

IBM Research Report

Experimental Investigation of Supplier-Retailer Contracts: The Wholesale Price Contract

Claudia Keser, Giuseppe A. Paleologo

IBM Research Division

Thomas J. Watson Research Center

P.O. Box 218

Yorktown Heights, NY 10598



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Experimental Investigation of Supplier-Retailer Contracts: The Wholesale Price Contract

Claudia Keser*

IBM T. J. Watson Research Center, NY, and CIRANO, Montreal

and

Giuseppe A. Paleologo

IBM T. J. Watson Research Center, NY

November 9, 2004

Abstract

We examine decision making in a simple supplier-retailer wholesale price contract in the experimental economics laboratory. We observe wholesale prices and order quantities below the game-theoretical predictions. The supply chain's efficiency is as predicted but profits are more equitably allocated.

Keywords: Supply chain coordination, industrial organization, experimental economics

JEL classification: C72, C91, L1

**Corresponding author:* Dr. Claudia Keser, IBM T. J. Watson Research Center, 1101 Kitchawan Road, Rte. 134, Yorktown Heights, NY, 10598, Phone +1-914-945 2473, ckeser@us.ibm.com

1 Introduction

The contractual relationship between the wholesale supplier and the retailer of goods has been the subject of intense theoretical study in recent years. This is because for a simple contract, known as the *wholesale price contract*, the predicted equilibrium is inefficient. In this contract, the supplier specifies a unit wholesale price for whatever number of units the retailer is going to order. The inefficiency arises because the decentralized decision making with this kind of contract results in a total profit that is smaller than what could have been obtained with a centralized decision. This result is of practical interest because the wholesale price contract is widely used in practice.

The existing literature takes inefficiency of the wholesale price contract as the starting point and tries to identify more complex contracts that achieve efficiency preserving decentralized decision making. In contrast, we suggest an empirical investigation of the wholesale price contract. There are three advantages to such an approach. First, it provides empirical validation to the theoretical prediction. Second, it reveals features of behavior induced by this kind of contract that are not dealt with in the theoretical literature but are relevant to actual instances of the contract. Finally, it provides guidance in the design of more sophisticated contracts aimed at increasing efficiency. Due to the difficulty in obtaining field data for the contract in its pure form, we pursue an experimental economics approach. This implies the examination of real human behavior in a controlled laboratory environment and with monetary incentives.

The inefficiency of simple supplier-retailer contracts was first noted by Spengler (1950). Subsequently, others extended his deterministic model to one in which the supplier decides the wholesale price, and the retailer decides the order quantity based on this price and sells units at a fixed market price in the face of uncertain demand. To alleviate the shortcomings of the wholesale price contract, Pasternack (1985) proposed a “buy-back” contract in which the supplier commits to buy back unsold items of the retailer. He showed that efficiency can be achieved if the wholesale price and the buy-back price satisfy a certain explicit relationship. Other contracts that can achieve efficient equilibrium outcomes are revenue sharing

(Cachon and Lariviere, forthcoming), quantity flexibility (Pasternack, 1985; Eppen and Iyer, 1997; Barnes-Schuster, Bassok and Anupindi, 2002; Tsay, 1999), sales rebate (Taylor, 2002), and quantity discount (Tomlin, 2003). These contracts are analyzed and compared in a survey by Cachon (2003).

In this article we present the results of an experimental investigation of a simple supplier-retailer wholesale price contract in a world of stochastic demand. We observe that participants in the role of a supplier charge lower wholesale prices than predicted by the subgame perfect equilibrium solution. Similarly, retailers order less than would be their best response to those wholesale prices. The resulting total profit is not significantly different from the equilibrium prediction, but the individual profits are more equitably allocated among suppliers and retailers. The latter is due to the suppliers' tendency to offer wholesale prices that would imply an equitable profit allocation if demand were deterministic. Retailers tend to anchor their quantity decision on the previous combination of wholesale price and order quantity (hereafter called price-quantity combination) and, from there, adapt to changes in the wholesale price, such that a wholesale price increase (decrease) leads to a decrease (increase) in the order quantity.

Our experiments are related to those by Schweitzer and Cachon (2000), which are confined to the examination of the retailers' decision making. Our results with respect to the retailers' decisions are only partially in keeping with theirs. This might be due, in part, to the retailers' direct interaction with their suppliers in our experiments, where we observe a predominant influence of the suppliers' change in pricing on the retailers' quantity decision.

The article is organized as follows. In the next section we introduce the game-theoretic model, characterize its solutions, and describe the experimental design. Section 3 presents the experimental results. Section 4 concludes the article.

2 The Model

In our simple game-theoretical model, a single supplier sells to a single retailer. The retailer faces a “newsvendor problem”: he sells a product in a short selling season with a stochastic demand D . He is assumed to be a price taker, and the unit retail price is denoted by p . In advance of the selling season, he decides on the number of units he orders from the supplier at the unit wholesale price w , which has been announced by the supplier. After receiving the retailer’s order, the supplier produces at a fixed unit production cost, c , with $c < p$, and delivers at the start of the selling season. The retailer has no additional replenishment opportunity.

The rules of the game are the following. In the first stage of the game, the supplier specifies the wholesale price, w , per unit purchased. In the second stage, the retailer is informed about w and then, based on w , decides on the number of units, q , he orders from the supplier. If he does not want to order from the supplier, he orders $q = 0$. In the third and final stage, the value of consumer demand, d , is observed. Sales are determined as the minimum of the ordered quantity and consumer demand.

The payoffs to the supplier and retailer, respectively, are given by

$$\pi_s = (w - c)q \tag{1}$$

$$\pi_r = p \min\{q, d\} - wq \tag{2}$$

The theoretical solution for this model for any continuous demand distribution function with decreasing hazard rate can be found in Lariviere and Porteus (2001) or, in a more condensed form, in Cachon’s (2003) survey on supply chain coordination with contracts.

In our experiments, demand is known by both parties to be uniformly distributed on $[a, b]$. In the following, we explicitly derive the theoretical solution for this specific demand function, under the assumption of a risk-neutral retailer.

2.1 Subgame perfect equilibrium

The game can easily be solved by backward induction. Given the wholesale price w , the expected profit to the retailer, $\Pi_r(\cdot)$, is a function of the order quantity q

$$\Pi_r(q, w) = pE \min\{q, D\} - wq \quad (3)$$

Maximization with respect to q and subject to the conditions $q \geq 0$ and $\Pi_r \geq 0$ yields the retailer's optimal order quantity, $q^*(w)$, as a function of w , or, the retailer's best reply function to any wholesale price w . The quadratic and concave objective function implies a unique maximum

$$q^*(w) = \begin{cases} b - w(b - a)/p & \text{if } w < p \\ 0 & \text{if } w \geq p \end{cases} \quad (4)$$

If $w < p$, q^* lies between a and b .

Let us now substitute the retailer's best reply function $q^*(w)$ for q in the supplier's profit function $\Pi_s(\cdot)$:

$$\Pi_s(q^*(w), w) = (w - c)q^*(w) \quad (5)$$

Maximization with respect to w and subject to $w \geq 0$ and $\Pi_s \geq 0$, yields the equilibrium wholesale price, w^* . The objective function is quadratic and concave in the interval $[0, p]$ and zero if $w > p$. The solution is given by

$$w^* = \min \left\{ p, \frac{c}{2} + \frac{p}{2} \frac{b}{b - a} \right\} \quad (6)$$

Note that $w^* > c$ and thus $\Pi_s \geq 0$. Consequently, in the subgame perfect equilibrium solution of the game, the supplier proposes a wholesale price w^* and the retailer orders $q^*(w^*) = b/2 - c(b - a)/2p$ units of the product.

This solution is based on the assumption of a risk-neutral retailer. Eeckhoudt, Gollier and Schlesinger (1995) have shown that for the newsvendor problem, corresponding to the second stage of our game, a risk-averse retailer will systematically order less than the best-reply order quantity, while a risk-seeking retailer will systematically order more than the best-reply order quantity.

2.2 Efficient Order Quantity

Let us define the supply chain's expected profit, $\Pi(q)$, as the sum of the two parties' expected profit functions. An efficient order quantity, q^o , is one that maximizes the supply chain's expected profit. The function

$$\Pi(q) = \Pi_r(q, w) + \Pi_s(q, w) = pE \min\{q, D\} - cq \quad (7)$$

depends on the order quantity q but is independent of the wholesale price w . Maximization with respect to q yields the supply chain optimal order quantity

$$q^o = b - \frac{c}{p}(b - a) \quad (8)$$

As $c < w^*$, it is obvious that $q^* < q^o$. The equilibrium order quantity is below the efficient level, and the supply chain's expected profit is increasing in q for $q \in [q^*, q^o]$.

2.3 Experimental design

In our experiments we use the following parameterization:

- Unit production cost, $c = 50$
- Retail price, $p = 250$
- Demand is uniformly distributed between $a = 40$ and $b = 230$

Note also that decisions can be made in integers only, and thus, demand is *discretely* uniformly distributed. Nevertheless, we use the preceding theoretical solutions as approximate benchmarks. Solving the discrete version of the game numerically leads to solutions within one integer of the analytical solutions.

The subgame perfect equilibrium involves a wholesale price of 176, an order quantity of 96, expected payoffs of 12,126 for the retailer and 5,011 for the supplier (see also Table 1 later). The efficient order quantity of 192 implies a total profit of 23,200. How this total profit is shared among the supplier and the retailer depends on the specific wholesale price, which is not determined by the efficient solution. The sum of payoffs in the subgame perfect equilibrium is 17,137, which implies an efficiency of 74 percent.

The participants in our experiments were informed that each of them was to play 30 repetitions of this game with one of the other participants. The finitely repeated game has a unique subgame perfect equilibrium solution that predicts in each of the 30 rounds the subgame perfect equilibrium solution of the baseline game presented earlier.

The experiments were run at the Center for Experimental Social Science (C.E.S.S.) computer lab at New York University. We developed the software based on z-Tree by Urs Fischbacher, University of Zurich. We organized three experimental sessions with ten participants each. Five of the participants in a session were randomly allocated the role of a supplier, and the remaining five were allocated the role of a retailer. Each participant in the role of a supplier was randomly and anonymously matched with one of the other participants in the role of a retailer, with whom he or she interacted over all 30 rounds of the experiment. This yields us five statistically independent observations per session, fifteen in total.

In the beginning of a session, written instructions (available upon request) were distributed and read aloud by the experimenter. Then, the participants had to go through a computerized questionnaire, which tested the understanding of the instructions. Only after each of the participants had answered to each of the questions correctly could the computerized experiment begin. Each participant was seated at a computer terminal that was isolated enough from the other participants' terminals so that participants could make their decisions in anonymity. None of the participants knew with whom they were interacting.

At the end of an experimental session, each participant was paid in US\$ based on his or her individual success in the experiment. A participant's profit in the experiment, that is, the sum of his or her profits over all 30 rounds, was converted into US\$ in the following way. The profit of each supplier (retailer) was compared to the average profit of the other suppliers (retailers) participating in the session:

$$\text{A supplier's payment} = \$23 + 0.001(\text{own profit} - \text{other suppliers' average profit})$$

$$\text{A retailer's payment} = \$23 + 0.001(\text{own profit} - \text{other retailers' average profit})$$

This payment scheme yields an amount larger than or equal to \$23 if the supplier's (retailer's) profit was higher than or equal to the other suppliers'(retailers') profit; otherwise, it yields an amount below 23 otherwise. To this amount, a \$7 show-up fee was added. Thus, the average total payment per participant was \$30.¹ An experimental session lasted about one and a half hours.

3 Results

In this section we compare the experimental results with the theoretical predictions. Moreover, we test several behavioral hypotheses in order to identify participants' decision heuristics. We present the results in three parts: the suppliers' pricing decision, the retailers' decision on the order quantity, and the resulting profits.

We use non-parametric statistics (Siegel and Castellan, 1988). All tests are two-sided. We require significance at the 10-percent level. We denote the Wilcoxon signed ranks test simply as the Wilcoxon test. For correlation analysis we use the Spearman rank-order correlation coefficient ρ .

Table 1 summarizes the experimental results, providing averages and standard deviations of the decisions and resulting profits by the fifteen supplier-retailer pairs over the 30 periods. It also compares these average to the respective theoretical predictions, the sub-game perfect equilibrium and the efficient order quantity solution. The last two columns of Table 1 show the averages and standard deviations when we exclude all those cases where the retailer ordered nothing from the supplier so that there was no effective contractual relationship between the supplier and the retailer. This happened in only 36 out of 450 decisions. Twenty-nine of the zero order decisions were made by the same two participants. The average wholesale price that led to a zero order was 225. In eight cases the wholesale

¹Note that our payment scheme is strategically equivalent to applying a conversion factor directly to a participant's profit in the experiment. We have chosen this payment scheme for two reasons. First, a conversion factor applied directly to a participant's profit would have to be significantly less important than the one we used. Second, we wanted to avoid extreme differences in the average payments to suppliers and retailers.

price was 250 or above, which would have led with certainty to a non-positive profit for the retailer. Thus, we will exclude the zero order cases in the data analysis below. Note that this exclusion does not qualitatively affect any of the statistical results.

Table 1: Theoretical evaluation of the game and experimental results

	Equilibrium	Efficient quantity	Average	(stdev)	Average w/o null orders	(stdev)
Wholesale price	176	N/A	151	(53)	144	(44)
Order quantity	96	192	90	(50)	98	(44)
Supplier's profit	12,126	N/A	7,845	(5,038)	8,527	(4,665)
Retailer's profit	5,011	N/A	6,786	(7,518)	7,376	(7,555)
Total profit	17,137	23,200	14,631	(8,336)	15,903	(7,434)

3.1 Wholesale prices

One immediate observation is that the observed wholesale prices are on average lower than in equilibrium. This difference is statistically significant at the 2-percent level (Wilcoxon test based on fifteen independent observations). The median and mode of the observed wholesale prices is 150.² At this wholesale price supplier and retailer equally split the total profit if the retailer can sell all ordered units to consumers. Figure 1 shows the average wholesale prices over all fifteen suppliers in each of the 30 periods; they are all below the equilibrium price (presented by the dashed line in Figure 1). The wholesale prices show no significant tendency to either increase or decrease from the first set of fifteen periods to the second set of fifteen periods (Wilcoxon test based on fifteen independent observation pairs).

²Eight of the fifteen suppliers offer an average retail price above 150, while seven retailers offer an average retail price below 150.




Figure 1: Average wholesale prices over time compared to the equilibrium (dashed line)

3.2 Order quantities

The observed order quantity is not significantly different from the subgame perfect equilibrium quantity of 96 (Wilcoxon test based on fifteen independent observations). We observe no significant increase or decrease in the order quantity from the first fifteen periods to the final fifteen periods (two-sided Wilcoxon test based on fifteen independent observation pairs).

It is important to note that the theoretically predicted order quantity of 96 is the best-reply only to the equilibrium wholesale price of 176. As the actual wholesale price tends to be significantly below the equilibrium level, we have to compare the observed order quantities to their actual best-reply quantities. The best-reply order quantity is a decreasing function in the wholesale price. This implies that the best-reply order quantities to the observed wholesale prices below the equilibrium level are larger than 96.

Figure 2 show for each of the fifteen retailers the average price-quantity combination over the 30 periods, as well as the retailers' theoretical best reply function to any potential wholesale price. The length of the dashed vertical line between each point, presenting one of the fifteen retailers, and the best reply function indicates the difference of the observed average order quantity from the retailer's average profit-maximizing order quantity to the actual wholesale prices. It is obvious in this figure that, on average, retailers tend to order less than their best reply to the given wholesale prices. This difference is statistically significant at the 2-percent level (Wilcoxon test based on fifteen observations). The retailers thus forgo profits by 6.7 percent of what they could have made by choosing best reply order quantities.




Figure 2: Comparison of retailers' average order quantities over all 30 periods to their respective average best reply order quantities

Schweitzer and Cachon (2000) hypothesize and find support in their data that retailers systematically order too few units with a high profit margin while they order too many with a low profit margin. According to their definition, a high profit margin is 50 percent or higher, while a low profit margin is below 50 percent. In our game, "too few" translates into less than best reply while "too many" translates into more than best reply. The subgame equilibrium solution of our game puts the retailer in a low-profit situation (30 percent profit margin). However, as we observed wholesale prices significantly below the equilibrium price, we need to consider the profit margin for each of the actually offered wholesale prices in order to evaluate the hypothesis. Of the fifteen retailers in our experiments, only five behave in the majority of periods according to the Schweitzer-Cachon hypothesis, while ten of the retailers behave in the opposite way. Thus, we cannot support their hypothesis. This result may be an artifact of our specific parameter choice. It might also be caused by the direct supplier-retailer interaction.

3.2.1 Orders related to wholesale prices

Theoretically, the retailer's best-reply order quantity decreases with the wholesale price charged by the supplier. In the experiments, considering for each retailer the average wholesale price at which he ordered positive quantities and the average order quantity (excluding zero orders) over all 30 periods, we observe that the correlation coefficient between wholesale prices and quantities is neither significantly negative nor positive ($\rho = -0.34$). At this level of aggregation, the predicted decrease of the order quantity in the wholesale price does not hold.

To examine the price-quantity relationship at the level of individual decisions over time, we begin considering the first period only. Similarly to the above result at the aggregate level, we observe no significantly negative or positive correlation between a retailer's order quantity and the respective supplier's wholesale price in the first period ($\rho = -0.14$). Note also that we observe no specific pattern of behavior in the first period, neither in the suppliers' pricing decision nor in the retailers' ordering decision.

Then, we investigate whether there is a dynamic relationship between wholesale prices

and order quantities, based on the previous price-quantity combination as an anchoring point. More concretely, we examine for each of the retailers whether he reacted to an increase or decrease in the wholesale price from the previous to the current period with an increase or a decrease in the order quantity, if he reacted at all. We observe that retailers significantly tend to react to a wholesale price increase (decrease) with a decrease (increase) in the order quantity: thirteen of the fifteen retailers reacted that way in the majority of cases in which there was a price change to which they reacted; two retailers reacted in the opposite direction (binomial test, 1-percent significance).

These two observations on the price-quantity relationship at the level of individual decisions give support to the following interpretation of retailers' behavior. In the first period, retailers choose their order quantities in an *ad hoc* way, for which we have no explanation. The first-period price-quantity combination will then be considered an anchor point around which the retailer's future decisions will be based. More specifically, in each of the following periods, the specific reference point will be the price-quantity combination of the previous period.³ This implies that, around the anchoring point, order quantities tend to decrease (decrease) with an increase (decrease) in the wholesale price. This result may be considered qualitatively in keeping with the theoretical prediction, although it is quantitatively different from the theoretical best reply.

3.2.2 Orders related to previous waste

An important question that comes to ones mind is whether and how retailers react to erroneous demand estimates. If a retailer orders less than the actual demand, he will face waste in terms of unsatisfied demand. On the other hand, if a retailer orders more than the actual demand, he will face waste in terms of overstock that will be thrown away at the end of the period.

³We limit our attention in this analysis to the previous period as we know from experimental investigations of participants' strategies that people typically condition their decisions not more than on one or two previous periods (Selten, Mitzkewitz and Uhlich, 1997, Keser, 2000).

Either unsatisfied demand or overstock might induce a reassessment of a retailer’s strategy. In the literature we find two opposing hypotheses that could apply. According to the availability hypothesis (Tversky and Kahneman, 1973), a recent waste is predominant in memory. This would imply an increase in the order quantity after unsatisfied demand and a decrease in the order quantity after overstock in order to avoid the same kind of waste. The opposite reaction relates to the gambler’s fallacy hypothesis (Camerer and Kunreuther, 1989), which would imply that, based on an erroneous belief in conditional probabilities, a retailer assumes that a specific event, such as a very high demand, is not likely to recur in the next period. The static equilibrium theory presented in Section 2 above does not account for this kind of dynamics, which could describe, however, potential decision heuristics.

Considering that each retailer’s change in order quantity depends on either unsatisfied demand or overstock in the previous period, we do not find significant evidence for either of the two hypotheses (binomial test). Eight of the fifteen retailers react in the majority of cases according to the availability hypothesis, while seven retailers react in the majority of cases according to the gambler’s fallacy hypothesis.

Schweitzer and Cachon (2000) in their newsvendor decision-making experiments find evidence for what they call the “chasing demand heuristic.” This heuristic, which relates to our application of the availability hypothesis above, is based on the anchoring on the previous order quantity and the adjustment toward prior demand. The reason why we do not find statistically significant evidence for this heuristic in our experiments might be that, due to the direct interaction with the supplier, the retailer’s attention is focused on the wholesale price proposed by the supplier rather than the demand.

3.3 Profits

Since the observed order quantity is not significantly different from the quantity prescribed by the subgame perfect equilibrium solution, it comes as no surprise that the sum of supplier-retailer profits is not significantly different from the equilibrium prediction (Wilcoxon test).

This implies that the observed efficiency of 69 percent is not significantly different from the 74 percent efficiency predicted by the subgame perfect equilibrium. Note also that there is no significant tendency of the total profit to increase or decrease from the first set of fifteen periods to the second set of fifteen periods (Wilcoxon test based on fifteen independent observation pairs).

We observe, however, that the suppliers' profit is significantly lower than predicted, while the retailers' profit is higher than predicted (Wilcoxon tests based on fifteen observations, 1- and 2-percent significance, respectively). Because in the subgame perfect equilibrium the supplier makes a much higher profit than the retailer, these observations imply that in the experiments profit is allocated more equitably than in equilibrium. Indeed, the retailers' profit is not significantly different from the suppliers' profit (Wilcoxon test): ten of the fifteen retailers make lower profits than their respective suppliers, while the opposite is true for the remaining five retailers.

The supplier's profit is significantly positively correlated with the wholesale price ($\rho = 0.83$, based on fifteen independent observations, 5-percent significance); whereas, the retailer's profit is significantly negatively correlated with the wholesale price ($\rho = -0.96$, based on fifteen independent observations, 5-percent significance). Thus, we observe a significantly negative correlation between the supplier's and the retailers profit ($\rho = -0.71$, based on fifteen independent observations, 5-percent significance). This means that the wholesale price tends to determine how the total profit is allocated among supplier and retailer. It does not influence the size of the total profit, though. The size of the expected total profit is, by the structure of the game, determined by the retailer's quantity decision. The latter shows, on the aggregate, no tendency to be affected by the wholesale price, given that order quantities and wholesale prices are not significantly correlated.

4 Conclusions

In our experiments, we observe that the wholesale price contract yields an efficiency that is not significantly different from the equilibrium prediction. The observed behavior of suppliers and retailers is very different, though, from the equilibrium solution. Suppliers charge lower wholesale prices than predicted by the subgame perfect equilibrium solution. More concretely, the median and mode of the prices are at 150, in the middle of the strategy space between unit production cost of 50 and retail price of 250. This implies an equal split of the net profit if the retailer can sell all of the units to consumers. However, as the latter is unlikely to happen, the burden of the risk of overstock has to be carried by the retailer. The wholesale price of 150 may be considered an obvious anchor point for fairness, although it is not obvious how one would precisely have to take the risk of overstock into account.

It can be easily seen in equation (6) that the subgame perfect equilibrium wholesale price is always above this anchor point midway between unit production cost and retail price. Thus, a focus in the suppliers' decision making on this anchor point will always imply a tendency to wholesale pricing below the equilibrium level.

As a consequence of pricing around the fair anchor point, suppliers make lower profits and retailers make higher profits than theoretically predicted. Consequently, the profits are more equitably allocated among suppliers and retailers than in the subgame perfect equilibrium solution. These results are very similar to those observed in ultimatum bargaining or principal agent experiments (e.g., Roth, 1995, Keser and Willinger, 2000), which is not surprising due to the ultimatum structure in our game. On the one hand, the supplier has the advantage of proposing the wholesale price and, thus, determining how the total profit will be allocated (if we ignore the risk to the retailer due to the stochastic demand). On the other hand, though, he has to fear rejection or very low orders from the retailer. In other words, the retailer can determine the size of the pie to be allocated. Thus, the suppliers tend to offer wholesale prices that lead to more or less equitable outcomes. This result is also in keeping with the theories of inequity aversion and fairness by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). Beyond such motivational limits to the participants' rationality, there

are, obviously, also cognitive limits: due to stochastic demand, the wholesale price contract decision situation is too complex to be solved analytically on the spot by participants in an experiment.

Another observation is that retailers order less than would be their best response to the wholesale prices they are offered. This could potentially be explained by risk-averse preferences of the retailers. Such an explanation was, however, ruled out by Scheitzer and Cachon (2000).

Our data provide support for a decision heuristic of the retailers that is based on an anchor point created by the price-quantity combination in the first period. In the following periods, the retailer will adapt his order quantity to changes in the wholesale price around this anchor point. He will decrease his order quantity as the wholesale price increases, and vice versa.

Interestingly, in spite of the observed behavioral differences from the equilibrium prediction, no efficiency is lost or gained in our experiments with respect to the theoretical efficiency. This might be a coincidence, due to the specific parameter choice, which should be subject of further investigation. Our future research agenda also involves the comparison of the observed behavior and efficiency in the wholesale price contract to participants' behavior and efficiency in more sophisticated contracts, as for example, the buy-back contract. The latter should, according to theory, lead to full efficiency.

5 References

- Barnes-Schuster, D., Y. Bassok and R. Anupindi, 2002, Coordination and flexibility in supply contracts with options, *Manufacturing & Service Operations Management* 4, 171-207.
- Bolton, G. E. and A. Ockenfels, 2000, ERC: a theory of equity, reciprocity and competition, *American Economic Review* 90, 166-193.
- Cachon, G., 2003, Supply chain coordination with contracts, in: A. G. de Kok and S. C. Graves, eds., *Handbooks in Operations Research and Management Science Volume 11: Supply Chain Management: Design, Coordination and Operation* (Elsevier: North-Holland, Amsterdam).
- Cachon, G. and M. Lariviere, forthcoming, Supply chain coordination with revenue sharing, *Management Science*.
- Eeckhoudt, L., C. Gollier and H. Schlesinger, 1995, The risk-averse (and prudent) newsboy, *Management Science* 41, 768-794.
- Eppen, G. and A. Iyer, 1997, Backup agreements in fashion buying – the value of upstream flexibility, *Management Science* 43, 1469-1484.
- Fehr, E. and K. M. Schmidt, 1999, A theory of fairness, competition and cooperation, *Quarterly Journal of Economics* 114, 817-868.
- Keser, C., 2000, Strategically planned behavior in public goods experiments, *CIRANO Scientific Series* 2000s-35.
- Keser, C. and M. Willinger, 2000, Principals' principles when agents' actions are hidden, *International Journal of Industrial Organization* 18, 163-185.
- Lariviere, M. and E. Porteus, 2001, Selling to the newsvendor: an analysis of price-only contracts, *Manufacturing & Service Operations Management* 3, 293-305.
- Pasternack, B., 1985, Optimal pricing and returns policies for perishable commodities, *Marketing Science* 4, 166-176.

- Roth, A. E., 1995, Bargaining Experiments, in: J. H. Kagel and A. E. Roth, eds., *The Handbook of Experimental Economics* (Princeton University Press, Princeton), 253-348.
- Schweitzer, M. E., and G. Cachon, 2000, Decision bias in the newsvendor problem with a known demand distribution: experimental evidence, *Management Science* 46, 404-420.
- Selten, R., M. Mitzkewitz and G.R. Uhlich, 1997, Duopoly strategies programmed by experienced players, *Econometrica* 65, 517-555.
- Siegel, S. and N. J. Castellan, Jr., 1988, *Nonparametric statistics for the behavioral sciences* (McGraw-Hill, USA).
- Spengler, J. , 1950, Vertical integration and antitrust policy, *Journal of Political Economy* 58, 347-352.
- Taylor, T., 2002, Coordination under channel rebates with sales effort effects, *Management Science* 48, 992-1007.
- Tomlin, B., 2003, Capacity investments in supply chains: sharing the gain rather than sharing the pain, *Manufacturing & Service Operations Management* 5, 317-333.
- Tsay, A., 1999, Quantity-flexibility contract and supplier-customer incentives, *Management Science* 45, 1339-1358.