University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

Supply Chain Management and Analytics Publications

Business, College of

2020

Inventory Sharing and Demand-Side Underweighting

Hui Zhao Pennsylvania State University, huz10@psu.edu

Liang Xu University of Nebraska at Lincoln, liang.xu@unl.edu

Enno Siemsen University of Wisconsin-Madison, esiemsen@wisc.edu

Follow this and additional works at: https://digitalcommons.unl.edu/supplychain

Part of the Business Administration, Management, and Operations Commons, Management Information Systems Commons, Management Sciences and Quantitative Methods Commons, Operations and Supply Chain Management Commons, and the Technology and Innovation Commons

Zhao, Hui; Xu, Liang; and Siemsen, Enno, "Inventory Sharing and Demand-Side Underweighting" (2020). *Supply Chain Management and Analytics Publications*. 17. https://digitalcommons.unl.edu/supplychain/17

This Article is brought to you for free and open access by the Business, College of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Supply Chain Management and Analytics Publications by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.



Inventory Sharing and Demand-Side Underweighting

Hui Zhao,¹ Liang Xu,² and Enno Siemsen³

1 Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802

2 College of Business, University of Nebraska–Lincoln, Lincoln, Nebraska 68588 3 Wisconsin School of Business, University of Wisconsin–Madison, Madison,

Wisconsin 53706

Correspondence — huz10@psu.edu (HZ); liang.xu@unl.edu (LX); esiemsen@wisc.edu (ES)

ORCID

Hui Zhao http://orcid.org/0000-0002-8995-6644 Liang Xu http://orcid.org/0000-0003-2218-9177 Enno Siemsen http://orcid.org/0000-0003-2033-3833

Abstract

- **Problem definition:** Transshipment/inventory sharing has been used in practice because of its risk-pooling potential. However, human decision makers play a critical role in making inventory decisions in an inventory sharing system, which may affect its benefits. We investigate whether the opportunity to transship inventory influences decision makers' inventory decisions and whether, as a result, the intended risk-pooling benefits materialize.
- *Academic/practical relevance:* Previous research in transshipment, which is focused on finding optimal stocking and sharing decisions, assumes rational decision making without any systematic bias. As one of the first to study inventory sharing from a behavioral perspective, we demonstrate a persistent stocking-decision bias relevant for inventory sharing systems.
- *Methodology:* We develop a behavioral model of a multilocation inventory system with transshipments. Using four behavioral studies, we identify, test, estimate, and
- Published as: Hui Zhao, Liang Xu, Enno Siemsen (2020) Inventory Sharing and Demand-Side Underweighting. *Manufacturing & Service Operations Management* (Articles in Advance), 20 May 2020. https://doi.org/10.1287/msom.2020.0875

Copyright © 2020, INFORMS. Used by permission.

Submitted May 21, 2018; revised August 23, 2019, December 6, 2019, & January 12, 2020; accepted January 15, 2020; published May 20, 2020.

mitigate a demand-side underweighting bias: although inventory sharing brings both a supply-side benefit and a demand-side benefit, players underestimate the latter. We show analytically that such bias leads to underordering. We also explore whether reframing the inventory sharing decision reduces this bias.

- **Results:** Our results show that subjects persistently reduce their order quantities when transshipments are allowed. This underordering, which persists even when a decision-support system suggests optimal quantities, causes insufficient inventory in the system, in turn reducing the risk-pooling benefits of inventory sharing. Underordering is evidently caused by an underweighting bias; although players correctly estimate the supply-side potential from transshipment, they only estimate 20% of the demand-side potential.
- *Managerial implications:* Although inventory sharing can profitably reduce inventory, too much underordering undermines its intended risk-pooling benefits. The demand-side benefits of transshipment need to be emphasized when implementing inventory sharing systems.
- **Keywords:** inventory sharing, behavioral operations, demand-side underweighting bias, structural estimation

1. Introduction

Inventory sharing, also called *transshipment*, is used in practice because of its potential to increase profitability and service levels through risk pooling. Example applications stem from different industries, such as automotive and machine tools (Narus and Anderson 1996), spare parts (Zhao et al. 2005, 2006), fashion (Dong and Rudi 2004), wholesale/retail (Gallagher 2002), and many online trading platforms (Zhao and Bisi 2010). The notion of a sharing economy has made sharing inventory or capacity more prevalent in everyday life as well.

The benefits of inventory sharing, for example, improving service levels, reducing system inventory, and improving profitability, ultimately depend on inventory stocking and sharing decisions. If these decisions are made optimally, the benefits of inventory sharing have been demonstrated through many analytical and simulation studies (e.g., Tagaras 1989, Evers 2001). There is, however, little academic research that examines how managers make such decisions. To gain a better understanding of how these decisions are made in practice, we conducted several interviews with managers of firms from different industries and further administered a survey in cooperation with the Association for Supply Chain Management on these topics. Fifty-four firms responded to the part of the survey focused on transshipments.¹ In the survey, we asked how inventory orders are created—whether these orders are automated or made through human judgment. Among the 54 respondents, none indicated that inventory orders are entirely made by software without any involvement of human judgment, three indicated that inventory orders are entirely made by human judgment, 35 indicated that inventory orders are based on software and then adjusted by human judgment, and another 10 indicated that orders represent averages between judgments and an automated order. The responses to this question highlight that inventory decisions in practice are mostly a combination of human judgment and software. This resonates with many documented cases where human planners regularly intervene

to adjust automated inventory orders (e.g., van Donselaar et al. 2010). Regarding transshipments, which occur if, for example, "one distribution center runs low on a certain stock and obtains extra inventory for replenishment from another distribution center instead of a manufacturing site," we asked respondents how often they transship inventory within different stages in the supply chain in the past 12 months. Responses for each stage could vary on a scale from never to extremely often. Respondents could also indicate that they do not know or that this supply chain stage does not apply to them. In general, inventory sharing does not occur frequently among the firms in our sample. For the firms indicating that some transshipments happen within their supply chains (n = 34), we asked what usually triggers these transshipments. Not a single respondent indicated that transshipments occur because of suggestions by inventory software. Eleven respondents indicated that transshipments are triggered by centralized inventory planners, and seven respondents indicated that these transshipments are triggered by managers at the site. Eleven firms said that transshipments are triggered by an immediate stockout for a critical item that needs to be addressed; this is a sign that transshipments are often less proactive but rather reactive to address "firefighting" needs. We asked firms where transshipments never or only rarely happen (n = 16) why this is the case. The most frequently chosen answers were a lack of strategy for transshipments (43%) and a sense of the cost of transshipping inventory generally outweighing the benefits (42%). People also indicated that their software does not consider transshipments (36%).

These data convey that although algorithms (e.g., inventory management software) are widely used in industry to support inventory decision making, inventory sharing opportunities are rarely considered by these algorithms. Consequently, store, distribution center, and manufacturing site managers must figure out the implications of inventory sharing on their own. How they make such decisions is a critical aspect influencing the profitability of the system. Underreacting to the inventory sharing opportunities would leave benefits of risk pooling on the table, whereas overreacting to them could lead to a dearth of inventory in the system. This motivates us to study the inventory sharing problem from a behavioral perspective.

Many experiments have documented biases in inventory decisions in the context of both single-location (e.g., Schweitzer and Cachon 2000) and multilocation newsvendor problems (e.g., Ho et al. 2010). Therefore, we aim to investigate two research questions: (1) Does the opportunity to transship inventory between decision makers alter their ordering behavior? And (2) if yes, what is the implication of such changes to the expected risk-pooling benefits?

To tackle these research questions, we examine the most studied inventory sharing model, that described by Rudi et al. (2001), from a behavioral perspective. Indeed, in addition to the familiar pull-to-center bias observed in a multilocation setting *without* inventory sharing, we find some persistent bias in human decision makers' stocking decisions when provided with inventory sharing opportunities. Note that transshipments bring both supply-side benefits (an opportunity to obtain extra supplies from the other player) and demand-side benefits (an opportunity to sell extra products to the other player). Although in theoretical models these two benefits are equally considered in the calculation of the optimal ordering quantities, we find that decision makers in our experiments consistently underweight the demand-side benefits. Such a bias leads to the systematic underordering of inventory because decision makers fail to perceive the extra revenue they may generate by selling to the other player. As a result, inventory sharing may be less beneficial in practice than in theory if there is too much underordering in the system. We further explore the potential root causes as well as potential intervention strategies to counter such a demand-side underweighting bias.

Our paper makes several important contributions to the field of inventory sharing as well as behavioral operations management. As one of the first papers to study the behavioral aspects of inventory sharing, we

5

identified, modeled, and tested a systematic demand-side underweighting bias leading to consistent underordering behavior. Specifically, first, we developed a new behavioral model of a multilocation inventory system with transshipments. This model assumes that in addition to the well-known psychological costs for underage and overage that are prevalent in the newsvendor setting, a decision maker also differs in his or her perceptions of the supply-side and demand-side benefits of inventory sharing. We show analytically that with demand-side underweighting, a decision maker will, in turn, underorder. Second, we designed and conducted experiments in the laboratory to examine order quantities, profits, and service-level outcomes in this context. Our experiments demonstrate that inventory sharing opportunities indeed alter our participants' ordering behavior-they persistently underorder when given the opportunity to transship inventory. This is true after the pull-to-center effect is taken into consideration. Such underordering persists even when decision makers are provided with a decision-support system that suggests inventory order quantities but permits their revision (a typical practice in industry). This underordering behavior causes a shortage of inventory in the system, which, in turn, reduces the potential for the players to benefit from risk pooling through inventory sharing. Third, using data from our experiments, we estimated the parameters in the proposed behavioral model through structural estimation and showed that the model captures inventory decisions in our context better than a model that incorporates only the psychological cost of underage and overage. The parameter estimates indicate that although the subjects' evaluations of the availability of extra supply through transshipments are close to the correct level, their evaluations of the demand-side potential are far off, representing only 20% of the true demand-side potential through transshipments. Fourth, we identified several potential root causes for such demand-side underweighting, including the desire for control, seeing reluctance in requesting in others but not in themselves, and valuing own markets more than others' markets. Finally, we explored several intervention strategies, such as reframing inventory sharing, automated transshipments, and advanced decision-support tools, to mitigate this demand-side underweighting bias.

From a practical point of view, our paper provides important guidance to decision makers of inventory sharing systems. Although inventory sharing has the potential to reduce inventory in the system because of risk pooling, too much underordering can undermine the intended benefits. In addition, managers should not underestimate the demandside benefits from inventory sharing because the demand-side serves as a secondary market for them and brings additional profit. Introducing inventory sharing systems should go together with emphasizing the extra revenue-generating potential of such a practice.

It is worthwhile to mention that our research is also relevant in the context of omnichannel retailing. Over the past decade, most traditional offline retailers have built an online presence. How much to stock for each channel and how to best fulfill orders originating from either channel remain among the top challenges for omnichannel retailers. Traditionally, many retailers held separate stocks for online and offline stores. However, because of the benefits of risk pooling, many of them are pooling such inventory, for example, using in-store inventory for online order fulfillment and vice versa (Alishah et al. 2017). Such an arrangement can lead to in-store inventory being reserved for either online or offline customers, and decisions to share inventory between these two reserved piles are often negotiated between the managers of online channels and stores. Indeed, one of the biggest fashion retailers in the nation that we interviewed regularly transships inventory between online and offline stockpiles and faces the challenge of making appropriate stocking decisions for these two channels. Our research suggests that channel managers may underestimate the value of using this system to sell in the other channel and instead primarily view this system as a way of "dipping" into the inventory reserved for the other channel.

We proceed with our paper as follows. In Section 2, we review the related literature. In Section 3, we briefly explain the rational theory underlying our context, serving as a benchmark. In Section 4, we present a corresponding behavioral model in which we consider the proposed demand-side underweighting bias; we also present experimental evidence for demand-side underweighting and its impact on the profits and service levels of the system and estimate the underweighting bias from our behavioral model. In Section 5, we explore the root causes for demand-side underweighting and propose interventions to mitigate this bias. In Section 6, we discuss insights from our findings and conclude our research.

2. Literature Review

The inventory sharing/transshipment problem has been studied extensively because the prospect of transshipments has the potential to both improve service levels and reduce system inventory. However, most of this research is focused on analytical modeling, characterizing optimal stocking, and sharing decisions. Early research assumed a centralized power that could decide how stockpiles were made available to different players. Although such a central authority may exist, supply chains are often more decentralized. There are ample examples of inventory sharing among *independent* players (e.g., Zhao et al. 2005). Correspondingly, a large amount of analytical modeling research has examined these decentralized players' optimal/equilibrium stocking and sharing decisions and how appropriate incentives can induce players to choose the optimal solution. Rudi et al. (2001), referred to as RKP hereafter, studied a single-period model and are among the first to study inventory sharing in such a *decentralized* system. After developing the equilibrium order quantities, the authors explore coordination of such a system through a linear transfer price. Many other papers have also studied decentralized inventory sharing systems in various settings, for example, Anupindi et al. (2001), Zhao et al. (2005), Rong et al. (2010), Shao et al. (2011), Yan and Zhao (2011).

Although this literature focuses on the optimal/equilibrium behavior to extract the risk-pooling benefit of inventory sharing, it also assumes perfectly rational decision makers who do not exhibit systematic biases. Behavioral experiments have generally challenged the assumption of perfect rationality in stocking decisions, both in single-location (e.g., Schweitzer and Cachon 2000) and in multilocation newsvendor problems (e.g., Ho et al. 2010). Our research investigates the potential bias in order quantity decisions in an inventory sharing system and its impact on the performance of such a system.

Although our research is among the first to examine the behavioral aspects of inventory sharing, some other works in progress are also examining inventory sharing systems. For example, Bostian et al. (2012) examine inventory orders when the system automatically makes transshipment decisions for the players and study whether decision makers in the experiment benefit from inventory sharing by estimating structural models to establish counterfactuals. In contrast, we examine this

question more explicitly through randomized treatments in which we allow 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 of transshipments: Villa and Katok (2018) focus on negotiation of the transfer price, and Chen and Li (2020) focus on the sequence in setting the transfer price and making stocking decisions. In our research, we treat the transfer price as exogenous in order to separate the two decisions (stocking and transfer prices) and examine how decision makers react to the transfer price in their ordering decisions.

3. Rational Theory

In this section, we briefly discuss the equilibrium solutions for rational players in the inventory sharing system, which serve as the normative benchmarks for our behavioral model and experiments, as well as provide basic insights into the system.

3.1. The Inventory Sharing Model

We adopt the basic framework from RKP, the single-period, two-location, decentralized inventory sharing system where transshipments lead to payments between parties through a transfer price. We chose this framework because (1) it is the simplest and most studied decentralized inventory sharing setting and (2) the transshipment decisions in this setting are simple, allowing us to examine decision makers' motivations without having to consider their bounded rationality.

Consider two independent players, one at each distinct location. Each player faces independent random demand D_i (i = 1, 2) in a single period and must place an order Q_i at price c_i at the beginning of the period before observing demand. After demand is realized, each player uses his or her own inventory to satisfy demand. For each unit of demand satisfied, a player receives revenue r_i . Any unsold inventory may be salvaged at the salvage value $s_i < c_i$. The decision maker also incurs a penalty cost 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. Let F_i represent the cumulative distribution function of D_i . Without inventory sharing, this framework corresponds to two independent newsvendors with their optimal order quantities calculated as

$$Q_{i,n}^* = F_i^{-1} \left(\begin{array}{c} v_i - c_i \\ v_i - s_i \end{array} \right)$$

In the inventory sharing model, after observing demand, inventory can be shared if one player has extra units available and the other has a shortage. Specifically, player *j* may request transshipment of a certain amount from player *i*, who may then decide how much, if any, to share with player *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, we set $c_{ij} \ge s_j + \tau_{ij}$ as in RKP to ensure that the source is willing to send (i.e., sending a unit is more profitable than salvaging the unit) and set $c_{ij} \le v_j$ to ensure that the recipient wants the transshipments (i.e., requesting a unit of transshipments is more profitable than losing the sale). To simplify the decisions in the experiments, we assume negligible transportation costs $(\tau_{ij} = 0)$ throughout the paper.

Any inventory sharing problem involves two essential decisions how much to stock initially and how much to share or request. In this model, because inventory sharing occurs at the end of the period after demand is realized, a player (say *i*) should always request the inventory that is needed to fill demand, that is, $(D_i - Q_i)^+$, and share what is left after filling demand, that is, $(Q_i - D_i)^+$. Here we define $x^+ = \max(0, x)$. Therefore, the total number of units shared from *i* to *j* is

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

Given T_{ij} , it can be shown that a unique set of equilibrium order quantities exists for the two players, $(Q_{i,s}^*, Q_{j,s}^*)$, which can be calculated by solving equation (10) in RPK.

3.2. Normative Benchmarks for Our Experiments

In Subsection 3.1, we provide general game-theoretic results. Next, we demonstrate and discuss the specific equilibrium solutions that result from the numerical parameters that we will later use in our experiments. These solutions will provide the intuition our behavioral experiments.

To be consistent with the literature, we adopted most applicable parameters from RKP. Specifically, we used a symmetric system in which



Figure 1. Theoretical Optimal Orders and Expected Profit Under Different Transfer Prices

both players face the same parameters. Hence, we will thereafter drop the subscripts *i* and *j* for the players. The cost parameters are r = 40, c = 20, s = 10, and p = 0. As mentioned earlier, we kept $\tau = 0$ to simplify the decisions in the experiments. For our numerical studies, we selected nine levels of the transfer price in the feasible transfer price interval $c_{ij} \in [s, v]$. For demand, we adopted normal demand as in RKP with $D_i \sim N(200, 70.71)$.² As in RKP, we also used a truncated normal distribution when generating demand realizations to avoid negative demand. Using the theoretical results, we calculated the equilibrium order quantities Q_{s}^* . **Figure 1** presents the equilibrium order quantities and expected profit as c_{ij} increases and compares these with the case in which sharing is prohibited ($Q_n^* = 230$).

The equilibrium order quantity Q_s^* increases as the transfer price increases, with $Q_s^* < Q_n^*$ at the lowest transfer price ($c_{ij} = 10$) and $Q_s^* > Q_n^*$ at the highest transfer price ($c_{ij} = 40$). This result follows directly from the theoretical results in RKP. The basic intuition of this result is that depending on the transfer price, transshipments can be viewed as an additional source of either uncertain supply or uncertain demand. Specifically, at the lowest transfer 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 the salvage value from transshipped inventory). Therefore, to maximize profit, players should order less (than the no-sharing model) to capitalize on this potential supply source. By contrast, when $c_{ij} = v = 40$, transshipments are a good source of demand but not of supply (a player buys a transshipped unit at the selling price, earning zero profit). Hence, players should order more (than the no-sharing model)

to take advantage of this possible extra demand source. As the transfer price increases, the tendency to take advantage of possible extra demand from the other player increases. At the same time, the tendency to take advantage of the possible supply from the other player decreases. Both these tendencies cause players to order more initial inventory as the transfer price increases. Therefore, for rational players, a right transfer price can achieve a desirable stocking level and thus coordinate the system. As Figure 1(b) shows, in our setting, the transfer price of 23 coordinates the system and maximizes the expected profit.

4. Demand-Side Underweighting in Inventory Sharing

Having defined the rational equilibrium order quantities in Section 3, in this section, we will discuss behavioral reasons that may lead to deviations from the rational benchmark, that is, demand-side underweighting (Section 4.1). We then demonstrate such deviations through a series of controlled experiments and examine their implication for the effectiveness of the inventory sharing system (Section 4.2). We further estimate the parameters of our behavioral model and quantify such behavioral bias through another set of experiments (Section 4.3).

4.1. Behavioral Theory

In this section, we propose a behavioral theory that explains why subjects' order decisions may deviate from the rational benchmark discussed in Section 3. Specifically, we propose that players suffer from a demand-side underweighting bias: although they intuitively understand that the opportunity to request transshipments from the other player provides them with an extra source of supply (supply-side benefits), they underestimate the potential for additional sales through transshipment requests from the other player (demand-side benefits). As the preceding section emphasizes, rational equilibrium strategies are based on both supply-side and demand-side benefits, with order quantities being below the newsvendor quantity for low transfer prices (when supply-side benefits are more salient) and above the news vendor quantity for high transfer prices (when demand-side benefits are more salient). Demandside underweighting points to an asymmetry—people focus more on the supply-side benefits than on the demand-side benefits of inventory sharing. To investigate how this asymmetry affects players' initial orders, we developed the following behavioral transshipment model. Recall that a rational player *j* expects the transshipment from player *i* as

$$T_{ii} = \min\left[(D_i - Q_i)^+, (Q_i - D_i)^+\right]$$
(1)

Equation (1) implies that player *j* is aware that the transshipments he or she can obtain from the other player are subject to how much the other player can supply, and player *j* has an unbiased expectation of the extra supply from the other player. But a biased player may not correctly estimate the supply available from the other player—player *j* may expect the transshipments he or she can obtain from player *i* as

$$T_{ij} = \min\left[(D_j - Q_j)^+, \alpha (Q_i - D_j)^+\right]$$
(2)

where $\alpha = 1$ indicates that player *j* has an unbiased estimation of the extra supply available to him or her from the other player, whereas $\alpha > (<)1$ indicates that player *j* will overestimate (underestimate) the inventory supply from the other player. Likewise, a biased player *j* may evaluate the amount he or she can sell to player *i* as

$$T_{ji} = \min\left[(Q_j - D_j)^+, \beta (D_i - Q_i)^+\right]$$
(3)

where $\beta = 1$ indicates that player *j* is unbiased in estimating the amount he or she could sell to the other player, whereas $\beta < (>)1$ indicates that player *j* underestimates (overestimates) the opportunity of selling to the other player. When β_0 , player *j* is totally ignorant of the demand-side opportunity to sell to the other player.

Both players decide their initial orders subject to the biased valuation of the supply/demand-side opportunities characterized by (2) and (3) rather than (1), thus leading to suboptimal ordering. Specifically, Proposition 1 shows the condition under which players exhibit underordering behavior. Recall that Q_s^* denotes the equilibrium order quantity under inventory sharing for rational players; that is, $\alpha = 1$ and $\beta = 1$. Let $Q_s(\alpha, \beta)$ be the equilibrium order quantity for biased players. Proposition 1 characterizes the relationship between $Q_s(\alpha, \beta)$ and Q_s^* .

Proposition 1. $Q_s(\alpha, \beta)$ decreases in α and increases in β . When ($\alpha = 1$ and $\beta < 1$) or ($\alpha > 1$ and $\beta = 1$), players exhibit underordering behavior,

that is, $Q_s(\alpha, \beta) < Q_s^*$.

The proof is in the supplement. Proposition 1 suggests that biased players may lower their initial orders when given the inventory sharing opportunity because of the asymmetry between the perceived supply-side and demand-side benefits either because of (1) overestimating sharing as an extra supply source (i.e., overweighting the supply-side benefit) or because of (2) underestimating sharing as an extra selling opportunity (i.e., underweighting demand-side benefits). In the following sections, we will design and report a series of experiments to explore a player's demand-side underweighting bias in inventory sharing and estimate the value of α and β using data from our experiments.

4.2. Evidence for Demand-Side Underweighting

We designed our first study to examine ordering and transshipment decisions under inventory sharing for products with different profit margins. We know from existing behavioral newsvendor research that decision makers—even without the opportunity to share inventory—place orders with less-than-optimal quantities in high-margin conditions and order more than the optimal quantities in low-margin conditions (the socalled pull-to-center effect; see, e.g., Zhang and Siemsen 2019). Because of the potential asymmetry between the perceived supply-side and demand-side benefits, we believe that participants may in addition experience a demand-side underweighting bias when given the opportunity to share inventory. With a low transfer price, the supply-side benefits of inventory sharing are very salient for decision makers. We thus expect that the opportunity to share inventory will lead to a reduction in initial orders.

Hypothesis 1a. Under inventory sharing with a low transfer price, subjects will lower their orders compared with the case without inventory sharing.

By contrast, with a high transfer price, the supply side of inventory sharing (i.e., getting supplementary inventory from the other player) is not very profitable, whereas the demand side (i.e., selling to the other player through transshipments) is highly profitable. Hence, rational players should increase their order quantity. However, because demand-side underweighting means that this profitable aspect of inventory sharing is not very salient for decision makers, we expect that inventory sharing with a high transfer price has little impact on initial order quantities. In other words, we expect initial orders to be like those in a situation without inventory sharing.

Hypothesis 1b. Under inventory sharing with a high transfer price, subjects will not increase their orders compared with the case without inventory sharing.

Hypotheses 1a and 1b require us to vary inventory sharing and the transfer price as treatments. The hypotheses are independent of the profit margin of the underlying newsvendor problem. Hence, we expect these hypotheses to hold across both high- and low-margin conditions. However, we include profit margin as an experimental factor for robustness.

4.2.1. Experimental Design. Our experimental task builds on a classic newsvendor design. We instructed participants to manage a product in a retail setting. In each period, they place their initial inventory orders before knowing their demand realization. As discussed in Section 3, revenues and costs followed a standard price/cost framing in which the sales price for the high-margin condition was set at \$40, the purchase cost at \$20, and the salvage value at the end of the period at \$10. We manipulated the experimental design through transfer prices at levels of \$15 and \$35, which are below and above the coordinating transfer price (i.e., \$23), respectively. We used these two transfer prices to test whether a higher transfer price induces a higher order quantity than a lower transfer price does (in accordance with Figure 1(a)), which is the premise for coordinating the system through the transfer price. We explicitly informed participants that the actual cost of a lost sale was \$0 to avoid ambiguity about the implications of stockouts. The critical fractile was 66% for the no-sharing newsvendor model. In the corresponding low-margin condition, the sales price was also set at \$40, but we increased the purchase cost to \$30, leading to a critical fractile of 33%. Demand in each period was normally distributed with a mean of 200 and a standard deviation of 70.71 and was truncated at zero. Without inventory sharing, the optimal order quantity is 230 units in the highmargin condition and 170 units in the low-margin condition.

Before starting the experiment, we provided participants with

Condition	Sharing	Critical fractile	Optimal order	Transfer price	Participants	Order quantity
С1	No	66%	230	_	26	209.74
С2	No	33%	170	_	24	194.07
СЗ	Yes	66%	212	\$15	26	197.79
<i>C</i> 4	Yes	66%	237	\$35	24	206.89
<i>C</i> 5	Yes	33%	162	\$15	26	185.80
С6	Yes	33%	186	\$35	26	187.50

Table 1. Overview of Experiments with/without Inventory Sharing

instructions about the normal distribution, and we depicted the density function of the demand distribution on their screens. Random draws of demand were different for each participant. As the experiment proceeded, these demand draws were plotted in a time-series graph as well as shown in a table. Demand draws across periods were independent, and participants were informed about that. Participants repeated the task for 30 time periods. They were compensated on 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 z-Tree (Fischbacher 2007) and conducted with a standardized subject pool at a large American public university. Participants were mostly undergraduate students at the university. Sessions were scheduled to last for 45 minutes. Participants were randomly assigned to different treatments (e.g., sharing or no sharing) in each session and were required to watch a five-minute video explaining their corresponding experimental task. They also had access to written instructions at any time during the experiment. We used a between-subjects design, varying three factors: (1) inventory sharing, (2) profit margin, and (3) transfer price. Because the transfer price has no meaning under no inventory sharing, this design resulted in six experimental conditions, summarized in Table 1. Summary statistics, including average order quantities and profits, for all experimental conditions in our paper are provided in Table A1 in the supplement.

Participants were grouped into channels of four or six players. Channels were unrestricted in size in no-sharing conditions. Each channel was fixed on one experimental condition. Subjects did not know that they were associated with a channel. We informed them that they would be matched to a random player after each period. This random matching was specific to channels. Under inventory sharing, subjects could request supplementary inventory transshipments from their matched player after observing their own demand. Participants could not see the other player's inventory orders or demand. The other player could then decide to send transshipments to the requester up to (but not exceeding) the number of units requested. In principle, both players could request and send transshipments to each other in the same period. There were no transshipment costs besides the transfer payment.

4.2.2. Analysis. We first examine order quantities as our dependent variable. The model we use to analyze our experimental data has the following structure:

$$Order_{i} = a_{1} + a_{2}C2_{i} + a_{3}C3_{i} + a_{4}C4_{i} + a_{5}C5_{i} + a_{6}C6_{i} + b_{1}D_{1i} + b_{2}D_{2i} + \varepsilon_{i}$$
(4)

Here Order, denotes the average order quantity of participant i across all 30 periods. We code the experimental conditions with dummy variables (e.g., C3 takes the value of one only if the observation is made in condition 3) and include them in our analysis, as Equation (4) demonstrated. The base case is the high-margin, no-sharing case (i.e., C1). Our only control variable in this analysis is the average demand faced by subjects. Demand is by design uncorrelated with treatments and thus does not bias our treatment effect estimates. Despite participants knowing that the underlying demand distribution stays the same across periods, they react strongly to the observed demand signals, thus introducing noise to their observed order decision. Adding this control variable thus creates econometric efficiency benefits (i.e., more precise coefficient estimates). We expect a nonlinearity in the effect of demand; in particular, there should be decreasing effects of observed demand on order quantities because the likelihood of obtaining profit (and thus the chance of positive reinforcement) decreases with higher orders. We estimate a spline regression on our combined experimental data with an estimated knot. Results indicate that this knot is very close to mean demand (b = 200.65, p = 0.90 for difference with 200); hence we fix the spline knot

	All 30 p	periods	Last 15 periods		
Condition	Coef.	SE	Coef.	SE	
<i>C</i> 2: Low margin, no sharing	-7.87	(5.95)	-17.76**	(6.78)	
<i>C</i> 3: High margin, low transfer price	-7.11	(5.82)	-14.89*	(6.63)	
C4: High margin, high transfer price	3.58	(5.91)	-2.05	(6.88)	
C5: Low margin, low transfer price	-18.69**	(5.76)	-26.31**	(6.63)	
<i>C</i> 6: Low margin, high transfer price	-15.20**	(5.81)	-23.89**	(6.66)	
Average demand (≤200)	1.07**	(0.30)	0.50*	(0.03)	
Average demand (>200)	0.64**	(0.22)	0.28	(0.17)	
Constant	-8.51	(58.82)	110.07*	(45.51)	
Ν	152		15	2	
<i>R</i> ²	3	1%	2	3%	

Table 2. Regression Estimates from Experiments with/Without Inventory Sharing

The omitted category is condition 1, that is, the high-margin, no-sharing condition. All coefficients measure differences in order quantities compared with this condition. Coef., coefficients; SE, standard error.

* $p \le 0.05$; ** $p \le 0.01$

at 200 throughout. Accordingly, we include $D_{1i} = \min(D_i, 200)$ and $D_{2i} = (D_i - 200)^+$ as controls.

Although our data are recorded as decisions, our unit of analysis is an individual; that is, we average decisions across time periods for an individual. Variance of our treatments occurs between subjects, which means that the treatment effects are only identified by between-subject variance. Further, this aggregation allows us to reduce the complexity of our model specification because we do not need to consider the details of adequately lagging control variables but can rely on averages instead. Because multiple individuals were nested in a z-Tree channel, we first estimated a random effects model to examine this potential for dependence within our data. This concern about individuals being nested in a channel during the experiment was recently highlighted as a methodological consideration in behavioral operations by Hyndman and Embrey (2018). However, in our case, the standard deviation of the channel random effects accounts for only 3% of the total error variance; we thus resorted to ordinary least squares regression analysis instead. Regression results are summarized in Table 2.

Table 2 suggests that with a low transfer price, subjects lower their orders when inventory sharing is allowed, with a decrease of 10.82 (= 18.69 - 7.87, *p* = 0.06) under a low-margin case and a decrease of 7.11 (*p* = 0.22) under a high-margin case. This is due to the salience of

supply-side benefits under a low transfer price. With a high transfer price, although subjects should increase their orders according to the normative theory, they do not do so, with a decrease of 7.33 (= 15.2 - 7.87, p = 0.21) under the low-margin case and an increase of only 3.58 (p = 0.55) under the high-margin case. This provides evidence that subjects underweight the demand-side benefits under a high transfer price.

To formally test Hypothesis 1a, we estimated a joint contrast between conditions 1 and 3 (for the high-margin case) as well as conditions 2 and 5 (for the low-margin case). Results from this contrast weakly support Hypothesis 1a by suggesting that the mean order quantities in the context of a low transfer price are less than those under the no-sharing setting ($F = 2.49, p \le 0.10$). Descriptively, 73% of subjects in conditions 3 and 5 (low transfer price setting across different margins) order less on average than the corresponding mean order quantities in conditions 1 and 2 (no-sharing cases across different margins). A similar contrast between conditions 1 and 4 as well as conditions 2 and 6 reveals that order quantities with a high transfer price do not exceed the order quantities without sharing (F = 0.98, p = 0.38). Only 36% of subjects in conditions 4 and 6 have average order quantities that exceed their corresponding average order quantities in conditions 1 and 2. These observations support Hypothesis 1b and are consistent with the idea that relative to the supply side, people underweight the demand-side benefit of inventory sharing. To examine whether our results are influenced by participant experience, we repeated our analysis using only the last 15 decision periods. Results from this analysis are also reported in Table 2. The results from our hypotheses tests become stronger in this analysis, providing support to both Hypothesis 1a (F = 3.33, p = 0.04) and Hypothesis 1b (F= 0.45, p = 0.63). Subjects also tend to be less influenced by demand observations in the last 15 periods.

We also examined subjects' transshipment decisions in all sharing conditions. Specifically, we calculated the optimal request for each subject and period (which is defined as *OptRequest* = [*Demand-Order*]⁺) and then established the difference between this optimal request and the actual request made. Similarly, we calculated the optimal send (which is defined as *OptSend* = min[*Request, Leftover*]) and contrasted the optimal send with the actual amount sent. Note that our subjects should have had no difficulties in calculating these optimal request/sends themselves because their shortages and leftovers were displayed on the

	Recipient has					
Request	SI (ordei	nortage r < demand)	No shortage (order ≥ demand)			
Nothing	132	8.2%	1378	94.9%		
Below shortage	249	15.5%	—	_		
Exactly shortage	1,200	74.6%	—	_		
More than shortage	27	1.7%	74	5.1%		
	1,608		1,452			

Table 3. Requests Made by Recipients in the Sharing Setting

screen during the experiment (see screen shots in the supplement). **Table 3** contains an overview of all 3,060 requests made under both highand low-margin sharing treatments.

We differentiate between situations in which the recipient had a shortage and situations in which no shortage existed. It is straightforward for participants to understand that they have enough inventory to meet demand if no shortage exists; indeed, we see that in 94.9% of the cases where there was sufficient inventory to meet demand, participants requested no additional inventory (subjects are allowed to request inventory even without shortages). However, when a shortage existed, it became interesting to examine whether participants requested what they needed. We see that recipients' requests exactly matched their shortages only 74.6% of the time. In the other 23.7% (= 8.2% + 15.5%) of situations, subjects generally requested less (instead of more) than they needed. Because the source is likely to use the recipient's requests as information to estimate expected demand from this secondary market, such underrequesting bias means that the source processes censored information. As a result, the source would generally lower the estimate of the potential sales volume in the secondary market, thus leading to the underweighting of the demand side of inventory sharing.

To examine how the source responds to inventory requests, we focused on the 1,550 (= 249 + 1,200 + 27 + 74; see Table 3) periods in which nonzero inventory was requested. Note that if a recipient requested no inventory, the source could not send any inventory, justifying the removal of such situations from the analysis of transshipments sent. A key to our analysis here was to interpret transshipment decisions both within the context of what was requested and what the source had

	Source Has					
Transshipment	Sufficien (request ≤	t inventory availability)	Insufficient inventory (request > availability)			
Below availability	_	_	50	4.2%		
At availability	—	—	972	80.9%		
Below requested	35	10.1%	86	7.2%		
At requested	313	89.9%	94	7.8%		
	348		1,202			

Table 4. Transshipments Made Upon Request in the Sharing Setting

available. **Table 4** examines how sources respond with transshipments to recipients' requests, depending on whether the request exceeds their available inventory.

It was apparent that most sources either fulfilled a request or sent what they had available. Strikingly, in 15% (= 7.2% + 7.8%) of situations in which inventory requests exceeded a source's availability (i.e., leftover supply), the source nevertheless complied with these requests and transshipped to the recipients more inventory than it had available. In doing so, the sources didn't satisfy all their own demand despite being able to sell this inventory at a higher margin in their own market. Note that the software we use in our experiment allows this to happen because profit is only realized after transshipments are completed. This indicates that participants tend to transship more than they should. This tendency to oversend inventory is stronger than the tendency to undersend. Because recipients are likely to use sources' response to their requests as information to estimate the extra supply through transshipments, such oversending bias means that the recipients process inflated information about the inventory availability at the source. This, in turn, may also lead to the underweighting of the demand-side of inventory sharing.

Our analysis provides additional arguments that support demandside underweighting. We found evidence that people persistently under-request inventory from their partners and oversend inventory to their partners; this further supports the notion of demand-side underweighting in inventory sharing. Decision makers lower their estimate of the demand-side benefit of sharing because they observe generally low requests.

4.2.3. Profitability Implications. To examine how inventory sharing influences the profitability of players, we estimated the effect of our treatments on participants' profits. The model we use is similar to Equation (4), with an additional control variable: the standard deviation of demand. We further add the average order quantities as a covariate to examine whether subjects who ordered more (or less) in particular conditions were able to obtain higher profits. Results from the analysis are summarized in Table A2 in the supplement. Inventory sharing generally increases profits. Effect sizes are fairly consistent, with profits in condition 1 (high margin, no sharing) being lower than profits in condition 3 $(b = 185.40, p \le 0.01)$ and condition 4 $(b = 223.91, p \le 0.01)$, that is, high margin with low and high transfer prices, respectively. Similarly, profits in condition 2 (low margin, no sharing)were lower than profits in condition 5 (b = 238.93, $p \le 0.01$) and condition 6 (b = 226.83, $p \le 0.01$), that is, low margin with low and high transfer prices, respectively. On average, those with higher order quantities in conditions 1, 3, and 4 (high margin) and those with lower order quantities in conditions 2 and 5 (low margin) could obtain more profit. Notably, the average order under condition 6 is 188 (as seen in Table A1 in the supplement), which is close to the optimal order 186 in this condition. Thus, a change in order quantities under condition 6 does not impact profits on average.

To further break down the profitability implications of inventory sharing, we take our observations from conditions 3 and 4 as an example (labeled as "Original" in Table 5) and compare the observed profit as well as the observed service level in these conditions with profits and service levels under several hypothetical counterfactuals. The counterfactual scenarios are (1) if there was no sharing of inventory in this condition (with the same order quantities as in "Original"), (2) if there was no sharing of inventory and hence optimal newsvendor order quantities were placed (= 230), (3) if transshipment decisions were made optimally (with the same order quantities as in "Original"), and (4) if transshipment decisions were made optimally (with optimal order quantities factoring in transshipments (= 212/237) for both players). We focus on conditions 3 and 4 (high-margin, low and high transfer price, respectively) because inventory sharing is typically implemented for high-margin products, and demand-side underweighting leads to forgone profits for the high-margin conditions specifically. Under low-margin conditions, subjects tend to overorder under no sharing because

	No Sharing			Sharing		
Condition	Outcome	(1) Original orders	(2) Optimal orders	Original	(3) Optimal sharing	(4) Optimal sharing & orders
СЗ	Service level	47.7%	66.7%	58.7%	59.9%	70.3%
	Profit	3,084	3,312	3,293	3,325	3,499
<i>C</i> 4	Service level	55.9%	68.8%	66.1%	68.1%	83.2%
	Profit	3,075	3,264	3,322	3,369	3,442

Table 5. Service Levels and Profitability in Conditions 3 and 4 with Counterfactuals

of the pull-to-center effect. Underordering behavior in inventory sharing thus counterbalances the pull-to-center effect and pushes the order closer to the optimal. Results from the analysis on conditions 3 and 4 are summarized in **Table 5**.

Profit and service level comparisons between the "Original" and our counterfactual scenarios (1) and (2) indicate that the ability to share inventory increases profitability (3,293 > 3,084 and 3,322 > 3,075) by a similar amount as ordering optimally would do in a no-sharing system (3,312 > 3,084 and 3,264 > 3,075). Service levels increase in condition 3 (58.7% > 47.7%) as well as in condition 4 (66.1% > 55.9%) because of the increased inventory availability through transshipments. Note that in condition 4, this makes service levels comparable to optimal news-vendor ordering (= 68.8%).

Examining scenarios (3) and (4), we can determine that suboptimal requesting and sending have an impact on profitability, both in condition 3 (3,325 > 3,293) and in condition 4 (3,369 > 3,322). By ordering the optimal quantity, we can see a further significant increase in profitability for condition 3 (3,499 > 3,325) as well as for condition 4 (3,442 > 3,369). Stated differently, in both conditions, on average, about \$163 of profits are foregone under inventory sharing because of ineffective requesting/transshipment decisions and suboptimal order quantities, where the latter aspect explains about 73% of this profit differential.

Our analysis highlights that although inventory sharing does provide significant benefits, even for nonrational decision makers, a significant portion of the benefits of inventory sharing are not realized because order quantities are too low in high-margin contexts and because sharing decisions are not made optimally. **4.2.4. Replication with Decision Support.** We replicated the high-margin treatments in a context with decision support. In practice, most inventory managers have access to a decision-support tool when preparing their initial order quantities, but this tool rarely considers transshipments. Hence, we implemented this aspect by offering players a recommended order quantity (the optimal newsvendor order quantities derived from forecasts made by a single exponential smoothing model with a very low smoothing parameter) that they could accept or revise. This approach is similar to the decision-support treatment in Lee and Siemsen (2017).

Our results (see Table A3 in the supplement) suggest that with decision support, average orders under the low transfer price with sharing are still less than the no-sharing average orders (b = -14.56, $p \le 0.10$), and the average orders are not significantly different between no sharing and the high transfer price setting (b = -7.42, p = 0.15). Thus, demandside underweighting bias persists even in the presence of decision support, where decision makers have the right to revise the system's suggested order quantity.

4.3. Fixed Roles and Estimation of Behavioral Model

In the preceding subsection, we provided evidence for a demand-side underweighting bias. In this subsection, we will quantify this bias with estimates (α and β) derived from the behavioral model developed in Section 4.1. We conducted additional experiments for this purpose because the behavioral parameters in the model cannot be uniquely identified using data from the previous experiments. Different values of α and β could lead to similar model-implied order quantities. We thus introduce a fixed-role treatment in which we separated the roles among participants. Specifically, instead of allowing participants to be both sources and recipients of transshipped inventory, we randomly assigned them roles as either a source or a recipient throughout the experiment. Recipient players saw only the supply side of inventory sharing, whereas source players saw only the demand side. This separation of roles allows us to identify values of α and β by using an experimental treatment.

4.3.1. Experiment Design. The experimental task in the fixed-role

treatment is the same as in our previous experiments. The key treatment we introduce here is the assignment of fixed roles. In addition to their random assignment as either a source or a recipient throughout the experiment, a source player was randomly matched to a recipient player in his or her channel in each period. We also varied the transfer price as an experimental factor. 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 previous treatments. We collected data from 83 subjects across both transfer price treatments (44 in the high transfer price treatment, 40 in the low transfer price treatment). One subject in the high transfer price treatment left the experiment before completion because of a medical emergency, and her data were removed from the analysis. Note also that the fixed-role treatment was run exclusively under the high-margin conditions because we have established that the different profit margin had little influence on the prevalence of supply-side and demand-side thinking. As a normative benchmark, recipient players should order 226 and 210 under transfer

prices of 35 and 15, respectively, whereas source players should order

241 and 232 under transfer prices of 35 and 15, respectively.

4.3.2. Analysis. Our data analysis strategy is similar to the analysis in Section 4.2. One key difference in the analysis lies in our coding approach; although in Section 4.2 experimental conditions were dummy coded, we use dummy codes to represent experimental factors (e.g., three dummy variables—*Recipient, Source,* and *TP* = 35 to indicate the transfer price) in the fixed-role treatment. The main results from this model estimation are shown in **Table 6.** The results show that although the transfer price affects the orders of recipient players (b = 14.41, t =2.20, $p \le 0.05$), source players seem unaffected by the transfer price in their ordering behavior (b = -0.16, t = 0.02, p = 0.98). Further, recipient players across transfer price conditions have lower overall order guantities than those of the recipient + source players (t = -1.63, $p \le 0.10$), whereas source players have similar overall order quantities as those of the recipient + source players (t = 0.45, p = 0.65). This observation is consistent with demand-side underweighting. Recipient players experience the supply side of inventory sharing, which leads them to lower order quantities; their order quantities increase as the potential of this supply-side benefit decreases with an increased transfer price. However,

	Coef.	SE
Recipient	-9.44	(6.39)
Source	7.14	(6.33)
<i>TP</i> = 35	10.08	(5.99)
Recipient × TP = 35	4.33	(8.92)
Source × TP = 35	-10.24	(8.89)
Average demand (≤ 200)	0.71*	(0.30)
Average demand (> 200)	0.53*	(0.27)
Constant	54.89	(59.63)
Ν	133	
R^2	18%	

Table 6. Regression Results for Fixed-Role Treatment

The omitted category is recipient + source. Coef., coefficient; SE, standard error. * $p \leq 0.05$

source players see only the demand side of inventory sharing and do not react to the increased potential of the demand-side benefit as the transfer price increases because of demand-side underweighting.

4.3.3. Estimation of the Behavioral Model. Fixed roles allow us to differentiate between the effects of overestimating the supply side of inventory sharing and underestimating its demand side. Essentially, the fixed-role experiment for the source players imposes α_0 and isolates the impact of β on source players' ordering behaviors. Similarly, the fixed-role experiment for recipient players essentially imposes $\beta = 0$ and isolates the impact of α on recipient players' ordering behaviors.

To estimate our behavioral model, we need to account for the pullto-center effect in newsvendor experiments. In our high-margin setting, the pull-to-center effect also leads to underordering behavior. If the pullto-center effect is not accounted for, underordering would be attributed entirely to demand-side underweighting bias, leading to incorrect estimates. Like Ho et al. (2010), we capture the pull-to-center effect in newsvendor experiments by modeling the psychological cost of overage and underage, that is, δ_o and δ_u . Together, our behavioral model has the following behavior parameters, denoted as $\theta = (\alpha, \beta, \delta_o, \delta_u)$.

We included three experimental designs under both high/low transfer prices (i.e., no sharing, sharing, fixed roles) for the estimation of $\theta = (\alpha, \beta, \delta_a, \delta_a)$. Let Q_n be the equilibrium order quantity from our behavioral model for the no-sharing treatment, $Q_{s,cij}$ for the sharing treatment with transfer price c_{ij} , and $Q^s_{s,cij}$ and $Q^r_{s,cij}$ the equilibrium order quantities with transfer price c_{ij} from our behavioral model for source and recipient players, respectively. Given parameter $\theta = (\alpha, \beta, \delta_o, \delta_u)$, we can solve for a unique set of equilibrium order quantities $(Q_n, Q_{s,cij}, Q^r_{s,cij}, Q^s_{s,cij})$. Specifically, Q_n is given by the modified critical fractile; that is,

$$Q_n = F^{-1} \begin{pmatrix} v - c + \delta_u \\ v - s + \delta_u + \delta_o \end{pmatrix}$$

The solution technique for $(Q_{s,cij}, Q^{r}_{s,cij}, Q^{s}_{s,cij})$ is given in the supplement.

The actual order quantities observed in the experiments may deviate from the prediction of our behavioral model because of noise in the players' decision process. We assume that the average order quantity made by each subject is normally distributed, with the mean order quantity across subjects specified by our behavioral model; that is,

$$q_{i,n} \sim N(Q_n(\theta), \sigma_n^2),$$

$$q_{i,s,cij} \sim N(Q_{s,cij}(\theta), \sigma_{s,cij}^2),$$

$$q_{i,s,cij}^r \sim N(Q_{s,cij}^r(\theta), \sigma_{s,cij}^r^2), \text{ and }$$

$$q_{i,s,cij}^s \sim N(Q_{s,cij}^s(\theta), \sigma_{s,cij}^s^2).$$

The noise parameters σ_{n}^{2} , $\sigma_{s,cij}^{2}$, $(\sigma_{s,cij}^{r})^{2}$, and $(\sigma_{s,cij}^{s})^{2}$ are specific to each treatment; the behavioral parameter θ is assumed to be common across all treatments. We use maximum likelihood to estimate the full behavioral model and two nested models. The details of the estimation procedure are included in the supplement.

The estimation results are shown in **Table 7**. Column (1) of the table shows the full behavioral model, which supports our expectation that $\beta < 1$ (p < 0.01); that is, subjects underestimate how much they can sell to the other players. In addition, the results do not support $\alpha < 1$ (p = 0.321); that is, subjects do not underestimate or overestimate the supply that is available from the other players. This finding implies that the observed underordering behavior can be mainly attributed to underestimation of the demand-side opportunity of selling through transshipments—players generally view transshipment as an additional source of supply instead. These parameter estimates further support the idea of demand-side underweighting bias when players have inventory sharing opportunities.

Parameters	(1) (2)		(3)
α	0.833	22.106	
	(0.455)	(31.239)	
β	0.198	4.931e-6	
	(0.245)	(5.01e-6)	
δ	12.540	34.243	
	(12.567)	(25.680)	
δ"	4.443	24.256	
-	(13.564)	(26.382)	
$-\ln L$	704.382	748.717	707.142
LR test		$\chi^2 = 88.67$	$\chi^2 = 5.52$
		$p \le 0.01$	$p \leq 0.10$

Table 7. Estimation Results of the Behavioral Model

LR, likelihood ratio.

The results of the two nested models are shown in columns (2) and (3) of Table 7. Column (2) shows the results of the model that does not allow for the psychological costs of overage/underage (i.e., $\delta_o = \delta_u = 0$), and column (3) shows the results of the model that does not allow for demand-side underweighting bias (i.e., $\alpha = \beta = 1$). The likelihood-ratio (LR) test shows that the behavioral model that allows for both the psychological cost of overage/underage and demand-side underweighting bias better explains the data than these two nested models. Without incorporating the psychological costs of overage/underage, which account for the pull-to-center effect, we will overstate the demand-side underweighting bias, as shown in column (2) of Table 7.

The ratio of the estimated values for the psychological cost factors in column (3) of Table 7 is similar to that reported in Ho et al. (2010) $(\delta_u/\delta_o = 0.65$ in Ho et al. (2010) versus $\delta_u/\delta_o = 0.71$ in this study). However, when we introduced our behavioral factors α and β into the estimation, this ratio changed. Column (1) in Table 7 shows that the ratio δ_u/δ_o drops to 0.35—in other words, the psychological costs of underage become even less salient once inventory sharing is factored into the decision-making process. It is possible that decision makers under inventory sharing no longer attribute some of the responsibility for an inventory shortage to themselves but to the other player.

5. Exploring Root Causes and Possible Mitigation Strategies

In this section, we examine several potential underlying reasons for demand-side underweighting and the resulting strategies to mitigate this bias. We begin by discussing existing theories for underordering behavior to see if they provide plausible explanations of such underordering associated with inventory sharing, followed by more novel proposed explanations (Section 5.1) combined with additional experiments to explore these explanations (Section 5.2). We then investigate several alternatives that may help mitigate demand-side underweighting (Section 5.3).

5.1. Root Causes

In this subsection, we provide a systematic discussion of existing behavioral theories to explain the observed underordering behavior (including psychological costs of overstocking/understocking, mean anchoring, overconfidence, and risk aversion), followed by three proposed novel root causes underpinning demand-side underweighting.

5.1.1. Existing Theories. Psychological costs of overstocking/ understocking—a key explanation for the pull-to-center effect—could lead to underordering. To account for the psychological underage and overage costs, we deliberately included these factors in the estimation of our behavioral model in the fixed-role treatment. Our estimation results indicate that the model incorporating demand-side underweighting bias better explains the data than the one with only psychological overage/underage costs, they will be hesitant to increase their orders in anticipation of potential transshipment demand. However, they should expect similar behavior from the other player; hence, they should not expect much supply-side benefit from inventory sharing either. Explaining demand-side underweighting ultimately requires incorporating a behavioral asymmetry between a player's own behavior and the other play-

Mean anchoring could also result in underordering for the high-margin treatments (Schweitzer and Cachon 2000). Given that the order decision under inventory sharing is more complex than the order decision under no sharing, subjects may be more inclined to anchor on mean demand as a heuristic, resulting in an order quantity under inventory sharing that is closer to the mean demand (e.g., 200), that is, lower than the order quantity under no sharing. However, this argument cannot explain why participants reduce their inventory orders further (i.e., order further away from mean demand) under the low-margin sharing treatments, as shown in Table 2.

Overconfidence can also cause underordering (Ren and Croson 2013). Overconfidence implies that decision makers underestimate the standard deviation of the underlying demand distribution. In the inventory sharing context, this could, in turn, imply that players overestimate their own (and the other player's) service level, thereby underestimating their own (and the other player's) propensity to request inventory and overestimating the other player's (and their own) propensity to have inventory available. Notice, though, that the effect of overconfidence also seems symmetrical— it would lead to an underestimation of the potential of both the demand side and the supply side of inventory sharing; hence it does not appear to be able to explain our situation.

Risk aversion could also lead to underordering because a risk-averse player may be reluctant to invest much into inventory for future uncertain demand either of his or her own or through transshipments. However, our theoretical analysis of risk aversion suggests the opposite: a risk-averse player will order more inventory under sharing than under no sharing (see Figure A1 in the supplement). Inventory sharing lowers the supply chain risk of not matching demand and supply, thus encouraging a higher order quantity. More fundamentally, classical behavioral newsvendor theories, such as risk aversion, mostly focus on how human decision makers balance the uncertain demand and *certain* supply (the order will be delivered for sure) and do not differentiate the risk from the uncertain demand and uncertain supply. However, the observed underordering behavior in our experiments points to an asymmetry: the risk of not being able to send excess inventory to the other player (i.e., uncertain demand) is more salient than the risk of not being able to receive inventory from the other player (i.e., uncertain supply). Therefore, any possible causes must explain such asymmetry and discuss how decision makers balance uncertain supply and uncertain demand.

5.1.2. Proposed Explanations. Having explored the preceding, we propose some additional causes for demand-side underweighting. Desire

for control over their environment (Leotti et al. 2010) can manifest in a preference for choice and sometimes even in a preference for the illusion of control through choice. In our context, transshipments are not under the direct control of either player but require the consent of both players. However, according to the sequence of events, it is the recipient who initiates a transshipment with his or her request, and the source can only respond to this request. As such, decision makers may construe their ability to request additional inventory from the source as a way of controlling their environment ("If more customers show up than expected, I can always do something about this by asking for more inventory from the other player!"). In contrast, the ability to sell inventory to the other player may appear to offer less control over the environment ("I surely can stock some extra units for the other player, but whether he or she requests them from me is out of my control"). In turn, this asymmetry could lead to greater salience of the supply-side benefit of inventory sharing during the decision-making process.

Bias in others is easier for people to see than bias in themselves. Transshipments are impossible if the potential recipient does not request inventory from the source. However, asking for a supplementary inventory shipment is akin to asking for help. Seeking help may imply a degree of admitting one's own incompetence and 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 depends on the other player to satisfy his or her demand. In other words, requesting units for transshipments imposes an extra psychological burden on the recipient; this would lead to less-than-optimal requests, which is consistent with our finding in Table 3. Although this psychological burden applies to both a focal player and his or her counterpart, people often spot and perceive bias in others more than in themselves (Pronin et al. 2002)—"Although I am going to request whenever I need in order to take advantage of the sharing opportunity, I suspect the other player may be reluctant to request from me." In turn, this implies that decision makers may expect the other player to be reluctant to request inventory from them, but they may not see themselves as reluctant to request inventory from the other player. This asymmetry may cause decision makers to underweight the demand-side benefit of inventory sharing.

Valuing one's own market more than the other's market may also lead to demand-side underweighting. Consumers value products in which they have invested effort more than those without their investment of effort (Norton et al. 2012). In Norton et al. (2012), subjects develop a particular valuation for the outcome of their manual labor; in our context, we expect that subjects develop a particular valuation for the outcome of their risky inventory investment. Although manual labor is different from inventory investment, both represent a form of effort. The logic underlying Norton et al. (2012) is built on the idea that "the more effort people put into some pursuit, the more they come to value it" and that the "successful completion of tasks [is] one crucial means by which people can meet their goal to feel competent and in control" (Norton et al. 2012, p. 454). In other words, people value the effort they put into a task, particularly if that effort leads to success. Fulfilling their own demand out of ordered inventory can be construed as expanding effort to succeed in a task. Transshipping inventory to a receiver thus may be seen by subjects as supporting another player who failed in this task rather than as succeeding in their own task. Such a construal would lead to different valuations.

5.2. Testing the Proposed Root Causes

To test these three proposed causes, we conducted three additional treatments: minimum send, automated requesting, and sending reward. Each treatment is designed to manipulate a corresponding root cause (i.e., desire for control; see reluctance in requesting in others but not in oneself and valuing one's own market more than the other's market). By contrasting subjects' orders under these three treatments with those under our original treatments from Section 4.2, we can explore which of the causes plays a role in subjects' order decisions.

5.2.1. Experiment Design. The original high margin sharing conditions from Section 4.2 are the baseline treatments underlying all factor variations in our new experiments. Under minimum send, a player who has extra units is guaranteed to be able to send up to five units that the other player must accept, even if the other player requests fewer than five units or does not request any units from the player. Such a treatment reduces the strategic uncertainty in the game and thereby increases a

subjects' control over their uncertain environments by enabling them to send a minimum amount of excess inventory.

Under *automated requesting*, our software will automatically request inventory at the optimal level for the player if he or she has a shortage instead of letting the player request himself or herself. Each player still decides whether and how much inventory he or she is willing to transship to the requesting player. This treatment tackles the second potential root cause of seeing a bias in others but not in oneself because requesting is now automated by software, and under-requesting will be eliminated. The concern that "others will not request" is now reduced.

Under *sending reward*, for each unit of inventory the player sends to the other player, the player will receive not only a transfer price from the other player but also an additional five units of profit (paid from an external source) as a reward. This simulates, for example, the situation where a dealer receives a discount for his or her next order from a manufacturer every time he or she shares inventory to another dealer in need, as in Caterpillar's dealer inventory sharing system (Zhao et al. 2005). This additional reward should make the benefit of selling inventory to the other player more salient.

We ran these additional treatments under both high and low transfer prices. We otherwise used the same subject pool and design as in our previous studies (although subjects who had participated in previous studies were not allowed to participate again). We collected data on 134 additional subjects.

5.2.2. Analysis. Our analysis in this section has to deal with two complications. First, the minimum send and sending reward treatments manipulate the salience of the discussed root causes of demand-side underweighting but also shift the rational optimal order quantities. In particular, our theoretical analysis suggests that under *minimum send* (*sending reward*), a rational player should order 240 (239) and 215 (216) under transfer prices of \$35 and \$15, respectively. In other words, compared with the baseline treatments in Section 4.2, minimum send and sending reward treatments should improve order quantities by roughly three units even among rational players. To disentangle the impact of the changed salience of the root causes from the rational response to incentives, we employ a different outcome measure: *OrderGap_i* = *Order_i* – *OptimalOrder_i*, where *Order_i* represents the observed order,

and *OptimalOrder*_i represents the corresponding optimal order for rational players. This dependent variable corrects observed order quantities for the differences that naturally exist between treatments because of changed optimal orders.

Second, these new treatments also influence subjects' transshipment decisions. Although in theory the actual transshipment decisions of another player should not influence a focal player, participants are likely to react to such observed information as an increased extra supply through transshipments, which could lower the order quantities of a focal player. Similarly, increased observed extra demand from the other player could increase order quantities. For example, we expect that under the sending reward treatment, subjects are more inclined to transship additional inventory because of the additional reward. Although our main interest is to examine whether this increased salience of the benefit of transshipments increases order quantities because of demand-side thinking, subjects will also likely lower their order quantities because they have received more extra supplies from the other player. To estimate direct treatment effects, these indirect effects need to be accounted for in the analysis. Therefore, to measure the extra demand ED, available through transshipments, we calculate the average transshipment requests received by a focal player. The extra supply ES, through transshipment is somewhat more intricate to measure because supply can only be sent upon request. We thus calculate the average transshipments received by a focal player over all time periods that he or she requested transshipments from his or her partner. Thus, ED, is averaged over all 30 time periods in the experiment, whereas ES, is averaged only over time periods where requests are made. Treatment averages for these two measures are summarized in Table A1 in the supplement.

Because ED_i and ES_i depend on players' order decisions, they are potentially endogenous. To address this potential endogeneity issue, we constructed the average excess supply XS_i (i.e., $XS_i = \text{mean}([Order_{-i, t} - Demand_{-i, t}]^+))$ of a focal player's partner as an instrument for ES_i , and the average excess demand XD_i (i.e., $XD_i = \text{mean}([Demand_{-i, t} - Order_{-i, t}]^+))$ of a focal player's partner as an instrument for ED_i , where -i represents focal player *i*'s partner. We believe that our instruments are exogenous because XS_i and XD_i are determined by the random demand draw of the partner, and the partner is also random matched every round. Furthermore, both instruments are relevant and strong because the first-stage regression shows that ED_i is predicted with an R^2 of 80% (minimum eigenvalue statistic of 550), and ES_i is predicted with an R^2 of 54% (minimum eigenvalue statistic of 20). Overidentification tests are not available because the equations are just identified. The results of two separate two-stage least squares regressions with either ED_i or ES_i being instrumented reveal that ES_i is not endogenous (Wu-Hausman F = 0.06, p = 0.81) but ED_i is endogenous (Wu-Hausman F = 5.45, $p \le 0.05$). Let the three binary variables MIN_i , $AUTOR_i$, and $REWARD_i$ represent the three experiment designs (namely, minimum send, automated requesting, and send reward) and let P_i represent the transfer price. We thus simultaneously estimate the following system of equations using the generalized structural modeling procedure in Stata:

$$OrderGap_{i} = a_{1}^{1}P_{i} + a_{2}^{1}MIN_{i} + a_{3}^{1}P_{i} \times MIN_{i}$$

$$+ a_{4}^{1}AUTOR_{i} + a_{5}^{1}AUTOR_{i} \times P_{i}$$

$$+ a_{6}^{1}REWARD_{i} + a_{7}^{1}REWARD_{i} \times P_{i}$$

$$+ \gamma_{1}D_{1i} + \gamma_{2}D_{2i} + \gamma_{3}ED_{i} + \gamma_{4}ES_{i} + \epsilon_{1,i}$$
(5)

$$ED_{i} = a_{1}^{2}P_{i} + a_{2}^{2}MIN_{i} + a_{3}^{2}P_{i} \times MIN_{i} + a_{4}^{2}AUTOR_{i}$$
$$+ a_{5}^{2}AUTOR_{i} \times P_{i} + a_{6}^{2}REWARD_{i}$$
$$+ a_{7}^{2}REWARD_{i} \times P_{i} + \delta_{1}XD_{i} + \epsilon_{2i}$$
(6)

$$ES_{i} = a_{1}^{3} P_{i} + a_{2}^{3} MIN_{i} + a_{3}^{3} P_{i} \times MIN_{i} + a_{4}^{3} AUTOR_{i} + a_{5}^{3} AUTOR_{i} \times P_{i} + a_{6}^{3} REWARD_{i} + a_{7}^{3} REWARD_{i} \times P_{i} + \delta_{2}XS_{i} + \epsilon_{3,i}$$
(7)

Notably, we did not employ this form of analysis in our previous experiment studies because their primary focus was to compare sharing with no-sharing treatments and the extra demand and extra supply variables are not defined under no-sharing treatments.

The estimation results are shown in **Table 8**. As expected, the extra supply available through transshipments decreases order quantities (b = -0.51, $p \le 0.01$), whereas the extra demand available through transshipments increases order quantities (b = 0.41, $p \le 0.01$). It is also apparent that our treatments increased the extra supply available through transshipments under the low transfer price, whereas extra demand remains mostly unaffected by our treatments. For instance, under a low transfer price, minimum send increases the average quantities received

	Ord	er	Extra su	Extra demand		
Dependent variable	Coef.	SE	Coef.	SE	Coef.	SE
<i>TP</i> = 35	0.42	(5.77)	23.41**	(2.72)	-0.21	(1.34)
Minimum Send	16.37**	(5.67)	18.46**	(2.79)	1.08	(1.37)
Minimum Send × TP = 35	-18.22*	(7.64)	-23.16**	(3.89)	-4.44*	(1.93)
Automatic Request	13.68*	(5.87)	20.85**	(2.82)	2.24	(1.40)
Automatic Request × TP = 35	-19.29**	(7.71)	-20.55**	(4.02)	-0.05	(2.00)
Sending Reward	18.78**	(5.78)	23.20**	(2.70)	0.69	(1.34)
Sending Reward × TP = 35	-27.32**	(7.38)	-15.92**	(3.89)	-1.81	(1.93)
Extra supply	-0.51**	(0.13)				
Extra demand	0.41**	(0.13)				
Average demand (≤200)	0.90**	(0.22)				
Average demand (>200)	0.33^{+}	(0.19)				
Partner excess supply	0.40**	(0.07)				
Partner excess demand	0.89**	(0.04)				
Constant	-206.61**	(43.42)	-11.49**	(2.76)	1.28	(1.52)
Ν	184					

Table 8. Estimation Results in Testing Root Causes

Coef., coefficient; SE, standard error.

** $p \le 0.01$; * $p \le 0.05$; † $p \le 0.10$

through transshipment by 18 units and does not impact the average quantities demanded by partner players. Such an indirect counter effect of increased received supplies through transshipments will by itself lead to a decrease in order quantities. Therefore, minimum send only marginally improves the average order quantities compared with the corresponding base treatment (204 - 198 = 6; see Table A1 in the supplement).

Once this indirect countereffect has been accounted for, we can establish the direct effects of our treatments. We observe a reduction in the gap between the actual and optimal orders for low transfer prices under minimum send (b = 16.37, $p \le 0.01$), automated requesting (b = 13.68, $p \le 0.05$), and sending reward (b = 18.78, $p \le 0.01$). There is no effect of our treatments under the high transfer price (the effect size for the three treatments under high transfer price can be obtained by adding the coefficients of the treatment dummy and its interaction with transfer price). Thus, our treatments succeeded both in increasing order quantities as well as transshipments supplied under a low transfer price, a situation where the natural benefit of sending transshipments is low. By manipulating the perceived control (minimum send), the reliance on the transshipment partner's making adequate requests (automated requesting) and increasing the value of the transshipment (sending reward), we succeeded in manipulating behavior in the expected direction. Our manipulations suggest that all three root causes we mention seem valid; demand-side underweighting thus does not appear as a phenomenon with a single root cause but is the result of several likely causes. This is not surprising, given that the pull-to-center effect in newsvendor decision making is also an empirical phenomenon with many likely causes.

It is interesting that the same manipulations were not successful under a high transfer price. This is surprising because, in principle, the same psychological barriers to recognizing the demand-side benefits of transshipments that are present under a low transfer price should also be present under a high transfer price. A possible explanation may be that ceiling effects kick in; given that order quantities are already higher under the high transfer price, participants are reluctant to cross a certain order quantity perceived as an upper limit of their orders.

5.3. Additional Treatments

To further explore ways to mitigate demand-side underweighting, we investigate three additional experimental treatments: the first one deals with the framing of the study; the second one deals with the *sharing* decisions by automating both the requesting and the sending decisions (hence eliminating oversending and under-requesting behaviors); and the third one directly deals with the *stocking* decisions through an advanced decision-support system.

First, it is plausible that the framing of our previous study has emphasized demand-side underweighting—subjects had to request inventory from the other players, which naturally emphasizes the supply-side benefit. We will thus examine whether changing this frame allows us to deemphasize the supply-side benefit and emphasize the demand-side benefit instead. We thus studied two additional alternative forms of framing the problem: customer-transfer framing and demand-side framing. Both frames are mathematically equivalent to our original frame and also practically relevant because they correspond to different ways to allow transshipments to happen. Under the customer-transfer frame, a decision maker directs his or her excess customers to the other player instead of requesting extra inventory from the other player for these customers. In other words, the recipient player does not receive inventory from the source player to satisfy his or her demand. Instead, he or she sends such customers to the source, who then satisfies the demand (with a commission, i.e., the transfer price, to the recipient player). Note that all financial transactions are the same as in our original experiment in study 1 (so that the rational solution remains the same). This is akin to retail stores from the same firm asking their customers to stop by a different store to pick up inventory. We believe that this frame deemphasizes the supply-side benefit (because the recipient now sends customers away). Such transferred customers have no long-term implications for the recipients' demand in our experiment (we told the subjects that sending away customers is a one-time occurrence with no implications for future demand). Nevertheless, this frame emphasizes the "lost business" aspect of insufficient inventory. Thus, we expect this frame to lead to higher order quantities.

Under the demand-side frame, transshipments are initiated when players propose to sell inventory to the other player instead of players requesting inventory to be transshipped. We believe that this frame emphasizes demand-side benefit because players now initiate the sales (instead of responding to a sales request from the other player). Hence, compared with the original frame, this should increase the players' order quantities.

We also designed a treatment with automated requesting and sending, under which the software will automatically make the optimal requesting/sending decisions. Under such a treatment, control over the demand side and supply side have been equally rendered to the software. Therefore, subjects do not have to worry about the other player's reluctance and irrationality in both requesting and sending. As a result, subjects can feed unbiased information about the inventory availability of the system into their ordering systems, which used to be inflated because of the underrequesting and oversending behaviors. We thus believe that underordering could be mitigated.

Finally, akin to Bolton et al. (2012), we designed a new decision-support treatment, under which we told our subjects the optimal order quantity. Specifically, given the transfer price of \$35 (\$15), the subjects are told that the optimal order quantity should be 237 (212). In addition, we also provided subjects with a plot of the expected profit under different order quantities to ensure that they understand the optimality of the suggested optimal order. Nevertheless, as in practice, subjects could

choose to accept the suggested optimal order quantity or to override it.

5.3.1. Experiment Design. Conditions 3 and 4 from Section 4.2 are again the baseline treatments underlying all factor variations here. Under the customer-transfer frame, instead of asking for supplementary inventory orders from a source, we ask players, "How many customers do you want to send to the other player?". Source players were then told, "Please enter the number of customers you are willing to serve for the other player." The transfer price was explained as "the amount you have to give to the other player for serving each unit of your own customer *demand*." Under the demand-side frame, we asked players, "How much inventory do you want to offer to sell to the other player?". Design and subject pool were otherwise similar to our previous work. We collected data from 186 additional subjects who participated across the eight (= 4×2) new treatments we ran.

5.3.2. Analysis. The analysis is like that in Section 5.2. **Table 9** contains our estimation results. When comparing the order quantities under customer transfer with order quantities under our original framing, it becomes clear that with the low transfer price, inventory orders are increased in the customer-transfer framing ($p \le 0.01$). They remain similar under the high transfer price (p = 0.71). This result is as expected. In the low transfer price scenario, in which supply-side benefits are most salient, reframing the setting clearly decreases this salience and leads to inventory order levels that are no different from newsvendor orders without sharing inventory. In the high transfer price setting, in which a supply-side benefit has little salience to begin with, customer transfer has no further effect.

Although our expectations for the customer-transfer setting are supported by the data from the new experiments, our expectations for increased order quantities under the demand-side frame are not supported. Order quantities under the low transfer price here do not change compared with the original frame (p = 0.17) and do not change under the high transfer price either (p = 0.97). This is again the result of the indirect countereffect we discussed in the preceding section. In particular, under the demand-side frame, participants offered too much inventory to the other player (whereas under the original frame, supplementary order requests were generally 4.5 units less than the optimum; under

	Orde	r	Extra su	pply	Extra demand	
Dependent variable	Coef.	SE	Coef.	SE	Coef.	SE
<i>TP</i> = 35	14.89**	(5.33)	23.73**	(3.45)	-0.89	(2.22)
Demand side	8.16	(6.06)	41.85**	(3.32)	-5.77**	(2.16)
Demand side \times TP = 35	-7.97	(6.94)	-20.12**	(4.76)	-3.53	(3.09)
Customer transfer	13.41**	(5.43)	23.96**	(3.50)	1.62	(2.26)
$Customer \times TP = 35$	-11.60†	(7.23)	-22.06**	(4.91)	-5.58^{\dagger}	(3.19)
Automatic Request & Send	10.58*	(5.17)	16.64**	(3.52)	1.79	(2.26)
Automatic Request & Send×TP=35	-6.17	(7.25)	-24.10**	(4.91)	-0.06	(3.18)
Decision support v.2	16.41**	(5.29)	20.56**	(3.51)	1.49	(2.29)
Decision support v.2 × TP = 35	-12.48†	(7.46)	-25.33**	(5.04)	-1.02	(3.27)
Extra supply	-0.13	(0.09)				
Extra demand	0.11	(0.11)				
Average demand (≤200)	1.15**	(0.18)				
Average demand (>200)	0.39*	(0.19)				
Partner excess supply			0.35**	(0.10)		
Partner excess demand					0.72**	(0.06)
Constant	-34.61	(36.38)	-10.03**	(3.63)	7.12	(2.49)
Ν			236			

Table 9. Estimation Results in Additional Treatments

Coef., coefficient; SE, standard error.

 $^{\dagger}p \le 0.10$; $^{*}p \le 0.05$; $^{**}p \le 0.01$

the demand-side frame, people, on average, offered 12 units more than optimal to the other player). In other words, the frame changed the setting from people under-requesting to people oversending inventory. This leads to the very rational perceptions of players that more than optimal inventory is available from the other player through transshipments, emphasizing a supply-side benefit, in turn, and hence reduced initial order quantities. This effect seems to offset any additional framing effects that may be present. Indeed, the average number of units transshipped under demand-side framing is 14, which is much higher than the average of nine units transshipped under the supply side framing.

The order quantities under automated requesting and sending increase by 10.58 units (p < 0.05) under the low transfer price and increase by 4.41 units (p = 0.37) under the high transfer price compared with the order quantities under our original framing in Section 4.2. In particular, the order quantity under the high transfer price is slightly higher than the order quantity under the no-sharing treatment in Section

4.2, which is a sign of overcoming the pattern of underordering. Thus, not only does automated requesting and sending correct the bias in subjects' transshipment decisions, but it also mitigates the underordering behavior in their initial stocking decisions.

It is also clear that the decision-support tool generally increases order quantities, with an increase of 16.41 units (p < 0.01) under the low transfer price setting and an increase of 3.93 units (p = 0.44) under the high transfer price setting compared with the original framing in Section 4.2. Nevertheless, the increase in order quantities is not significant under the high transfer price setting despite the fact that subjects are explicitly told that the optimal order quantity should be 237 units. This indicates that the bias is so strong that the subjects are willing to leave money on the table (possibly because of the ceiling effect we mentioned that subjects are not willing to order above a certain level); hence, some intervention or education is needed to increase subjects' compliance with such a decision-support tool.

6. Conclusion

The analytical modeling literature on inventory sharing has emphasized that this practice has the potential to simultaneously reduce inventory and increase service levels (e.g., Tagaras 1989, Evers 2001). However, this literature has always assumed rational decision makers and has not considered that human decision makers can exhibit biased decision making. Using four behavioral studies, we found persistent and significant underordering behavior among decision makers when provided with inventory sharing opportunities. Such behavior persisted even when decision makers had a decision-support system that suggested inventory order quantities but permitted them to make revisions (the typical practice in industry). We identified demand-side underweighting as one evident underlying bias behind this underordering behavior. Specifically, inventory sharing should confer a supply-side benefit (that others may become an extra source of supply) and a demand-side benefit (that others may also become a secondary market to whom to sell extra inventory). However, its demand-side benefit is largely underestimated by decision makers. Inventory sharing is perceived as rather a supply-side than a demand-side phenomenon. For high-margin products, such underordering exacerbates the already low orders caused by pull-to-center

behavior and can lead to much foregone benefit from risk pooling, as evidenced by limited improvement of service levels and profits. For *lowmargin products*, such underordering counterbalances the pull-to-center behavior and may drive the order quantity closer to optimal.

Our work provides some important managerial insights: First, unless human behavior is considered, inventory sharing systems may forego a portion of their promised risk-pooling benefits. Thus, during the implementation of such systems, the demand-side benefits of the system need to be emphasized for decision makers. This may be achieved through (1) appropriate emphasis to overcome demand-side underweighting bias, (2) more robust decision-support systems with transshipment consideration to suggest stocking level, and (3) providing guidelines or automation of the inventory requesting and sending decisions to signal the correct inventory availability of the system. 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. Third, the unwillingness of participants to share their inventory does not appear to be a major cause of the breakdown of inventory systems. More blame seems to lie with overall insufficient inventory in the system coupled with players' reluctance to request adequate inventory from others through transshipments. To improve the system, equal importance is warranted to guiding the initial ordering process as well as to helping encourage reasonable requests.

As one of the first studies to consider inventory sharing from a behavioral perspective, our study opens many directions to explore. For example, to separate the impact of demand uncertainty, we look at a setting in which there is no residual demand uncertainty when transshipments are requested. In reality, decision makers often need to decide whether they will share with others when their demand is not yet fully realized, making a trade-off between the sharing opportunity now and potential lost sales in one's market in the future because of demand uncertainty. Further, this research context provides a clear opportunity to study social preferences as well as social comparison effects. The act of sharing may make people sensitive to the other players' outcomes, whereas the mere ability to observe decisions and outcomes from the other players on a platform may lead to explicit or implicit performance comparisons. In addition, future research may further explore the root causes of demand-side underweighting by surveying participants through verbal protocol analysis such as the one reported in Gavirneni and Isen (2010).

We hope that our research will lead to more behavioral studies in this interesting area.

Acknowledgments The authors thank the review and editorial team for constructive and useful comments. Further gratitude goes to Shan Li and Kay-Yut Chen for their comments and feedback on earlier drafts of this manuscript.

Funding The authors thank the Smeal Small Research Grant at Smeal College of Business, Pennsylvania State University, for supporting their research. The authors also acknowledge the support from the Association for Supply Chain Management (ASCM) for providing the funds to conduct the survey.

Supplemental Material follows the References.

Endnotes

- 1 The survey has parts focused on inventory management, forecasting, supply chain management, and sales and operations planning. To reduce time requirements, respondents answered only one of these four parts.
- 2 Note that this choice of parameters for the normal demand distribution scales the parameters from RKP by a factor of two. This deviation from RKP occurred because our original design included a model in which demand revelation was split into two stages. The resulting experiments were removed from this paper to reduce complexity.

References

- Alishah EJ, Moinzadeh K, Zhou Y (2017) Store fulfillment strategy for an omnichannel retailer. Working paper, Microsoft Corporation, Redmond, WA.
- Anupindi R, Bassok Y, Zemel E (2001) A general framework for the study of decentralized distribution systems. *Manufacturing Service Oper. Management* 3(4):349–368.
- Bolton GE, Ockenfels A, Thonemann UW (2012) Managers and students as newsvendors. *Management Sci.* 58(12):2225–2233.
- Bostian AA, Holt CA, Jain S, Ramdas K (2012) Is transshipment a behaviorally-robust risk-pooling strategy? Working paper, University of Tampere, Tampere, Finland.
- Chen K-Y, Li S (2020) The commitment conundrum of inventory sharing. *Production Oper. Management* 29(2):353–370.
- Dong L, Rudi N (2004) Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Sci.* 50(5):645–657.
- Evers PT (2001) Heuristics for assessing emergency transshipments. *Eur. J. Oper. Res.* 129(2):311–316.
- Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments.

Experiment. Econom. 10(2):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 Bus. Rev.* December:1–6.
- Gavirneni S, Isen AM (2010) Anatomy of a news vendor decision: Observations from a verbal protocol analysis. *Production Oper. Management* 19(4):453–462.
- Ho T-H, Lim N, Cui TH (2010) Reference dependence in multilocation newsvendor models: A structural analysis. *Management Sci.* 56(11):1891–1910.
- Hyndman K, Embrey M (2018) The econometrics of experiments. Katok E, Leider S, Donohue K, eds. *Handbook of Behavioral Operations Management* (John Wiley & Sons, Hoboken, NJ), 35–88.
- Lee F (1997) When the going gets tough, do the tough ask for help? Help seeking and power motivation in organizations. *Organ. Behav. Human Decision Processes* 72(3):336–363.
- Lee YS, Siemsen E (2017) Task decomposition and newsvendor decision making. *Management Sci.* 63(10):3226–3245.
- Leotti LA, Iyengar SS, Ochsner KN (2010) Born to choose: The origins and value of the need for control. *Trends Cognitive Sci.* 14(10):457–463.
- Narus JA, Anderson JC (1996) Rethinking distribution: Adaptive channels. *Harvard Bus. Rev.* 74(4):112–120.
- Norton MI, Mochon D, Ariely D (2012) The IKEA effect: When labor leads to love. *J. Consumer Psych.* 22(3):453–460.
- Pronin E, Lin DY, Ross L (2002) The bias blind spot: Perceptions of bias in self vs. others. *Personality Soc. Psych. Bull.* 28(3):369–381.
- Ren Y, Croson R (2013) Overconfidence in newsvendor orders: An experimental study. *Management Sci.* 59(11):2502–2517.
- Rong Y, Snyder LV, Sun Y (2010) Inventory sharing under decentralized preventive transshipments. *Naval Res. Logist.* 57(6):540–562.
- Rudi N, Kapur S, Pyke DF (2001) A two-location inventory model with transshipment and local decision making. *Management Sci.* 47(12):1668–1680.
- Schweitzer ME, Cachon GP (2000) Decision bias in the newsvendor problem with a known demand distribution: experimental evidence. *Management Sci.* 46(3):404–420.
- Shao J, Krishnan H, McCormick ST (2011) Incentives for transshipment in a supply chain with decentralized retailers. *Manufacturing Service Oper. Management* 13(3):361–372.
- Tagaras G (1989) Effects of pooling on the optimization and service levels of twolocation inventory systems. *IIE Trans.* 21(3):250–257.
- van Donselaar KH, Gaur V, Van Woensel T, Broekmeulen RA, Fransoo JC (2010) Ordering behavior in retail stores and implications for automated replenishment. *Management Sci.* 56(5):766–784.
- Villa S, Katok E (2018) Negotiating transfer prices for improving supply chain transshipments. Working paper, University of Texas at Dallas, Dallas.

- Yan X, Zhao H (2011) Technical note—Decentralized inventory sharing with asymmetric information. *Oper. Res.* 59(6):1528–1538.
- Zhang Y, Siemsen E (2019) A meta-analysis of newsvendor experiments: Revisiting the pull-to-center asymmetry. *Production Oper. Management* 28(1):140–156.
- Zhao H, Bisi A (2010) Optimal operating policies in a commodity trading market with the manufacturer's presence. *Naval Res. Logist.* 57(2):127–148.
- Zhao H, Deshpande V, Ryan JK (2005) Inventory sharing and rationing in decentralized dealer networks. *Management Sci.* 51(4):531–547.
- Zhao H, Deshpande V, Ryan JK (2006) Emergency transshipment in decentralized dealer networks: When to send and accept transshipment requests. *Naval Res. Logist.* 53(6):547–567.

Appendix I: Tables and Figures

Treatment	Transfer Price	Ν	Order	Profit	Extra Supply	Extra Demand
No Sharing - High Margin	-	26	209.74	3220.00	0.00	0.00
No Sharing - Low Margin	-	24	194.07	1040.31	0.00	0.00
Sharing - High Margin	35	24	206.89	3321.79	25.93	27.97
Sharing - High Margin	15	26	197.79	3293.09	15.41	31.71
Sharing - Low Margin	35	26	187.50	1285.23	23.42	33.10
Sharing - Low Margin	15	26	185.80	1295.06	18.89	39.47
No Sharing - Dec. Support	-	20	225.30	3155.78	0.00	0.00
Sharing - Dec. Support	35	20	216.22	3333.22	24.80	21.07
Sharing - Dec. Support	15	20	209.43	3374.50	24.69	26.21
Fixed Role	35	43	201.28	3051.42	30.94	15.63
Fixed Role	15	40	194.96	3100.98	36.85	21.72
Minimum Send	35	24	208.07	3231.78	22.31	25.69
Minimum Send	15	22	203.84	3226.20	21.49	31.80
Auto Request	35	20	199.69	3208.97	26.07	31.64
Auto Request	15	20	198.75	3268.60	22.09	34.14
Sending Reward	35	24	201.60	3443.90	31.04	30.12
Sending Reward	15	24	201.84	3339.01	25.14	28.38
Customer Transfer	35	24	206.66	3116.69	28.80	24.84
Customer Transfer	15	22	208.01	3221.50	25.55	31.21
Demand Side	35	26	200.09	3231.44	47.02	26.52
Demand Side	15	26	197.02	3244.62	42.31	28.39
Auto Request & Send	35	24	213.53	3376.50	19.54	28.93
Auto Request & Send	15	22	203.29	3285.76	18.75	32.39
Decision Support v.2	35	20	212.94	3391.38	21.27	26.52
Decision Support v.2	15	22	210.07	3405.44	22.42	28.39

Table A1: Summary of Experimental Data

Notes. While in general, the average extra supply is below the extra average demand, the opposite can be true as well. Extra supply is averaged only over time periods where requests were made, since extra supply must be zero otherwise. Extra demand is averaged over all time periods, since requests can be made in any period. This difference in the denominator can lead to average extra supply being greater than average extra demand. Further, the sequencing of extra demand/supply is revered in the Demand Side treatment, leading to extra demand being less than extra supply.

	Model 1		Model 2		
	Coef.	SE	Coef.	SE	
C2: Low Margin – No Sharing	-2019.42**	(54.85)	-2032.76**	(48.02)	
C3: High Margin – Low Transfer Price	185.40**	(53.54)	174.74**	(46.75)	
C4: High Margin – High Transfer Price	223.91**	(54.36)	211.51**	(47.52)	
C5: Low Margin – Low Transfer Price	-1780.49**	(53.46)	-1787.76**	(46.65)	
C6: Low Margin – High Transfer Price	-1792.59**	(53.56)	-1805.65**	(46.75)	
Order Quantity @ C1			4.34**	(1.66)	
Order Quantity @ C2			-2.72*	(1.38)	
Order Quantity @ C3			5.43**	(1.74)	
Order Quantity @ C4			7.79^{**}	(1.71)	
Order Quantity @ C5			-3.95**	(1.27)	
Order Quantity @ C6			0.31	(1.46)	
Average Demand (≤200)	11.17**	(2.77)	9.66**	(2.57)	
Average Demand (>200)	16.52**	(2.05)	14.96**	(1.88)	
Demand Standard Deviation	-11.45**		-12.21**	(1.41)	
Constant	1593.91**	(555.6)	1963.08**	(515.6)	
N	15	2	152		
R ²	97%		98%		

Table A2: The Effects of Treatment and Order Quantity on Profitability

Notes. ** p≤0.01. The omitted category is Condition 1, i.e. the high margin, no sharing condition.

Table A3:	Robustness	Test	with	Decision	Support
					~~ppo-t

	Coef.	SE		
C3*: High Margin – Low Transfer Price	-14.56**	(5.18)		
C4*: High Margin – High Transfer Price	-7.42	(5.09)		
Average Demand (≤200)	0.73**	(0.35)		
Average Demand (>200)	0.42^{**}	(0.30)		
Constant	78.96	(67.11)		
N	60			
R ²	29%			

Notes. ** $p \le 0.01$. C3* indicates that this condition is similar to C3 with the addition of decision support. The omitted category is C1*. All coefficients measure differences in order quantities compared to this condition.

Figure A1: Comparison of Order Quantities in the Sharing and No-sharing Situations under Risk-Aversion



Figure A2: Screen Shots of the Base Study - Initial Stocking Decision



Revenue Information:			Cost Information:						You		
ltem	Description	Revenue Per Unit	Units applicable	Subtotal	ltem	Description	Cost Per Unit	Units applicable	Subtotal	Unite Ordered	210
Retail Price	The amount you will recieve for selling a unit	40	210	8400	Procurement Cost	The amount you will pay to procure a unit from wholesaler.	20	210	4200	onto ordered.	210
Transfer Price	The amount you will recieve for transfering a unit	35			Lost Sales Cost	The amount you will pay for any units you neither sell nor transfer to the other player.	0	33	0	Demand	243
	to the other player. The amount you				Transshipment Cost	The per unit amount it will cost to ship to the other player.	0			Revenue	8400
Salvage Price	will recieve for any unit that you do not sell or transfer to the other player.	10	0	0	Transfer Price	The per unit amount you will pay to receive a unit from the other player.	35			Projected Cost	4200
Total Revenue:				8400	Total Costs:				4200	Projected Profit	4200
You Ordered 33 units UNDER the current demand.											
Your supplementary order from other retailor (This order may be filled, partially filled, or not filled at all):											
Salant .											

Figure A3: Screen Shots of the Base Study - Requesting Decision



Revenue Information: Cost Information:								You					
ltem	Description	Revenue Per	Units	Subtotal	Item	Description	Cost Per Unit	Units applicable	Subtotal				
	The amount you		appreubio			The amount you have					Units Ordered:	200	
Retail Price	will receive for each unit sold to consumers.	40	105	4200	Procurement Cost	to pay per unit originally ordered.	20	200	4000				
Transfer Price	The amount you will receive from the other player for each unit transferred.	35	-		Lost Sales Cost	The amount you have to pay per customer demand not met.	0	0	0		Demand	105	
Salvage Price	The amount you will receive for each unit that you do not sell or transfer.	10	95	950	Transshipment Cost	The cost you have to pay to transfer a unit to the other player.	0				Projected Revenue	5150	
-	-			-	Transfer Price	The per unit amount it you pay to recieve a unit from the other player.	35	-	-		Projected Cost	4000	
Total Revenue:	-	-		5150	Total Costs:			-	4000		Projected Profit	1150	
	Image:												

Appendix II: Representative Instructions for Experiment (Base Study, Condition 3)

This is an experiment in inventory decision making. During the experiment, you will play a game from which you will receive cash earnings based on your performance. Upon completion of the game, you will be paid your total earnings in cash plus a \$5 show-up fee. If you have any questions, feel free to raise your hand and we will assist you. Please do not communicate with other participants in the game,

and please refrain from using your cell phones.

Description of the Game. You are a retailer selling a product. To be able to sell the product, you must first place inventory orders in advance of knowing the exact customer demand for that period. If the realized customer demand during the selling period is less than the number of units you order (Order > Demand), there will be some units that you ordered but cannot sell. While leftover units can be salvaged at a price which is below the cost to order units, these units do not carry over to the next selling period. Conversely, if the customer demand is greater than the number of units ordered (Demand > Order), there will be some consumer demand that cannot be met.

There will be a total of 30 selling periods for your product. You will be randomly matched with a new player in each selling period. After demand in a selling period is revealed, you will get a chance to request additional inventory from that other player, and the other player will be provided with a similar opportunity. Your goal is to maximize your total profits over all 30 selling periods.

Revenue/Cost Information

Retail Price (the amount you will receive for each unit sold to consumers):	40
Transfer Price (the amount you receive from the other player for each unit transferred):	15
Salvage Price (the amount you will receive for each unit that you do not sell or transfer):	10
Procurement Cost (the amount you have to pay per unit originally ordered):	20
Lost Sales Cost (the amount you have to pay per customer demand not met):	0
Transhipment Cost (the cost you have to pay to transfer a unit to the other player):	0

The other players with whom you will be matched face the same price/cost parameters.

Demand Information. Demand in each selling period is drawn from a Normal Probability Distribution. Average demand is 200, with a standard deviation of 70.71. Demand in one selling period is uncorrelated with demand in the next period.

Payoffs. The computer program will calculate your profits in every selling period. Your cash earning in each selling period will be approximately

Real Cash Earning () = Profits x 0.0095%.

For example, if your Profits in a turn are equal to 3500, your Real Cash Earnings that period will be about 33 cents. Note that your cash earnings will be negative if your profit for that round goes below 0. Your total cash earnings in this experiment will be \$6 plus your total cash earnings for the 30 decisions rounds. The computer program will calculate your cash earnings for each selling period and your total cash earnings for all the periods you have completed. We will pay your total earnings upon the completion of the game by rounding them to the highest dollar amount and making sure that you receive at least \$6 (show-up fee). The maximum Earnings that we will pay out per participant is \$24.

Appendix III: Proof of Proposition 1

To solve the optimal order $Q_s(\alpha, \beta)$, we partition the demand space (D_1, D_2) as shown below.



Given realized demand (D_1, D_2) in each of the regions of the above figure, the corresponding transshipment (T_{ij}, T_{ji}) , shortage (Z_i, Z_j) , leftover (L_i, L_j) and sales (R_i, R_j) are summarized below.

	Transship(T_{ij}, T_{ji})	Shortage (Z_i, Z_j)	Leftover (L_i, L_j)	Sales (R_i, R_j)
\varOmega_1	0,0	0,0	$Q_i - D_i, Q_j - D_j$	D_i, D_j
Ω_2	$0, D_i - Q_i$	0,0	$0, Q_i + Q_j - D_i - D_j$	D_i, D_j
\varOmega_{41}	$0, D_i - Q_i$	$0, D_i + D_j - Q_i - Q_j$	0,0	$D_i, Q_j + Q_i - D_i$
Ω_{42}	$0, \alpha(D_j - D_j)$	$D_i - Q_i - \alpha (Q_j - D_j),$	0,0	$Q_i + \alpha (Q_j - D_j),$
		$D_j - Q_j + \alpha (Q_j - D_j)$		$Q_j - \alpha(Q_j - D_j)$
$\varOmega_{3,51}$	$\beta(D_j - Q_j), 0$	$0, (1-\beta)(D_j - Q_j)$	$Q_i - D_i - \beta (D_j - Q_j), 0$	$D_i, Q_j + \beta (D_j - Q_j)$
Ω_{52}	$Q_i - D_i, 0$	$0, D_j + D_i - Q_j - Q_i)$	0,0	$D_i, Q_j + Q_i - D_i$
Ω_6	0,0	$D_i - Q_i, D_j - Q_j$	0,0	Q_i, Q_j

Hence, player *i*'s expected profit is

$$\Pi_i(Q_i,Q_j;\alpha,\beta) = \int_{\Omega} (r_i R_i + s_i L_i + c_{ij} T_{ij} - c_{ji} T_{ji}) dD_i dD_j - c_i Q_i$$

Thus, we have the following first-order condition with respect to Q_i ,

$$\frac{\partial \Pi_i (Q_i, Q_j; \alpha, \beta)}{\partial Q_i} = v_i Pr(\Omega_{42} + \Omega_6) + c_{ij} Pr(\Omega_2 + \Omega_{41} + \Omega_{52}) + s_i Pr(\Omega_1 + \Omega_3 + \Omega_{51}) - c_i$$

Since both players are symmetric, we have $Q_i = Q_j = Q_s(\alpha, \beta)$ at equilibrium. Thus, we can solve $Q_s(\alpha, \beta)$ through the above first-order condition $\partial \Pi_i(Q_i, Q_j; \alpha, \beta)/\partial Q_i = 0$. Next, we show that $Q_s(\alpha, \beta)$ decreases as α increases and decreases as β decreases.

Based on the derivative of implicit function, we have $\frac{\partial Q(\alpha,\beta)}{\partial \alpha} = -\frac{\partial^2 \Pi}{\partial Q \partial \alpha} / \frac{\partial^2 \Pi}{\partial Q^2}$. Since $\frac{\partial^2 \Pi}{\partial Q^2} < 0$ at the optimal $Q_s(\alpha,\beta)$, it suffices to show that $\frac{\partial^2 \Pi}{\partial Q \partial \alpha} < 0$. We have

$$\frac{\partial^2 \Pi}{\partial Q \partial \alpha} = v_i \frac{\partial Pr(\Omega_{42})}{\partial \alpha} + c_{ij} \frac{\partial Pr(\Omega_{41})}{\partial \alpha} = -(v_i - c_{ij}) \frac{\partial Pr(\Omega_{41})}{\partial \alpha} < 0,$$

Therefore, we have proved that $Q_s(\alpha, \beta)$ decreases as α increases.

Likewise, we can show that

$$\frac{\partial^2 \Pi}{\partial Q \partial \beta} = s_i \frac{\partial Pr(\Omega_{51})}{\partial \beta} + c_{ij} \frac{\partial Pr(\Omega_{52})}{\partial \beta} = (c_{ij} - s_i) \frac{\partial Pr(\Omega_{52})}{\partial \beta} > 0$$

Thus, we have proved that $Q_s(\alpha, \beta)$ decreases as β decreases. By definition, for rational players, $Q_s^* = Q_s(\alpha = 1, \beta = 1)$. Hence, when $\alpha = 1$ and $\beta < 1$, we have $Q_s^* > Q_s(\alpha, \beta)$; and when $\alpha > 1$ and $\beta = 1$, we have $Q_s^* > Q_s(\alpha, \beta)$.

Appendix IV: Structural Estimation of the Behavioral Model

Given a set of behavioral parameters $\theta = (\alpha, \beta, \delta_o, \delta_u)$, we can solve the equilibrium order quantity predicted by the behavioral model. In particular, for the no-sharing case, we can solve Q_n through the modified critical fractile formula $Q_n = F^{-1} \left(\frac{v - c + \delta_u}{v - s + \delta_u + \delta_o} \right)$.

For the inventory sharing case, we can solve the predicted order quantity $Q_s(\theta)$ through the following first-order condition,

$$\frac{\partial \Pi_s}{\partial Q_s} = (\nu + \delta_u) \operatorname{Pr}(\Omega_{42} + \Omega_6) + c_{ij} \operatorname{Pr}(\Omega_2 + \Omega_{41} + \Omega_{52}) + (s - \delta_o) \operatorname{Pr}(\Omega_1 + \Omega_3 + \Omega_{51}) - c = 0$$

For the fixed role treatment, we can solve the predicted order quantity for the recipient, i.e., $Q_s^r(\theta)$, through the following first-order condition,

$$\frac{\partial \Pi_s^r}{\partial Q_s^r} = (v + \delta_u) \operatorname{Pr}(\Omega_{42} + \Omega_6) + c_{ij} \operatorname{Pr}(\Omega_2 + \Omega_{41}) + (s - \delta_o) \operatorname{Pr}(\Omega_1 + \Omega_3 + \Omega_5) - c = 0,$$

and solve the predicted order quantity for the source, i.e., $Q_s^s(\theta)$, through the following first-order condition,

$$\frac{\partial \Pi_s^s}{\partial Q_s^s} = (v + \delta_u) \operatorname{Pr}(\Omega_2 + \Omega_4 + \Omega_6) + c_{ij} \operatorname{Pr}(\Omega_{52}) + (s - \delta_o) \operatorname{Pr}(\Omega_1 + \Omega_3 + \Omega_{51}) - c = 0$$

Note that $\Omega_4 = \Omega_{41} \cup \Omega_{42}$, $\Omega_5 = \Omega_{51} \cup \Omega_{52}$. Essentially, given a set of parameters $\theta = (\alpha, \beta, \delta_o, \delta_u)$, we can solve predicted order quantity (Q_n, Q_s, Q_s^r, Q_s^s) . The observed order quantities in the experiments, however, do not perfectly coincide with the model predictions due to various noises at players' decision processes. To capture the variation of the observed order quantities, we assume that the average order quantity placed by a player *i* in the experiments is normally distributed with mean specified by the behavioral model predication, i.e.,

$$q_{i,n} \sim N(Q_n(\delta_o, \delta_u), \sigma_n^2)$$

$$q_{i,s,c_{ij}} \sim N\left(Q_{s,c_{ij}}(\alpha, \beta, \delta_o, \delta_u), \sigma_{sc_{ij}}^2\right)$$

$$q_{i,s,c_{ij}}^r \sim N\left(Q_{s,c_{ij}}^r(\alpha, \delta_o, \delta_u), \left(\sigma_{s,c_{ij}}^r\right)^2\right)$$

$$q_{i,s,c_{ij}}^s \sim N\left(Q_{s,c_{ij}}^s(\beta, \delta_o, \delta_u), \left(\sigma_{s,c_{ij}}^s\right)^2\right)$$

where $c_{ij} \in \{15, 35\}$ indicates the high or low transfer price.

The noise parameters $\sigma^2 = (\sigma_n^2, \sigma_{sc_{ij}}^2, (\sigma_{s,c_{ij}}^r)^2, (\sigma_{s,c_{ij}}^s)^2)$ are specific to each treatment, while the behavioral parameters $\theta = (\alpha, \beta, \delta_o, \delta_u)$ are specified to be common across all treatments because they affect all players across these treatments. We estimate parameters θ and σ^2 through maximum likelihood estimation (MLE), as described below.

For shorthand, let *t* denote the index for different treatments and N_t denote the number of players in treatment *t*. Let $i = 1, ..., N_t$ be the index for player in each treatment. Thus, the MLE is to find θ and σ^2 so as to maximize the log-likelihood function for the observed order quantity, i.e.,

$$\max_{\theta,\sigma^2} \ln(L(\theta,\sigma^2)) = -\frac{1}{2} \ln(2\pi) \sum_t N_t + \sum_t \sum_{i=1}^{N_t} \left(-\frac{1}{2} \ln(\sigma_t^2) - \frac{(q_{it} - Q_t(\theta))^2}{2\sigma_t^2} \right)$$
$$= -\frac{1}{2} \ln(2\pi) \sum_t N_t - \frac{1}{2} \sum_t N_t \left(\ln(\sigma_t^2) + \frac{\sigma_{q_t}^2 + (\bar{q}_{it} - Q_t(\theta))^2}{\sigma_t^2} \right)$$

where $\bar{q}_{it} = \frac{\sum_{i=1}^{N_t} q_{it}}{N_t}$, i.e., the average observed order for treatment *t*, and $\sigma_{q_t}^2 = \frac{\sum_{i=1}^{N_t} (q_{it} - \bar{q}_t)^2}{N_t}$, i.e., the variance of the observed order for treatment *t*.

Taking the first derivative of the log-likelihood against σ^2 gives

$$\frac{\partial \ln L}{\partial \sigma_t^2} = -\frac{N_t}{2} \left(\frac{1}{\sigma_t^2} - \frac{\sigma_{q_t}^2 + \left(\bar{q}_{it} - Q_t(\theta)\right)^2}{\sigma_t^2} \right)$$

Thus, given θ , we can solve the corresponding MLE estimator of σ_t^2 as

$$\sigma_t^2 = \sigma_{q_t}^2 + \left(\bar{q}_{it} - Q_t(\theta)\right)^2$$

Substituting σ_t^2 back into the log-likelihood function, we have

$$\ln L(\theta) = -\frac{1}{2}(\ln(2\pi) + 1)\sum_{t} N_{t} - \frac{1}{2}\sum_{t} N_{t} \left(\ln \left(\sigma_{q_{t}}^{2} + \left(\bar{q}_{it} - Q_{t}(\theta) \right)^{2} \right) \right)$$

Thus, we can obtain the MLE estimator of θ through $\max_{\theta} \ln L(\theta)$, which can be done using commercial optimization software such as Matlab.