq \hat{\delta}_i(u_i)$, and if $s_i > s_{i0}$ then $\hat{\delta}_i(\hat{k}_i(u_i, s_i)) < \hat{\delta}_i(u_i)$. Here $s_{i0} = (\bar{K}_i - u_i)g_{i0}(0)/G_{i0}(0)$.*

Theorem 4 points out that, if the uncertainty about the retailer's administrative cost is small then the optimal discount is *higher* than in the case of no uncertainty; conversely, if the uncertainty is large then the firm should set a *lower* discount than in the case of no uncertainty.

This discussion offers the firm an important practical guideline for setting the discount when it faces a retailer about whom it has limited knowledge. The key to setting the appropriate discount lies not only in the estimate of the unobservable administrative cost but also in the uncertainty around it. In essence, when the uncertainty about the retailer's administrative cost is small, a higher discount should be given than in the case of no uncertainty; when the uncertainty is large, a lower discount is optimal. In our experience with implementing APD schemes at real firms, we find that very often the firm does not know the retailer's administrative cost but does have some sense of its typical level and of the uncertainty around this estimate. Armed merely with these two intuitive measures and the insights from our foregoing discussion, the firm has a strong guideline on how to proceed in offering these discounts within its supply chain. We shall illustrate one such experience of applying APDs in Section 5. Next, we examine equilibrium outcomes under limited knowledge.

4.2. Equilibrium Outcome. The equilibrium outcome follows from the strategy profiles just described for the retailers and the firm. If $\hat{k}_i \leq \bar{K}_i$ then the firm finds it profitable to offer an Advance

Purchase Discount, and if $\hat{k}_i \geq K_i$ then retailer i finds it in her interest to purchase some quantities in advance. This purchase signals the retailer's private information. If $\hat{k}_i > \bar{K}_i$ or $\hat{k}_i < K_i$ then either the firm does not find it profitable to offer the APD scheme or retailer i won't participate given an insufficient discount. As a result, the usual wholesale price contract outcome ensues. The following theorem formalizes this equilibrium outcome.

Theorem 5. *Equilibrium Outcome*

- (1) **Actions:** The tuple of actions $\{\hat{\delta}_i^*, \hat{Q}_i^*, \hat{P}_i^*\}$ characterizes a perfect Bayesian equilibria of the setup just described. For all i ,

$$\begin{aligned} \hat{\delta}_i^* &= \hat{\delta}_i, & \hat{Q}_i^* &= 0, & \hat{P}_i^* &= \mu_i + \sigma_i z_\gamma & \text{if } \hat{k}_i \leq \bar{K}_i \text{ and } \hat{k}_i < K_i, \\ \hat{\delta}_i^* &= \hat{\delta}_i, & \hat{Q}_i^* &= \mu'_i + \sigma'_i z_{\delta_i}, & \hat{P}_i^* &= \mu'_i + \sigma'_i z_\gamma & \text{if } \hat{k}_i \leq \bar{K}_i \text{ and } \hat{k}_i \geq K_i, \\ \hat{\delta}_i^* &= 0, & \hat{Q}_i^* &= 0, & \hat{P}_i^* &= \mu_i + \sigma_i z_\gamma & \text{if } \hat{k}_i > \bar{K}_i. \end{aligned}$$

- (2) **Profits:** The equilibrium expected profits for retailer i are $\mathbb{E}\hat{\Pi}_i^R$, and the equilibrium expected profits for the firm $\mathbb{E}\hat{\Pi}^F$ are $\sum_{i=1}^N \mathbb{E}\hat{\Pi}_i^F$, where

$$\begin{aligned} \mathbb{E}\hat{\Pi}_i^R &= (R_i - w_i) \mu_i + \hat{k}_i - K_i, & \mathbb{E}\hat{\Pi}_i^F &= (w_i - \gamma C_L) \mu_i - C_L \phi(z_\gamma) \sigma'_i - \hat{k}_i & \text{if } K_i \leq \hat{k}_i \leq \bar{K}_i, \\ \mathbb{E}\hat{\Pi}_i^R &= (R_i - w_i) \mu_i, & \mathbb{E}\hat{\Pi}_i^F &= (w_i - \gamma C_L) \mu_i - C_L \phi(z_\gamma) \sigma_i & \text{otherwise,} \end{aligned}$$

where \hat{k}_i and $\hat{\delta}_i$ are as defined in Theorem 3.

When the retailer's administrative cost is unobservable, the firm treats the retailer as if her true cost is \hat{k}_i . The firm then offers the equilibrium discount in the same fashion as when K_i is observable, and all the resulting expected profits and actions retain the same forms as in Theorem 2 but with K_i replaced by its "certainty equivalent" \hat{k}_i . However, two differences should be noted here. First, since \hat{k}_i might be lower than the retailer's true administrative cost, the retailer's participation is not ensured even when the firm finds it profitable to incentivize the retailer's participation with an appropriate discount $\hat{\delta}_i^*$. This means that each retailer's participation condition must be considered even when a sufficiently deep discount is offered, which leads to the first inequality of the new condition $K_i \leq \hat{k}_i \leq \bar{K}_i$. Similarly, there is also a chance that \hat{k}_i is greater than the retailer's true administrative cost, in which case the retailer could earn a positive benefit under equilibrium; this leads to the term $\hat{k}_i - K_i$ in the retailer's new expected profits.

4.3. Benefits of Advance Purchase Discounts under Unobservability. In this section, we investigate how the firm's limited knowledge about retailer parameters influences the benefits arising from APDs. We describe the dynamics formally in the next theorem.

Theorem 6. *Equilibrium Benefits:* *If the retailer's administrative cost K_i is unobservable to the firm then, on the equilibrium path, the participating retailers, the firm, and the supply chain are all better off under APDs. The benefits are given as*

$$\begin{aligned}\Delta \mathbb{E} \hat{\Pi}_i^R &= \mathbf{1}_{K_i \leq \hat{k}_i \leq \bar{K}_i} (\hat{k}_i - K_i) \geq 0, \\ \Delta \mathbb{E} \hat{\Pi}^F &= \sum_{i=1}^N \mathbf{1}_{K_i \leq \hat{k}_i \leq \bar{K}_i} (\bar{K}_i - \hat{k}_i) \geq 0, \\ \Delta \mathbb{E} \hat{\Pi}^S &= \sum_{i=1}^N \mathbf{1}_{K_i \leq \hat{k}_i \leq \bar{K}_i} (\bar{K}_i - K_i) \geq 0,\end{aligned}$$

where \hat{k}_i is defined as in Theorem 3.

Theorem 6 shows that, when the APD scheme is implemented, every participating agent in the supply chain is better off. This result differs from the case where retailers' administrative costs are observable to the firm in two important ways. First, when administrative cost is unobservable to the firm, the retailer has some further private information that allows her to extract some *information rents* $\hat{k}_i - K_i$ from the firm, which consist of the difference between her "certainty equivalent" administrative cost and her real one. As a result, instead of setting the discount at each retailer's MAD to appropriate all the benefits accruing to the supply chain from the information sharing $\bar{K}_i - K_i$, the firm now gets $\bar{K}_i - \hat{k}_i$ and the remaining benefits $\hat{k}_i - K_i$ go to the retailer. Second, given unobservability, even when the firm wants the retailer to participate in the APD scheme and sets the appropriate discount $\hat{\delta}_i^*$, it remains possible that this discount will turn out to be insufficient from the retailer's perspective. In such cases, the retailer does not participate and the supply chain does not appropriate the potential benefits from the APD. The indicator $\mathbf{1}_{K_i \leq \hat{k}_i \leq \bar{K}_i}$ takes this into account. The incidence of such outcomes reduces the expected gains to the supply chain from employing the APD scheme, though they still remain positive. In short, the unobservability leads to a more egalitarian distribution of the gains from the APD scheme, but the total magnitude of the gains is somewhat reduced.

Note here that, under unobservability, the firm's incentive to set the discount is not aligned with the supply chain's incentives. The supply chain prefers the highest possible discount so as to include all the retailers with nonnegative $\bar{K}_i - K_i$. Yet the firm has an additional goal in setting the discount: it cares not only about ensuring the participation of all information-rich retailers and maximizing the total surplus generated by the APD contract but also about the share of supply chain benefits that it can appropriate. This second concern is driven by the uncertainty about the retailers' administrative costs. The following corollary illustrates the effect of this uncertainty.

Corollary. *Under the implementation of the APD scheme with discount $\hat{\delta}_i^*$ characterized in Theorem 4 and 5:*

- (i) if $s_i \leq s_{i0}$, then $\Delta\mathbb{E}\hat{\Pi}^R(\hat{k}_i(u_i, s_i)) \geq \Delta\mathbb{E}\hat{\Pi}^R(u_i)$ and $\Delta\mathbb{E}\hat{\Pi}^F(\hat{k}_i(u_i, s_i)) \leq \Delta\mathbb{E}\hat{\Pi}^F(u_i)$;
- (ii) if $s_i > s_{i0}$, then $\Delta\mathbb{E}\hat{\Pi}^R(\hat{k}_i(u_i, s_i)) < \Delta\mathbb{E}\hat{\Pi}^R(u_i)$ and $\Delta\mathbb{E}\hat{\Pi}^F(\hat{k}_i(u_i, s_i)) > \Delta\mathbb{E}\hat{\Pi}^F(u_i)$.

Recall that $\Delta\mathbb{E}\hat{\Pi}^R(\hat{k}_i(u_i, s_i))$ and $\Delta\mathbb{E}\hat{\Pi}^F(\hat{k}_i(u_i, s_i))$ represent the retailer's and firm's respective benefits from the APD scheme under unobservability, and $\Delta\mathbb{E}\hat{\Pi}^R(u_i)$ and $\Delta\mathbb{E}\hat{\Pi}^F(u_i)$ represent those benefits under full observability. The corollary shows that, when the firm's uncertainty about the retailer's administrative cost is small, its benefits from the implementation of an APD scheme are lower than if it knew her costs with certainty. At the same time, the retailer's benefits increase when the uncertainty about her administrative cost is small. However, when the uncertainty becomes larger than a threshold value, the retailer's benefits from the APD scheme are lower than in the full observability case while the firm's benefits become higher. At this point, the firm is too uncertain about the retailer and essentially loses interest in attempting to induce her participation.

This result offers two insights into the strategic behavior of both the firm and the retailer. From the firm's point of view, it is beneficial to collaborate with a retailer about whom it has either full knowledge or no knowledge at all. In either case, the optimal discount offered would be a very modest one and thus would not hurt the firm much in benefits. This all-or-nothing policy shares the same intuition—though in a different setting—regarding how a manufacturer should choose its forecasting newsvendor described in Taylor and Xiao [2008].

On the retailer's side, the corollary provides the retailer an incentive to strategically communicate information about her administrative cost. The retailer will receive the deepest discount when the firm's uncertainty about her administrative cost is neither too big nor too small. In order to encourage the firm to set a high discount, the retailer should not completely hide her administrative

cost information from the firm, since doing so would lead the firm to react with a safe strategy: setting a modest discount to ensure that it gains a large part of the APD benefits. Nor should the retailer communicate her cost information fully, since then the firm would take full advantage and set the discount as if the retailer's MAD were observable. A conceptually similar dilemma is observed in a different context by Kim and Netessine [2009]. In their paper, a manufacturer and a supplier collaborate in the product design stage to reduce the uncertainty about component production cost; the supplier fears revealing his proprietary cost information and thus is reluctant to fully collaborate.

5. IMPLEMENTATION AT COSTUME GALLERY

In this section, we report on the implementation of our proposed APD scheme at a New Jersey-based fashion apparel wholesaler that sources extensively from Asia and retails through an expanding U.S. network. This allows us to (i) validate the assumptions and setup of our model, (ii) illustrate an implementation of APDs, and (iii) enumerate the benefits of applying an APD scheme. The interested reader is referred to Girotra and Tang [2009] and Girotra and Tang [2010] for a detailed description of Costume Gallery's business context and the benefits from implementation of APDs.

5.1. Background of Costume Gallery. Costume Gallery is a privately held, New Jersey-based wholesaler of dance costumes. In the U.S. market, it is among the top three wholesalers of dance costumes; sales in 2005 amounted to about US\$30 million.⁹ Costume Gallery has been family run since its inception in 1957; the third generation of the family took over in 1997 and has been instrumental in bringing scientific management principles to the enterprise.

Costume Gallery was founded on the premise of excellent customer service. The strategic focus at Costume Gallery is on fully satisfying its customers at all costs. This strategic focus translates into a wide choice of customizable costumes and a target 100% fill rate. At any time, Costume Gallery has as many as 500 styles in its catalog; if a requested style is not in stock, then Costume Gallery can produce it to spec for the customer.

Costume Gallery's supply chain is illustrated in Figure 4. Costume gallery sells most of its merchandise through dance schools. Typically, the end customer is a student enrolled in dance classes at dance schools. The instructor at a dance school plans a dance production and then decides on

⁹Costume Gallery is privately held and does not disclose its financial performance. This number is the authors' estimate.

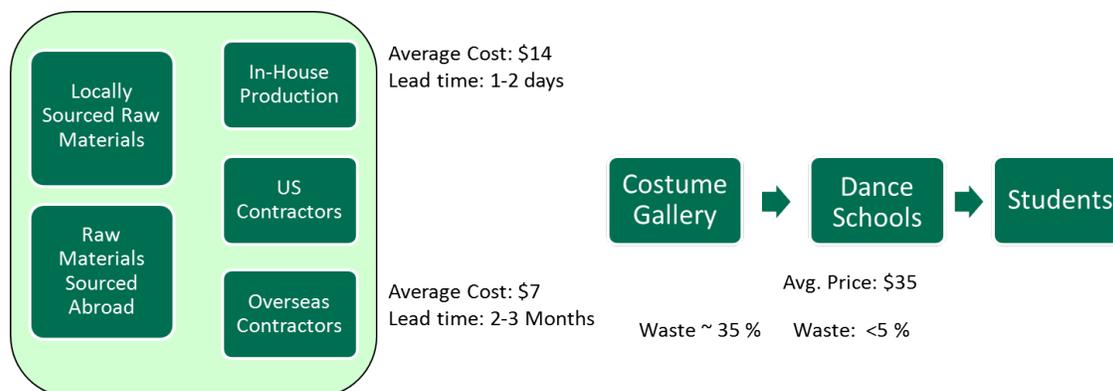


FIGURE 4. The Dance Costume Business

the appropriate costumes. The school then procures the costumes from Costume Gallery and sells them to the students at a small markup. Despite this markup, dance schools do not usually hold any inventory or bear any of the supply chain risks; all inventory costs and risks are traditionally borne by Costume Gallery.

Costume Gallery can produce the dance costumes in-house and can source their production (domestically or overseas). Costumes produced in-house cost about \$14 on average and can be produced with a lead time of 1–2 days. Very small lots can be produced at Costume Gallery. As an alternative, costumes can be sourced—usually from Asia. Such fashion apparel is typically sourced at a total cost that is just 30%–40% of the western production cost; however, the offshore production and shipping lead time can be 2–3 months. Because there is a substantial cost saving in sourcing from Asia, Costume Gallery would prefer to source as much of its production from Asia as possible. Costume Gallery operates two business units: one deals in generic dance products and the other procures products that are customized to meet dance school requirements. Both generic and customized products can be manufactured at the in-house or Asian facilities. Because of product design constraints and the difficulty of finding skilled labor, Costume Gallery does not use any form of postponement for the customized products.

The dance costume business is highly seasonal: 90% of the annual demand occurs in the second half of April, which coincides with the end-of-school-year dance performances. The timeline of the dance costume business is illustrated in Figure 5. Costume Gallery typically starts designing costumes in July of the previous year. These costumes are profiled in a catalog that is sent out to dance schools in early August. Dance schools finalize their enrollment in November and, by December, they have a good idea of the theme for the dance performance in April. Therefore, the schools have potential



FIGURE 5. Timeline of the Dance Costume Business

access to information on the demand for different dance costumes: they can restrict the size of their classes and fix the theme of the dance. This may come at the cost of additional business, but if the schools wanted to, they can do this. However, there is still some residual uncertainty caused by changing sizes of students and by additions to or dropouts from classes. Given this residual uncertainty, schools behave strategically and do not place any orders with Costume Gallery at this time.

To meet the April demand, Costume Gallery must place its overseas orders by January in view of the 2–3-month lead time required. Typically, these orders must be placed before the Chinese New Year holiday, which occurs in late January or early February. These orders are placed on the basis of the best information that Costume Gallery has about the possible demand from each dance school. Dance schools have a far better idea of the likely demand but, because of their residual uncertainty about this demand, they have no incentive to place any orders early in the season. Therefore, Costume Gallery’s overseas orders are placed without using the demand-relevant information available to dance schools. From January to April, Costume Gallery obtains increasingly accurate demand information and uses its in-house production (and local contractors) to meet demand above the quantities already sourced from Asia.

5.2. Application of APDs at Costume Gallery. The supply chain and business cycle of Costume Gallery display the essence of our stylized model. Costume Gallery corresponds to the firm in our model and the dance schools to its network of retailers. Like the firm, Costume Gallery has two production modes: one cheaper with longer lead time (Asia); the other more expensive with shorter lead time (in-house). Like the retailers, dance schools have access to superior knowledge about future demand. In the absence of advance purchase discounts, however, the schools have

no incentive to either gather this superior knowledge or place an early order; thus the information is never shared and utilized to improve supply chain efficiency. Our theoretical model suggests that, by implementing Advance Purchase Discounts, Costume Gallery can elicit accurate, timely information from the dance schools that can then be used to improve sourcing decisions.

5.3. Estimation of the Benefits from APDs. In this section, we provide a rough estimate of the potential benefits that Costume Gallery can obtain if the advance purchase discount scheme suggested by our theory is implemented. To estimate this number, we used actual data from the operations of Costume Gallery to the greatest extent possible. We have made every effort for these numbers to be representative and applicable to a wide set of apparel businesses. Nonetheless, as with any case-based data, the numbers estimated in this section should only be viewed as indicative of the potential gains possible; they should not be construed as a rigorous empirical estimate for a typical firm in this industry.

For ease of illustration, we restrict our attention to the customized products business unit at Costume Gallery. Typically, Costume Gallery does not have extensive knowledge about the operations of its retailers, the dance schools. In particular, it is difficult for Costume Gallery to assess the exact administrative order costs that each dance school would incur in placing an early order. Thus, the theoretical results from Section 3 on customized products, and from Section 4 on limited information are most applicable to our setup.

In order to estimate the benefits from our proposed strategy, we consider a subset of costumes and dance schools. There is almost no difference in the operational economics of different costumes in this subset (or between the vast majority of costumes). Hence we build this rough estimate using the economics of a representative costume, “Dream Ballet”. The design details, required workmanship, level of customization, and costs of materials for this costume are all fairly representative of a typical product at this firm.

Each Dream Ballet costume is offered to a dance school at a price of \$35 (w_i in our model), whereafter the dance school normally charges a \$5 commission and sells it to each student at a price of \$40 (R_i). It costs Costume Gallery \$14 (C_L) to produce such a costume in-house; however, if Costume Gallery outsources the production to Asia, then the total cost—including shipping and all other incidentals—is estimated to be \$7 ($\gamma = 0.5$).

The demand D_i from our model corresponds to the number of units demanded, by a particular dance school, of a particular costume (style) in a particular size. This corresponds to demand for one SKU to be sourced. We refer to this as an order, and we refer to D_i as the “order size”. To estimate a forecast distribution for D_i , we employ the A/F ratio method discussed in Cachon and Terwiesch [2006]. We estimate Costume Gallery’s order demand forecast to be normally distributed, with a mean of 44.3, and a standard deviation of 19.4 units. Next we must obtain the demand forecasts that the dance schools construct, which are needed to estimate the posterior distribution of demand that Costume Gallery could construct if they obtained early order information. This, in turn, requires the forecasts that schools would construct if they could avail themselves of Advance Purchase Discounts. In absence of APDs, the schools do not construct any forecasts and so the data needed to estimate this distribution do not exist. However, based on the authors’ conversations with the management at Costume Gallery, we estimate that the forecast errors can probably be reduced by about 50% if it used advanced information on each school’s enrollment size and dance theme (both of which are school-level private information). Hence, we assume that the posterior demand distribution has a standard deviation that is half that of the prior demand distribution—in other words, $\sigma'_i = 10$.

In order to implement the proposed APD scheme, Costume Gallery must calculate the administrative costs for dance schools to participate. The administrative costs depend on dance schools’ internal mechanisms, organizational inertia, and other costs incurred at their end to acquire the information. These include lost revenues due to restricting the class size, fixing the dance theme, the costs of capital required, labor costs for measuring students an additional time (students sizes often change during the school year). Costume Gallery does not have a good idea of these costs. Moreover, they believe that these costs differ widely across different schools.

Some schools are located in low-cost regions and are staffed by cheaper teachers. Other schools employ professional dancers as teachers, and their costs are expected to be higher. In addition, Costume Gallery has a better sense of the costs at schools with which it’s had a long collaboration than at schools that are new clients. To deal with this heterogeneity in school types, we consider three different classes of schools. Type A are high cost schools with which Costume Gallery has a long-standing relationship, type B are high-cost schools that have recently begun purchasing from Costume Gallery, and type C schools are low-cost schools. Costume Gallery estimates the administrative cost per order for each of the three school types separately, as summarized in Table 2.

School Type	Characteristics	Estimate of Order Costs u_i	Standard Deviation of Estimate of Order Costs s_i
A	High Cost, Long Relationship	40	4
B	High Cost, Short Relationship	40	25
C	Low Cost	20	25

TABLE 2. Dance Schools' Administrative Costs

School Type	# of Orders from Schools in This Category	Discount	Percentage of Schools that Avail APDs	Expected Increase in Profits for Costume Gallery (US\$)
A	25	5.14%	75.78%	183.84
B	29	3.74%	33.65%	224.95
C	102	2.76%	50.93%	1658.09

Total Increase in Profits (US\$)	2,066.88
Pre-APD Net Profits (US \$)	12,093.90 [†]
Increase in Net Profits (%)	17.09%

TABLE 4. Estimated Benefits of Advance Purchase Discounts

[†]Profits on sales of US\$0.24 million covered by this illustration (computed assuming 5% net profitability of the firm).

Using the results from Theorems 3 and 6, we estimate the benefits for Costume Gallery of employing Advance Purchase Discounts; the results are reported in Table 4. Our estimates indicate that Costume Gallery could increase its net profits by about 17% if it utilized APDs. Although we provide this analysis for only a subset of schools and products, we believe that these numbers are representative of Costume Gallery's product lines and also of other firms in similar industries.

6. CONCLUSION

This study proposes the mechanism of an Advance Purchase Discount contract that enhances the effectiveness of a firm's global sourcing strategy by incentivizing *accurate information communication* from downstream tiers in its supply chain *at a time chosen* by the firm. The strategy involves offering a discount to downstream tiers for orders that they place before a specific time, which is usually chosen by the firm to coincide with the last possible chance for overseas orders. Acting in their own best interests, the downstream retailers use the best information accessible to them in making the early orders. We show that observing these orders enables the upstream firm to infer the information available to the downstream retailers; then the firm can make its global sourcing

decisions using the demand information inferred. The principal effect of the early order opportunity is to synchronize the timeline of orders placed by retailers with the timeline of decisions that the global sourcing firm must make. This alignment drives the timely communication of accurate demand information from the downstream retailers to the firm. We demonstrate the benefits of this scheme, identify the conditions under which it is most potent, and isolate the benefits arising from improved information sharing and redistribution of risk. We then demonstrate that, taken together, these benefits can improve the performance of each agent in the supply chain and also of the supply chain as a whole. We demonstrate the robustness of these results by considering two different product designs: one where products are customized for each market and one where the same generic product is sold in all markets.

We also examine the application of these contracts in a realistic setting. We analyze our proposal for employing APD contracts in the context of practical constraints, as when a firm has limited knowledge about its retailers' operations. We characterize the optimal strategy in terms of a "certainty-equivalent" value of the unobserved parameter, a value that the firm can use as if it knew the unobserved parameter for sure. We remark that this value leads to an asymmetric and "degree-of-unobservability dependent" departure from the full-observability design of the Advance Purchase Discount. When the uncertainty in the unobserved parameter is small, the optimal discount is *higher* than in the case of full observability; and when the uncertainty is large, the firm should set a *lower* discount. It is significant that this analysis affirms that the benefits of APDs (improving each agent's profits as well as the supply chain's profits as a whole) persist even under this practical constraint.

To validate the applicability of our analytical setup, estimate the benefits accruing from our prescriptions, and develop practical guidelines for application, we close this study by reporting on the implementation of our APD scheme at Costume Gallery, a small-scale U.S.-based firm in the apparel business. We find that our model accurately captures the business context for this firm. The proposed contract shares its structure with commonly used payment terms and thus found limited resistance in implementation. Further, using data derived from the operations at Costume Gallery, we estimate that applying the APD scheme could yield a 17% increase in the firm's net profits.

Our proposal for APD contracts suffers from several key limitations. Designing an optimal APD contract requires the estimation of parameters related to retailers' operations to an infeasible degree of accuracy. Our analysis of the case with limited knowledge of retailer operations addresses these

concerns to a certain extent; however, for a productive implementation of an APD scheme, the firm must have extensive archival data on its own and its retailers' forecasting abilities, as well as good cost accounting systems, in order to measure accurately the benefit of sourcing overseas. This may not be possible for many firms. Further more, computing and administering these discounts require managerial skills that may not be available to many small organizations. Finally, the benefits of this scheme are highly sensitive to the extent of the discount offered. Yet practical constraints arising from industry practice and long-held relationships with partners prevent the firm from implementing an optimal version of the APD contract, and this may significantly reduce the accrued benefits.

This study provides several new directions for future research. Our preliminary analysis assumed that the demand for multiple products and from multiple retailers is uncorrelated. However, in practice we find demand to be correlated significantly across different products and across different markets. In such a situation, early order information would be useful for improving the demand forecasts of multiple products and markets. In the absence of modeling such dynamics, our current analysis underestimates the value of information acquired from Advance Purchase Discounts and consequently underestimates their benefits. A study taking into account the correlation structure of demand and the consequent dynamics holds significant promise.

Our current Bayesian information structure with conjugate demand-signal pairs, while theoretically appealing and analytically tractable, does not capture an important facet of information acquisition: the dependence of the quality of information obtained on the extent of the discount. An alternate modeling setup that captures this real-life feature is likely to provide useful insights into the science of administering APDs.

Another avenue for future research is in estimating parameters that allow for effective administration of Advance Purchase Discounts. As we have emphasized, knowledge of many retailer side parameters is critical in this analysis. As in our illustration of the Costume Gallery, some parameters (such as the retailers' forecasting abilities) are hard to observe directly even with extensive archival databases. Perhaps a sophisticated empirical strategy can be found to impute some of these parameters from other demand-related actions by the retailers. Such a strategy holds significant promise and could make widespread implementation of APDs a reality.

REFERENCES

- M. Akan, B. Ata, and J. Dana. Revenue management by sequential screening. *Working Paper*, 2009.
- K. Anand and K. Girotra. The strategic perils of delayed differentiation. *Management Science*, 53(5):697, 2007.
- K. Anand and M. Goyal. Strategic information management under leakage in a supply chain. *Management Science*, 55(3):438–452, 2009.
- K. Anand and H. Mendelson. Information and organization for horizontal multimarket coordination. *Management Science*, 43(12):1609–1627, 1997.
- A. Arya and B. Mittendorf. Using return policies to elicit retailer information. *The RAND Journal of Economics*, 35(3):617–630, 2004.
- Y. Aviv. The effect of collaborative forecasting on supply chain performance. *Management Science*, 47(10):1326–1343, 2001.
- Y. Aviv. Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing & Service Operations Management*, 4(1):55, 2002.
- F. Bernstein, P. Guo, and Y. Wang. Advance purchase contracts in a multi-location system with transshipment. *Working Paper*, 2009.
- H. Block, T. Savits, and H. Singh. The reversed hazard rate function. *Probability in the Engineering and Informational Sciences*, 12:69–90, 1998.
- G. Cachon. The Allocation of Inventory Risk in a Supply Chain: Push, Pull, and Advance-Purchase Discount Contracts. *Management Science*, 50(2):222–238, 2004.
- G. Cachon and M. Fisher. Supply chain inventory management and the value of shared information. *Management science*, 46(8):1032–1048, 2000.
- G. Cachon and M. Lariviere. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science*, 47(5):629–646, 2001.
- G. Cachon and C. Terwiesch. *Matching supply with demand: an introduction to operations management*. McGraw-Hill, 2006.
- N. Chandra and D. Roy. Some results on reversed hazard rate. *Probability in the Engineering and Informational Sciences*, 15(01):95–102, 2001.
- L. Chen and H. Lee. Information Sharing and Order Variability Control Under a Generalized Demand Model. *Management Science*, 55(5):781–797, 2009.

- S. Cho and C. Tang. Advanced booking programs for managing supply, demand, and price risks in a two-Level supply chain. *Working Paper*, 2009.
- C. Corbett, D. Zhou, and C. Tang. Designing supply contracts: Contract type and information asymmetry. *Management Science*, 50(4):550–559, 2004.
- J. Dana, Jr. Advance-Purchase Discounts and Price Discrimination in Competitive Markets. *Journal of Political Economy*, 106(2):395–422, 1998.
- K. Donohue. Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Science*, 46(11):1397–1411, 2000.
- M. Fisher and A. Raman. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations research*, 44(1):87–99, 1996.
- M. Fisher, J. Hammond, W. Obermeyer, and A. Raman. Making Supply Meet Demand in an Uncertain World. *Harvard Business Review*, 72:83–83, 1994.
- M. Fisher, K. Rajaram, and A. Raman. Optimizing Inventory Replenishment of Retail Fashion Products. *Manufacturing & Service Operations Management*, 3(3):230, 2001.
- K. Girotra and W. Tang. Strategic Behavior in Supply Chains: Information Acquisition. In C. Tang and S. Netessine, editors, *Consumer-Driven Demand and Operations Management Models*, volume 131, chapter 18, pages 459–477. Springer, 2009.
- K. Girotra and W. Tang. Costume gallery inc. *INSEAD Teaching Case*, 2010.
- P. Gundepudi, N. Rudi, and A. Seidmann. Forward vs. spot buying of information goods on Web: analyzing the consumer decision process. *Journal of Management Information Systems*, 18(2): 107–131, 2001.
- A. Ha and S. Tong. Contracting and information sharing under supply chain competition. *Management science*, 54(4):701, 2008.
- A. Ha, S. Tong, and H. Zhang. Sharing imperfect demand information in competing supply chains with production diseconomies. 2008.
- J. Hammond and A. Raman. Sport Obermeyer, Ltd. *Harvard Business School Case*, pages 9–695, 1994.
- W. H. Hausman, H. L. Lee, and U. Subramaniam. Global logistics indicators, supply chain metrics and bilateral trade patterns. *World Bank Policy Research Working Paper*, (3773), November 2005.
- E. Kayis, F. Erhun, and E. Plambeck. Delegation vs. Control of Component Procurement. Technical report, Department of Management Science and Engineering and Graduate School of Business,

- Stanford University, Stanford, CA., 2007.
- S. Kim and S. Netessine. Supply chain collaboration under information asymmetry. *Working Paper*, 2009.
- H. Lee, K. So, and C. Tang. The value of information sharing in a two-level supply chain. *Management science*, 46(5):626–643, 2000.
- C. Li and F. Zhang. Advance demand information, price discrimination, and pre-order strategies. *Working Paper*, 2009.
- L. Li. Information sharing in a supply chain with horizontal competition. *Management Science*, 48(9):1196–1212, 2002.
- L. Li and H. Zhang. Confidentiality and Information Sharing in Supply Chain Coordination. *Management Science*, 54(8):1467–1481, 2008.
- K. McCardle, K. Rajaram, and C. Tang. Advance Booking Discount Programs Under Retail Competition. *Management Science*, 50(5):701, 2004.
- J. Milner and P. Kouvelis. On the complementary value of accurate demand information and production and supplier flexibility. *Manufacturing & Service Operations Management*, 4(2):99–113, 2002.
- J. Milner and P. Kouvelis. Order quantity and timing flexibility in supply chains: The role of demand characteristics. *Management Science*, 51(6):970–985, 2005.
- J. Miyaoka and W. Hausman. How Improved Forecasts Can Degrade Decentralized Supply Chains. *Manufacturing & Service Operations Management*, 10(3):547, 2008.
- O. Ozer and W. Wei. Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Science*, 52(8):1238, 2006.
- O. Ozer, O. Uncu, and W. Wei. Selling to the newsvendor with a forecast update. *European Journal of Operational Research*, 182(3):1150–1176, 2007.
- M. Parlar and Z. Weng. Designing a firm’s coordinated manufacturing and supply decisions with short product life cycles. *Management Science*, 43(10):1329–1344, 1997.
- B. Pasternack. Optimal pricing and return policies for perishable commodities. *Marketing science*, pages 166–176, 1985.
- G. Perakis and G. Roels. Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188, 2008.

- Z. Ren, M. Cohen, T. Ho, and C. Terwiesch. Sharing forecast information in a long-term supply chain relationship: the role of customer review strategy. *Operations Research*, 2008.
- X. Su. Consumer returns policies and supply chain performance. *Manufacturing & Service Operations Management*, 11(4):595–612, 2009.
- C. Tang, K. Rajaram, A. Alptekinoglu, and J. Ou. The Benefits of Advance Booking Discount Programs: Model and Analysis. *Management Science*, 50(4):465, 2004.
- W. Tang, J. Bearden, and I. Tsetlin. Ultimatum Deadlines. *Management Science*, 55(8):1423–1437, 2009.
- T. Taylor. Channel coordination under price protection, midlife returns, and end-of-life returns in dynamic markets. *Management Science*, 47(9):1220–1234, 2001.
- T. Taylor. Sale timing in a supply chain: When to sell to the retailer. *Manufacturing & Service Operations Management*, 8(1):23, 2006.
- T. Taylor and W. Xiao. Does a manufacturer benefit from selling to a better-forecasting retailer. 2008.
- T. Taylor and W. Xiao. Incentives for retailer forecasting: rebates versus returns. *Management Science*, 8, 2009.
- R. Winkler. Combining probability distributions from dependent information sources. *Management Science*, 27(4):479–488, 1981.
- J. Xie and S. Shugan. Electronic tickets, smart cards, and online prepayments: When and how to advance sell. *Marketing Science*, 20(3):219–243, 2001.

APPENDIX

7. BENEFITS OF ADVANCE PURCHASE DISCOUNTS: GENERIC PRODUCTS

Here we present an alternate version of the model of Section 3. In this model, each retailer sells an identical good and so, unlike in Section 3, the goods are not customized for each retailer/market. This introduces several changes in the model setup. First, since the product sold to each retailer is identical, the wholesale price for each retailer must be the same; we denote this price by w . And instead of tailoring a different discount δ_i for each retailer, the firm must offer the same discount δ to all retailers. Second, since the same product is sold through all retailers, the firm can procure and inventory the same product for all retailers. Essentially, instead of ordering P_i for each retailer, the firm places an overseas order of P units for all retailers. Third, we also relax the assumption of Sections 3 and 4 that the retailers cannot return or cancel the advance orders. We now allow *a return policy*: after demand D_i is realized, retailer i can either place additional orders $(D_i - Q_i)^+$ at w or cancel $(Q_i - D_i)^+$ from purchased quantity

and get back v for each canceled unit, where $0 < v < \delta w$ to ensure that the retailers would still be penalized for over-ordering. To focus on the benefits of Advance Purchase Discounts, we assume here that v is exogenously given such that the firm's expected profits are greater under a return policy than otherwise.¹⁰

As in previous sections, we first provide the retailers' and firm's strategy, followed by the equilibrium outcome and the benefits of the APD scheme.

Lemma. Retailers' Strategy Profile: *Each retailer orders an advance purchase quantity Q_i and earns expected profits $\mathbb{E}\pi_i^R$, where*

$$(7.1) \quad \begin{aligned} Q_i &= \mu'_i + \sigma'_i z_\delta, & \mathbb{E}\pi_i^R &= (R_i - (1 - \delta)w)\mu_i - (w - v)\phi(z_\delta)\sigma'_i - K_i & \text{if } \delta \geq \underline{\delta}_i, \\ Q_i &= 0, & \mathbb{E}\pi_i^R &= (R_i - w)\mu_i & \text{otherwise,} \end{aligned}$$

$$\underline{\delta}_i = \left\{ \delta \mid \delta w \mu_i - (w - v)\phi(z_\delta)\sigma'_i - K_i = 0 \right\},$$

$$\text{and } z_\delta = \Phi^{-1}\left(\frac{w\delta}{w-v}\right).$$

Regardless of the product characteristics, retailers are still independent decision makers and so their strategy profile remains the same as in Section 3. The only difference stems from our introduction of the return policy: because retailers can now cancel their orders and recover v for each canceled unit, essentially they share less risk than in the customized product case. This reduced risk share is captured by $(w - v)\phi(z_\delta)\sigma'_i$.

At the same time, several changes occur at the firm's end. First, owing to the generic nature of the product, the same wholesale price and discount are set for all retailers; moreover, the firm can procure and inventory the same products for all retailers, which has a risk pooling effect on demand uncertainty. Second, the introduction of a return policy further enlarges this risk pooling effect. In particular, after demand is realized, instead of re-ordering the shortfall $(D_i - P_i)^+$ for each retailer i , the firm now can redirect the returned units and so need only re-order $\left(\sum_{i=1}^N D_i - P\right)^+$ units. Finally, recall that the firm in our original model could offer a tailored discount δ_i to each retailer and so the decision to offer the APD scheme depends solely on the characteristics of the retailer in question. However, when the firm must use a single level of discount δ , this decision is influenced by the entire population of retailers. In particular, for a given discount δ , all retailers whose MADs (from Equation 7.1) are less than δ will participate in this APD scheme.

The choice of δ thus translates into a choice of selecting $m \in \{0, 1, 2, \dots, N\}$ such that the firm's expected profits are maximized when the retailers with the m lowest MADs participate in the scheme.¹¹ Unfortunately, for an unrestricted

¹⁰There is a large literature on supply chain coordination via return contracts (see Pasternack 1985, Donohue 2000, Taylor 2001). In those settings, the return policy is adopted to mitigate *retail* risk of *overstock*, thereby incentivizing retailers to order the supply chain optimal quantity. There is a certain range of return prices, closely related to the wholesale price, that achieve this goal. We remark that a consumer return policy plays an important role in resource allocation, and the optimal return price maximizes social welfare [Su, 2009]. In our model, however, the return policy is used to mitigate the *supply* risk of *understock*, and the supply chain is coordinated under any value of return price $v > 0$. This allows us to separate the decision of v and δ and to develop our results for an exogenously specified value of $v > 0$.

¹¹If $m = 0$ then $\delta = 0$; that is, the firm does not offer any discount.

set of retailer characteristics $(\mu_i, \sigma_i, t_i, K_i)$, the firm's expected profits are not monotonic in m ; hence there is no closed-form solution to the choice of m .¹² The algorithm to choose m is relatively simple, though. First, the firm sorts the retailers in the order of their MADs (Equation 7.1). Next, the firm compares its expected profits for different choices of m , where the expected profits are computed by using Equation (7.2) and assuming the retailers with the m lowest MADs will participate.

As Equation (7.2) illustrates, the firm's expected profits depend on each participating retailer's characteristics: the prior demand distribution $N(\mu_i, \sigma_i^2)$ faced by the retailer, her private information quality t_i , and her administrative cost K_i , each of which has a different effect on the firm's expected profits. This is why we do not provide, beyond the algorithm already proposed, a general conclusion or a structural property that drives the final choice. In the following analysis, we assume that the firm has already applied that algorithm and has chosen the $m > 0$ desirable retailers. Our next lemma formally captures the changes just described in the firm's strategy and explains the intuition behind the second part of the firm's strategy profile: the order quantity.

Lemma. *Firm's Strategy Profile:* *Suppose the firm offers the APD scheme to N retailers with MADs $\underline{\delta}_{(1)} \leq \underline{\delta}_{(2)} \leq \dots \leq \underline{\delta}_{(m)} \leq \dots \leq \underline{\delta}_{(N)}$, where $1 \leq m \leq N$ and, for all i , $\underline{\delta}_i$ is defined as Equation (7.1). Then the firm sets the discount $\delta = \underline{\delta}_{(m)}$, to make the first m retailers participate, and places an overseas order $P = \mu' + \sigma' z_\gamma$. The resulting expected profits become*

$$(7.2) \quad \begin{aligned} \mathbb{E}\pi^F &= \sum_{i=1}^m \left[((1-\delta)w - (1-\gamma))\mu_i + (w-v)\phi(z_\delta)\sigma'_i \right] \\ &+ \sum_{i=m+1}^N (w - (1-\gamma)C_L)\mu_i - C_L\phi(z_\gamma)\sigma', \end{aligned}$$

where $\mu' = \sum_{i=1}^m \mu'_i + \sum_{i=m+1}^N \mu_i$, $\sigma' = \sqrt{\sum_{i=1}^m \sigma_i'^2 + \sum_{i=m+1}^N \sigma_i^2}$, $z_\delta = \Phi^{-1}\left(\frac{w\delta}{w-v}\right)$, and $z_\gamma = \Phi^{-1}(\gamma)$.

Instead of ordering a unique P_i for each retailer, now the firm considers the demand from all N retailers as a whole and orders P units for all of them. The total demand consists of two parts: for participating retailers, the firm infers demand signals from their order quantities and updates its posterior distribution of total demand for m retailers to $N(\sum_{i=1}^m \mu'_i, \sum_{i=1}^m \sigma_i'^2)$; for nonparticipating retailers, the posterior distribution remains $N(\sum_{i=m+1}^N \mu_i, \sum_{i=m+1}^N \sigma_i^2)$. Therefore, the total demand faced by the firm has a posterior normal distribution with mean $\mu' = \sum_{i=1}^m \mu'_i + \sum_{i=m+1}^N \mu_i$, and standard deviation $\sigma' = \sqrt{\sum_{i=1}^m \sigma_i'^2 + \sum_{i=m+1}^N \sigma_i^2}$. The firm's optimal order quantity and expected profits follow directly from the same newsvendor problem under customized products as in Section 3 but with the new demand distribution.

In the case of generic products, the firm offers the same discount to all N retailers instead of setting a separate discount at each retailer's MAD; this offered discount is equal to the highest MAD among all m retailers. Thus, in equilibrium, retailers with the m lowest MADs would participate in APD and the rest would be better off staying out. The following theorem formally describes the equilibrium.

¹²It is relatively easy to derive closed-form results when all retailers are identical or by restricting retailer heterogeneity in other ways. Yet because we feel this has limited practical appeal, the most general model is reported here.

Theorem 7. Equilibrium Outcome

- (1) **Actions:** The tuple of actions $\{\delta^*, Q_i^*, P^*\}$ characterizes a perfect Bayesian equilibria of the setup just described. For all i ,

$$\delta^* = \underline{\delta}_{(m)}, \quad Q_i^* = \mathbf{1}_{i \leq m} \left(\mu'_i + \sigma'_i z_{\delta^*} \right), \quad P^* = \mu' + \sigma' z_\gamma.$$

- (2) **Profits:** The equilibrium expected profits for retailer i are $\mathbb{E}\Pi_i^R$, and the equilibrium expected profits for the firm $\mathbb{E}\Pi^F$ are $\sum_{i=1..N} \mathbb{E}\Pi_i^F$, where

$$\begin{aligned} \mathbb{E}\Pi_i^R &= \mathbf{1}_{i \leq m} \left[(R_i - (1 - \delta^*)w) \mu_i - (w - v) \phi(z_{\delta^*}) \sigma'_i - K_i \right] + \mathbf{1}_{m < i \leq N} \left[(R_i - w) \mu_i \right], \\ \mathbb{E}\Pi^F &= \sum_{i=1}^m \left[((1 - \delta)w - (1 - \gamma)) \mu_i + (w - v) \phi(z_\delta) \sigma'_i \right] + \sum_{i=m+1}^N (w - (1 - \gamma)C_L) \mu_i \\ &\quad - C_L \phi(z_\gamma) \sigma'. \end{aligned}$$

We next examine the supply chain performance under optimal policy.

Theorem 8. Benefits of Advance Purchase Discounts: With an APD scheme using δ^* as defined in Theorem 7, participating retailers, the firm, and the supply chain are all better off. The benefits are given as follows:

$$\begin{aligned} \Delta \mathbb{E}\Pi_i^R &= \mathbf{1}_{i \leq m} \left[\delta^* w \mu_i - (w - v) \phi(z_{\delta^*}) \sigma'_i - K_i \right] \geq 0, \\ \Delta \mathbb{E}\Pi^F &= \sum_{i=1}^m \left[-\delta^* w \mu_i + (w - v) \phi(z_{\delta^*}) \sigma'_i \right] + \bar{K} \geq 0, \\ \Delta \mathbb{E}\Pi^S &= \bar{K} - K \geq 0, \end{aligned}$$

where $\bar{K} = C_L \phi(z_\gamma) \left(\sum_{i=1}^N \sigma_i - \sigma' \right)$ and $K = \sum_{i=1}^m K_i$.

Although with generic products the firm adopts a different discount strategy, the description of the APD benefits in Theorem 8 is the equivalent of Theorem 1 for customized products. In fact, the intuition follows along exactly the same lines as before.

To participate in the APD scheme, each retailer incurs an additional administrative cost K_i (Effect Admin). By offering the APD to retailers, the average transfer price between the firm and each participating retailer is reduced by $\delta^* w \mu_i$, which leads to higher expected profits for each retailer but to lower ones for the firm (Effect Discount). At this cost, the firm invites each participating retailer into sharing demand uncertainty $(w - v) \phi(z_{\delta^*}) \sigma'_i$ (Effect Risk). Furthermore, under the APD scheme, superior demand information flows from the downstream tier (the retailers) to the upstream tier (the firm); hence the firm derives benefits \bar{K} from all participating retailers by using its information to make sourcing decisions that utilize long-lead time suppliers (Effect Information). Despite the internal effects (Effect Discount and Effect Risk), the supply chain benefits \bar{K} from superior information even though it comes with extra administrative costs K . As before, the combination of these effects drives the benefits for all agents in the supply chain. However, there is now one notable difference. With customized products, we recall, the firm is able to tailor

a different discount for each retailer and thus can appropriate all supply chain benefits. Yet with generic products, the firm offers all N retailers the same discount $\underline{\delta}_{(m)}$, which is the lowest MAD among m retailers. Therefore, while retailers with higher MADs would have no incentive to participate, retailers with lower MADs can obtain strictly positive benefits from participating in the APD scheme. This makes all the participating agents, except retailer m , strictly better off under Advance Purchase Discounts.

8. PROOFS OF RESULTS IN PRECEEDING SECTIONS

8.1. Proofs of Results in Section 3.

Proof of Lemma: Retailers' Strategy Profile

Proof. Given δ_i , retailer i 's profit if participating in the APD is $\pi_i^R = R_i D_i - (1 - \delta_i) w_i Q_i - w_i (D_i - Q_i)^+ - K_i$, which has a typical newsvendor solution with optimal order quantity $Q_i = \mu'_i + \sigma'_i z_{\delta_i}$ and corresponding expected profits $\mathbb{E}\pi_i^R = (R_i - (1 - \delta_i) w_i) \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i - K_i$, where $z_{\delta_i} = \Phi^{-1}(\delta_i)$.¹³ On the other hand, recall that retailer i 's expected profits when *not* participating in the APD are $\mathbb{E}\pi_{i0}^R = (R_i - w_i) \mu_i$. The two expected profits are equal when $\delta_i = \underline{\delta}_i$, where $\underline{\delta}_i = \left\{ \delta_i \mid \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i - K_i = 0 \right\}$. The difference between the two expected profits is $\mathbb{E}\pi_i^R - \mathbb{E}\pi_{i0}^R = \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma'_i - K_i = w_i \Phi(z_{\delta_i}) \left(\mu_i - \frac{\phi(z_{\delta_i})}{\Phi(z_{\delta_i})} \sigma'_i \right) - K_i$ (recall that $\delta_i = \Phi(z_{\delta_i})$). The standard normal distribution is log-concave and thus has decreasing reversed hazard rate (DRHR); that is, $\frac{\phi(z_{\delta_i})}{\Phi(z_{\delta_i})}$ is decreasing in z_{δ_i} . Therefore, if $\delta_i \geq \underline{\delta}_i$, then $\mathbb{E}\pi_i^R - \mathbb{E}\pi_{i0}^R$ is nonnegative and increasing in z_{δ_i} (and hence increasing in δ_i); if $\delta_i \leq \underline{\delta}_i$, then $\mathbb{E}\pi_i^R < \mathbb{E}\pi_{i0}^R$ and retailer i does not participate in the offered APD scheme. This concludes the proof. \square

Proof of Lemma: Firm's Strategy Profile

Proof. Given δ_i and retailer i 's participation, the firm's profit is $\pi_i^F = (1 - \delta_i) w_i Q_i + w_i (D_i - Q_i)^+ - (1 - \gamma) C_L P_i - C_L (D_i - P_i)^+$. This expression, which has a typical newsvendor solution with optimal order quantity $P_i = \mu'_i + \sigma'_i z_\gamma$ and corresponding expected profits $\mathbb{E}\pi_i^F = ((1 - \delta_i) w_i - (1 - \gamma) C_L) \mu_i + w_i \phi(z_{\delta_i}) \sigma'_i - C_L \phi(z_\gamma) \sigma'_i$ for retailer i 's corresponding order quantity $Q_i = \mu'_i + \sigma'_i z_{\delta_i}$, where $z_\gamma = \Phi^{-1}(\gamma)$ and $z_{\delta_i} = \Phi^{-1}(\delta_i)$. On the other hand, recall that the firm's expected profits when retailer i does not participate in the APD scheme are $\mathbb{E}\pi_{i0}^F = (w_i - (1 - \gamma) C_L) \mu_i - C_L \phi(z_\gamma) \sigma_i$. The difference between the two expected profits is $\mathbb{E}\pi_i^F - \mathbb{E}\pi_{i0}^F = -\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma'_i + C_L \phi(z_\gamma) (\sigma_i - \sigma'_i)$, which from the previous proof, is decreasing in δ_i when $\delta_i \geq \underline{\delta}_i$. To ensure retailer i 's participation, the firm's optimal discount is therefore $\delta_i = \underline{\delta}_i$, which makes $\mathbb{E}\pi_i^F - \mathbb{E}\pi_{i0}^F = -K_i + C_L \phi(z_\gamma) (\sigma_i - \sigma'_i)$. Thus, the firm would offer an Advance Purchase Discount to retailer i if and only if $-K_i + C_L \phi(z_\gamma) (\sigma_i - \sigma'_i) \geq 0$. This concludes the proof. \square

¹³Unless specifically indicated otherwise, all expectations of profits in this paper include expectation over both demand and available private signal; all expectations of benefits include expectation over demand, private signal, and administrative cost. Benefits alone refer to the difference between expected profits with and without APDs.

Proof of Theorem 1

Proof. The benefits that accrue to retailer i , the firm, and the supply chain under Advance Purchase Discounts stem directly from the difference between the expected profits with and without the APD scheme. For $\delta_i \geq \underline{\delta}_i$, the firm's expected profits are decreasing in δ_i ; therefore, the maximum discount it could offer is the $\bar{\delta}_i$ that yields the firm zero benefits from offering APD scheme—that is, when $\Delta \mathbb{E} \pi_i^F = 0$ for $\bar{\delta}_i = \left\{ \delta_i \mid \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' - \bar{K}_i = 0 \right\}$. Comparing the definition of $\underline{\delta}_i$ and $\bar{\delta}_i$: if $K_i < \bar{K}_i$, then $\underline{\delta}_i < \bar{\delta}_i$ and $K_i < \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' < \bar{K}_i$ for a $\delta_i \in (\underline{\delta}_i, \bar{\delta}_i)$. \square

Proof of Theorem 2

Proof. In equilibrium, when it is beneficial to employ the APD scheme (i.e., when $K_i \leq \bar{K}_i$), the firm offers a discount that simultaneously maximizes its expected profits and ensures retailer i 's participation (i.e., $\delta_i^* = \underline{\delta}_i$); in the second stage, each agent orders the newsvendor optimal quantity. In contrast, when it is not beneficial to employ the APD scheme (i.e., when $K_i > \bar{K}_i$), the firm offers no discount. This concludes the proof. \square

8.2. Proofs of Results in Section 4.

Proof of Lemma

Proof. The firm's expected benefits from offering the APD scheme with a discount δ_i are

$$\mathbb{E} \Delta \hat{\pi}_i^F = \begin{cases} 0, & \delta_i < \underline{\delta}_{i2}, \\ (1-b) \left(-\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma_i' + C_L \phi(z_\gamma) (\sigma_i - \sigma_i') \right), & \underline{\delta}_{i2} \leq \delta_i < \underline{\delta}_{i1}, \\ -\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma_i' + C_L \phi(z_\gamma) (\sigma_i - \sigma_i'), & \delta_i \geq \underline{\delta}_{i1}. \end{cases}$$

For $\underline{\delta}_{i2} \leq \delta_i < \underline{\delta}_{i1}$ and $\delta_i \geq \underline{\delta}_{i1}$, $\mathbb{E} \Delta \hat{\pi}_i^F$ is decreasing in δ_i ; hence the optimal discount reduces to either $\underline{\delta}_{i1}$ or $\underline{\delta}_{i2}$.

The lemma follows from comparing

$$(1-b) \left(-\underline{\delta}_{i2} w_i \mu_i + w_i \phi(z_{\underline{\delta}_{i2}}) \sigma_i' + C_L \phi(z_\gamma) (\sigma_i - \sigma_i') \right)$$

and

$$\left(-\underline{\delta}_{i1} w_i \mu_i + w_i \phi(z_{\underline{\delta}_{i1}}) \sigma_i' + C_L \phi(z_\gamma) (\sigma_i - \sigma_i') \right).$$

\square

Proof of Theorem 3

Proof. First we consider the optimal discount to set if the firm offers the APD scheme. (i) $B_i \leq \bar{K}_i$. Recall that $\bar{K}_i = C_L \phi(z_\gamma) (\sigma_i - \sigma_i')$. In response to an offer of discount δ_i , retailer i would participate in the APD if and only if

she would have nonnegative benefits under this discount (i.e., when $\delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' \geq K_i$). Taking into account retailer i 's participation, the firm's expected benefits are

$$\mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i) \right] = \left(-\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma_i' + \bar{K}_i \right) G_i \left(\delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' \right).$$

Define $k_i = \delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i'$, $\mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i(k_i)) \right] = (\bar{K}_i - k_i) G_i(k_i)$, and

$$\frac{\partial \mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i(k_i)) \right]}{\partial k_i} = -G_i(k_i) \left(1 - \frac{(\bar{K}_i - k_i) g_i(k_i)}{G_i(k_i)} \right),$$

where $k_i \in [0, \bar{K}_i]$. Since $g_i(\cdot)$ has a DRHR, it follows that $\frac{g_i(k_i)}{G_i(k_i)}$ decreases in k_i and that $\frac{\partial \mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i(k_i)) \right]}{\partial k_i}$ decreases in k_i when $1 - \frac{(\bar{K}_i - k_i) g_i(k_i)}{G_i(k_i)} \geq 0$. Let $k_{i0} = \left\{ k_i \mid \frac{(\bar{K}_i - k_i) g_i(k_i)}{G_i(k_i)} = 1 \right\}$. For $k_i \in [0, \bar{K}_i]$, $\frac{\partial \mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i(k_i)) \right]}{\partial k_i}$ is positive when $k_i \in [0, k_{i0}]$ but is decreasing and negative when $k_i \in [k_{i0}, \bar{K}_i]$. Therefore, if $A_i \leq k_{i0} \leq B_i$ then $k_i^* = k_{i0}$; if $k_{i0} < A_i$ then $k_i^* = A_i$; and if $k_{i0} > B_i$ then $k_i^* = B_i$. (ii) $B_i > \bar{K}_i$, the firm's expected benefits become¹⁴

$$\mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i) \right] = \left(-\delta_i w_i \mu_i + w_i \phi(z_{\delta_i}) \sigma_i' + \bar{K}_i \right) G_i \left(\delta_i w_i \mu_i - w_i \phi(z_{\delta_i}) \sigma_i' \right) G_i(\bar{K}_i),$$

which under the same transformation becomes $\mathbb{E}_{K_i} \left[\Delta \mathbb{E} \hat{\pi}_i^F(\delta_i(k_i)) \right] = (\bar{K}_i - k_i) G_i(k_i) G_i(\bar{K}_i)$. Therefore, if $A_i \leq k_{i0} \leq \bar{K}_i$ then $k_i^* = k_{i0}$; if $k_{i0} < A_i$ then $k_i^* = A_i$.

Next, we consider the firm's decision to offer the APD scheme or not. By offering the APD scheme with discount $\hat{\delta}_i(\hat{k}_i)$, the firm's expected benefits become $(\bar{K}_i - \hat{k}_i) G_i(\hat{k}_i)$; therefore, the firm would only offer the APD scheme when it gets nonnegative expected benefits, i.e., when $\hat{k}_i \leq \bar{K}_i$. This concludes the proof. \square

Proof for Theorem 4

Proof. Observe that $g_i(k_i | u_i, s_i) = \frac{1}{s_i} g_{i0} \left(\frac{k_i - u_i}{s_i} \right)$. By Equation (4.2),

$$(8.1) \quad \frac{\frac{1}{s_i} (\bar{K}_i - k_{i0}) g_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)}{G_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)} = 1.$$

When $s_i = s_{i0} = \frac{(\bar{K}_i - u_i) g_{i0}(0)}{G_{i0}(0)}$, we have $k_{i0} = u_i$. First we show that, for a given u_i , if $s_i < s_{i0}$ then $k_{i0} > u_i$ and therefore $\hat{\delta}_i(k_{i0}) > \hat{\delta}_i(u_i)$. Suppose to the contrary that $s_i < s_{i0}$ but $k_{i0} \leq u_i$. Then, by Equation (8.1), we must have $\frac{\frac{1}{s_i} (\bar{K}_i - k_{i0}) g_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)}{G_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)} = 1$. However, the DRHR of $g_{i0}(\cdot)$ yields that

$$\frac{\frac{1}{s_i} (\bar{K}_i - k_{i0}) g_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)}{G_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)} > \frac{\frac{1}{s_{i0}} (\bar{K}_i - u_i) g_{i0}(0)}{G_{i0}(0)} = 1,$$

a contradiction. Therefore, if $s_i < s_{i0}$ then $k_{i0} > u_i$. Similarly, suppose that $s_i > s_{i0}$ but $k_i > u_i$. Then

$$\frac{\frac{1}{s_i} (\bar{K}_i - k_{i0}) g_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)}{G_{i0} \left(\frac{k_{i0} - u_i}{s_i} \right)} < \frac{\frac{1}{s_{i0}} (\bar{K}_i - u_i) g_{i0}(0)}{G_{i0}(0)} = 1,$$

contradicting Equation (8.1). Therefore, if $s_i > s_{i0}$ then $k_{i0} < u_i$. This concludes the proof.

¹⁴If $K_i > \bar{K}_i$ then the APD is not offered and so the firm's benefits are zero.

□

Proof of Theorem 5

Proof. By Theorem 3, the firm would treat retailer i as if her administrative cost were \hat{k}_i . If $\hat{k}_i > \bar{K}_i$ and $\hat{k}_i \leq \bar{K}_i$, it would not be beneficial for the firm to offer the APD scheme so $\hat{\delta}_i^* = 0$. If $\hat{k}_i < K_i$, it would be beneficial for the firm to offer the APD scheme and so $\hat{\delta}_i^* = \hat{\delta}_i(\hat{k}_i)$; however, in this case it would not be beneficial for retailer i to participate. In both cases, the APD scheme cannot be implemented. In contrast, when $K_i \leq \hat{k}_i \leq \bar{K}_i$, the firm would offer the APD scheme with discount $\hat{\delta}_i(\hat{k}_i)$ and retailer i would participate. The optimal order quantities and expected profits thus follow.

□

Proof of Theorem 6

Proof. The benefits follow directly from the difference between the two expected profits in Theorem 5, bearing in mind that $\bar{K}_i = C_L \phi(z_\gamma) (\sigma_i - \sigma'_i)$.

□

Proof of Corollary

Proof. The corollary follows from the benefits listed in Theorem 6 and the statements in Theorem 4 that $\hat{k}_i(u_i, s_i) \geq u_i$ when $s_i \leq s_{i0}$ and that $\hat{k}_i(u_i, s_i) < u_i$ when $s_i > s_{i0}$.

□

8.3. Proofs of Results in Section 7.

Proof of Lemma: Retailers' Strategy Profile

Proof. Given δ , retailer i 's profit if participating in the APD is $\pi_i^R = RD_i - (1 - \delta)wQ_i - w(D_i - Q_i)^+ + v(Q_i - D_i)^+ - K_i$, which has a typical newsvendor solution with optimal order quantity $Q_i = \mu'_i + \sigma'_i z_\delta$ and corresponding expected profits $\mathbb{E}\pi_i^R = (R - (1 - \delta)w)\mu_i - (w - v)\phi(z_\delta)\sigma'_i - K_i$, where $z_\delta = \Phi^{-1}\left(\frac{w\delta}{w-v}\right)$. On the other hand, recall that retailer i 's expected profits if not participating in the APD are $\mathbb{E}\pi_{i0}^R = (R - w)\mu_i$. The two expected profits are equal when $\delta = \underline{\delta}_i$, where $\underline{\delta}_i = \left\{ \delta \mid \delta w \mu_i - (w - v)\phi(z_\delta)\sigma'_i - K_i = 0 \right\}$. The difference between the two expected profits is $\mathbb{E}\pi_i^R - \mathbb{E}\pi_{i0}^R = \delta w \mu_i - (w - v)\phi(z_\delta)\sigma'_i - K_i = (w - v)\Phi(z_\delta)\left(\mu_i - \frac{\phi(z_\delta)}{\Phi(z_\delta)}\sigma'_i\right) - K_i$; recall that $\delta = \frac{(w-v)\Phi(z_\delta)}{w}$. The standard normal distribution is log-concave and thus has DRHR; in other words, $\frac{\phi(z_\delta)}{\Phi(z_\delta)}$ is decreasing in z_δ . Therefore, if $\delta \geq \underline{\delta}_i$ then $\mathbb{E}\pi_i^R - \mathbb{E}\pi_{i0}^R$ is nonnegative and increasing in z_δ (and hence increasing in δ); if $\delta \leq \underline{\delta}_i$ then $\mathbb{E}\pi_i^R < \mathbb{E}\pi_{i0}^R$ and retailer i does not avail herself of the offered Advance Purchase Discount. This concludes the proof.

□

Proof of Lemma: Firm's Strategy Profile

Proof. Given δ and the participation of m retailers, the firm's profit is

$$\begin{aligned}\pi^F &= (1-\delta)w \sum_{i=1}^m Q_i + w \sum_{i=1}^m (D_i - Q_i)^+ - v \sum_{i=1}^m (Q_i - D_i)^+ + w \sum_{i=m+1}^N D_i \\ &\quad - (1-\gamma)C_L P - C_L \left(\sum_{i=1}^m (D_i - Q_i)^+ + \sum_{i=1}^m Q_i + \sum_{i=m+1}^N D_i - \sum_{i=1}^m (Q_i - D_i)^+ - P \right)^+ \\ &= -\delta w \sum_{i=1}^m Q_i + (w-v) \sum_{i=1}^m (Q_i - D_i)^+ + w \sum_{i=1}^N D_i - (1-\gamma)C_L P - C_L \left(\sum_{i=1}^N D_i - P \right)^+.\end{aligned}$$

This expression has a typical newsvendor solution with optimal order quantity $P = \mu' + \sigma' z_\gamma$ and corresponding expected profits

$$\mathbb{E}\pi^F = \sum_{i=1}^m \left[((1-\delta)w - (1-\gamma))\mu_i + (w-v)\phi(z_\delta)\sigma'_i \right] + \sum_{i=m+1}^N (w - (1-\gamma)C_L)\mu_i - C_L\phi(z_\gamma)\sigma',$$

given each participating retailer i 's corresponding order quantity $Q_i = \mu'_i + \sigma'_i z_\delta$ for $\mu' = \sum_{i=1}^m \mu'_i + \sum_{i=m+1}^N \mu_i$, $\sigma' = \sqrt{\sum_{i=1}^m \sigma_i'^2 + \sum_{i=m+1}^N \sigma_i^2}$, $z_\delta = \Phi^{-1}\left(\frac{w\delta}{w-v}\right)$, and $z_\gamma = \Phi^{-1}(\gamma)$. Since $\left((1-\delta)w + (w-v)\phi(z_\delta)\sigma'_i\right)$ is decreasing in δ when $\delta \geq \underline{\delta}_i$, it follows that $\mathbb{E}\pi^F$ is decreasing in δ when $\delta \geq \underline{\delta}_{(m)}$. To maximize the firm's expected profits $\mathbb{E}\pi^F$ and to ensure the first m retailers' participation, the optimal discount for the firm becomes $\underline{\delta}_{(m)}$. This concludes the proof. □

Proof of Theorem 7

Proof. In equilibrium, when it is beneficial to offer the APD scheme to the first m retailers, the firm offers a discount that simultaneously maximizes its expected profits and ensures the participation of those retailers (i.e., $\delta^* = \underline{\delta}_{(m)}$). In the second stage, each agent orders the newsvendor optimal quantity, and the corresponding expected profits follow. □

Proof of Theorem 8

Proof. Recall that, when no APD is implemented, the respective expected profits of retailer i , the firm, and the supply chain are $\mathbb{E}\pi_{i0}^R = (R-w)\mu_i$, $\mathbb{E}\pi_0^F = \sum_{i=1}^N ((w - (1-\gamma)C_L)\mu_i - C_L\phi(z_\gamma)\sigma_i)$, and $\mathbb{E}\pi_0^S = \sum_{i=1}^N ((R - (1-\gamma)C_L)\mu_i - C_L\phi(z_\gamma)\sigma_i)$. The benefits of APDs are a direct consequence of the difference between the expected profits in Theorem 7 and the corresponding expected profits just described. □