Sales Information Transparency and Trust in Repeated Vertical Relationships *

By Noam Shamir†and Yaron Yehezkel‡

January 2023

Abstract

Problem definition: We study a repeated interaction between a manufacturer and a retailer, where the retailer may share with the manufacturer past sales information. In our model, such information cannot improve the latter's predictive capabilities of future demand, but it does allow him to infer past demand. Our main research question is under what conditions the retailer and the manufacturer benefit from sharing such past sales information, and how dynamic interaction and past sales information affect the efficiency of the distribution channel. ***Methodology:*** We model a repeated relationship between a manufacturer and a retailer, where demand fluctuates in an i.i.d manner between periods. In each period, the retailer privately observes the current demand, and the manufacturer offers a menu of contracts to elicit the retailer to reveal its private information. The manufacturer may observe sales information that reveals past demand at the end of each period, if the retailer chooses to share such information. ***Results:*** We find that even without sharing sales information, repeated interaction by itself enhances efficiency and profits for both firms. Past sales information further improves the channels' efficiency and increases the manufacturer's expected profit. Yet, past sales information increases (decreases) the retailer's per-period expected profit when the retailer places a low (high) value on its future profits. ***Managerial implications:*** Our results provide a new strategic reasoning for sharing past-sales information – as a way to increase trust in repeated vertical relationships. Furthermore, when the retailer can share a noisy signal regarding past demand, this may facilitate the exchange of sales information. We also consider the case of a financially constrained retailer, and demonstrate that financial constraints may benefit the retailer as they limit the market power of the manufacturer. In contrast, the manufacturer and the channel's efficiency are always worse-off when the retailer is financially constrained.

*The authors thank the Department Editor, Wedad Elmaghraby, the anonymous associate editor, and two anonymous referees for their detailed review of this work and for suggesting key ideas that led to major improvements in the paper. For helpful comments and discussions we thank Chiara Fumagalli, David Gilo, Kobi Glazer, Sarit Markovich, Yossi Spiegel, Lior Trachterman, Nikos Vettas and seminar participants at the CRESSE 2018 and the EAIRE 2019. The authors gratefully acknowledge financial support from the Coller Foundation, the Eli Hurvitz Institute and the Henry Crown Institute.

†Coller School of Management, Tel Aviv University (email: nshamir@tauex.tau.ac.il)
‡Coller School of Management, Tel Aviv University (email: yehezkel@tauex.tau.ac.il)

1 Introduction

The vision of Supply-chain 4.0, supported by the rapid advancement in IT systems and data science, calls for better collaboration between firms that comprise a distribution channel (Alicke et al. 2017). One of the key ways to achieve this collaboration is exchange of past sales information (Ferrantino and Koten 2019). The general consensus is that sales information improves the efficiency of the distribution channel by lowering inventory related costs, the cost of lost sales and improving the firms' predictive capabilities (e.g., Grean and Shaw 2002; Chen 2003 and Ha and Tang 2017 provide an extensive survey of the research about information sharing). Indeed, in a recent survey conducted by Coresight Research (2020), 85% of the retailers and 92% of the suppliers reported an improvement in their data sharing plans; however, this survey also reported that data sharing is one of the main tension points between firms due to reluctance to provide closely held sales information, because of fear that such information can be used against the firm that provides it. When such collaborative efforts do materialize, they involve heavy monetary investments, and, thus, represent a commitment on behalf of the firms in the distribution channel to develop long-term relationships (Kalwani and Narayandas 1995, Kumar 1996, Prajogoand Olhager 2012).

In this work we set to better understand the effects of repeated interaction and the ability to exchange past sales information on the trust built between the parties in the distribution channel and subsequently the performance of the channel. We focus on the way past sales information affects the strategic contract design in a repeated interaction and an environment characterized by information asymmetry. To achieve this goal, we consider a retailer that sources a product from a manufacturer over an infinite horizon, and during each period demand fluctuates in an i.i.d manner between two possible states. Being closer to the market, the retailer is able to observe at the beginning of each selling season the actual demand state while the manufacturer cannot observe such information ex-ante. While sharing past sales does not affect the predictive capabilities of the manufacturer it does affect the level of trust formed between the retailer and the manufacturer and, thus, have a strategic effect on the contract design.

In order to understand the effect of sales information exchange in repeated interaction and the retailer's incentive to provide such information, we examine several scenarios. In a benchmark case, the firms interact over one selling season; the manufacturer designs a menu of contracts that leaves the retailer with positive information yet exhibits the known quantity distortion which reduces the joint profit of the distribution channel. We then move to an infinitely repeated game. To study the retailer's incentive to share sales information, we first study a setting in which the firms interact repeatedly, but no sales information is exchanged. The introduction of repeated interaction improves the expected payoffs of both parties in the distribution channel. Repeated interaction allows the manufacturer to increase the offered quantity during periods of low demand, and, thus, to improve the overall efficiency of the distribution channel. When increasing the offered quantity during low demand, the manufacturer charges a higher payment than the retailer's revenues, thus leaving the retailer with a negative payoff during such periods. However, during periods of high demand, the manufacturer more than offsets the retailer's negative payoff during the low demand states, such that the overall expected per-period payoff of the retailer is higher compared with the static settings. Therefore, repeated interaction, even without any information exchange, can improve the expected payoff of both parties in the distribution channel.

We then proceed to examine the retailer's incentive to offer past sales information to the manufacturer. Recall that in our model, demand fluctuates in an independent and identical manner between periods, such that learning past demand does not improve the predictive capabilities of the manufacturer regarding future demand. This modeling choice allows us to mute the effect of information as a way to improve forecasting capabilities and focus on the strategic role of sharing past sales information as a way to affect the level of trust formed between the parties. When the manufacturer observes past demand the retailer is perceived as more trustworthy (Ozer et al. 2011), because the manufacturer can "punish" him for choosing a contract not according to the actual market state – an outcome that cannot happen if the retailer does not provide the manufacturer with this information. Having the ability to observe past demand allows the manufacturer, for a given set of offered quantities, to decrease the information rents paid to the retailer – an outcome that hurts the retailer. However, in spite of this effect, we show that there are conditions in which the retailer prefers to share past sales information with the manufacturer, due to an additional effect: the exchange of sales information allows the manufacturer to further increase the offered quantity during periods of low-demand, thus mitigating the problem of quantity distortion. This effect improves the overall performance of the distribution channel, and this improvement may also propagate to the retailer. Specifically, we show that when the discount factor of the retailer is not too high, the latter effect outweighs the former, and the retailer finds it beneficial to exchange past sales information with the manufacturer.

Our findings that sharing past sales information can improve trust in repeated interactions is supported by some previous empirical and anecdotal evidence. For example, Kwon and Suh (2004) conduce a large scale survey study to understand the factors that affect trust in supply-chains; they conclude that information sharing contributes to trust formation. In an interview to Harvard Business Review, Michael Dell described the vision of achieving "virtual integration" between Dell and its suppliers (Magretta 1998). One of the ways to achieve such virtual integration is the establishment of a web portal for supplier collaboration that includes (among other information) detailed data regarding product demand (Dyer and Hatch 2004). Dyer and Ouchi (1993) show that although American automakers are more vertically integrated, it is the Japanese automakers that enjoy lower production costs, and they attribute this success to the formation of long-term relationships that alleviate the problem of information asymmetry. We contribute to the understanding of these examples by proposing a new mechanism for why sharing past sales information can not only enhance trust but can also be beneficial to both the manufacturer and the retailer.

After exploring the motivation to exchange past sales information, we extend the basic model in two ways. In the main model, we assume that past sales information allows the manufacturer to perfectly infer the state of demand during the previous period. We relax this assumption, and study a setting in which past sales information provides an imperfect signal for the manufacturer regarding past demand. We study the way the precision level affects the payoffs of the two firms, and demonstrate that there are cases in which the retailer benefits from information ambiguity – i.e., from sharing imperfect information regarding the previous period demand state. We argue that ambiguity may facilitate information exchange, since there are cases in which the retailer will refrain from exchanging information when information fully reveals the past demand, but will choose to do so when it imperfectly reveals this information. The second extension we consider is when the retailer has limited liability; in this case, the retailer's payoff cannot be negative during any period. In the main model, absent limited liability, the optimal contract designed by the manufacturer entails charging the retailer a high transfer price during periods of low demand such that the retailer earns negative payoff. We show that limited liability can sometimes improve the situation for the retailer since it reduces the bargaining power of the manufacturer in the distribution channel. In contrast, limited liability always hurts the manufacturer.

To summarize, the main contributions of this work are as follows: 1) Repeated interaction increases the set of possible contracts and benefits both the retailer and the manufacturer; 2) Sharing past sales information increases the trust level; it benefits the retailer when his discount factor is not too high, and it always makes the manufacturer and the distribution channel better-off; 3) Partial information sharing can be beneficial to the retailer when the discount factor is intermediate. Furthermore, partial information sharing may facilitate information exchange; 4) Financial constraints can benefit the retailer because it limits the market power of the manufacturer.

The remainder of the paper is organized as follows. Section 2 provides a brief literature review. Section 3 describes the model and the static benchmark. Section 4 solves for the dynamic mechanism and compares between sharing of past sales information and no information sharing. In Section 5, we extend the model to study the cases of partial information sharing and limited liability. We discuss the managerial implications of our findings in Section 6. Finally, we conclude in Section 7. In Appendix A, we study the way some of our model assumptions affect the main insights.

2 Related Literature

This research is relevant to the following research strands: the value and incentives for information sharing in a distribution channel, the effect of repeated interaction on relational contracts and trust formation between a retailer and a manufacturer.

The value of information sharing between different parties in a distribution channel has been well documented (e.g., Lee and Whang 2000) and received considerable attention by operations management (OM) scholars. Chen (2003) and Ha and Tang (2017) provide an extensive review of the effect of information exchange on the performance of a distribution channel. In spite of the general consensus that information improves the efficiency of the distribution channel, when the channel is comprised from independent profit-maximizing firms, achieving an information sharing equilibrium means that each firm must find sharing information to be preferable over the concealment of information. Li (2002), Zhang (2002) and Li and Zhang (2008) are a few examples for papers that study the ex-ante incentives to share information in a distribution channel under the assumption that if information is shared, it is shared in a truthful manner. We contribute to this line of research but also differ from it across a few important dimensions. First, we assume that the firms in our model interact repeatedly as opposed to the static models outlined above; second, we assume that the retailer can choose strategically a contract not according to the state of demand, if he finds it beneficial; finally, our work is centered around the value of sharing past sales information – information that does not affect the forecasting capabilities or the future capacity decisions of the manufacturer.

Other scholars have also studied the incentives for credible exchange of information, and have shown that information sharing may result in information manipulation. For example, Cohen et al. (2003) and Oh and Ozer (2013) provide several examples of overly optimistic forecasts provided to an upstream firm in a distribution channel in several industries. In order to solve the problem of information manipulation two types of solutions have emerged – designing a menu of contracts (known as screening games) or choosing a signaling action. In the former case, the uninformed firm designs a menu of contracts that elicits the informed party to reveal the hidden information upon the choice of a specific contract out of this menu. Some examples for this kind of information revelation mechanism in the context of a relationship between suppliers and retailers include Gal-Or (1991a) and (1991b), Martimort (1996), Cachon and Lariviere (1999), Cachon and Lariviere (2001), Ha (2001), Corbett (2001) and Ozer and Wei (2006), Yehezkel (2008), Acconcia, Martina and Piccolo (2008) and Yehezkel (2014). In the second alternative (a signaling game), the informed party takes an action that conveys information to the uninformed party (for example, Gal-Or et al. (2008), Dukes et al. (2017) and Jiang et al. (2016)). In this paper, we adopt the idea of designing a screening contract, but we differ from the above mentioned line of research by assuming that firms in the distribution channel interact repeatedly and that the shared information is related to the past sales and not future forecast.

As opposed to the extensive research that studies mechanisms for information sharing in a static setting, relatively little attention has been devoted to examining the effect of repeated interaction on this issue. Plambeck and Taylor (2007) and Taylor and Plambeck (2007) illustrate in a series of papers that the potential for renegotiation of supply contracts has important implications for the firms' investments in innovation and capacity, the allocation of capacity, and the resulting profits. Zhang and Zenios (2008) study a dynamic contractual relationship between a principal and an agent in which the principal has a primary stake in the performance of a system, but delegates its control to an agent. Ren et al. (2010) study a problem similar to ours: the ability to share demand information in a distribution channel comprised of a manufacturer and his supplier; Ren et al. (2010) emphasize the ability to share information via cheap-talk (Crawford and Sobel 1982) when the supply chain relationship spans an infinite horizon, and they demonstrate that this outcome can be achieved by adopting a sophisticated review policy. Additional papers that study long-term relationships include Lobel and Xiao (2017) and Zhang et al. (2010). Lobel and Xiao (2017) characterize the optimal long-term contract in a setting with hidden demand information. Zhang et al. (2010) study a model in which a single supplier sells to a downstream retailer. The retailer does not share his inventory level with the supplier, and they investigate the sequence of short-term contracts offered by the supplier. Similar to Zhang et al., we focus on a scenario in which a manufacturer offers a sequence of short-term contracts, as opposed to Lobel and Xiao (2017) that characterize a long-term contract. Calzolari and Spagnolo (2017) and Martimort et al. (2017) consider an infinitely repeated game when retailers have private information and the supplier designs a menu of contracts that elicits information revelation. We contribute to these papers by considering the case where the retailer can reveal the past demand by sharing past sales information.

While in the main model, we assume that when information is shared, it completely reveals past demand, we also investigate the effect of sharing imperfect sales information. We demonstrate that imperfect information may actually facilitate the exchange of information. This finding is also supported by Li et al. (2014) and Amornpetchkul et al. (2015) who studied the issue of imperfect information exchange in a static setting.

One of our main findings is that information transparency (i.e., sharing past sales information between members in the distribution channel) builds trust and improves channel efficiency. This finding contributes to the recent literature that studies the issue of trust formation in similar relationships. Özer et al. (2011)

introduced to the newsvendor setting the notion of trust and trustworthiness, and showed that the incentive to distort information weakens in the presence of trust. Özer et al. (2014) study the way cultural differences affect trust formation and subsequently information sharing; Brinkhoff et al. (2015) evaluated the effect of trust on successful project completion; Choi et al. (2020) study the role of trust and trustworthiness in impacting high-ranking executives' decisions in supply chain interactions. A survey of the topic of trust in a supply chain context can be found in Özer and Zheng (2017) and Özer and Zheng (2018). Within a firm-customers relationship, a series of papers focused on the way operational transparency, and specifically information sharing, can improve the way customers value the product or the service (e.g., Buell and Norton (2011), Buell et al. (2017), Mohan et al. (2020), Buell and Kalkanci (2021)).

3 The Model

We study a distribution channel comprised from a manufacturer selling to a retailer. The two firms interact repeatedly over an infinite number of periods. We first describe of the events in a single period and we then embed the single-period model into an infinite repeated game.

3.1 The one-period model

During every selling period, market demand can be either high (H) or low (L), with probabilities p and $1-p$, respectively, and demand realizations are i.i.d across periods. The retailer's revenue from selling a quantity q of the manufacturer's product is $\pi_R(q;\theta)$, where $\theta \in \{L, H\}$ captures the state of demand. Suppose that $\pi_R(q;H) > \pi_R(q;L)$ for any q, hence for a given sold quantity q, the retailer's revenue is higher during the high market condition than during the low market condition. We further assume that the gap $\Delta(q) \equiv \pi_R(q;H) - \pi_R(q;L)$ is (weakly) increasing in q for all q, and that $\pi_R(q;\theta)$ is concave in q and have a unique maximization at the quantities q_H^* and q_L^* for the high and low states, respectively.

To better illustrate our main results, it would be useful to think of a specific functional form (all of our results carry over to the general specification described above). Suppose that the inverse demand function during state θ is captured using the linear form $P(q;\theta) = V_\theta - q$, where $V_H > V_L$. For the linear demand case, $\pi_R(q;\theta) = P(q;\theta)q = (V_\theta - q)q$, hence $q_\theta^* = V_\theta/2,\ \theta \in \{L, H\}$.

Due to the proximity of the retailer to the consumer market, we assume that the retailer, prior to the actual selling season, can observe the realization of θ.[1] We refer to a retailer observing the high market demand as a high-type retailer, and to a retailer observing the low market demand as a low-type retailer. The per-period profit of the retailer depends on the contract terms between the retailer and

[1]Many papers that study the topic of information asymmetry in a vertical distribution-channel also adopt this assumption regarding the retailer's better forecasting capabilities compared with those of the manufacturer (e.g., Cachon and Lariviere 2001, Li 2002, Ozer and Wei 2006, Li and Zhang 2008).

the manufacturer. We further follow the traditional literature that studies vertical relationships under information asymmetry and assume that the manufacturer offers a screening contract to the retailer. A screening contract is a menu of contracts that upon the retailer's choice of a specific contract out of this well-crafted menu reveals to the manufacturer the state of the demand. For the case of two possible demand states, and without loss of generality, the screening menu that we study takes the form of $\{(q_H, T_H), (q_L, T_L)\}$. For the quantity of q_H the retailer pays the sum of T_H, and for the quantity of q_L, the retailer pays the sum of T_L. Consequently, the per-period profit of the retailer during demand state $\theta \in \{L, H\}$ and choosing a contract $(q_{\tilde{\theta}}, T_{\tilde{\theta}})$ is $\pi_R(q_{\tilde{\theta}}; \theta) - T_{\tilde{\theta}}$. In equilibrium, based on the revelation principle, the manufacturer designs a screening menu such that a retailer of type $\theta \in \{L, H\}$ chooses the contract (q_θ, T_θ). We further denote the ex-ante per-period profit of the retailer when choosing the contract that corresponds to the true state by $\Pi_R = p\,(\pi_R(q_H; H) - T_H) + (1-p)\,(\pi_R(q_L; L) - T_L)$.

In our model, the manufacturer does not face capacity constraints, and can produce any quantity. We further assume that the manufacturer incurs a linear production cost; for exposition clarity, we normalize this production cost to zero. Therefore, the per-period profit of the manufacturer is the transfer price of T_i, and the ex-ante per-period profit of the manufacturer when the retailer chooses, based on the revelation principle, the contract that corresponds to the true state is $\Pi_M = pT_H + (1-p)T_L$. Finally, the per-period ex-ante value of the distribution-channel when the contract is chosen according to the true state of the demand is $\Pi_M + \Pi_R = p\pi_R(q_H; H) + (1-p)\,\pi_R(q_L; L)$. Note that the quantities that maximize the overall performance of the distribution channel are also q_H^* and q_L^*.

To be explicit, the timing and information structure of each period are as follows. At the beginning of each period, the retailer privately observes whether the demand is H or L in the current period. The manufacturer offers a take-it-or-leave-it menu of contracts $\{(q_H, T_H), (q_L, T_L)\}$ from which the retailer chooses a specific contract. If the retailer rejects both contracts, no trade takes place in the current period. Otherwise, the retailer chooses a contract from the menu.

3.2 The infinitely repeated game

We embed the single period model described above into an infinitely repeated game. The market demand in each period is independent of the market demand during previous periods. We adopt this assumption in order to highlight the strategic effect of information sharing regarding past demand when this information does not provide any knowledge about future demand. While previous research emphasized information sharing as a tool to achieve better forecasting capabilities (Li 2002, Li and Zhang 2008, Gal-Or et al. 2008, Ha et al. 2011, Shamir and Shin 2018), in our model, learning about past-periods demand does not have any value for predicting future demand. This assumption enables us to focus on how such information affects the relational contracts and trust between the retailer and the manufacturer. The total payoff of each firm is the sum of the per-period payoffs discounted using a common discount factor of δ ($0 \leq \delta \leq 1$), such that the discounted profit of the retailer is given by $\sum_{t=1}^{\infty}\left[\delta^{t-1}\Pi_R\right]$, and similarly the discounted profit of the manufacturer is given by $\sum_{t=1}^{\infty}\left[\delta^{t-1}\Pi_M\right]$.

3.3 Benchmark - A one-period model

As a benchmark, consider the static equilibrium, where firms interact over one selling period. This benchmark is the simple principal-agent problem (see, for example, Laffont and Martimort (2002)). We use the superscript "S" to denote the static setting. In this problem, the manufacturer designs a menu $\{(q_H, T_H), (q_L, T_L)\}$ that maximizes its payoff while satisfying the following set of constraints:

$$IR_\theta^S: \qquad \pi_R(q_\theta;\theta) - T_\theta \geq 0, \quad \text{for } \theta \in \{L, H\}, \tag{1}$$

$$IC_\theta^S: \quad \pi_R(q_\theta;\theta) - T_\theta \geq \pi_R(q_{\widetilde{\theta}};\theta) - T_{\widetilde{\theta}}, \quad \text{where } \theta, \widetilde{\theta} \in \{L, H\} \text{ and } \theta \neq \widetilde{\theta}. \tag{2}$$

These constraints ensure that the retailer accepts the contract (q_θ, T_θ) in state θ (IR_θ^S) over an outside option (normalized to zero), and prefers the contract (q_θ, T_θ) over the contract $(q_{\widetilde{\theta}}, T_{\widetilde{\theta}})$ in state θ (IC_θ^S). As is well-known, in this problem the binding constraints are IR_L^S and IC_H^S. Let T_H^S and T_L^S denote the solutions to IR_L^S and IC_H^S in equality. The manufacturer's expected profit, $pT_H^S + (1-p)T_L^S$, is:

$$\Pi_M^S(q_H, q_L) = p\pi_R(q_H; H) + (1-p)\pi_R(q_L; L) - p\Delta(q_L), \tag{3}$$

where recall that $\Delta(q) \equiv \pi_R(q; H) - \pi_R(q; L) > 0$. The first two terms are the expected profits of the distribution channel. The third term is the "information rents" that the manufacturer needs to leave to the retailer in state H, for motivating the retailer to reveal its type. Let q_H^S and q_L^S denote the quantities offered by the manufacturer in the static case. Given the linear demand function, the optimal contract offered by the manufacturer is[2]

$$q_H^S = q_H^* = \frac{V_H}{2}, \quad q_L^S = \frac{V_L}{2} - \frac{p}{(1-p)}\frac{(V_H - V_L)}{2} < q_L^*. \tag{4}$$

$$T_H^S = \frac{V_H^2}{4} - \frac{(V_L - pV_H)(V_H - V_L)}{2(1-p)}, \quad T_L^S = \frac{V_L^2}{4} - \frac{p^2(V_H - V_L)^2}{4(1-p)^2}. \tag{5}$$

In this case, $q_H^S = q_H^*$ and $q_L^S < q_L^*$. The manufacturer decreases q_L below its optimal level in order to

[2]This solution is based on the the assumption that the manufacturer wishes to sell during both demand states, such that $q_L^S > 0$. This means that $pV_H < V_L$. Intuitively, if the probability of state H, p, or if the ratio between the demand in the two states, V_H/V_L, are too high, the manufacturer prefers to forgo sales during state L, and to collect the entire profits from the retailer in state H.

reduce the high-type retailer's incentive to mimic the low-type; it consequently enables the manufacturer to reduce the high-type retailer's information rent. This contract allows the manufacturer to extract the entire surplus from the low-type retailer and leave the high-type retailer with positive information rents. The profit of the low-type retailer is 0 and the profit of the high-type retailer is $\Delta(q_L^S)$. In expectation, the retailer earns $\Pi_R^S = p\Delta(q_L^S)$.

4 The Dynamic Menu

In this section we solve the model in which the retailer and the manufacturer interact repeatedly, and information about the previous period's demand can be shared. The dynamic model introduces two main effects: repeated interaction and the ability to share past sales information. To distinguish between these two effects we first (in subsection 4.1) analyze the case when no information is shared, and then (in subsection 4.2) we add the possibility to share past sales information.

Notice first that the static equilibrium in Section 3.3 is an equilibrium in the dynamic case as well, under both sharing and no-sharing scenarios. It is supported by the firms' beliefs that the manufacturer offers in every period the static contract. Yet, the presence of dynamics and potentially sales information support the possibility of other equilibria as well.[3]

In every period, the manufacturer offers a *dynamic* menu; on the equilibrium path, the retailer accepts a specific contract out of the menu and the manufacturer continues to offer the same menu in all future periods. The equilibrium can break down in two ways. First, if the retailer rejects any contract offered in the menu such that no trade takes place during this period. The second option concerns the case when past sales information is shared. If the manufacturer infers that the retailer has chosen a contract not according to the true state of the demand, he will stop offering the dynamic menu in all subsequent periods. We assume that each of the above deviations triggers the static equilibrium described in Section 3.3, which is a subgame perfect equilibrium in the dynamic game as well. This trigger strategy provides the retailer with the lowest possible profit that constitutes a subgame perfect equilibrium, and is aligned with the market power that the manufacturer has in our model (we later, in Section **??** discuss the implications of this assumption on our main findings).[4]

The mechanism described above is characterized by the following set of constraints. The first set of constraints, IR_θ^D (which is the same under sharing and no information sharing), is the retailer's partic-

[3]We focus on stationary mechanisms because it is a repeated game with i.i.d demand observations.

[4]In Appendix D, we consider an alternative mechanism in which the manufacturer pays the retailer at the beginning of each period a fixed fee given that the retailer reported truthfully at the previous period. Then, the manufacturer offers a menu that specifies the centralized quantities and fully collects the retailer's profits from sales. We show that such a mechanism is inferior from the view point of the manufacturer to the mechanism proposed in our base model, and may result in negative profits for the manufacturer.

ipation constraint in state $\theta \in \{L, H\}$. These constraints ensure that the retailer prefers accepting the contract (q_θ, T_θ) in state θ, given that doing so maintains the equilibrium, over rejecting the contract and receiving the static menu in all future periods. We use the superscript "D" to denote the dynamic setting. Formally:

$$IR_\theta^D: \quad \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\big[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\big] \geq \frac{\delta}{1-\delta} p\Delta(q_L^S). \tag{6}$$

The left-hand-side (LHS) denotes the discounted payoff of the retailer when accepting the dynamic contract. In the current period, the retailer earns the equilibrium profit in state θ, i.e., $\pi_R(q_\theta, \theta) - T_\theta$, followed by the discounted sum of expected profits from the dynamic contract. In the right-hand-side (RHS) the retailer rejects the offered contract during the current period (and hence earns the outside option that is normalized to zero), followed by the discounted sum of the expected profits from the static contract (as defined in Section 3.3).

The second set of constraints, IC_θ^D, is the retailer's incentive compatibility constraints in state $\theta \in \{L, H\}$. These constraints ensure that the retailer prefers accepting the contract (q_θ, T_θ) in state θ given that doing so maintains the equilibrium, over accepting the contract $(q_{\widetilde{\theta}}, T_{\widetilde{\theta}})$, where $\widetilde{\theta} \neq \theta$. When past sales information is not exchanged between periods, the static IC_θ^S, ensures that the retailer does not make such a deviation, because the manufacturer cannot detect it at the end of the period. Yet, when sales information is exchanged, the manufacturer detects the deviation at the end of the period and then offers the static menu in all future periods. Hence:

$$IC_\theta^D: \; \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\big[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\big] \geq \tag{7}$$

$$\pi_R(q_{\widetilde{\theta}};\theta) - T_{\widetilde{\theta}} + \frac{\delta p}{1-\delta}\Delta(q_L^S).$$

The LHS captures the payoff of the retailer when choosing the contract according to the true state of the demand. This payoff is comprised from the current period payoff and the discounted sum of expected profits given that the manufacturer will continue to offer the dynamic menu. The RHS denotes the payoff when the retailer chooses a contract not according to the true demand state, given that sales information is shared. In this case, the retailer first chooses the contract $(q_{\widetilde{\theta}}, T_{\widetilde{\theta}})$ although the demand is $\theta \neq \widetilde{\theta}$. Based on the exchanged information, the manufacturer infers this deviation, and offers during all subsequent periods the static menu. Note that if sales information is not exchanged, the manufacturer cannot detect such deviation. In this case, the IC_θ^D constraints are identical to the constraints in the static model. Any menu satisfying conditions (6) - (7) can be a dynamic equilibrium. We focus on the menu that maximizes the manufacturer's expected profit. The rationale for doing so is that the manufacturer has the market power to make a take-it-or-leave offer to the retailer.

4.1 Repeated interaction without past sales information sharing

To understand the role of past-sales information and its potential benefit to the retailer and the distribution channel, this sub-section considers the case where the manufacturer and the retailer engage in a dynamic, infinitely repeated game, but do not share sales information. This allows us to first understand the effect of repeated interaction and disentangle it from the effect of past sales information. The main conclusion of this subsection is that even when the manufacturer does not learn the state ex-post - repeated interaction increases the profits of the manufacturer and the retailer, as well as improves the efficiency of the distribution channel. This improvement entails the retailer receiving negative payoffs during periods of low demand.

Because the retailer and the manufacturer engage in a dynamic game, the dynamic participation constraint, IR^D_L (as in equation (6)), is binding. However, the manufacturer does not observe the demand realization at the end of the period, so the binding incentive compatibility constraint is identical to the static IC^S_H (as given in equation (2)).[5] The manufacturer's problem is to maximize $pT_H + (1-p)T_L$ subject to:

$$IR^D_\theta: \qquad \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right] \geq \frac{\delta p}{1-\delta}\Delta(q^S_L), \tag{8}$$

$$IC^S_\theta: \qquad \pi_R(q_\theta;\theta) - T_\theta \geq \pi_R(q_{\widetilde{\theta}};\theta) - T_{\widetilde{\theta}}, \quad \theta, \widetilde{\theta} \in \{L, H\}, \quad \theta \neq \widetilde{\theta}. \tag{9}$$

We use the superscript "D, ne" to denote the dynamic settings when no past sales information is exchanged. Let $T^{D,ne}_H(q_H, q_L)$ and $T^{D,ne}_L(q_H, q_L)$ denote the transfer prices charged by the manufacturer in this case. The manufacturer's profit as a function of these quantities is $pT^{D,ne}_H(q_H, q_L) + (1 - p)T^{D,ne}_L(q_H, q_L)$; it can be expressed as:

$$\Pi^{D,ne}_M(q_H, q_L) = p\pi_R(q_H;H) + (1-p)\pi_R(q_L;L) - p\Delta(q_L) + \left[\delta p\left(\Delta(q_L) - \Delta(q^S_L)\right)\right]. \tag{10}$$

The first three terms in equation (10) are the static profit: it is comprised from the expected joint profit of the distribution channel minus the static information rents paid to the retailer. The last term in the squared brackets is the additional manufacturer's profit due to the dynamic interaction between the two firms. The term in the brackets represents the difference between the information rents paid

[5]The proof that the retailer's participation constraint in state H and the the retailer's incentive compatibility constraint in state L are not binding is given in Appendix C.

to the high-type retailer in the dynamic contract and the information rents paid in the static contract. Notice that when the information rents in the dynamic contract are higher than in the static contract, the manufacturer's profit is increasing with δ. Hence, the dynamic consideration introduces the following trade-off for the manufacturer: on one hand, the manufacturer has the standard incentive to reduce the retailer's information rents similar to the analysis of the static model. Yet, the dynamic considerations (when $\delta > 0$) adds an opposite incentive to increase the dynamic information rents above the static information rents.

Let $q_H^{D,ne}$ and $q_L^{D,ne}$ denote the quantities that maximize $\Pi_M^{D,ne}(q_H, q_L)$. Based on the linear demand model, we present the following result (all proofs are provided in Appendix B).

Proposition 1. ***(Contract in the dynamic game without sales information).*** *Suppose that the game is infinitely repeated without sales information sharing.*

(i) *The manufacturer sets the quantities* $q_H^{D,ne} = q_H^S = q_H^*$ *and*

$$q_L^{D,ne} = \frac{V_L}{2} - \frac{p}{(1-p)}\frac{(V_H - V_L)}{2} + \delta\frac{p}{(1-p)}\frac{(V_H - V_L)}{2} = q_L^S + \delta\frac{p}{(1-p)}\frac{(V_H - V_L)}{2}, \qquad (11)$$

such that during high demand the quantity is identical to the static setting, and during low demand the quantity is higher compared with the static setting.

(ii) *The manufacturer charges:*

$$T_H^{D,ne} = T_H^S - \frac{\delta p(1-\delta p)(V_H - V_L)^2}{2(1-p)}, \quad T_L^{D,ne} = T_L^S + \frac{\delta p^2\left(2(1-\delta p)+\delta\right)(V_H - V_L)^2}{4(1-p)^2}. \qquad (12)$$

Recall that in the static setting, the manufacturer offers a lower quantity to the low-type retailer compared with the quantity that maximizes the overall performance of the distribution channel, and he also needs to pay information rents to the high-type retailer. The introduction of repeated interaction allows the manufacturer to increase the quantity offered during low demand, compared with the static setting. Notice that $q_L^{D,ne}$ is comprised of three terms. The first term is the first-best quantity (i.e., $\frac{V_L}{2}$); the second term is the distortion due to asymmetric information in the static setting; the third term is the net effect of the repeated interaction – this last term is positive, implying that repeated interaction mitigates the adverse effect of asymmetric information. This result holds even though the repeated interaction does not have any informational value for the manufacturer (as the manufacturer can not observe the state ex-post). The following corollary describes the effect of the repeated interaction on the quantity offered to the low-type.

Corollary 1. ***(repeated interaction mitigates the quantity distortion).*** *When* $\delta = 0$, $q_L^{D,ne} = q_L^S$. *When* $\delta > 0$, $q_L^{D,ne}$ *is increasing in* δ *and approaches the first best level as* $\delta \to 1$.

To understand the intuition why repeated interaction even without sales information sharing enhances efficiency during the low demand state it would be useful to also examine the way repeated interaction affects the information rents paid to the high-type retailer. The retailer earns in states H and L:

$$\pi_R(q_H^{D,ne};H) - T_H^{D,ne} = \frac{(V_H - V_L)(V_L - pV_H)}{2(1-p)} + \delta(1-\delta p)\frac{p(V_H - V_L)^2}{2(1-p)}, \tag{13}$$

$$\pi_R(q_L^{D,ne};L) - T_L^{D,ne} = 0 - \delta^2 \frac{p^2(V_H - V_L)^2}{2(1-p)} < 0, \tag{14}$$

and in expectation, the retailer earns:

$$\Pi_R^{D,ne} = \frac{p(V_H - V_L)(V_L - pV_H)}{2(1-p)} + \delta(1-\delta)\frac{p^2(V_H - V_L)^2}{2(1-p)} = \Pi_R^S + \delta(1-\delta)\frac{p^2(V_H - V_L)^2}{2(1-p)}. \tag{15}$$

The first terms in each equation (13) - (15) are the retailer's profits in the static game, and the second terms are the additional (positive or negative) profits due to the repeated interaction. Notice that for all $\delta > 0$, compared with the static case, the retailer earns higher profits in the repeated game when demand is high, but lower (hence negative) profits when demand is low. In expectation, the retailer earns higher profits when $0 < \delta < 1$ and the identical profits as in the static game when $\delta = 0, 1$.

The following corollary summarizes the effect of the repeated interaction (without information sharing) on the profits of the retailer.

Corollary 2. ***(repeated interaction without sales information and the retailer's payoff).***

(*i*) *The retailer's profit in state H, $\pi_R(q_H^{D,ne};H) - T_H^{D,ne}$, is higher than his profit in the static game for any $\delta > 0$.*

(*ii*) *The retailer's profit in state L, $\pi_R(q_L^{D,ne};L) - T_L^{D,ne}$, is lower than his profit in the static game (and hence negative) for any $\delta > 0$.*

(*iii*) *The retailer's expected profit, $\Pi_R^{D,ne}$, is higher than his expected profit in the static game for $0 < \delta < 1$, and equals the static profit when $\delta = 0, 1$. Furthermore, $\Pi_R^{D,ne}$ has an inverse U-shape with respect to δ.*

The intuition for these results is the following. Recall that in the static case, the manufacturer faces a trade-off because setting a high quantity during state L forces the manufacturer to pay high information rents in state H. Hence, the manufacturer distorts the quantity offered in state L. The repeated interaction enables the manufacturer to mitigate this trade-off: The manufacturer can raise the quantity offered in state L (compared with the static case), and, thus, increase the overall efficiency of the distribution channel. Doing so increases the incentives of the high-type retailer to mimic the low-type, and should have forced the manufacturer to pay high information rents in all future periods. Indeed, as shown above, the high-type retailer earns higher information rents in the dynamic game compared with the static one. However, the manufacturer is able to limit the amount of information rents paid to the high-type retailer by charging a high fixed fee from the low-type retailer such that the latter earns a negative payoff. In this way, the manufacturer is able to increase the efficiency of the distribution channel without paying too much information rents to the high-type retailer. The low-type retailer agrees to incur negative profits due to the anticipation of higher expected profits during subsequent periods of high-demand compared with the option of refusing to accept this contract and reverting to the static contract option.

Turning to the manufacturer. His expected per-period profit in the repeated interaction is given by:

$$\Pi_M^{D,ne}(q_H^{D,ne}, q_L^{D,ne}) = \frac{p(V_H - V_L)^2 + (1-p)V_L^2}{4(1-p)} + \delta^2 \frac{p^2(V_H - V_L)^2}{4(1-p)}. \tag{16}$$

The first term in (16) is the manufacturer's static profit, and the second term is the (positive) benefit from the repeated interaction. It is possible to see that the manufacturer's profit is increasing with the discount factor, as summarized in the next corollary.

Corollary 3. ***(the manufacturer benefits from the repeated interaction)****. The manufacturer's profit, $\Pi_M^{D,ne}(q_H^{D,ne}, q_L^{D,ne})$, is strictly increasing in δ.*

The repeated interaction introduces a few effects on the manufacturer's expected profit. First, the manufacturer pays a higher information rents to the high-type retailer However, during the low demand state, the manufacturer is able to increase the sold quantity compared with the static setting, and furthermore to leave the retailer with negative payoff during these periods of low demand. The benefits that the manufacturer sees during the low-demand state more than compensates the manufacturer for the additional information rents paid during the high-demand state.

Notice that when $\delta = 1$, the manufacturer offers the first-best quantities during both demand states, such that the distribution channel maximizes its efficiency. Yet, the manufacturer does not earn the entire industry profit. In this case, the manufacturer leaves the retailer with an expected profit equals to the static setting, such that the manufacturer is able to extract the entire joint surplus that stems from the repeated interaction (i.e., the distribution channel's payoff on top of its payoff in the static setting).

4.2 Repeated interaction with past sales information sharing

Suppose that now, at the end of each period, the manufacturer observes all relevant data concerning past demand during this period (sold quantities and market prices). This information enables the manufacturer to perfectly infer the state of the demand. Recall that knowledge about past demand does not improve the manufacturer's forecasting capabilities regarding future demand, since we assume that demand fluctuations are i.i.d. We further assume that sales information is not contractible. Yet, the main conclusion of this subsection is that such information does affect the contract designed by the manufacturer and, if shared, can improve the overall market outcome. In particular, by sharing sales information, incentives for opportunistic behavior on behalf of the retailer are weakened and he is perceived as more trustworthy. Consequently, the manufacturer is motivated to offer a higher quantity during periods of low-demand, compared with the case of no such information; thus, the joint expected profit of the distribution channel increases. In the following section, we will study the incentives of the retailer to share such information.

When the manufacturer observes the retailer's sales information, he can "punish" the retailer for choosing a contract not according to the true state of the demand, by reverting back to offering the static contract during all future periods. Hence, now both IC_θ^D and IR_θ^D are dynamic, according to equations (6) and (7). The manufacturer's problem is to maximize $pT_H + (1-p)T_L$ subject to:

$$IR_\theta^D: \qquad \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right] \geq \frac{\delta}{1-\delta} p\Delta(q_L^S), \quad (17)$$

$$IC_\theta^D: \; \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right] \geq \quad (18)$$

$$\pi_R(q_{\tilde{\theta}};\theta) - T_{\tilde{\theta}} + \frac{\delta}{1-\delta} p\Delta(q_L^S).$$

The superscript "D,e" denotes the case in which sales information is exchanged in this repeated relationship. Let $T_H^{D,e}(q_H,q_L)$ and $T_L^{D,e}(q_H,q_L)$ denote the solution to IR_L^D and IC_H^D for T_H and T_L.[6] The manufacturer's ex-ante profit as a function of the quantities is:

$$\Pi_M^{D,e}(q_H,q_L) = p\pi_R(q_H;H) + (1-p)\pi_R(q_L;L) - p\Delta(q_L) + \left[\frac{1+p}{1+\delta p}\right]\delta p\left(\Delta(q_L) - \Delta(q_L^S)\right). \quad (19)$$

As in the dynamic game without sales information sharing, the manufacturer's profit is comprised from two parts. The first three terms in equation (19) are identical to the static profit – it is the total expected profit of the distribution channel minus the retailer's static information rents. The last term is the dynamic information rents, which is positive if the information rents paid to the high-type retailer are higher than in the static setting, and is increasing in δ. Comparing the last term in equation (19) with the equivalent term in the manufacturer's profit without sales information in equation (10) reveals that sales information introduces the additional term of $\frac{1+p}{1+\delta p} > 1$. Hence, given the same quantities offered by the manufacturer as in the no information sharing case, sales information enables the manufacturer to increase his profit

[6] The proof that IR_H^D and IC_L^D are not binding with sales information is given in Appendix C.

compared with the dynamic setting without information sharing. The intuition for this result is that sales information enables the manufacturer to "punish" the retailer for choosing the "wrong" contract, which enables the manufacturer to reduce the information rents paid to the high-type retailer. The manufacturer takes this into consideration as he designs the optimal quantities offered when information is shared. We next characterize the contract designed by the manufacturer when sales information is shared.

Proposition 2. ***(Contract in the dynamic game with past sales information)*** *Suppose that the game is infinitely repeated with past sales information sharing.*

(*i*) *The manufacturer sets the quantities* $q_H^{D,e} = q_H^S = q_H^*$, *and*

$$q_L^{D,e} = q_L^S + \frac{(1+p)}{(1+\delta p)} \frac{\delta p}{(1-p)} \frac{(V_H - V_L)}{2}. \tag{20}$$

(*ii*) *The manufacturer charges:*

$$T_H^{D,e} = T_H^S - \frac{(1+p)}{(1+\delta p)^2} \frac{\delta p(1-\delta p)(V_H - V_L)^2}{2(1-p)}, \tag{21}$$

$$T_L^{D,e} = T_L^S + \frac{(1+p)}{(1+\delta p)^2} \frac{\delta p^2 \left(2+\delta(1-p)\right)(V_H - V_L)^2}{4(1-p)^2}.$$

As in the case without sales information, the manufacturer sets a higher quantity in state L than the static quantity: $q_L^{D,e} > q_L^S$, and this quantity increases with δ and approaches the first-best level as $\delta \to 1$. Moreover, the fixed fees in state L (H) are higher (lower) than in the static case. The effect of an increase in δ on the firms' expected profits is qualitatively similar to the case of no sales information: the manufacturer's expected profit is increasing in δ while the retailer's expected profit has an inverse U-shape. We therefore move in the next subsection to compare between the equilibrium quantities and profits when the manufacturer observes sales information with the previous case of no sales information, and examine whether such information will be shared in equilibrium.

4.3 Should the retailer share past sales information with the manufacturer?

We now move to our main research question: when does sharing past sales information benefit the two firms? Naturally, the manufacturer cannot be worse off by receiving this information because he can always ignore it and offer the no information contract. Yet, exchange of information should involve both the manufacturer and the retailer. Only if the retailer is also better-off by sharing sales information, the retailer will collaborate with the manufacturer. The main conclusion of this subsection is that sales information can be mutually beneficial, when the retailer's discount factor, δ, is not too high.

To answer the above questions, we first show how past sales information affects the offered quantities. The following corollary compares between the quantity in state L with and without sales information.

Corollary 4. ***(past sales information increases the quantity in state*** L***).*** *For all* $0 < \delta < 1$, $q_L^* > q_L^{D,e} > q_L^{D,ne}$.

The corollary shows that sharing sales information increases the quantity offered in state L compared with the setting of repeated interaction without information exchange. Hence, sales information further mitigates the problem of quantity distortion during the low demand state and, thus, increases the joint profits of the distribution channel as well as the consumer surplus. The intuition for this result is the following. Recall that the manufacturer faces the trade-off between reducing the retailer's information rents and maximizing the total profits of the distribution channel. As the manufacturer decreases the quantity in state L, he reduces the high-type retailer's incentives to mimic the low-type, and can, consequently, offer the retailer lower information rents in state H; however, such an action reduces the total profits in state L. Sales information relaxes this trade-off: by sharing information the retailer is more trustworthy, and hence, for a given q_L, the manufacturer can elicit the retailer to reveal the high state while reducing the retailer's information rents. Information sharing introduces a new lever for motivating the retailer to reveal a high state: the manufacturer observes the state ex-post based on the sales information and can "punish" the retailer when deviation is observed. The reduction in information rents for a given q_L, enables the manufacturer to raise this quantity, in comparison with the case of no sales information.

In spite of the manufacturer's ability to reduce the retailer's information rents for a given set of offered quantities, when sales information is exchanged, the following proposition shows that information may actually increase the retailer's expected profits.

Proposition 3. ***(past sales information can benefit the retailer).*** *There is a threshold,* $\widetilde{\delta}$, *such that sharing past sales information benefits the retailer,* $\Pi_R^{D,e} \geq \Pi_R^{D,ne}$, *if and only if* $\delta \leq \widetilde{\delta}$, *where:*

$$\widetilde{\delta} = \frac{1}{1+\sqrt{1+p}}. \tag{22}$$

Figure 1 illustrates the result of Proposition 3. As discussed above, for a given q_L, sales information reduces the retailer's expected profit. Yet, Corollary 4 reveals that this allows the manufacturer to increase the equilibrium quantity in state L and the overall value of the distribution channel. The increase in the offered quantity during the low demand state results in an increase in the retailer's information rents in state H. Proposition 3 shows that when $\delta \leq \widetilde{\delta}$ the latter effect may outweigh the former, resulting in higher expected profits for the retailer when information is exchanged.

This result raises the question why the second effect dominates the first when δ is low. Recall from

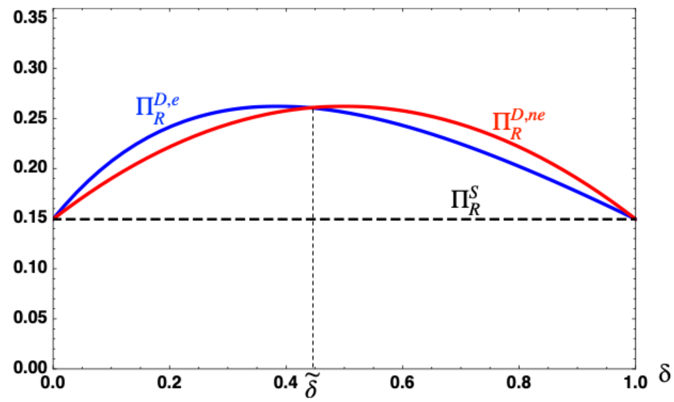

Figure 1: Retailer's expected profit with and without ex-post information as a function of δ ($V_H = 3,\ V_L = 2,\ p = 0.6$).

Corollary 1 that identifies the features of $q_L^{D,ne}$ that when information is not shared, the quantity distortion, $q_L^* - q_L^{D,ne}$, is high (low) when δ is low (high). This implies that when δ is low the joint profits are substantially lower than the potential of the distribution channel. Once information is shared, if δ is low, the manufacturer uses the information sharing more for alleviating the quantity distortion problem than for decreasing the retailer's information rents. Hence, the effect described above, of increasing the quantity during the low state demand, is stronger than the effect of reducing the retailer's information rents for low values of δ. The opposite case occurs when δ is high. In this case, even in the absence of sales information, the quantity distortion is relatively small (due to the repeated interaction, as was shown in Corollary 1). Hence, in this case, the main effect of information sharing is to decrease the retailer's information rents and the secondary effect is to further alleviating the quantity distortion. Therefore, in this case, the retailer is hurt when sales information is shared.

Another way to understand the incentive of the retailer to share sales information is through the lenses of trust. Absent information sharing, the retailer is not perceived as trustworthy, and the manufacturer uses the contract design to elicit the retailer to choose a contract according to the realized demand state. This is achieved by distortion of the offered quantity during periods of low demand – a distortion that also reduces the overall channel efficiency. Sharing sales information ensures that the retailer chooses a contract according to the realized demand state because any deviation is immediately observed and "punished" by reverting to the static menu of contracts. Consequently, the retailer is perceived as more trustworthy, and the manufacturer can limit the quantity distortion during periods of low demand. Part of the increase in the overall channel efficiency propagates to the retailer.

It is interesting to compare the results presented above against classic results obtained when firms interact repeatedly. In many games of infinitely repeated interaction, when the discount factor is high enough, the first best solution can be achieved (known as the folk Theorem). For example, competing firms can form a horizontal cartel and coordinate on the monopoly prices. In our setting, it is possible to interpret the mechanism as "vertical" coordination within the distribution channel. Yet, for any $\delta < 1$, although the introduction of repeated interaction improves channel's efficiency, the channel's performance is still sub optimal compared with the centralized outcome (see Proposition 1) due to the incentive of the manufacturer to maximize its own payoff and not the channel's performance (see Appendix D for a discussion of this issue). Therefore, to further increase efficiency, past sales information comes into play. By sharing such information, the distortion level, during periods of low demand, is further mitigated and channel performance is improved. In our setting it is the combination of repeated interaction and sharing sales information that results in this improvement; however, even in this case, the channel cannot achieve first best (for any $\delta < 1$). Finally, note that while in a horizontal coordination, the competing firms increase their payoffs at the expense of the consumers, in our case repeated interaction and sharing past sales information increase the channel's efficiency and improve consumer welfare, due to the increased sold quantity during periods of low-demand.

5 Extensions

Our base model makes two assumptions that we relax in this section. First, we consider the case where the manufacturer imperfectly observes the state of the demand. Second, we solve for the equilibrium contract when the retailer is financially constrained and, thus, cannot accept a contract that results in a negative payoff during any period.

5.1 Imperfect sales information

So far, we have assumed that when sales information is shared, it provides the manufacturer with the ability to perfectly infer the last period's demand state. In this section, we relax this assumption, and we do it for two main reasons. First, in many instances, the shared information cannot provide the manufacturer with the ability to perfectly infer the state of the demand due to problems in the information exchange process.[7] Second, we have shown, in Section 4, that the retailer benefits (hurts) from information sharing when the discount factor δ is low (high). This finding raises the question whether there are cases in which the retailer prefers to share partial information. In this section, we first characterize the contract under partial information transmission, and we then proceed by showing that there are indeed cases in which the retailer will prefer to share imperfect information, over sharing perfect information or no information.

[7]See for example, https://www.rapidionline.com/blog/data-integration-error-handling.

Suppose now that at the end of each period, the manufacturer observes an imperfect signal $s \in \{\phi, l, h\}$ related to the state of demand. Table 1 characterizes the probability of observing a signal s given the state θ. If the state of demand is H the manufacturer observes the signal h with the probability α, and a non-informative signal ϕ with the complement probability of $1-\alpha$. In a similar manner, if the market demand is low, such that $\theta = L$, the manufacturer observes the signal $s = l$ with the probability of α, and the non-informative signal ϕ with the complement probability of $1-\alpha$. According to this information structure, the value of α captures the precision level of the shared information; when $\alpha = 1$, the model is identical to the information sharing scenario studied in Section 4.2, and when $\alpha = 0$ it is identical to the case of no-information sharing (see Section 4.1). Therefore, according to this information structure, the probability to observe the non-informative signal is $P(\phi) = P(\phi|H) = P(\phi|L) = 1-\alpha$, and with the probability of α the manufacturer observes an informative signal that perfectly reveals the past demand. It is also possible to view the value of α as the ex-ante probability that sales information will be shared in a specific period. We assume that this probability is identical across all periods, and it is independent of the actual demand realization.

		State (θ)	
		H	L
Signal (s)	h	α	0
	l	0	α
	ϕ	$1-\alpha$	$1-\alpha$

Table 1: The probability of observing signal $s \in \{h, l, \phi\}$ in state $\theta \in \{H, L\}$.

Based on this information structure, when the retailer chooses the contract $(q_{\widetilde{\theta}}, T_{\widetilde{\theta}})$ in state $\theta \neq \widetilde{\theta}$ – meaning deviates from the prescribed equilibrium contract, the manufacturer is able to detect such a deviation only with the probability of α. When such a deviation is observed, the manufacturer punishes the retailer by offering the static contract during all subsequent periods (consistent with Section 4). However, with probability $1-\alpha$, the manufacturer does not detect this deviation (since the non-informative signal, ϕ, is observed). We assume that in the latter case, the manufacturer continues to offer the dynamic contract since the manufacturer cannot infer that the deviation has occurred. Hence, imperfect information affects the retailer's incentive-compatibility constraint in the following manner (we use the superscript "D,α" to denote the case of repeated interaction with imperfect information transfer):

$$IC_{\theta}^{D,\alpha}: \quad \pi_R(q_\theta;\theta) - T_\theta + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right] \geq \tag{23}$$

$$\pi_R(q_{\widetilde{\theta}};\theta) - T_{\widetilde{\theta}} + \alpha\frac{\delta}{1-\delta}p\Delta(q_L^S) + (1-\alpha)\frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right].$$

The LHS denotes the retailer's expected profit when choosing the contract according to the realized demand state. The RHS denotes the retailer's expected payoff when deviating and choosing a contract not according to the state of the demand. In this case, with the probability α the manufacturer observes this deviation since the informative signal is observed, and, thus, "punishes" the retailer by repeatedly offering the static contract. With the probability of $(1-\alpha)$ this deviation is not detected and the manufacturer keeps offering the same dynamic equilibrium contract.

Solving equations (23) and (6) for T_H and T_L defines $T_L^{D,\alpha}$ and $T_H^{D,\alpha}$. Substituting into the manufacturer's profit, yields that the optimal quantities offered by the manufacturer are $q_H^{D,\alpha} = \frac{V_H}{2}$ and:

$$q_L^{D,\alpha} = \frac{V_L - pV_H}{2(1-p)} + \frac{\delta p(1+\alpha p)(V_H - V_L)}{2(1-p)(1+\alpha\delta p)} = q_L^S + \frac{\delta p(1+\alpha p)(V_H - V_L)}{2(1-p)(1+\alpha\delta p)}. \tag{24}$$

The first term in $q_L^{D,\alpha}$ is the static quantity, q_L^S, and the second term is due to the dynamic and imperfect information features of this model. We now turn to evaluate the way the precision level α affects the firm's profits:

Proposition 4. ***(Information precision and firms' profits).*** *The manufacturer's profit is increasing with α. The retailer's optimal level of α, denoted as $\hat{\alpha}(\delta)$, is:*

$$\hat{\alpha}(\delta) = \begin{cases} 1; & \text{if } \delta \in (0, \frac{1}{2+p}]; \\ \frac{1-2\delta}{\delta p}; & \text{if } \delta \in (\frac{1}{2+p}, \frac{1}{2}); \\ 0; & \text{if } \delta \in \left[\frac{1}{2}, 1\right], \end{cases}$$

where $\hat{\alpha}(\delta)$ is decreasing with δ when $\delta \in (\frac{1}{2+p}, \frac{1}{2})$.

Proposition 4 shows that if the retailer can choose the extent of information sharing, he will choose full information sharing for low values of the discount factor (i.e., when $\delta \leq \frac{1}{2+p}$); no information sharing for high values of the discount factor (i.e., when $\delta \geq 1/2$) and an intermediate level of α when the value of the discount factor δ is intermediate as well ($\delta \in (\frac{1}{2+p}, 1/2)$). An increase in the precision level α increases the manufacturer's ability to detect a deviation by the retailer and consequently "punish" the retailer by reverting to the static contract; this results in the manufacturer offering, for given quantities q_L and q_H, lower information rents to the retailer. At the same time, an increase in α provides the manufacturer with the ability to alleviate the quantity distortion problem during periods of low demand (note that $q_L^{D,\alpha}$ is increasing with α; this quantity equals the quantity offered in the case of no sales information sharing when $\alpha = 0$ and it equals the quantity offered in the model with full information sharing when $\alpha = 1$). This quantity increase improves the efficiency of the distribution channel and the payoff of the manufacturer, and may also benefit the retailer when the discount factor is low.

While the manufacturer prefers to always receive more precise information, there are cases in which the manufacturer will be better-off by allowing the retailer to choose strategically the level of α. Recall that when information is shared perfectly, the retailer will choose to share such information only when $\delta \leq \widetilde{\delta}= \frac{1}{1+\sqrt{1+p}}$ (see Proposition 3), where $\frac{1}{2+p} <\widetilde{\delta}< 1/2$. Therefore, there is a range, $\delta \in (\widetilde{\delta}, 1/2)$ in which the retailer will decline to share perfect information, but if allowed to strategically choose the level of α, partial information will be shared. Hence, having the opportunity to share partial information results in a higher amount of information being shared. Likewise, there is a range $\delta \in (\frac{1}{2+p}, \widetilde{\delta})$, in which the retailer prefers to share perfect information over not sharing any information, but if possible, prefers to share partial information. Hence, having the opportunity to share partial information results in lower amount of information being shared. We summarize this understanding in the following corollary.

Corollary 5. ***(Retailer's incentive to share partial information)***

(*i*) *For any $\delta \in (\widetilde{\delta}, 1/2)$, where $\widetilde{\delta} = \frac{1}{1+\sqrt{1+p}}$, the retailer prefers to share no information over sharing perfect information, while some level of partial information can be shared. Thus, the ability to share partial information results in higher amount of information being shared compared with the case of sharing perfect or no information.*

(*ii*) *For any $\delta \in (\frac{1}{2+p}, \widetilde{\delta})$, the retailer prefers to share perfect information over not sharing information, while some level of partial information can be shared. Thus, the ability to share partial information results in lower amount of information being shared compared with the case of sharing perfect or no information.*

5.2 Dynamic contracts with limited liability

In the main model we have assumed that in the dynamic settings it is possible to offer a contract that leaves the retailer with negative payoff during periods of low demand. This assumption is reasonable when the retailer is not financially constrained, and hence can recover these losses during periods of high demand. Yet, a financially constrained retailer may not be able to accept a contract with a negative payoff, as the one analyzed above. In this section, we consider the case of limited liability: the retailer cannot accept a contract that results in a negative payoff during any period. The main conclusion of this section is that while the manufacturer is worse-off by limited liability, the retailer is better off when sales information is shared and the discount factor is high, and worse-off otherwise. Moreover, limited liability reduces the efficiency of the distribution channel.

We start with the simple case of limited liability with no sales information sharing. Under limited liability, the retailer's participation constraint is identical to the case of the static settings, IR_L^S. Furthermore, recall, that when no sales information is shared, the incentive compatibility constraint is also identical to the case of static settings, IC_H^S. Therefore, the scenario of limited liability and no sales information sharing reduces to the static setting. We summarize this observation in the following Corollary.

Corollary 6. ***(Limited liability with no sales information sharing).*** *Suppose that the retailer cannot incur negative profits. Then, with no sales information sharing, the manufacturer offers in every period the static contract.*

Next, we study the case of limited liability when sales information is shared. In this case, the participation constraint is identical to the static setting, while the form of the incentive compatibility constraint is identical to the case of past sales information sharing, $IC_H^{D,e}$. We denote this case using the superscript "D,l". Comparing the profits with and without limited liability yields the following result:

Proposition 5. ***(Limited liability hurts the manufacturer).*** *The manufacturer prefers no-limited liability over sales information sharing with limited liability:* $\Pi_M^{D,l} < \Pi_M^{D,ne} < \Pi_M^{D,e}$.

Proposition 5 shows that not only that limited liability hurts the manufacturer, but the manufacturer prefers a scenario without limited liability and without information sharing over a scenario with limited liability and sales information sharing. As the previous sections show, the manufacturer can gain more from the repeated interaction the more he can increase the gap between the retailer's information rents during the two states of the demand. Recall that this high gap can be achieved by offering the retailer negative payoffs during the low demand state, and compensating the retailer with a higher payoff during periods of high demand. Limited liability precludes the manufacturer from offering this negative payoff and, thus, decreases the value of repeated interaction from the manufacturer's perspective.

Turning to the retailer, Figure 2 compares the retailer's expected profit between the following three cases: (i) no limited liability and no information sharing ($\Pi_R^{D,ne}$); (ii) no limited liability and information sharing ($\Pi_R^{D,e}$); (iii) limited liability and information sharing, defined by $\Pi_R^{D,l}$.[8] The figure illustrates that for low values of the discount factor the retailer prefers to share information with no limited liability; for high values of the discount factor, the retailer prefers to share information but with financial constraints; finally, for intermediate values of the discount factor, the retailer prefers not to exchange information and no-limited liability. We formalize this comparison in the following proposition.

Proposition 6. ***(when does the retailer prefer limited liability?)*** *There is a threshold,* $\widehat{\delta}$*, where* $\widetilde{\delta} < \widehat{\delta} < 1$*, such that if and only if* $\delta \geq \widehat{\delta}$*, limited liability with past sales information sharing benefits the*

[8]Recall that the case of limited liability and no ex-post information sharing is identical to the static settings.

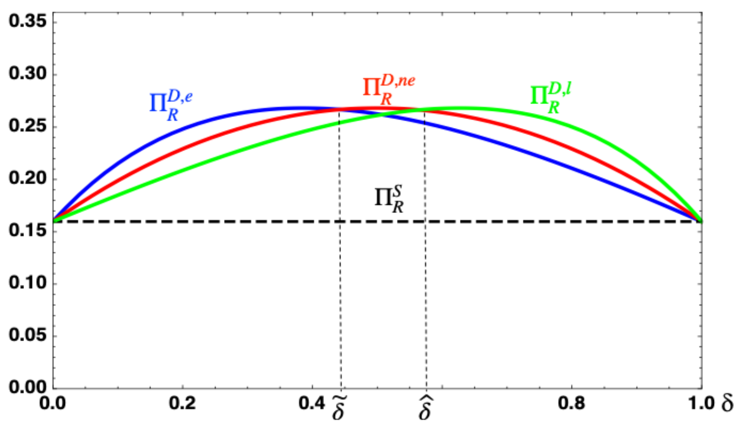

Figure 2: Retailer's expected profit with unlimited and limited liability as a function of δ ($V_H = 3,\ V_L = 2,\ ,p = 0.6$).

retailer in comparison with no limited liability (with and without information sharing): $\Pi_R^{D,l} \geq \Pi_R^{D,ne} \geq \Pi_R^{D,e}$*, where:*

$$\widehat{\delta} = \frac{1}{1 - \sqrt{p}}. \tag{25}$$

Proposition 6 demonstrates that the retailer may actually benefit from being financially constrained. Specifically, when the discount factor is sufficiently high, the retailer's profits in the case of limited liability with past sales information are higher compared with the case of no limited liability (regardless of the information exchange policy). Limited liability results in two opposing effects. First, fixing q_L, limited liability increases the retailer's expected profits since the manufacturer cannot charge a high fee as in the case of no limited liability; this effect benefits the retailer. However, limited liability also weakens the incentives of the manufacturer to increase the offered quantity during the low demand state. This effect hurts the retailer because a low q_L allows the manufacturer to pay lower information rents during periods of high demand, and it also limits the overall value of the distribution channel. Proposition 6 shows that when the discount factor is high enough, the first effect outweighs the second effect; thus, limited liability improves the payoff of the retailer.

Another way to understand this result is related to the trade-off between channel efficiency and bargaining power. In our model, the manufacturer has the bargaining power in the channel, and he makes a take-it-or-leave-it offer to the retailer. In addition, as the discount factor increases, the bargaining power of the manufacturer increases as well, because the retailer is willing to accept a reduction in current period profits in exchange for a promise for a higher future reward. The limited liability setting reduces the bargaining power of the manufacturer as it precludes the ability of the latter to offer contracts with one period negative payoffs. On the other hand, the limited liability setting also results in channel inefficiency due to the inability of the manufacturer to increase the offered quantity during periods of low demand (i.e., alleviate the quantity distortion effect); this channel efficiency loss affects the retailer as well. When the discount factor is very high, the main effect of limited liability is the reduction in the bargaining power of the manufacturer, and, thus, the retailer is better-off with limited liability. However, when the discount factor is low, the main effect of limited liability is the loss of channel efficiency; therefore, in this case, the retailer prefers a contract with no financial constraints.

6 Managerial Implications

6.1 The value of long-term relationships

The paper demonstrates that repeated interaction can benefit both parties in the distribution channel. The introduction of repeated interaction allows the manufacturer to increase the offered quantity during periods of low-demand compared with the static setting. This increased quantity, along with offering the optimal (from the distribution channel's perspective) quantity during periods of high-demand, results in higher overall channel performance. This additional payoff that stems from the repeated interaction is allocated between the manufacturer and the retailer based on the latter's discount factor. In the optimal contract, designed by the manufacturer, the retailer receives a negative payoff during periods of low-demand; the retailer is willing to incur this loss due to the anticipation that during periods of high-demand he will be compensated such that his overall expected payoff is higher than in the static setting. While the overall performance of the distribution channel is increasing with the level of the discount factor δ, the retailer's expected payoff is uni-modal with respect to δ. This implies that it is better for the manufacturer to form long-term relationships with a retailer that has a high discount factor.

Dyer and Ouchi (1993) study the Japanese Auto-manufacturing industry. In their study, they compare American and Japanese automakers and show that although American automakers are more vertically integrated (see Figure 1 in Dyer and Ouchi (1993)), it is the Japanese automakers that enjoy lower production costs and higher customer satisfaction. These authors provide a set of explanations as to how such decentralized supply-chains outperform the more vertically integrated American supply-chains in spite of the double marginalization effect (Spengler (1950)) and the challenge to exchange information under information asymmetry (Cachon and Lariviere (2001)). One of the main explanations these authors provide is that the Japanese style partnerships are characterized by long-term relationships, and that such long-term relationships can solve the inefficiency that stems from information asymmetry. The results presented in our model are aligned with the empirical findings presented by Dyer and Ouchi (1993): long-term relationships can increase the efficiency of a decentralized supply-chain by mitigating the problem of information asymmetry.

6.2 Past sales information sharing and trust formation

After establishing the effect of long-term relationship, this work also explores the incentives to exchange past sales information. In our model, sales information does not provide the manufacturer with better forecasting capabilities, since demand fluctuates between periods in an i.i.d manner and the exchanged information provides only knowledge about past demand. Nevertheless, we show that there are scenarios in which the retailer will find it beneficial to share such information and the manufacturer always benefits from receiving this information.

In our model, information exchange allows the manufacturer to verify that the retailer has chosen the contract according to the true market condition. If the exchanged information reveals that this is not the case, the manufacturer is able to "punish" the retailer by reverting to the contract menu offered in the static setting. Due to this effect, the manufacturer is able to offer the retailer, for a given set of quantities, lower information rents. This information rent reduction makes the retailer worse-off, and may suggest that the latter will choose not to share such information. However, another effect of sales information exchange is that the manufacturer is able to offer a higher quantity during periods of low demand – an outcome that improves the overall performance of the supply-chain and may also increase the expected information rents paid to the retailer. We show that, from the retailer's perspective, when the discount factor is not too high, the latter effect dominates the former, implying that the retailer will choose to share past sales information. In this case, the increase of the offered quantity during periods of low demand more than compensates the retailer for the former effect of information-rents reduction.

These results can also be understood using the concepts of trust and trustworthiness. Absent sales information sharing, the retailer is perceived as less trustworthy, and the manufacturer uses quantity distortion during low-demand periods to encourage the retailer to choose the appropriate contract. By committing to share information, the level of trust increases in this relationship because the incentives for opportunistic behavior are weakened due to the transparency regarding the actions of the retailer. Consequently, the magnitude of quantity distortion diminishes, thus, improving the overall channel efficiency.

6.3 Partial information exchange

In the main model, the focus was on a setting where information can fully reveal the past market demand to the manufacturer. We extend this basic model in Section 5 and examine the implications of sharing partial information on the equilibrium outcome and the incentives of the retailer to share such information. Our results show that when the discount factor is low, the retailer will choose to share perfect information (i.e., perfectly revealing the past demand). In this case, the retailer's payoff is increasing with the level of information exchanged, since the main effect of information exchange is an increase in the sold quantity during periods of low demand. When the discount factor is high, the retailer prefers not to share any information, because in this case the main effect of information sharing is reduction in the information rents offered to the retailer. Since this outcome hurts the retailer, he will choose not to exchange any information.

For an intermediate value of the retailer's discount factor, the retailer will choose to exchange partial information. In this case, the information sometimes enables the manufacturer to infer the state of the demand, and in other cases it leaves the manufacturer with no knowledge about past demand beyond the prior belief. Our results suggest that the ability to exchange imperfect information may facilitate information exchange - there is a region in which the retailer will prefer to exchange imperfect information, but if forced to choose between perfect information exchange and no-information exchange the retailer would choose not to exchange any information. This implies that although the manufacturer always prefers to receive more precise information, he should sometimes allow the retailer the flexibility to choose the precision level of the exchanged information in order to avoid an equilibrium in which no information is shared. At the same time, there is another region in which the retailer prefers to exchange imperfect information, but if forced to choose between perfect information exchange and no-information exchange, the retailer would choose to exchange perfect information. In this region, allowing the retailer to choose the precision level of the exchanged information hurts the manufacturer.

6.4 Financially constrained retailer

The optimal contract designed by the manufacturer entails the retailer receiving a negative payoff during periods of low-demand. We further examine a scenario in which such an option is not viable, and the retailer cannot receive a negative payoff during any period (coined limited liability). Under limited liability, we first show that when no information is exchanged, the outcome is reduced to the static equilibrium. This implies that under limited liability and no sales information sharing, the long-term relationship between the retailer and the manufacturer does not carry any value compared with a relationship that lasts over only one selling period.

When the retailer has limited liability and sales information is shared, the manufacturer and the distribution channel are both worse-off compared with a scenario in which no financial constraints exist. In contrast, the retailer may actually benefit from financial constraints when his discount factor is high. In this case, absent limited liability, the manufacturer is able to extract much of the surplus generated due to the repeated interaction and information exchange. The manufacturer extracts this surplus by offering a negative payoff to the retailer during periods of low demand. Financial constraints limit the bargaining power of the manufacturer in the distribution channel, and, thus, benefit the retailer when the discount factor is high.

Since the manufacturer benefits when the retailer does not suffer from financial constraints, the former may offer financing options to the retailer to ensure that the retailer does not suffer from limited liability. Recent research has also highlighted the strategic effects of trade credit awarded in a distribution channel (Yang et al. 2021, Yang and Birge 2018, Peura et al. 2017 are some recent examples). Our paper highlights the effect of financial constraints on the optimal dynamic contract that is designed by the manufacturer, and the way relaxing these constraints can improve the payoff of the manufacturer and the overall performance of the distribution channel. Interestingly, while liquidation constraints usually hurt the firm under financial stress, this is not the situation in our model – the retailer may benefit from financial constraints since it limits the bargaining power of the manufacturer in the distribution channel.

7 Summary

This paper considers a long-term relationship between a manufacturer and a retailer when in every period past sales information can be transmitted from the retailer to the manufacturer. We focus on the strategic effect of repeated interaction and information exchange and analyze the incentives of the retailer and the manufacturer to send and receive this information.

Asymmetric information typically results in efficiency loss in a distribution channel due to quantity distortion. Repeated interaction and the exchange of sales information allow the manufacturer to alleviate this problem compared with a scenario of a one-period interaction. This increase in channel efficiency always benefits the manufacturer and may also increase the retailer's payoff. On the other hand, information exchange also results in weaker incentives for the retailer to try and mislead the manufacturer; this is because the retailer is concerned from being offered a less favorable contract during all future periods if the manufacturer infers a deviation from the prescribed equilibrium based on the shared information. This concern, of losing future rewards, results in the manufacturer offering lower information rents to the retailer. We demonstrate that when the retailer's discount factor is low, the former positive effect outweighs the latter negative one, and the retailer is willing to share sales information.

We further show that sharing imperfect information may facilitate information exchange. There are cases in which the retailer will refuse to exchange perfect information, but will agree to sharing information that does not provide the manufacturer with the ability to perfectly observe past demand. In this case, the ability to share imperfect information facilitates information exchange. In addition, we discuss the effect of financial constraints on the strategic contract design and we demonstrate that being financially constrained may actually benefit the retailer. The benefit to the retailer, of financial constraints, are that it weakens the bargaining power of the manufacturer by limiting the space of possible contracts.

While the effect of information sharing and the incentives to share it were extensively studied in static settings, less attention has been devoted to understanding these issues in repeated interaction. This research sheds some light on the market conditions that support and hinder such information exchange and can explain the concerns of retailers from taking part in such information exchange projects.

References

[1] Acconcia, A. Martina, R. and Piccolo, S. (2008). Vertical restraints under asymmetric information: On the role of participation constraints. *Journal of Industrial Economics*, Vol. 20, pp. 1389–1406.

[2] Alicke, K., Rexhausen, D., and Seyfert, A. (2017). Supply Chain 4.0 in consumer goods. Mckinsey & Company, 1-11.

[3] Amornpetchkul, T., Duenyas, I., and Şahin, Ö. (2015). Mechanisms to induce buyer forecasting: Do suppliers always benefit from better forecasting? *Production and Operations Management*, 24(11), 1724-1749.

[4] Anand, K., Anupindi, R., & Bassok, Y. (2008). Strategic inventories in vertical contracts. *Management Science*, 54(10), 1792-1804.

[5] Brinkhoff, A., Özer, Ö., & Sargut, G. (2015). All you need is trust? An examination of inter-organizational supply chain projects. *Production and Operations Management*, 24(2), 181-200.

[6] Buell, R. W., & Kalkanci, B. (2021). How transparency into internal and external responsibility initiatives influences consumer choice. *Management Science*, 67(2), 932-950.

[7] Buell, R. W., Kim, T., & Tsay, C. J. (2017). Creating reciprocal value through operational transparency. *Management Science*, 63(6), 1673-1695.

[8] Buell, R.W. and Norton, M.I. (2011). The labor illusion: How operational transparency increases perceived value. *Management Science*, 57(9), pp.1564-1579.

[9] Cachon, G. P., and Lariviere, M. A. (1999). Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Science*, 45(8), 1091-1108.

[10] Cachon, G. P., and Lariviere, M. A. (2001). Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science*, 47(5), 629-646.

[11] Calzolari, C. and Spagnolo, G. (2017). Relational contracts, procurement competition, and supplier collusion. Working paper.

[12] Chen, F. (2003). Information sharing and supply chain coordination. Handbooks in operations research and *Management Science*, 11, 341-421.

[13] Choi, E. W., Özer, Ö., & Zheng, Y. (2020). Network trust and trust behaviors among executives in supply chain interactions. *Management Science*, 66(12), 5823-5849.

[14] Cohen, M. A., Ho, T. H., Ren, Z. J., and Terwiesch, C. (2003). Measuring imputed cost in the semiconductor equipment supply chain. *Management Science*, 49(12), 1653-1670.

[15] Corbett, C. J. (2001). Stochastic inventory systems in a supply chain with asymmetric information: Cycle stocks, safety stocks, and consignment stock. *Operations Research*, 49(4), 487-500.

[16] Crawford, V. P., and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 1431-1451.

[17] Dukes, A., Gal-Or, E., and Geylani, T. (2011). Who benefits from bilateral information exchange in a retail channel?. *Economics Letters*, 112(2), 210-212.

[18] Dukes, A., Gal-Or, E., and Geylani, T. (2017). Bilateral information sharing and pricing incentives in a retail channel. In Handbook of information exchange in supply chain management (pp. 343-367). Springer.

[19] Dyer, J. H., and Ouchi, W. G. (1993). Japanese-style partnerships: giving companies a competitive edge. *MIT Sloan Management Review*, 35(1), 51.

[20] Dyer, J.H. and Hatch, N.W. (2004). Using supplier networks to learn faster. MIT Sloan Management Review, 45(3), p.57.

[21] Ferrantino, M. J., and Koten, E. E. (2019). Understanding Supply Chain 4.0 and its potential impact on global value chains. Global Value Chain Development Report 2019, 103.

[22] Gal-Or, E. (1991a). A common agency with incomplete information. *The RAND Journal of Economics*, Vol. 22 , pp. 274 - 286.

[23] Gal-Or, E. (1991b). Vertical restraints with incomplete information. *Journal of Industrial Economics*, Vol. 39 , pp. 503-516.

[24] Gal-Or, E., Geylani, T., and Dukes, A. J. (2008). Information sharing in a channel with partially informed retailers. *Marketing Science*, 27(4), 642-658.

[25] Grean, M. and Shaw, M.J. (2002). Supply-chain partnership between P&G and Wal-Mart. In E-business management, 155-171. Springer, Boston, MA.

[26] Guan, H., Gurnani, H., Geng, X., & Luo, Y. (2019). Strategic inventory and supplier encroachment. *Manufacturing & Service Operations Management*, 21(3), 536-555.

[27] Ha, A. Y. (2001). Supplier-buyer contracting: Asymmetric cost information and cutoff level policy for buyer participation. Naval Research Logistics (NRL), 48(1), 41-64.

[28] Ha, A. Y., & Tang, C. S. (2017). Handbook of information exchange in supply chain management. Springer.

[29] Ha, A. Y., Tong, S., & Zhang, H. (2011). Sharing demand information in competing supply chains with production diseconomies. *Management Science*, 57(3), 566-581.

[30] Jiang, B., Tian, L., Xu, Y., & Zhang, F. (2016). To share or not to share: Demand forecast sharing in a distribution channel. *Marketing Science*, 35(5), 800-809.

[31] Kalwani, M.U. and Narayandas, N. (1995). Long-term manufacturer-supplier relationships: Do they pay off for supplier firms?. *Journal of Marketing*, 59(1), 1-16.

[32] Kwon, I. W. G., & Suh, T. (2004). Factors affecting the level of trust and commitment in supply chain relationships. *Journal of supply chain management*, 40(1), 4-14.

[33] Kumar, N. (1996). The power of trust in manufacturer-retailer relationships. *Harvard Business Review*, 74(6), p.92.

[34] Laffont, J and Martimort, D. (2002). The theory of incentives. Princeton University Press.

[35] Lee, H. L., and Whang, S. (2000). Information sharing in a supply chain. *International Journal of Manufacturing Technology and Management*, 1(1), 79-93.

[36] Li, L. (2002). Information sharing in a supply chain with horizontal competition. *Management Science*, 48(9), 1196-1212.

[37] Li, T., Tong, S., & Zhang, H. (2014). Transparency of information acquisition in a supply chain. *Manufacturing & Service Operations Management*, 16(3), 412-424.

[38] Li, L., & Zhang, H. (2008). Confidentiality and information sharing in supply chain coordination. *Management Science*, 54(8), 1467-1481.

[39] Lobel, I., & Xiao, W. (2017). Optimal long-term supply contracts with asymmetric demand information. *Operations Research*, 65(5), 1275-1284.

[40] Magretta, J. (1998). The power of virtual integration: An interview with Dell Computer's Michael Dell. Harvard Business Review, 76(2), pp.72-83.

[41] Martimort, D. (1996). Exclusive dealing, common agency, and multiprincipals incentive theory. *The RAND Journal of Economics*, Vol. 27, pp. 1 - 31.

[42] Martimort, D., Semenov, A and Stole, L. (2017). A theory of contracts with limited enforcement. *Review of Economic Studies*, Vol. 84, pp. 816 - 852.

[43] Mohan, B., Buell, R. W., & John, L. K. (2020). Lifting the veil: The benefits of cost transparency. *Marketing Science*, 39(6), 1105-1121.

[44] Oh, S., & Özer, Ö. (2013). Mechanism design for capacity planning under dynamic evolutions of asymmetric demand forecasts. *Management Science*, 59(4), 987-1007.

[45] Özer, Ö., & Wei, W. (2006). Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Science*, 52(8), 1238-1257.

[46] Özer, Ö., & Zheng, Y. (2017). Establishing trust and trustworthiness for supply chain information sharing. In Handbook of information exchange in supply chain management (pp. 287-312). Springer.

[47] Özer, Ö., & Zheng, Y. (2018). Trust and Trustworthiness. In Handbook of Behavioral Operations (pp. 489-524). Wiley Press.

[48] Özer, Ö., Zheng, Y., & Chen, K. Y. (2011). Trust in forecast information sharing. *Management Science*, 57(6), 1111-1137.

[49] Özer, Ö., Zheng, Y., & Ren, Y. (2014). Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. *Management Science*, 60(10), 2435-2460.

[50] Plambeck, E. L., & Taylor, T. A. (2007). Implications of breach remedy and renegotiation design for innovation and capacity. *Management Science*, 53(12), 1859-1871.

[51] Prajogo, D. and Olhager, J., 2012. Supply chain integration and performance: The effects of long-term relationships, information technology and sharing, and logistics integration. *International Journal of Production Economics*, 135(1), pp.514-522.

[52] Peura, H., Yang, S. A., & Lai, G. (2017). Trade credit in competition: A horizontal benefit. *Manufacturing & Service Operations Management*, 19(2), 263-289.

[53] Ren, Z. J., Cohen, M. A., Ho, T. H., & Terwiesch, C. (2010). Information sharing in a long-term supply chain relationship: The role of customer review strategy. *Operations Research*, 58(1), 81-93.

[54] Roy, A., Gilbert, S. M., & Lai, G. (2019). The implications of visibility on the use of strategic inventory in a supply chain. *Management Science*, 65(4), 1752-1767.

[55] Shamir, N., & Shin, H. (2018). The perils of sharing information in a trade association under a strategic wholesale price. *Production and Operations Management*, 27(11), 1978-1995.

[56] Taylor, T. A., & Plambeck, E. L. (2007). Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. *Management Science*, 53(10), 1577-1593.

[57] Yang, S. A., Bakshi, N., & Chen, C. J. (2021). Trade credit insurance: Operational value and contract choice. *Management Science*, 67(2), 875-891.

[58] Yang, S. A., & Birge, J. R. (2018). Trade credit, risk sharing, and inventory financing portfolios. *Management Science*, 64(8), 3667-3689.

[59] Yehezkel, Y. (2008). Retailer's choice of product variety and exclusive dealing under asymmetric information. *The RAND Journal of Economics*, 2008. Vol. 39, pp. 115 - 143.

[60] Yehezkel, Y. (2014). Motivating a supplier to test product quality. J*ournal of Industrial Economics*, Vol. 62, pp. 309 - 345.

[61] Zhang, H. (2002). Vertical information exchange in a supply chain with duopoly retailers. *Production and Operations Management*, 11(4), 531-546.

[62] Zhang, H., Nagarajan, M., & Sošić, G. (2010). Dynamic supplier contracts under asymmetric inventory information. *Operations Research*, 58(5), 1380-1397.

[63] Zhang, H., & Zenios, S. (2008). A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. *Operations Research*, 56(3), 681-696.

Sales Information Transparency and Trust in Repeated Vertical Relationships

By Noam Shamir and Yaron Yehezkel

Online Supplement

We start this online supplement with a discussion of the main model assumptions. Then, Appendix B provides the proofs of propositions 1 - 7. We then, in Appendix C, show that the retailer's individual rationality constraint in state H and incentive compatibility constraint in state L are not binding in the dynamic game. Appendix D considers an alternative mechanism under past-sales information.

Appendix A: Discussion of Main Model Assumptions

In the main model, a few important assumptions were adopted. In this section, we discuss these assumptions, describe the way they can be interpreted to practical market conditions and highlight which assumptions are critical to our results.

Multiple states of demand

For simplicity, our main model assumes two possible states of demand. Yet, asymmetric information may involve additional states, which can be drawn from a discrete or a continuous distribution. In this section, we discuss the implications of such a setting on the results reported above.

A standard result in a setting with multiple states (i.e., more than two), discrete or continuous, is that the manufacturer offers a menu that specifies more than two options. In a static game, the manufacturer sets the quantity that maximizes the profits of a centralized channel for the highest possible state, which is captured by state H in our model, and distorts the quantities below the monopoly level for all other states of demand, which are captured by state L in our model. Such quantity distortion intensifies as the state of demand decreases. For a formal treatment of the case of multiple states under the static setting see Laffont and Martimort (2002).[9] In particular, when the demand parameter of our model, V, is continuously distributed along the interval $V \sim [V_L, V_H]$ according to a distribution function $f(V)$ with a CDF $F(V)$, the quantity in the static game under asymmetric information is $q_V^S = V/2 - (1 - F(V))/2f(V)$.[10] At the

[9] Pages 86 - 92 cover the case of more than two discrete types and pp. 134 - 140 analyze a continuum of types.

[10] To see why, the marginal information rents, equivalent to $\pi_R(q_L; H) - \pi_R(q_L; L) = (V_H - V_L)q_L$ in our two-states framework, is $\partial \pi_R(q_V; V)/\partial V = q_V$. Hence the expected information rents is: $\int_{V_L}^{V_H} \left(\int_{V_L}^{V} q_V dV \right) f(V)dV = \int_{V_L}^{V_H} \frac{1-F(V)}{f(V)} q_V f(V)dV$.

highest possible state of demand, $q^S_{V_H} = V_H/2$ as in the two-states case analyzed above. At the lowest possible state of demand and for a uniform distribution $q^S_{V_L} = V_L/2 - (V_H - V_L)/2$, which qualitatively corresponds to q^S_L in the two-states model in equation (4). Hence, the results of our model concerning the quantity distortion in state L qualitatively capture the predicted quantity distortion of all states (below the highest possible state) in the context of a model with multiple states. Applying the intuition behind the results of our model to a setting with multiple states, we expect that a repeated game without past sales information sharing reduces the quantity distortion in all states of demand (below the highest state). According to the same logic, past sales information is expected to further reduce the quantity distortion.

As for our main result – that the retailer may benefit or hurt by past-sales information sharing – notice that this result (as illustrated by Figure 1) depends on the *expected* profit of the retailer among all states, and it does not directly depend on the number of states. We therefore expect that the two main effects of past sales information sharing that were identified above: manufacturer's ability to detect a deviation (reduces the retailer's expected information rents) and the manufacture's incentive to raise the quantity in low states (increases the retailer's information rents), should follow to more than two states of demand. Hence, along the lines of our model, the retailer may find it optimal to share past sales information or hold such information, depending on the strength of these two conflicting effects.

The retailer has the market power

In the main model, we assumed that the manufacturer – the uninformed firm – has the market power to design a take-it-or-leave-it menu of contracts to the retailer: the privately informed firm. The assumption regarding the distribution of market power in the channel is crucial to our results. To understand why, consider the opposite extreme case in which the retailer is endowed with both the private information regarding demand and has the market power to make a take-it-or-leave offer to the manufacturer. Then, the retailer can offer the contract that specifies the monopoly quantity and compensates the manufacturer for his production costs. In this case, the retailer extracts the entire surplus from the distribution channel, and leaves zero payoff to the manufacturer. A similar outcome happens if it is the manufacturer that has both the market power and is endowed with the private information. Therefore, our model is relevant in settings in which one firm has superior information, while the other has the market power. Much of the literature that studies models of asymmetric information with contract design adopts similar assumptions (e.g., Ha 2001, Özer and Wei 2006, Lobel and Xiao 2017).

Demand is correlated between periods

In the main model, it was assumed that the state of demand fluctuates in an i.i.d manner between periods. Under this assumption, past sales information that is shared cannot improve the predictive capabilities of the manufacturer, thus, allowing us to focus on the strategic effects of sharing past sales information and mute the role of such information in forecasting future demand. We now relax this assumption, and explore the other extreme case – when demand is perfectly correlated across periods, and we then discuss the case of imperfect demand correlation between periods.

Suppose that the state of demand in all periods is identical and equals to $V_\theta \in \{V_H, V_L\}$ with probabilities p and $1-p$, respectively. In the first period, the manufacturer offers the retailer a menu $\{(q_H, T_H), (q_L, T_L)\}$ and the retailer chooses a specific contract. When past sales information is not shared, in equilibrium, the chosen contract reveals the demand state.Therefore, in all future periods, the manufacturer offers a contract $(q_{\widetilde{\theta}}, T_{\widetilde{\theta}}) = (q^*_{\widetilde{\theta}}, \pi_R(q^*_{\widetilde{\theta}}; \widetilde{\theta}))= (V_{\widetilde{\theta}}/2, V^2_{\widetilde{\theta}}/4)$, where $\widetilde{\theta}$ is the inference made by the manufacturer based on the retailer's chosen contract during the first period and $q^*_\theta = V_\theta/2$ is the quantity that maximizes the centralized profit given θ. When the retailer shares past sales information, the manufacturer observes the state of demand of all future periods at the end of the first period and offers in each period the contract that implements the centralized outcome given the true state, $(q_\theta, T_\theta) = (q^*_\theta, \pi_R(q^*_\theta; \theta))= (V_\theta/2, V^2_\theta/4)$. Note that in both of these cases, when no-past sales information is shared and when it is shared, in equilibrium the manufacturer earns in the second period onward the centralized profits while the retailer earns a payoff of 0. Therefore, the difference between these two cases stems from the contract offered by the manufacturer during the first period. The following proposition summarizes the preference of the retailer with respect to information sharing in this setting.

Proposition 7. ***(with perfectly correlated demand, the retailer never shares information).*** *Suppose that demand is identical in all periods. Then, under both past sales and no past sales information, in the first period the manufacturer sets the static quantities and then the centralized quantities in all future periods. Moreover, the retailer is always hurt by sharing past sales information while the manufacturer always benefits from such information.*

The main conclusion is that when past sales information provides perfect knowledge regarding future demand, information losses its strategic role, and the retailer will choose not to share it with the manufacturer. This result highlights the main contribution of our paper: identifying the strategic effect of past sales information as a way to form trust and showing that because of this effect, the retailer may want to share such information. Intuitively, recall that sharing past sales information in our model has two effects on the retailer. First, given a fixed quantity, past sales information decreases the retailer's information rents because the manufacturer can detect and punish a retailer's deviation from choosing the correct contract. Second, past sales information motivates the manufacturer to raise the offered quantity during periods of low demand, which increases the overall efficiency of the distribution channel and also the retailer's information rents. When demand is perfectly correlated across periods, the second effect vanishes and the manufacturer does not increase the quantity following information sharing: the offered quantity is identical to the quantity in the static setting (during the first selling period) and is independent of whether there is information sharing or not. Consequently, past sales information always hurts the retailer.

Based on the analysis of the two extreme cases, of i.i.d demand and perfectly correlated demand, we conjecture that for the case of non-perfect (but positive) correlation between demand periods, there is a threshold, based on the level of correlation, that determines whether the retailer will choose to share with its manufacturer past sales information.

Sales information manipulation by the retailer

So far, we have assumed that the retailer cannot misrepresent the sales information: if information is shared, it is truthfully transmitted to the manufacturer, and in Section 5.1 the assumption was that with a certain ex-ante probability information is not transmitted to the manufacturer. In both of these cases, the retailer was unable to manipulate the content of the shared information.

We now discuss the implications of information manipulation. Suppose that the retailer can choose to transmit past-sales information in a manipulative manner such that when demand is high, the retailer can transmit the wrong information as to create the impression that demand is actually low. In this case, if the retailer wishes to choose a contract not according to the actual demand state, he will manipulate the shared information such that observing past sales cannot allow the manufacturer to infer deviation. Anticipating this, the manufacturer ignores past-sales information. Since such information cannot be used by the manufacturer, the model will be identical to the one without information sharing. We summarize this discussion using the following Corollary.

Corollary 7. *When past sales information can be manipulated, the model is identical to the one without past sales information sharing.*

It is also worth noting that in such a case, a signaling game may emerge in which the retailer signals the true state of demand to the manufacturer. We leave the analysis of such a setting for future research.

Retailer's ability to carry inventory between periods

Another aspect that was muted in the main model is the ability of the retailer to carry inventory between periods. In the main model, we have assumed that units that are not sold during a specific period cannot be used in future periods.

The ability of the retailer to carry inventory between periods changes the solution outlined above. Due to the complexity of this issue, such an extension deserves a separate paper. Below we explain the way such an ability influences the dynamics between the retailer and the manufacturer. Consider a retailer observing a low market demand; in the main model, we show that in the optimal solution such a retailer will be strictly better-off not mimicking a retailer observing a high demand state; thus, the incentive compatibility constraint of the low-type retailer is non-binding. However, when the low type can carry inventory it is possible that for the solution outlined in the main model and when no-information is shared, he will prefer to mimic a retailer observing a high demand state. In this case, the retailer receives a high number of units from the manufacturer and sells in the market only partial quantity (based on the low market demand). In a future period, when demand is high, the retailer may choose the contract designed for the low market demand, and supplement the quantity offered for this contract with the units held in inventory. Moreover, when demand is low, the retailer may choose to reject the contract and sell only the units that he already holds in inventory. Therefore, the contract outlined above for the case with no-information sharing may not constitute an equilibrium when units can be held in inventory and sold in future periods. Yet, note that when past sales information is shared, such an outcome cannot happen because the manufacturer will observe a deviation in the form of discrepancy between the sold quantity and the purchased quantity. While in this paper we do not analyze the case of carrying inventory between periods, the strategic role of inventory in dynamic models has been recognized before. Some examples include Anand et al. (2008), Guan et al. (2019) and Roy et al. (2020), where the last paper also discusses the effect of information transparency on the ability to use inventory in a strategic manner.

Appendix B: Proofs

Below are the proofs of propositions 1 - 7.

Proof of Proposition 1:

Differentiating (10) with respect to q_L and q_H yields the following first-order conditions:

$$\frac{\partial \Pi_M^{D,ne}(q_H, q_L)}{\partial q_H} = p\pi'_R(q_H; H) = 0.$$

$$\frac{\partial \Pi_M^{D,ne}(q_H, q_L)}{\partial q_L} = -p\pi'_R(q_L; H) + \pi'_R(q_L; L) + \delta p \Delta'(q_L) = 0. \tag{26}$$

Substituting $\pi'_R(q_{\widetilde{\theta}}; \theta) = V_\theta - 2q_{\widetilde{\theta}}$ (where $\theta, \widetilde{\theta} \in \{H, L\}$) and $\Delta'(q_L) = V_H - V_L$ into the first-order conditions and solving for q_L and q_H yields $q_H^{D,ne} = q_H^S = q_H^*$ and $q_L^{D,ne}$ as defined in (11). Turning to

$T_L^{D,ne}$ and $T_H^{D,ne}$, solving IR_L^D (in (6)) and IC_H^S (in (2)) for T_H and T_L given q_L and q_H yields:

$$T_L^{D,ne} = \pi_R(q_L; L) + \delta p\big(\Delta(q_L) - (\Delta(q_L^S)\big), \tag{27}$$

$$T_H^{D,ne} = \pi_R(q_H; H) - \Delta(q_L)$$

$$+\delta p\big(\Delta(q_L) - \Delta(q_L^S)\big).$$

Substituting $\pi_R(q_{\widetilde{\theta}}; \theta) = (V_\theta - q_{\widetilde{\theta}})q_{\widetilde{\theta}}$, $q_H = q_H^{D,ne}$, $q_L = q_L^{D,ne}$ and rearranging yields (12).

Remark: Notice that the first line of $T_H^{D,ne}$ in (27) is identical to T_H^S (evaluated at a given q_L), and the second line is the dynamic element. This additional dynamic term is positive, yet the proposition reveals that $T_H^{D,ne} < T_H^S$. The reason why $T_H^{D,ne} < T_H^S$ even though the second line in $T_H^{D,ne}$ is positive is that $q_L^{D,ne} > q_L^S$. The increase in $q_L^{D,ne}$ above q_L^S decreases $T_H^{D,ne}$ because it increases the static information rents (represented by the term: $-\Delta(q_L)$ of the first line), but increases the dynamic information rents (represented by the second line). As $\delta p < 1$, the second effect is weaker than the first effect, resulting in a decrease in $T_H^{D,ne}$ below T_H^S.■

Proof of Proposition 2:

Differentiating (19) with respect to q_L and q_H yields the following first-order conditions:

$$\frac{\partial \Pi^{D,e}(q_H, q_L)}{\partial q_H} = p\pi_R'(q_H; H) = 0,$$

$$\frac{\partial \Pi_M^{D,e}(q_H, q_L)}{\partial q_L} = -p\pi_R'(q_L; H) + \pi_R'(q_L; L) + \left[\frac{1+p}{1+\delta p}\right]\Delta'(q_L) = 0. \tag{28}$$

Substituting $\pi'_R(q_{\widetilde{\theta}}; \theta) = V_\theta - 2q_{\widetilde{\theta}}$ (where $\theta, \widetilde{\theta} \in \{H, L\}$) and $\Delta'(q_L) = V_H - V_L$ into the first-order conditions and solving for q_L and q_H yields $q_H^{D,e} = q_H^S = q_H^*$ and $q_L^{D,e}$ as defined in (20). Turning to $T_L^{D,e}$ and $T_H^{D,e}$, solving IR_L^D (in (6)) and IC_H^D (in (7)) for T_H and T_L given q_L and q_H yields:

$$T_L^{D,e} = \pi_R(q_L; L) + \frac{\delta p}{1+\delta p}\left(\Delta(q_L) - \Delta(q_L^S)\right), \tag{29}$$

$$T_H^{D,e} = \pi_R(q_H; H) - \Delta(q_L) + \frac{2\delta p}{1+\delta p}\big(\Delta(q_L) - \Delta(q_L^S)\big).$$

Notice that as in the case of no ex-post information, the dynamic element in $T_H^{D,e}$ (the last term in $T_H^{D,e}$) is positive for a given q_L. Yet, $T_H^{D,e} < T_H^S$, because $q_L^{D,e} > q_L^S$ and $\frac{2\delta p}{1+\delta p} < 1$, hence we can apply the same intuition as in the proof of Proposition 1. Finally, substituting $\pi_R(q_{\widetilde{\theta}}; \theta) = (V_\theta - q_{\widetilde{\theta}})q_{\widetilde{\theta}}$, $q_H = q_H^{D,e}$, $q_L = q_L^{D,e}$ and rearranging yields (21).■

Proof of Proposition 3:

The retailer's expected profit in the ex-post information case is: $\Pi_R^{D,e} \equiv p(\pi_R(q_H^{D,e};H) - T_H^{D,e}) + (1-p)(\pi_R(q_L^{D,e};L) - T_L^{D,e})$, or:

$$\Pi_R^{D,e} = \frac{p(V_H - V_L)(V_L - pV_H)}{2(1-p)(1+\delta p)^2} \tag{30}$$

$$+\frac{\delta p^2(V_H - V_L)}{2(1-p)(1+\delta p)^2}\left((1-p-\delta(1+p+p^2))V_H + (1+\delta-p+2\delta p)V_L\right).$$

Recalling that $\Pi_R^{D,ne}$ is given by (15), we have that the gap in the retailer's expected profit with and without ex-post information is:

$$\Pi_R^{D,e} - \Pi_R^{D,ne} = \frac{\delta p^3(1-\delta)(V_H - V_L)^2}{2(1-p)(1+\delta p)^2}\left[1-\delta(2+\delta p)\right]. \tag{31}$$

The first term in (31) is strictly positive, hence the sign of $\Pi_R^{D,e} - \Pi_R^{D,ne}$ is determined according to the sign of the term in the squared brackets in (31). We have that $1-\delta(2+\delta p) > 0$ iff $\delta < \widetilde{\delta} = \frac{1}{1+\sqrt{1+p}}$. ■

Proof of Proposition 4:

Solving $IC_H^{D,\alpha}$ from (23) and IR_L^D from (6) for T_L and T_H yields:

$$T_L^{D,\alpha}(q_H,q_L) = \pi_R(q_L;L) + \frac{\delta p}{1+\alpha\delta p}\left(\Delta(q_L) - \Delta(q_L^S)\right), \tag{32}$$

$$T_H^{D,e}(q_H,q_L) = \pi_R(q_H;H) - \Delta(q_L) + \frac{(1+\alpha)\delta p}{1+\alpha\delta p}\big(\Delta(q_L) - \Delta(q_L^S)\big).$$

Substituting $T_L^{D,\alpha}(q_H,q_L)$ and $T_H^{D,\alpha}(q_H,q_L)$ into the manufacturer's expected profit, $pT_H + (1-p)T_L$, yields:

$$\Pi_M^{D,\alpha}(q_H,q_L) = p\pi_R(q_H;H) + (1-p)\pi_R(q_L;L) - p\Delta(q_L) \tag{33}$$

$$+\left[\frac{1+\alpha p}{1+\alpha\delta p}\right]\delta p\big(\Delta(q_L) - \Delta(q_L^S)\big).$$

Notice that given q_L and q_H, $\Pi_M^{D,\alpha}(q_H,q_L) \to \Pi_M^{D,ne}(q_H,q_L)$ as $\alpha \to 0$, $\Pi_M^{D,\alpha}(q_H,q_L)$ is increasing with α, and $\Pi_M^{D,\alpha}(q_H,q_L) \to \Pi_M^{D,e}(q_H,q_L)$ as $\alpha \to 1$. This proves the first part of the proposition. Differentiating $\Pi_M^{D,\alpha}(q_H,q_L)$ with respect to q_L and q_H yields $q_H^{D,\alpha}$ and $q_L^{D,\alpha}$ as defined by (24).

Turing to the retailer's expected profit, substituting $T_L^{D,\alpha}(q_H,q_L)$ and $T_H^{D,\alpha}(q_H,q_L)$ to the retailer's expected profit $p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)$, and evaluating at $q_H^{D,\alpha}$ and $q_L^{D,\alpha}$ yields the retailer's expected profit as a function of α:

$$\Pi_R^{D,\alpha}(\alpha) = \Pi_R^S + \delta(1-\delta)\frac{p^2(1+\alpha p)(V_H - V_L)^2}{2(1-p)(1+\alpha\delta p)^2}. \tag{34}$$

The first term, Π_R^S, is the static expected profit and is independent of α. The second term is maximized at: $\alpha = \frac{1-2\delta}{\delta p}$ (it is straightforward to verify that the second-order condition holds). Because $\frac{1-2\delta}{\delta p} < 0$ when $\delta < 1/2$ and $\frac{1-2\delta}{\delta p} > 1$ when $\delta > \frac{1}{2+p}$, we have that the optimal α for the retailer when taking into account that $0 \leq \alpha \leq 1$ is given by $\hat{\alpha}(\delta)$ as defined in the proposition. ■

Proof of Proposition 5:

Under limited liability, the retailer's participation constraint is identical to the static setting, IR_L^S, while the form of the incentive compatibility constraint is identical to the case of past sales information sharing, $IC_H^{D,e}$. Solving IR_L^S (equation (1)) and IC_H^D (equation (7)) for T_H and T_L, the manufacturer's payoff under limited liability with past sales information sharing, denoted as $\Pi_M^{D,l}$, is:

$$\begin{aligned}\Pi_M^{D,l}(q_H, q_L) &= p\pi_R(q_H; H) + (1-p)\pi_L(q_L; L) - p\big[\pi_R(q_L; H) - \pi_R(q_L; L)\big] \\ &+ \left[\frac{p}{1-\delta(1-p)}\right] \delta p\big(\pi_R(q_L; H) - \pi_R(q_L; L) - (\pi_R(q_L^S; H) - \pi_R(q_L^S; L))\big)\end{aligned} \tag{35}$$

As in the two previous cases of dynamic game without information sharing (equation (10)) and with information sharing (equation (19)), the manufacturer's profit is comprised from the static profit (the first line in equation (35)) and the dynamic information rents (the second line), which is positive if the information rents paid to the high-type retailer are higher than in the static setting. The difference between these three settings is captured in the last term of each profit. Recall from equation (10), that in the dynamic setting without information sharing, the equivalent to the term in the square brackets is 1, while with information sharing, the equivalent term is $\frac{1+p}{1+\delta p} > 1$. Now, under information sharing with limited liability, the equivalent term is $\frac{p}{1-\delta(1-p)}$. As $\frac{p}{1-\delta(1-p)} < 1 < \frac{1+p}{1+\delta p}$, limited liability with information sharing provides the manufacturer with lower expected profits compared with the previous two cases.

Proof of Proposition 6:

We first solve for the optimal contract under limited liability. Differentiating (35) with respect to q_L and q_H yields the following first-order conditions:

$$\begin{aligned}\frac{\partial \Pi^{D,l}(q_H, q_L)}{\partial q_H} &= p\pi_R'(q_H; H) = 0, \\ \frac{\partial \Pi_M^{D,l}(q_H, q_L)}{\partial q_L} &= -p\pi_R'(q_L; H) + \pi_R'(q_L; L) + \left[\frac{p}{1-\delta(1-p)}\right] \Delta'(q_L) = 0.\end{aligned} \tag{36}$$

Substituting $\pi'_R(q_{\widetilde{\theta}}; \theta) = V_\theta - 2q_{\widetilde{\theta}}$ (where $\theta, \widetilde{\theta} \in \{H, L\}$) and $\Delta'(q_L) = V_H - V_L$ into the first-order conditions and solving for q_L and q_H yields $q_H^{D,l} = q_H^S = q_H^*$ and $q_L^{D,l}$, where:

$$q_L^{D,l} = q_L^S + \frac{p}{1-\delta(1-p)} \frac{\delta p}{(1-p)} \frac{(V_H - V_L)}{2} < q_L^{D,ne} \tag{37}$$

Turning to $T_L^{D,l}$ and $T_H^{D,l}$, solving IR_L^S (in (1)) and IC_H^D (in (7)) for T_H and T_L given q_L and q_H yields:

$$T_L^{D,l} = \pi_R(q_L; L), \tag{38}$$

$$T_H^{D,l} = \pi_R(q_H; H) - \Delta(q_L) + \frac{\delta p}{1-\delta(1-p)} \left(\Delta(q_L) - \Delta(q_L^S)\right).$$

Because $q_L^{D,l} > q_L^S$ and $\frac{\delta p}{1-\delta(1-p)} < 1$, we have that $T_H^{D,l} < T_H^S$ even though the second line in $T_H^{D,l}$ is positive. Substituting $\pi_R(q_{\widetilde{\theta}}; \theta) = (V_\theta - q_{\widetilde{\theta}})q_{\widetilde{\theta}}$, $q_H = q_H^{D,l}$, $q_L = q_L^{D,l}$ and rearranging yields

$$T_H^{D,l} = T_H^S - \frac{1-\delta}{(1-\delta(1-p))^2} \frac{\delta p^2 (V_H - V_L)^2}{2(1-p)}, \tag{39}$$

$$T_L^{D,l} = T_L^S + \frac{p\,(2-\delta(2-p))}{(1-\delta(1-p))^2} \frac{\delta p^2 (V_H - V_L)^2}{4(1-p)^2}.$$

Next, we move to the retailer's expected profit in the case of ex-post information with limited liability. The retailer earns: $\Pi_R^{D,l} \equiv p(\pi_R(q_H^{D,l}; H) - T_H^{D,l}) + (1-p)(\pi_R(q_L^{D,l}; L) - T_L^{D,l})$, or:

$$\Pi_R^{D,l} = \Pi_R^{D,ne} - \frac{\delta(1-\delta)p^2(V_H - V_L)^2}{2(1-\delta(1-p))^2} \left[1 - \delta(2 - \delta(1-p))\right], \tag{40}$$

where $\Pi_R^{D,ne}$ is given by (15). Hence, the sign of the gap $\Pi_R^{D,l} - \Pi_R^{D,ne}$ is given by the sign of the term in squared brackets in (40). We have that $1 - \delta(2 - \delta(1-p)) > 0$ iff $\delta < \widehat{\delta} = \frac{1}{1+\sqrt{p}}$. ■

Proof of Proposition 7:

No past - sales information. Given the first period menu, $\{(q_H, T_H), (q_L, T_L)\}$, if the state in all periods is L and the retailer chooses the contract that reveals this state, the retailer receives in the first period the contract (q_L, T_L) followed by the contract $(q_L^*, \pi_R(q_L^*; L))$ in all future periods which yields to the retailer 0 future profits. Hence, the retailer's individual rationality constraint in state L is identical to the static:

$$IR_L^S : \pi_R(q_L; L) - T_L + 0 \cdot \delta/(1-\delta) \geq 0.$$

If the state is H the retailer chooses the contract (q_H, T_H) in the first period based on the revelation principle. This contract is followed by the contract $(q_H^*, \pi_R(q_H; H))$ in all future periods which yields 0 future profits to the retailer. If the retailer misrepresents a high state as a low state, the retailer receives at the first period the contract (q_L, T_L) followed by the contract $(q_L^*, \pi_R(q_L^*; L))$ in all future periods, which yields positive profit because $\pi_R(q_L^*; H) - \pi_R(q_L^*; L) = \Delta(q_L^*) > 0$. Given that the retailer already misrepresented the state at the first period, the retailer will accept the state L centralized contract in all future periods because this is the only contract that will be offered. Hence, the retailer's dynamic incentive compatibility constraint in state H under perfect correlation is:

$$\pi_R(q_H; H) - T_H + \frac{\delta}{1-\delta} \cdot 0 \geq \pi_R(q_L; H) - T_L + \frac{\delta}{1-\delta} \Delta(q_L^*).$$

Solving the two constraints for T_H and T_L and substituting into the manufacturer's expected profit in the first period, $pT_H + (1-p)T_L$, yields (we use the notation $\widehat{\Pi}$ to capture the case of correlated states of demand):

$$\begin{aligned}\widehat{\Pi}_M^{D,ne}(q_H, q_L) = {} & p\pi_R(q_H; H) + (1-p)\pi_R(q_L; L) - p\Delta(q_L) \\ & - \frac{\delta p}{1-\delta} \Delta(q_L^*).\end{aligned}$$

Notice that the first line is identical to the static profits under asymmetric information. The second line is negative and represents the additional information rents that the manufacturer needs to leave the retailer due to the retailer's potential profits from misrepresenting the type in all future periods. This last line is independent of the quantities of the first period, hence, the manufacturer sets in the first period the static quantities under asymmetric information, $q_\theta = q_\theta^S$ as defined by (4), and earns $\widehat{\Pi}_M^{D,ne}(q_H^S, q_L^S)$ while the retailer earns the expected profit of:

$$\widehat{\Pi}_R^{D,ne}(q_H^S, q_L^S) = p\Delta(q_L^S) + \frac{\delta p}{1-\delta} \Delta(q_L^*).$$

Past - sales information. Assume that sales information is exchanged, and the manufacturer observes it at the end of the first period. In the second period onward the manufacturer offers the contract that maximizes the centralized profit based on the observable state, regardless of the retailer's chosen contract during the first period. Therefore, in the first period, the manufacturer offers the static menu, $\{(q_H^S, T_H^S), (q_L^S, T_L^S)\}$, because the retailer's chosen contract in the first period does not affect future periods. Hence, the manufacturer earns in the first period $\widehat{\Pi}_M^{D,e}(q_H^S, q_L^S) = \Pi_M^S(q_H^S, q_L^S)$ as given by (3), while the retailer earns $\widehat{\Pi}_R^{D,e}(q_H^S, q_L^S) = p\Delta(q_L^S)$.

Comparison between no past sales and past sales information. Under both no-past sales and past sales information, the manufacturer earns in the second period onward the centralized profits while the retailer earns a payoff of 0. In the first period, under both cases, the offered quantities are identical to the static quantities. Moreover, the manufacturer always benefits from past-sales information while the retailer is always hurt by past sales information because:

$$\widehat{\Pi}_M^{D,e}(q_H^S, q_L^S) - \widehat{\Pi}_M^{D,ne}(q_H^S, q_L^S) = \widehat{\Pi}_R^{D,ne}(q_H^S, q_L^S) - \widehat{\Pi}_R^{D,e}(q_H^S, q_L^S) = \frac{\delta p}{1-\delta}\Delta(q_L^*),$$

where the last term is positive because $\Delta(q_L^*) > 0$.■

Appendix C: Proof that IR_H and IC_L are not binding

Below we show that in the dynamic setting, given that IR_L and IC_H are binding, IR_H and IC_L are satisfied.

Repeated game with no ex-post information

We prove that $IR_H^{D,ne}$ and $IC_L^{D,ne}$ are not binding given any arbitrary q_L and q_H, as long as $q_L < q_H$, which holds in equilibrium because $q_L^{D,ne} \leq q_L^* < q_H^* = q_H^{D,ne}$. We start with $IR_H^{D,ne}$. Rearranging IR_H^D from (6), IR_H^D holds when:

$$\pi_R(q_H; H) - T_H + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H; H) - T_H) + (1-p)(\pi_R(q_L; L) - T_L)\right] - \frac{\delta}{1-\delta}\left[p\Delta(q_L^S)\right] > 0 \qquad (41)$$

Substituting $T_L^{D,ne}$ and $T_H^{D,ne}$ from (27) into (41) and rearranging, condition $IR_H^{D,ne}$ becomes $\Delta(q_L) > 0$, which holds because of our assumption that $\pi_R(q; H) > \pi_R(q; L)$ for any q.

Turning to $IC_L^{D,ne}$, this condition is identical to the static constraint in (2). Substituting $T_L^{D,ne}$ and $T_H^{D,ne}$ from (27) into (2), we have that $\pi_R(q_L; L) - T_L - (\pi_R(q_H; L) - T_H) > 0$ if $\Delta(q_H) - \Delta(q_L) > 0$. This condition holds because of our assumption that the gap $\Delta(q) > 0$ is increasing in q and because $q_L^{D,ne} < q_H^{D,ne}$.

Repeated game with ex-post information

We prove that $IR_H^{D,e}$ and $IC_L^{D,e}$ are not binding given any arbitrary q_L and q_H, as long as $q_L < q_H$ and $q_L > q_L^S$, which holds in equilibrium because $q_L^S < q_L^{D,e} < q_H^{D,e}$. We start with $IR_H^{D,e}$. Given T_L and T_H, the IR_H^D is identical to the case of no ex-post information described above. Substituting $T_L^{D,n}$ and $T_H^{D,n}$ from (29) into (41) and rearranging, condition $IR_H^{D,ne}$ holds when:

$$\frac{\Delta(q_L) + \delta p \Delta(q_L^S)}{1 + \delta p} > 0,$$

which holds for any arbitrary q_L and q_H because by assumption, $\Delta(q_L) > 0$.

Next consider $IC_L^{D,e}$. Rearranging IC_L^D from (7), $IC_L^{D,e}$ holds when:

$$\pi_R(q_L;L) - T_L + \frac{\delta}{1-\delta}\left[p(\pi_R(q_H;H) - T_H) + (1-p)(\pi_R(q_L;L) - T_L)\right] \tag{42}$$

$$-\left(\pi_R(q_H;L) - T_H + \frac{\delta}{1-\delta}\left[p\Delta(q_L^S)\right]\right) > 0.$$

Substituting $T_L^{D,n}$ and $T_H^{D,n}$ from (29) into (42) and rearranging, condition $IC_L^{D,ne}$ holds when:

$$\frac{\Delta(q_H) - \Delta(q_L)}{1+\delta p} + \frac{\delta p\left(\Delta(q_H) + \Delta(q_L) - 2\Delta(q_L^S)\right)}{1+\delta p} > 0. \tag{43}$$

The first term is positive because $\Delta(q_H) - \Delta(q_L) > 0$ when $q_H > q_L$. The second term is also positive because $\Delta(q_H) + \Delta(q_L) - 2\Delta(q_L^S) > \Delta(q_H) + \Delta(q_L) - 2\Delta(q_L) = \Delta(q_H) - \Delta(q_L) > 0$, where the first inequality follows because $q_L > q_L^S$ and $\Delta(q)$ is increasing in q, and the second inequality follows because $\Delta(q_H) - \Delta(q_L) > 0$ when $q_H > q_L$.

Repeated game with limited liability

Consider first $IR_H^{D,l}$. Substituting $T_L^{D,l}$ and $T_H^{D,l}$ from (38) into IR_H^D as defined by (41) and rearranging, IR_H^D is: $\pi_R(q_L;H) - \pi_R(q_L;L) > 0$, which holds by assumption. As the retailer is financially constraint in both states, we need to verify that $\pi_R(q_H;H) - T_H^{D,l} > 0$. Substituting $T_H^{D,l}$ from (38) into $\pi_R(q_H;H) - T_H^{D,l}$, we have:

$$\pi_R(q_H;H) - T_H^{D,l} = \frac{(1-\delta)\Delta(q_L) + \delta p\Delta(q_L^S)}{1-\delta(1-p)},$$

which is always positive.

Next, consider $IC_L^{D,l}$. To prove that the financially constraint retailer will not deviate in state L to the contract of state H, it is sufficient to show that $\pi_R(q_H;L) - T_H^{D,l} < 0$. Substituting $T_H^{D,l}$ from (38) into $\pi_R(q_H;L) - T_H^{D,l}$, we have:

$$\pi_R(q_H;L) - T_H^{D,l} = \frac{\delta p\Delta(q_L^S) - (1-\delta(1-p))\Delta(q_H) + (1-\delta)\Delta(q_L)}{1-\delta(1-p)} <$$

$$\frac{\delta p\Delta(q_L) - (1-\delta(1-p))\Delta(q_H) + (1-\delta)\Delta(q_L)}{1-\delta(1-p)} =$$

$$-\left(\Delta(q_H) - \Delta(q_L)\right) < 0,$$

where the first inequality follows because $q_L = q_L^{D,l} > q_L^S$ and $\Delta(q)$ is increasing in q, and the last inequality follows whenever $q_H < q_L$.

Appendix D: An alternative mechanism under past-sales information

Consider the scenario in which past-sales information is shared. Suppose that the manufacturer offers the following mechanism. In each period, as long as the retailer reported the true state during the previous period:

Stage 1: The manufacturer pays F

Stage 2: The manufacturer offers a menu of contracts $\{(q_H^*, T_H^*), (q_L^*, T_L^*)\}$ where recall that q_θ^* is the centralized selling quantity in state $\theta \in \{H, L\}$ and $T_\theta^* = \pi_R(q_\theta^*; \theta)$.

Stage 3: The retailer observes the true realization of $\theta \in \{H, L\}$ and chooses a contract from the menu.

Stage 4: At the end of the period, the manufacturer observes the realization of demand.

The manufacturer continues to offer the above menu as long as the retailer chose the contract according to the actual demand. If the retailer deviated, then the manufacturer stops to pay F and instead offers the static menu in all future periods (which is also the punishment in our paper).

Below we show the following results:

1. This alternative mechanism does not result in positive profits for the manufacturer for all values of δ. In particular, it result negative profit for the manufacturer when δ is small, and positive profit if δ is high.

2. This alternative mechanism always (for all $\delta < 1$) provides lower profits to the manufacturer than the profit of the mechanism in our paper. Hence, the manufacturer would never find it optimal to offer it.

3. If the manufacturer wishes to implement a mechanism with the quantities that maximize total industry profits, the manufacturer would rather do so using the mechanism in our model (and setting the quantities q_H^* and q_L^* instead of $q_H^{D,e}$ and $q_L^{D,e}$).

The above mechanism satisfies the retailer's individual rationality in state L because $T_L^* = \pi_R(q_L^*; L)$, and on top of that the retailer receives a fixed sum of F. We solve for the lowest F that motivates the retailer to choose the contract according to the realized demand. To this end, the incentive compatibility constraint in state H is the following:

$$IC_H^D: \quad 0 + \frac{\delta}{1-\delta} F \geq \pi_R(q_L^*; H) - \pi_R(q_L^*; L) + \frac{\delta}{1-\delta} \left[p\Delta(q_L^S) \right]. \tag{44}$$

To see why, suppose that in a certain period the demand is H. If, at stage 3 of this period, the retailer truthfully chooses the contract that was designed for state H, then the retailer earns 0 at the current period, because $T_H^* = \pi_R(q_H^*; H)$ (note that because the retailer already received the fixed sum of F at the beginning of the period, this payment is viewed as sunk). At stage 4, based on the sales information, the manufacturer observes that the retailer chose the correct contract, and hence the manufacturer will continue to offer this contract during all future periods. Thus, the LHS of IC_H^D captures the payoff of the retailer from choosing the correct contract during periods of high-demand.

Next, consider the possibility that the retailer mis-represents the type (i.e., chooses the contract not according to the realized demand). When the retailer chooses the contract that was designed for state L, he earns at the current period $\pi_R(q_L^*; H) - T_L^* = \pi_R(q_L^*; H) - \pi_R(q_L^*; L) = \Delta(q_L^*) > 0$. When sales information is shared, the manufacturer infers that the retailer has deviated from the prescribed contract, and the manufacturer offers during all future periods the static contract (see Section 3). Again, notice that we omit the fixed payment of F from the current period payoff as it is sunk at the point the retailer decides whether to choose the contract according to the realized demand or the alternative one (furthermore, note that adding the fixed payment of F of the current period to the two sides of IC_H^D has no effect on the results). This payoff constitutes the right-hand-side of IC_H^D.

Solving IC_H^D for F in equality (which is the best situation for the manufacturer), we have that the lowest F that motivates the retailer to choose the appropriate contract according to the demand realization is:

$$F = p\Delta(q_L^S) + \frac{1-\delta}{\delta}\Delta(q_L^*). \tag{45}$$

The manufacturer's expected per-period profit from this mechanism is

$$\widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*) = -F + pT_H^* + (1-p)T_L^*.$$

Substituting $T_\theta^* = \pi_R(q_\theta^*; \theta)$ and F, we have:

$$\widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*) = \pi_R(q_L^*; H) - p(\Delta(q_L^S) - \Delta(q_L^*)) \tag{46}$$

$$-\frac{\Delta(q_L^*)}{\delta}.$$

Notice that the manufacturer's profit from this alternative mechanism is negative for low values of δ. To see why, substituting $\delta \to 0$ at the second line, this second line is: $-\infty$, because $\Delta(q_L^*) > 0$. Moreover, it is possible to confirm that $\widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*)$ is increasing in δ.

Comparing $\widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*)$ with the manufacturer's profit from the mechanism in our base model, $\Pi_M^{D,e}(q_H^{D,e}; q_L^{D,e})$, as given by equation (15) in our base model, we have that for all $\delta < 1$:

$$\Pi_M^{D,e}(q_H^{D,e}; q_L^{D,e}) > \Pi_M^{D,e}(q_H^*; q_L^*) > \widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*),$$

where the first inequality is obtained by revealed preferences: the quantities $q_H^{D,e}$ and $q_L^{D,e}$ maximize $\Pi_M^{D,e}(q_H^{D,e}; q_L^{D,e})$. The second inequality follows because:

$$\Pi_M^{D,e}(q_H^*; q_L^*) - \widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*) = \frac{(1-\delta)\left(\delta p \Delta(q_L^S) + \Delta(q_L^*)\right)}{\delta(1+\delta p)} > 0, \tag{47}$$

where this inequality follows because $\pi_R(q; H) > \pi_R(q; L)$ for all q. Hence, the manufacturer prefers to offer the mechanism in our base model over this alternative mechanism. The two mechanisms provide identical profit only when $\delta = 1$ because then $q_L^{D,e} = q_L^*$ and thus $\Pi_M^{D,e}(q_H^*; q_L^*) = \widetilde{\Pi}_M^{D,e}(q_H^*; q_L^*)$.

Finally, notice that the analysis above (see Equation (43)) implies that if the manufacturer would like to implement the quantities that maximize industry profits, q_H^* and q_L^*, the manufacturer would rather do so using the mechanism in our base model and set $q_\theta^{D,e} = q_\theta^*$, $\theta \in \{H, L\}$, and not using a fixed fee that is independent of the realized demand.