Negotiating Transfer Prices for Improving Supply Chain Transshipments

Sebastián Villa

School of Management, University of los Andes, Bogotá, Colombia, s.villab@uniandes.edu.co

Elena Katok

Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080, ekatok@utdallas.edu

Abstract

When a retail channel includes multiple retailers, transshipment can be utilized to better match supply and demand. We study retailers' behavior in a channel with one supplier and two independent retailers. The two retailers in our setting can negotiate a transfer price either before or after making their ordering decisions under demand uncertainty. After demand realizes, retailers transship units from one another at the negotiated price (when possible). In our controlled laboratory experiment, human subjects play the role of retailers. We manipulate the product profitability (the critical ratio) as well as the timing of the transfer price negotiation, to study how deviations in transshipment decisions (ordering decisions and transfer price negotiations) affect supply chain coordination. According to the theory, the optimal transfer price should depend on the critical ratio, but our results show that participants almost completely ignore the critical ratio when setting transfer prices. Moreover, we find that the sequence of decisions matters: the likelihood of closing a negotiation over the transfer price is significantly decreased when the negotiation takes place after the orders have been placed. We generalize the pull-to-center effect, documented in previous studies of newsvendor behavior, to a setting with transshipment. Finally, even though our participants deviate from optimal behavior, the possibility of transshipment does increase profit, and there is a positive relationship between transfer prices and ordering quantities. Both observations are consistent with standard analytical predictions.

1. Introduction

The newsvendor problem is one of the most studied problems in Operations Management. This problem captures a single-period single-agent problem, in which a retailer makes an inventory ordering decision before the realization of an uncertain demand. In this kind of problem, both leftovers and shortages are costly (Eeckhoudt et al., 1995; Schweitzer & Cachon, 2000) and result in the mismatch between supply and demand. In a setting with multiple retailers in the same industry, it can be profitable for the retailers to reallocate products among themselves after the stochastic demand realizes. Retailers who are short can purchase units from retailers who have excess inventory, improving the match between supply and demand. This exchange of product among retailers is called *transshipment*.

Traditional analytical models of inventory management assume that retailers make optimal decisions. In the context of transshipment, this would imply that retailers place orders and set transfer prices that maximize their joint expected profit in equilibrium. However, human behavior on inventory decisions often deviates from standard theoretical predictions (Bendoly et al., 2010; Cui et al., 2007; Loch & Wu, 2008). We investigate the behavior in the transshipment setting by conducting a behavioral study of two transshipment decisions: placing orders and setting transfer prices. We capture these decisions and their effect on performance in a setting with two retailers and one supplier. The supplier has unconstrained capacity and both retailers place orders simultaneously prior to knowing random demand realization.

Transshipments are common in industries characterized by long lead times, short selling seasons and high demand uncertainty, as a mechanism to alleviate the problem of localized demand shocks (Herer et al., 2006). Some examples of the use of transshipments include industries such as semi-conductor manufacturing (Kranenburg & van Houtum, 2009), fashion apparel (Dong & Rudi, 2004), financial and electricity markets (Werdigier & Dougherty, 2007), among others. For example, such manufacturers as General Motors and Bosch, distribute their products using different independently-owned retailers. However, if one of the retailers is out of stock, it can obtain the required units from another retailer in the same country in a relatively short period of time and for a fair transfer price, instead of waiting for a delayed supply from the central distributor (Rudi et al., 2001). Similarly, in the humanitarian sector, the International Federation of the Red Cross, distributes vehicles among different regional hubs that operate independently. When a disaster occurs, some of these vehicles might be transferred from one hub to another according to the needs (Stauffer et al., 2016). Many industries are also using e-market places, such as Inventory Locator System, in which retailers trade online with other retailers in addition to their own supplier (Zhao & Bisi, 2010)

Transshipments encourage independent agents to coordinate, share their inventories, and negotiate proper prices to achieve a better supply chain performance (Dong & Rudi, 2004; Rudi et al., 2001; Sošić, 2006). Agents require both accurate ordering decisions and aligned transfer prices, to achieve good coordination. If the ordering decisions in all locations are made centrally, then transfer prices do not matter for coordination. However, in decentralized channels, each retailer orders its own inventory, and therefore, the process of negotiating transfer prices becomes important because optimal order quantities depend on transfer prices.

Negotiations define the conditions under which the deal takes place. In these processes, agents sign bilateral contracts in which prices are negotiated freely between the two parties. For example, for the energy market in Asia and Europe, Russia and Turkey closed a deal for energy exchange between the two continents, aiming both at improving the supply conditions, and at increasing the energy supply control over the region (Arsu, 2009). Similarly, in the energy market in Latin America, non-regulated users use transshipment as a special form of "risk pooling" to satisfy several classes of demands (CREG, 2009).

Recent behavioral studies of the transshipment problem have considered the benefits of transshipments, the effects of communication, inventory policies and rules. For example, Bostian et al. (2008) studied how retailers benefit from transshipments when they make ordering decisions with automated transshipments. Villa and Castañeda (2017) study typical biases in a Newsvendor problem with automated transshipments and describe the effect of communication and best response strategies on subjects' performance. Zhao et al. (2016) examine inventory sharing effectiveness and find that retailers tend to under-stock inventory, more so when there is transshipment opportunity, and therefore allowing transshipments does not generally improve efficiency. Chen and Li (2018) consider a setting with voluntary transshipments and examine the effect of the timing of transfer price negotiations. They find that prices and order quantities are generally set below coordinating levels, and the efficiency is lowest when transfer prices are negotiated before random demand is realized, and transshipments are voluntary. Chen and Li (2018) propose a behavioral model incorporating fairness and inventory overage aversion to organize their data.

There are several ways in which our work is distinct from the above-mentioned behavioral studies of transshipment. First, ours in the only study that varies the critical ratio (CR). Second, transfer prices in our study, are negotiated through a free form bargaining process, either before or after orders are placed. Thus, transfer prices in our study are closer to the prices that emerge from real decentralized transactions (Rudi et al., 2001) than prices that are either exogenous or are set unilaterally by one of the retailers. We also directly examine the effect of transfer prices on ordering decisions, which is important because appropriate evaluation and understanding of the ordering quantities and transfer prices can improve stock polices, increase overall profits, create cost reduction and improve service levels (Herer et al., 2006).

We report on three behavioral regularities in our laboratory experiments. First, order quantities in transshipment settings are affected by the CR but not by whether prices are negotiated before or after the orders are placed. Due to a pooling effect, transshipments cause the optimal order quantity to be closer to mean demand, and because order quantities we observe are closer to the mean demand, the order quantities with transshipments are closer to optimal. Second, negotiated transfer prices are not affected by the CR. This means that participants do not learn to use transfer prices to improve channel coordination. Lastly, negotiation success (i.e. agreeing on a transfer price) is more likely when transfer prices are negotiated before orders are placed rather than after. This is in contrast to analytical models that do not predict any sequence-dependent differences in performance.

In the next section we summarize the analytical background for the newsvendor problem with transshipments, present our experimental design and protocol, and derive analytical predictions for our experimental setting. In section 3, we use analytical background we presented, along with what we know from existing behavioral literature, to develop experimental hypothesis. We report results of our experimental study in section 4, and in section 5 we summarize our findings and discuss their implications.

2. Standard Theory and Experimental Design

2.1. Standard Theory

Rudi et al. (2001) present a simple model with one supplier and two retailers. In our setting, 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 while the other is short, the number of units transshipped will be the minimum of the excess and the shortage

(Tagaras, 1989). We are interested in understanding the way in which the two independent and symmetric retailers R_i (with i = 1, 2) negotiate their transfer prices (t_{ij}, with $i, j = 1, 2$) and make ordering decisions (q_i) in a decentralized setting. Each retailer R_i buys his items from the supplier at a cost of c_i, 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 addition value to the retailers ($b_i = 0$). The transfer price t_{ij} is paid by the retailer receiving the transshipped units to the other retailer.

To provide some analytical intuition about the standard theory behind our system, we build on the typical newsvendor problem, extend it to a *centralized* system with transshipments, and finally summarize the results for a system with independent (*decentralized*) retailers.

2.1.1 Isolated Retailers: Newsvendor Problem

Initially, we consider a system in which two isolated retailers place their orders to a supplier with unlimited capacity. In this system, transshipments among retailers are not possible. Therefore, the optimal ordering quantity for each retailer R_i will be determined by the well-known critical ratio (CR).

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

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

2.1.2 Centralized System

The centralized setting 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 orders, (ii) supplier fills these orders, (iii) final customer demand takes place, (iv) demand is satisfied, and (v) potential transshipments among retailers and additional final customer demand satisfaction take place. Notice that transshipments are feasible only when one retailers faces a shortage and the other retailer has excess inventory (Krishnan & Rao, 1965; Rudi et al., 2001). The first state of the world in a Transshipment Problem 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 of is the complement to the other three states and it 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 possible states of the world (and their associated probabilities) that take place in a supply chain with transshipments for R1 and R2.

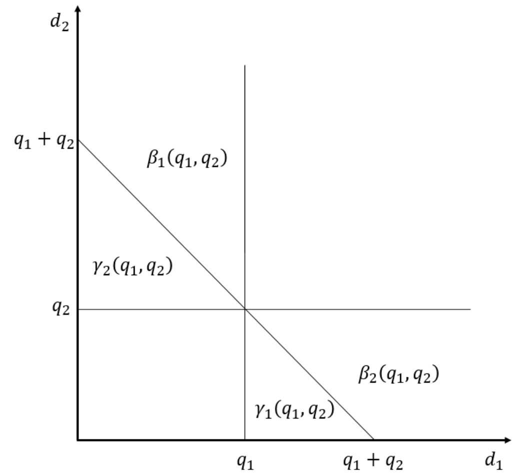

Figure 1. States of the world and probabilities for a Transshipment Problem (Rudi et al., 2001).

Building on these states of the world and probabilities, and assuming that the joint distribution over the demand is continuously differentiable, the ordering quantities (q_i^c, q_j^c) that maximize the total expected profit for a centralized system can be obtained by solving eq. (2), such that it holds for $i, j = 1, 2$.

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

Notice that eq. (2) is just an adjustment of eq. (1). The second term on the left-hand side tries to increase q_i to account for 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 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 (eq. (1)). The intuition is that with one supplier and multiple retailers who individually face random demand that is not perfectly correlated, the possibility of transshipment between the retailers causes the centralized solution to approach mean demand as the number of retailers increases due to inventory pooling.

2.1.3 Decentralized System

Decentralized settings consider transshipments at an interfirm level in which decisions are locally defined to maximize individual profits. Decentralized systems are common in practice and they involve more operational challenges to decision makers than a centralized system (Zhao et al., 2005). Retailers need to coordinate their operations while guarantee their individual profitability.

Given the nature of this decentralized system, R_i's profit depends on the transfer prices t_{ij}. Rudi et al. (2001) consider a simplified model of a single-supplier two-retailer supply chain. They show that there is a unique set of transfer prices t_{ij}^* and ordering quantities (q_1^*, q_2^*) that can coordinate the whole system to optimize the total supply chain 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 optimal inventory policy for each retailer R_i. The Nash equilibrium $\left(q_i^d, q_j^d\right)$ for a decentralized system at t_{ij} is obtained by solving the reaction function presented on equation (3), such that it holds both for $i, j = 1, 2$.

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

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

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

$$t_{ij}^* = \frac{r_j\beta_i(q_i^c, q_j^c)\beta_j(q_j^c, q_i^c) - r_i\beta_j(q_j^c, q_i^c)\gamma_i(q_i^c, q_j^c)}{\beta_i(q_i^c, q_j^c)\beta_j(q_j^c, q_i^c) - \gamma_i(q_i^c, q_j^c)\gamma_j(q_j^c, q_i^c)} \tag{4}$$

Using the unique set of transfer prices t_{ij}^* to solve eq. (3) leads to the estimation of the optimal ordering 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 (eq. (2)). That is, when there are multiple retailers that can transship, the optimal order approaches mean demand as the number of retailers increases. This centralized solution can be implemented in a decentralized setting with a particular set of transfer prices.

We provide an analytical interpretation of the relationship between the optimal transfer prices and the critical ratio by performing an additional analysis of the reaction function for the decentralized system (eq.(3)) evaluated at the optimal ordering 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 $\beta_i(q_i^*, q_j^*), \gamma_i(q_i^*, q_j^*)$ and r_i are expected to be always positive for a system with transshipments, Eq. (5) shows that there is a negative relationship between the coordinating transfer prices (t_{ij}^*) and the critical ratio $(\frac{r_i - c_i}{r_i})$: the larger the *CR*, the smaller the coordinating transfer price.

2.2. Experimental Design

Our laboratory experiment includes a 3 × 3 full factorial design: we manipulate three factors at three levels. First, we consider product profitability by varying the critical ratio (CR) at three levels (Low CR = 0.25, Medium CR = 0.5, and High CR = 0.75). Second, we consider three decision sequences that we display in Figure 2: (1) No transshipment (the baseline), (2) transfer price negotiation followed by the ordering decision (Price First), and (3) ordering decision followed by transfer price negotiation (Quantity First).

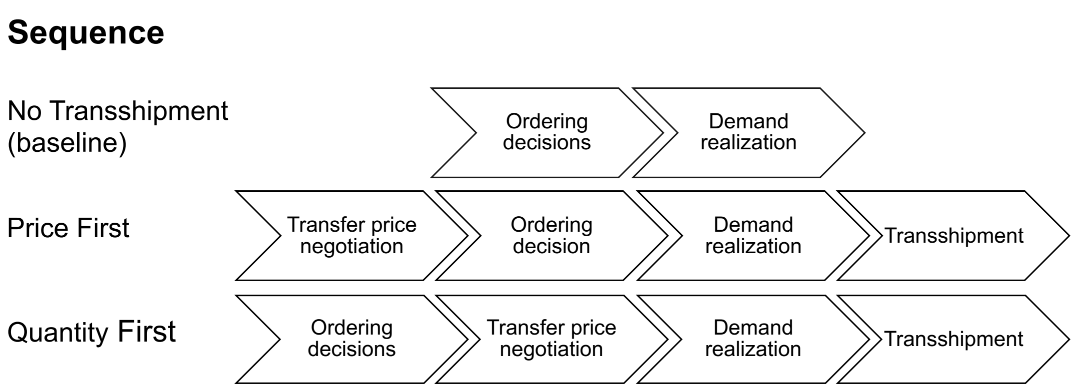

Figure 2. Decision sequence for each experimental treatment

Table 1 lists the nine treatments with their coordinating transfer prices (where applicable) and order quantities. During the ordering decision stage of the experiment, each participant places an order for the number of units they wish to purchase. The Price First and Quantity First conditions also include 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 ordering decision, and where applicable transfer price negotiations have been completed, customer demands are randomly and independently drawn. Following demand realization, the transfer units are automatically shipped and the profits are calculated.

Table 1. Experimental Design.

		Sequence		
		No Transshipment	Price First	Quantity First
Critical Ratio	0.25	$q_i^* = 50$	$t_{ij}^* = 29$ $q_i^* = 70$	$t_{ij}^* = 29$ $q_i^* = 70$
	0.50	$q_i^* = 100$	$t_{ij}^* = 20$ $q_i^* = 100$	$t_{ij}^* = 20$ $q_i^* = 100$
	0.75	$q_i^* = 150$	$t_{ij}^* = 11$ $q_i^* = 130$	$t_{ij}^* = 11$ $q_i^* = 130$

Note: Each treatment included 4 cohorts of 8 subjects (32 participants per treatment).

We implemented the experiment using the Software Platform for Human Interaction Experiments (SoPHIE) (Hendricks 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 288 participants in total. Our participants were master students in Management from a public U.S. University. Once participants entered the lab, they were directed to a computer terminal and had 10 minutes to read the instructions. After ten 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 (r) and the unit purchase cost (c). In addition, participants were given a table with information about the relationship between transfer prices and order amounts that lead to the highest expected profit (see Appendix for experimental instructions and the information table). All sessions lasted 30 rounds and, in each round, participants were randomly re-matched with another person within their own cohort. Participants were informed about this random re-matching process. 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, ordering quantities and, if applicable, agreed transfer prices. Based on their decisions and demand, 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 pre-specified rate and paid out in cash along with a $5 show-up fee. The average dollar earning was $15.84 for an average of 75-minute session.

2.3. Theoretical Benchmarks

In our experiments, we use a unit revenue r_i of $40, a uniform demand distribution *U*[0,200] for the final customer demand (d_i), and a unit cost of $30, $20 and $10 for the low (*L*), medium (*M*) and high (*H*) CR conditions, respectively. Based on these sets of parameters, and using equations (3) and (4), we get that the per unit transfer prices (t_{ij}^*) that lead to the joint optimal solution in our decentralized systems are $29, $20 and $11 for the *L*, *M* and *H* CR conditions, respectively. In addition, solving equation (3) at t_{ij}^*, we get that corresponding ordering quantities (q_i^*) that coordinate the supply chain are 70, 100, 130 for the *L*, *M* and *H* CR conditions, respectively. Figure 3 shows the reaction functions (equation (3)) and the coordinating ordering quantities evaluated at t_{ij}^*, for each of the different CR conditions defined in our study.

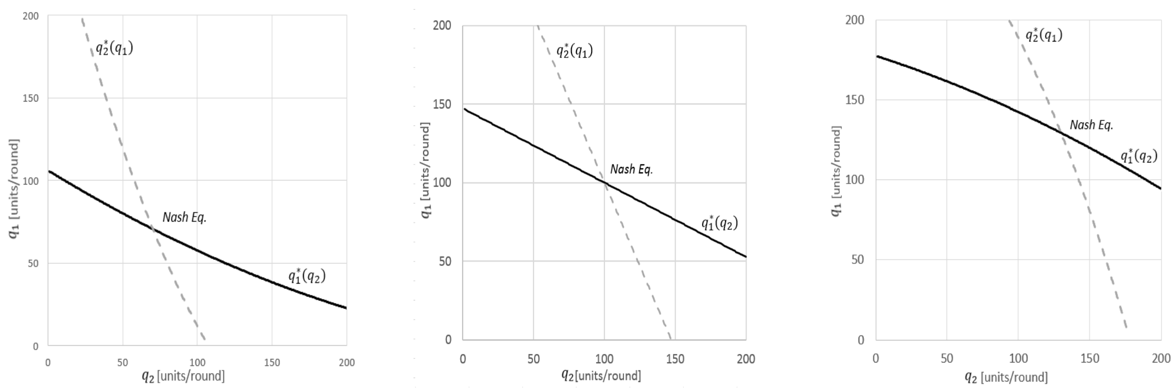

(a) CR = 0.25 (Low) (b) CR = 0.5 (Medium) (c) CR = 0.75 (High)

Figure 3. Reaction curves and Nash equilibria for (a) Low, (b) Medium, and (c) High CR conditions at t_{ij}^*.

Even if t_{ij}^* and q_i^* are the set of decisions that coordinate the overall supply chain, different transfer prices lead to different equilibria in the ordering quantities. By solving equation (3) for different values of t_{ij}, we obtain different values for the equilibrium order quantity. Figure 4 shows the optimal relationship between orders and transfer prices for our three different CR conditions. In general, we observe a linear and increasing relationship between the transfer price and the equilibrium order quantity, such that if subjects negotiate a higher (lower) price, it would be better to order more (fewer) units (Rudi et al., 2001). Therefore, for each particular transfer price a unique ordering quantity (the *Best Reply*) exists, which optimizes the supply chain. Notice that t_{ij}^* and q_i^* give the only equilibrium that makes the expected profit in the decentralized system equal to the profit of a centralized supply chain. We observe that t_{ij}^* decrease and q_i^* increase as CR increases.

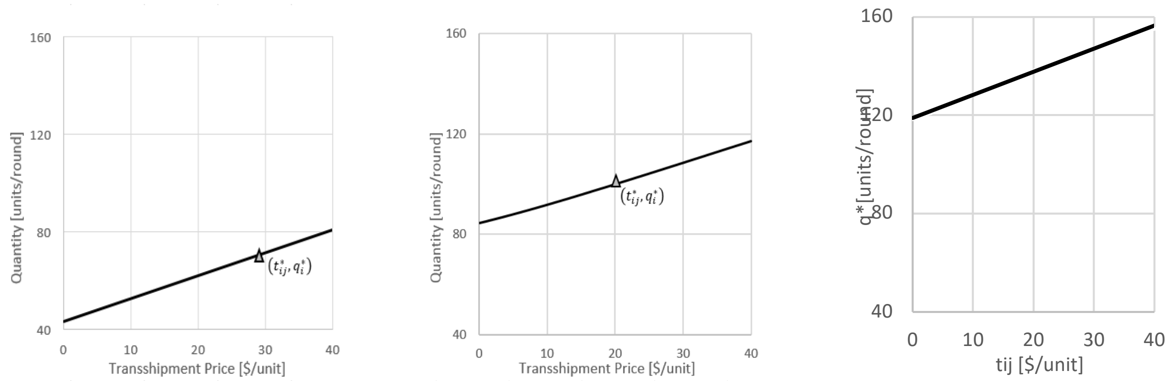

(a) CR = 0.25 (Low) (b) CR = 0.5 (Medium) (c) CR = 0.75 (High)

Figure 4. Optimal relationship between order and transfer prices for (a) Low, (b) Medium, and (c) High CR conditions.

Finally, equation (1) provides the optimal ordering quantities for a supply chain with no thransshipment involved. Using this equation, we obtain that 50, 100, and 150, are the optimal ordeing quantities for our low, medium and high profitability treatments, respectively. Table 1 summarizes the experimental predictions for all treatments in our experiment.

3. Hypotheses

From a rational perspective, self-interested subjects should make decisions aiming at maximizing their own expected profits (Berg et al., 1995). Therefore, we would expect subjects to make decisions based on the theoretical benchmarks that we presented in the previous section. However, evidence about the newsvendor problem has shown evidence for the pull-to-center behavior, in which individuals order less than the optimal order quantity for high-profit products and more than the optimum for low-profit products (Bolton & Katok, 2008; Bostian et al., 2008; Schweitzer & Cachon, 2000). In these situations, both in the high- and low-profit conditions, retailers' orders move closer to the mean of the customer demand. However, due to the risk pooling in the system with transshipments, participants in the High CR treatments with transshipment might be more prone to place higher orders than participants in the High CR treatment without transshipment. Therefore, we formulate two competing hypotheses about the effect of transshipments on average orders. H1a is based on predictions from the standard model.

Hypothesis 1a: *Average orders in High and Low CR treatments with transshipment will be closer to mean demand than average orders in corresponding treatments without transshipment, while average orders in the Medium CR treatments will be unaffected by transshipments.*

H1b is based on behavioral predictions in which risk aversion plays a role in the pull-to-center behavior in high CR conditions. Risk aversion generally decreases orders for any CR (Eeckhoudt et al., 1995; Schweitzer & Cachon, 2000). Allowing transshipments decreases risk, and therefore, allowing transshipments should cause orders to increase for all CRs.

Hypothesis 1b: *Average orders with transshipment will be higher than average orders without transshipment.*

If participants are interested in increasing the overall efficiency of the system, then they should set transfer prices in line with theoretical benchmarks as shown in Figure 4. Therefore, we formulate the following hypotheses.

Hypothesis 2b: *Average transfer prices will increase as CR decreases.*

We also formulate a competing hypothesis that links to operations work on bargaining and uses cooperative game theory ideas (see Leider & Lovejoy, 2016; Lovejoy, 2010). The notion of coordinating the channel by setting the appropriate transfer price is challenging and counterintuitive. Considering that the two retailers are symmetric, there is an obvious focal point for the transfer price. As the *zone of possible agreement* (ZOPA) is between zero and forty, buying a product from the other retailer at a transfer price above $40/unit will be unprofitable for the buyer, while selling the units at a transfer price below $0/unit will be unprofitable for the seller. Therefore, we formulate an alternative hypothesis claiming that transfer prices will be around the mean of 0 and 40.

Hypothesis 2c: *Average transfer prices will not be significantly different from 20 and will not be affected by the CR or by the decision sequence.*

In terms of the relationship between transfer prices and order quantities, Figure 4 shows that optimal orders increase as transfer prices increase regardless the CR condition.

Hypothesis 3: *There is a positive correlation between average transfer prices and average orders.*

Our next hypothesis considers the retailer's ability to transship and its effect on profits. Considering that transshipments improve the match between supply and demand regardless of the transfer price, the ability to transship should improve profits.

Hypothesis 4: *Average retailer profits will be higher in treatments with transshipment than in corresponding treatments without transshipment.*

Our last set of hypotheses deals with the effect of sequence on behavior. Standard theory implies that sequence should not affect behavior:

Hypothesis 5a: *Average order quantities, transfer prices and retailer profits will not be affected by the decision sequence.*

An alternative hypothesis is that the Quantity First condition may cause an *endowment effect* (Kahneman et al., 1991) because retailers who ordered more units now value these units more than retailers who ordered fewer units. This endowment effect may create asymmetries in some negotiations. In particular, the endowment effect can decrease the likelihood of an agreement.

Hypothesis 5b: *Successful agreements are less likely and average retailer profits are lower in the Quantity First treatments than in the corresponding Price First treatments.*

4. Experimental Results

In Table 2, we provide the averages and standard errors for order quantities in all treatments. The unit of analysis is cohort. To compute the average of the cohort, we follow two steps. First, we take the average order of an individual subject during the 30 simulated rounds. Second, we compute the average of the individual averages within each cohort. We use the Mann-Whitney U (Wilcoxon rank sum) Tests to make the comparisons that we report.

Table 2. Average order quantities and optimal order quantities conditional on observed transfer prices (best response order quantities)

		Sequence				
			Average Orders Observed		Average Best Reply Orders	
		No Transshipment	Price First	Quantity First	Price First	Quantity First
Critical Ratio	0.25	78.640** (4.433)	79.219†† (2.871)	82.145†† (0.853)	66.577 (2.206)	64.605 (1.465)
	0.50	95.257** (1.533)	95.794†† (0.748)	97.908†† (1.334)	103.820 (0.796)	102.920 (0.806)
	0.75	104.159** (2.607)	116.112†† (1.977)	112.694†† (2.334)	142.643 (0.428)	143.415 (1.012)

Standard Errors are in parentheses
Ho: $q = q^*$ *p<0.1, **p<0.05 (comparing no transshipment orders to optimal)
Ho: $q = q(t_{ij})$ †p<0.1 ††p<0.05 †††p<0.01 (comparing transhhipment orders to best reply)

Results for the *No Transshipment* treatments confirm the pull-to-center behavior. Differences between observed and optimal orders (see Table 1) are strongly significant in the High, Medium and Low treatments. Significance in the Medium treatment is consistent with risk aversion.

For treatments with transshipment, the proper order quantity comparison is between the observed averages, and the best reply given the observed transfer prices. We display these average best reply order quantities in the last two columns of Table 2. Differences between observed and best reply orders are strongly significant in all three CR conditions[1]. Note that average orders in the Medium CR treatments are significantly lower than the best reply, which is consistent with risk aversion.

When comparing observed orders in the treatments without transshipments to the orders in corresponding treatments with transshipment, we do not find significant differences for the Low and Medium CR treatments. However, observed orders in the High condition are significantly lower without transshipments ($p = 0.01$ in the Price First treatment and $p = 0.05$ in the Quantity First treatment). Thus, the data is not consistent with H1a, but it is partially consistent with H1b.

The coordinating transfer prices are 29, 20 and 11 for low, medium, and high CR conditions. Average observed transfer prices in the treatments with transshipment are summarized in Table 3. We find that the observed transfer prices do not differ for the different CR, contrary to H2a. The results partially support H2b. On the one hand, the average prices are slightly above twenty[2]. On the other hand, we find that there are no differences due to CR or to the decision sequence. Subjects' negotiations are not affected by the value of the CR reducing any possibility of coordinating the supply chain.

Table 3. Average transfer prices.

		Price First	Quantity First
Critical Ratio	0.25	25.988 (2.261)	25.69 (1.831)
	0.50	24.530 (0.958)	23.863 (1.104)
	0.75	24.940 (0.480)	25.255 (1.158)

Standard Errors are in parentheses

[1] The coordinating order quantities for treatments with transshipments are 70, 100, and 130 for low, medium, and high CR respectively. When comparing the average observed orders to coordinating order quantities, we observe the same qualitative differences as with the best-reply quantity comparisons. Average order quantities are also significantly different for the different CR conditions.

[2] The evolution of the average negotiated transfer prices are shown in Appendix 2.

4.1. Relationship between Transfer prices and Ordering Quantities

Figure 5 shows the average prices and the corresponding average orders for each treatment, grouped by CR condition. Each dot represents a cohort average (an independent observation). The figure shows that average transfer prices are in general between \$20/unit and \$30/unit regardless the CR condition. We can see that price/quantity combinations are below the best reply line for medium and high CR conditions and above the best reply line for the low CR condition. Deviations (vertical distance to the line) are larger for high and low CR conditions than for the medium CR condition. For an average transfer price of 24, it would be optimal to order around 104 units (for the medium CR condition) and 142 (for the high CR condition). However, subjects order on average less than 100 and 125 for medium and high CR, respectively. For low CR condition, for an average transfer price of 26, it would be optimal to order 67 units, while observed orders are closer to 80 units.

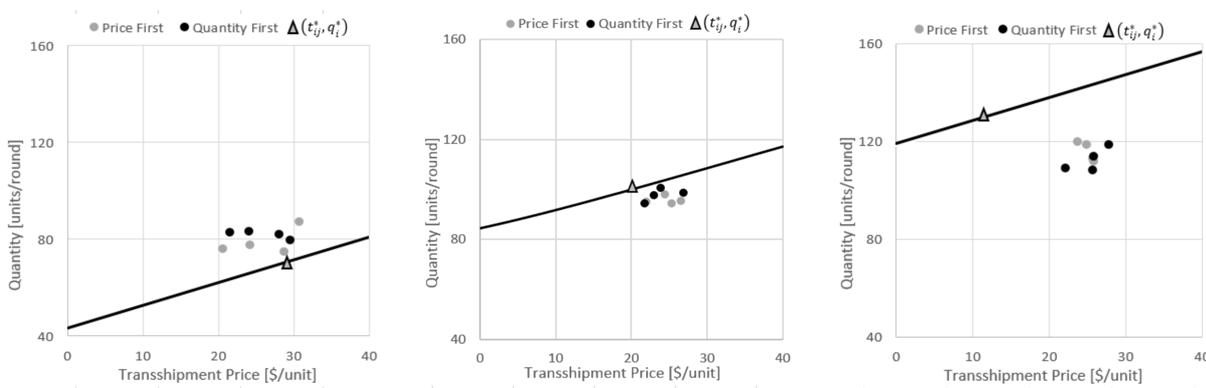

(a) Low CR (b) Medium CR (c) High CR

Figure 5. Observed relationship between average orders and transfer prices.

We now formally test for the relationship between prices and orders (H3). First, in the Price First condition, we wish to establish whether higher transfer prices result in higher order quantities. Then, in the Quantity First condition, we test whether higher order quantities result in higher transfer prices. We estimate two panel data regression models with individual random effects and cohort fixed effects, separately for the Price First and Quantity First treatments. Table 4 provides the results for both models[3]. The Price First model in Table 4 shows that higher transfer prices result in higher average orders. The Quantity First model shows that higher order quantities result in higher transfer prices. These results are consistent with the theoretical predictions presented in Figure 4, which shows a positive relationship between transfer price and ordering quantities, consistent with H3.

[3] Results persist for fixed effect models and for models including time-fixed effects instead of period number.

Table 4. Panel data estimations for relation between transfer prices and orders

	Price First Orders	Quantity First Transfer price
Constant	73.292*** (6.218)	17.981*** (1.272)
Transfer price	0.914*** (0.084)	-
Orders	-	0.063*** (0.004)
Low CR	-10.802 (8.171)	1.173 (1.342)
High CR	27.318*** (8.146)	0.985 (1.332)
Period Number	-0.113* (0.062)	2.039*** (0.298)
Wald Chi^2	209.35***	465.29***

Standard Errors are in parentheses
*refers to $p< 0.1$; ***refers to $p< 0.01$

4.2. Average Profits

Table 5 provides the averages and standard errors for average retailer profits in all treatments. Average retailer profits are higher in treatments with transshipment than in corresponding treatments without transshipments, giving support to H4. Average retailer profits are higher in the Price First than in the Quantity First condition, although the differences are not statistically significant. Next, we examine the effect of decision sequence further.

Table 5. Average retailer profits

		Sequence		
		No Transshipment	Price First	Quantity First
Critical Ratio	0.25	325.698 (15.965)	516.917** (19.576)	483.290** (22.610)
	0.50	1,151.964 (22.118)	1,400.583** (15.742)	1,365.972** (26.152)
	0.75	2,156.466 (49.398)	2,504.594** (26.053)	2,446.198** (22.515)

Standard Errors are in parentheses
Ho: Average profit is higher for treatments with transshipment than the corresponding no transshipment treatment** $p<0.05$.

4.3. Sequence of Decisions

Table 2 and Table 3 show that sequence has no significant effect either on order quantities or on transfer prices, which is consistent with H5a. Now we consider the effect of sequence on the probability of a successful negotiation. Interestingly, we observe that in some cases participants abort the negotiation process in disagreement, even though failing to agree on a transfer price is not rational because it reduces the expected profit for both retailers. In line with H5b, we find that the probability of a successful negotiation is lower in the Quantity First condition than in the Price First condition (Table 6).

Table 6. Average agreement probabilities

		Agreement Probability	
		Price First	Quantity First
Critical Ratio	0.25	0.938* (0.026)	0.855 (0.045)
	0.50	0.979** (0.005)	0.876 (0.010)
	0.75	0.968* (0.012)	0.900 (0.046)

Standard Errors are in parentheses
Ho (one sided): Agreement probability is lower for Price First than for Quantity First: * $p<0.1$, ** $p<0.05$

We further analyze the effect of decision sequence on the probability of agreement by fitting a logit regression with random effects for individuals and fixed effects for cohorts. The period number is included in the regression to control for learning, and "Price First" is used as a binary variable. We fit this model for each CR condition separately, and provide estimates in Table 7. The Price First variable is positive and significant for all the CR conditions, confirming that the agreement on a transfer price is more likely in the Price First condition. These results also support H5b. The sequence of decisions does not affect the actual orders or negotiated prices, but it does affect the outcome of the negotiation. When the negotiation over the transfer price takes place after orders have been placed, the probability of a successful negotiation decreases.

Table 7. Logit estimations for probability of agreement.

	Low CR	Medium CR	High CR
Constant	2.100*** (0.344)	2.054*** (0. 333)	2.657*** (0.440)
Price First	1.177*** (0. 299)	1.922*** (0. 313)	1.159*** (0. 355)
Period Number	-0. 014 (0. 009)	-0.004 (0.010)	0. 030** (0. 011)
-LL	581.363	487.108	409.952

Standard Errors are in parentheses
*** refers to $p< 0.01$; ** refers to $p< 0.05$

In

Table 8 we present the results of regressions in which we measure the effect of sequence on profits. We see that the Price First indicator variables are positive and significant for the low CR (strongly significant), and medium CR (weakly significant) conditions. These results confirm that the lowered efficiency that is due to lower agreement rate also can translate into lower profits, partially confirming H5b.

Table 8. Panel data estimations for profits

	Low CR	Medium CR	High CR
Constant	245.649*** (47.991)	910.821*** (85.158)	1941.121*** (130.879)
Price First	112.421*** (39.960)	136.839* (78.631)	203.483 (124.990)
Period Number	11.337*** (1.714)	20.583*** (2.010)	24.874** (2.504)
Wald Chi^2	53.090***	108.270***	101.400***

Standard Errors are in parentheses
*** refers to $p< 0.01$; ** refers to $p< 0.05$; * refers to $p< 0.1$

5. Conclusion

Channel coordination is a central question in the study of inventory management. There exists a large body of analytical inventory management literature which assumes that inventory managers place optimal orders. But behavioral studies have shown that this is not the case— decision makers frequently deviate from theoretical benchmarks. To the best of our knowledge, this paper is one of the first to systematically analyze the behavioral biases observed in a supply chain with transshipments, and we are the first to jointly analyze transfer price negotiation order placing decisions.

We investigate a setting in which retailers, following a realization of stochastic demand, can transship units from locations with excess inventory to locations with excess demand. When the system consists of autonomous retailers, Rudi et al. (2001) showed that there exists a set of transfer prices that induce these retailers to place system-optimal orders. We report on a set of laboratory experiments that we designed to gain insight into how human decision-makers set transfer prices and place orders. We examine settings in which retailers negotiate transfer prices through free form bargaining, and compare the performance of a system with and without transshipments. We also compare a setting in which decision makers negotiate transfer prices before they place orders, to one in which they negotiate transfer prices after they placed orders.

We find that transfer prices that come out of negotiations have little relationship with coordinating prices. In fact, unlike coordinating transfer prices that are a function of the critical ratio, transfer prices we observe do not differ across systems with different critical ratios. Instead, these transfer prices are roughly between the retail and wholesale price.

Allowing transshipments does improve profitability, even though ordering behavior systematically differs from optimal. Retailers continue to order more than they should in the low critical ratio condition, and less than they should in the medium and high critical ratio conditions. Therefore, transshipment improves the efficiency of the system, not by inducing better orders, but by better matching supply and demand.

Negotiations over the transfer price do not always successd—occasionally negotiations fail, in which case transshipment does not take place. We find that negotiation failure happens more often when retailers negotiate transfer prices after they placed their orders. We conjecture that this may have to do with the so-called endowment effect—retailers who placed high orders anticipate valuing their excess inventory more than retailers who placed lower orders.

We draw three practical conclusions from our research. First, transshipment is unambiguously beneficial, resulting in higher profitability, and therefore, when transshipment is feasible, it should be encouraged. Second, when retailers negotiate transfer prices they do not consider how these prices affect the system as a whole, but rather they seem to use a simple heuristic by roughly speaking, splitting the difference between the retail price and the wholesale price. Therefore, transfer prices should, when possible, be set centrally. Our third conclusion is that if retailers are left to their own devises in negotiating transfer prices, then it is better to conduct this negotiation before, rather than after, the order has been placed. This way it is more likely that the negotiation will succeed.

References

Arsu, S. (2009). Turkey And Russia Conclude Energy Deals. *New York Times*, A7. Retrieved from http://www.nytimes.com/2009/08/07/world/europe/07turkey.html

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, Reciprocity, And Social History. *Games and Economic Behavior*, 10(1), 122–142.

Bolton, G. E., & Katok, E. (2008). Learning By Doing In The Newsvendor Problem: A Laboratory Investigation Of The Role Of Experience And Feedback. *Manufacturing & Service Operations Management*, 10(3), 519–538.

Bostian, A. A., Holt, C. A., & Smith, A. M. (2008). Newsvendor “Pull-to-Center” Effect: Adaptive Learning In A Laboratory Experiment. *Manufacturing & Service Operations Management*, 10(4), 590–608.

Chen, K., & Li, S. (2018). *Modelling Behaviors in Inventory Sharing Among Decentralized Retailers*.

CREG. (2009). Colombia Aumenta Exportaciones De Energía A Ecuador. Retrieved from http://www.creg.gov.co/index.php/noticias/749-noticias-270

Dong, L., & Rudi, N. (2004). Who Benefits From Transshipment? Exogenous Vs. Endogenous Wholesale Prices. *Management Science*, 50(5), 645–657.

Eeckhoudt, L., Gollier, C., & Schlesinger, H. (1995). The Risk-Averse (And Prudent) Newsboy. *Management Science*, 41(5), 786–794.

Herer, Y. T., Tzur, M., & Yücesan, E. (2006). The Multi-Location Transshipment Problem. *IIE Transactions*, 38(3), 1–33.

Hu, X., Duenyas, I., & Kapuscinski, R. (2007). Existence Of Coordinating Transfer prices In A Two-Location Inventory Model. *Management Science*, 53(8), 1289–1302.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies : The Endowment Effect , Loss Aversion , And Status Quo Bias. *American Economic Association*, 5(1), 193–206.

Kranenburg, A. A., & van Houtum, G. J. (2009). A New Partial Pooling Structure For Spare Parts Networks. *European Journal of Operational Research*, 199(3), 908–921. Retrieved from http://dx.doi.org/10.1016/j.ejor.2009.01.057

Krishnan, K. S., & Rao, V. R. K. (1965). Inventory Control In N Warehouses. *Journal of Industrial Engineering*, 16, 212–215.

Leider, S., & Lovejoy, W. S. (2016). Bargaining In Supply Chains. *Management Science*, 62(10), 3039–3058. Retrieved from http://pubsonline.informs.org/doi/10.1287/mnsc.2015.2273

Lovejoy, W. S. (2010). Bargaining Chains. *Management Science*, 56(12), 2282–2301. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1100.1251

Rudi, N., Kapur, S., & Pyke, D. F. (2001). A Two-Location Inventory With Transshipment And A Local Decision Making. *Management Science*, 47(12), 1668–1680.

Schweitzer, M. E., & Cachon, G. P. (2000). Decision Bias In The Newsvendor Problem With A Known Demand Distribution: Experimental Evidence. *Management Science*, 46(3), 404–420.

Sošić, G. (2006). Transshipment Of Inventories Among Retailers: Myopic Versus Farsighted Stability. *Management Science*, 52(10), 1493–1508.

Stauffer, J. M., Pedraza-Martinez, A. J., & Van Wassenhove, L. N. (2016). Temporary Hubs For The Global Vehicle Supply Chain In Humanitarian Operations. *Production and Operations Management*, 25(2), 192–209.

Tagaras, G. (1989). Effects Of Pooling On The Optimization And Service Levels Of Two-location Inventory Systems. *IIE Transactions*, 21(3), 250–257.

Villa, S., & Castañeda, J. A. (2017). *Transshipments in Supply Chains: A Behavioral Investigation*.

Werdigier, J., & Dougherty, C. (2007). ECB Acts To Dispel Clouds Over European Credit Markets. *International Herald Tribune*, 23.

Zhao, H., & Bisi, A. (2010). Optimal Operating Policies In A Commodity Trading Market With The Manufacturer's Presence. *Naval Research Logistics*, 57(2), 127–148.

Zhao, H., Deshpande, V., & Ryan, J. K. (2005). Inventory Sharing And Rationing In Decentralized Dealer Networks. *Management Science*, 51(4), 531–547. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1040.0321

Zhao, H., Xu, L., & Siemsen, E. (2016). *A Behavioral Perspective on Inventory Sharing*.

Appendices

Appendix 1. Experimental Instructions for treatment with High Critical Ratio and Price First

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you will earn money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will answer it. We ask you not to talk with one another for the duration of the experiment.

The experiment consists of **30** rounds. Based on your decisions, every round you will earn some experimental revenues measured in experimental currency units (ECU). At the end of the session, your ECU earnings will be converted to US dollars at a pre-specified rate and paid to you in cash along with a $5 show-up fee. All earnings are confidential.

In this session, you will be in the role of a retailer. Each round, you and another participant, also in the role of a retailer, will first negotiate over a **transfer price** of a product you are both selling and then each of you will decide an **order quantity**. Each round, you will be matched with a different person.

Overview of the Game

You are a retailer who sells a single (fictional) product. In each round, you order units of the product from a supplier at a cost of **10** ECU per unit, and sell units to your customers at a price of **40** ECU per unit. The other retailer with whom you are matched in this round is doing the same thing. The **customer demand** (D) for your units and for the other retailer's units is an integer from 0 to 200, with each integer equally likely. The two demands will be determined randomly and independently, so they will in general be different.

You will have **two tasks** each round:

i. Your task at the start of each round is to 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.

ii. Once you agree on the transfer price, your task is to decide **how many units to order** to the supplier to satisfy the uncertain customer demand (see attached page that shows the order amounts that lead to highest expected profit corresponding to different transfer prices, however, you are free to order any number between 0 and 200).

Then, the customer demands will be randomly drawn, the transfer of units, if any will be automatically made, and finally the profits will be calculated and displayed on your screen.

Example:

Suppose the transfer price you negotiated is 25. Now, suppose you decide an order of 90 and the other retailer decide an order of 105. Finally, suppose your demand was 50 and the other retailer's demand was 150.

Your profit from customer sales: 40 x 50 – 10 x 90 = 1100.

The other retailer's profit from customer sales: = 40 x 105 – 10 x 105 = 3150

You have excess inventory of 90 – 50 = 40

The other retailer has excess demand of 150 – 105 = 45.

So you transfer 40 units to the other retailer, who pays you 25 per unit.

Your profit from transfer = 40 x 25 = 1000

The other retailer's profit from transfer = 40 x (40-25) = 600.

Your profit for the round: 1100 + 1000 = 2100

The other retailer's profit for the round: 3150 + 600 = 3750.

Your *goal* is to *maximize* the profit you make totaled over all the rounds of the game.

How to Negotiate

Every round you negotiate the transfer price with the other retailer over a computer interface. The system will allow you to make and receive price offers. Once you agree to one specific price, click the "Accept Offer" button. In case you do not make an agreement, the system will move to the ordering stage but no transfers will take place in that round. Therefore, you will not benefit from the extra potential profit from the transfers.

How to Order

Every round you will have to decide how many units you would like to order to the supplier over the computer interface. You make your decision by typing the order quantity into the text box and clicking "Accept" button.

Additional Information you will see

Every round you will have access to the main cost parameters during the negotiation and ordering phases. During the round results, you will receive information from all previous periods, including your demand, units shipped and received by transfers, profit, ordering quantities and agreed transfer prices.

Relation Price and Ordering Quantities

Based on a negotiated price, the order amounts that lead to highest expected profit are shown in the following table:

Price	Order
0	118
1	119
2	120
3	121
4	122
5	123
6	124
7	125
8	126
9	127
10	129

Price	Order
11	129
12	130
13	131
14	132
15	134
16	135
17	136
18	136
19	137
20	138

Price	Order
21	139
22	140
23	141
24	141
25	142
26	143
27	144
28	146
29	147
30	147

Price	Order
31	148
32	149
33	150
34	151
35	152
36	152
37	154
38	154
39	155
40	155

Screenshot for the negotiation screen

The following screen shot is used as an example to illustrate how you use the computer interface to make and accept offers.

1. Timer: On the top right side of the screen, you will see a timer that shows how many minutes and seconds are remaining in this round of negotiation.

2. Main parameters: On the top left side of the screen, you will see the round number, the main cost parameters of the system and information about the demand distribution.

3. History: This box will show the history of the price offers place by you and the other retailer during the current round.

4. New Offer: You can make offers to the other players by typing the proposed price into the text box and clicking "*Submit*" button. This price offer will appear in the *History* area.

5. Current Offer: An agreement will be reached when one of the players accepts other player's offer by clicking the "*Accept Offer*" button.

6. Terminate now: Clicking this button immediately ends the round; aborting the bargaining process in disagreement and therefore no transfers will take place that round.

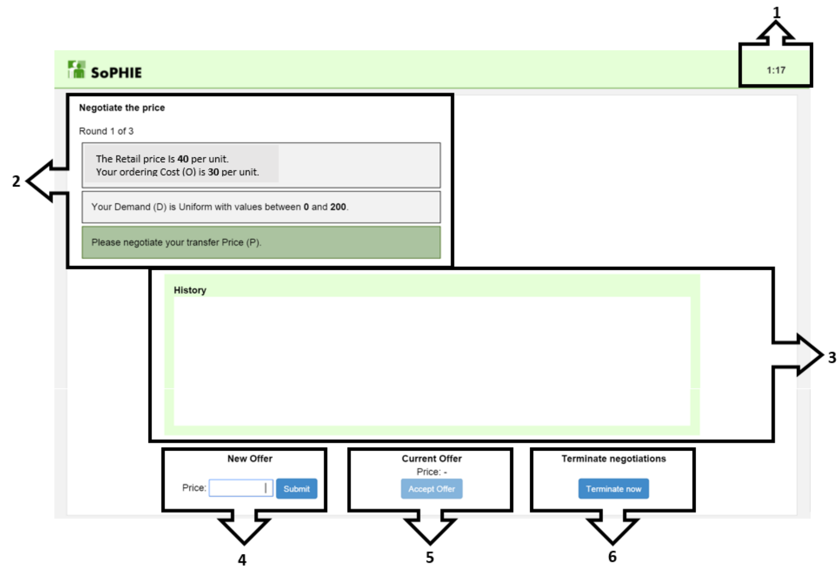
1
SoPHIE
1:17
2
Negotiate the price
Round 1 of 3
The Retail price is 40 per unit.
Your ordering Cost (O) is 30 per unit.
Your Demand (D) is Uniform with values between 0 and 200.
Please negotiate your transfer Price (P).
History
3
New Offer
Price:
Submit
Current Offer
Price: -
Accept Offer
Terminate negotiations
Terminate now
4
5
6

Appendix 2: Evolution of average decisions

Figure A1. Evolution of average ordering quantities for (a) Low, (b) Medium and (c) High CR conditions

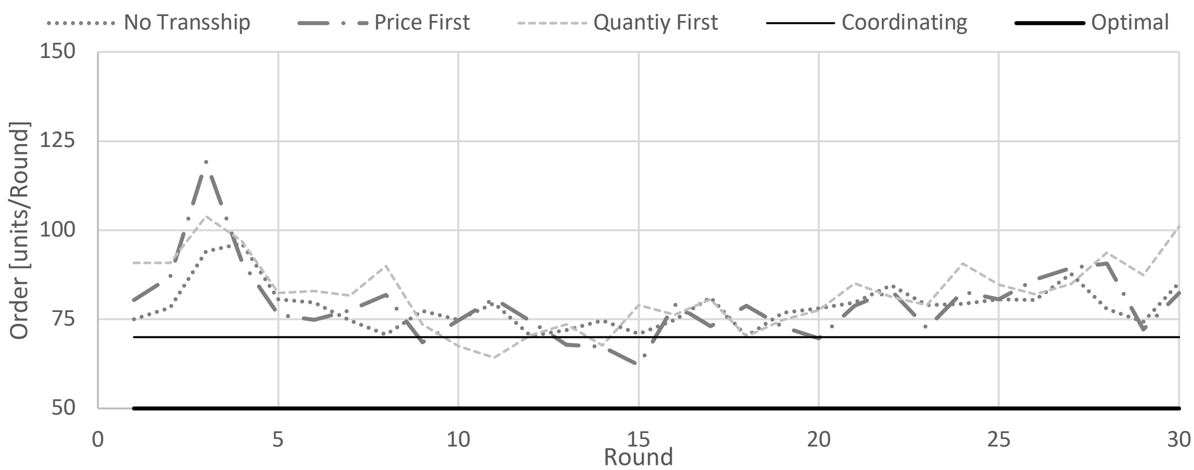

(a)

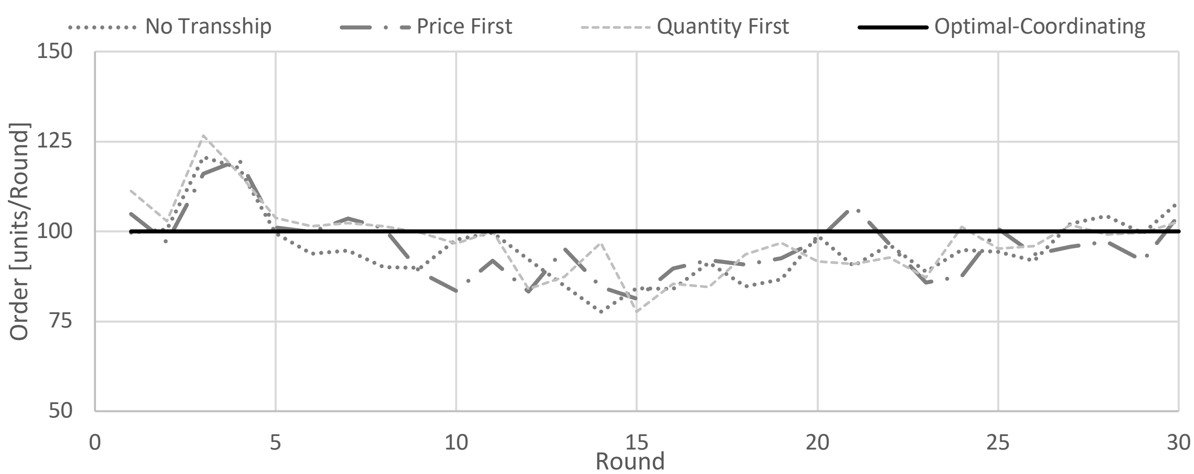

(b)

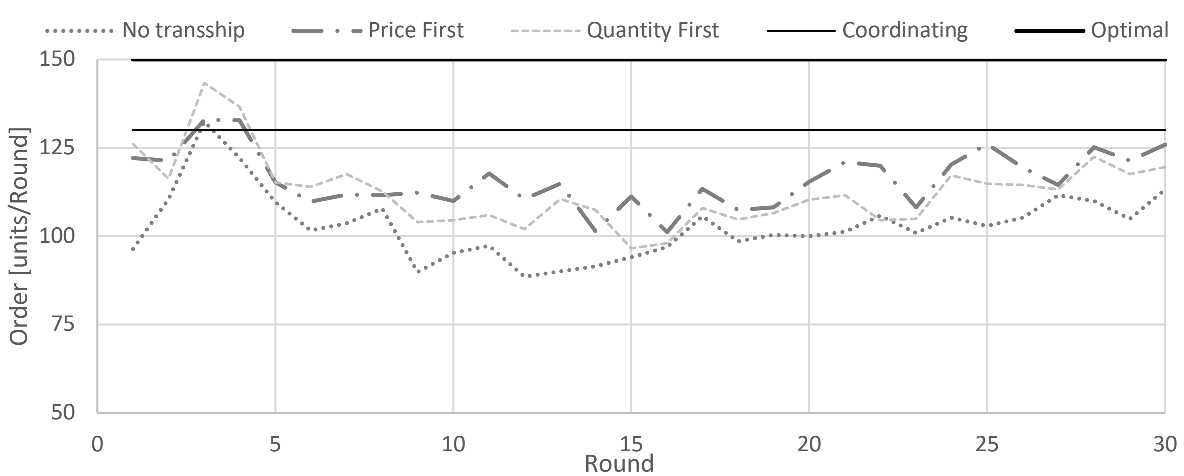

(c)

Figure A2. Evolution of average transfer price for (a) Low, (b) Medium and (c) High CR conditions

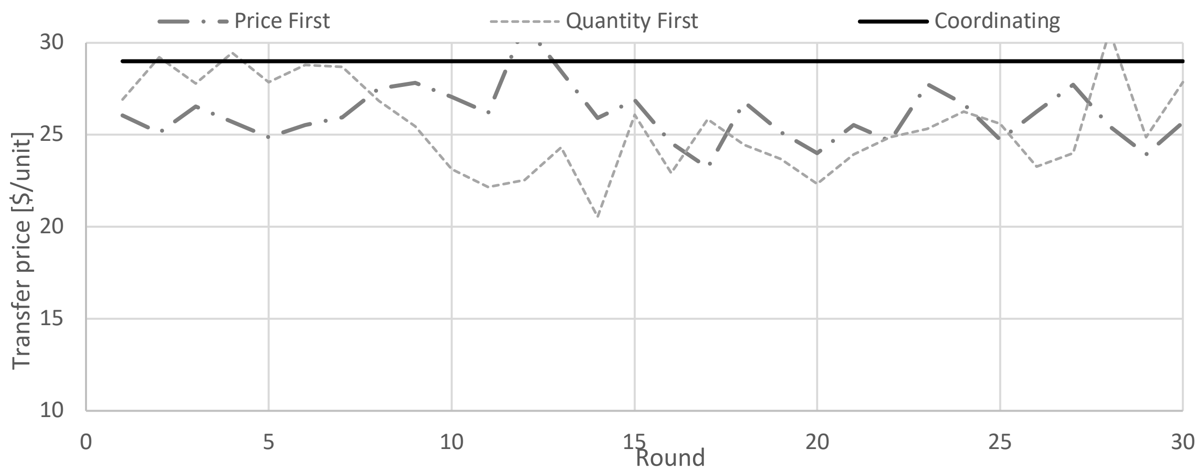

(a)

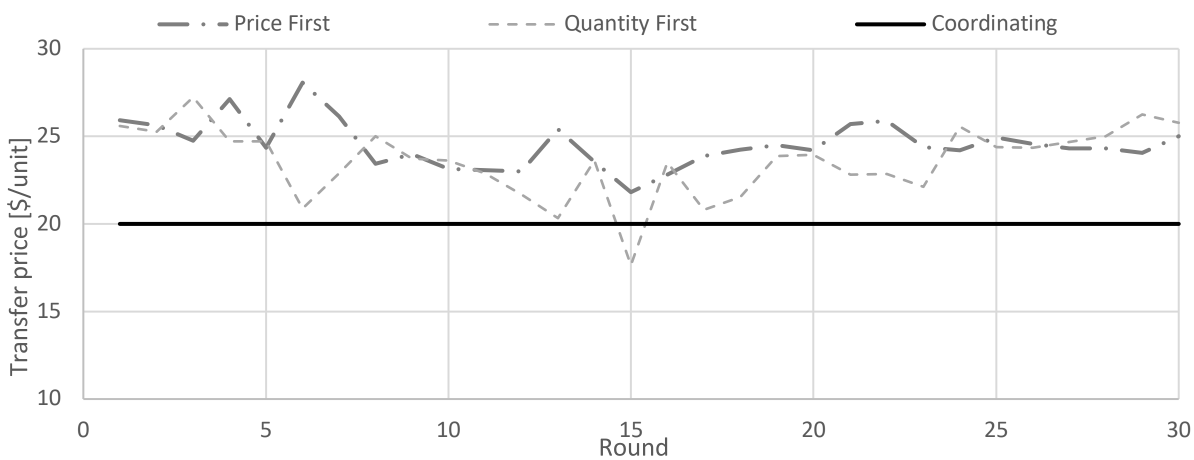

(b)

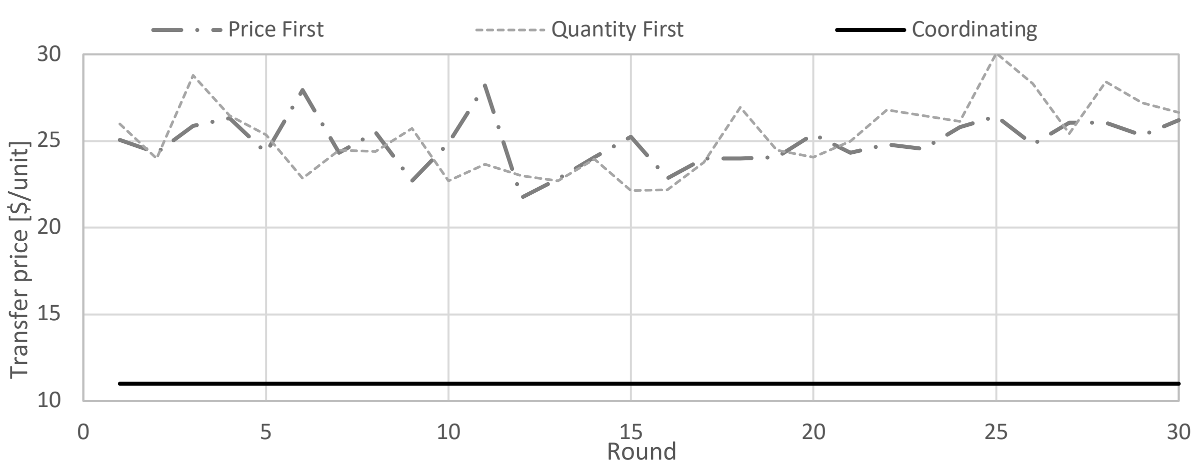

(c)