March, 2013

How can we improve the performance of supply chain contracts? An experimental Study

Ehsan Elahi, *University of Massachusetts Boston*
Narasimha Lamba, *University of Massachusetts Boston*
Chinthana Ramaswamy, *University of Massachusetts Boston*

How can we improve the performance of supply chain contracts? An experimental Study

Ehsan Elahi, Narasimha Lamba, Chinthana Ramaswamy
University of Massachusetts, Boston

To appear in the International Journal of Production Economics (2013)

Abstract

Although optimal forms of supply chain contracts have been widely studied in the literature, it has also been observed that decision makers fail to make optimal decisions in these contract setups. In this research, we propose different approaches to improve the performance of supply chain contracts in practice. We consider revenue sharing and buyback contracts between a rational supplier and a retailer who, unlike the supplier, is susceptible to decision errors. We propose five approaches to improve the retailer's decisions which are in response to contract terms offered by the supplier. Through laboratory experiments, we examine the effectiveness of each approach. Among the proposed approaches, we observe that offering free items can bring the retailer's effective order quantity close to the optimal level. We also observe that the retailer's learning trend can be improved by providing him with collective feedbacks on the profits associated with his decisions.

Keywords: Supply Chain Contracts; Revenue Sharing; Buyback; Behavioral Operations Management

1. Introduction

Supply chain contracts have been extensively studied by researchers. A large stream of research in this field considers a two echelon supply chain consisting of a supplier (seller) and a retailer (buyer) who sells a seasonal (fashion) product to a market with random demand. Due to usually lengthy production and distribution lead times (Fisher & Raman, 1996), the retailer has to decide about the order quantity (initial inventory level) long before the start of the selling season. Under this setup, the retailer faces a classical Newsvendor inventory problem. That is, if the retailer's order quantity is less than the realized demand, the retailer faces with inventory shortage (unmet demand), while if the order quantity is more than the realized demand the retailer is left with unsold inventory, which should be discarded or salvaged with a very low price. The classical Newsvendor solution identifies the optimal order quantity which maximizes the retailer's expected profit.

In a simple wholesale price contract, the retailer faces all the risk and the wholesale price that maximizes the supplier's profit causes the retailer to order a quantity less than the value that maximizes the channel profit (Spengler, 1950). To avoid this situation, the supplier can offer a contract in which she provides the retailer with proper economic incentives to order the quantity that maximizes the supply chain profit (a coordinating contract). In this research, we consider two types of coordinating contracts: revenue sharing and buyback. In a revenue sharing contract, the supplier offers a relatively low wholesale price but asks the retailer to share part of the revenue of every item sold. Revenue sharing contracts have been used successfully (among other industries) in the video-rental industry (Cachon & Lariviere, 2005). In a buyback contract, the supplier buys back any unsold item from the retailer with a price lower than the wholesale price. Buyback contracts are common practice in the publishing, software, and pharmaceutical industries (Padmanabhan & Png, 1995). In both contracts, the supplier shares part of the retailer's risk of facing a random demand.

Although the theoretical benefits of optimal Newsvendor solutions and coordinating contracts have been widely studied, it is also known that retailers fail to place the optimal order quantities in practice. Fisher & Raman (1996) and Corbett & Fransoo (2007) show industry evidence that managers' inventory decisions systematically deviate from the optimal quantities. Fisher & Raman (1996) show that managers' less-than-optimal production quantity, at a ski apparel manufacturer, resulted in a profit which was 60% less than their calculated optimal

profit. Corbett & Fransoo (2007) study inventory decisions of 51 small businesses. They show that the inventory decisions deviates from the optimal decisions calculated by a Newsvendor model. They show that the deviations are consistent with the prospect theory predictions.

Almost all the research papers in this field have focused on finding how and why decision makers' order quantities deviate from the optimal values (we will briefly review these papers in section 2). The more important question of how this deviation could be avoided, however, has received little attention in the existing literature. As an attempt to fill this gap, we explore possible ways through which we can improve the performance of a supply chain by inducing the retailer to choose order quantities close to the channel's optimal order quantity.

Here, we consider an ideal supplier whose decisions are rational and sets the parameters of the contract according to their theoretical optimal values. The retailer, however, is assumed to be prone to behavioral misjudgments and errors. Therefore, the order quantities chosen by the retailer can systematically deviate from the optimal values. The retailer's suboptimal decision has a negative impact on his profitability as well as the supplier's and the channel's profitability. Hence, the supplier tries to design the contract terms or offer additional information to address the inefficiency in the retailer's decision and increase her (and consequently channel's) profit.

We explore five approaches which could possibly improve the performance of a revenue sharing or buyback contracts. We first identify the concept or logic behind each approach and then verify its effectiveness through laboratory experiments. Three of these approaches concern the contract terms which the supplier offers the retailer. The other two approaches concern providing the retailer with additional information or feedback that might help him to make better decisions. In our first approach we consider a new type of contract which is a combination of revenue sharing and buyback contracts. The second approach examines the possibility that risk-aversion is the source of suboptimal decisions. If this is the case, then a coordinating contract that is designed for a risk-averse (not a risk-neutral) retailer should result in an optimal order quantity. The third approach considers the offering of free items by the supplier. If the number of free items offered increases with the size of the order, the retailer might be encouraged to increase his order quantity. Moreover, these free items increase the number of items in the supply chain. In our fourth approach we examine the impact of providing the retailer with visual information about the nature of demand uncertainty. This could possibly discourage the retailer to follow shortsighted strategies such as demand chasing. In our last approach we provide the

retailer, in each decision round, with a new performance measure that shows the collective impact of last decision if the current order quantity were the decision for previous decision rounds as well. This new piece of information should also discourage the retailer to follow a demand chasing strategy.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the theoretical background of the problem, explains the general experimental setup, and shows the results of our benchmark experiments. Sections 4 to 8 present the five studies through which we explain and investigate the effectiveness of each of our approaches to improve the performance of the supply chain. Section 9 concludes the paper with a summary of our results.

2. Related Literature

In this research we study a two echelon supply chain consisting of a supplier and a retailer, in which the retailer faces a classical Newsvendor problem. In a Newsvendor problem a decision maker, who faces a random demand for a single selling period, has to decide about the quantity (inventory level) he needs to order/manufacture before the beginning of the period. Optimal order quantity is a trade-off between overage and underage inventory costs (Arrow et al, 1951). In its basic form, the Newsvendor problem has an elegant solution which can be applied to many applications other than single period inventory problems (e.g. multi-period inventory problems, capacity selection, choice of staffing level, time should be allocated to a given task, etc.). A review of different extensions of this widely studied problem is beyond the scope of this paper. We refer the interested readers to Olivars et al (2005) and Qin et al (2011) for reviews of this literature.

Although the elegant structure of the Newsvendor problem has let researchers develop analytical solutions for different variants of this problem, it has been known for a while that decision makers facing this problem deviate from the theoretical optimal solution in practice. Fisher & Raman (1996) and Corbett & Fransoo (2007) provide industry evidence for this deviation. These observations have attracted many researchers' attention as to how and why this deviation occurs. There are many research papers that try to explore this behavior through laboratory experiments.

Schweitzer & Cachon (2000), in a set of laboratory experiments, observe that the subjects' order quantity always fall between the average demand and the optimal value. That is, for a high profit margin product, for which the optimal order quantity is higher than the average demand, the subjects' average order quantity is also higher than the average demand, but lower than the optimal value. For low profit margin products, for which the optimal order quantity is lower than the average demand, the subjects' average order quantity is lower than the average demand, but higher than the optimal value. This behavior is known as "pull to center." The authors attribute this behavior to *ex post* inventory error, anchoring, and insufficient adjustment. Through their experimental analysis, they rule out the influential impacts of other factors like risk aversion, loss aversion, prospect theory preferences, waste aversion, and stock-out aversion. Our research is different from Schweitzer & Cachon (2000) since we study coordinating contracts between a supplier and a retailer, while they study a single Newsvendor decision maker (retailer). Similar to their results, our subjects (retailers) demonstrate the pull to center behavior. We also rule out the influential role of risk aversion in retailers' suboptimal decisions, which is similar to what they conclude (using a completely different method).

Building on Schweitzer & Cachon's (2000) model, Bostian et al (2008) use an adaptive learning algorithm to justify the pull to center behavior. Unlike Schweitzer & Cachon (2000), Bostian et al (2008) find that subjects' average order quantity is very close to the mean demand in the first round of decisions. However, order quantities diverge from the mean demand in successive decision rounds. The authors' adaptive learning model explains the pull to center behavior and shows that subjects respond to recent gains and losses. They also show that payoff insensitivity to order quantity in the vicinity of the optimal order quantity could not explain the pull to center behavior. One of the approaches that we propose in this paper (the collective feedback approach) is partly based on Bostian et al (2008) observation that subjects respond mostly to recent gains and losses.

Using a model based on the quantal choice theory, Kremer et al (2010) show that decision makers' random errors cannot be the main source of deviation from the optimal order quantity. They show that context dependent decision strategies such as anchoring, chasing, or inventory error minimizing play more influential roles. The conclusion that context dependent and systematic biases play the influential role in subjects' suboptimal decisions (rather than their random errors) suggest that there should be ways to counter these systematic biases. In this

research, we propose approaches to work against these systematic biases and bring the supply chain profit close to its optimum level.

Bolton & Katok (2008) study the impact of experience and feedback on the subjects' behavior. The authors show that subjects' decisions improve over the 100 rounds of decisions in their experiments. However, they report a very slow rate of improvement. They also show that restricting subjects' decisions to 10 rounds of standing orders can improve the quality of decisions (they increased the number of order quantities to 1000 rounds in this experiment). Among other results, the authors show that limiting the number of options from 100 possible order quantities in each decision round to 9 or 3 options cannot improve the quality of decisions. Their other results include examining the impacts of providing the subjects with extra information such as the payoff for the foregone options or providing payoff statistics for different decision options at the beginning of the experiment. They show that none of the mentioned information and feedback can improve the outcome.

Lurie & Swaminathan (2009) also use laboratory experiments to study the impact of feedback frequency on the quality of decisions in a Newsvendor problem. More specifically, they examine the performance of a Newsvendor when an order quantity decision is standing for a set of rounds and the profit feedback is provided at the end of each set of rounds. They show that the Newsvendor's profit can increase with a decrease in feedback frequency. They also find that introducing costs to make changes in successive decisions does not improve the Newsvendor performance when the feedback frequency is high. The authors show when the feedback frequency is high, decision makers tend to limit their information access to the most recent set of presented data, hence, they are more prone to overreacting to noisy feedback. They also show feedback frequency plays a more influential role than decision frequency. Our collective feedback approach provides the subjects with a feedback similar to what Lurie & Swaminathan (2009) provide in their experiment with standing orders. As we show in section 8, our collective feedback does not have any of the practical limitations that exist when we use standing orders.

Different types of supply chain contracts have been studied under different experimental settings. Keser & Paleologo (2004) and Loch & Wu (2008) study wholesale price contracts. Coordinating contracts are studied by Ho & Zhang (2008), Katok & Wu (2009), and Davis (2010). Two-part tariffs and quantity discount contracts are studied by Ho & Zhang (2009). Katok & Wu (2009) study buyback and revenue sharing contracts. Davis (2010) investigates pull

contracts (both wholesale price and coordinating). The common result in all these papers is that these contracts fail to coordinate the supply chain in experimental setups.

Katok & Wu (2009) separate the interaction of suppliers and retailers by letting subjects play the role of retailers against computerized (fully rational) suppliers, or the role of suppliers against computerized retailers. In this way, they can avoid the fairness effect which appears when human retailers interact with human suppliers. They find that the way demand distribution is presented (framed) to subjects affects their decision quality. The authors also show that in a high demand situation, the retailer performs better under a buyback contract than under a revenue sharing contract. The difference, however, decreases and disappears with experience. Similar to Katok & Wu (2009), in this research, we separate the interaction of suppliers and retailers by asking subjects, who play the role of retailers, to respond to contracts offered by computerized suppliers.

All the above-mentioned papers try to explain the reasons behind retailers' suboptimal decisions, which lead to less-than-optimal profits for all parties. The existing literature, however, fails to address how we can improve these suboptimal decisions. To fill this gap, we try to identify approaches to improve the retailer's order quantity decisions, which could lead to higher supply chain profits.

Becker-Peth et al (2011), through laboratory experiments, study the performance of buyback contracts. They show a Newsvendor retailer responds differently to different contract parameters even if these parameters result in the same critical ratio[1]. They build a behavioral model that depends on the buyer's anchoring to mean demand, loss aversion, and different valuation of income. The authors first estimate the parameters of the model through subjects' responses to a wide range of contract parameters and then find a contract that could result in the channel's optimal solution. They also show that the contract can be customized for each subject. Similar to our paper, Becker-Peth et al (2011) try to find a contract that results in an order quantity close to the optimal value. Their approach, however, cannot control the share of supplier's profit from the total channel's profit. Therefore, the supplier cannot aim for a target profit level when she offers a contract in this approach. A detailed review of experimental studies on other forms of contracts is beyond the scope of this paper. A recent review of this literature can be found in Katok (2011).

[1] The Newsvendor critical ratio is defined by (p-c)/p, where p is the unit selling price and c is the unit cost. Theoretically, the optimal order quantities of two Newsvendors are the same when the critical ratios are equal.

In one of our studies in this research, we consider the impact of offering free products by the supplier to induce the retailer to place higher order quantities. Through a series of experiments, Shampanier et al (2007) show people usually perceive the benefits associated with free products to be higher than what classical economics predicts. They attribute this behavior to people's difficulty in mapping their utility. Hence, they are more inclined toward a free product since it is an option with no downside.

3. Preliminaries

3.1. Theoretical Background

In this subsection, we describe our supply chain model and present the theoretical formulation of the basic contracts that governs the supply chain. The theory behind each approach (to improve retailer's decisions) will be presented separately in the corresponding sections. We consider the supply chain of a seasonal product which consists of a supplier and a retailer. The supplier offers the retailer a contract that specifies the payment scheme between the supplier and the retailer. We will study three forms of contracts: (*a*) wholesales price, (*b*) revenue sharing, and (*c*) buyback. The retailer faces a random demand in each selling season. The distribution of this demand is common knowledge. Based on the received contract and the demand distribution, the retailer chooses how much to order from the supplier (the order quantity). Therefore, the retailer faces a classical Newsvendor problem. The optimal order quantity of a Newsvendor decision-maker can be found from (Silver et al, 1998)

$$q^* = F^{-1}\left(\frac{c_u}{c_u + c_o}\right), \tag{1}$$

where c_u is the unit inventory underage cost, c_o is the unit inventory overage cost, and $F^{-1}(.)$ denotes the inverse of cumulative distribution function of random demand. We will show how the values of c_u and c_o can be identified in each contract type.

We assume the demand is uniformly distributed between *A* and *B*, $D \sim U(A,B)$. The retailer sells each unit of the product with a price *p*. The unit production cost for the supplier is *c*. The salvage-price of unsold items is assumed to be zero. This happens when the excess inventory cannot be carried to the next selling season (either because it is too costly or because the product expires). Considering these assumptions is a common practice in this field. Almost all the papers

that study either Newsvendor problem or supply chain contracts through experimental approaches use these assumptions to keep the problem parameters in their simplest form (uniform demand distribution, constant selling price and production cost, and zero salvage-price). Considering the problem setup in its simplest form lets the researcher focus on decision maker's basic behavioral errors. To be consistent with the earlier studies in this field, we use these assumptions too. Considering a uniform demand distribution, equation (1) can be rewritten as

$$q^* = A + (B - A)\frac{c_u}{c_u + c_o}. \tag{2}$$

Before presenting the optimal forms of the wholesale price, revenue sharing and buyback contracts, we want to identify the order quantity that maximizes the supply chain profit as a whole. This would be the order quantity chosen by a centralized decision maker who controls both the retailer and the supplier. For such a decision maker, the overage and underage inventory costs would be $c_u = p - c$ and $c_0 = c$. Replacing these values in (2), we can calculate the optimal order quantity for a centralized decision maker as

$$q_c = A + (B - A)\frac{p - c}{p}. \tag{3}$$

This order quantity results in the maximum supply chain expected profit. Therefore, any contract that results in q_c coordinates the supply chain. For a coordinated supply chain (or equivalently for a supply chain with a centralized decision maker) the expected sales volume is (Cachon, 2003)

$$ES(q_c) = q_c - \frac{B - A}{2}\left(\frac{p - c}{p}\right)^2. \tag{4}$$

The corresponding supplier's expected profit can be found from

$$E\pi_S(q_c) = (1 - \lambda)\left[(p - c)q_c - \frac{B - A}{2}p\left(\frac{p - c}{p}\right)^2\right]. \tag{5}$$

As we mentioned earlier, the transaction between the supplier and the retailer is defined by a contract which is offered by the supplier. In a wholesale price contract, the only payment between the two parties is the wholesale price, w_{ws}, which should be paid to the supplier for each

unit of product ordered by the retailer. The value of w_{ws} that maximizes the supplier's expected profit can be calculated from the following equation (Katok & Wu, 2009).

$$w_{ws} = \begin{cases} p & if \quad 2p > c + pB/(B-A) \\ \dfrac{c}{2} + \dfrac{p}{2}\dfrac{B}{B-A} & otherwise \end{cases}$$

It is easy to verify that $c < w_{ws} \le p$. The underage and overage inventory costs for this type of contract is $c_u = p - w_{ws}$ and $c_0 = w_{ws}$. Replacing these values in (2), we can then calculate the retailer's optimal order quantity in a wholesale price contract as

$$q_{ws} = A + (B-A)\frac{p - w_{ws}}{p}.$$

Since $c < w_{ws}$, it is easy to verify that $q_{ws} < q_c$. Therefore, a wholesale price contract cannot coordinate the supply chain, which means the supply chain expected profit under this contract is lower than the maximum achievable expected profit for the supply chain.

One of the contracts that coordinate the supply chain is the revenue sharing contract. To provide the retailer with the proper incentive to choose an optimal order quantity, the supplier offers a low wholesale price w_{rs} which is smaller than her production cost, *c*. In return, the retailer has to pay the supplier *r* for any unit that the retailer manages to sell. In this way, the supplier shares the risk of overstocking with the retailer. The optimal wholesale price and shared revenue that coordinate the supply chain are (Cachon & Lariviere, 2005)

$$w_{rs} = \lambda c \quad and \quad r = (1-\lambda)p, \tag{6}$$

where $0 \le \lambda \le 1$ is the percentage share of retailer from the total supply chain profit. We can then have a revenue sharing contract for each value of λ. In practice, the value of λ is determined by the relative power of the supplier and the retailer or retailer's alternative opportunities. The underage and overage inventory costs for this type of contract is $c_u = p - w_{rs} - r = \lambda(p-c)$ and $c_0 = w_{rs} = \lambda c$. Replacing these values in (2), we can then calculate the retailer's optimal order quantity in a revenue sharing contract as

$$q_{rs} = q_c = A + (B-A)\frac{p-c}{p}.$$

Since the order quantity under a revenue sharing contract is equal to q_c, the supply chain can achieve its maximum expected profit under this contract.

Another contract that can coordinate the supply chain is the buyback contract. In this contract, the supplier offers a wholesale price w_{bb} which is larger than the supplier's production cost, *c*, but smaller than the retailer's selling price, *p*. In addition, in order to share the risk of inventory overage, the supplier offers to buy back the retailer's unsold products at a price *b*. The optimal wholesale and buyback prices that coordinate the supply chain are (Cachon & Lariviere, 2005)

$$w_{bb} = (1-\lambda)p + \lambda c \quad and \quad b = (1-\lambda)p\,. \tag{7}$$

The underage and overage inventory costs for this type of contract is $c_u = p - w_{bb} = \lambda(p-c)$ and $c_0 = w_{bb} = \lambda c$. Replacing these values in (2), we can then calculate the retailer's optimal order quantity in a revenue sharing contract as

$$q_{bb} = q_c = A + (B-A)\frac{p-c}{p}\,.$$

Therefore, a buyback contract can also achieve the maximum expected profit for the supply chain.

3.2- Experimental Design

We use laboratory experiments to investigate the effectiveness of the different approaches that we propose to improve the performance of revenue sharing and buyback contracts. In all these experiments, we assume the supplier's production cost is $c = 4$, the retailer's selling price is $p = 20$, and the demand is uniformly distributed between 100 and 300 units, $D \sim U(100,300)$. This choice of *p* and *c* represent a high profit margin product, $(p-c)/p > 0.5$. Since the benefit of coordination is larger for high profit margin products (Katok & Wu, 2009), here we focus only on this type of products. Moreover, for a low profit margin product, subjects' more than optimal order quantities (as Schweitzer & Cachon, 2000 show) can in fact increase the supplier's profit. So, there is no incentive for the supplier to try to lower the order quantities to the supply chain optimal level. Using a cost structure that results in a high profit margin is consistent with all the papers that study coordinating contracts in a supply chain through laboratory experiments (Ho & Zhang, 2008, Katok & Wu, 2009, and Davis, 2010).

In our experiments the subjects responded to the contracts offered by a (computerized) supplier. We assume the supplier is rational and risk-neutral. As a result, supplier's decisions are always consistent with the theoretical optimal solutions. By letting human subjects (retailers)

interact with computerized suppliers, we can avoid fairness concerns (Katok & Wu 2009). This will let us focus on subjects' behavioral error. For the revenue sharing and buyback contracts we always set the contract parameters such that the retailer's theoretical share of total profit is $\lambda = 1/4$. This value of λ let both parties benefit from the coordination. In the absence of fairness concerns (responding to computerized suppliers), the choice of λ should not change subjects' decision patterns. Katok & Wu (2009), for instance, arbitrarily choose $\lambda = 1/3$.

All of our subjects were College of Management students at the University of Massachusetts, Boston. We conducted the experiments in different management classes. The instructors of selected courses let us run the experiments in their classes as a required class activity. We conducted the experiments in a mix of graduate and undergraduate classes in four semesters during academic years of 2010-11 and 2011-12. To ensure that the results from undergraduate and graduate classes were comparable, we conducted the experiment on simple revenue sharing contract in a graduate and in an undergraduate class. The results were statistically equivalent. Katok & Wu (2009) observe the same results about the equivalence of the responses from undergraduate and graduate students in their experiments.

To incentivize students, we presented each experiment as a contest through which the students can find out how good they were at making decisions under an uncertain environment (random demand). In addition, we offered cash prizes ($40, $30, and $20) to the first three students with the best total performance. Subjects played the role of retailer's purchasing manager who decided about the order quantities for different selling seasons (rounds). Therefore, each subject's performance was measured by the retailer's total profit after 50 rounds of decision making.

At the beginning of each experiment session, the supply chain setup was explained to subjects using a PowerPoint presentation. The presentation, which usually took around 20 minutes, included simple numerical examples and how the subjects can interact with the software. A summary of the numerical values of the experiment parameters was visible on top of the screen at all times during the experiment. After a subject chose the order quantity for each selling season (round), the demand realized (a draw from a uniformly distributed random variable). Then, this demand along with the profit for the selling season, the accumulated profit so far, the cost of overstocking, and the cost of under-stocking for that round were shown to the subject. Two graphs on the screen showed the history of decisions made (order quantities)

accompanied by the realized demand and the history of profits in previous rounds. A screenshot of the user interface can be found in appendix B.

As a validation step and to ensure that our experimental setup is consistent with the existing results in the literature, we conducted a series of experiments to see if we can observe the pull to center phenomena[2] which has been reported repeatedly in the literature. We conducted an experiment with a wholesale price contract which is comparable to a simple Newsvendor problem. We also conducted one experiment with a simple revenue sharing contract and one experiment with a simple buyback contract. The details of these experiments and their results are presented in subsection 3.3. We consistently observed the pull to center phenomena in all these experiments similar to what is reported by Schweitzer & Cachon (2000), Bostian et al (2008), Kremer et al (2010), Bolton & Katok (2008), Lurie & Swaminathan (2009) , Katok & Wu (2009), Becker-Peth et al (2011). We will also use the results of these experiments as benchmarks for the results of other experiments which investigate the effectiveness of our proposed approaches.

Bolton et al (2012), through an experimental study, compare the decisions made by students with those made by experienced managers when they play the role of Newsvendor decision makers. The authors show managers exhibit ordering behavior similar to students, including biased ordering towards average demand. A similar result is reported by Katok et al (2008). These observations suggest experimental papers that rely on responses from students can provide useful insights about what we can expect from managers in practice.

3.3- Benchmark Experiments

We first conducted three experiments on wholesale price, simple revenue sharing, and simple buyback contracts[3]. The contract parameters and the sample sizes of these experiments are presented in table 1. Contract parameters are calculated based on the theoretical results in subsection 3.1 and the cost structure mentioned in subsection 3.2, with $\lambda = 1/4$ for the coordinating contracts.

Our first three hypotheses verify the existing results in the literature under our experimental setup. They also work as benchmarks for our proceeding results. Hypothesis 1 verifies whether

[2] We decided to verify the presence of pull to center phenomena in our experiments since it is the basic behavior which defines how the retailer's decisions deviate from the optimal order quantity.

[3] Since we conducted different variations of revenue sharing and buyback contracts, we call the traditional versions of these contracts simple revenue sharing and simple buyback contracts.

the retailer's decision deviates from the optimal theoretical value in a wholesale price contract. In analyzing the hypotheses, throughout this research, we use Wilcoxon rank sum test (Levine et al 2011, pp. 447-451). The unit of our analysis is the average order quantity of each subject, except when we want to investigate subjects' learning pattern for which we use the average of subjects' decisions in each round (see section 8).

Hypothesis 1. The average order quantity placed by the retailer in the wholesale price contract will be 125.

Contract Type	Sample Size	*w*	*r*	*b*	Optimal Order Quantity
Wholesale price	14	17.00	--	--	125
Simple Revenue sharing	14	1.00	15.00	--	260
Simple Buyback	20	16.00	--	15.00	260

Table 1 – Parameters of benchmark experiments

Subject's average order quantity in the wholesale price experiment is 178.8, which is considerably higher than the optimal value. Hence, we can strongly reject hypothesis 1 ($p < 0.001$). In hypothesis 2 we want to see if the coordinating contracts can improve the performance of the supply chain by increasing the retailer's order quantity in comparison to that of the wholesale price contract.

Hypothesis 2. The average order quantity of revenue sharing and buyback contracts will be higher than the average order quantity of wholesale price contract.

Subjects' average order quantities in simple revenue sharing and simple buyback experiments are 228.9 and 225.6, respectively. The differences between these order quantities and the average order quantity of the wholesale price contract are statistically significant. Therefore, the experiment results support hypothesis 2. Hypothesis 3 looks at the performance of the revenue sharing and buyback contracts. This hypothesis checks whether these contracts are able to coordinate the supply chain as it is promised by the standard theory.

Hypothesis 3. The average order quantities placed by the retailer in both simple revenue sharing and simple buyback contracts will be 260.

The average order quantities under simple revenue sharing and buyback contracts are considerably smaller than the optimal value of 260. We can, therefore, strongly reject hypothesis 3 $(p < 0.001)$. These results confirm the known pull to center behavior.

4. Approach 1: Combined Contract

4.1. Theoretical Results

Hypothesis 2 shows that revenue sharing and buyback contracts can indeed improve the performance of the supply chain, even if the improvement is not as much as the theory predicts. The idea behind this approach comes from the observation that each of these two contracts can individually improve the performance of the supply chain to some extent. So, an interesting question could be whether a combination of these two contracts could improve the performance even further? A combined revenue sharing and buyback contract is a contract in which the supplier offers a relatively low wholesale price and in return asks the retailer to share part of the revenue of the sold items. In addition, the supplier promises to buy back the unsold items at a price lower than the wholesale price.

Cachon & Lariviere (2005) show buyback and revenue sharing contracts are theoretically equivalent. That is, they result in the same profits for the retailer and the supplier for any realization of the random demand. Although the literature reports the theoretical equivalence of the two contracts, they are always treated as two distinct contracts. Here, we show that these two contracts are the two ends of a spectrum of combined contracts as we defined above.

Let w_{com}, r_{com}, and b_{com} be the wholesale price, shared revenue, and buyback price of a combined contract, respectively. As before, we denote the percentage share of the retailer from the supply chain profit with λ. It is not very difficult to verify that the overage and underage inventory costs of this contract are $c_o = w_{com} - b_{com}$ and $c_u = p - w_{com} - r_{com}$. Replacing these quantities in (2), we can calculate the optimal order quantity of the combined contract as

$$q_{com} = A + (B - A)\left(\frac{p - w_{com} - r_{com}}{p - b_{com} - r_{com}}\right). \tag{8}$$

As we can see, different combinations of contract parameters (w_{com}, r_{com}, b_{com}) result in different values for q_{com}. However, for any chosen wholesale price, w_{com}, if the supplier sets the values of the shared revenue and buyback prices as

$$r_{com} = (1-\lambda)p + \lambda c - w_{com} \quad \text{and} \quad b_{com} = w_{com} - \lambda c, \tag{9}$$

then equation (8) simplifies to

$$q_{com} = q_c = A + (B-A)\left(\frac{p-c}{p}\right). \tag{10}$$

This means, for any chosen wholesale price, w_{com}, the choice of shared revenue and buyback prices as identified in (9) coordinates the supply chain. From (9), it is evident that the chosen wholesale price cannot be larger than $(1-\lambda)p + \lambda c$ or smaller than λc. That is, $w_{com} \in [\lambda c, \lambda c + (1-\lambda)p]$.

Note that the combined contract turns into a pure revenue sharing contract if we choose the lowest range of wholesale prices, $w_{com} = \lambda c$. Similarly, the combined contract turns into a pure buyback contract, if we choose the highest range of wholesale prices, $w_{com} = \lambda c + (1-\lambda)p$.

4.2. Experimental Results

The idea behind this approach is that a contract which has the appealing features of both revenue sharing and buyback contracts might inspire more confidence in subjects and encourage them to place higher order quantities. This could be in spite of the theoretical results, which predicts the same order quantities for the family of combined contracts with the same value of λ. The following hypothesis is to verify this conjecture.

Hypothesis 4. All the contracts in a family of combined contracts with the same value of λ results in the same average order quantity.

To examine this hypothesis, we compare four contracts in a family of combined contracts with $\lambda = 1/4$. We consider the two simple revenue sharing and buyback contracts from subsection 3.3 as the two ends of the spectrum. We also consider the results of another two experiments with new combined contracts. One combined contract has $w_{com} = 4.75$, which is closer to the revenue sharing end of the spectrum, and another with $w_{com} = 12.25$, which is closer to the buyback end of the spectrum. The experiment parameters along with the observed average order quantities are shown in table 2. We compare the average order quantities of these four contracts pairwise. We could not find any significant difference between the average order quantities ($p > 0.05$ in all cases). Therefore, we do not have enough evidence to reject hypothesis 4.

Contract Type	Sample Size	*w*	*r*	*b*	Observed Ave. Order Quantity
Simple Revenue sharing	14	1.00	15.00	--	228.9
Combined 1	18	4.75	11.25	3.75	227.5
Combined 2	18	12.25	3.75	11.25	226.1
Simple Buyback	20	16.00	--	15.00	225.6

Table 2 – Parameters and results of combined contract experiments

Katok & Wu (2009) observe differences between simple buyback and simple revenue sharing average order quantities in their experiments. They observe that depending on the demand range and the way it is framed for the subjects, the buyback contract can result in higher or lower average order quantities. They attribute this observation to subjects' loss aversion behavior. Their results, however, show that the difference between the average order quantities of the two contract types decreases and disappears with subjects' experience. This means that in general subjects do not react to different forms of combined contracts. Therefore, using a combined contract is not an effective approach to induce suppliers to place higher order quantities.

5. Approach 2: Risk Averse Contract

5.1. Theoretical Results

The results of our benchmark experiments show that retailers tend to order less than the optimal order quantity in both revenue sharing and buyback contracts. Katok & Wu (2009) observe similar results. One possible explanation for retailer's less-than-optimal order quantity could be retailer's risk-averse behavior. It has been long argued that decision makers in the business world tend to be risk-averse. Eeckhoudt et al (1995) show that the optimal order quantity decreases with an increase in risk-aversion in a Newsvendor problem. If risk-aversion is the reason behind the less-than-optimal order quantity, then the supplier should be able to rectify this problem by designing contracts not for a risk-neutral but for a risk-averse retailer.

To model retailer's risk-aversion, we consider an exponential utility function for the retailer, i.e. $u(x) = -e^{-\phi x}$, where ϕ is the constant risk-aversion coefficient. The following proposition characterizes the optimal parameters of a revenue sharing contract for a risk-averse retailer.

Proposition 1. *A coordinating revenue sharing contract for a retailer with a utility function of* $u(x) = -e^{-\phi x}$ *is characterized by the wholesale price* $\hat{w}_{rs}$ *and shared revenue* $\hat{r}$*, where the shared revenue is the unique solution to*

$$\frac{\phi(p-\hat{r})(B-q_c)(p-\hat{r})}{e^{\phi(p-\hat{r})(q_c-A)}-1+\phi(p-\hat{r})(B-q_c)} = \frac{E\pi_S(q_c)-\hat{r}ES(q_c)+cq_c}{q_c}, \tag{11}$$

and the wholesale price can be calculated from

$$\hat{w}_{rs} = \frac{E\pi_S(q_c)-\hat{r}ES(q_c)+cq_c}{q_c}, \tag{12}$$

where $E\pi_S(q_c)$ *and* $ES(q_c)$ *can be found from (4) and (5).*

Moreover, for a retailer with a risk-aversion coefficient of ϕ*, the optimal order quantity,* q_ϕ*, under a revenue sharing contract with wholesale price w and shared revenue r is the unique solution to*

$$e^{\phi(p-r)(q_\phi-A)} = 1+\frac{p-r}{w}\phi(p-r-w)(B-q_\phi), \tag{13}$$

The proof of this and other propositions can be found in Appendix A.

5.2. Experimental Results

To design a revenue sharing contract for a risk-averse retailer, we first need to find the value of risk-aversion coefficient ϕ. We can estimate the value of ϕ from the results of our benchmark experiments. In other words, we can numerically solve equation (13) for ϕ with $q_\phi = 228.9$, $w = 1$, and $r = 15$. The result will be $\phi = 0.0022$. This value shows the level of risk-aversion that theoretically results in an order quantity of 228.9 (our observed average order quantity under a simple revenue sharing contract).

Knowing the risk-aversion coefficient, we can then use (11) and (12) to design a revenue sharing contract for our risk-averse subjects. We conducted an experiment with such a revenue sharing contract. The parameters of this contract and the subjects' average order quantity are shown in table 3. We can see that for a risk-averse retailer, the supplier should lower the wholesale price and in turn increase the shared revenue. Hypothesis 5 verifies the performance of a revenue sharing contract that is designed for a risk-averse retailer.

Hypothesis 5. The average order quantity of a revenue sharing contract that is designed for a risk-averse retailer ($\phi = 0.0022$) will be 260.

Contract Type	Sample Size	w	r	Observed Ave. Order Quantity
Simple Revenue Sharing	14	1	15	228.9
Risk Averse Revenue sharing	17	0.42	15.77	228.0

Table 3 – Parameters and results of risk-averse revenue sharing experiments

The result of our experiment does not show any improvement in the average order quantities and there is still a large gap between the observed average order quantity and the optimal value. We can then strongly reject hypothesis 5 ($p < 0.001$). This means risk-aversion does not play an influential role in the subjects' behavior in a revenue sharing contract and hence we cannot use it to improve the contract performance. Since we did not find any influential impact of risk-aversion, we did not repeat a similar experiment for a buyback contract.

Schweitzer & Cachon (2000) look at the possible impact of risk-aversion on subjects' behavior too. In their experiments they examine a simple Newsvendor problem (not contracts between suppliers and retailers). Their observation, however, is consistent with ours. They conclude risk-aversion cannot play an influential role in Newsvendor's pull to center behavior. They use the contrast between the subjects' behavior when they face high profit margin and low profit margin products to draw this conclusion.

6. Approach 3: Offering Free Items

6.1. Theoretic Results

To provide the retailer with more incentive to increase the order quantity, we consider forms of revenue sharing and buyback contracts in which the supplier offers one free unit of product to the retailer for any N products ordered. Offering free items could have two impacts on the performance of the supply chain. First, it increases the effective order quantity, which we define as the sum of actual order quantity and free items, $q_{eff} = q(1+1/N)$. Second, the lure of receiving free items (Shampanier et al, 2007) might encourage the retailer to increase his order quantity. These increases mean the total number of items in the supply chain can get closer to the optimal level.

Once a free item is delivered to the retailer, it is treated similar to a regular paid item. That is, in a revenue sharing contract, the retailer has to share the revenue of all sold items (free and paid). Similarly, in a buyback contract, the supplier buys back all unsold items (free and paid). This condition is required to keep the percentage share of the supplier in the contracts with free items similar to her percentage share in simple coordinating contracts (see the proof of proposition 2). The following proposition characterizes the coordinating contracts with free items.

Proposition 2. *When the supplier offers one free item for every N items ordered by the retailer, coordinating contracts are characterized by*

(a) revenue sharing contract: $w_{rs}^{F} = (1+1/N)\lambda c$ *and* $r^{F} = (1-\lambda)p$,

(b) buyback contract: $w_{bb}^{F} = (1+1/N)[(1-\lambda)p+\lambda c]$ *and* $b^{F} = (1-\lambda)p$.

Moreover, in these two contracts, the percentage share of the supplier from the total channel profit is $(1-\lambda)$.

Proposition 2 states that in a coordinating contract with free items, the supplier maintains the same shared revenue and buyback price as in the case of a simple coordinating contract. However, to make up for the cost of free items, the supplier has to increase the wholesale price. Note, from a theoretical point of view, simple coordinating contracts and the corresponding contracts with free items perform equivalently. That is, the latter results in the same expected profit for the supplier and the retailer as the former. However, for a retailer who is prone to behavioral error and misjudgment, the two contracts might perform differently.

6.2. Experimental Results

Hypotheses 6A and 6B verify the effectiveness of revenue sharing and buyback contracts with free offerings.

Hypothesis 6A. The average effective order quantity of a revenue sharing contract with free offering (N=7) will be 260.

Hypothesis 6B. The average effective order quantity of a buyback contract with free offering (N=6) will be 260.

We chose the values of N based on the subjects' average order quantities in simple revenue sharing and buyback contracts (benchmark experiments) in a way that the number of free items

brings the effective order quantity to a level close to the optimal value (260). This means *N*=7 for the revenue sharing contract and *N*=6 for the buyback contract. We conducted two experiments for revenue sharing and buyback contracts with free items. In these experiments the monetary value of free items received by the retailer is presented to the subjects after they entered their order quantities in each round. In our instructions we emphasized that this monetary value turns into profit only when the retailer manages to sell them. The parameters and results of these experiments are shown in table 4.

In the revenue sharing experiment, the actual average order quantity (221.3) remains almost the same as the order quantity in the simple revenue sharing experiment (228.9). There is no significant difference between the two values ($p > 0.05$). As a result, the effective order quantity, $q_{eff} = 252.7$, becomes very close to the optimal value (no significance difference, $p > 0.05$). Therefore, we cannot reject hypothesis 6A. In other words, offering free items in a revenue sharing contract can bring the effective order quantity to a coordinating level.

Contract Type	Sample Size	*w*	*r*	*b*	*N*	Ave. Order Quantity Actual	Ave. Order Quantity Effective*
Revenue Sharing							
• Simple	14	1.00	15.00	--	--	228.9	228.9
• With Free Items	17	1.14	15.00	--	7	221.3	252.9
• With Adjusted Free Items	17	Initially 1.14 then adjusted	15.00	--	Initially 7 then adjusted	221.4	258.9
Buyback							
• Simple	20	16.00	--	15.00	--	225.6	225.6
• With Free Items	19	16.79	--	15.00	6	202.1	235.8
• With Adjusted Free Items	19	Initially 16.79 then adjusted	--	15.00	Initially 6 then adjusted	209.1	249.6

* Effective order quantity = Actual order quantity + Free items

Table 4 – Parameters and results of contracts with free items

It is interesting to note that the actual average order quantity is statistically equivalent to that of a simple revenue sharing contract. It seems that, in this experiment, the tendency to order more because of free items is cancelled out by the tendency to order less because of a slightly higher wholesale price. As a result, the extra free items can coordinate the supply chain. In the case of buyback contract with free items, the subjects' actual average order quantity (202.1) is significantly ($p < 0.01$) less than the average orders in a simple buyback contract (225.6), which

in turn means that the effective average order quantity, $q_{eff} = 235.8$, is significantly less than the optimal value (260). This means that the offering of free items in the buyback contract does not work similar to the revenue sharing contract. This could be due to the fact that in the buyback contract the amount of increase in the wholesale price is higher than the similar increase in a revenue sharing contract (considering the same value of *N*). This is because, in the revenue sharing contract, part of the revenue of the sold free items returns to the supplier. Therefore, the supplier needs to increase the wholesale price only by a small amount. This is not the case in the buyback contract. Therefore, the supplier has to increase her price by a larger amount to keep her expected profit similar to a simple buyback contract. As a result, the tendency to increase the order quantity to receive more free items cannot balance the tendency to reduce the order quantity because of higher wholesale price. Hence, offering free items cannot coordinate the supply chain and we can strongly reject hypothesis 6B ($p < 0.001$).

To improve the performance of this form of contract, we can customize the contract terms for each individual subject. We will see that this approach can significantly improve the performance of both contracts. In this approach (adjusted free), we try to tailor the value of *N* for each subject. That is, instead of choosing one value of *N* for all subjects, we use a separate value of *N* for each subject based on the subject's history of orders. Through the following two hypotheses we investigate the performance of revenue sharing and buyback contracts with adjusted free items.

Hypothesis 7A. The average order quantity of a revenue sharing contract with adjusted free offering will be 260.

Hypothesis 7B. The average order quantity of a buyback contract with adjusted free offering will be 260.

To verify these two hypotheses we conducted two experiments (a revenue sharing and a buyback contract) with free items in which the value of *N* for each subject is adjusted after the first 25 rounds based on the average of each subject's order quantities in in the first 25 rounds. The wholesale prices were also changed according to the new values of *N*. The new values of *N* and wholesale price after the first 25 rounds were highlighted on the screen to attract subjects' attention.

The results of the experiments are quite interesting. In the revenue sharing contract, the value of N adjusts from 7 to an average of 11.2 after the 25th round. Note that the higher the value of N the fewer the number of free items in the supply chain. The resulting effective average order quantity of the second 25 rounds increases to $q_{eff} = 257.8$. For the buyback contract, the values of N adjusts from 6 to an average of 8.3 after the 25th round. The effective average order quantity of the second 25 rounds increases to $q_{eff} = 249.6$. In both contracts, there is no significant difference between the effective order quantities and the optimal values. Therefore, we cannot reject hypotheses 7A and 7B ($p > 0.05$). The adjusted values of N suggest (again) that we need fewer numbers of free items in the revenue sharing contract than what we need in the buyback contract.

Note that the improved performance in the second half of the experiment with the adjusted number of free items cannot be associated with a learning process in the subjects' ordering behavior. This is due to the fact that the actual order quantities in the adjusted free experiments remain almost the same as the actual order quantities of the experiments with the fixed number of free items. Therefore, the improvement can only be due to the adjusted number of free items. We will have more discussion about subjects' learning in section 8.

7. Approach 4: Showing the Demand Pattern

One of the reasons behind subjects' suboptimal decisions is argued to be subjects' focus on the most recent demand which could in turn lead to a demand chasing pattern (Schweitzer and Cachon 2000, Bostian et al, 2008, and Lurie & Swaminathan, 2009). Subjects' focus on the most recent demands could be due to their inability to comprehend the true nature of demand uncertainty. In our benchmark experiments we informed the subjects about the demand distribution. This is the case in almost all other similar research papers. This information, however, might not effectively be involved in the subjects' decision making process.

The idea behind our fourth approach is to provide visual information about the demand pattern to help the subjects to better understand the random nature of the demand and discourage them to chase the demand. The following hypothesis investigates the impact of providing visual information about the demand pattern.

Hypothesis 8. The average order quantity of a revenue sharing contract with additional (visual) information about the demand pattern will be 260.

In an experiment with the revenue sharing contract, we added a graph on the software user-interface. The graph showed a sample history of demand in 50 consecutive selling seasons (rounds). Figure 1 shows such a graph. This graph was visible throughout the experiment. All other conditions were the same as our benchmark experiment.

The resulting average order quantity (228.3) is almost the same as the average order quantity in the simple revenue sharing contract. We can, therefore, strongly reject hypothesis 8 $(p < 0.001)$. This means, trying to create a better understanding about the demand behavior, through visualizing the demand pattern as shown in figure 1, cannot improve the quality of retailer's decisions. Since the pull to center phenomena still prevails, there can be only two explanations for this observation. Either subjects' lack of comprehension of the demand behavior is not the main source of subjects' suboptimal decisions, or being exposed to the visual demand pattern, as shown in figure 1, is not enough to create a better comprehension of demand behavior. Either way, the results of this experiment suggest that the subjects' attentions still remain on the last realized demand and the inventory error that it creates. This means showing the demand pattern in not an effective approach to improve subjects' decisions.

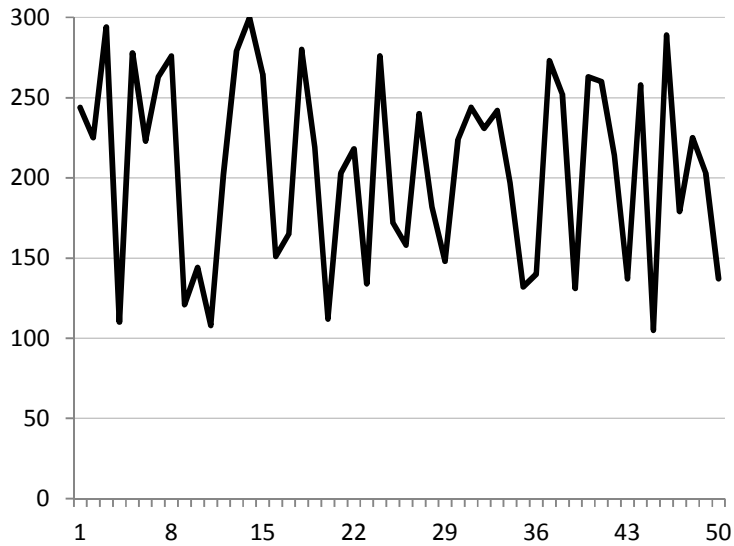

Figure 1 – Visualization of demand pattern (uniform distribution)

8. Approach 5: Providing Collective Feedback

Both Bolton & Katok (2008) and Lurie & Swaminathan (2009) show that the performance of a Newsvendor can be improved by restricting a decision to stand for a set of rounds. That is, when the retailer makes a decision in a selling season (round), then the same decision is applied to a set of successive selling seasons. The result of this decision (feedback) is revealed to the retailer only after all these selling seasons are over.

Standing order reduces the frequency of orders. Therefore, the subjects know that each order quantity decision impacts more than one round. This could encourage them to look at the random demand in a more collective way, which in turn might reduce their tendency for demand chasing. On the other hand, standing orders reduce the feedback frequency too. Hence, each feedback contains the collective impacts of an order quantity on the profit of more than one realized demands. This collective measure, in a sense, reduces the randomness in demand and show a more accurate value of each order quantity. The experiments by Lurie & Swaminathan (2009) suggest the improvement in a standing order setup is mainly due to a reduction in feedback frequency (not due to a reduction in order frequency).

Although standing orders can result in average order quantities that are closer to the optimal value, it has the practical limitation of preventing the retailer to place an order for each selling season. It has also the limitation of preventing the retailer to access the result of a decision at the end of each season. Hence, applying standing orders might not be practical in many business situations.

To address these restrictions, we propose a new approach. In this approach, after a subject makes an order quantity decision, he is presented with the total profit that would be earned if the chosen order quantity were chosen for all previous rounds. Therefore, this *would-be* total profit is similar to an imaginary total profit of a standing order from the beginning of the experiment (using the current order quantity). For example, if a subject chooses an order quantity of 230 in the 20th round, then the *would-be* total profit will show the total profit if 230 were the chosen order quantity for all the first 20 rounds. As a result, this *would-be* total profit provides a feedback in every round which is very much similar to the feedback provided in a regular standing order. In the new approach, however, the retailer does not face the limitations of a standing order. That is, the retailer can make a decision for every round and access the feedback at the end of each round.

The value of this *would-be* total profit is negligible in the starting rounds of the experiment. However, as the number of rounds increases, the value of information provided by this number also increases. In other words, in the higher rounds, this *would-be* total profit (and its comparison with the actual total profit) is a good measure for the real value of the selected order quantity. Since the decision maker usually focuses on the feedbacks of the latest round, this piece of information should be under retailer's attention range. As a result, we can expect a learning pattern in the retailer's decision process and observe better decisions toward the final rounds. Through the following two hypotheses we try to verify this conjecture.

Hypothesis 9A. The average order quantity of a revenue sharing contract with *would-be* total profit feedback will be 260.

Hypothesis 9B. The average order quantity of a buyback contract with *would-be* total profit feedback will be 260.

We conducted two experiments for the revenue sharing and buyback contracts while providing the *would-be* total profit feedback. To make sure that the subjects fully understood this new piece of information, we asked them to write a sentence or two about the meaning of the *would-be* total profit before they started making their order quantity decisions. All other experimental conditions were the same as our benchmark experiments.

Before we examine hypothesis 9A, we use linear regression to verify if there is an increasing trend in subjects' average order quantities. We find a significant increasing trend (0.58 units per round, $p < 0.001$) in the revenue sharing experiment. Figure 2 shows the average order quantities across the 50 rounds of our experiment. Such an increasing trend does not exist in the simple revenue sharing (benchmark) experiment. This lack of considerable learning trend in simple contracts is consistent with the prior research papers. Schweitzer & Cachon (2000) do not observe a learning trend in their Newsvendor experiment with 15 rounds of decisions. Bolton & Katok (2008) observe a learning trend in the Newsvendors' decisions in their extended experiment with 100 rounds of decisions. However, they report a very low rate of increase in the average order quantities (0.13 units per round).

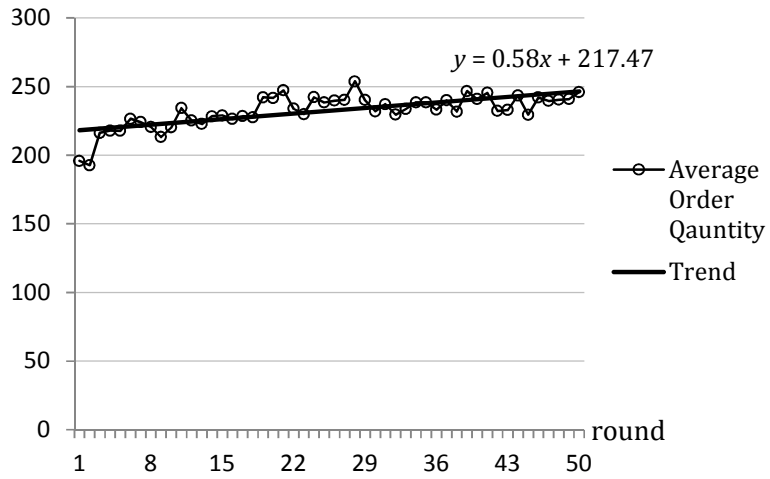

Figure 2 – Average order quantities with *would-be* total profit feedback (revenue sharing)

Considering the learning trend in the revenue sharing experiment with the *would-be* total profit feedback, we verify hypothesis 9A for the last 10 rounds of the experiment to observe the impact of the learning process. The parameters and results of our experiments are presented in table 5. It is evident that the average order quantity (last 10 rounds) of the experiment with the *would-be* total profit feedback (239.1) has increased significantly ($p < 0.05$) compared to the corresponding value for a simple revenue sharing experiment (214.8). Although average of the last 10 rounds is still short of the optimal value (260), the difference is not statistically significant ($p > 0.05$). So, we do not have enough evidence to reject hypothesis 9A. This learning trend suggests that, in a revenue sharing contract, the retailer could eventually choose order quantities very close to the optimal value when he is provided with this type of feedback.

The result of the buyback experiment with the *would-be* total profit feedback does not show any learning pattern. The regression analysis shows no significant slope in the average order quantities placed by the subjects through the 50 rounds of the experiment. The comparison of the last 10 rounds of this experiment with the corresponding value for a simple buyback contract does not show any improvement either (see table 5). We can, therefore, strongly reject hypothesis 9B ($p < 0.001$).

It is interesting to observe that the same type of feedback results in different outcomes in the revenue sharing and buyback contracts. In this approach, the subjects respond more positively to the collective feedback in the revenue sharing than what we can see in the buyback contract. One possible explanation could be the fact that in the buyback contract the higher wholesale price

means a higher initial payment (hence a higher prospect of loss). This prospect of loss could encourage the decision maker to focus more on the last demand value and not on the collective feedback. This behavior could be attributed to subjects' loss-aversion (Katok & Wu, 2009).

Contract Type	Sample Size	*w*	*r*	*b*	Ave. Order Quantity (last 10 rounds)
Revenue Sharing					
• Simple	14	1.00	15.00	--	214.8
• Would-be total profit	18	1.00	15.00	--	239.1
Buyback					
• Simple	20	16.00	--	15.00	221.1
• Would-be total profit	15	16.00	--	15.00	225.3

Table 5 – Parameters and results of contracts with *would-be* total profit feedback

9. Concluding Remarks

In this research, we examine a two-echelon supply chain consisting of a rational supplier and a retailer who is prone to behavioral errors. We show (like others before us), when the supplier offers a coordinating contract (either revenue sharing or buyback), the retailer systematically fails to place an order with optimal quantity. This sub-optimal behavior, in turn, results in less-than-optimal profits for all parties and the supply chain as a whole.

We contribute to the existing literature by proposing five approaches to improve the decisions made by the retailer. We verify the effectiveness of each approach through laboratory experiments. The first three approaches concern the contract terms offered by the supplier. These approaches are (1) combined contracts, (2) contracts designed for risk-averse retailers, and (3) contracts with free-item offering. Among these approaches, we show, only the contracts with free-item offering can actually bring the order quantities close to the optimal level and coordinate the supply chain. The next two approaches concern extra information and feedback for the retailer. These are (4) providing the visual pattern of demand randomness and (5) providing a collective feedback on each decision. We show that the collective feedback (*would-be* total profit) can create a stronger learning process in the revenue sharing contract, which means decision makers can learn from their prior decisions and eventually place close-to-optimal order quantities. This approach is not effective in the buyback contract.

A general takeaway from this research is that it is possible to improve the results of decisions made by the retailer either through a contract mechanism or through carefully designed feedback. It is interesting to note that the two effective approaches that we find in this research improve the performance of the contract in two very different ways. The free-item approach (or its adjusted counterpart) does not improve the retailer's decisions. Instead, it adds a proper number of items to the items ordered by the retailer. Therefore, it increases the total number of items in the supply chain. The required change in the wholesale price is so small that it does not change the decision maker's ordering behavior. Hence, the resulting total number of items (effective order quantity) increases to a number very close to the optimal order quantity.

On the other hand, the collective feedback approach improves the retailer's decisions by weakening the demand chasing behavior. One of the reasons behind the retailer's suboptimal decisions is argued to be the decision maker's limited attention span. Having a limited attention span, the decision maker mostly focuses on the feedback from the latest decision which is either a shortage or excess of inventory. This shortsightedness results in the demand chasing behavior. Providing a collective feedback (the potential impact of a decision on all previous selling seasons) can help the decision maker to overcome the tendency to chase the random demand. The collective feedback, in fact, shows a more realistic value of each decision in each selling season.

It is also interesting to note that these two approaches are less effective on Buyback contracts. The reason behind this behavior can be attributed to the decision makers' loss aversion behavior (Katok & Wu, 2009). In a buyback contract the wholesale price is higher than the wholesale price in a revenue sharing contract. In the free-item approach, higher initial wholesale price means that the price increase due to offering free items is more noticeable by the decision maker. Therefore, the retailer tends to order fewer items to reduce the risk of loss due to inventory overage. Similarly, in the collective feedback approach, when the wholesale price is high, the risk of loss due to inventory overage attracts the decision maker's attention. This prevents the retailer from paying enough attention to the collective feedback. Therefore, the demand chasing behavior prevails.

This research also contributes to the supply chain contracting literature by introducing theoretical forms of three new contracts, which are extensions of revenue sharing and buyback contracts. These are (*a*) revenue sharing contracts for risk-averse retailers, (*b*) combined

contracts (combination of revenue sharing and buyback), and (*c*) revenue sharing and buyback contracts with free items. For each type of contract, we derive contract parameters that theoretically coordinate the supply chain.

Having the results of this research, it would be interesting to explore the possible approaches that can improve other forms of contracts (such as two-part tariffs and quantity discount) in practice. This could be a possible avenue for future research.

References

Arrow, K. J., Harris, T., Marschak, J., 1951. Optimal inventory policy. Econometrica 19 (3). 250-272.

Becker-Peth, M., Katok, E. , Thonemann, U.W., 2011. Designing Contracts for Irrational but Predictable Newsvendors. Working Paper.

Bolton, G., Katok, E., 2008. Learning by Doing in the Newsvendor Problem: A Laboratory Investigation of the Role of Experience and Feedback. Manufacturing and Service Operations Management 10 (3), 519-538.

Bolton, G., Ockenfels, A., Thonemann, U., 2012. Managers and Students as Newsvendors. Accepted for publication in Management Science.

Bostian, A.A., Holt, C.A., Smith, A. M., 2008. Newsvendor "Pull-to-Center" Effect: Adaptive Learning in a Laboratory Experiment. Manufacturing and Service Operations Management 10 (4), 590-608.

Cachon, G.P., 2003. Supply Chain Coordination with Contracts. In: Graves, S., Kok, T.d., (Eds). Handbooks in Operations Research and Management Science: Supply Chain Management, Vol. 11. Cambridge: North-Holland.

Cachon, G.P., Lariviere, M.A., 2005. Supply Chain Coordination with Revenue-Sharing Contracts: Strengths and Limitations. Management Science 51 (1), 30-44.

Corbett, C.J., Fransoo, J.C., 2007. Entrepreneurs and Newsvendors: Do small Businesses Follow the Newsvendor Logic When Making Inventory Decisions? Working Paper, University of California, Los Angeles.

Eeckhoudt, L., Gollier, C., Schlesinger, H., 1995. Risk averse (and prudent) newsboy. Management Science 41 (5) 786-794.

Davis, A.M., 2010. An Experimental Investigation of Pull Contracts. Working paper, The Pennsylvania State University.

Fisher, M., Raman, A., 1996. Reducing the Cost of Demand Uncertainty Through Accurate Response to Early Sales. Operations Research 44 (1) , 87-99.

Ho, T., Zhang, J, 2008. Designing Pricing Contracts for Boundedly Rational Customers: Does the Framing of the Fixed Fee Matter? Management Science 54 (4), 686-700.

Katok, E., Thomas, D., Davis, A., 2008. Inventory Service Level Agreements as Coordination Mechanisms: The effect of review periods. Manufacturing and Service Operations Management 10 (4), 609-624.

Katok, E., 2011. Using Laboratory Experiments to Build Better Operations Management Models. Foundations and Trends in Technology, Information and Operations Management 5 (1), 1-86.

Katok, E., Wu, D.Y., 2009. Contracting in Supply Chains: A Laboratory Investigation. Management Science 55 (12), 1953- 1968.

Keser, C., Paleologo, G.A., 2004. Experimental Investigation of Supplier-Retailer Contracts: The Wholesale Price Contract. CIRANO, Working Paper, IBM Research.

Kremer, M., Minner, S., Wassenhove, V.L., 2010. Do Random Errors Explain Newsvendor Behavior? Manufacturing and Services Operations Management 12 (4), 673-81.

Levine, D.M., Stephan, D.F., Krehbiel, T.C., Berenson, M.L., 2011. Statistics for Managers. 6th edition. Prentice Hall.

Luch, C.H., Wu, Y., 2008. Social Preferences and Supply Chain Performance: An Experimental Study. Management Science 54 (11), 1835-1849.

Lurie, N. H., Swaminathan, J. M., 2009. Is Timely Information Always Better? The Effect of Feedback Frequency on Decision Making. Organizational Behavior and Human Decision Processes 108 (2), 315-329.

Olivares, M., Terwiesch, C., Cassorla, L., 2008. Structural Estimation of the Newsvendor Model: An Application to Reserving Operating Room Time. Management Science 54 (1), 41-55.

Padmanabhan, V., Png, I.P.L., 1995. Returns policies: Make money by making good. Sloan Management Review 37(1), 65–72.

Qin, Y., Wang, R., Vakharia, A.J., Chen, Y., Seref M.H., 2011. The Newsvendor Problem: Review and Directions for Future Research, European Journal of Operational Research 213 (2), 361-374.

Schweitzer, M., Cachon, G., 2000. Decision Bias in the Newsvendor Problem: Experimental Evidence. Management Science 46 (3), 404-420.

Shampanier, K., Mazar, N., Ariely, D., 2007. Zero as a Special Price: The True Value of Free Products. Marketing Science 26 (6), 742-57.

Sliver, E. A., Pyke, D. F., Peterson, R., 1998. Inventory Management and Production Planning and Scheduling. 3rd edition. John Wiley and Sons.

Spengler, J., 1950. Vertical integration and antitrust policy. Journal of Political Economy 58 (4), 347-52.

Appendix A: Proofs

Proof of proposition 1: The order quantity which maximizes the expected utility of the retailer can be calculated as follows.

$$\pi(q)=\begin{cases}\pi_{-}(q)=(p-r)D-wq & if \quad D\le Q\\ \pi_{+}(q)=(p-r-w)q & if \quad D>Q\end{cases}$$

$$Eu\left(\pi(q)\right)=\int_A^q u\left(\pi_{-}(q)\right)f(x)dx+\int_q^B u\left(\pi_{+}(q)\right)f(x)dx$$

$$Eu\left(\pi(q)\right)=\frac{1}{B-A}\left(\int_A^q -e^{-\phi\left((p-r)x-wq\right)}dx-e^{-\phi\left((p-r-w)q\right)}(B-q)\right)$$

$$\frac{\partial Eu\left(\pi(q)\right)}{\partial q}=\frac{-1}{B-A}\left(\frac{w}{p-r}\left[e^{-\phi\left((p-r)A-wq\right)}-e^{-\phi(p-r-w)q}\right]-\phi(p-r-w)e^{-\phi\left((p-r-w)q\right)}(B-q)\right)$$

$$\frac{\partial Eu\left(\pi(q)\right)}{\partial q}=0 \;\Rightarrow\; \frac{w}{p-r}\left[e^{-\phi\left((p-r)A-wq_\phi\right)}-e^{-\phi(p-r-w)q_\phi}\right]=\phi(p-r-w)e^{-\phi\left((p-r-w)q_\phi\right)}(B-q_\phi)$$

$$e^{\phi(p-r)(q_\phi-A)}=1+\frac{p-r}{w}\phi(p-r-w)(B-q_\phi) \tag{A1}$$

Now, for a given risk aversion coefficient ϕ, we want to find a new revenue sharing contract $(\hat{w}_{rs},\hat{r})$ that can result in the same order quantity and supplier's expected profit that risk neutral retailer generates with (w,r), that is

$$q_c=A+(B-A)\frac{p-r-w}{p-r} \qquad \text{and} \qquad E\pi_S(q_c)=(w+r-c)q_c-r\frac{B-A}{2}\left(\frac{p-r-w}{p-r}\right)^2.$$

In designing $(\hat{w}_{rs},\hat{r})$, we know the value of q_c that we want to achieve therefore we can calculate the value of w from equation (A1) for given values of ϕ and $\hat{r}$.

$$\hat{w}_{rs}=\frac{\phi(p-\hat{r})(B-q_c)(p-\hat{r})}{e^{\phi(p-\hat{r})(q_c-A)}-1+\phi(p-\hat{r})(B-q_c)} \tag{A2}$$

On the other hand, supplier's expected profit under the contract $(\hat{w}_{rs},\hat{r})$ can be written as:

$$E\pi_S(q_c)=(\hat{w}_{rs}-c)q_c+\hat{r}ES(q_c), \text{ where } ES(q_c)=q_c-\frac{B-A}{2}\left(\frac{p-r-w}{p-r}\right)^2$$

Therefore, $\hat{w}_{rs}$ can be calculated as follows:

$$\hat{w}_{rs} = \frac{E\pi_S(q_c) - \hat{r}ES(q_c) + cq_c}{q_c} \tag{A3}$$

By equating equations (A2) and (A3) we can calculate $\hat{r}$. Then we can replace this value to either (A2) or (A3) to calculate $\hat{w}_{rs}$.

Proof of Proposition 2:

Let $\bar{q} = (1+1/N)q$. Then, the retailer's profit for a revenue sharing contract can be written as:

$$\pi_R = \begin{cases} (p-r)\bar{q} - wq & if \quad D \geq \bar{q} \\ (p-r)D - wq & if \quad D < \bar{q} \end{cases} = \begin{cases} \left(p - r - \dfrac{w}{1+1/N}\right)\bar{q} & if \quad D \geq \bar{q} \\ (p-r)D - \dfrac{w}{1+1/N}\bar{q} & if \quad D < \bar{q} \end{cases}$$

This is similar to a retailer's profit in a simple revenue sharing contract in which the wholesale price is $w/(1+1/N)$ and the shared revenue is *r*. Therefore, the parameters of a coordinating revenue sharing contract with free items can be related to the parameters of a coordinating simple revenue sharing contract as follows.

$$w_{rs}^F = (1+1/N)w_{rs} = (1+1/N)\lambda c \qquad \text{and} \qquad r^F = r = (1-\lambda)p$$

The channel expected profit can be written as

$$E\pi_c = pES - \bar{q}c,$$

where *ES* is the expected sales. Supplier's expected profit can be written as

$$E\pi_c = r^F ES + w_{rs}^F q - c\bar{q} = (1-\lambda)pES + (1+1/N)\lambda cq - c\bar{q} = (1-\lambda)pES - (1-\lambda)\bar{q}c = (1-\lambda)E\pi_C,$$

Similarly for a buyback contract we have:

$$\pi_R = \begin{cases} p\bar{q} - wq & if \quad D \geq \bar{q} \\ pD + wq + b(\bar{q} - D) & if \quad D < \bar{q} \end{cases} = \begin{cases} (p - \dfrac{w}{1+1/N})\bar{q} & if \quad D \geq \bar{q} \\ (p-b)D - \left(\dfrac{w}{1+1/N} - b\right)\bar{q} & if \quad D < \bar{q} \end{cases}$$

This is similar to a retailer's profit in a simple buyback contract in which the wholesale price is $w/(1+1/N)$ and the buyback price is *b*. Therefore, the parameters of a coordinating buyback

contract with free items can be related to the parameters of a coordinating simple buyback contract as follows.

$$w_{bb}^{F} = (1+1/N)w_{bb} = (1+1/N)[(1-\lambda)p+\lambda c] \qquad \text{and} \qquad b^{F} = b = (1-\lambda)p$$

Supplier's expected profit can be written as

$$\begin{aligned} E\pi_c &= w_{bb}^{F}q - b^{F}(\bar{q}-ES) - c\bar{q} = (1+1/N)[(1-\lambda)p+\lambda c]q - (1-\lambda)p[\bar{q}-ES] - c\bar{q} \\ &= (1-\lambda)pES - (1-\lambda)\bar{q}c = (1-\lambda)E\pi_C \end{aligned},$$

Appendix B: User interface of the experimental software

Demand is random with an average of **200** and uniformly distributed between 100 and 300.

If the Realized Demand is more than your Order Quantity, then you have to incur the costs of unsold items ($1 per unsold item).

If the Realized Demand is less than your Order Quantity then you lose the opportunity to make a sale ($4 per item short).

Selling Price (p) =	$20.00
Buying Price (w) =	$1.00
Shared Revenue (r) =	$15.00

Sl.No	Order Quantity(Q)	Realized Demand(D)	You Ordered	Profit($)	Total Profit($)	Cost of Unsold Items($)	Lost Profit($)
1	214	141	More	491	491	73	0
2	100	235	Less	400	891	0	540
3	150	290	Less	600	1491	0	560
4	200	295	Less	800	2291	0	380
5	232	236	Less	928	3219	0	16
6	121	149	Less	484	3703	0	112
7	156	113	More	409	4112	43	0
8	186	124	More	434	4546	62	0
9	250	265	Less	1000	5546	0	60
10	210	192	More	750	6296	18	0
11	220	163	More	595	6891	57	0
12	300	190	More	650	7541	110	0
13		0	---	ENTER Q	0	0	0
14		0	---	ENTER Q	0	0	0
15		0	---	ENTER Q	0	0	0
16		0	---	ENTER Q	0	0	0
17		0	---	ENTER Q	0	0	0
18		0	---	ENTER Q	0	0	0
19		0	---	ENTER Q	0	0	0
20		0	---	ENTER Q	0	0	0
21		0	---	ENTER Q	0	0	0

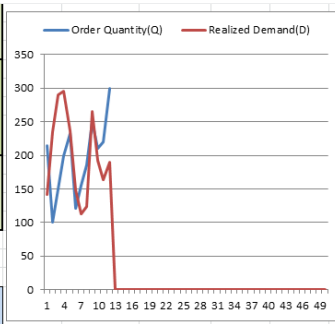

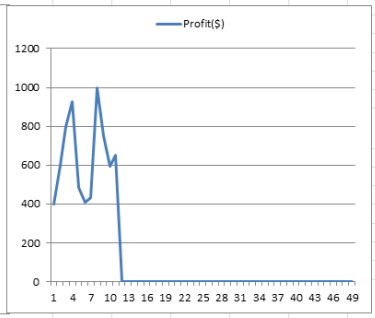

REVENUE SHARING CONTRACT