A Behavioral Perspective on Inventory Sharing

Hui Zhao
Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802

Liang "Leon" Xu
Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802

Enno Siemsen
Wisconsin School of Business, University of Wisconsin-Madison, Madison Wisconsin 53706

huz10@psu.edu • lzx103@psu.edu • esiemsen@wisc.edu

December 7, 2016

Abstract: Effective transshipment of inventory is a way to take advantage of aggregation effects in decentralized inventory systems. Previous research in the inventory sharing area has been mostly analytical in nature, focused on optimal stocking and sharing decisions and the coordinating role of transfer prices. We examine inventory sharing effectiveness from a behavioral perspective. Four conditions are necessary to create an effective inventory sharing system: adequate inventory in the system, sufficient requests for sharing, reasonable transshipments upon requests, and responsiveness of initial orders to the transfer price. Using behavioral experiments, we examined these four conditions. Our results show that without a carefully designed system, inventory sharing may be neither effective nor a route to enhanced profitability. In general, decision makers do not stock enough to benefit from sharing, and the opportunity to share inventory leads to an even further reduction in the initial inventory stock. Further, decision makers tend not to react to the transfer price unless inventory is not transparent in the supply chain. In addition, decision makers are more likely to under-request inventory, leading to reduced demand signal processing among decision makers. Finally, we provide evidence that by reducing inventory transparency in the system and providing decision support, inventory sharing can become effective and increase profitability for participants.

Keywords: *transshipment*, *inventory sharing*, *behavioral operations*, *newsvendor networks*

1. Introduction

Risk pooling is a fundamental principle in the design of supply chains. The more one stockpile of inventory can serve multiple markets, the lower is the combined demand uncertainty, and thus, the less safety stock is needed to provide the same service level. This scale economy underlies many practices and much research in modern supply chain management, from newsvendor networks (Van Mieghem and Rudi 2002) to the exploration of product component commonality effects (Baker et al. 1986). Inventory sharing, also called transshipment, is a widely-used industry practice because of its promises to reduce inventory and/or improve service levels because of risk pooling. Many different industries have experimented with inventory sharing systems, including automotive and machine tools (Narus and Anderson 1996), heavy machinery (Zhao et al. 2005, 2006), fashion (Dong and Rudi 2004), wholesale/retail (Gallagher 2002), and many online traders (Zhao and Bisi 2010). Further, many technology-enabled services now exist that allow consumers to share their available resources and capacity either with a platform or directly with each other. Understanding when and how such systems can work is an important research topic.

An abundance of analytical modeling literature exists on inventory sharing, much of which focuses on characterizing optimal stocking and inventory sharing decisions. Early research usually assumed a form of centralized power that could decide which stockpiles were made available to which actors. Findings centered on the optimal decisions to maximize system profit or to minimize system cost (e.g., Robinson 1990). Although such a central authority can exist, the operations of many organizations and supply chains are more decentralized. There are ample examples in practice in which inventory is shared among *independent* actors (Zhao et al. 2005). However, benefiting from aggregation effects in decentralized contexts relies on the commitment of individual actors to share their stockpiles with each other. Unsurprisingly, a large amount of analytical modeling research has therefore examined the decentralized actors' optimal/equilibrium stocking and sharing decisions and how to use incentives to promote the right behavior to make inventory sharing effective. Rudi et al. (2001), referred to as RKP hereafter, was one of the first papers to study inventory sharing in such a *decentralized* system. Specifically, RKP looked at a single-period model in which two retailers, after satisfying their independent demand with their own stock, have the opportunity to share inventory if one retailer has extra supply and the other one needs extra supply. After developing the equilibrium order quantities, the authors explored the use of linear transfer prices to coordinate the system and characterized the coordinating transfer price. Later papers such as Hu et al. (2007) and Shao et al. (2011) continued to explore the existence of coordinating transfer prices in various settings. Although other coordinating mechanisms besides the transfer price have been proposed (Anupindi et al. 2001, Yan and Zhao 2011), the transfer price is still the simplest and most studied coordinating mechanism in the context of inventory sharing.

The above studies all assume rational actors. The behavioral aspects of inventory sharing have so far received little attention. However, inventory sharing is not only a profit-driven action, but also a social behavior, related to, for example, the development of norms of reciprocity within a social context. Therefore, our research examines inventory sharing from a behavioral perspective, aiming to establish whether and when inventory sharing systems can be effective if decision makers are not purely profit-maximizing rational actors. Answering such a research question requires looking at decision makers' stocking decisions as well as examining their requesting and transshipping decisions.

Some other work-in-progress in the field of behavioral operations is beginning to examine inventory sharing systems. A working paper by Bostian et al. (2012) examined inventory orders under automated transshipments. Their research examined whether decision makers in the experiment benefited from inventory sharing by estimating structural models to establish counterfactuals; we examined this question more explicitly through randomized treatment, also allowing subjects to make their own decisions on requesting inventory and accepting/rejecting such requests for transshipment. Two other works are concerned with setting the transfer price: Villa and Katok (2016) focus on the negotiation of the transfer price, and Chen and Li (2016a) focus on the sequence of setting the transfer price and making the stocking decisions. We treat the transfer price as exogenous in our research and examine how decision makers react to the transfer price, among other things.

Inventory sharing systems requires at least two participants – a source, and a recipient. For sharing to take place, a recipient must signal a need to a source who must be willing to fulfill the need of the recipient. For this process to be effective in terms of leading to higher profitability for both participants, four requirements must be met: (1) sufficient inventory must be available in the system, (2) the recipient must request adequate amounts of inventory, (3) the source honors these requests based on inventory availability, and (4) inventory orders must react to the transfer price. The fourth condition is essential to achieve supply chain coordination between both participants if the transfer price serves as a coordinating mechanism.

Unless Condition (1) is met, the benefits of inventory sharing will be minimal. Without sufficient stock, the source is unlikely to be able to share many units of inventory profitably. The complexity of inventory sharing lies in the fact that it has a demand side perspective and a supply side perspective. From a demand side perspective, transshipping inventory to another participant essentially corresponds to selling product to another market. On the other hand, from a supply side perspective, requesting supplementary inventory from another participant constitutes an extra supply source. The relative salience of the participants' perception of these two perspectives will determine whether the opportunity to share inventory leads to higher or lower initial orders compared to a context without inventory sharing. If inventory sharing reduces initial order quantities, too little inventory may be available to allow participants to effectively benefit from inventory sharing.

If Condition (2) cannot be met, i.e., the recipient does not request enough inventory, the source is constrained as to how much inventory she or he can share; further, this also leads to inadequately signaling of needs which will likely cause the source to lower expectations as to how much additional future demand (i.e., requests) the recipient will generate, thus further reducing initial order quantities. From a rational perspective, a participant would not withhold inventory requests when the price is right and the need is there; however, from a behavioral perspective, insufficient requests may be rooted in two different psychological phenomena. First, recipients may want their requests to be reasonable in order not to create a situation in which the source appears unkind by having to violate norms of reciprocity (Falk and Fischbacher 2006). Second, asking for supplementary inventory is essentially the equivalent of asking for help, and power concerns often prohibit such help-seeking requests (Lee 1997). Help-seekers essentially acknowledge their incompetence and admit their dependence on the other person. These factors can diminish a help-seeker's power in relation to the provider of help. In turn, this behavioral aspect may diminish the tendency of recipients to request inventory from a source, even if such a request, if granted, would increase the recipient's profitability.

If Condition (3) is not met, requests from the source have little implication. Sharing inventory in response to a request in a bilateral relationship is to some degree an act of kindness that could trigger future reciprocity (Falk and Fischbacher 2006). The interactions we model and test in a lab are essentially "one-shot" transactions and thus may lack the concrete expectation of a future benefit. Nevertheless, we believe that norms of reciprocity will generally trigger a favorable response by the source, even when it is not profitable to do so. This effect could lead to "oversharing," i.e., more inventory being transshipped between partners than what may directly benefit the source. Conversely, though, one could also argue that source sites may be prone to hoarding and therefore less willing to share their inventory with a recipient. Although hoarding sometimes has the connotation of a serious medical condition, in its milder form of a hoarding orientation (Maycroft 2009), it is a natural phenomenon that has been observed in the context of the beer game (Sterman and Dogan 2015b). In our particular context, the endowment effect, which emphasizes that the owner of inventory may value it higher than the recipient making the request, would also predict that sources of inventory are potentially reluctant to share (Morewedge and Giblin 2015). This would mean that we could observe undersharing, i.e. less transshipment of inventory between participants than what would directly benefit the source and the recipient.

The effectiveness of inventory sharing also benefits from participants being responsive to some coordinating instruments that can be used to increase the profit of the system. In the analytical literature, the transfer price is a key coordinating instrument, leading us to propose Condition (4). A higher transfer price makes sharing with others more profitable in general, thus leading, from a rational perspective, to increased initial order quantities. If the transfer price is right, the decentralized solution will be equivalent to the centralized solution, and the (rational) players in the sharing context are perfectly coordinated. If participants react to the transfer price accordingly, then even if ordering and sharing behavior is irrational, the transfer price may at least be a lever to pull to create a better functioning system. However, there are reasons to think that the transfer price may be an inadequate lever to pull. According to social exchange theory, as the transfer price increases, the transshipment-payment transaction may at some point transition from a social exchange to an economic one (Blau 1964). This implies that for the source, the underlying motivation to share inventory transitions from long-term reciprocity to short-term gain. The effectiveness of incentives to induce sharing and ordering behavior thus depends on whether the decision makers perceive inventory sharing as an economic exchange or a social exchange. Therefore, we also ask the question whether, from a behavioral perspective, the transfer price can adequately influence people's behavior to improve outcomes across participants.

Throughout our paper, we focus essentially on two inventory sharing models. The first is a single-stage model, which is equivalent to the model used in RKP in which sharing happens after demand is realized. The other is a two-stage model in which decision makers can share inventory after the demand in the first stage has been realized and fulfilled, but before the second-stage demand is realized. We introduced the two-stage model for several reasons. One reason is realism. In the single-stage model, sharing inventory carries no risk because the source has no further use for the leftover inventory other than salvage. In the two-stage model, decision makers face the risk that if they share inventory with the other player, this inventory will not be available for them during the second stage of demand. This concern that sharing inventory may disadvantage them is often mentioned by managers in real inventory sharing contexts. Thus, we considered it important to recognize and model it in our experiment. The other reason is to better understand decision heuristics. Optimal requests and transshipments are easy to understand in the single-stage setting. Thus, individual motivations are relatively easy to assess when looking at experimental data. But the disadvantage is that we do not see the impact of boundedly rational behavior and decision heuristics on both requests and transshipments. Optimal requests and transshipments are much more difficult to determine in the two-stage model, making the impact of decision heuristics more observable compared to the single-stage model.

Our behavioral experiments reveal some surprising and interesting results. First, in stark contrast to rational equilibrium behavior, decision makers do not seem to respond at all to the transfer price in terms of their initial order quantities. Second, regardless of the transfer price, compared with a no-sharing scenario, decision makers always seem to order less initial inventory when presented with an inventory sharing opportunity. Third, decision makers appear less willing to ask for/request transshipments but are quite willing to share inventory with a recipient when asked.

The remainder of this manuscript is organized as follows. In Section 2, we present the decision makers' solutions in the single-stage and two-stage models from a rational perspective. In Section 3, we develop behavioral theory related to our context. In Sections 4-6, we present three behavioral experiments designed to test our hypotheses as well as to test different strategies to improve decision making in our context. Finally, we conclude with a summary of managerial and theoretical insights in Section 7, and discuss possible future studies.

2. Rational Theory

In this section, we set up our model framework and discuss the theoretical results for both the single-stage and the two-stage models for rational players. We also numerically calculate the theoretical results for both models under the experiment parameters we use. These results serve as the normative benchmark for analyzing the human behavior in our experiments.

2.1 The Single-stage Sharing Model

We adopted the basic framework and cost parameters from RKP for the single-stage model. Consider two independent decision makers (e.g., retailers) at two distinct locations, indexed *i*, *j*. Each decision maker (player) *i* faces independent random demand D_i in a single period. Players independently order Q_i (*i*=1,2) at price c_i at the beginning of the period before observing their demand. After demand is realized, each player uses his or her own stock to satisfy demand. For each unit of demand satisfied, a player receives revenue of r_i. Any unsold inventory may be salvaged at $s_i < c_i$. The decision maker also incurs a penalty of p_i for each unit of unmet demand. As in RKP, we define $v_i = r_i + p_i$ as the marginal value of additional sales at location *i*. Without inventory sharing, this framework corresponds to two independent newsvendors with their order quantities calculated as $Q_i^N = F_i^{-1}\left(\frac{v_i - c_i}{v_i - s_i}\right)$.

In the single-stage model, inventory can be shared after each player fulfills demand with his or her own stock. Specifically, player *j* may request transshipment of a certain amount from *i*, who may then decide how much, if any, to share with *j*. For each unit of inventory shared from *i* to *j*, player *i* receives a unit transfer price of c_{ij} from *j*, but pays a transportation cost τ_{ij}. To avoid trivial cases (e.g., that a decision maker will just rely on sharing from the other party and never stock his or her own inventory), we assume, as in RKP, that $c_{ij} \in [s_j + \tau_{ij}, v_j]$. To simplify the decisions in the experiments, we assume here and throughout the paper that the transportation cost is negligible ($\tau_{ij} = 0$).

Any inventory sharing problem involves two essential decisions — how much to stock initially and how much to share or request. Under the single-stage model, because inventory sharing occurs at the end of the period after demand is realized, a player (e.g., *i*) should always request how much is needed, $(D_i - Q_i)^+$, and share what is left, $(Q_i - D_i)^+$. Therefore, the total number of units shared from *i* to *j* is

$$T_{ij} = \min\left[\left(D_j - Q_j\right)^+, (Q_i - D_i)^+\right]$$

Given T_{ij}, it can be shown that a unique set of equilibrium order quantities exists for the two players, $\left(Q_i^{SS}, Q_j^{SS}\right)$, which can be calculated by following equation (10) in RPK, i.e.

$$\alpha_i(Q_i) - \beta_i\left(Q_i, Q_j\right)\left(\frac{c_{ij} - s_i}{v_i - s_i}\right) + \gamma_i\left(Q_i, Q_j\right)\left(\frac{v_i - c_{ji}}{v_i - s_i}\right) = \frac{v_i - c_i}{v_i - s_i},$$

where

$$\alpha_i(Q_i) = \Pr(D_i < Q_i),$$

$$\beta_i\left(Q_i, Q_j\right) = \Pr(Q_i + Q_j - D_i < D_i < Q_i), \text{ and}$$

$$\gamma_i\left(Q_i, Q_j\right) = \Pr\left(Q_i < D_i < Q_i + Q_j - D_j\right).$$

2.2 The Two-stage Sharing Model

In the single-stage model, decision makers do not face any remaining demand uncertainty when making their decisions to share inventory. To introduce such uncertainty in its simplest form into the two-stage model, we allow decision makers to share inventory in the middle of the period instead of at the end. This creates a fair comparison between the two models because in both models, decision makers have one opportunity to order and one opportunity to share. Essentially, we can think of the period as split into two identical stages. At the end of the first stage, players have each satisfied their first-stage demand from their own stock and then share inventory in anticipation of the demand in the second stage. For simplicity, we assume that demand across these two stages is independent and identically distributed (*i.i.d.*), and that players do not update their demand estimate after observing the first stage.

The sequence of events is as follows. At the beginning of the period (containing two stages now), players place their orders Q_i simultaneously and independently. Then, demand for the first stage, D_{i1}, is realized. After each player satisfies the first-stage demand using his or her own inventory, a player, *j*, may request transshipment from the other player, *i*. Player *i* then decides how much to share with *j*. Again, player *i* receives a unit transfer price of c_{ij} from *j* for each unit of inventory shared with *j*. After transshipment takes place (we assume no transshipment lead time), each player (i) will face his or her second-stage demand, D_{i2}. D_{i2} is realized at the end of the second stage and each player satisfies the second-stage demand with his or her own inventory. Corresponding revenues and costs are tallied for satisfied demand, salvaged inventory, and unmet demand.

Rong et al. (2010) have examined this two-stage model. The equilibrium solution is characterized by two components, i.e., the transshipment policy at the end of the first stage and the stocking policy at the beginning, which can be obtained through backward induction. Specifically, when τ_{ij}=0, in anticipation of the second-stage demand, each player approaches the transshipment decision with a base-stock conserving transshipment policy. That is, there exists an optimal stock level $q_i^* = F_{i2}^{-1}\left(\frac{v_i - c_{ij}}{v_i - s_i}\right)$ for each player *i* so that if *i*'s inventory level (after she or he satisfies first-stage demand) is below q_i^*, the player will buy up to q_i^* through transshipment if possible; if the inventory level is above q_i^*, the player will sell down to q_i^* through transshipment if possible. Note that when the transportation cost τ_{ij} is not negligible, a player will have different buy-up-to and sell-down-to levels, making the decisions much harder. Hence, to simplify the context, we focus on a case in which τ_{ij}=0 in the experiments.

Given the above transshipment policy, the transshipment amount can be written as

$$T_{ij} = \min\left[(Q_i - D_{i1} - q_i^*)^+, \left(q_j^* - \left(Q_j - D_{j1}\right)\right)^+\right].$$

Following Theorem 2 from Rong et al. (2010), there exist unique Nash equilibrium order quantities for the two players, $\left(Q_i^{TS}, Q_j^{TS}\right)$, that can be calculated from the following equations:

$$\frac{v_i - c_{ji}}{v_i - s_i}\rho_{i1}(Q_i, Q_j) + \frac{v_i - c_{ij}}{v_i - s_i}\rho_{i2}(Q_i, Q_j) + \int_{\Omega_{i1}} F_{i2}(Q_i - D_{i1}) dF_{i1} dF_{j1}$$

$$+ \int_{\Omega_{i2}} F_{i2}(Q_i + Q_j - D_{i1} - D_{j1} - q_j^*) dF_{i1} dF_{j1} \quad = \frac{v_i - c_i}{v_i - s_i},$$

where

$$\rho_{i1}(Q_i, Q_j) = \Pr(Q_i - q_i^* < D_{i1} < Q_i + Q_i - D_{j1} - q_i^* - q_j^*),$$

$$\rho_{i2}(Q_i, Q_j) = \Pr(Q_i + Q_j - D_{j1} - q_i^* - q_j^* < D_{i1} < Q_i - q_i^*),$$

$$\Omega_{i1} = \left\{(Q_i, Q_j): \{D_{i1} < Q_i - q_i^* \ \& \ D_{j1} < Q_j - q_j^*\} \cup \{Q_i - q_i^* < D_{i1} < Q_i \ \& \ D_{j1} > Q_j - q_j^*\}\right\}, \text{ and}$$

$$\Omega_{i2} = \left\{(Q_i, Q_j): \{D_j < Q_j - q_j^* \ \& \ Q_i + Q_j - D_{j1} - q_i^* - q_j^* < D_{i1} < Q_i + Q_j - D_{j1} - q_j^*\} \cup \right.$$

$$\left. \{Q_j - q_j^* < D_{j1} < Q_i + Q_j - D_{i1} - q_i^* - q_j^*\}\right\}.$$

2.3 Normative Benchmarks for Our Experiments

In the above subsections, we provide general theoretical results. Next, we will calculate and discuss the specific equilibrium solutions under the experiment parameters for the no-sharing, single-stage, and two-stage models. These solutions will provide the intuition and benchmarks necessary for analyzing the results from our behavioral experiments.

To obtain validation from the literature, we adopt all applicable parameters from RKP. Specifically, we used a symmetric system in which both players face the same parameters. The cost parameters are $r = 40$, $c = 20$, $s = 10$, and $p = 0$. As mentioned, we kept $\tau = 0$ so that the player will have a single threshold (q^*) to decide whether to buy or sell through transshipment at the end of the first stage in the two-stage model. For our numerical studies, we selected 14 levels of the transfer price in the feasible transfer price interval, $c_{ij} \in [s, v]$. For demand, we adopted the same normal demand in RKP as the demand facing each player *in each stage*. This means that for the two-stage model, $D_{i1}, D_{i2} \sim N(100, 50)$. For a fair comparison, when convoluting the two normal demands together, we have $D_i \sim N(200, 70.71)$ as the demand function for the whole period for the no-sharing and single-stage models. As in RKP, we also used a truncated normal when generating demand to avoid negative demand.

Using the theoretical results, we calculated the equilibrium order quantities (Q^*), the threshold for sharing (q^*) in the two-stage model, and the amount transshipped under the optimal policy (T_{ij}^*) for each player under the no-sharing, single-stage, and two-stage models. Figure 1 presents the optimal order quantities as c_{ij} increases under the three models.

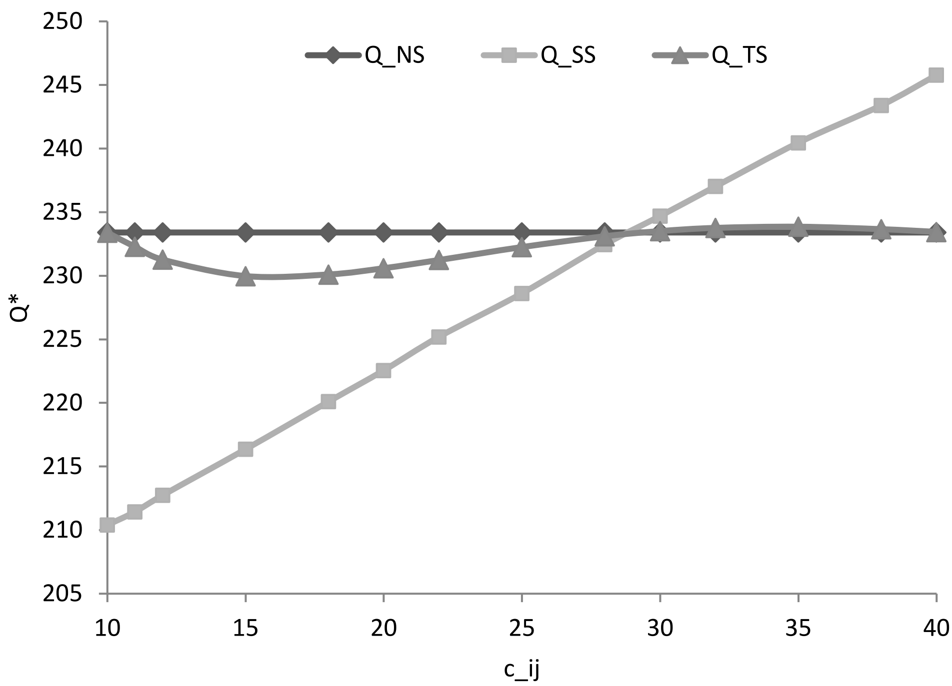

Figure 1: Theoretical Optimal Order Quantities under Different Transfer Prices

First, under the single-stage model, Q^* increases as the transfer price increases, with $Q^{SS*} < Q^{NS*}$ at the lowest transfer price ($c_{ij} = 10$) and $Q^{SS*} > Q^{NS*}$ at the highest transfer price ($c_{ij} = 40$). This follows directly from the theoretical results in RKP. The basic intuition of this result is that, depending on the transfer price, transshipment can be viewed as an additional source of either uncertain supply or uncertain demand. Specifically, at the lowest transshipment price, $c_{ij} = s = 10$, transshipments are a good source of supply but not a good source of demand (a player earns no more than a salvage value from transshipped inventory). Therefore, players would like to order less (than the no-sharing model) to take advantage of this potential supply source. On the other hand, when $c_{ij} = v = 40$, transshipments are a good source of demand but not a good source of supply (a player buys a transshipped unit at the selling price, earning zero profit). Hence, players would like to order more (than the no-sharing model) to take advantage of this possible extra demand source. As the transfer price increases, the tendency to want to take advantage of the possible extra demand from the other player increases. At the same time, the tendency to want to take advantage of the possible supply from the other player decreases. Both of these two forces cause players to order more initial inventory as the transfer price increases.

Under the two-stage model, in addition to the above effect, players in making sharing decisions also face the uncertainty of their own demand in the upcoming second stage. They must take this uncertainty into consideration when deciding the optimal order quantity Q^*. Such uncertainty would have two impacts. The first is the reservation impact. Compared to *no sharing* in the no-sharing model and *complete sharing* in the single-stage model, this uncertainty of the second-stage demand gives the players some cause for reservation when sharing inventory. The two-stage model lies between the no-sharing and the single-stage models, i.e., the players' optimal order quantities and the amount transshipped would be between those under the no-sharing and single-stage models. The second is the extreme value impact, i.e., that players will not share any units when the transfer price is either too low or too high. Specifically, with the lowest transfer price ($c_{ij} = 10$), at the end of the first stage, no transshipment will ever occur because one player gains no extra profit through sharing a unit with the other: He or she only obtains a transfer price equal to the salvage value if he or she sends a unit away but risks being unable to satisfy a unit of demand in the second stage. Therefore, at $c_{ij} = 10$, no sharing will ever occur: Each player orders based on the no-sharing model. Similarly, no transshipment would occur and each player orders based on the no-sharing model at the highest transfer price, $c_{ij} = 40$.

3. Behavioral Theory

Having defined a normative benchmark, we now discuss behavioral reasons that may lead to deviations from this benchmark. As mentioned in the introduction, four conditions must be met to make inventory sharing systems effective: (1) there should be enough inventory in the system, i.e., decision makers need to stock sufficient inventory, (2) transfer prices need to influence how much inventory they stock, (3) decision makers need to request sufficient inventory when in need, and (4) decision makers need to transship sufficient inventory when inventory is requested and available. We next develop behavioral hypotheses related to these four different conditions.

Are Initial Inventory Orders Sufficiently High?

We know from the existing behavioral newsvendor literature that average order quantities are less than optimal for the high-margin condition that we examined in our research (e.g., Zhang and Siemsen 2016). This behavioral regularity implies that inventory will be relatively scarce. In turn, scarcity implies that recipients are more likely to request, but sources are less likely to share. Suppose that in the single-stage model, two players order to obtain a service level of p. After demand is revealed, three scenarios are possible: (1) one player has excess inventory but the other player doesn't, making sharing profitable (likelihood of $2 \times p \times (1-p)$); (2) both players have sufficient inventory so sharing is unnecessary (likelihood of p^2); and (3) both players need extra inventory, but neither has inventory to share (likelihood of $(1-p)^2$). If the possibility to share inventory creates a context where players place initial orders that lead to a service level of p below 50%, the likelihood of scenario (3) occurring increases at the expense of both scenarios (1) and (2). In other words, with little inventory in the system, despite the likelihood of an increased need for sharing, the likelihood of its happening decreases. Scarcity of inventory makes sharing ineffective.

As we mentioned, the complexity of inventory sharing resides in the fact that decision makers can experience it as either an extra source of emergency supply (supply side) or a profitable extra demand opportunity through transshipments (demand side). A rational player considers both sides equally important in his or her stocking decision because both sides may bring profits. From a behavioral perspective, however, one of these two perspectives may be perceived more salient to a player, influencing his or her stocking decisions. For example, if supply-side thinking is perceived more salient for a player (i.e., the player perceives transshipment as more of a supply-side opportunity), he or she will order less initial stock, thinking extra units are available from the other player (which may not be true if the other player lacks sufficient stock). On the other hand, if the demand-side is perceived more salient for a player, she or he will order a higher quantity of initial stock with the hope to sell to the other player.

We expect players to overweight the supply side, leading to lower initial order quantities. First, it is well-known that sharing in centralized systems leads to lower stock requirements (Evers 2001) and higher service levels (Tagaras 1989); hence players may believe that sharing "should" provide "extra supply" to them and lead to less stock, even though the pull-to-center effect already puts them into a situation in which further inventory reduction would not be beneficial. Second, decision makers may value different revenue streams differently, depending on whether these revenues originate in their own market or in sales in secondary markets through transshipments. Specifically, they may value sales in their own market more than sales (through transshipments) in the other players' markets. This is because decision makers will consider sales in their own markets a personal accomplishment, but secondary market sales come through the intermediation of others. This is somewhat akin to the IKEA effect (Norton et al. 2012), i.e., the increased valuation consumers attach to products they assemble themselves. In turn, such increased valuation of sales in one's own market makes the demand side of inventory sharing less salient and can be shown mathematically to lead to reduced initial orders. Based on the above discussion, we hypothesize:

> HYPOTHESIS 1: *Initial orders will decrease when inventory sharing opportunities are presented, such that the initial order quantity under inventory sharing for the rationally neutral transfer price*[1] *will be below the average order quantity observed under no sharing.*

Does the Transfer Price Influence Initial Order Quantities?

An important condition that the transfer price must meet for coordinating decision making is whether it can influence initial order quantities. As discussed in the previous section, for rational actors, the relationship between equilibrium order quantities and transfer prices increases in the single-stage model, but becomes more complex in the two-stage model (See Figure 1 for details). Despite these complexities, higher transfer prices (except for those close to extreme values) are generally associated with higher initial order quantities in both models. This is because in general a low transfer price makes supply-side thinking more salient, but a high transfer price makes demand-side thinking more salient. Specifically, under a lower transfer price, the other player increasingly turns into a cheap source of emergency supply, making initial stocking of less inventory a better tactic to capitalize on this alternative procurement opportunity. On the other hand, under a

[1] I.e., the transfer price at which rational decision makers would stock as much inventory as they would without inventory sharing.

high transfer price, the other player’s market becomes a valuable secondary market for a focal player through inventory sharing. Stocking more initial inventory to capitalize on this secondary market thus becomes the more profitable approach.

From a behavioral perspective, there may be a reason to expect the role of the transfer price to be less pronounced than rationally expected. According to social exchange theory (Blau 1964), exchanges can be seen either in purely economic terms between actors or as social exchanges. In a social exchange, one party does a favor with the expectation that the other party may reciprocate at some future time. Increasing the transfer price clearly moves the transaction from a social exchange to an economic exchange. If decision makers frame inventory sharing under the mindset of an economic exchange, then different transfer prices should entice them to stock different quantities. But if the transfer price merely changes their mindset from a social to an economic exchange, the higher price itself may have less of an effect on initial order quantities. Nevertheless, we believe that transfer prices and initial order quantities remain inherently linked and hypothesize:

> HYPOTHESIS 2: *A higher transfer price leads to a higher initial order quantity than a lower transfer price.*

Are Supplementary Inventory Requests Sufficient?

Inventory requests are a simple form of communication. In our setting, a request per se does not incur any cost and only limits the amount of supplementary inventory that the source can send to the recipient. Therefore, rational players should be free to request what they need.

There is no economic theory of reasonable requests. Economic theories of fairness relate to outcomes, but because requests are not binding on the source, outcomes are determined by the source’s decision to honor them, and not by the requests per se. Nevertheless, reciprocity concerns may apply. Reciprocity in economic exchanges basically means that actors reward kindness by others and punish unkindness (Falk and Fischbacher 2006). A reasonable request is one that allows a source to respond kindly without much cost. An unreasonable request is one that, if honored, would lead the source to incur a significant loss. In the single-stage model, any request by the recipient that does not exceed the source’s demand shortfall is reasonable because the source can simply transship the extra units from storage, reaping the benefits of a transfer price without incurring any cost. Requests exceeding inventory availability may be perceived as unreasonable and unkind because they put the source in a position of either sacrificing profits or appearing unkind and selfish. Thus, recipients may be reluctant to issue such a request to avoid potential future punishment by the source.

One could legitimately ask why it would be a problem if people avoid making unreasonable requests. After all, if a source is unlikely to ship more units than he or she has available, why is it problematic if the recipient under-requests supplementary inventory? The answer is that under-requesting may lead to communication of reduced demand signals. If the recipient persistently under-requests, the source will get the impression that the extra possible demand originating from the recipient is less than it is. In turn, this will further reduce the salience of the demand side of inventory sharing, leading to lower initial inventory levels and subsequently to lower probability of fulfilling future requests.

Another factor may explain why inventory requests may not be made rationally. Asking for a supplementary inventory shipment is akin to asking for help from the source. Such seeking of help implies to some degree an admission of incompetence and also creates a form of dependence on the other person (Lee 1997). Incompetence when making a request means that the recipient admits having placed an insufficient initial order; dependence when making a request implies that the recipient acknowledges his or her inability to serve his or her market without help from the source. Both aspects decrease the relative power of the recipient in relation to the source. Therefore, recipients may tend to avoid making requests in order not to lose their real or perceived comparative power. Considering the above behavioral concerns, we hypothesize:

HYPOTHESIS 3: *Recipients will under-request inventory from their sources.*

Are Transshipments Sufficient?

If the recipient requests inventory, the source decides how much inventory to transship to the recipient. This decision is constrained by the total inventory requested as well as the total inventory available at the source. Rationally, in the single-stage model, the source would transship up to all his or her excess inventory, i.e. the source would satisfy his or her own demand and transship any remaining items upon request to the recipient to earn the transfer price on each unit. Behaviorally, two other factors come into play. First, norms of reciprocity may lead sources to share their inventory and transship upon request, even if this means foregoing profits. Second, hoarding behavior may become prevalent as sources become reluctant to share their scarce inventory.

Reciprocity, as discussed previously, implies that participants will have an incentive to act kindly. The actual game implementation we use is not repeated, and we rematch players anonymously after each round, so direct reputational concerns and social preferences should not affect decision making. Nevertheless, established norms of reciprocity, although no direct benefit can be expected, will still guide decision making toward complying with requests.

Hoarding is essentially an emotional response to scarcity (Sterman and Dogan 2015a). In the single-stage model, scarcity is not in effect for requests made below availability. In the two-stage model, scarcity is more salient, because any inventory retained may be useful in the second stage of the game. Do reciprocity concerns generally outweigh hoarding concerns in this context, or vice versa? In the single-stage model, reciprocity will likely work in the direction of rationality. Requests up to inventory availability will be fulfilled, both to earn profits as well as to avoid appearing unkind. Requests made above inventory availability may be easily dismissed as unreasonable and only filled up to availability. Thus, we do not expect reciprocity to lead to transshipments that are too high. However, hoarding behavior can drive transshipments below their optimum, especially in the two-stage model.

There is another behavioral reason for expecting insufficient transshipments. An important behavioral bias is the so-called "endowment effect," i.e. the tendency of a decision maker to put a higher value on goods he or she owns than the value non-owners would put on them (Morewedge and Giblin 2015). The endowment effect has historically been associated with loss aversion, but current research suggests more basic cognitive roots. In our context, the source owns his or her inventory when making the decision to transship. The endowment effect would lead the source to overvalue the alternative of keeping the inventory in question compared with transshipping it to a supply chain partner, resulting in less than optimal transshipments.

> HYPOTHESIS 4: *(A) Sources will under-transship inventory to recipients. (B) This effect is particularly pronounced in the two-stage model.*

4. Study 1 – Effective Inventory Sharing

Our first study was designed to test our four hypotheses. Hypothesis 1 requires the comparison of inventory sharing to no-sharing settings. Hence, we explicitly vary this factor in our experiment. Hypothesis 2 requires varying the transfer price, which becomes the second factor we manipulate. Hypotheses 3 and 4A are observational hypotheses that do not require a factor manipulation. The third factor we manipulate in Study 1 is the model used (i.e., single-stage vs. two-stage). This allows us to clearly test Hypothesis 4B. Further, contrasting these two models gives us an interesting lens on our other hypotheses as well. Although the single-stage model makes calculations comparatively easy for participants and provides for a stronger effect of the transfer price, the two-stage model is more realistic, and does not provide obvious anchors for decision making. Specifically, although optimal requests and transshipments are easy to calculate in the single-stage model, their calculation becomes challenging in the two-stage model. Thus, while the observed behavior allows easier characterization of the intentions of decision makers under the single-stage model, the two-stage model allows us to better examine the effect of the decision heuristics being used.

There is one important caveat. Introducing the two-stage model not only changes the timing of requests and transshipments, but also changes how demand is revealed. Thus, to allow for a full comparison of treatments, we also need to examine a two-stage version of the no-sharing factorial variation. As we shall see in the analysis, revealing demand over two stages per se decreases initial order quantities because it anchors order quantities on lower numbers.

Experimental Design

Our experimental task followed a classic newsvendor design. Decision makers were instructed to manage a product in a retail setting. In each period, they had to place initial inventory orders before knowing their demand realization. As discussed in Section 2, demand in each period was normally distributed with a mean of 200 and a standard deviation of 70.71. Random draws were different for each subject and were plotted in a time series graph as well as shown in a table as the experiment proceeded. Revenues and costs followed a standard price/cost framing in which the sales price was set at $40, the purchase cost at $20, and the salvage value at the end of the period at $10. We explicitly informed participants that the actual cost of a lost sale was

$0 to avoid ambiguity about the implications of stockouts. The resulting critical fractile was 66%. Demand was truncated at zero, leading to an optimal order quantity of 233 units (without inventory sharing). Subjects were explicitly instructed about this demand distribution and were also informed that demand draws across periods were independent. Subjects repeated the task for 30 time periods. They were compensated upon completion based on their profits obtained across all periods. Their total compensation was based on $6 for showing up in addition to 0.0095% of their total profits earned across all experimental conditions. Average payment per participant was $15.51. The experiment was implemented in zTree (Fischbacher 2007), and conducted with a standardized subject pool at a large American public university. Participants were mostly undergraduate and graduate students at the university. Sessions were scheduled to last for 45 minutes. Participants were required to watch a five-minute video explaining the experimental task. They also had access to written instructions they could refer to at any time during the experiment.

We used a between-subjects design, varying three factors: (1) the experiment mode, (2) the model used, and (3) the transfer price. The experiment mode was varied as either sharing or no-sharing, with the no-sharing mode representing classic newsvendor decision making without the opportunity to share inventory between participants. The sharing mode represented a setting in which subjects had the opportunity to share inventory with each other. There was complete transparency about the other player in the sharing mode – subjects could see the other player's inventory orders and demand. Further, there were no transshipment costs (besides a transfer payment) to either player. The model used was either single-stage or two-stage. In the latter model, under the sharing mode, subjects had the opportunity to share inventory with each other between stages. In the no-sharing mode within the two-stage model, demand stages would simply be revealed directly without an opportunity to share. In both stages, demand was an independent draw from a normal distribution with a mean of 100 and a standard deviation of 50. The transfer price was varied to be either "high" (=$35) or "low" (=$15). These values approximated the transfer prices maximizing/minimizing equilibrium initial order quantities for rational players in the two-stage model. Note that the transfer price was only varied if the experiment mode was sharing, resulting in 2+2*2=6 experimental conditions. A total of 167 subjects participated in the study. See Table 1 below for an overview of all experimental conditions. Participants were grouped into channels of four or six players (channels were unrestricted in no-sharing conditions), and each channel was fixed on one experimental condition. Allocation of subjects to channels was at random within each experimental session. Subjects did not know that they were associated with a channel. They were told that they would be matched to a random player after each period, and they were matched to a random player in their channel after each period.

Table 1: Experimental Conditions in Study 1

	Sharing	Demand	Transfer Price	Participants
Condition 1	No	N(200,70.71)	-	26
Condition 2	No	2*N(100,50)	-	27
Condition 3	After Period	N(200,70.71)	$15	28
Condition 4	After Period	N(200,70.71)	$35	30
Condition 5	After Stage 1	2*N(100,50)	$15	26
Condition 6	After Stage 1	2*N(100,50)	$35	30

Analysis

We began our analysis by examining order quantities. Our unit of analysis is an individual decision. Because decisions in period *t* were nested in individuals and multiple individuals were nested in a zTree channel, we estimated a multilevel mixed-random effects model. Note that the concern of individuals *i* being nested in channel *j* during the experiment was recently highlighted as a methodological consideration in behavioral operations by Hyndman and Embrey (2017); specific to our case, individuals in a channel interact and can affect each other's performance during the experiment. Hence, controlling for this level of nesting is reasonable. We used 5,010 decisions for the analysis; they were nested in 167 individuals and in turn, nested in 39 channels. Standard errors require further adjustments. Our experimental treatments were likely to create differences in order quantity variability across treatments, and thus Huber-White corrected standard errors were used in the analysis. Further, at this level of analysis, observations within a subject constitute a time-series with likely sequential dependence; thus, an AR(1) model was imposed on the error structure. All tests used are two-sided. Independent variables used are sharing (=S) and two-stage (=M) and their interaction; transfer price (=P) effects were estimated in a separate estimation that excluded data from conditions 1 and 2. Thus, Model (1) can be written as follows:

$$Order_{tij} = a_0 + a_1 S_i + a_2 M_i + a_3 S_i \times M_i + \epsilon_j + \epsilon_{ij} + \epsilon_{tij} + \varphi\epsilon_{(t-1)ij}$$

Model (1) was estimated using the mixed procedure in Stata 14.1. Model estimates of the error structure confirm that random effects were present at the individual level (σ=24.8, standard error (s.e.)=2.16). No random effects were observed at the channel level (σ=0.00, s.e.=0.04). The autoregressive component was significant (ρ=0.18, $p \leq 0.01$). The estimated marginal mean for no-sharing and single-stage i.e., classic newsvendor decision making, was equal to 209.74 (s.e.=3.89). This average order quantity was as expected, due to the pull-to-center effect (Schweitzer and Cachon 2000). To estimate marginal effects, we used the *margins* command in Stata with the *asbalanced* option to counter the slight imbalances in sample size between our experimental conditions.

Results from this analysis indicate that the opportunity to share inventory decreases order quantities on average by 9.26 units (p≤0.05), providing support for Hypothesis 1. Being in a two-stage model instead of a single-stage model decreases order quantities by an additional 14.87 units (p≤0.01). This decrease occurs for both sharing models (b=-17.86, p≤0.01) and no-sharing models (b=-11.87, p≤0.05). These two marginal effects do not differ significantly from each other (χ^2=0.55, p=0.46); these estimates suggest that allowing people to share and splitting demand across two stages lead to two separate reductions in initial inventory orders. It is important to emphasize that combining sharing with the two-stage revelation of demand pushes average order quantities considerably below mean demand.

Model (2) is like Model (1), but focuses on a subsample that excludes conditions 1 and 2. This data restriction allowed us to replace the factor variable of sharing with a factor variable for the transfer price. It can be written as follows:

$$Order_{tij} = a_0 + a_1 M_i + a_2 P_i + a_3 M_i \times P_i + \epsilon_j + \epsilon_{ij} + \epsilon_{tij} + \varphi \epsilon_{(t-1)ij}$$

The resulting estimation used data from 3,420 observations nested in 114 individuals nested in 24 channels. The transfer price had no significant overall effect on order quantities (b=-3.83, p=0.45). When estimated differently across single-stage and two-stage models, the transfer price had no effect in the two-stage model (b=-0.08, p=0.99) and only a weakly significant and negative effect in the single-stage model (b=-7.58, p=0.08). Therefore, the expectation that order quantities should increase with the transfer price was not confirmed, and Hypothesis 2 was rejected.

Model (3) repeats the analysis from Model (1), but focuses on a different dependent variable: profit earned. Note that newsvendor research generally uses expected profit as a performance criterion, whereas we used the noisier actual profit made as a criterion. We did so because calculating expected profit in the sharing conditions is nontrivial, whereas actual profits can always be observed. More formally,

$$Profit_{tij} = a_0 + a_1 S_i + a_2 M_i + a_3 S_i \times M_i + \epsilon_j + \epsilon_{ij} + \epsilon_{tij} + \varphi \epsilon_{(t-1)ij}$$

Results indicate that the marginal effect of sharing was not significant (b=49.87, p=0.27), and that the marginal effect of two-stage was negative and significant (b=-258.27, p≤0.01). Thus, participants were unable to reap any benefits from inventory sharing in our experiment, and their performance suffered significantly when we split demand over two stages.

In Model (4), we repeated the estimation of Model (1). This time we used the outcome variable of whether a stockout occurred. Because this dependent variable is dichotomous, we needed to change our model to a random effects probit model (xtprobit procedure in Stata). This change required us to drop both the clustering of individuals within channels as well as autocorrelation in our error term in the analysis. The resulting specification can be described as follows:

$$P(Order_{tij} < Demand_{tij}) = a_0 + a_1 S_i + a_2 M_i + a_3 S_i \times M_i + \epsilon_i + \epsilon_{ti}$$

Marginal effects from this analysis show that sharing per se has no effect on the likelihood of stockouts (b=-0.03, p=0.18), but the two-stage model increases the likelihood of stockouts (b=0.13, p≤0.01).

In summary, participants in our experiment had incredible difficulty leveraging the potential of inventory sharing. Inventory sharing was ineffective in the sense that participants earned similar profits and achieved similar service levels as under the no-sharing setting. Additionally, the moment that we revealed demand across two stages, order quantities decreased by a comparatively large amount (to a point below the mean demand), with corresponding implications for profits and stockouts. The high transfer price appears to fail as an incentive for participants to order more.

To further examine the effectiveness of inventory sharing, we next focused on both the inventory sharing requests made by the recipient and the transshipments sent by the source. To understand requests, we needed to establish a normative benchmark by discussing an optimal request. In the single-stage setting, an optimal request is the difference between demand and inventory ordered — if there is a shortfall, participants can try to make up for it by requesting additional inventory from their source. Similarly, optimal transshipment decisions in the single-stage model are guided by inventory availability. If the source has inventory not needed to satisfy his or her focal demand, he or she can transship this inventory profitably to the recipient. Defining a normative benchmark is slightly more challenging in the two-stage setting. As discussed in Section 2, there is an optimal stock level (q^*) for the second stage that depends on the transfer price. This optimal stock level equals 149 for the low transfer price and 55 for the high transfer price. The difference between these optimal stock levels reflects the difference in profit margins obtainable from selling shared inventory in the market. Consequently, the optimal request for a recipient in the two-stage model is to request the difference between his or her optimal stock level for the second stage and his or her available inventory after satisfying Stage 1 demand. The optimal transshipment by the source follows a similar logic: Upon request, the source only transships any inventory that he or she has available beyond his or her optimal stock level after satisfying Stage 1 demand. We did not inform participants about these optimal stock levels, and thus making an optimal response was challenging in the two-stage model. Participants likely resorted to decision heuristics in this task.

Table 2 contains an overview of all 1,740 requests made in the single-stage model. We differentiated between situations in which the recipient had a shortage and situations in which no such shortage existed. We believe that it is straightforward for participants to understand that they have enough inventory to meet demand if no shortage exists; indeed, we see that participants requested additional inventory in only 6% of cases despite having enough inventory available to meet demand. However, when a shortage existed, it became interesting to examine if participants requested what they need. We saw that recipients' requests exactly matched their shortages only half the time. In the other 50% of situations, subjects mostly requested less (rather than more) than they needed. This observation supports Hypothesis 3.

Table 2: Requests Made by Recipients in the Single-stage Setting

Request	Recipient Has: Shortage (Order < Demand)		No Shortage (Order ≥ Demand)	
Nothing	174	21%	873	94%
Below Shortage	207	25%	-	-
Exactly Shortage	409	50%	-	-
More than Shortage	23	3%	54	6%
	813		927	

To further examine the motivation underlying this tendency to request too little inventory from sources, it was useful to scrutinize the 174+207=381 decisions in which participants under-requested despite having a shortfall. These situations are further analyzed in Table 3. Many recipients requested less than or equal to the amount that the source had available; apparently, the actual inventory availability of the source became an important anchor. If the source had inventory available, 42% of requests were made exactly at that inventory availability. When the source had no inventory available, recipients requested shipments of inventory in only 38% of cases. Thus, as discussed previously, appearing reasonable in their requests seems indeed to be a major factor in the overall under-requesting by recipients.

Table 3: Requests and Source Inventory Availability

Request	Source Has: Excess Inventory (Order > Demand)		No Excess Inventory (Order ≤ Demand)	
Nothing	52	28%	122	62%
Less than Source Availability	36	20%	-	-
Exact Source Availability	78	42%	-	-
More than Source Availability	18	10%	75	38%
	184		197	

Requesting only what is available from a source may be seen as rational. If the other player faces a shortage, then why request inventory from him or her? Or why make a request for inventory that if fulfilled would put the other player into a shortage situation? Following such a rule-of-thumb appears reasonable. However, there are two important caveats. First, although the source may be in a shortage situation, she or he may still decide to ship inventory to a requester. The game was designed in such a way that profits are only resolved after sharing, and a source could choose to prioritize responding to a recipient's request rather than using inventory to satisfy her or his own primary demand. Ultimately, the source must decide how much inventory to ship upon request. Thus, requesting additional inventory does no direct harm to the source.

Second, through the availability heuristic, the source is likely to use the recipient's requests as information to estimate expected demand from this secondary market. Thus, under-requesting means that the source processes censored information, leading the source to generally underestimate the potential sales volume in the secondary market, and hence undervaluing the demand side of inventory sharing.

A more pronounced picture emerges if we look at the two-stage setting. As noted, the optimal requests and transshipments are less obvious for participants in this setting. Nevertheless, participants seem to understand when not to make a request — in 95% (=690/726) of situations in which not making a request was optimal, participants requested no inventory. However, in 97% (922/954) of situations in which subjects should request additional inventory, they requested nothing or less than the optimal amount, clearly emphasizing that under-requesting is more prevalent than over-requesting.

To examine how source players respond to such requests with transshipments, we first focused again on the single-stage setting because optimal decisions are easy to calculate in this context. A key to our analysis here was to interpret transshipment decisions both within the context of what was requested as well as what the source had available. We examined the 693 decisions in which nonzero inventory was requested (=813+927-174-873, see Table 2). Note that if a recipient requested no inventory, the source had no choice and could not send any inventory. Thus, removing such situations from an analysis of transshipment decisions was reasonable. Table 4 below examines how sources respond with transshipments to recipient requests, depending on whether the request exceeds their available inventory. It was apparent that most sources either fulfilled a request or sent what they have available. Strikingly, in 22% of situations in which requests for more inventory exceeded what was available, sources nevertheless complied with these requests and transshipped to the recipient more inventory than they had available despite being able to sell this inventory at a higher margin in their own market. This indicates that participants tend to ship more than they should rather than hold on to their inventory when they should not. Therefore, Hypothesis 4A was rejected – over-transshipping inventory was more likely than under-transshipping.

Table 4: Transshipments Made Upon Request in the Single-stage Setting

	Source Has			
Transshipment	**Sufficient Inventory (Request ≤ Availability)**		**Insufficient Inventory (Request > Availability)**	
Below Availability	-	-	14	4%
At Availability	-	-	294	75%
Below Requested	42	14%	40	10%
At Requested	257	86%	46	12%
	299		394	

Again, a more pronounced picture emerges when looking at the two-stage setting. Note that in this setting, because of the complexity of calculating the optimal stock levels, it may not be immediately clear to either player how much inventory a source has available. When recipients made requests at or below the availability of sources, 74% (=82/111) of transshipments fully honored such requests. Transshipments were less than the request only 26% of the time. Conversely, when recipients placed requests above what a source had available, 59% (185/315) of the time, the source sent a larger transshipment than what was optimal for a source. Clearly, the tendency to over-transship is stronger than the tendency to under-transship. This provides further evidence against Hypothesis 4A, and rejects Hypothesis 4B. Despite a clear shortage situation, sources in the two-stage setting are more eager to transship their inventory upon request.

Profitability Implications

To assess the profitability implications of the above behavioral barriers to effective inventory sharing, we examined observed profits across four different scenarios: (1) *original as observed in the experiments* in which we simply calculated observed average profits, (2) *optimal requests and transshipments* in which we calculated observed average profits, assuming the inventory requests and transshipments were optimal but initial stock levels were as observed in the experiments, (3) *newsvendor* in which in addition to (2) we assumed that subjects had ordered according to the newsvendor critical fractile policy, and (4), *optimal order* in which in addition to (2) we assumed that subjects had ordered the optimal order quantity taking inventory sharing into account. For this analysis, we grouped both price treatments together under the single-stage model (i.e., high and low transfer prices). We did the same with the two-stage model. We did this (i.e., combined these samples) because we know that order quantities do not change as a function of transfer price, and transshipment payments cancel out among players over a treatment. Similarly, we grouped both no-sharing treatments (single- and two-stage demand realization) under scenarios 3 and 4 because optimal order quantities across these two treatments are equivalent. Results from our analysis are shown in Table 5 below.

Table 5: Observed Profits Under Different Scenarios

Scenario	(1) Original		(2) Opt. Req.		(3) Newsvendor		(4) Optimal Order	
	Avg. Profit	SE	Avg. Profit	SE	Avg. Profit	SE	Avg. Profit	SE
No Sharing – Single	3,220	(46)	3,220	(46)	3,356	(37)	3,356	(37)
No Sharing – Two	2,942	(45)	2,942	(45)				
Sharing – Single	3,250	(32)	3,328	(32)	3,416	(38)	3,432	(37)
Sharing - Two	3,011	(28)	3,063	(28)	3,367	(36)	3,367	(35)

Notes. Values in brackets are standard errors (SE).

Several observations can be made from Table 5. First, the benefits of inventory sharing in Scenario (1) are relatively small – 3,250-3220=30 in the single-stage model and 3,011-2,942=69 in the two-stage model.

Replacing observed requests and transshipments with optimal decisions increases these benefits in Scenario (2) to 3,328-3,220=108 in the single-stage model and to 3,063-2,942=121 in the two-stage model. Increasing order quantities to the critical fractile in Scenario (3) leads to a further increase in profits. In the single-stage model, profits under no sharing are increased by 3,356-3,220=136, and profits under sharing are increased by 3,416-3,328=88 units. Much more pronounced increases are recorded in the two-stage model, where profits under no sharing increase by 3,356-2,942=414, and profits under sharing increase by 3,367-3,063=394 units. This highlights that especially in the two-stage model (when stocking decisions become more complex), a lack of sufficient inventory affects profitability much more than the insufficient requests and transshipments. Finally, if optimal order quantities could be induced by changing the transfer price, profits would only increase slightly – by 3,432-3,416=16 units in the single-stage case, and by 3,367-3,367=0 units in the two-stage case. In terms of effect size, problems in creating sufficient inventory and ensuring that the requests and transshipments are made optimally appear as much more important to solve than getting decision makers to adequately react to a transfer price. We illustrate these effects in the two waterfall diagrams in Figure 2.

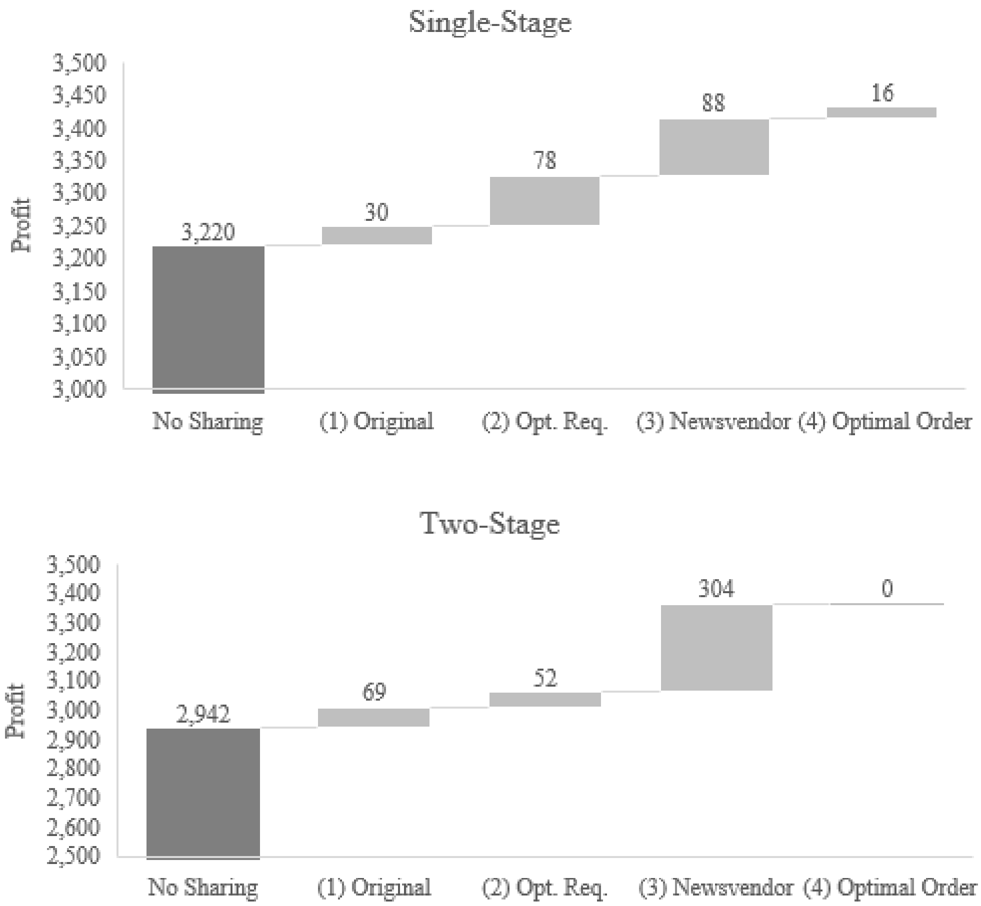

Figure 2: Decision Makers' Profitability in an Inventory Sharing System

Discussion

Our analysis in Study 1 indicates that without careful design, inventory sharing systems may be ineffective from a behavioral perspective. In our experiments, participants in the sharing condition failed to benefit from inventory sharing: Their performance was essentially equivalent to (if not worse than) performance in the no-sharing condition. One key reason is that players do not order enough initially, causing an overall inventory shortage in the system at the outset. The transfer price, which has been analyzed in the modeling literature as a coordinating tool for inventory sharing systems, appeared ineffective at inducing higher order quantities. Further, players were reluctant to request inventory from their peers, but, somewhat surprisingly, were more than happy to share their inventory if requested to do so.

Our finding that transfer prices do not influence order quantities is somewhat at odds with some other ongoing research. Two studies on this topic both find that inventory orders and transfer prices correlate (Bostian et al. 2012, Chen and Li 2016). A key difference is that in our study, transfer prices are exogenously given; in these other papers, transfer prices were endogenously set by players. Such an endogenous setting makes the inventory order the only response to the transfer price set by the other player – almost like a wholesale price influencing order quantities. Many large firms such as Caterpillar and GM set transfer prices centrally but give their independent dealers decision rights on requests/transshipments – our results speak about such settings.

Our analysis also uncovered an interesting if somewhat unrelated behavioral phenomenon: If the demand realization was split into two stages, inventory quantities declined (even if no inventory sharing is allowed). We included the two-stage version of our model for theoretical reasons related to inventory sharing, yet one very relevant insight we developed from including this factor in our experiment is that two-stage demand revelation per se leads to a sizable drop in initial inventory orders. This is an important phenomenon because most inventory systems will see demand split into multiple stages, particularly if inventory review and order cycles differ. The reasons for this phenomenon may be rooted either in anchoring effects (splitting demand means that participants see lower numbers) or in the concept of subadditivity, i.e., the notion that decision makers have difficulties summing up the estimated probabilities of constituent events (see e.g., Bearden et al. 2007). Investigating this behavioral phenomenon further is beyond the scope of our research. However, we do believe that future research should examine this phenomenon more closely.

5. Study 2 – Demand and Supply Side Thinking

One key conceptual insight in our theoretical development is the idea that inventory sharing has two perspectives: a demand side and a supply side. As we hypothesized in Study 1, a player's relative focus on supply-side thinking compared with demand-side thinking would lead to lower initial orders if inventory sharing were allowed. This study was designed to test this assertion more directly by separating roles among participants. This separation allowed us to examine whether players who act as recipients lower their initial inventory more than players who act as sources.

Experimental Design

The experimental task in Study 2 was like the task in Study 1. The key treatment we introduced in Study 2 was fixed roles. Although in Study 1, participants could be both a source and a recipient, in Study 2, we fixed these roles for each participant: Participants were randomly assigned as either a source or a recipient throughout the experiment; moreover, in each period, source players were randomly matched to recipient players in each channel. We used the single-stage model throughout because our previous studies established that, despite the greater realism and more pointed results of the two-stage model, the single-stage model is more readily understood and does not lead to excessively low inventory levels. We also varied the transfer price as an experimental factor. Note that because our variable of interest now becomes the role assignment within this treatment, more subjects were recruited within the fixed role treatment than in the Study 1 treatments. We collected data from 84 subjects across both transfer price treatments (44 in the high transfer price, 40 in the low transfer price). One subject in the high transfer price treatment left the experiment before completion because of a medical emergency and her data was removed from the analysis. The resulting data set has 2,490 observations from 83 subjects.

Analysis

Our data analysis strategy followed Model (1) from Study 1. Results indicate that source players stock higher on average than recipient players (b=8.19, p=0.10). In addition, the transfer price had no effect on source players (b=0.09, p=0.99). Recipient players, however, reacted to the transfer price. Order quantities for recipient players in the low transfer price condition (=188, s.e.=4.58) were less than those for recipient players in the high transfer price condition (=200, s.e.=6.04). This difference was only weakly significant (b=12.57, $p \leq 0.10$) but provided some evidence that the transfer price does affect recipients much more than source players. In other words, players recognized that their emergency supplies through transshipments were becoming more expensive as the transfer price increased, although they failed to realize that a decrease in the transfer price implied more requests for transshipment.

These differences became more pronounced when we compared order quantities of source and recipient players in the single-stage treatment from Study 1. We give an overview of estimated marginal means across treatments in Figure 3. Note that we labelled the data from Study 1 as *Recipient & Source* because participants in the single-stage treatment there could request inventory sharing and respond to requests from the other player. When subjects had both roles, the effect of the transfer price was negative, but not significant (b=-7.59, p=0.14). In contrast, when subjects played only as recipients, the transfer price seemed to increase order quantities (b=12.57, $p \leq 0.10$). When they played as sources, the transfer price again had no effect on behavior (b=0.09, p=0.99). When contrasting these three different effect sizes, we saw that the effect size of the transfer price was different if participants played just as recipients ($p < 0.05$, compared with playing both roles), but not different if they played as sources (p=0.36, compared with playing both roles).

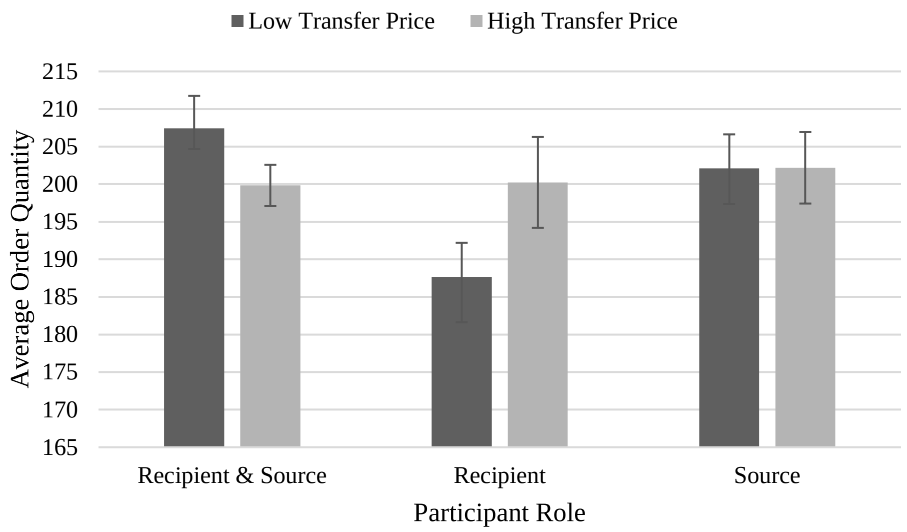

Figure 3: Average Order Quantities across Roles and Transfer Prices

Discussion

Results from our analysis show that participants overweight the supply side of inventory sharing and largely ignore the demand side. Recipients who only care about the supply side of sharing (i.e., recipients), stock less than demand-side players (i.e., sources) and clearly react to a transfer price by increasing their order quantities as the transfer price increases to become less reliant on the source players under such conditions. In theory, source players should have reacted similarly to the transfer price, but instead they did not react to it at all. This provides evidence that the rationale underlying Hypothesis 1 was valid. It also provides evidence that the transfer price does impact behavior – but more on the recipient side than on the side of sources.

6. Study 3 – Improving Inventory Sharing Effectiveness

Our previous two studies were designed to understand whether inventory sharing systems can be effective from a behavioral perspective and why the effectiveness of such systems may break down. Study 3 examined how we can design a system to create more effective inventory sharing. Specifically, we examined the role of framing, transparency, and decision support in improving inventory sharing effectiveness.

Experimental Design

The single-stage treatment from Study 1 was the baseline treatment underlying all factor variations in Study 3. We examined three important variations of this treatment as to their effectiveness in improving inventory sharing systems. The first variation is the *demand-side framing* treatment, which is designed to counter the players' overweighting of the supply-side impact of inventory sharing and help them to better realize the demand-side impact as well. Our base treatment used a *supply-side framing* because we asked players to request supplemental inventory from their partners. In contrast, in the demand-side framing treatment, we asked players to offer their inventory for sale to a recipient, who could then decide how much to buy from the units offered. We expected this to help emphasize the demand side of inventory sharing and potentially lead to higher order quantities.

The second variation we studied related to reducing transparency in our experiment. We observed in Study 1 that under-requesting was one of the reasons for ineffective sharing, and one of the reasons for under-requesting was that players knew of other players' inventory and this knowledge led to concerns about requesting unreasonable amounts. Therefore, although the base setting allows all players to see orders and demand realizations for both players, in our no transparency treatment, we did not inform players about the orders and demand realizations of the other players. This treatment allowed us to examine whether removing inventory transparency from the system would increase the amounts requested.

Finally, our decision support treatment expands on the no-transparency treatment by not only reducing visibility across the supply chain, but also offers players a recommended order quantity that they accept or revise. This approach is similar to the decision support treatment in Lee and Siemsen (2016) and would potentially increase the low inventory that players stock in the system, a key element for the inventory sharing system to be effective. The suggested order quantities are newsvendor quantities derived from forecasts made by a single exponential smoothing model with a very low smoothing parameter and optimal service levels. These order quantities are not the *rational* equilibrium order quantities under the inventory sharing system presented in Section 2, but instead focus exclusively on a participant's focal market. The potential to share inventory was not considered because the simple newsvendor quantity does not change with the transfer price, making it easier to interpret for the players. All these treatments were tested for both transfer prices, resulting in six additional experimental conditions. The resulting data set contained 4,260 decisions from 142 participants.

Analysis

Our analysis here again followed Models (1) and (2). We first estimated a mixed-effects model that estimated the impact of all treatments under different transfer prices on order quantities. Results from the analysis are shown in Figure 4. Our demand-side framing did not sufficiently influence participants' behavior; inventory quantities in this treatment did not differ significantly from our base treatment in Study 1 ($p=0.32$). While on average, the transparency treatment also failed to influence order quantities ($p=0.71$), we note that the transfer price in this treatment starts to have an impact on ordering behavior – the higher transfer price now was clearly associated with higher order quantities ($b=8.96$, $p\leq0.05$), and the contrast to the same effect in the base treatment showed a significant difference ($\chi^2=6.07$, $p\leq0.05$). Finally, added decision support increased the order quantity overall ($b=10.52$, $p\leq0.01$), and the transfer price retained an influence on order quantities under decision support without transparency ($b=6.81$, $p\leq0.01$).

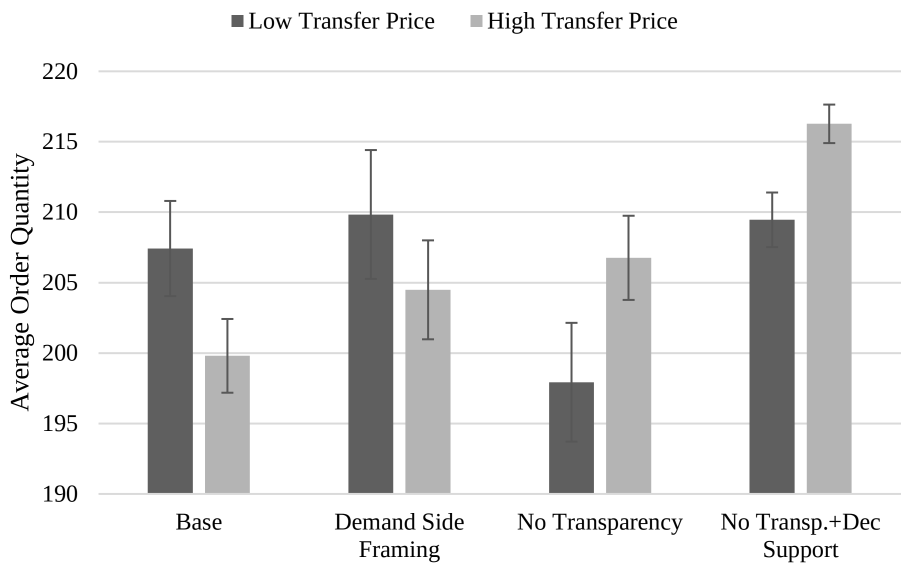

Figure 4: Average Order Quantities Across Treatments and Transfer Prices

To further examine whether the no-transparency treatment encouraged more inventory sharing, we next analyzed inventory requests and transshipments in this treatment as well as in the decision support treatment, compared with the base treatment. The average inventory sharing request in the base treatment (i.e. with transparency) was lower (21/20 units in low/high transfer price settings) than that in the no-transparency treatment (32/28 units in low/high transfer price settings). This indicated that removing transparency indeed leads to higher requests. No-transparency with decision support again leads to a slight decrease in requests (26/21 units in low/high transfer price settings), which may not be surprising, given that more inventory was available in this treatment (hence fewer requests were needed).

To test these observations more formally, we estimated a hurdle model in Stata. Independent variables for both the selection and main equations were limited to treatment variables and interaction terms. Situations in which the recipient had sufficient inventory to meet demand without making a request were excluded from the analysis. The resulting analysis was based on 2,046 decisions made by 148 individuals. Standard errors were cluster corrected to account for the nested structure of the data. The results confirmed that participants in the no-transparency treatment were more likely to both request inventory ($b=0.89$, $p\leq 0.10$) and request more inventory in their requests ($b=26.37$, $p\leq 0.05$). Decision support has no further effect on inventory requests, except to lower the number of units requested under the high transfer price treatment ($b=-36.06$, $p<0.10$).

These increased requests led to a slight increase in transshipments. Across all observations in the Base treatment, the average transshipment was 8.96 units. This number increased to 9.50 units in the no-transparency treatment and to 10.37 units in the decision support treatment. A similar hurdle model confirmed that this increase in transshipments was mostly because of an increased propensity to respond to recipient requests under no-transparency ($b=0.26$, $p<0.05$).

We next examined the implications for profitability under these modified inventory sharing regimes. Average profits in our newsvendor no-sharing setting from Study 1 were about 3,220, and average profits under the single-stage base setting were 3,352/3,155 for low/high transfer prices. Under no-transparency, these profits changed to 3,293/3,321 for low/high transfer prices – an average marginal effect across both transfer prices that is positive but not statistically significant (b=54, p=0.32). However, if we added decision support to the no-transparency setting, average profits increased further to 3,375/3,333 for low/high transfer prices, which now corresponded to a significant combined marginal effect of no-transparency and decision support (b=101, p<0.05). Thus, it appears that an inventory sharing system with non-transparency (to encourage higher requests) and decision support (to counter the lower inventory stocks resulting from allowing participants to share)[2] would reap the benefits of inventory sharing.

Discussion

Study 3 revealed several important insights. First, it is interesting that our framing manipulation is not strong enough to change behavior; despite pushing participants to think about the demand side of inventory sharing, we still did not see a change in their initial ordering behavior. To some extent, this is in line with our rationale for Hypothesis 1: Players still thought inventory sharing should "reduce" their initial stock (without realizing they had too much reduction). Further, players still valued revenues generated in their own market more than those from the secondary market through transshipment, despite making this market opportunity more salient. Second, reducing transparency enabled the transfer price to become effective; without transparency, participants requested more inventory and increased their initial orders with an increased transfer price. This suggests that reducing transparency in inventory sharing systems seems a good approach to make these systems more effective, especially if combined with decision support. These two combined effects lead to not only higher inventory availability in the system, but also decision makers' greater responsiveness to the transfer price.

7. Conclusion

The analytical modeling literature on inventory sharing has been abundant because of its potential to simultaneously reduce inventory and increase service levels (Tagaras 1989 and Evers 2001). Using three behavioral studies, we demonstrate that without careful design, inventory sharing systems may not be particularly effective in achieving their claimed benefit of improved profitability for participants. Key reasons include players' over-reduction in initial stock when provided inventory sharing opportunities leading to a scarcity of inventory in the system, an unresponsiveness to the transfer price, as well as a reluctance to request adequate amounts of inventory from supply chain partners. One apparent underlying mechanism behind these barriers is that inventory sharing seems to be perceived as more of a supply-side phenomenon – i.e., other players become a source of emergency supply – rather than a demand-side phenomenon *as well* – i.e., other

[2] We also ran a small treatment in which we offered decision support under no sharing (N=20). Profits under decision support without sharing are significantly lower than profits under decision support with sharing (b=-198, p≤0.05).

players can also become a secondary market to sell inventory through transshipments. This mechanism, together with the perception that inventory sharing should lead to lower inventory and higher service levels at the same time, causes players to order too little initial inventory to make inventory sharing effective. We also provide evidence that by reducing inventory transparency and offering decision support, inventory sharing can become effective and increase profitability for participants. It is clear from our analysis that from a behavioral perspective, designing an effective inventory sharing system requires careful consideration, with the transfer price being a relatively minor lever in terms of making the system more effective.

Our work has led to some important theoretical and managerial insights: First, the transfer price may not be a very effective tool to coordinate the supply chain in this context. In many of our settings, decision makers appeared to ignore the transfer price in their stocking decisions. This irresponsiveness to transfer price was only mitigated when we reduced system transparency. This finding is in stark contrast to the modeling literature on inventory sharing systems that has largely focused on exploring the transfer price as a coordinating tool to increase system performance.

Second, although the modeling literature emphasizes that to take advantage of risk pooling, initial stock levels should be revised after inventory sharing opportunities are provided, preventing decision makers from dramatic reductions in their inventory appears more important from a behavioral perspective. Even just sticking to the original stock levels (used in a no-sharing system) in a sharing system can be more profitable than an unguided revision of inventory that can quickly destroy the purpose of inventory sharing.

Third, the unwillingness of participants to share their inventory does not appear to be a major concern for the breakdown of inventory systems. The greater culprits appear to be overall insufficient inventory in the system coupled with players' reluctance to request adequate inventory from others through transshipments. If a more centralized form of organization were imposed, then guiding the initial ordering process and providing guidance to encourage reasonable requests appear equally important in providing people with the proper incentives to share their inventory.

As one of the first studies to consider the important inventory sharing problems from a behavioral perspective, our experiments can serve as a basis for further experimental work in this area, and several variations of our experiments could be explored. For example, it would be interesting to examine whether it is better to direct extra customers to another player instead of shipping extra supplies. This may further emphasize the demand side of inventory sharing and may also increase the salience of revenue generated from this additional demand. Such a setting could also serve as a basis to explore the design of omni-channel retailing systems. We hope that our research will lead to more behavioral studies in this area.

References

Anupindi, R., Y. Bassok, E. Zemel. 2001. A General Framework for the Study of Decentralized Distribution Systems. *Manufacturing & Service Operations Management* **3** 349–368.

Baker, K. R., M. J. Magazine, H. L. W. Nuttle. 1986. The Effect of Commonality on Safety Stock in a Simple Inventory Model. *Management Science* **32** 982–988.

Bearden, J. N., T. S. Wallsten, C. R. Fox. 2007. Contrasting stochastic and support theory accounts of subadditivity. *Journal of Mathematical Psychology* **51** 229–241.

Blau, P. M. 1964. Exchange and Power in Social Life. Wiley, New York, NY.

Bostian, A. A., C. A. Holt, S. Jain, K. Ramdas. 2012. Is Transshipment a Behaviorally-Robust Risk-Pooling Strategy? Working Paper.

Chen, K.-Y., S. Li. 2016. Coordinating transfer price in inventory sharing? Evidence from experiments. Working Paper.

Dong, L., N. Rudi. 2004. Who Benefits from Transshipment? Exogenous vs. Endogenous Wholesale Prices. *Management Science* **50** 645–657.

Evers, P. T. 2001. Heuristics for Assessing Emergency Transshipments. *European Journal of Operational Research* **129** 311–316.

Falk, A., U. Fischbacher. 2006. A theory of reciprocity. *Games and Economic Behavior* **54** 293–315.

Fischbacher, U. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10** 171–178.

Gallagher, E. 2002. One stop co-op, discover how Johnstone Supply, a lead wholesale distribution company, uses a cooperative business model and succeeds. *US Business Review* 1–6.

Hyndman, K., M. Embrey. 2017. The Econometrics of ExperimentsPage *in* E. Katok, S. Leider, and K. Donohue, editors. *Handbook of Behavioral Operations Management*.

Lee, F. 1997. When the Going Gets Tough, Do the Tough Ask for Help? Help Seeking and Power Motivation in Organizations. *Organizational Behavior and Human Decision Processes* **72** 336–63.

Lee, Y. S., E. Siemsen. 2016. Task Decomposition and Newsvendor Decision Making. *Management Science, forthcoming*.

Maycroft, N. 2009. Not moving things along: hoarding, clutter and other ambiguous matter. *Journal of Consumer Behaviour* **364** 354–364.

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor Networks: Inventory Management and Capacity Investment with Discretionary Activities. *Manufacturing & Service Operations Management* **4** 313–335.

Morewedge, C. K., C. E. Giblin. 2015. Explanations of the endowment effect: An integrative review. *Trends in Cognitive Sciences* **19** 339–348.

Narus, J. A., J. C. Anderson. 1996. Rethinking distribution. *Harvard Business Review* **74** 112–120.

Norton, M. I., D. Mochon, D. Ariely. 2012. The IKEA effect: When labor leads to love.

Robinson, L. 1990. Optimal and Approximate Policies in Multiperiod, Multilocation Inventory Models with Transshipments. *Operations Research* **38** 278–295.

Rong, Y., L. V. Snyder, Y. Sun. 2010. Inventory sharing under decentralized preventive transshipments. *Naval Research Logistics* **57** 540–562.

Rudi, N., S. Kapur, D. F. Pyke. 2001. A Two-Location Inventory Model with Transshipment and Local Decision Making. *Management Science* **47** 1668–1680.

Schweitzer, M. E., G. P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: experimental evidence. *Management Science* **46** 404–420.

Sterman, J. D., G. Dogan. 2015a. "I'm not hoarding, I'm just stocking up before the hoarders get here.": Behavioral causes of phantom ordering in supply chains. *Journal of Operations Management* **39–40** 6–22.

Sterman, J. D., G. Dogan. 2015b. "I'm not hoarding, I'm just stocking up before the hoarders get here." *Journal of Operations Management* **39–40** 6–22.

Tagaras, G. 1989. Effects of Pooling on the Optimization and Service Levels of Two-location Inventory Systems. *IIE Transactions* **21** 250–257.

Villa, S., E. Katok. 2016. Transshipments in supply chains: Beyond the analytical models. Working Paper.

Yan, X., H. Zhao. 2011. Technical Note—Decentralized Inventory Sharing with Asymmetric Information. *Operations Research* **59** 1528–1538.

Zhang, Y., E. Siemsen. 2016. A meta-analysis of newsvendor experiments: Revisiting the pull-to-center asymmetry. Working Paper.

Zhao, H., A. Bisi. 2010. Optimal operating policies in a commodity trading market with the manufacturer's presence. *Naval Research Logistics* **57** 127–148.

Zhao, H., V. Deshpande, J. K. Ryan. 2005. Inventory Sharing and Rationing in Decentralized Dealer Networks. *Management Science* **51** 531–547.

Zhao, H., V. Deshpande, J. K. Ryan. 2006. Emergency transshipment in decentralized dealer networks: When to send and accept transshipment requests. *Naval Research Logistics* **53** 547–567.