THE IMPACT OF HUMAN DECISION MAKERS' INDIVIDUALITIES ON THE WHOLESALE PRICE CONTRACT'S EFFICIENCY: SIMULATING THE NEWSVENDOR PROBLEM

Stavrianna Dimitriou
Stewart Robinson
Kathy Kotiadis

Operational Research and Management Sciences Group

Warwick Business School

Coventry, CV4 7AL, UK

ABSTRACT

Suppliers and retailers in the newsvendor setting need to submit their pricing and inventory decisions respectively, well before actual customer demand is realized. In the literature they have both been typically considered as perfectly rational optimizers, exclusively interested in their own respective benefits. Under the above set of conditions the wholesale price-only contract has long been analytically proven as inefficient. We asked real human subjects to act as suppliers or retailers in simulation games performed in the laboratory. We found their decisions to significantly deviate from the perfectly rational decisions. By using Agent Based Simulation as the evaluation tool, we investigated the effect of their varying individual preferences on the contract's efficiency. In doing so we established sufficient evidence that the contract can emerge as efficient, in spite of the underlying strategies' under-performances. This counter-intuitive result fully supports the contract's long observed wide popularity.

1 INTRODUCTION

In the typical newsvendor problem there is only one supplier and one retailer that interact with each other. Although this setting is not complex at all, it constitutes the fundamental building block of a plethora of realistic network configurations. In addition, it is sufficiently rich to address the question of whether and to which particular degree can any type of contract become efficient. A contract is said to be *efficient*, if the set of partners' decisions produces an aggregate channel profit that is not inferior to the *first-best case* profit, achieved by an 'integrated' newsvendor, who has access to all available information (Lariviere 1999, Tsay, Nahmias, and Agrawal 1999). In case the *wholesale price-only contract* is in place, there is only one incentive exchanged between the two partners: the wholesale price w that the supplier decides to charge to the retailer for every unit of product that the first delivers to the latter. In response to this price, the retailer determines his preferred order quantity q for the time period.

This problem has been extensively studied over decades and has been analytically proven as *inefficient*. Namely, a rationally optimizing supplier that interacts with a retailer in turn exclusively interested in maximizing his respective profit, would never make the *wholesale price only* contract efficient. The reason is that neither partner takes into account the externality its decision imposes on the other's profit (this phenomenon is known as the "double marginalization" problem: Spengler, 1950). In contrast, a number of other contracts, usually more costly to administer, have been suggested as *efficient* or at least worth adopting. In this light, the *wholesale price contract*'s practical prevalence over other contractual forms remains paradoxical (Cachon 2003, Keser and Paleologo 2004). For detailed surveys of contracts and reviews of the analytical results acquired so far we refer the interested reader to Tsay, Nahmias, and Agrawal (1999), Cachon (2003), Simchi-Levi, Kaminsky, and Simchi-Levi (2008).

In this paper, in an attempt to explain the afore-mentioned dominance of the *wholesale price contract* over other contractual forms, we adopt the view that human suppliers and retailers do not necessarily make perfectly rational decisions. Because of their varied limited knowledge and finite cognitive abilities they might not be in a position of searching the entire solution space and, thus, identifying their optimal decision. It is not only because they might lack perfect information that they cannot make perfectly rational choices, but also because they "experience limits in formulating and solving complex

problems and in processing information" (Simon 1957 in Williamson 1981: pp. 553). This is the reason why we perceive them as 'boundedly rational', having to settle at reasonable, *satisficing*, choices (Simon 1996, North and Macal 2007, Gilbert 2008).

In order to provide the relevant empirical evidence, we asked real human subjects to act as suppliers or retailers in simulation games in the laboratory. We observed them to make decisions that are significantly different from their rationally optimizing counterparts'. But, what is even more important, their decisions greatly varied: namely their individual preferences, emotions, motivations and cognitive abilities differed to a degree that could not justify the choice of a 'typical' decision maker. For this reason, we explored the effect of the combination of their varying decision making strategies on the *wholesale price contract*'s efficiency.

The purpose of this paper is to provide an added explanation to the *wholesale price contract's* wide practical popularity by demonstrating that this contract can become *efficient* when particular human suppliers and retailers interact with each other over prolonged periods of time. Since we recognized at least one such human supplier-retailer interaction in the newsvendor setting, we can safely argue that the *wholesale price contract* can become *efficient*, in contrast to non-empirically validated beliefs.

This result is of equal significance to both academics and practitioners. Academics will find an interest in the conditions under which the *wholesale price contract*'s *efficiency* can be achieved. The methodological differences of this study from prior empirical work on the newsvendor problem will also be of interest. As for practitioners, this paper will hopefully help them understand that instead of solely investing in implementing and administering complex, yet efficient, contract types, they could alternatively consider effective management training that focuses on overall efficiency. The reason is that in spite of partners' poor individual decisions, global efficiencies can be achieved and, so, it is really important to train decision makers how to reach these decisions. This is where our simulation games along the lines of 'business flight simulators' could help (Sterman 1989).

The paper starts by reviewing existing empirical work on the *newsvendor problem*. Our approach is subsequently described. Later on, the experimental results are outlined. The paper concludes by summarizing the results acquired and recommending possible directions for future research.

2   RELATED EMPIRICAL RESEARCH

Carlson and O'Keefe (1969) were the first to perform a controlled human experiment of the newsvendor problem and they acknowledged participants' tendency to make "almost every kind of mistake" (p. 483). Fisher and Raman (1996) conducted an industrial experiment at a fashion apparel manufacturer's site and ascertained that managers ordered quantities that were systematically inferior from their respective cost minimizing quantities. Nevertheless, they were occupied neither with explaining this observed bias nor reflecting on its potential persistence at different settings. The reason is it was not until relatively recently that a systematic behavioral perspective has been brought to the newsvendor problem (Bendoly, Donohue, and Schultz 2006; Gino and Pisano 2006). To this end, researchers have mainly resorted to closed laboratory experiments that offered the possibility of complete isolation of the experimental factors in question from all remaining environmental factors (Croson and Donohue 2002, Wu and Katok 2006).

Sterman's (1989) seminal paper described the first behavioral experiment that demonstrated individuals' bounded rationality' and, hence, limited ability to understand and control systems with lagged, indirect and non-linear feedbacks. Although there is no direct association between Sterman's paper and the newsvendor problem (this behavioral experiment was conducted within the Beer Game setting), it was the first that systematically explored the effect of a multitude of behavioral complexities on human decisions. It was also the first paper that suggested the application of Kahneman, Slovic and Tversky (1982) "anchoring and insufficient adjustment" heuristic in the supply chain context. In accordance with this heuristic, Sterman's participants were found to initially choose their order quantities ("anchor") based on the current stock levels and subsequently make insufficient adjustments ("insufficiently adjust") towards desired stock levels.

Most subsequent laboratory investigations confirmed similar results for the newsvendor problem. When human participants were asked to determine their retail order quantity decisions, facing uniformly distributed customer demand, they were found to "anchor" on mean demand and "insufficiently adjust" towards optimal order quantities (Schweitzer and Cachon 2000; Benzion et al 2008; Bostian, Holt, and Smith 2008). Although human retailers were proven to use their previous rounds' results to improve their later rounds' decisions (Benzion et al 2008; Bolton and Katok 2008; Bostian, Holt, and Smith 2008), their decisions consistently diverged from the theoretical optimum quantities (Schweitzer and Cachon 2000; Benzion et al 2008; Bolton and Katok 2008; Bostian, Holt, and Smith 2008).

Analogous results about the limited predictive power of analytical models were corroborated by recent experimental studies on supply chain contracts, applied to the newsvendor setting. Most of these studies applied a somewhat similar experimental protocol. Participants were asked to act as suppliers or retailers facing uniformly distributed customer demand, either interactively with each other or against computer pre-automated responses that represented their corresponding partner. The

resulting aggregate channel efficiency was subsequently assessed in comparison to the predictions of standard contracting theories.

In this regard, Katok and Wu (2006) and Ho and Zhang (2008) turned their attention to types of contracts that had been analytically demonstrated as efficient. Katok and Wu (2007) discovered that the buy-back and revenue sharing contracts could not achieve channel efficiency, as analytically anticipated. What was even more important was that Katok and Wu concluded that the buy-back and revenue sharing contracts were found not to offer significant respective improvements over the analytically proven inefficient *wholesale price contract.* Ho and Zhang (2008) found that, in contrast to standard contracting theories' predictions, neither two-part tariffs nor quantity discounts could achieve channel efficiency.

Keser and Paleologo (2004) and Loch and Wu (2008) concentrated on empirically refuting the *wholesale price contract's* analytically proven in-efficiency. Keser and Paleologo drew the conclusion that the wholesale price contract's overall efficiency coincided with its theoretical prediction, despite the fact that suppliers charged significantly lower prices than their theoretical counterparts. Loch and Wu could witness increase of the achieved overall efficiency, when human suppliers and retailers were interested in maintaining positive relationships with their partners. As for the cases they sought to improve their individual social status, overall efficiencies aggravated below theoretical predictions. Nevertheless, improved efficiency scores remained significantly lower than 100%, leaving the question of whether efficiency can be experimentally achieved or not open to further exploration.

3 THE APPROACH

This paper adds further support to the *wholesale price contract's* potential efficiency. In greater detail, we adopt the view that participants have their individual, boundedly rational preferences and limitations which drive their decision making. In addition, human suppliers and retailers have been assumed to face a normally distributed customer demand, because it more closely reflects reality (Keser and Paleologo 2004, Son and Sheu 2008). From the preliminary analysis of the data, we found that this added complexity complicated our human subjects' task. The human subjects showed that they needed more time to get used to their roles and, hence, effectively implement their preferred decision making strategies. In order to enable the long run performances of their decisions to be understood, we chose to use Agent Based Simulation in conjunction with laboratory experiments. This enabled us to draw conclusions on efficiency scores achieved after sufficiently long periods of time that would not be possible in the laboratory alone.

Figure 1 presents the steps that we followed to elicit knowledge on how human subjects make their decisions and assess the overall performance of their possible interactions. As a first step (*Stage 1*), we had to understand how suppliers and retailers make their decisions in the newsvendor context; namely identify the decisions each of them is entrusted with (*decision variables*) and the factors (or *decisions attributes*) they take into consideration to make their decisions. Based on this (*Stage 2*), we ran laboratory experiments, where we asked human subjects to play the role of the supplier or the retailer in simulation games. In *Stage 3* we used multiple regression analysis to determine the participants' decision models, namely the type of relationship that exists between their decision variables and the corresponding attributes. We also attempted to better understand the underlying reasoning of different participants' strategies. In *Stage* 4 we simulated all possible combinations of participants' decision making strategies to compare their respective outcomes. In this way, we studied their corresponding efficiency scores and established whether and under which conditions efficiency could be achieved.

3.1 The Decision Making Process

When the *wholesale price-only* contract is repetitively applied in the *newsvendor* setting, the supplier is assumed to be the one who leads the Stackelberg game. So, in advance of every selling season (t) he needs to specify the wholesale price w that he wishes to charge the retailer for every unit of product that he delivers. In response to that, the retailer must choose an order quantity q. The supplier is assumed to have infinite capacity and is in a position to immediately provide to the retailer any quantity that he places an order for. The order is assumed to be instantaneously delivered to the retailer. The retailer is, in turn, responsible for satisfying the customer demand (*D*), based on his inventory availability. The retailer sells each unit of product at the price of p that is determined by market competition, as is usual for commodity products (Chopra and Meindl 2007). The supplier has to incur a unitary production cost c that is $c<p$. The product under study can only last for one selling season and no left-over inventories at the end of a season can be carried over from one period to the next. Figure 2 illustrates the problem.

According to the above description, the supplier and retailer have exactly one decision to make in each and every round (or selling season). The supplier's decision variable is the wholesale price w, while the retailer's is the order quantity q.
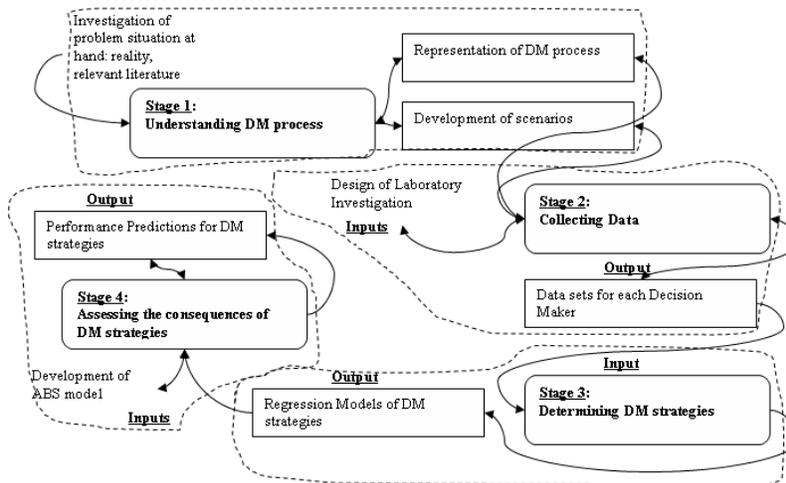
Figure 1: The Approach used (adapted from Robinson et al. 2005)

Let f be the probability density function of stochastic demand D. F is differentiable, strictly increasing and $F(0) = 0$. Let $\overline{F}(x) = 1 - F(x)$; $\mu = E(D)$ and $\sigma^2 = Var(D)$. The demand observed is represented by d. For each unit of demand the retailer does not satisfy the retailer incurs a goodwill penalty cost of g_r. For reasons of simplicity we assume that the analogous cost for the supplier is $g_s=0$, so that $g=g_r$. The retailer earns at the end of the selling season $s<c$ from the supplier per unit unsold, where s stands for the net of any salvage expenses. Let $I(q)$ be the expected leftover inventory $I(q) = (q-d)^+$, while $L(q)$ is the lost sales function, given by $L(q) = (d-q)^+$.
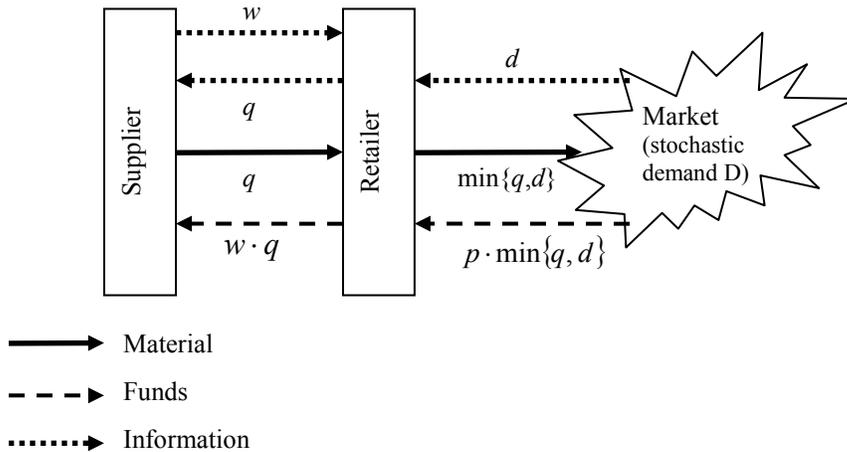

Material

Funds

Information

Figure 2: The Decentralized Newsvendor Problem

A rationally optimizing supplier would charge as much (w^*) as maximizes his expected profit $\Pi_s = (w - c) \cdot q$. Previous analytical investigations of the supplier's problem have proven that w^* is given by the following closed-form mathematical expression (Pasternack 1985, Lariviere and Porteus 2001, Debo and Sun 2004):

$$w^* : F^{-1}(\frac{p+g-w^*}{p+g-s}) - \frac{w^*-c}{(p+g-s)\cdot\overline{F}(F^{-1}(\frac{p+g-w^*}{p+g-s}))} = 0 \tag{1}$$

But in the laboratory human suppliers (*i*) have been found to base each period's wholesale price decision $w(t)$ (Keser and Paleologo 2004, Katok and Wu 2007) on the:
i. previously charged wholesale price $w(t-1)$,
ii. previously placed order quantity $q(t-1)$, and
iii. previously realized profit $Ps(t-1)$: $\langle w(t) \rangle_i = f_i[w(t-1), q(t-1), P_s(t-1)]$.

A rationally optimizing retailer would order as many (q^*) products as would maximize his expected profit: $\Pi_r = p \cdot E\{\min(q, D)\} + s \cdot E\{(q-d)^+\} - g \cdot E\{(d-q)^+\} - w \cdot q$. Analytical research has shown that q^* would depend on the charged wholesale price w, according to the following mathematical expression (Lariviere 1999, Lariviere and Porteus 2001, Cachon 2003):

$$q^*(w) = F^{-1}(\frac{p + g - w}{p + g - s}) \tag{2}$$

Nevertheless, human retailers (j) have been observed to make their order quantity decisions $q(t)$ taking into consideration (Schweitzer and Cachon 2000, Benzion et al 2008):

i.   the currently charged wholesale price $w(t)$,
ii.   the last period's order quantity $q(t$-1$)$,
iii. the previously observed demand $d(t$-1$)$ and
iv.   the previously realized profit $Pr(t$-1$)$: $\langle q(t) \rangle_j = f_j [w(t), q(t-1), d(t-1), P_r(t-1)]$.

These are the decision variables we use in our experiments.

3.2   The Gaming Sessions

Volunteers were recruited from a pool of 2007 graduate students at the University of Warwick on the basis of convenience. The only requirement set is that all participants, prior to the experiment, had received, as part of their curriculum requirement, formal classroom training on the newsvendor problem. They were randomly assigned to play either the role of the supplier or the retailer. Written instructions on the required task were distributed to them prior to their allocated session, so they could work with the game as quickly as possible during the gaming sessions. The instructions informed them that the product under study is a perishable widget of general nature with cost of c=50 monetary units and sold at p=250 monetary units. Without loss of generality we set net salvage expenses s=0 and retailer's goodwill penalty cost g=1. Participants were also informed that demand is random and each round's demand is independent of any previous round's.

We ran the game for 50 consecutive rounds for each participant. We found all human subjects to need some time to get used to their new role, during which the 'initialization bias' was still under effect (Pidd 2004, Robinson 2004, Law 2007). For example, RET$_3$, illustrated in Figure 3, only converged to his decision making strategy after the first couple of rounds. In this figure RET$_3$ decisions are compared with the respective decisions that would have been made by the rationally optimizing retailer RET$_{OPT}$ $q^*(w)$ in response to the wholesale price w. In order, thus, to avoid the risk of making inferences while the 'initialization bias' was still under effect, we eliminated from each gaming session the first 10 observations. We kept 10 samples for each decision variable, which both met the minimum sample size requirements and ensured sufficient statistical power for both the supplier's and retailer's derived decision models (Weisberg 2005, Hair et al 2006).
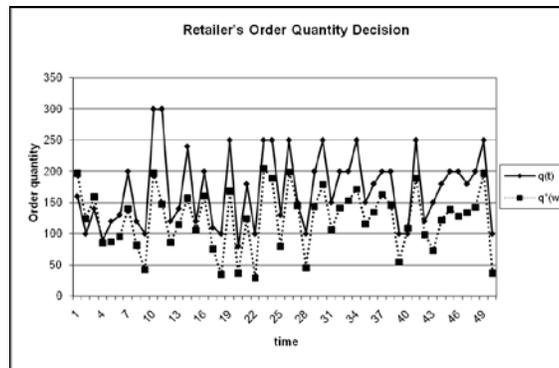

Figure 3: RET$_3$ q-decisions as observed in the laboratory

Participants acting as suppliers were asked to play against computer-enabled scenarios where possible order quantities closely followed true customer demand. Customer demand followed the truncated at zero normal distribution with μ=140 and $\sigma = 80$ (although the participants were not aware of this). Special focus has been given to the scenarios where the automated retailer behaved like the rationally optimizing retailer and ordered q^* widgets. As for participants acting as retailers, they were asked to play against automated supplier responses that took all possible values in the range of acceptable w- prices [c=50,p=250].

3.3   The Decision Making Strategies

Bowman's managerial coefficient theory (Bowman 1963) has already been widely used to model decision making in experimental work (Remus 1978, Croson and Donohue 2006). Since the data we collected complied with all linear regression requirements (Weisberg 2005, Hair et al 2006), we portrayed each supplier's (*i*) and retailer's (*j*) decision making strategies as a simple linear model:

$$\langle w(t)\rangle_i = \alpha_0^i + \alpha_w^i \cdot w(t-1) + \alpha_q^i \cdot q(t-1) + \alpha_P^i \cdot P_s(t-1) \tag{3}$$

$$\langle q(t)\rangle_j = \beta_o^j + \beta_w^j \cdot w(t) + \beta_q^j \cdot q(t-1) + \beta_d^j \cdot d(t-1) + \beta_P^j \cdot P_r(t-1) \tag{4}$$

The value of each coefficient α^i and β^j reflects the importance that each supplier *i* and retailer *j* respectively assigned to each of their decision attributes for the decision $\langle w(t)\rangle_i$ or $\langle q(t)\rangle_j$. The high adjusted values of R^2 that we achieved increased our confidence in the models' explanatory power (Weisberg 2005, Hair et al 2006).

The suppliers' results are shown in Table 1. From the regression coefficients it becomes evident that no human supplier seemed to take into account the profit he previously realized (all α_P^i=0). In contrast, all human suppliers seemed to assign the greatest importance to the wholesale price *w(t-1)* they charged during the last period. They also assigned little importance to the retailer's response quantity, because they lacked knowledge and control over retailers' underlying decision making logic.

Table 1: Human suppliers' linear regression decision making models

	Dependent Decision Variable: $\langle w(t)\rangle_i$		
	Coefficients		
	i=1 SUP$_1$	*i*=2 SUP$_2$	*i*=3 SUP$_3$
α_0^i	115.851	43.929	11.733
α_w^i	0.506	0.769	0.921
α_q^i	-0.014	0.011	-0.002
α_P^i	0	0	0
Adj. R^2	0.852	0.889	0.958
Significance level *p*	<0.01	<0.01	<0.01

Table 2 outlines the retailers' decision models.

Table 2: Human retailers' linear regression decision making models

	Dependent Decision Variable: $\langle q(t)\rangle_j$			
	Coefficients			
	j=1 RET1	*j*=2 RET2	*j*=3 RET3	*j*=4 RET4
β_0^j	246.807	258.416	246.067	32.589
β_w^j	-0.945	-1.030	-0.952	-0.048
β_q^j	-0.033	0.180	0.035	0.455
β_d^j	-0.045	0.262	0.173	0.029
β_P^j	0	-0.001	-0.001	0.002
Adj. R^2	0.867	0.778	0.881	0.724
Significance level *p*	<0.01	<0.01	<0.01	<0.01

RET$_4$ appeared to share some of the suppliers' internal decision making logic. In contrast, the other three human retailers preferred to concentrate on the wholesale price *w* that was charged to them, although an external event to their own operation.

3.4 The Resulting Outcomes

We used Agent Based Simulation as the evaluation tool, because it provides a natural test-bed for modeling heterogeneous decision rules and interactions (North and Macal 2007, Gilbert 2008). Given the problem's small size and mostly for reasons of speed of model build and use, we developed the model in Excel VBA (Robinson 2004), although later we developed a version of the model in AnyLogic® Version 6.2.2 (XJ Technologies 2007). We obtained true customer demand instances via normal distribution variates truncated at zero, according to Barr and Sherrill's (1999) recommendations. We produced the variates by using the Mersenne-Twister pseudo-random number generator to ensure repeatability of results (Matsumoto and Nishimura1998).

At the beginning of any simulation run the decision makers did not hold sufficient information. As, however, the simulation progressed, the decision makers gradually built an understanding of the effect of their own decisions on their profit, which in turn led their corresponding decision models to converge to their respective steady state means. In order, thus, to avoid making inferences while the "initialization bias" phenomenon was still under effect (Pidd, 2004; Robinson, 2004; Law, 2007), we implemented the following run strategy: *i.* we established an estimate of the *warm-up* length amounting to 160 time periods, according to the MSER-5 method (White 1997, White and Spratt 2000) and *ii.* we ran each trial for 1,800 time periods, according to Banks et al's (2005) recommendation. In order, finally, to obtain accurate estimates of mean performances we replicated each trial for 100 times, following Hoad et al's (2007) replications algorithm.

Although we asked participants to determine their *w-* and *q-* decisions in the laboratory, what we were mainly interested in exploring was the effect of the interaction of their varying decision making strategies on the contract's efficiency. In order to objectively assess the participants' performance in comparison to their rationally optimizing counterparts, we also included the rationally optimizing supplier and retailer in the experimental design (*i.e.* their decision making strategies are given by mathematical expressions (1) and (2), respectively). In this regard, we considered the various decision making strategies as the treatment factors (F_1: supplier, F_2: retailer) with F_1 appearing at $s_1=4$ levels (SUP_i, *i*=1, 2, 3, *OPT*) and F_2 at $s_2=5$ levels (RET_j, *j*=1, 2, 3, 4, *OPT*). Since the resulting asymmetrical factorial design consists of a totally non-prohibitively high number of possible combinations ($s_1 \times s_2 =20$), we applied the full factorial 'two way layout' of all $F_1 - F_2$ possible combinations (Robinson 2000, Toutenburg 2002, Mukerjee and Wu 2006).

4 EXPERIMENTAL RESULTS

The key outcomes of interest that we report in this section are: *i.* the *w*-decisions, *ii.* the *q*-decisions and *iii.* the efficiency scores achieved. *Efficiency scores* are defined according to: $Eff. = \overline{P}_c / \overline{P}_{int}$, where $\overline{P}_c$ is the mean channel profit $\overline{P}_c = \overline{P}_s + \overline{P}_r$, and $\overline{P}_{int}$ is the mean profit that would have been realized by the *integrated newsvendor*. The latter orders q^{INT}:

$$q^{INT}(c) = F^{-1}(\frac{p+g-c}{p+g-s}) \tag{5}$$

in order to acquire his maximum expected profit: $\Pi_{int} = p \cdot E\{\min(q^{INT}, D)\} + s \cdot E\{(q^{INT} - d)^+\} - g \cdot E\{(D - q^{INT})^+\} - c \cdot q^{INT}$ (Cachon 2003). The only cost that the *integrated newsvendor* has to incur is his manufacturing cost $c=w^{INT}=50$, which gives rise to $q^{INT}(c) = 202.56$, according to (5).

4.1 The *w*-decisions

In line with previous empirical research on the *wholesale price contract* (*e.g.* Keser and Paleologo 2004, Loch and Wu 2007, Katok and Wu 2007) we would not expect human suppliers to charge prices that are consistent with the profit maximizing price *w**. This is what our first research hypothesis suggests.

Hypothesis 1 Human suppliers charge *w*-prices that are significantly different than the rationally optimizing supplier's price *w** ($w \neq w^*$)

In order to statistically accept the above alternative hypothesis we would need to reject the corresponding null hypothesis ($w=w^*$).

Table 3 presents all suppliers' steady state mean $\overline{w}$-decisions over $n = 100$ simulated replications for all 20 treatment combinations studied ($\overline{w} = \dfrac{\sum_{i=1}^{n} w}{n}$). Between parentheses () in *italics font* the standard deviation of observations among different replications is given, while between brackets [] in **bold font** the half widths of the corresponding 99% confidence intervals are provided. The reason that we based all our inferences on the low significance level of $\alpha = 0.01$ is because we preferred to reduce the probability of rejecting a true null hypothesis (*i.e.* committing a *Type I* error) as much as possible.

Table 3: The *w*-decisions

F_2 \ F_1	RET$_1$	RET$_2$	RET$_3$	RET$_4$	RET$_{OPT}$
SUP$_1$	228.78 (*0*) [**±0**]	228.78 (*0*) [**±0**]	228.78 (*0*) [**±0**]	228.78 (*0*) [**±0**]	228.78 (*0*) [**±0**]
SUP$_2$	199.79 *(0)* [**±0**]	199.79 *(0)* [**±0**]	199.79 *(0)* [**±0**]	199.79 *(0)* [**±0**]	199.79 *(0)* [**±0**]
SUP$_3$	143.85 (*0*) [**±0**]	143.85 (*0*) [**±0**]	143.85 (*0*) [**±0**]	143.85 (*0*) [**±0**]	143.85 (*0*) [**±0**]
SUP$_{OPT}$ (*w**)	247.86 (*0*) [**±0**]	247.86 (*0*) [**±0**]	247.86 (*0*) [**±0**]	247.86 (*0*) [**±0**]	240.83 (*0*) [**±0**]

We can see from Table 3 that human suppliers' *w*-decisions did not vary, when asked to interact with different retailers. The reason is that there was insufficient change in retailers' response quantities between different treatment combinations for suppliers to change their mean prices. As for SUP$_{OPT}$ he steadily charged *w**=247.86 monetary units (according to expression (1)), which is independent of the retailer's response. Table 3 also demonstrates that all above standard deviations are exactly equal to 0, turning all corresponding half-width 99% Confidence Intervals for all $\overline{w}$-decisions reported to zero. This is so because, suppliers ignored demand observations in their *w*-decisions, which is the only factor that varied across different replications.

From Table 3 we can also observe that SUP$_{OPT}$ charged the highest mean $\overline{w}$-prices. He was followed by SUP$_1$, whose, however, $\overline{w}$-decisions were statistically significantly lower than *w**. The reason that SUP$_1$ was the one that charged the highest prices is that for every *w(t)* decision he made he considered the least his previous decision *w(t-1)*, but did take into account to some extent the retailer's response quantity *q(t-1)*. In contrast, SUP$_3$ assigned the greatest importance to his previously charged *w*-prices in every new decision he made. This is why he charged distinctively lower $\overline{w}$-prices than the other suppliers. As a result, we can safely reject the corresponding to hypothesis 1 null hypothesis and, thus, accept hypothesis 1 (*w≠w**: at *p<0.01*), which demonstrates that, as expected, the assumptions of existing analytical research cannot be confirmed in the laboratory.

4.2 The *q*-decisions

In line with previous empirical research on the *wholesale price contract* (*e.g.* Keser and Paleologo 2004, Loch and Wu 2007, Katok and Wu 2007) we would not expect human retailers to order quantities that are consistent with the profit maximizing respective quantities *q**. This is exactly what our second hypothesis suggests.

Hypothesis 2 Human retailers place orders of *q*-quantities that are significantly different than the rationally optimizing retailer's respective quantities *q** (*q≠q**).

In order to statistically accept the above hypothesis we would need to reject the corresponding null hypothesis (*q=q**).

Table 4 presents all retailers' steady state mean $\overline{q}$-decisions over $n = 100$ simulated replications for all 20 treatment combinations studied. From the last column of Table 4 we can detect that the standard deviations of all observations and resulting half width 99% confidence intervals of all combinations of the rationally optimizing retailer amount to zero. The reason is again that the rationally optimizing retailer ordered quantities that followed expression (2) and, thus, did not depend at all on the demand observations that varied from one period to the next.

From Table 4 we can observe that when interacting with the most expensively charging suppliers, SUP$_{OPT}$ and SUP$_1$, it was RET$_4$ who ordered the highest quantities. When interacting with the remaining suppliers (SUP$_2$ and SUP$_3$) it was RET$_2$

who ordered the highest quantities. The reason is that RET_4 appeared to be almost indifferent to the supplier's *w*-prices (Table 2). In contrast, the wholesale price that RET_2 was charged is the most important factor that he took into account before making his $\overline{q}$-decision. The same also applied for RET_1 and RET_3. In spite of the similarities between RET_1, RET_2 and RET_3 underlying decision making strategies, their resulting $\overline{q}$-quantities were very different: RET_1 was the only one who constantly ordered quantities that were significantly lower than the rationally optimizing retailer; RET_2 ordered significantly higher quantities than the rationally optimizing retailer ($q^*(\overline{w})$); RET_3 ordered quantities that very closely followed the rationally optimizing retailer's order quantities: they were higher than $q^*(\overline{w})$, when asked to interact with the most expensively charging suppliers (SUP_{OPT} and SUP_1) and lower than $q^*(\overline{w})$, when asked to interact with the other two suppliers (SUP_2 and SUP_3).

Table 4: The *q*-decisions

F_2 ＼ F_1	RET_1	RET_2	RET_3	RET_4	RET_{OPT} ($q^*(\overline{w})$)
SUP_1	17.58 (*0.06*) [**±0.028**]	97.74 (*0.29*) [**±0.133**]	59.97 (*0.21*) [**±0.09**]	124.71 (*0.38*) [**±0.15**]	58.96 (*0*) [**±0**]
SUP_2	44.98 (*0.06*) [**±0.03**]	125.87 (*0.25*) [**±0.11**]	85.82 (*0.18*) [**±0.08**]	118.13 (*0.32*) [**±0.71**]	92.68 (*0*) [**±0**]
SUP_3	97.84 (*0.06*) [**±0.03**]	178.12 (*0.18*) [**±0.08**]	133.40 (*0.13*) [**±0.06**]	107.82 (*0.24*) [**±0.11**]	135.07 (*0*) [**±0**]
SUP_{OPT}	1.19 (*0.03*) [**±0.01**]	78.78 (*0.31*) [**±0.14**]	42.41 (*0.22*) [**±0.10**]	129.66 (*0.43*) [**±0.20**]	32.89 (*0*) [**±0**]

In summary, all human retailers RET_j (*j*=1, 2, 3, 4) ordered statistically significantly different quantities than the rationally optimizing retailer ($q^*(\overline{w})$). Therefore, we can safely reject the corresponding to hypothesis 2 null hypothesis and, thus, accept hypothesis 2 ($q \neq q^*$: at *p*<0.01). This demonstrates that, as expected, the assumptions of existing analytical research cannot be confirmed in the laboratory.

4.3 Efficiency Scores

Since our ultimate objective in this research is to identify whether and under which particular interactions efficiency of the *wholesale price contract* can be achieved, we seek for the treatment combinations for which the resulting efficiency score can become equal to one. To this end, the third research hypothesis tests whether the efficiency scores achieved by all human supplier-retailer interactions are strictly less than one, as expected. The human supplier-retailer interactions for which this hypothesis is rejected will illustrate the conditions under which the *wholesale price contract* can become efficient, that is provided there are any. Since the *wholesale price-only* contract has long been analytically proven inefficient, ensuring that the third hypothesis cannot be rejected for the interaction of the rationally optimizing supplier and retailer would provide some evidence for the validity of our model, according to the 'black-box' approach (Robinson 2004).

Hypothesis 3: The efficiency score achieved by each supplier-retailer interaction is strictly less than 1. (*Eff<1*).

In order to statistically accept the above hypothesis we would need to reject the corresponding null hypothesis (*Eff=1*). Table 5 presents the mean efficiency scores over *n*=100 replications achieved by all 20 treatment combinations studied. It becomes evident that the efficiency score achieved, when the rationally optimizing supplier (SUP_{OPT}) and the rationally optimizing retailer (RET_{OPT}) interacted with each other, is strictly less than one. Hence, we can safely reject the corresponding to the third hypothesis null hypothesis and, thus, accept the third hypothesis for the rationally optimizing supplier and retailer interaction (*Eff*<1: at *p*<0.01). Although the performance of this particular interaction was very poor, it is interesting to note that there were other interactions that led to even worse efficiency scores: The interaction of RET_1 with SUP_{OPT} gave rise to the lowest efficiency score, followed by the interaction of RET_1 with SUP_1. It is interesting that in both of these cases, RET_1 came into play. This is so because RET_1 was the only retailer who systematically ordered less than the already low quantities of the rationally optimizing retailer q^*.

From Table 5 we can also see that all other interactions generated improved efficiency scores. Still, we can only identify one interaction with a 99% efficiency score confidence interval that included the value of one: the interaction of SUP_3 with RET_2. For this reason, we accept the third hypothesis (*Eff<1:* at *p*<0.01) for all human supplier-retailer combinations, except for SUP_3 with RET_2.

Table 5: The efficiency scores

F_2 \ F_1	RET$_1$	RET$_2$	RET$_3$	RET$_4$	RET$_{OPT}$
SUP$_1$	0.132 (*0.055*) [**±0.026**]	0.812 (*0.081*) [**±0.037**]	0.572 (*0.110*) [**±0.050**]	0.911 (*0.059*) [**±0.027**]	0.572 (*0.110*) [**±0.051**]
SUP$_2$	0.428 (*0.120*) [**±0.055**]	0.918 (*0.060*) [**±0.027**]	0.756 (*0.089*) [**±0.041**]	0.892 (*0.064*) [**±0.029**]	0.798 (*0.084*) [**±0.039**]
SUP$_3$	0.822 (*0.080*) [**±0.037**]	0.998 (*0.020*) [**±0.009**]	0.941 (*0.053*) [**±0.024**]	0.857 (*0.072*) [**±0.033**]	0.946 (*0.051*) [**±0.024**]
SUP$_{OPT}$	0.004 (*0.007*) [**±0.003**]	0.705 (*0.100*) [**±0.044**]	0.387 (*0.106*) [**±0.049**]	0.923 (*0.055*) [**±0.025**]	0.27 (*0.001*) [**±0**]

This implies that when the least expensively charging supplier SUP$_3$ interacted with RET$_2$, then the *wholesale price contract* became efficient. RET$_2$ was found to be the retailer who placed the highest order quantities, when interacting with SUP$_3$. It is interesting that although RET$_2$ and RET$_4$ placed orders of overall comparable sizes, it was RET$_2$ who generated efficiency. This can be attributed to RET$_2$ being very conscious of the interacting supplier's *w*-prices and subsequent ability to exploit the low prices offered to him. The latter further demonstrated that suppliers may be the Stackelberg leaders, but it was the retailers who tended to have a stronger impact on overall efficiency scores. We consider this worthy of further exploration.

It is really interesting that, in contrast to our prior expectations, we found the *wholesale price contract* to be efficient. Although we found all suppliers and retailers to make decisions that systematically deviated from their rationally optimizing counterparts', when combined, they gave rise to an efficient interaction. Therefore, it was from the interaction between the suppliers' and retailers' distinct decision making strategies that the *wholesale price contract's* efficiency could 'emerge' (Casti 1997, Holland 1998, Axelrod 2005). We consider this result a valuable addition to the existing empirical research on supply chain contracts.

5   CONCLUSIONS

This paper attempts to provide an additional explanation to the *wholesale price contract's* wide practical popularity, that is despite its long proven inefficiency. It adopts the view that the contract's potential *efficiency* can emerge from the decision makers' underlying 'boundedly rational' preferences and cognitive limitations. It is for this reason that we asked human subjects to act as suppliers or retailers in laboratory-enabled simulation games of the newsvendor setting. For the same reason, we also did not eliminate their individualities via averages but, instead, treated them as the main experimental factor of analysis. In order to identify *efficient* human supplier-retailer interactions, we used Agent Based Simulation as the evaluation tool, because of its appropriateness to model heterogeneous decision makers. Each participant was assumed to make decisions that followed the respective model that was fitted to the decisions previously made in the laboratory.

We gained a qualitative understanding of each participant's underlying decision making reasoning, in addition to a quantitative insight on the efficiency of pairing the decision makers in every possible combination. We found all participants' decisions to significantly deviate from each other's. We also found them to fail to resemble to their rationally optimizing counterparts' decisions. Nevertheless, one of these possible combinations' aggregate profit approximated the *integrated newsvendor's first-best case* channel profit, leading, thus, to near global *efficiency* (Lariviere 1999, Tsay, Nahmias, and Agrawal 1999). Therefore, this paper provides evidence as to the *wholesale price contract's* potential *efficiency* in a practical setting. We see this empirical divergence as the result of the interplay between participants' varying individual preferences and cognitive limitations.

As a result, we consider it worthwhile to pursue a more systematic investigation of the different types of interactions that can exist in the supply chain context. For our part, our research will continue by investigating a more complex supply chain with three interacting decision makers through the use of the Beer Distribution Game (Sterman 2000).

REFERENCES

Axelrod, R. 2005. Advancing the art of simulation in the social sciences. In *Handbook of Research on Nature Inspired Computing for Economy and Management*, ed. J. Rennard. Idea Group, Hersey.

Banks, J., J. Carson, B. Nelson, and D. Nicol. 2005. *Discrete-Event Simulation*. 4th ed. Upper Saddle River: Pearson Prentice Hall.

Barr, D., and T. Sherrill. 1999. Mean and Variance of Truncated Normal Distributions. *The American Statistician* 53 (4): 357-361.

Bendoly, E., K. Donohue, and K. Schultz. 2006. Behavior in operations management: Assessing recent findings and revisiting old assumptions. *Journal of Operations Management* 24: 737-752.

Benzion, U., Y. Cohen, R. Peled, and T. Shavit. 2008. Decision-making and the newsvendor problem: an experimental study. *Journal of the Operational Research Society* 59: 1281-1287.

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

Bostian, A., C. Holt, and A. Smith. 2008. The Newsvendornewsvendor pull-to-center effect: Adaptive learning in a laboratory experiment. *Manufacturing Service Operations Management* 10(4): 590-608.

Bowman, E. 1963. Consistency and optimality in managerial decision making. *Management Science* 9 (2): 310-321.

Cachon, G. 2003. Supply chain coordination with contracts. In *Supply chain management: Design, coordination and operation*, ed. A. Kok, and S. Graves, 229-341. Amsterdam: Elsevier.

Carslon, J., and T. O' Keefe. 1969. Buffer stocks and reaction coefficients: An experiment with decision making under risk. *The Review of Economic Studies* 36(4): 467-484.

Chopra, S., and P. Meindl. 2007. *Supply chain management: Strategy, planning and Operation*. Upper Saddle River: Pearson Prentice Hall.

Croson, R., and K. Donohue. 2002. Experimental economics and supply-chain management. *Interfaces* 32 (5): 74-82.

Croson, R., and K. Donohue. 2006. Behavioural causes of the bullwhip effect and the observed value of inventory information. *Management Science* 52(3): 323-336.

Fisher, M., and A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* 44 (1): 87-99.

Gilbert, N. 2008. *Agent-Based Models*. Los Angeles: Sage Publications.

Gino, F., and G. Pisano. 2006. Toward a theory of behavioral operations. *Manufacturing Service Operations Management* 10 (4): 676-691.

Hair, J., W. Black, R. Anderson, and R. Tatham. 2006. *Multi-variate data analysis*. 6th ed. New Jersey: Pearson Prentice Hall.

Ho, T.-H., and J. Zhang. 2008. Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter?. *Management Science* 54(4): 686-700.

Hoad, K., S. Robinson, and R. Davies. 2007. Automating DES output analysis: How many replications to. In *Proceedings of the 2007 Winter Simulation Conference*, ed. S. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. Tew, and R. Barton, 505-512.

Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgement under uncertainty: Heuristics and biases*. New York: Cambridge University Press.

Katok, E., and D. Wu. 2007. Contracting in supply chains: A laboratory investigation. *Penn State University Working Paper*.

Keser, E., and G. Paleologo. 2004. Experimental investigation of supplier-retailer contracts: The wholesale price contract. *CIRANO Working Paper*.

Lariviere, M. 1999. Supply chain contracting and coordination with stochastic demand. In *Quantitative Models For Supply Chain Management*, ed. S. Tayur, R. Ganeshan, and M. Magazine, 233-268. Boston: Kluwer Academic Publishers.

Lariviere, M., and E. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing Service Operations Management* 3(4): 293-305.

Law, A. 2007. *Simulation Modeling and Analysis*. New York: McGraw Hill.

Loch, C., and Y. Wu. 2008. Social preferences and supply chain performance: An experimental study. *Management Science* 54(11): 1835-1849.

Matsumoto, M., and T. Nishimura. 1998. Mersenne-Twister: A 623-dimensionally equi-distributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation* 8 (1): 3-30.

Mukerjee, R., and C. Wu. 2006. *A Modern Theory of Factorial Designs*. New York: Springer.

North, M., and C. Macal. 2007. *Managing business complexity*, New York: Oxford University Press.

Pasternack, B. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Science* 4 (2): 166-176.

Pidd, M. 2004. *Computer simulation in management science*. Chichester: John Wiley and Sons.

Remus, W. 1978. Testing Bowman's managerial coefficient theory using a competitive gaming environment. *Management Science* 24(8): 827-835.

Robinson, G. 2000. *Practical Strategies for Experimenting*. Chichester: John Wiley and Sons.

Robinson, S. 2004. *Simulation: The practice of model development and use*, Chichester: John Wiley & Sons.

Robinson, S., T. Alifantis, J. Edwards, J. Ladbrook, and A. Waller. 2005. Knowledge-based improvement: Simulation and artificial intelligence for identifying and improving human decision-making in an operations system. *Journal of the Operational Research Society* 56: 912-921.

Schweitzer, M., and G. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science* 46 (3): 404-420.

Simchi-Levi, D., P. Kaminsky, and E. Simchi-Levi. 2008. *Designing and managing the supply chain: Concepts, strategies and case studies*. 3rd ed. New York: McGraw-Hill Irwin.

Simon, H. 1957. A behavioural model of rational Choice. In *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behaviour in a Social Setting*, New York: Wiley.

Simon, H. 1996. *The sciences of the artificial*. 3rd ed. Cambridge: The MIT Press.

Son, J., and C. Sheu. 2008. The impact of replenishment policy deviations in a decentralized supply chain *International Journal of Production Economics* 113: 785-804.

Spengler, J. 1950. Vertical integration and anti-trust policy. *Journal of Political Economy* 58(4): 347-352.

Sterman, J. 1989. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science* 35(3): 321-339.

Sterman, J. 2000. *Business dynamics: Systems thinking and modeling for a complex world*. Boston: Irwin/McGraw-Hill.

Toutenburg, H. 2002 *Statistical Analysis of Designed Experiments*. New York: Springer.

Tsay, A., S. Nahmias, and N. Agrawal. 1999 Modeling supply chain contracts: A Review. In *Quantitative Models For Supply Chain Management*, ed. S. Tayur, R. Ganeshan and M. Magazine, 233-268. Boston: Kluwer Academic Publishers.

Weisberg, S. 2005. *Applied linear regression*. 3rd ed. New Jersey: John Wiley and Sons.

White, P. 1997. An effective truncation heuristic for bias reduction in simulation output. *Simulation* 69 (6): 323-334.

White, P., and S. Spratt. 2000. A comparison of five steady-state truncation heuristics for simulation. In *Proceedings of the 2000 Winter Simulation Conference*, ed. J. Joines, R. Barton, K. Kang, and P. Fishwick, 755-760.

Williamson, O. 1981. The economies of organization: the transaction cost approach. *American Journal of Sociology* 87 (3): 548-577.

Wu, D., and E. Katok. 2006. Learning, communication and the bullwhip effect. *Journal of Operations Management* 24: 839-850.

XJ Technologies. 2007. AnyLogic® v6.2.2. Available from the world wide web: <http://www.xjtek.com/>

AUTHOR BIOGRAPHIES

STAVRIANNA DIMITRIOU is a Doctoral Researcher at Warwick Business School. She holds a Diploma of Mechanical Engineering from National Technical University of Athens, Greece and a MSc in Management Science and Operational Research from Warwick Business School. Her research interests include simulation modeling, agent based simulation, supply chain coordination, contracts, experimental research. She is a member of the OR Society. Her e-mail address is: <Stavrianna.Dimitriou06@wbs.ac.uk> .

STEWART ROBINSON is a Professor of Operational Research at Warwick Business School and Associate Dean for Specialist Masters Programmes. He holds a BSc and PhD in Management Science from Lancaster University. Previously employed in simulation consultancy, he supported the use of simulation in companies throughout Europe and the rest of the world. He is author/co-author of three books on simulation. His research focuses on the practice of simulation model development and use. Key areas of interest are conceptual modeling, model validation, output analysis and modeling human factors in simulation models. His email address is <stewart.robinson@warwick.ac.uk> and his Web address is <www.btinternet.com/~stewart.robinson1/sr.htm>.

KATHY KOTIADIS is an assistant professor at Warwick Business School (UK) and co-chair of the UK Simulation Special Interest Group. She holds a BSc and PhD from the University of Kent. Her main research interests include health service modelling and combining problem structuring approaches with DESM. Her email address is <Kathy.Kotiadis@wbs.ac.uk>.