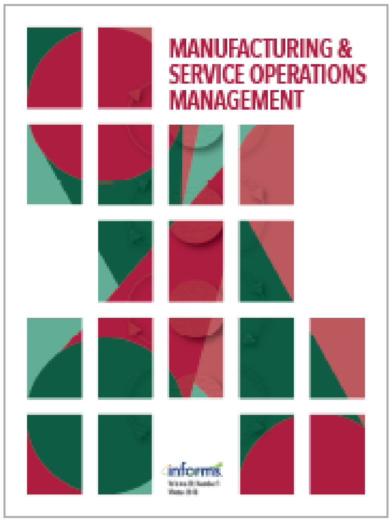

Manufacturing & Service Operations Management

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

Centralized or Decentralized Transfer Prices: A Behavioral Approach for Improving Supply Chain Coordination

Elena Katok, Sebastián Villa

informs
http://pubsonline.informs.org/journal/msom
MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–16
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Centralized or Decentralized Transfer Prices: A Behavioral Approach for Improving Supply Chain Coordination

Elena Katok,[a] Sebastián Villa[b,c]

[a] Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; [b] Operations and Decision Technologies, Kelley School of Business, Indiana University, Bloomington, Indiana 47405; [c] School of Management, University of Los Andes, Bogotá 110311, Colombia

Contact: ekatok@utdallas.edu, https://orcid.org/0000-0002-7037-7896 (EK); svillab@iu.edu, https://orcid.org/0000-0002-8708-2576 (SV)

Received: February 23, 2019
Revised: May 24, 2020; August 19, 2020
Accepted: September 1, 2020
Published Online in Articles in Advance: March 8, 2021

https://doi.org/10.1287/msom.2020.0957

Abstract. ***Problem definition***: We study supply chain coordination in a setting with transshipment. We use centralized and decentralized transfer prices as a way to increase supply chain coordination. ***Academic/practical relevance***: The ability to transship can improve channel efficiency by improving the match between supply and demand. We study how human decision makers behave in this setting and provide clear insights to improve coordination. ***Methodology***: We use controlled laboratory experiments with financially incentivized human subjects. We study a broad set of critical ratios under both decentralized and centralized transfer-price settings. In the decentralized transfer-price setting, retailers negotiate a transfer price. In the centralized transfer-price setting, we use two different approaches: theoretical and behavioral transfer price. Both approaches suggest opposite recommendations. ***Results***: Analytically, the optimal transfer price should depend on the critical ratio; but results from the decentralized setting show that participants set prices as if they ignore the critical ratio and instead focus on splitting potential profit from transshipped units in half. However, there is a positive relationship between transfer prices and ordering decisions. Moreover, generalizing the pull-to-center effect, we find that subjects do not place orders that coordinate the supply chain. For the centralized setting, we find that using the theoretical approach does not coordinate ordering decisions and does not improve decisions compared with the decentralized setting. The behavioral approach suggests a transfer price close to product selling price for a high critical ratio and a transfer price below product cost for a low critical ratio. These recommendations lead to coordinating ordering decisions. ***Managerial implications***: We draw two practical conclusions from our research. First, transshipments are unambiguously beneficial, resulting in higher profitability, and, when feasible, should be encouraged. Second, when possible, transfer prices should be set centrally but taking into account subjects' behavior. Otherwise, price negotiation might lead to better performance.

Funding: Villa gratefully acknowledges support from the Swiss National Science Foundation (SNSF) [Grant P1TIP1_164985].
Supplemental Material: The online appendices are available at https://doi.org/10.1287/msom.2020.0957.

Keywords: behavioral operations • Transshipments • pricing and revenue management • inventory theory and control

1. Introduction

Most products are sold by multiple retailers. These retailers can be independently operated, such as car dealerships, or centrally owned, such as large retail chains with multiple locations (e.g., Wal Mart, Kruger, Walgreens, CVS, etc.). Uncertain customer demand, coupled with long production and distribution lead times, generally causes costly mismatch between supply and demand for each individual retailer. The costs of such mismatch became exposed during the 2020 COVID-19 pandemic, when retailers experienced repeated stockouts of such staples as toilet paper and flour and hospitals were unable to purchase personal protective equipment. But costly mismatches are not confined to times of disruptions. They are costly during normal times as well, affecting service levels and creating excess inventory (Bolton et al. 2016, Castañeda et al. 2019).

Fortunately, today retailers have access to information technology that facilitates tracking inventory and other communication (Axsäter 2003), making it both profitable and feasible to reallocate inventory when possible. This exchange of product among retailers is called *transshipment* (Krishnan and Rao 1965, Rudi et al. 2001). Transshipments are common in many industries, such as semiconductor manufacturing (Kranenburg and van Houtum 2009), fashion apparel (Dong and Rudi 2004), and financial and electricity markets (Werdigier and Dougherty 2007), just to name a few examples. The motivation of our work is to understand how to effectively design transshipment mechanisms. This is important because transshipment arrangements can increase service levels, reducing excess inventory, and generally improve not only the profitability of retailers and manufacturers and the satisfaction of consumers but also the welfare of society as a whole.

Efficiency of a retail channel with transshipment is achieved when retailers place coordinating orders, that is, orders that maximize the profit of the channel as a whole (Rudi et al. 2001, Dong and Rudi 2004, Sošić 2006). Even when retail locations are centrally owned, it is common to leave ordering decisions up to individual stores, because local managers are likely to have location-specific information.

When orders are placed by local retailers, there is a relationship between order quantities and transfer prices—the amount the retailer with excess inventory charges the retailer with excess demand. Therefore, we aim to understand how the orders that human decision makers place are affected by transfer prices and how transfer prices should be set centrally. When retailers are independent, even transfer prices are likely to require retailers to negotiate; therefore, we also intend to understand how centralized and decentralized pricing arrangements differ in their performance.

When multiple retail locations are centrally owned, transshipment policies are usually set centrally. Shao et al. (2011, p. 361) refer to this arrangement as a "chain store" and contrast it to more decentralized arrangements that some manufacturers encourage among their independent retailers. Companies such as Caterpillar, John Deere, and General Motors actively promote these arrangements among their independent dealers (Zhao et al. 2005) as do some tool manufacturers, such as Bosch (Rudi et al. 2001) and Okuma America Corporation (Narus and Anderson 1996). Even though these manufacturers try to influence transfer prices, these prices are usually set by the retailers themselves (Narus and Anderson 1996, Zhao et al. 2005, Shao et al. 2011). A very different example of transfer prices that are set through negotiations include the energy markets in Asia and Europe (Arsu 2009) as well as in Latin America (CREG 2009). Regardless of whether transfer prices are set centrally or not, they are typically set prior to ordering decisions, so orders are affected by transfer prices.

Analytical models on transshipment differ in their assumptions about the structure of the channel and the role of the manufacturer. When firms are vertically integrated, the focus is primarily on inventory management (Krishnan and Rao 1965, Tagaras 1989, Robinson (1990). A model that includes a manufacturer's wholesale price decision when selling to a retailer with multiple locations is analyzed by Dong and Rudi (2004) and extended by Zhang (2005). There are several papers that model transshipment among independent retailers, either symmetric or asymmetric (Rudi et al. 2001, Hu et al. 2007) in a single shot game and in a multiperiod game (Zhao et al. 2005), focusing on equilibrium inventory levels. Shao et al. (2011) analyze the case with a manufacturer who sets the wholesale price when selling to independent retailers and specifically focus on the effect of the transfer price on profits. They find that a high transfer price benefits the manufacturer.

These analytical models, like most analytical models in operations management, assume that decision makers behave optimally. Testing analytical models in a controlled laboratory setting is a good first step toward better understanding whether these models are, or can be, useful in practice, because laboratory experiments can be designed to make sure that most of the model's assumptions are fulfilled. There is, however, a growing body of behavioral operations literature reporting that, even in a controlled laboratory setting, people do not order inventory optimally, when faced with a newsvendor problem, which is the simplest inventory management problem with stochastic demand (see, for example, Schweitzer and Cachon 2000, Benzion et al. 2008, and Bolton and Katok 2008 for a review). These behavioral deviations from optimality are quite systematic; even though Schweitzer and Cachon (2000) show that these deviations are not consistent with most standard behavioral explanations (risk aversion, loss aversion, etc.), there are several relatively new behavioral models that have been proposed that do a good job in organizing these data. Ho et al. (2010) develop a model with asymmetric disutility from supply and demand mismatch; Becker-Peth et al. (2013) propose a model that combines anchoring and aversion to leftovers (Becker-Peth et al. 2013 refer to it as loss aversion); Long and Nasiry (2015) extend prospect theory to include a reference point that depends on the order quantity.

Because our objective is to provide insight into designing transshipment arrangements, we start with a laboratory experiment with human subjects in a setting with independent retailers that compares channel performance with and without transshipments. When transshipments are not allowed, participants only place orders; but when transshipments are allowed, participants negotiate transfer prices prior to placing orders. We find that ordering decisions qualitatively resemble behavior without transshipments; allowing transshipments significantly improves profitability, although the transfer prices that our participants negotiate do not depend on the critical ratio (CR). Rudi et al. (2001) show that there exist transfer prices and order quantities that should, in equilibrium, coordinate the retail channel; those coordinating transfer prices decrease in the CR—a regularity that we fail to observe in our data.

Given that decentralized transshipment, although significantly improving the retailers' profits, does not deliver either coordinating transfer prices or coordinating order quantities, we next ask whether a centralized setting can produce better results. To this end, we extend the models of the newsvendor behavior to a setting with transshipments and with transfer prices that are set centrally. To calibrate these models out-of-sample, we assemble a large data set of newsvendor decisions from four separate studies and use these data to estimate parameters of two behavioral models from the literature (Ho et al. 2010, Becker-Peth et al. 2013). We then use these models to estimate the optimal centralized transfer prices for a wide range of critical ratios. It turns out that the two behavior models result in similar optimal transfer prices; these behaviorally informed transfer prices increase in the CR, which is the opposite of the relationship between optimal transfer prices and the CR based on the standard equilibrium model.

We proceed to validate these behavioral findings using a set of experiments under the centralized transshipment setting, using new treatments that allowed us to compare the performance of centralized transfer prices determined using a behavioral model to the performance under transfer prices determined using a standard equilibrium model. We found that behaviorally informed transferred prices deliver average orders that are significantly closer to coordinating than equilibrium transfer prices do.

There is a growing interest in behavioral studies to understand the transshipment problem. Bostian et al. (2008) study how retailers benefit from transshipments when they make ordering decisions with automated transshipments. Villa and Castañeda (2018) study the effect of communication and best response strategies in a newsvendor problem with automated transshipments. Zhao et al. (2016) examine inventory sharing effectiveness and find that retailers tend to understock inventory, more so when there is transshipment opportunity, and therefore allowing transshipments does not generally improve efficiency. Chen and Li (2020) consider a setting with voluntary transshipments. They find that prices and order quantities are generally set below coordinating levels.

Our work is distinct from the abovementioned papers and contributes to the understanding of transshipment in several ways. First, we analyze and compare coordination in settings in which transfer prices are set centrally and in which they are negotiated through a free bargaining process. Second, we build on analytical and behavioral models to provide empirical recommendations about how transfer prices should be set to elicit coordinated order quantities under a broad set of CRs. Third, we directly examine the effect of transfer prices on ordering decisions, which is important because appropriate understanding of this relationship can improve stock polices, increase profits, and improve coordination (Herer et al. 2006, Li and Ryan 2017).

In the next section, we summarize the analytical background for the newsvendor problem with transshipments and derive standard equilibrium predictions. In Section 3, we explain our experimental design, protocol, and theoretical benchmarks. Section 4 shows results for a decentralized setting. In Section 5, we provide behavioral and analytical predictions for the centralized setting and report experimental results. In Section 6, we summarize our findings and discuss their implications.

2. Standard Theory

We consider a simple model with one supplier and two retailers. Retailers face the same stochastic demand distribution (d_i) and sell an identical product under a complete pooling policy. Complete pooling means that if at the end of the selling season one retailer has excess stock, whereas the other is short of units, the number of units transshipped will be the minimum of the excess and the shortage (Tagaras 1989). Each retailer R_i (with $i = 1, 2$) buys his or her items (q_i) from the supplier at a cost of c_i and receives a revenue r_i for each unit sold to the final customers, and leftovers at the end of the whole season do not generate any additional value to the retailers ($b_i = 0$). The unit transfer price (t_{ij}, with $i, j = 1, 2$ and $i \neq j$) is paid by the retailer receiving the transshipped units to the other retailer.

We build on Rudi's et al. (2001) study to provide some analytical intuition about the standard theory for situations in which transfer prices are set before the ordering decision. We build on the typical newsvendor problem and summarize the most relevant analytical results for a *centralized* system with transshipments and for a system with independent (*decentralized*) retailers.

2.1. Isolated Retailers: Newsvendor Problem

We consider a system in which two isolated retailers place ordering decisions to a supplier who has unlimited production capacity. In this system, transshipments among retailers are not possible. Therefore, the optimal ordering decision for each retailer R_i will be determined by the well-known CR.

$$P(d_i < q_i) = \alpha_i(q_i) = \left(\frac{r_i - c_i}{r_i}\right), \qquad (1)$$

where $\alpha_i(q_i)$ represents the probability of the R_i of facing excess inventory at the end of the selling period.

2.2. Centralized System

The centralized system considers transshipments at an intrafirm level in which order quantities for both retailers are centrally set to maximize the aggregate profits. Following the model in Rudi et al. (2001), the sequence of events in this transshipment problem is as follows: (i) retailers place ordering decisions, (ii) supplier fills these orders, (iii) final customer demand takes place, (iv) demand is satisfied, and (v) potential transshipments among retailers and satisfaction of additional final customer demand take place. Notice that transshipments are feasible only when one retailers faces a shortage and the other retailer has excess inventory (Krishnan and Rao 1965, Rudi et al. 2001). There are four states of the world in a transshipment problem. The first state of the world considers only the situation in which each R_i faces excess inventory $\alpha_i(q_i) = \Pr(d_i < q_i)$. This is the typical analysis for a newsvendor problem. The second state of the world considers the situation in which R_j has unsatisfied demand and R_i has excess inventory but not enough to satisfy R_j's unsatisfied demand. The probability of occurrence of this state is represented by $\beta_i(q_i, q_j) = \Pr(q_i + q_j - d_j < d_i < q_i)$. The third state of the world considers the situation in which R_i has unsatisfied demand and R_j has excess inventory, but this excess inventory is enough to satisfy R_i's unsatisfied demand. The probability of occurrence of this state is represented by $\gamma_i(q_i, q_j) = \Pr(q_i < d_i < q_i + q_j - d_j)$. Finally, the fourth state considers only the situation in which each R_i faces excess demand ($\Pr(d_i > q_i)$). Observe that in situations in which both retailers face either surplus or shortage, transshipments do not materialize. Figure 1 provides a graphical explanation about the different states of the world (and their associated probabilities) that take place in a supply chain with transshipments for R_1 and R_2.

Building on these states of the world and probabilities, and assuming that the joint distribution over the demand is continuously differentiable, the order quantities (q_i^c, q_j^c) that maximize the total expected profit for a centralized system can be obtained by solving Equation (2), such that it holds for $i, j = 1, 2$. Please refer to the Rudi et al. (2001) paper for further explanation about this model.

$$\alpha_i(q_i) - \beta_i(q_i, q_j) + \gamma_i(q_i, q_j) = \left(\frac{r_i - c_i}{r_i}\right) \tag{2}$$

Notice that Equation (2) is just an adjustment of Equation (1). The second term on the left-hand side attempts to increase q_i to account for the possibility of transshipments from i to j. Then, when R_i perceives a higher probability of selling units to R_j at the end of the selling period, R_i has more incentives to place higher q_i as an additional source of revenue. Similarly, the third term represents the adjustment in q_i due to the probability of transshipping units from j to i. In this case, when R_i perceives a higher probability of receiving units from R_j, R_i has an incentive to place lower q_i as a way of reducing potential excess inventory at the end of the selling period. Finally, notice that for a system with no transshipments involved ($\beta_i(q_i, q_j) = 0$ and $\gamma_i(q_i, q_j) = 0$), the solution is that of the standard newsvendor problem (Equation (1)).

Figure 1. States of the World and Probabilities for a Transshipment Problem (Rudi et al. 2001)

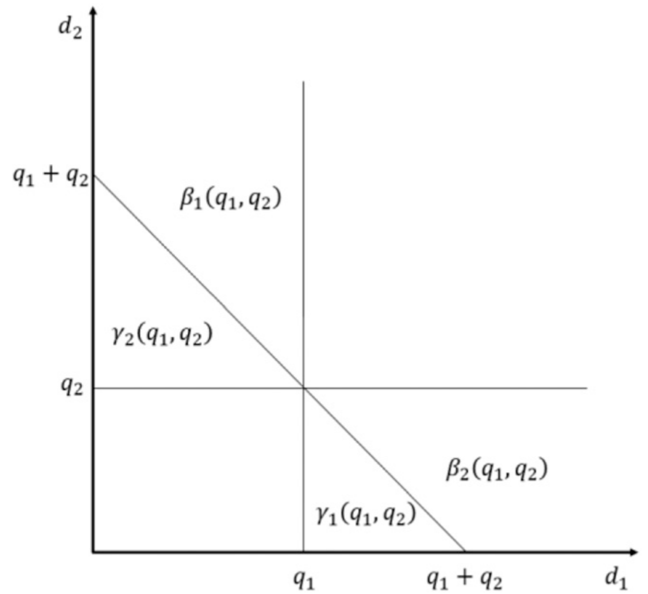

2.3. Decentralized System

A decentralized system considers transshipments at an intra- or interfirm level in which retailers' decisions are locally defined to maximize individual profits. Retailers need to coordinate their operations while guaranteeing their individual profitability. Given the nature of this decentralized system, R_i's profit depends on the transfer prices. Rudi et al. (2001) show that there is a unique set of transfer prices t_{ij}^* and order quantities (q_1^*, q_2^*) that can coordinate the whole system to optimize the total channel profit (Hu et al. 2007).

We start by accounting for exogenously defined transfer prices (t_{ij}, t_{ji}) to define a profit function for each retailer. The maximization of this profit function leads to a reaction function $q_i(q_j)$ defining the desired inventory policy for each retailer R_i. The Nash equilibrium (q_i^d, q_j^d) for a decentralized system at t_{ij} is obtained by solving the reaction function presented on Equation (3), such that it holds both for i, $j = 1, 2$.

$$\alpha_i(q_i) - \beta_i(q_i, q_j)\left(\frac{t_{ij}}{r_i}\right) + \gamma_i(q_i, q_j)\left(\frac{r_i - t_{ji}}{r_i}\right) = \left(\frac{r_i - c_i}{r_i}\right). \quad (3)$$

Notice that Equation (3) is quite similar to Equation (2). However, Equation (3) accounts for the fraction of additional profit coming from the units sold to (second term on the left-hand side of Equation (3)) or bought from (third term on the left-hand side of Equation (3)) the other retailer.

To define the set of transfer prices that maximize retailers' joint profits, we equate the left-hand sides of Equation (2) and Equation (3) and isolate t_{ij}. In this way, Rudi et al. (2001) estimate the transfer price that makes the solution of the decentralized system equal to the solution of the centralized system. The coordinating transfer price for a decentralized system that yields to the joint optimal solution is given by Equation (4).

$$t^*_{ij} = \frac{r_j\beta_i(q^c_i, q^c_j)\beta_j(q^c_j, q^c_i) - r_i\beta_j(q^c_j, q^c_i)\gamma_i(q^c_i, q^c_j)}{\beta_i(q^c_i, q^c_j)\beta_j(q^c_j, q^c_i) - \gamma_i(q^c_i, q^c_j)\gamma_j(q^c_j, q^c_i)}. \quad (4)$$

Using the unique set of transfer prices t^*_{ij} to solve Equation (3) leads to the estimation of the optimal order quantities (q^*_1,q^*_2). This solution optimizes the aggregate profit of a decentralized system, which is equal to the solution to a centralized system. Therefore, it is sufficient to derive transfer prices that induce retailers to choose optimal order quantities, as defined in the centralized system (Equation(2)). This centralized solution can be implemented in a decentralized system with a particular set of transfer prices.

We provide an analytical interpretation of the relationship between the coordinating transfer prices and the CR by performing an additional analysis of the reaction function for the decentralized system (Equation (3)) evaluated at the optimal order quantities (q^*_1,q^*_2) and the optimal transfer prices (t^*_{ij}). We assume symmetric retailers ($t^*_{ij} = t^*_{ji}$) and isolate t^*_{ij} to find the following mathematical relationship:

$$t^*_{ij} = \frac{r_i\,\alpha_i(q^*_i)}{\beta_i(q^*_i, q^*_j) + \gamma_i(q^*_i, q^*_j)} + \frac{r_i\,\gamma_i(q^*_i, q^*_j)}{\beta_i(q^*_i, q^*_j) + \gamma_i(q^*_i, q^*_j)} - \frac{r_i}{\beta_i(q^*_i, q^*_j) + \gamma_i(q^*_i, q^*_j)}\left(\frac{r_i - c_i}{r_i}\right). \quad (5)$$

Given that $\beta_i(q^*_1,q^*_2)$, $\gamma_i(q^*_1,q^*_2)$ and r_i are expected to be always positive for a system with transshipments, Equation (5) shows that there is a negative relationship between the coordinating transfer prices (t^*_{ij}) and the CR ($(r_i - c_i)/r_i$): the larger the CR, the smaller the coordinating transfer price.

The analytical models presented in this section provide theoretical benchmarks that can be used to compare against subjects' behavior and to improve decision-making strategies.

3. Experimental Design and Protocol

Our study includes two parts. In the first part, we test two transshipment conditions: the baseline with no transshipment and the decentralized setting in which participants negotiate the transfer price prior to placing orders. The second part includes the centralized setting, in which the transfer price is set centrally and announced to participants before they place their orders. We display the decision sequence in all three conditions graphically in Figure 2. We describe the baseline and decentralized treatments below and will provide further detail on centralized treatments, when we describe the second part of our study.

In all treatments, we use a unit revenue r_i of \$40, a uniform demand distribution $U[0,200]$ for the final customer demand (d_i). In the no transshipment and the decentralized conditions, we vary unit cost at \$30, \$20, and \$10 for the low (*Low*), medium (*Medium*), and high (*High*) CR conditions, respectively. Based on these sets of parameters and using Equations (3) and (4), we compute the per unit transfer prices (t^*_{ij}) that coordinate the system to be \$29, \$20, and \$11 for the *Low*, *Medium* and *High* CR conditions, respectively. In addition, the corresponding ordering quantities (q^*_i) that coordinate the channel are 70, 100, 130 for the *Low*, *Medium*, and *High* CR conditions, respectively. Figure 3 shows the reaction functions (Equation (3)) and the coordinating order quantities evaluated at t^*_{ij}, for each of the different CRs. Analytically, coordinating order quantities depend on CR.

By solving Equation (3) for different values of t_{ij}, Figure 4 shows the relationships between transfer prices and their corresponding equilibria for orders. We observe a linear and increasing relationship, such that if subjects negotiate a higher (lower) price, the equilibria suggest that they should order more (fewer) units (Rudi et al. 2001). Therefore, for each particular transfer price, a unique equilibrium order quantity (the *Best Reply*) exists. Notice that t^*_{ij} and q^*_i (shown as triangles in Figure 4) give the unique equilibrium that

Figure 2. Decision Sequence for Each Experimental Treatment

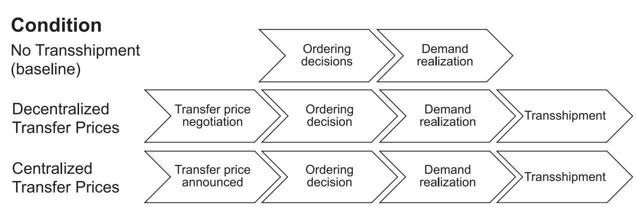

Figure 3. Reaction Curves and Nash Equilibria for (a) 0.25 (*Low*), (b) 0.5 (*Medium*), and (c) 0.75 (*High*) CR Conditions at t_{ij}^*

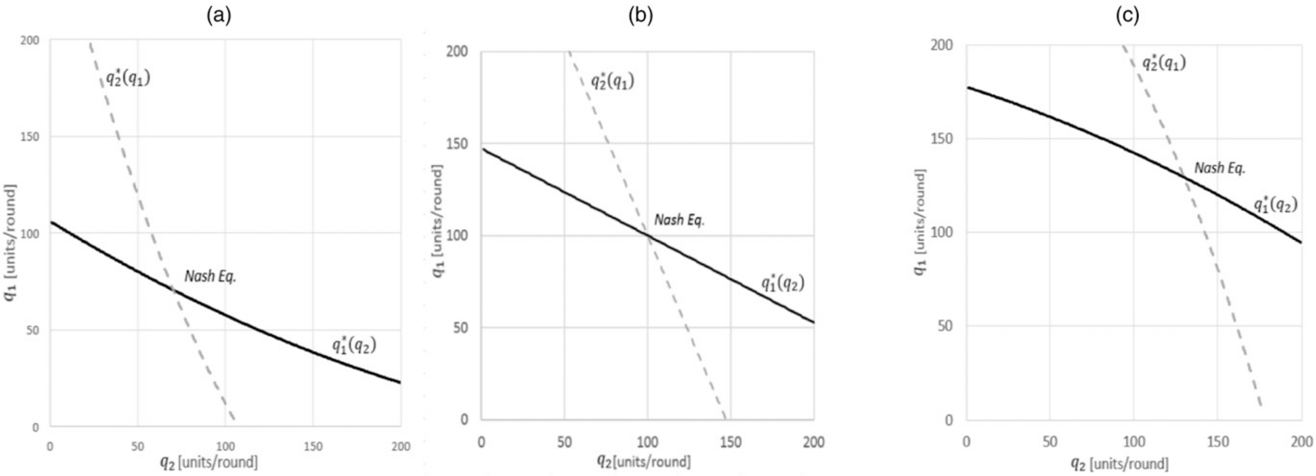

makes the expected profit in the decentralized system equal to the expected profit of a centralized system.

Finally, Equation (1) gives us the well-known optimal ordering quantities with no transshipments involved, which are 50, 100, and 150, for our *Low, Medium* and *High* CR treatments, respectively.

Table 1 lists the six treatments in the first part of our study, with their optimal transfer prices and order quantities. In our experiment, each subject assumes the role of an independent retailer in our one-supplier two-retailer supply chain. During the ordering decision stage of the experiment, each participant orders the number of units they wish to purchase. The decentralized condition also includes a negotiation stage, in which participants negotiate a transfer price per unit to be paid at the end of the round by the retailer with excess demand to the retailer with excess inventory in exchange for additional units. We implement these negotiations as free bargaining: participants make price offers and counteroffers over time until either both of them agree to a price or abort the bargaining process in disagreement. After the transfer price negotiations have been completed, customer demands are randomly and independently drawn. Following demand realization, the transfer units are automatically shipped in the decentralized condition and the profits are calculated.

We designed the second part of our study—the centralized condition—based in part on the results we observed in the first part of the study. We explain the design of the centralized treatments in Section 5.3.

We implemented the experiment using the Software Platform for Human Interaction Experiments (SoPHIE) (Hendriks 2012), which is an internet-based system for implementing laboratory experiments, conducting sessions, and storing data. We recruited participants through the online recruitment system SONA and offered them cash as an incentive to participate. Each treatment included four cohorts of 8 participants, for the total of 32 participants per treatment and 192 participants in in the first part of our study. Our participants were master students in management from a

Figure 4. Optimal Relationship Between Order and Transfer Prices for (a) 0.25 (*Low*,) (b) 0.5 (*Medium*), and (c) 0.75 (*High*) CR Conditions

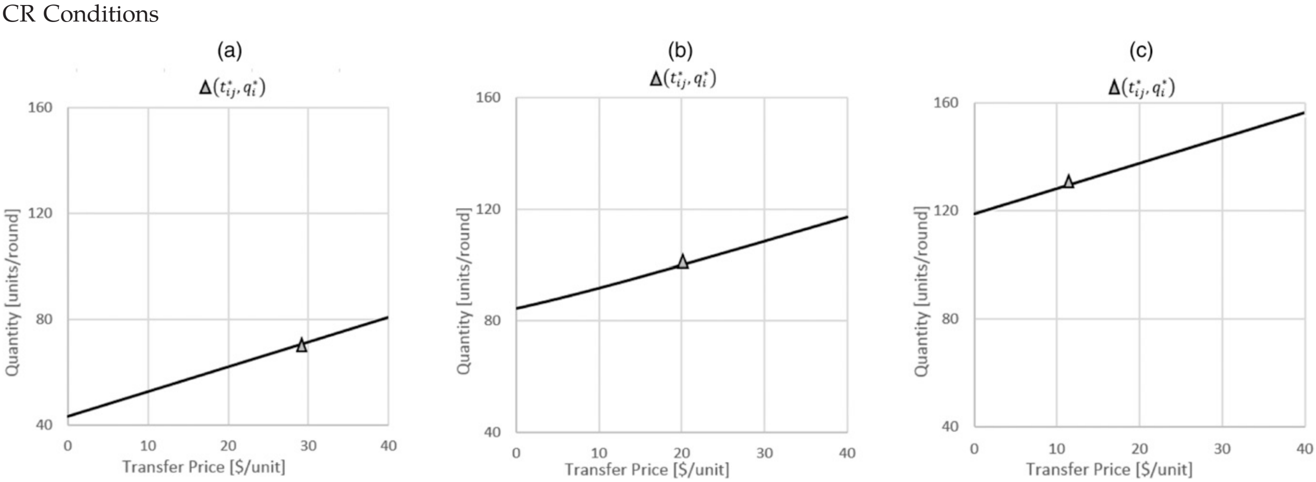

Table 1. Experimental Design Baseline and Decentralized Treatments

		No transshipment	Decentralized
Critical ratio	0.25	$q_i^* = 50$	$t_{ij}^* = 29$
			$q_i^* = 70$
	0.50	$q_i^* = 100$	$t_{ij}^* = 20$
			$q_i^* = 100$
	0.75	$q_i^* = 150$	$t_{ij}^* = 11$
			$q_i^* = 130$

Note. Each treatment included four cohorts of eight subjects (32 participants per treatment).

public U.S. university. Once participants entered the laboratory, they were directed to a computer terminal and had 10 minutes to read the instructions (see Online Appendix 1). After 10 minutes, the experimenter read the instructions aloud to ensure common knowledge about the rules of the game. Any clarification questions were also answered at this point. The same experimenter conducted all sessions.

In the instructions, participants received full information about the distribution of the final customer demand, the retail price ($\boldsymbol{r}$), and the unit purchase cost ($\boldsymbol{c}$). In addition, participants were given the information about the relationship between transfer prices and equilibrium orders.

We used a between-subjects design for the baseline and the decentralized treatments. All sessions lasted 30 rounds; in each round, participants were randomly rematched with another person within their own cohort. Participants were informed about this random rematching process. In every round, participants had access to the main cost parameters during both the negotiation and the ordering stages. At the end of each round, participants saw the outcome of the current round as well as the results of all previous rounds, including the realized demand, units shipped and received through transshipments (where applicable), profit, order quantities, and, if applicable, agreed transfer prices. Based on their decisions and demand, in every round, subjects earned some profit measured in experimental currency units (ECU). At the end of the session, these ECU earnings were converted to U.S. dollars at a prespecified rate and paid out in cash along with a \$5 show-up fee. The average dollar earning was \$15.84 for an average 75-minute session.

4. Results for the Baseline and Decentralized Treatments

4.1. Transfer Prices and Orders

In Table 2, we provide the averages and standard errors for order quantities in the baseline and decentralized conditions. The unit of analysis is cohort. We use the Mann-Whitney U (Wilcoxon rank sum) tests to make the comparisons that we report.

Average orders for the no transshipment treatments are consistent with the pull-to-center behavior. Differences between average and optimal orders (see Table 1) are significant for all three CRs ($p < 0.05$). The lower-than-optimal orders in the *Medium* CR treatment are consistent with risk aversion.

Average transfer prices in the decentralized treatments are also summarized in Table 2. The coordinating transfer prices are 29, 20, and 11 for *Low*, *Medium*, and *High* CR treatments, respectively. This is, theory predicts, a negative relationship between transfer prices and CRs. However, we find that the observed transfer prices do not differ across CRs ($p >$ 0.1 for all comparisons among treatments). In other words, when negotiating, subjects behave as if they were completely ignoring the effect of the CR on their overall expected profits. In fact, the average prices are slightly above 20 ($p < 0.05$ for all comparisons). Subjects do not quite split transshipment revenue equally (which would imply the transfer price of $r/2$), but instead the retailer with excess inventory earns about 60% of the transshipment revenue. Online Appendix 2 provides some descriptive results about the negotiation process.

For decentralized treatments, a fair comparison for the ordering decisions is between average orders and

Table 2. Average Order Quantities and Optimal Order Quantities Conditional on Observed Transfer Prices

		No transshipment	Decentralized		
		Average orders	Average transfer prices	Average orders	Best reply
Critical ratio	0.25	78.640**	25.988	79.219††	66.577
		(4.433)	(2.261)	(2.871)	(2.206)
	0.50	95.257**	24.530	95.794††	103.820
		(1.533)	(0.958)	(0.748)	(0.796)
	0.75	104.159**	24.940	116.112††	142.643
		(2.607)	(0.480)	(1.977)	(0.428)

Note. Standard errors are in parentheses.
H_o: $q = q^*$; $^{**}p < 0.05$ (comparing no transshipment orders to optimal);
H_o: $q = q(t_{ij})$; $^{††}p < 0.05$ (comparing transshipment orders to best reply).

best reply orders (optimal orders conditional on the observed transfer price). Average orders are significantly higher than the best reply orders in the *Low* CR treatment (which is not consistent with risk aversion but is consistent with pull-to-center). Average orders are significantly lower than best reply orders in the *Medium* and *High* CR treatments (consistent with pull-to-center for *High* CR). Results are also similar if we make the comparison between the average orders and q_i^* (see Table 1). In addition, average deviations from the theoretical benchmarks are smaller in magnitude in the decentralized treatments than in the Baseline treatments for high and low CR ($p < 0.05$ for both CR comparisons). Deviations from the theoretical benchmarks are not significant for a CR of 0.50 ($p = 0.248$).

Finally, average retailer profits are higher in the decentralized treatments than in corresponding decentralized treatments without transshipments ($p < 0.05$ for all CR conditions), as we show in Table 3. So, transshipments do improve efficiency, but subjects fail to coordinate the channel. Our analysis of centralized treatments that we present in Section 5 incorporate behavioral findings to design centralized pricing policies that would improve coordination.

4.2. Relationship Between Transfer Prices and Order Quantities

Figure 5 shows the average transfer prices and the corresponding average order quantities for each treatment, grouped by CR condition. Each dot represents a cohort average. The figure shows that average transfer prices are in general between \$20/unit and \$30/unit regardless the CR condition. We see that transfer price and order quantity combinations are below the best reply line for *Medium* and *High* CR treatments and above the best reply line for the *Low* CR condition. Deviations (vertical distance to the line) are larger as we move away from the CR of 0.5.

We next test for the relationship between transfer prices and order quantities more formally. We test whether higher transfer prices result in higher order quantities. We estimate a panel data regression model with order quantity as dependent variable, and transfer price as independent variable. We control for the round number and use individual random effects and cohort fixed effects in our estimations. Table 4 shows that higher transfer prices result in higher average orders. This result is consistent with the theoretical predictions presented in Equation (5), which shows a positive relationship between transfer price and order quantity.

Table 3. Average Retailer Profits

		No transshipment	Decentralized
Critical ratio	0.25	325.698	516.917**
		(15.965)	(19.576)
	0.50	1,151.964	1,400.583**
		(22.118)	(15.742)
	0.75	2,156.466	2,504.594**
		(49.398)	(26.053)

Notes. Standard errors are in parentheses.
H_o: Average profit is higher for treatments with transshipment than the corresponding no transshipment treatment. $^{**}p < 0.05$.

5. Centralized Transfer Prices

How should we set transfer prices centrally to improve supply chain coordination? To answer this question, we conducted several centralized treatments using a range of CRs. We follow a two-step process. First, we use a standard analytical model that provides equilibrium transfer prices that should result in coordinating order quantities. Second, with the aim of improving channel coordination by incorporating behavioral findings, we use behavioral models from the literature that describe newsvendor behavior, and estimate their parameters using several large datasets of newsvendor decisions without transshipment from previously published and working papers. We then extend these behavioral models to include transshipment, and derive predictions for transfer prices that could maximize channel profit. This analysis guides our design for testing system performance in the centralized setting.

5.1. Behavioral Models

Results from Section 4 show evidence to the pull-to-center behavior, which is a robust regularity that has been reported in the literature. There are several behavioral models that explain the pull-to-center effect that have been proposed in the literature (for example, Becker-Peth et al. 2013, Ho et al. 2010, Long and Nasiry 2015, Ren and Croson (2013), Ockenfels and Selten (2014)). The Long and Nasiry (2015) model is mathematically equivalent to Ho et al. (2010) because according to both models, behavioral order is linear in the CR (the interpretation of the behavioral parameters are different, however). The Ren and Croson (2013) model is based on overprecision, so we cannot use it with the datasets that did not include over-precision measurements. Ockenfels and Selten (2014) model is based on impulse balance equilibrium and has no behavioral parameters that can be estimated. Therefore, the two models that we can use are the Becker-Peth et al. (2013) and the Ho et al. (2010).

The model by Ho et al. (2010) assumes that people have disutility from the mismatch between supply and demand, so that $\delta_o \geq 0$ is the psychological per-unit cost of overordering when the order exceeds the realized demand, and $\delta_u \geq 0$ is the psychological per-unit cost of underordering when the realized demand exceeds the order. Therefore, the expected utility of a

Figure 5. Observed Relationship Between Average Orders and Transfer Prices, (a) 0.25 (*Low*), (b) 0.5 (*Medium*), and (c) 0.75 (*High*) CR Conditions

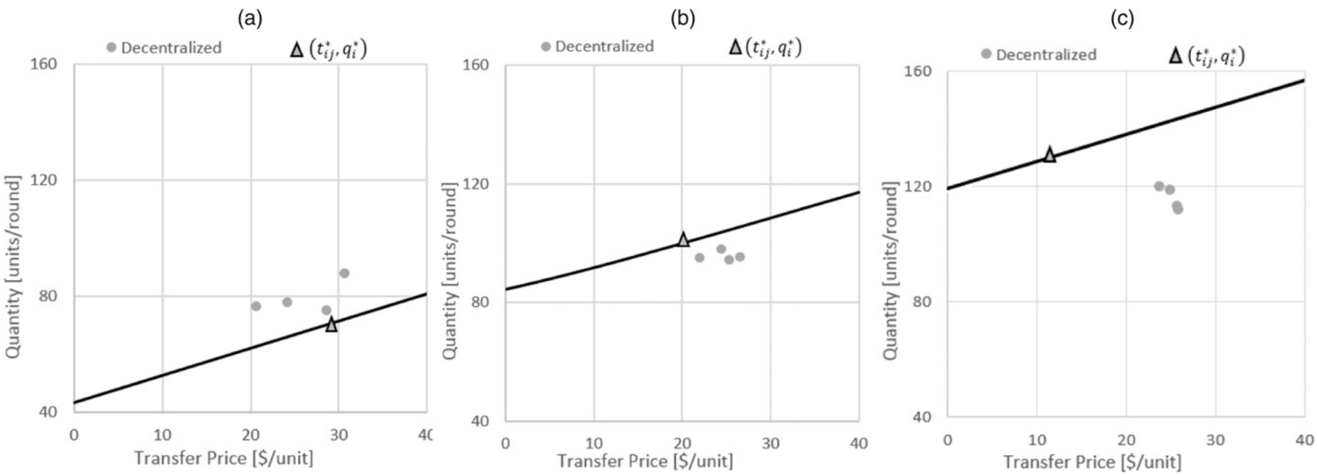

retailer from ordering q units when facing a customer demand with cumulative distribution $D(\cdot)$ and probability density function $d(x)$ is

$$u(q) = \int_0^q [rx - \delta_o(q - x)]\, d(x)dx + \int_q^\infty [rq - \delta_u(qx - q)]d(x)dx - cq. \quad (6)$$

The order quantity that maximizes the above expected utility function turns out to be given by

$$q = D^{-1}(\mathrm{CR}) = D^{-1}\left(\frac{r - c + \delta_u}{r + \delta_u + \delta_o}\right). \quad (7)$$

The Becker-Peth et al. (2013) model assumes people dislike each unit of excess inventory at λc and place orders that are a weighted average of the order implied by this loss-averse utility function and mean demand μ. The utility-maximizing order quantity in this setting is

$$q = (1 - \omega)D^{-1}\left(\frac{r - c}{r - c + \lambda c}\right) + \omega\mu, \quad (8)$$

where $\omega \in [0,1]$ represents the level of anchoring to the mean demand. We estimate behavioral parameters out of sample, using several large data sets of newsvendor decisions without transshipment from previously published papers (Bolton and Katok 2008, Ockenfels and Selten 2014) as well as working papers (Katok et al. 2019 and our previously reported baseline treatments). To make all the data comparable for the estimations, we normalize all these data sets so that their parameters and decisions have the same magnitude: Because all data sets use demand that is uniform from $\underline{d}$ to $\bar{d}$, we normalize the order quantities to be between 0 and 100 ($q = 100 \times \hat{q} - \underline{d}/\bar{d} - \underline{d}$), where $\hat{q}$ is the order quantity from the data set. We also set selling prices to 100 and normalize the unit costs so as to preserve the original CR ($c = 100 \times \hat{c}/r$), where $\hat{c}$ and r are the production cost and the retail price from the data set. We have a total of 124,878 observations from four data sets that we summarize in Table 5.

Table 6 provides the estimates of the parameters of the behavioral models. Results show that there is evidence of the anchoring and loss aversion behavior. Also, the psychological per-unit cost of overordering and underordering are close to 50% of the selling price, with a higher and significant ($p < 0.01$) value for the per-unit cost of overordering.

Figure 6 shows an overview of how the two behavioral models we estimate fit the data on aggregate,

Table 4. Panel Data Estimations for Relationship Between Transfer Prices and Order Quantities

Independent variables	Estimation
Constant	73.292***
	(6.218)
Transfer price	0.914***
	(0.084)
Low CR	−10.802
	(8.171)
High CR	27.318***
	(8.146)
Round number	−0.113*
	(0.062)
Wald χ^2	209.35***

Note. Standard errors are in parentheses.
*$p < 0.1$; ***$p < 0.01$.

Table 5. Summary of Data Sets Used for Structural Estimation of Behavioral Models

	Katok et al. (2019)	Bolton and Katok (2008)	This paper	Ockenfels and Selten (2014)
Sample size (N)	620	38	113	340
Retail price (r)	100	12	40	100
Minimum/maximum demand ($\underline{d}/\bar{d}$)	0/100	0/100 50/150	0/200	0/300
Critical ratio	Average order quantities (q) and standard errors in parentheses			
0.01	28.23 (1.81)			
0.10	34.04 (1.60)			18.37 (1.49)
0.20				29.45 (1.59)
0.25		37.81 (2.65)	39.24 (1.37)	
0.30	46.59 (1.35)			33.16 (1.07)
0.40				36.93 (1.61)
0.50	50.69 (1.24)		47.48 (0.86)	50.56 (1.82)
0.60				56.34 (2.61)
0.70	55.26 (1.10)			60.26 (2.17)
0.75		60.76 (2.48)	51.57 (1.99)	
0.80				64.44 (2.51)
0.90	63.33 (1.51)			72.66 (2.35)
0.99	71.24 (1.97)			

for a range of CRs. Boxplots represent the experimental data. The black continuous line represents the CR solution. The dotted-black and dashed-gray lines represent the predictions of the models of Becker-Peth et al. (2013) and Ho et al. (2010), respectively. An important observation is that the two behavioral models deliver very similar predictions. Even though the Becker-Peth et al. (2013) model is nonlinear in CR, the nonlinearity is slight. Therefore, we expect the coordinating transfer prices based on these two models to be qualitatively similar, if not identical.

Table 6. Structural Estimations of Behavioral Models

Parameter	Becker-Peth et al. (2013)	Ho et al. (2010)
ω	0.475*** (0.002)	
λ	1.215*** (0.007)	
δ_u		41.323*** (0.370)
δ_o		49.189*** (0.396)
σ	25.831 (0.102)	25.853 (0.102)
-LL	556,019	555,511
AIC	1,111,385	1,111,042

Note. Standard errors are in parentheses.
***$p < 0.01$.

5.2. Behavioral and Analytical Predictions

We build on the results of the out-of-sample behavioral estimations to provide predictions about subjects' behavior in a system with transshipments. Notice that the cumulative distribution of demand for a newsvendor problem with transshipment (F_i) is a function of the cost and revenue, behavioral parameters, and the probability functions α_i, β_i, γ_i. The mean-demand anchoring component for the Becker-Peth et al. (2013) model is summarized in Equation (11).

$$\hat{q}_i = (1-\omega)\, D_i^{-1}(r_i, c_i, t_{ij}, t_{ji}) + \omega\, \mu \tag{11}$$

To derivate the loss aversion component of a behavioral model, we compute the expected utility for a retailer as a function of (i) the underage cost $(r_i - c_i + \delta_u)$ in the Ho et al. (2010) model or $(r_i - c_i)$ in the Becker-Peth et al. (2013) model and (ii) the overage cost given the behavioral parameter measuring

Figure 6. An Overview of Order Quantities from the Data Sets, with Behavioral and Analytical Predictions

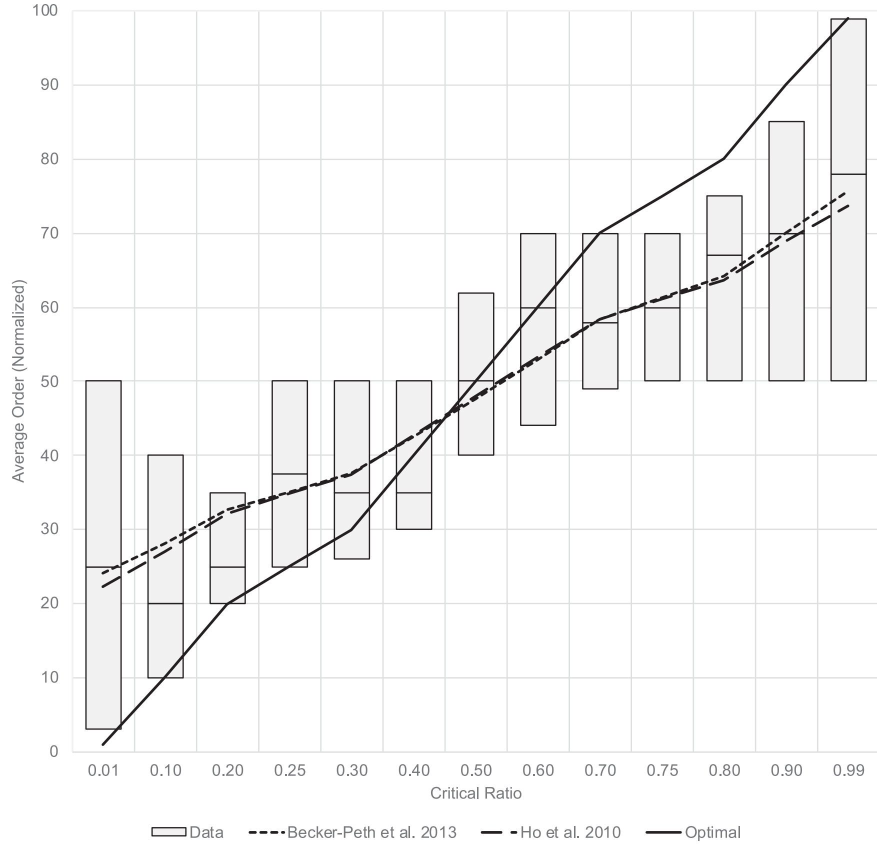

disutility from excess inventory, which is δ_o in the Ho et al. (2010) model and $(\lambda\, c_i)$ in the Becker-Peth, et al. (2013) model. Equation (12) shows the loss-aversion component of the Becker-Peth, et al. (2013) behavioral model in the setting with transshipment.

$$\alpha_i(\hat{q}_i) - \beta_i(\hat{q}_i, \hat{q}_j)\left(\frac{t_{ij}}{r_i - c_i + \lambda\, c_i}\right) + \gamma_i(\hat{q}_i, \hat{q}_j)\left(\frac{r_i - t_{ji}}{r_i - c_i + \lambda\, c_i}\right) = \left(\frac{r_i - c_i}{r_i - c_i + \lambda\, c_i}\right) \tag{12}$$

The analogous expression for the Ho et al. (2010) model is

$$\alpha_i(\hat{q}_i) - \beta_i(\hat{q}_i, \hat{q}_j)\left(\frac{t_{ij} + \delta_o}{r_i + \delta_u + \delta_o}\right) + \gamma_i(\hat{q}_i, \hat{q}_j)\left(\frac{r_i - t_{ji} + \delta_u}{r_i + \delta_u + \delta_o}\right) = \left(\frac{r_i - c_i + \delta_u}{r_i + \delta_u + \delta_o}\right). \tag{13}$$

Replacing δ_u with $(1-b)r$ and $\delta_u + \delta_o$ with $(\lambda - 1)b^2 r + r$ gives us the Long and Nasiry (2015) model.

Using the behavioral model, we can solve the system to determine the transfer price that should lead subjects to place coordinating order quantities. Figure 7 shows coordinating transfer prices for the analytical model and for the two behavioral models. The analytical model shows a negative relationship between CR and the coordinating transfer price (in other words, a high transfer price is required in order to coordinate a system with a low CR, and vice versa). But both behavioral models indicate the opposite relationship between the transfer price needed to coordinate the system and the CR—a low transfer price is needed to coordinate a system with a low CR, and a high transfer price is needed to coordinate the system with a high CR. This is intuitive because both models imply pull-to-center, so a high transfer price is required to compensate for this when the CR is high, whereas a low transfer price is required to compensate for this when the CR is low. In fact, there is a range of high CRs for which behavioral models indicate that the coordinating transfer price should be above the retail price. For example, for CR = 0.75, the model suggests a transfer price of 52 is needed to induce retailers to order 130 units, whereas the retail price is only 40. Of course, any transfer price above r is not reasonable because it would result in losses for the retailer buying the units. Therefore, the strongest incentive that a centralized planner can provide to the retailers to increase their orders is to set $t^* = r$, which is

Figure 7. Behavioral and Analytical Predictions to Elicit Optimal Orders

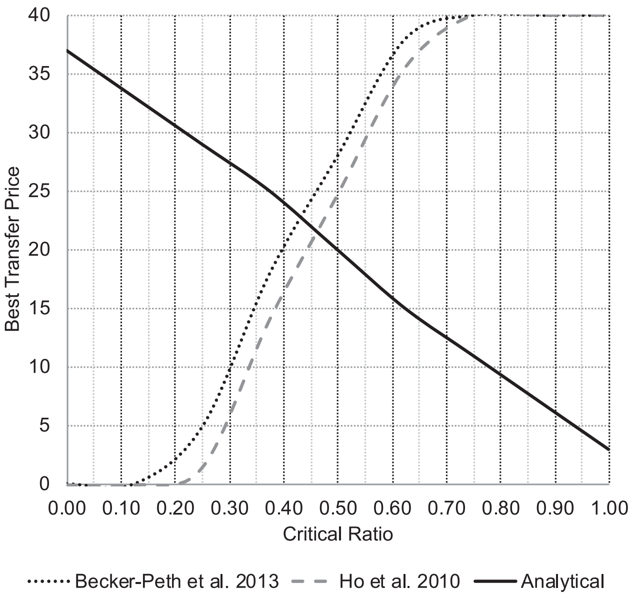

what we do in our experiments: for CRs for which $t^* > r$, we use $t^* = r$.

In summary, the key difference in the guidance for setting centralized transfer prices between the standard model and the behavioral models is that, although standard model's advice is to set high transfer prices for low CRs and low transfer prices for high CRs, behavioral models provide the opposite guidance. We designed our experiments to determine which policy results in better performance.

5.3. Experimental Design

In this section, we explore subjects' decisions under a centralized setting in which participants place their orders after the transfer price has been set centrally. As in our previous treatments, we use a unit revenue r_i of \$40 and a uniform demand distribution $U[0,200]$ for the final customer demand (d_i). To capture how participants respond to a range of transfer prices for different CRs, our experimental design includes a number of different combinations of unit costs and transfer prices. Our design builds on both the behavioral and analytical predictions shown in Figure 7. Table 7 summarizes our centralized treatments. In the first two columns of Table 7, we show CRs we used in our experiment and corresponding coordinating order quantities. As we can see from Figure 7, the optimal behavioral transfer price for CRs below approximately 0.25 is 0 and for CRs above approximately 0.75 it is 40, so behavioral regularities we observe for those two CRs we expect to continue to hold for more extreme CRs. Thus, we restrict our design to CRs that are between 0.25 and 0.75. Each populated cell in the last five columns of Table 7 corresponds to a treatment we conducted. The top number in each cell corresponds to the predicted behavioral order quantity (q_B) based on the Ho et al. 2010 model (the Becker-Peth et al. 2013 model predictions are similar), and the numbers in square brackets correspond to the orders predicted by the standard theory (q_T).

We selected combinations of CR and the transfer price in a way that would test behavioral predictions most efficiently. For the CR of 0.75, we examine behavior with the transfer price of 11 (the optimal transfer price based on the standard model) and the transfer price of 40 (the optimal transfer price based on the behavioral standard models). Treatments for the CR of 0.75 with transfer prices below 11 are unnecessary because the order quantity with $t_{ij} = 11$ that the behavioral models predicts is 112, which is already substantially below the coordinating order of 130. For the CR of 0.25, we examine transfer prices of 0 (the optimal transfer price based on the Ho et al. 2010 behavioral model), 7 (the optimal transfer price based on the Becker-Peth et al. 2013 behavioral model), 30 (the optimal transfer price based on the standard model), and 40 (as a robustness check). We omitted a

Table 7. Experimental Design and Optimal Orders for Centralized Treatments

Critical ratio (CR)	Coordinating order (q^*)	Transfer price (t_{ij}): 0	7	11	30	40
0.25	70	69 [59]	70 [62]		80 [76]	85 [82]
0.38	86					96 [100]
0.63	114					115 [135]
0.75	130			112 [130]		126 [155]

treatment with the CR of 0.25 and $t_{ij} = 11$ because 11 is very close to 7 and we would not expect to detect a difference in behavior between those two treatments. Finally, we also consider two less extreme CRs, 0.38 and 0.63, with the transfer price of 40, to check the robustness of the policy of setting the transfer price to the retail price r, because this policy seems both promising for a wide range of CRs and simple to implement. We did not conduct treatments with CRs of 0.63 with transfer prices below 40 because we found (details in Section 5.4) that the average order with the transfer price of 40 was already significantly below coordinating (108 versus 115), indicating that lower transfer prices would only further decrease order quantities. We also elected not to conduct a treatment with CR of 0.38 and transfer prices below 40 because the average order of 92 with the transfer price of 40 was quite close to the coordinating order of 86. A transfer price between 15 and 20 would likely have resulted in orders closer to coordinating, but we would have likely not been able to detect a statistically significant difference.

In the five right-most columns of Table 7, cells enclosed by a border correspond to a separate set of participants, and shaded cells within the same rectangle are treatments that we conducted within subjects (so we conducted the CR = 0.25 condition within subjects for $t = 0$, 7, and 30, whereas the $t = 40$ condition within subjects for CR = 0.25, 0.38, and 0.63; we varied the order in those sessions, to check for order effects, and we did not find any). The two unshaded cells (CR = 0.75 transfer prices of 11 and 40) are treatments that we conducted between subjects. Sessions of those two between-subject treatments lasted for 30 rounds, whereas the within-subject sessions lasted for 45 rounds.

All sessions included four cohorts of 8 participants, so the total of $32 \times 4 = 128$ participants were included in the centralized treatments. Within each cohort, participants were randomly rematched each round. Sessions for the within-subjects treatment included three blocks of 15 rounds during which parameters remained constant. The order of those blocks was randomized for each participant (to control for order effects). The rest of the protocol remained identical to the one we used in the decentralized treatments.

5.4. Results

Table 8 provides the average orders for the centralized treatments, along with comparison with the coordinating order quantity (q^*), the order quantity predicted by the standard theory (q_T), and the order quantity predicted by the behavioral model (q_B). Generally, we cannot reject the null hypothesis that the observed data are not different from behavioral prediction. The two exceptions, marked in bold in Table 8, are the treatments with the transfer price of 40 and the CRs of 0.25 ($q > q_B$) and 0.63 ($q < q_B$). Considering that the behavioral model predictions were obtained out-of-sample and using a different task, we can conclude that behavioral models predict orders quite well.

In contrast to the good predictive performance of the behavioral models, the standard theory only predicts well for the low CR of 0.25. For the other CRs, including even a slightly higher CR of 0.38, average orders are significantly lower than predicted ($q < q_T$).

We now consider the practical question of how the transfer price should be set centrally. We can use the data from the CR = 0.25 and CR = 0.75 conditions to answer this question. We see from Table 8 that order quantities are quite insensitive to the transfer price when the CR is 0.25—for this low CR, to reject the null hypothesis that the average order is different from 70, the transfer price has to go as high as 40. Low transfer prices do cause average orders to decrease relative to the no transshipment condition ($p = 0.03$ when $t = 0$ and $p = 0.07$ when $t = 7$) and the decentralized condition ($p = 0.008$ for $t = 0$ and $p = 0.03$ when $t = 7$). But in the $t = 30$ treatment, average

Table 8. Average Orders for Centralized Treatments

		Transfer price (t)				
Critical ratio (CR)	q^*	0	7	11	30	40
0.25	70	64.385	66.270		76.347	**90.913**$^{***/\dagger\dagger\dagger}$
		(2.582)	(3.442)		(3.242)	(0.711)
0.38	86					92.640$^{**/\dagger\dagger}$
						(1.937)
0.63	114					**108.415**$^{*/\dagger\dagger\dagger}$
						(1.994)
0.75	130			109.399$^{**/\dagger\dagger}$		126.150$^{**/\dagger\dagger\dagger}$
				(4.770)		(1.199)

H_o: $q = q^*$ (coordinating); $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$; H_o: $q = q_T$ (standard equilibrium); $^{\dagger}p < 0.1$; $^{\dagger\dagger}p < 0.05$; $^{\dagger\dagger\dagger}p < 0.01$;
H_o: $q = q_B$ (behavioral);
bold $p < 0.05$.

Table 9. Average Profits

			Transfer price (t)				
Critical ratio (CR)	No transshipment (π_N)	Decentralized transfer price (π_D)	0	7	11	30	40
0.25	325.70	516.92	426.27	426.13		445.06$^{**/\dagger\dagger}$	315.79††
	(15.97)	(19.58)	(41.02)	(53.58)		(12.78)	(60.12)
0.75	2,156.47	2,504.59			2,459.25**		2,550.07**
	(49.40)	(26.05)			(31.72)		(13.17)

H_o: $\pi = \pi_N$; *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$;
H_o: $\pi = \pi_D$; $^{\dagger}p < 0.1$; $^{\dagger\dagger}p < 0.05$; $^{\dagger\dagger\dagger}p < 0.01$.

orders are not different from either the centralized order of 70 ($p = 0.14$) or the no transshipment order of 78.64 ($p = 0.69$) or the order with decentralized prices of 79.22 ($p = 0.53$). Therefore, we can conclude that when CR is low, to mitigate the behavioral propensity to order too much because of the pull-to-center effect, transfer prices should be set rather low—at the wholesale price or even slightly below the wholesale price.

But orders are more sensitive to transfer prices for the CR of 0.75. For the transfer price of 40, which is the optimal transfer price based on the behavioral prediction, even though average orders are slightly below the optimal level of 130 ($p = 0.048$), they are significantly above the average orders of 104.16 that we observed without transshipment ($p = 0.0003$) or the average orders of 116.11 that we observed with decentralized transfer prices ($p = 0.005$). For the transfer price of 11, which is the optimal transfer price based on the standard theory, performance is much worse. The average orders when using the optimal transfer price are not only below the optimal level of 130 ($p = 0.02$) but are also neither different from the average orders without transshipment ($p = 0.37$) nor the decentralized orders ($p = 0.24$). Therefore, we can conclude that to effectively improve performance in a setting with a high CR, it is important to set the transfer price high—at or near the retail price—to counteract the behavioral propensity of not ordering enough.

In order to complete the picture, we also present results describing the effect of transfer pricing policy on profitability. Table 9 presents the average profits our participants earned in the centralized treatments, along with comparison with the profits with no transshipments (π_N), and profits with decentralized transshipments (π_D).

We already observed that merely allowing transshipments increases profitability, and we confirm this conclusion in the settings with centralized transfer prices. But, with our data, it was not possible to confirm that the use of a centralized transfer price can improve profitability relative to a decentralized setting. When the CR is low, even the best centralized transfer price results in average profits that are slightly lower than profits under the decentralized setting ($p = 0.02$). When the CR is high, the best pricing policy, which is to set the transfer price at the retail price level, delivers slightly higher (although not statistically significantly so with $p = 0.17$) average profitability than the decentralized setting. Failing to detect significant differences in profitability is most likely because of the specifics of our experimental parameters, including specific demand realizations, because the expected profit functions are quite flat. Also, we can see from standard errors in Table 9, that there is quite a bit of heterogeneity in the behavior of our subjects in some of the treatments.

6. Conclusion

Transshipment provides a way for retailers to share inventory. Taking advantage of the demand pooling effect of transshipment is an effective way to decrease the mismatch between supply and demand, potentially increasing service levels and reducing excess inventory. We use a set of controlled behavioral laboratory experiments to compare systems with and without transshipment, under a broad set of scenarios. In our experiments, we vary the critical ratio (which represents product profitability) and the process of setting the transfer price (centralized or through negotiation among the retailers). We find that transshipment usually improves profitability (one exception is a treatment with low critical ratio and high transfer price, which causes retailers to overorder so much that transshipment becomes detrimental to their profits); but when retailers negotiate transfer prices, they behave as if they do not pay attention to the critical ratio, focusing instead on the profit allocation from transshipped units, and average orders are usually different from coordinating orders. Interaction between transfer price negotiation in decentralized settings and the critical ratio may be an interesting direction for future research.

To find out whether a system in which transfer prices are set centrally can outperform the decentralized system, we conduct additional experiments that compare the performance of retail channels with standard equilibrium transfer prices to the performance with optimal behaviorally informed transfer prices. To estimate optimal behaviorally informed transfer prices, we extended two of the established behavioral models that explain newsvendor ordering behavior, to a setting with transshipments, and estimated their behavioral parameters out-of-sample using a data set that included newsvendor decision data from four separate studies. Under the standard equilibrium model, coordinating transfer prices decrease with the critical ratio; but we found that behaviorally informed transfer prices increase with the critical ratio. Our experiments confirmed that average orders are closer to coordinating when we use behaviorally informed transfer prices. We conclude that when designing transshipment arrangements, it is important to account for behavioral regularities, because failing to do so degrades the performance of a centralized channel below that of a decentralized channel.

Our work has several limitations. One is that directly comparing orders under the centralized and the decentralized setting is not entirely fair, because in the centralized setting transfer prices are not only different but are also exogenous. Only about 3% of decentralized decisions in the high CR condition involve transfer prices around 11 or above 35. This amount of data is not enough for a formal comparison; but average decentralized orders with the transfer price of about 11 are 105.15, which is quite close to an average centralized order of 109.4 with the transfer price of 11. And average orders are much higher with high transfer prices (123.53 with decentralized versus 126.15 centralized). But a formal analysis would require a different set of experiments, and we leave it for future research. Another limitation is that we consider a case with symmetric retailers and complete pooling of inventory. The analysis of the case with asymmetric retailers is more complicated, because equilibrium transfer prices may not be unique (Hu et al. 2007) and there may well be additional behavioral implications because of the lack of symmetry.

We draw three conclusions from our research. First, transshipment is unambiguously beneficial, resulting in higher profitability; therefore, when transshipment is feasible, it should be encouraged. Second, having retailers negotiate transfer prices performs quite well, so if a manufacturer wishes to influence transfer prices (as in the Shao et al. 2011 model) or if prices are set centrally as in a "chain store model," transfer prices should increase, rather than decrease, with the critical ratio. Our third conclusion is that a transfer price that is close to the retail price works well for products with a high critical ratio, and a transfer price that is below cost works well for products with a low critical ratio, so behaviorally informed transfer prices may be easy to implement in practice.

Acknowledgments
Katok gratefully acknowledges support from the Jindal School of Management at the University of Texas at Dallas Mago Professorship. The authors thank seminar participants at the Indiana University for their helpful comments.

References
Arsu S (2009) Turkey and Russia conclude energy deals. *New York Times* (August 6), http://www.nytimes.com/2009/08/07/world/europe/07turkey.html.

Axsäter S (2003) A new decision rule for lateral transshipments in inventory systems. *Management Sci.* 49(9):1168–1179.

Becker-Peth M, Katok E, Thonemann UW (2013) Designing buyback contracts for irrational but predictable newsvendors. *Management Sci.* 59(8):1800–1816.

Benzion U, Cohen Y, Peled R, Shavit T (2008) Decision-making and the newsvendor problem: An experimental study. *J. Oper. Res. Soc.* 59(9):1281–1287.

Bolton GE, Katok E (2008) Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing Service Oper. Management* 10(3):519–538.

Bolton GE, Stangl T, Thonemann UW (2016) Decision making under service level contracts – An experimental analysis. Preprint, submitted September 14, https://ssrn.com/abstract=2838645.

Bostian AA, Holt CA, Smith AM (2008) Newsvendor "pull-to-center" effect: Adaptive learning in a laboratory experiment. *Manufacturing Service Oper. Management* 10(4):590–608.

Castañeda JA, Villa S, Urrea G (2019) *Leveraging Service Level and Understock Costs to Improve Ordering Decisions* (University of Colorado Boulder, Boulder).

Chen K-Y, Li S (2020) The commitment conundrum of inventory sharing. *Production Oper. Management* 29(2):353–370.

CREG (2009) Colombia aumenta exportaciones de energía a Ecuador .com. *Regul. Energía y Gas*:1. Accessed May 23, 2020, https://www.creg.gov.co/node/8389.

Dong L, Rudi N (2004) Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Sci.* 50(5):645–657.

Hendriks A (2012) *SoPHIE-Software Platform for Human Interaction Experiments* (Osnabrueck, Germany).

Herer YT, Tzur M, Yücesan E (2006) The multi-location transshipment problem. *IIE Trans.* 38(3):1–33.

Ho TH, Lim N, Cui TH (2010) Reference dependence in multilocation newsvendor models: A structural analysis. *Management Sci.* 56(11):1891–1910.

Hu X, Duenyas I, Kapuscinski R (2007) Existence of coordinating transshipment prices in a two-location inventory model. *Management Sci.* 53(8):1289–1302.

Katok E, Stangl T, Thonemann U (2019) Drivers of the pull-to-center effect. Working paper, University of Texas, Dallas.

Kranenburg AA, van Houtum GJ (2009) A new partial pooling structure for spare parts networks. *Eur. J. Oper. Res.* 199(3):908–921.

Krishnan KS, Rao VRK (1965) Inventory control in N warehouses. *J. Indust. Engrg.* 16:212–215.

Li R, Ryan JK (2017) Coordination in a single-supplier. *Multi-Retailer Distribution System: Supplier-Facilitated* 26(5):784–801.

Long X, Nasiry J (2015) Prospect theory explains newsvendor behavior: The role of reference points. *Management Sci.* 61(12):3009–3012.

Narus JA, Anderson JC (1996) Rethinking distribution: Adaptive channels. *Harvard Bus. Rev.* 74(4):112–120.

Ockenfels A, Selten R (2014) Impulse balance in the newsvendor game. *Games Econom. Behav.* 86:237–247.

Ren Y, Croson R (2013) Overconfidence in newsvendor orders: An experimental study. *Management Sci.* 59(11):2502–2517.

Robinson L (1990) Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Oper. Res.* 38(2):278–295.

Rudi N, Kapur S, Pyke DF (2001) A two-location inventory with transshipment and a 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):361–372.

Shao J, Krishnan H, McCormick ST (2011) Incentives for transshipment in a supply chain with decentralized retailers. *Manufacturing Serv. Oper. Management* 13(3):590–608.

Sošić G (2006) Transshipment of inventories among retailers: Myopic vs. farsighted stability. *Management Sci.* 52(10):1493–1508.

Tagaras G (1989) Effects of pooling on the optimization and service levels of two-location inventory systems. *IIE Trans.* 21(3):250–257.

Villa S, Castañeda JA (2018) Transshipments in supply chains: A behavioral investigation. *Eur. J. Oper. Res.* 269(2):715–729.

Werdigier J, Dougherty C (2007) ECB acts to dispel clouds over European credit markets. *New York Times* (August 22), https://www.nytimes.com/2007/08/22/business/worldbusiness/22iht-banks.4.7214405.html.

Zhang J (2005) Transshipment and its impact on supply chain members' performance. *Management Sci.* 51(10):1534–1539.

Zhao H, Deshpande V, Ryan JK (2005) Inventory sharing and rationing in decentralized dealer networks. *Management Sci.* 51(4):531–547.

Zhao H, Xu L, Siemsen E (2016) A behavioral perspective on inventory sharing. Preprint, submitted December 9, http://dx.doi.org/