This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Addressing the sample size problem in behavioural operational research: Simulating the newsvendor problem

PLEASE CITE THE PUBLISHED VERSION

<http://dx.doi.org/10.1057/s41274-016-0016-3>

PUBLISHER

© Operational Research Society. Published by Palgrave Macmillan

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This work is made available according to the conditions of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) licence. Full details of this licence are available at:
<https://creativecommons.org/licenses/by-nc-nd/4.0/>

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Robinson, Stewart, Stavrianna Dimitriou, and Katherine Kotiadis. 2016. "Addressing the Sample Size Problem in Behavioural Operational Research: Simulating the Newsvendor Problem". Loughborough University.
<https://hdl.handle.net/2134/21840>.

Addressing the Sample Size Problem in Behavioural Operational Research: Simulating the Newsvendor Problem

Stewart Robinson, Stavrianna Dimitriou, Kathy Kotiadis

Abstract

Laboratory based experimental studies with human participants are beneficial for testing hypotheses in behavioural operational research. However, such experiments are not without their problems. One specific problem is obtaining a sufficient sample size, not only in terms of the number of participants, but also the time they are willing to devote to an experiment. In this paper we explore how agent-based simulation can be used to address the sample size problem and demonstrate the approach in the newsvendor setting. The decision-making strategies of a small sample of individual decision-makers are determined through laboratory experiments. The interactions of these suppliers and retailers are then simulated using an agent-based simulation (ABS) to generate a large sample set of decisions. With only a small number of participants, we demonstrate that it is possible to produce similar results to previous experimental studies that involved much larger sample sizes. We conclude that agent-based simulation provides the potential to extend the scope of experimental research in behavioural operational research.

Key Words: behavioural operational research, experimental research, agent-based simulation (ABS), supply chain management, newsvendor problem

Addressing the Sample Size Problem in Behavioural Operational Research: Simulating the Newsvendor Problem

1. Introduction

With the developing interest in behavioural operational research (OR) (Hämäläinen et al, 2013) there is a need to adopt experimental research methods such as those long used in behavioural psychology, and more recently adopted in fields such as behavioural economics and behavioural operations management. Laboratory experiments are beneficial for testing hypotheses while controlling for the conditions under which participants work; something that cannot easily be achieved in real world experiments. As an example that has been extensively investigated, work in behavioural operations management uses laboratory experiments to study the actual performance of supply chain contracts. In these studies decision-makers, normally students, make decisions in supply chain games under different contractual arrangements (e.g. wholesale price, buyback and revenue sharing contracts). These studies consistently demonstrate that human decision-makers perform very differently to the rational-optimising decision-makers assumed in mathematical algorithms. Katok and Wu (2009), for instance, demonstrate that the improvement obtained from using the buyback and revenue sharing contracts are not as great as expected in the presence of ‘real’ decision-makers.

Although valuable, laboratory experiments are not without their problems. One difficulty arises in the selection of participants. Convenience often leads to the use of students. Whilst there is evidence that they can outperform managers on experimental tasks (Bakken et al, 1994), is the use of students always appropriate? In operational research our ultimate aim is to work with real decision-makers working on real problems in order to bring about an improvement. In this case, student participants would almost certainly not suffice. Beyond the choice of participants, Wynder (2004) identifies a series of factors that can affect a participant’s performance in laboratory experiments: providing a problem that interests and engages the participants; ensuring that participants have the necessary skills for the task; establishing criteria for participants to self-assess and validate their performance during the task; use of extrinsic motivators, especially whether monetary incentives are beneficial.

Obtaining a sufficient sample size can be particularly problematic at three levels: the number of participants, the length of time the participants spend on the task (number of decisions made), and, where required, the number of pairings of participants. It can be very challenging and time-consuming to obtain a large number of participants for a study, even if the participants are students. Also, large sample sizes can be costly to obtain when financial incentives are provided. Meanwhile, if the experiment is to be performed with real decision-makers, then the sample size is dictated by the

size of the available workforce. For instance, Robinson et al. (2012) only managed to obtain a sample of eight decision-makers when working in a Ford manufacturing plant.

However, the sample size problem is not only manifested in the number of participants, but also in the length of time they work on the task. Some experimental research requires the participants to make a series of consecutive decisions. There are limits to the time that a participant is willing to devote to an experimental task and also limits to their level of concentration. Hence, the size and quality of a sample of decisions from a participant is constrained. Added to this, it is sometimes useful to ask participants to play against each other in a ‘gaming’ environment, for instance, in the ‘beer game’ (Senge, 1990). This further adds to the complexity of the experiment and the number of groups of partnerships that can be tested. Given differences in decision-making strategies, it could be useful to play every participant against every other participant, but practically impossible to do so due to the time commitment involved.

In this paper we explore the use of agent-based simulation (ABS) to address the issue of sample size in experimental research. Our focus is on the case where there are only a few participants, as would often occur in real decision-making situations. We then use the ABS to address the second and third sample size issues: the length of time the participants spend on the task and the number of pairings of participants. Because extensive experimental research on supply chain contracts already exists, we base our work in this problem domain, focusing specifically on the well studied newsvendor problem and the wholesale price contract (Whitin, 1955). The difference with our approach is that we only require a small number of participants. By learning the decision-making strategies of these participants, we are then able to simulate every (retailer and supplier) participant playing with each other over a long-run game. Our ABS generates similar results to previous experimental studies which require much larger numbers of participants. We believe that this feature of requiring fewer participants makes the approach attractive, especially in cases where the input of only a few (real world) decision-makers can be obtained.

Because the focus of this paper is on the newsvendor problem, the next section outlines the newsvendor problem, the wholesale price contract and previous experimental research on this and related problems. In section 3, we provide an overview of the design of the study. We describe the implementation of the study, including the development of the ABS, in section 4. In section 5, we report the results from the ABS. In section 6 we discuss the benefits of the ABS, the limitations of the work and the potential for future research.

2. The Newsvendor Problem and Prior Experimental Research

As already noted, supply chain contracts have been extensively studied using laboratory-based experiments. Examples include the work of Schweitzer and Cachon (2000), Croson and Donohue (2006), Loch and Wu (2008), Wu (2013) and Schiffels et al (2014). These studies typically use students as participants, sometimes working with an automated supply chain partner and sometimes working with another student. Financial inducements are often provided both for participating in the study and also in relation to performance on the task.

In this section we briefly review this work and its key findings, as well as identify the sample sizes these studies employ. First we set out the theoretical background to the newsvendor problem and the wholesale price contract.

2.1 Theoretical Background: The Newsvendor Problem and Wholesale Price Contract

Consider the typical newsvendor setting under the wholesale price contract, as illustrated in Figure 1 (Whitin, 1955). In advance of each time period t , the supplier specifies the wholesale price w that he/she wishes to charge to the retailer. In response, the retailer chooses an order quantity q . It is assumed that the supplier instantaneously delivers the full order quantity to the retailer. The retailer is then responsible for satisfying the market demand d , which is stochastic in nature. The product only lasts for one selling season and so inventory is not carried over from one time period to the next. Therefore, based on the quantity ordered from the supplier in the current time period, the retailer either fully satisfies the market demand (d), or sends the maximum amount of product available (q). The retailer sells each unit of product at price p , while the supplier incurs a unitary production cost c . For each unit of demand the retailer does not satisfy, the retailer incurs a goodwill penalty cost of g .

Figure 1 The decentralised newsvendor problem, under the wholesale price contract

Under the assumption that the two supply chain partners are self-interested, rational-optimising decision-makers, they would make the following decisions (denoted by $*$). In every time period t , a rational-optimising supplier would charge a wholesale price w^* , so as to maximise his/her expected profit P_s (Cachon, 2003):

$$P_s^* = \max\{P_s(w, q)\} = (w^* - c)q_r^* \quad (1)$$

In (1), q_r^* represents the order quantity that the rational-optimising retailer would place in response to the price w^* , such that $q_r^* = \arg[\max P_s(w, q)]$. Since the rational-optimising retailer is exclusively interested in maximising his/her expected profit P_r in response to the price w^* , he/she would order q_r^* products (Lariviere, 1999; Lariviere and Porteus, 2001; Cachon, 2003):

$$q_r^* = F^{-1}\left(\frac{p + g - w^*}{p + g}\right) \quad (2)$$

The resulting aggregate channel profit $P_c = P_s + P_r$ where:

$$P_r = (p + g)S(q) - g\mu - wq \quad (3)$$

$S(q)$ is the expected sales of the retailer (Cachon, 2003) and μ is the mean level of the demand.

An alternative strategy would be to centralise the decision-making, creating an integrated supply chain. The rational-optimising integrated newsvendor would order as many products (q_{int}^*) so as to maximise his/her expected profit P_{int} (Khouja, 1999; Lariviere, 1999):

$$q_{int}^* = F^{-1}\left(\frac{p + g - c}{p + g}\right) \quad (4)$$

As an example, which we use as the basis of our experimental study in the rest of this paper, assume that $p=250$, $c=50$, $g=1$, and that customer demand follows the truncated at zero normal distribution with $\mu_n=140$ and $\sigma = 80$ (where μ_n refers to the mean of the non-truncated normal). Under these circumstances the rational-optimising integrated newsvendor would order $q_{int}^* = 209.87$ units and only incur the production cost $w_{int}^* = c$, leading the entire channel to the first-best case aggregate profit of $P_{int}^* = 24,098.74$ (Halkos and Kevork, 2011).

Under decentralised control, the rational-optimising supplier would charge $w^*=174.75$, and the rational-optimising retailer would place an order of $q^*=105.18$. This would result in individual profits of $P_s^* = 13,120.77$ and $P_r^* = 4,742.02$ respectively, and an aggregate channel profit of $P_c^* = 17,862.79$. As expected, the profits under decentralised control are much lower than in the integrated case. The efficiency score for decentralised control is $Eff = \frac{P_c^*}{P_{int}^*} = 0.741 < 1$, demonstrating the theoretical inefficiency of the wholesale price contract.

2.2 Findings from Experimental Studies on Supply Chain Contracts

In reality human decision-makers take different decisions to their rational-optimising counterparts, leading to very different overall performances in the newsvendor setting from those that are theoretically expected. Indeed, previous experimental research has confirmed that human newsvendors' order quantities and price decisions consistently diverge from the rational-optimising levels predicted by the standard normative models (Schweitzer and Cachon, 2000; Keser and Paleologo, 2004; Loch and Wu, 2008; Katok and Wu, 2009; Kalkanci et al, 2011; Elahi et al, 2013; Wu, 2013; Moritz et al, 2013). For instance, there is experimental evidence of the 'pull-to-centre' effect, in which human retailers order too few of high profit products and too many of low profit products (Schweitzer and Cachon, 2000; Bostian et al, 2008). This is contradicted, however, by Elahi et al (2013), who find that human retailers order considerably more than the optimum quantity of a high profit product. In the experimental studies by Keser and Paleologo (2004), Loch and Wu (2008) and Wu (2013), the human suppliers charge lower prices than would a rational-optimising supplier.

With respect to supply chain efficiency, in the experimental study by Keser and Paleologo (2004) the wholesale price contract with human decision-makers leads to a very similar efficiency to that predicted for the rational-optimising decision-makers. However, Katok and Wu (2009) find that the wholesale price contract performs better (higher efficiency) in their laboratory experiment than theoretically predicted due to an initial over ordering behaviour; although they do note that this may have been due to the specific game parameters and so this is not necessarily a general result. They also show that the buy-back and revenue sharing contracts do not perform as well, in terms of efficiency, in a laboratory setting as theoretically predicted.

Unlike Keser and Paleologo, Katok and Wu control for social preference by studying suppliers and retailers separately. They believe that Keser and Paleologo's findings may have emerged because their results are subject to social preference, that is, decision-makers are not solely interested in their own profit, but include social considerations in their decision-making, such as, reciprocity, fairness,

status seeking and group identity (Loch and Wu, 2008; Wu and Niederhoff, 2014). To address this potential bias, Loch and Wu (2008) perform an experimental study to investigate the impact of social preference on supply chain performance. They show that relationship preference improves system efficiency, while status preference reduces efficiency.

Recent work has focused on providing explanations for deviations in performance between human decision-makers and theoretical predications, with an emphasis on developing behavioural models. Moritz et al (2013) investigate the relationship between cognitive reflection (Frederick, 2005) and decision-making in the newsvendor context through an experimental study involving students and supply chain managers. They find that cognitive reflection is a better predictor of performance than level of education or real world experience. Kalkanci et al (2011) adopt an experience-weighted attraction learning model to data from their experimental study. The model shows that as contract complexity increases, humans rely more on simple heuristics for decision-making. Meanwhile, through behavioural modelling, Becker-Peth et al (2013) show that decision-makers anchor on mean demand, aim to avoid losses and place different values on alternative income streams. They argue that contracts should be designed to take these factors into account.

These studies show varying performance levels of human decision-makers in terms of supply chain efficiency. What they demonstrate is that human decision-makers can perform at least as well as their rational-optimising counterparts when working under the wholesale price contract, but they fall short of obtaining the aggregate profit that is achievable in a fully coordinated supply chain.

2.3 Sample Size in Experimental Studies on Supply Chain Contracts

We now explore the approach taken in the literature to the three sample size problems, namely: number of participants, number of decision rounds and combinations of participants playing against each other. The reported studies generally involve over 100 participants. For instance, Bolton and Katok (2008) used 234, mostly undergraduate, students; Loch and Wu (2008) worked with 168 subjects – students of varying backgrounds; Katok and Wu (2009) involved 200 participants; de Véricourt et al (2013) worked with 102 MBA students; and Schiffels et al (2014) have a sample of 148 students. There are some studies that use a smaller number of participants, for instance, Schweitzer and Cachon (2000) use 34 and then 44 participants in their two experiments.

In terms of the number of decisions participants are asked to make, this is normally quite small; at most 30. For example, Schweitzer and Cachon (2000), Keser and Paleologo (2004) and Schiffels et al (2014) all work with 30 decisions. Meanwhile, de Véricourt et al (2013) asked participants to play

only 20 rounds and Loch and Wu (2008) obtained only 15 decisions from each participant. Bolton and Katok (2008) and Wu (2013) do record more extensive data sets, asking participants to make 100 decisions as retailers, and Katok and Wu (2009) asked their participants to play for 200 rounds.

Most of the reported studies do not involve multiple participants playing against one another. Among the few studies that do are Keser and Paleologo (2004), Loch and Wu (2008), Ho and Zhang (2008) and Wu (2013). In each case, the experiment only involves a single pairing of players and it does not involve observing participants working with different partners.

3. Study Approach

In this study a small number of participants (seven) act as suppliers and retailers in the newsvendor setting. We use regression models to represent their decision-making strategies and then model those strategies in an ABS. This allows us to play every retailer participant in a pairing with every supplier participant. Our aim is to determine whether the use of ABS can address the sample size problem by enabling us to model multiple partnerships over a large number of decisions, based on data collected from a small sample of participants. In so doing, we compare our results with those from earlier studies that involve many more participants, but do not employ ABS.

Figure 2 presents the four stages in the experimental study that we performed: understanding the decision-making process, conducting the gaming sessions, fitting the decision-making strategies, and running the ABS model. This approach is based on the knowledge based improvement methodology of Robinson et al (2005). Supplier and retailer simulation games are designed out of stage 1. These games are then run with human agents in stage 2 to generate separate sets of decision-making data for each agent. In stage 3 those data sets are used to create regression based decision models, which are then run with a simple Excel-VBA agent based model of the supplier-retailer interactions (stage 4). The pricing and ordering decisions, and profits achieved, are recorded as the key outcomes. For this study, the parameters used for the newsvendor problem are those set out in the example presented in section 2.

Figure 2 The experimental study: developing and using the agent-based simulation

The purpose of stage 1 is to understand the factors (‘decision attributes’) which inform the decisions (‘decision variables’) that the two agents (supplier and retailer) make. We know from the structure of the newsvendor problem that the supplier makes decisions about the wholesale price and in response the retailer determines the order quantity. To determine the potential decision attributes a series of pilot sessions were run with volunteers. In each session the volunteers made their decisions with full information available to them, that is, all attribute values from every round. The sessions were followed by interviews in which we asked which attributes the participants had used to make their decisions. In this way we were able to isolate a set of potential decision attributes.

Stage 2 consists of gaming sessions based on the decision variables and attributes identified in stage 1. These sessions were performed separately with each participant. A participant assumed a particular role (supplier or retailer) and proceeded to make a series of decisions (enters a decision variable) in response to a predetermined set of scenarios (set of decision attributes). The same scenarios were presented to each supplier participant, and another set of scenarios were used with each retail participant. The decisions made by each participant in each round were recorded, creating a separate dataset of decisions variables and decision attributes for each participant.

Each participant is assumed to follow his/her own decision-making strategy, and so there is a unique relationship between each participant’s decision variables and their decision attributes. In order to determine each participant’s specific decision-making strategy, separate multiple regression models

were fitted in Stage 3 using the datasets of decision variables (dependent variable) and decisions attributes (independent variables).

The ABS model of the newsvendor problem was then developed around the decision variables (inputs) and decision attributes (outputs). The model, developed in Excel-VBA, includes these inputs and outputs, and the sequence in which material, funds, and information are moved at each time step as described in Figure 1. The decision models for each supplier and retailer can be implanted into the model so as to represent any combination of supplier-retailer interaction. Excel-VBA was chosen, rather than specialist ABS software (e.g. NetLogo, Repast or AnyLogic), because the structure of the model is relatively simple. Indeed, VBA is only required to control the multiple replications of the model runs.

Finally, in stage 4, the model was run for an extended period (1,800 time periods and 100 replications) with different combinations of supplier and retailer decision-making strategies, as represented by the multiple regression models. This enabled the combined performance of each supplier working with each retailer to be predicted in terms of price and quantity decisions, and in terms of the profit attained.

4. Developing and Using the Agent-Based Simulation (ABS)

We now describe, in detail, how each stage of the study was implemented, including the details of the regression based decision models of the human suppliers and the retailers.

4.1 Stage 1: Understanding the Decision-Making Process

From the pilot sessions and interviews we ascertained that the human suppliers ($i = SUP$) base each period's wholesale price decision $w(t)$ on the following three decision attributes:

- i.* The previously charged wholesale price $w(t - 1)$
- ii.* The previously placed order quantity $q(t - 1)$
- iii.* The previously realised profit $P_s(t - 1)$

Meanwhile, the retailers ($j = RET$) made their order quantity decisions $q(t)$ taking into consideration:

- i.* The currently charged wholesale price $w(t)$
- ii.* The last period's order quantity $q(t - 1)$

- iii. The previously observed demand $d(t - 1)$
- iv. The previously realised profit $P_r(t - 1)$

Despite the availability of a full history of decision attributes, in the interviews the participants stated that they ignored all information except for that pertaining to the previous ($t - 1$) or current (t) round. This reliance on recent information relates to the individual bias of ‘immediacy’ identified by Camerer (1995) and Loch and Wu (2007). It also typifies ‘salience’ in which decision-makers are selective about the information they use.

Those that acted as suppliers explained that they had relied heavily on the retailers’ previous order quantities for their price decisions. Meanwhile the retailers had relied on the previous round’s demand for their decisions. The reason was that they could not predict with certainty the incoming order quantities or demand, respectively. Camerer (1995) and Loch and Wu (2007) identify this tendency to rely on the information that is available as a means for dealing with the underlying uncertainty in the decision situation.

The participants reported difficulty in understanding how their decisions would affect their profits and the system’s overall performance. Therefore, in order to make their decisions, they used their profit from the previous round. This use of simple heuristics arises from the complexity that is inherent within the newsvendor problem (Kalkanci et al, 2011).

Therefore, supplier i ’s decision function is expressed as:

$$\langle w(t) \rangle_i = f_i[w(t-1), q(t-1), P_s(t-1)] \quad (5)$$

and retailer j ’s decision function as:

$$\langle q(t) \rangle_j = f_j[w(t), q(t-1), d(t-1), P_r(t-1)] \quad (6)$$

4.2 Stage 2: The Gaming Sessions

Volunteers were recruited from a pool of graduate students at the University of Warwick. The only requirement set was that all participants had received formal classroom training in the newsvendor problem prior to the experiment as part of their curriculum. The reason is that an earlier empirical

study has confirmed the superior performance of well-trained students acting as newsvendors compared to experienced supply chain managers (Bolton et al, 2008).

Volunteers were randomly assigned to play either the role of the supplier or the retailer against an automated retailer or supplier, respectively. They worked with a computer interface, written in Excel VBA, that simulated the interacting partner's responses. Provision of automated players may not have been as realistic as direct interaction of human suppliers and retailers, but it mitigated the effects of social preference and reputation, such as, for example, players' possible concern regarding fairness, reciprocity, status seeking and group identity (Loch and Wu, 2008; Katok and Wu, 2009).

In total, four participants acted as retailers (denoted RET_1 to RET_4) and three as suppliers (denoted SUP_1 to SUP_3). Written instructions on the required task were distributed to all participants well in advance of the allocated session so that they could familiarise themselves with the task and the software as quickly as possible. The instructions informed them that the product under study was a perishable widget with random customer demand. They were also made aware that each round's demand was independent of any previous round's, but they were not informed about the exact type of distribution that the customer demand followed. This follows Bearden and Rapoport's (2005) contention that decision-makers are unlikely to know the nature of the 'operative' distribution in a decision situation. Although the instructions asked participants to make decisions that, to the best of their knowledge, would maximise supply chain profit, they were only supplied with information on their own profit at the end of each round. As such, they could only focus on their own profit.

In case the participants required any clarifications, they could address questions both before the start of the session and during its course. Nevertheless, the game could not be re-started at any time. We ran the game for 50 consecutive rounds. In order to give participants some time to get used to their new roles, the first 10 rounds were used for practice. This gave a total of 40 periods of actual data for each participant, providing at least 10 samples per decision attribute. This complied with the minimum sample size recommended by Weisberg (2005) and Hair et al (2006). After every period participants received feedback on their previous decisions and their realised profit; the retailer also received feedback on the previous round's demand. The participants were not aware of the session's duration so that end-of-game effects could be eliminated (Schweitzer and Cachon, 2003; Steckel et al, 2004; Loch and Wu, 2008).

All participants acting as suppliers were asked to play against the same automated retailer. In response to the wholesale price set by the human supplier, the automated retailer placed orders according to a truncated at zero normal distribution, $N(\mu_n, \sigma)$. The sample value was determined

using the probability $1 - \left(\frac{w(t-1)}{p+g}\right)$. In order to ensure consistency across the different subjects, the automated retailer used the same ordering strategy in all the gaming sessions. Figure 3 provides an indicative example of one of the participants' w -decisions over time (SUP_2), each period representing the next decision in the sequence of 50 rounds. It is evident that, with the exception of period 4, SUP_2 systematically set higher prices than would the rational-optimising supplier w^* . It appears that in period 4 SUP_2 was investigating the impact of a lower wholesale price. This occurred during the 'practice' period and so did not form part of SUP_2 's sample of decisions.

Figure 3 SUP_2 w -decisions, as observed in the laboratory

All participants acting as retailers were asked to play against the same automated supplier that exhibited the same pricing strategy, charging prices between $c=50$ and $p=250$, according to the uniform distribution. Figure 4 presents RET_1 q -decisions over time. In this example RET_1 on average orders almost exactly the same as the rational-optimising retailer (mean order is 105.4). There is, however, a lot of variability in the order quantities of RET_1 , and there may have been some attempt to follow customer demand, 'demand chasing' (Benzion et al, 2008; Bostian et al, 2008), at least in a few periods. Interestingly, RET_1 chooses to order nothing in period 18 despite a high demand in the previous period. This is likely to be either a mistake or a response to the price set by the supplier.

Figure 4 RET₁ q -decisions, as observed in the laboratory

4.3 Stage 3: The Decision-Making Strategies

All participants' recorded w and q decisions satisfied the linearity, normality and hetero-skedasticity requirements of linear regression (Weisberg, 2005; Hair et al, 2006). As a result, we portrayed each supplier's (i) and retailer's (j) decision-making strategies as first order auto-regressive time-series models according to equations (7) and (8) respectively. In these equations, a time-lagged dependent variable is included as one of the explanatory variables (Mills, 1990; Box et al, 1994; Greene, 2003).

$$\langle w(t) \rangle_i = \alpha_0^i + \alpha_w^i \cdot w(t-1) + \alpha_q^i \cdot q(t-1) + \alpha_p^i \cdot P_s(t-1) \quad (7)$$

$$\langle q(t) \rangle_j = \beta_0^j + \beta_w^j \cdot w(t) + \beta_q^j \cdot q(t-1) + \beta_d^j \cdot d(t-1) + \beta_p^j \cdot P_r(t-1) \quad (8)$$

The value of each intercept α_0^i , β_0^j represents the corresponding initial prices or quantities on which each subject i and j anchored all his/her subsequent price or quantity decisions. The intercepts α_0^i and β_0^j also represent the significance that each subject i and j assigned to the pre-selected prices and quantities that were given to him/her at the beginning of each simulation game. 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 the decision attributes of his/her decision $\langle w(t) \rangle_i$ or $\langle q(t) \rangle_j$.

For all the human suppliers' and retailer RET_1 's decision models the corresponding profits $P_s(t-1)$ and $P_r(t-1)$ were removed from the list of independent variables due to the high multi-collinearity that existed between them and the remaining independent variables. High multi-collinearity was exhibited by tolerance levels lower than 0.10. Tolerance levels are defined as the amount of variability of $P_s(t-1)$ and $P_r(t-1)$, respectively, that cannot be explained by the remaining independent variables (Hair et al, 2006).

Since the lagged dependent variable constituted one of the explanatory variables in the decision models, there was auto-correlation in all collected data-sets. This was confirmed by the Breusch-Godfrey test (Breusch, 1978; Godfrey, 1978). For this reason, the appropriate quasi-differences data transformations were applied. Given the relatively small sample sizes and low values of correlation ρ , we preferred the respective ordinary least squares estimators over the feasible generalised least squares that are tailored to time-series processes (Rao and Griliches, 1969).

The linear regression models that have been fitted to the human suppliers and retailers' decisions, along with their corresponding t -values and p -values are presented in Tables 1 and 2. The t -values show how significant the effect of each of the decision attributes was on the actual decisions that participants made. The p -values demonstrate the lowest significance level for which the corresponding decision attributes would be taken into account for subjects' i and j respective decisions $\langle w(t) \rangle_i$ and $\langle q(t) \rangle_j$.

Table 1 Human suppliers' linear regression decision models

	SUP_1			SUP_2			SUP_3		
	<i>Coef.</i>	<i>t</i> -value	<i>p</i> -value	<i>Coef.</i>	<i>t</i> -value	<i>p</i> -value	<i>Coef.</i>	<i>t</i> -value	<i>p</i> -value
α_0^i	115.851	14.710	<0.001	43.929	4.919	<0.001	11.733	2.596	0.015
α_w^i	0.506	15.941	<0.001	0.769	19.098	<0.001	0.921	32.955	<0.001
α_q^i	-0.014	-0.708	0.485	0.011	0.404	0.689	-0.002	-0.097	0.923
α_p^i	-			-			-		
<i>Adj. R</i> ²		0.852			0.889			0.958	

Table 2 Human retailers' linear regression decision models

	RET_1			RET_2		
	<i>Coef.</i>	<i>t-value</i>	<i>p-value</i>	<i>Coef.</i>	<i>t-value</i>	<i>p-value</i>
β_0^j	246.807	18.564	<0.001	258.416	12.294	<0.001
β_w^j	-0.945	-17.686	<0.001	-1.030	-13.110	<0.001
β_q^j	-0.033	-0.449	0.656	0.180	2.311	0.027
β_d^j	-0.045	-0.852	0.400	0.262	3.018	0.005
β_P^j	-			-0.001	-3.146	0.003
<i>Adj. R</i> ²	0.867			0.778		
	RET_3			RET_4		
	<i>Coef.</i>	<i>t-value</i>	<i>p-value</i>	<i>Coef.</i>	<i>t-value</i>	<i>p-value</i>
β_0^j	246.067	14.492	<0.001	32.589	2.938	0.006
β_w^j	-0.952	-18.690	<0.001	-0.048	-1.048	0.301
β_q^j	0.035	0.469	0.642	0.455	5.797	<0.001
β_d^j	0.173	2.285	0.028	0.029	0.794	0.432
β_P^j	-0.001	-1.591	0.120	0.002	6.785	<0.001
<i>Adj. R</i> ²	0.881			0.724		

We can see from Table 1 that all human suppliers ($i=1,...,3$) assigned significant importance to the wholesale price that they charged during the last period $w(t-1)$ (all corresponding p -values<0.001). Although human suppliers did assign some marginal consideration to the retailer's order quantity in the last period (q_{t-1}), the corresponding p -values ≥ 0.485 indicate that this effect might not statistically differ from zero. Most probably it was because the suppliers lacked knowledge and control over the way that retailers' order quantities would respond to their own prices so they tended to only base their w -decisions on their own previous w -prices. Overall the decision models that we fitted to each human supplier were statistically significant at the 1% level and explained more than 85% of the total variation that existed in their recorded decisions (*adj. R*²).

From Table 2 we can see that RET_1 , RET_2 and RET_3 concentrated on the wholesale price $w(t)$ that they were charged (p -values<0.001). RET_4 adopted a very different strategy and chose to ignore this exogenously set price (p -value=0.301). Instead he concentrated on his previous order quantity decision q_{t-1} and realised profit $P_r(t-1)$ (p -values<0.001). RET_2 also took into account previous

demand $d(t-1)$ and realised profit $P_r(t-1)$ for his order quantity decision (p -values of 0.005 and 0.003 respectively). Meanwhile, RET_3 referenced demand in the last period $d(t-1)$ when reaching an order quantity decision (p -value=0.028). Overall the decision models that we fitted to human retailers' decisions were statistically significant at the 1% level, and explained at least 72% of the total variation that was inherent in their recorded decisions ($adj. R^2$).

4.4 Stage 4: The Agent-Based Simulation (ABS) and the Model Runs

The ABS of the newsvendor problem was developed in Excel-VBA. The two agents (supplier and retailer) involved in the newsvendor problem satisfy the following characteristics of an agent (Macal and North, 2010):

- *Self-contained*: they are uniquely identifiable individuals with a clear boundary across which they receive information and make decisions.
- *Autonomous*: they are independent with their behaviour defined by unique individual linear regression models.
- *State that varies over time*: the agents' state varies over time, specifically with respect to the profit achieved in each period.
- *Social*: both the supplier-agent and the retailer-agent are social with the ability to communicate with each other. The wholesale price contract specifies the terms of trade and any exchange that occurs between them.
- *Goal directed*: both the supplier-agent and the retailer-agent have separate goals to achieve and clearly defined internal logic rules that govern their actions. They focus on maximising their own profit.
- *Heterogeneity*: the agents follow their own intentions and adopt individual decision strategies. Those strategies apply different weights to the information available to the agents.

The only characteristic not satisfied by the agents in the newsvendor ABS model is to be *adaptive*. Their decision-making strategies are fixed, as represented by the linear regression models, and so the agents have no specific capability to learn. Macal and North (2010) describe this as a useful, but not essential, characteristic of an agent.

Figure 5 describes the logic of the ABS using state charts. The interactions between the agents are shown by the dashed arrows. In each round (time period, t) of the simulation the supplier agent sets the wholesale price (w), waits for the order from the retailer, delivers the order (q) to the retailer, waits for payment and then receives the payment ($w \cdot q$) from the retailer. Meanwhile, the retailer waits for

the price to be set by the supplier, then determines the order quantity (q), waits for the delivery from the supplier, satisfies the customer demand as much as is possible ($\text{Min}(q, d)$) and receives payment from the customer ($p \cdot \text{Min}(q, d)$). The decision-making strategies of the supplier and retailer agents with respect to setting w and q are represented by the regression equations set out in Table 1 and Table 2 respectively. The profits for the supplier (P_s) and retailer (P_r) are calculated as follows:

$$P_s = q(w - c) \quad (9)$$

$$P_r = p \cdot \text{Min}(q, d) - qw - \text{Max}((d - q)g, 0) \quad (10)$$

The aggregate channel profit in each time period is $P_c = P_s + P_r$.

Figure 5 Agent-based simulation: state charts for supplier and retailer agents, and their interactions

In the simulation, demand values (d) are sampled from a truncated at zero normal distribution, $N(140, 80)$. In order to ensure the efficacy and repeatability of the results, these variates are produced by using the Mersenne-Twister pseudo-random number generator (Matsumoto and Nishimura, 1998) with the same seeds for each scenario.

Table 3 shows the layout of the Excel ABS model with explanations of the calculations in each column given in Table 4. The model is initialised with $w=0$, $q=0$ and $P_r=1000$ in order to seed the calculations in period 1. Any bias incurred from this initial condition is dealt with through the warm-up period in the run strategy explained below.

Table 3 Excel ABS Model: First 4 Periods of Simulation

(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
Period	Random number	Demand (d)	Wholesale price (w)	Order quantity (q)	Orders satisfied	Retailer profit (P_r)	Supplier profit (P_s)	Channel profit (P_c)
0	0.523087	148.48	0	0		1000	-	
1	0.570733	157.78	11.73	38.33	38.33	9013.72	1466.84	7546.88
2	0.236518	90.27	22.46	71.55	71.55	16262.66	1970.44	14292.22
3	0.721902	189.76	32.28	98.74	98.74	21406.92	1749.90	19657.03
4	0.049460	31.51	41.26	123.85	31.51	2767.98	1082.05	1685.93
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.

Table 4 Excel ABS Model Calculations

Column	Calculation
(b)	Random number generated from Mersenne-Twister pseudo-random number generator (Matsumoto and Nishimura, 1998)
(c)	Demand sampled from truncated at zero normal distribution, $N(140, 80)$
(d)	Wholesale price (w) set by supplier's decision model, Table 1
(e)	Order quantity (q) set by retailer's decision model, Table 2
(f)	Orders satisfied = $\text{Min}(d, q)$
(g)	P_r from equation 10
(h)	P_s from equation 9
(i)	$P_c = P_s + P_r$

We then used the ABS model to explore the overall performance of the wholesale price contract under all possible combinations of the inferred supplier and retailer decision-making strategies. The following run strategy was adopted. For all simulations the initial values of w and q were set to zero. Initialisation bias was detected using the MSER-5 heuristic (White, 1997; White et al, 2000; Hoad et

al, 2010) and a warm-up period of 160 time periods was selected. This was based on the longest warm-up for all of the outputs. We ran the model for 1,800 time periods. In order to obtain accurate estimates of mean performance we replicated each simulation 100 times. Excel-VBA code was used to perform the multiple replications using a simple for-next loop.

In order to explore the wholesale price contract's efficiency under different, realistic interactions, we treated the interacting supplier's and retailer's decision-making strategies as the two treatment factors of analysis (F_1 : supplier, F_2 : retailer), with F_1 appearing at 4 levels ($SUP_i, i=1, 2, 3, OPT$) and F_2 at 5 levels ($RET_j, j=1, 2, 3, 4, OPT$). The rational-optimising supplier and retailer (OPT) were kept in the experimental design so that the human suppliers' and retailers' decisions could be directly compared with their rational-optimising counterparts'. It also made it possible to play each human decision-maker with a rational-optimising counterpart. Since the total number of all possible factor combinations ($F_1F_2=20$) was not prohibitively high, a full factorial 'two way layout' design was employed. Had there been a larger number of human participants, it may have been necessary to employ a fractional factorial design. For instance, a doubling in both the number of retailers and suppliers, to 8 and 10 respectively, would have increased the number of factor combinations fourfold ($F_1F_2=80$). The results of these simulation runs are presented in the next section.

5. Simulation Results

We now discuss the results obtained from the simulation runs with the ABS model. The results for the suppliers' pricing decisions, the retailers' order quantity decisions, and the efficiency attained by each supplier-retailer combination are reported. For validation purposes we also discuss how these results compare to the results from previous experimental studies which have employed a much larger number of human participants.

5.1 The Suppliers' Pricing Decisions

Table 5 presents all suppliers' steady-state mean $\bar{w}$ -decisions over $n = 100$ simulated replications for all 20 treatment combinations studied. Between parentheses () in *italics font* the standard deviation of the results across the replications is given, while between brackets [] in **bold font** the half widths of the corresponding 99% confidence intervals are provided.

Table 5 Suppliers' $\bar{w}$ -decisions

$F_1 \backslash F_2$	RET_1	RET_2	RET_3	RET_4	RET_{OPT}
SUP_1	233.99 (0.002) [±0.001]	232.53 (0.015) [±0.004]	233.06 (0.008) [±0.002]	233.15 (0.006) [±0.002]	233.36 (0) [±0]
SUP_2	192.84 (0.003) [±0.001]	195.59 (0.020) [±0.005]	194.30 (0.012) [±0.003]	192.02 (0.015) [±0.004]	194.35 (0) [±0]
SUP_3	146.01 (0.002) [±0.001]	144.15 (0.008) [±0.002]	145.24 (0.005) [±0.001]	146.53 (0.014) [±0.004]	145.26 (0) [±0]
SUP_{OPT}^*	174.75 (0) [±0]	174.75 (0) [±0]	174.75 (0) [±0]	174.75 (0) [±0]	174.75 (0) [±0]

* w decision fixed at rational-optimising level of 174.75

It is clear that a supplier's w -decisions do not differ much when interacting with different retailers. This is because they almost solely focus on their previous w -decision with little cognisance of the retailer's response (Table 1). The rational-optimising supplier SUP_{OPT} is set to charge $w^*=174.75$ in all periods and so there is no difference in this figure between simulation runs. Because the supplier sets w before the retailer chooses q , and independently of previous values of q , we assume that this supplier consistently believes the retailer will choose to order q_r^* . As a result, the supplier chooses the wholesale price $w^*=174.75$ according to expression 1. There is no variance in the suppliers' price decisions during the simulation when interacting with the rational-optimising retailer. Following an initial transient, the interacting supplier and retailer reach a constant equilibrium value for w and q , since variations in demand have no impact on their decisions (expression 2 and Table 1).

The simulated human suppliers seem to adopt two different strategies for maximising their individual profits. SUP_1 and SUP_2 adopt a 'profit margin driven' pricing strategy in which they attempt to maximise their individual profits by charging high prices, above the rational-optimising supplier. SUP_1 follows a more extreme version of this strategy by charging the highest price. Meanwhile, SUP_3 adopts the completely opposite 'demand driven' strategy, charging lower prices than would a rational-optimising supplier, to stimulate demand. All simulated human suppliers charge prices that are significantly different to the rational-optimising supplier at $p<0.01$. This accords with previous experimental research which shows that human suppliers charge wholesale prices that are not consistent with the profit maximising price w^* . Keser and Paleologo (2004), Loch and Wu (2008)

and Wu (2013) all find that, similar to SUP_3 , human suppliers charge lower prices than would a rational-optimising supplier.

5.2 The Retailers' Order Quantity Decisions

Table 6 presents all retailers' steady-state mean $\bar{q}$ -decisions over $n = 100$ simulated replications for all 20 treatment combinations studied. There are much greater differences between the retailers' $\bar{q}$ -decisions than between the suppliers' $\bar{w}$ -decisions, both between retailers and for individual retailers working with different suppliers. RET_1 , for instance, orders on average 18.45 items when interacting with SUP_1 , but orders on average more than five times the items (98.94) when working with SUP_3 .

As discussed in section 5.1, the rational-optimising retailer's order quantities reach an equilibrium value with no variance. However, suppliers' individual pricing strategies do lead to different order quantities when working with the rational-optimising retailer. This is because RET_{OPT} makes decisions in response to the price set by the supplier. When the rational-optimising retailer interacts with the rational-optimising supplier, $\bar{q} = 105.18$, which is as predicted.

Table 6 Retailers' $\bar{q}$ -decisions

$F_2 \backslash F_1$	RET_1	RET_2	RET_3	RET_4	RET_{OPT}
SUP_1	18.45 (0.077) [±0.020]	70.23 (0.533) [±0.140]	51.25 (0.299) [±0.078]	47.99 (0.209) [±0.055]	40.82 (0) [±0]
SUP_2	56.10 (0.071) [±0.019]	113.81 (0.414) [±0.109]	86.78 (0.244) [±0.064]	59.93 (0.313) [±0.082]	87.72 (0) [±0]
SUP_3	98.94 (0.076) [±0.020]	172.43 (0.311) [±0.082]	129.63 (0.196) [±0.052]	78.61 (0.549) [±0.144]	128.83 (0) [±0]
SUP_{OPT}	72.65 (0.075) [±0.020]	137.81 (0.382) [±0.100]	104.11 (0.231) [±0.061]	66.56 (0.390) [±0.102]	105.18 (0) [±0]

RET_1 and RET_4 order quantities that are consistently lower than the rational-optimising retailer in response to each supplier, with the exception of the interaction RET_4 - SUP_1 . It seems that they are employing a strategy of 'minimising excess stock'. RET_2 consistently orders much higher quantities from each supplier than RET_{OPT} . His ordering behaviour seems to be driven by his strong preference

to ‘maximise sales’ by avoiding stock-outs. RET_3 ’s order quantities do not differ substantially from those of the rational-optimising retailer when working with each supplier. The only exception is for the case where he interacts with SUP_1 , who charges the highest prices. It appears that RET_3 is attempting to balance excess stock minimisation with sales maximisation.

As predicted by previous experimental research, all the retailers’ order quantities are significantly different to the rational-optimising retailer (i.e. $q \neq q^*$: at $p < 0.01$). Cases where the human retailers under order are highlighted in Table 6. Schweitzer and Cachon (2000) and Bostian et al (2008) similarly found cases of under ordering in their experimental research. The unshaded cells represent cases where the retailers over ordered, a result that is consistent with the findings of Elahi et al (2013).

5.3 Supply Chain Efficiency Scores

Table 7 presents the mean efficiency scores (Eff) over $n=100$ replications achieved by all 20 treatment combinations studied. First we note that when the two rational-optimising decision-makers interact in the simulation ($SUP_{OPT}-RET_{OPT}$), the efficiency score is 0.741 with no variance, as is expected from the analytical result.

Table 7 Supplier-retailer efficiency scores

$F_2 \backslash F_1$	RET_1	RET_2	RET_3	RET_4	RET_{OPT}
SUP_1	0.150 (0.001) [±0.000]	0.527 (0.006) [±0.002]	0.397 (0.003) [±0.001]	0.375 (0.003) [±0.001]	0.320 (0.001) [±0.000]
SUP_2	0.434 (0.002) [±0.001]	0.778 (0.008) [±0.002]	0.636 (0.005) [±0.001]	0.461 (0.005) [±0.001]	0.644 (0.005) [±0.001]
SUP_3	0.708 (0.005) [±0.001]	0.968 (0.012) [±0.003]	0.851 (0.009) [±0.002]	0.585 (0.008) [±0.002]	0.849 (0.008) [±0.002]
SUP_{OPT}	0.550 (0.003) [±0.001]	0.878 (0.010) [±0.003]	0.734 (0.007) [±0.002]	0.507 (0.006) [±0.002]	0.741 (0.006) [±0.002]

With the exception of the interaction $SUP_{OPT}-RET_{OPT}$, in every case the efficiency scores differ significantly from 0.741 ($Eff=0.741$: at $p < 0.01$). In 14 of the interactions the efficiency is less than expected from the interaction of rational-optimising decision-makers ($Eff < 0.741$: at $p < 0.01$). However, in five of the interactions, highlighted in Table 7, the efficiency is greater than expected ($Eff > 0.741$: at $p < 0.01$). We characterise these as ‘high-performing’ partnerships, since they achieve

efficiency scores above those achieved by self-interested rational-optimising partners. This finding is in line with that of Katok and Wu (2009) who similarly found instances where higher than predicted efficiency scores were achieved by human decision-makers.

In order to help identify which supplier and retailer strategies lead to the highest efficiency scores, their strategies are summarised in Table 8. The first row in each cell identifies the supplier's strategy and the second row shows the retailer's strategy. The cells are un-shaded for partnerships with $Eff < 0.741$; they are light shaded for partnerships with $Eff = 0.741$; and they are shaded dark for partnerships with $Eff > 0.741$. It is immediately obvious that the demand driven strategy of SUP_3 is the most consistently successful among the suppliers. Meanwhile the most successful retailer is RET_2 who follows a sales maximisation strategy. This highlights that the best strategy for maximising channel profit in the newsvendor problem presented in this paper is for the supplier to stimulate demand by charging low wholesale prices and for the retailer to order high quantities in order to try and maximise sales. It is when these two strategies meet through the interaction of SUP_3 with RET_2 that the highest level of profit is achieved ($Eff = 0.968$).

Table 8 Supplier-retailer strategies (row 1 shows supplier strategy, row 2 shows retailer strategy)

$F_2 \backslash F_1$	RET_1	RET_2	RET_3	RET_4	RET_{OPT}
SUP_1	Profit margin Min. Stock	Profit margin Max. Sales	Profit margin Balanced	Profit margin Min. Stock	Profit margin Optimise
SUP_2	Profit margin Min. Stock	Profit margin Max. Sales	Profit margin Balanced	Profit margin Min. Stock	Profit margin Optimise
SUP_3	Demand driven Min. Stock	Demand driven Max. Sales	Demand driven Balanced	Demand driven Min. Stock	Demand driven Optimise
SUP_{OPT}	Optimise Min. Stock	Optimise Max. Sales	Optimise Balanced	Optimise Min. Stock	Optimise Optimise

Key

	$Eff < 0.741$
	$Eff = 0.741$
	$Eff > 0.741$

5.4 Comparison to Results from Previous Studies

The efficacy of our approach is, in-part, determined by whether the ABS reproduces results from previous studies that employed a much larger number of participants. Table 9 summarises the findings from previous studies and then compares these with the findings from the ABS. The comparison is made across the three key results reported in this and other newsvendor studies: the decisions made by the suppliers and the retailers, and their performance, reported here as supply chain efficiency.

Table 9 Comparison of results from this study with predictions from previous work

Result	Prediction from previous work	Findings from this study
Suppliers' pricing decisions	Human suppliers charge lower prices than a rational-optimising supplier. Keser and Paleologo (2004), Loch and Wu (2008) and Wu (2013)	SUP_3 charges lower prices than a rational-optimising supplier. SUP_1 and SUP_2 charge higher prices than a rational-optimising supplier.
Retailers' order quantity decisions	Human retailers order less than a rational-optimising retailer. Schweitzer and Cachon (2000) and Bostian et al (2008) Human retailers order more than a rational-optimising retailer. Elahi et al (2013)	RET_1 consistently under orders. RET_2 consistently over orders. RET_3 and RET_4 under or over order depending on which supplier they work with.
Supply chain efficiency	Human decision-makers can generate higher efficiency scores than rational-optimising decision-makers. Katok and Wu (2009)	Three instances identified of a human supplier and human retailer generating higher efficiency scores than the pairing of a rational-optimising supplier and retailer.

Using our approach, we are able to identify similar outcomes to those of previous studies for suppliers' pricing decisions, retailers' order quantity decisions and the efficiency of the supply chain. We have not been able to identify a previous study in which suppliers charge higher prices than the rational-optimising supplier, as did SUP_1 and SUP_2 . However, it is not unlikely that such a result can occur with decision-makers that have limited knowledge and information. It is interesting that RET_3 and RET_4 under or over order depending on which supplier they work with; something that, as far as we know, has not been identified in previous studies. This finding has emerged because, unlike other studies, the ABS has allowed us to play all participants with each other.

6. Discussion

Our purpose is to explore the use of ABS as a means for addressing issues of sample size in experimental research. The sample size issue is seen from three perspectives: the number of participants, the number of decisions each participant makes, and the number of pairings of participants. The contribution of our approach has been to address all three issues.

Previous studies of the newsvendor problem utilise at least 100 participants, with the exception of Schweitzer and Cachon (2000) who use a smaller sample of 34 and 44 in their two experiments. Our experiment only involves seven participants (four retailers and three suppliers). We are not claiming to be able to improve the statistical power of the results if we wish to draw conclusions about the range of human behaviours when interacting with the automated retailer or supplier. Nor are we claiming to be able to map the full set of possible decision-making strategies. These requirements would need a much larger sample size. Our sample of seven participants, however, does enable us to learn the decision-making strategies of each of these individuals.

Our review of supply chain experimental research shows that participants are asked to complete between 15 and 200 rounds of decision-making. In order to collect the required information to learn a participant's decision-making strategy in our experiment, a sample of 50 decisions is taken. However, once the decision model is implemented in the ABS there is no effective limit to the number of decisions that can be played out. In our example, the model is run for 1,800 decisions and replicated 100 times. As such, in each pairing of retailers and suppliers, we are able to generate 180,000 decisions for each participant, and so 360,000 decisions. Given that we then run 20 combinations of participant pairs, the total number of decisions generated is 7.2 million; and this only takes a few minutes to generate with the ABS. This is orders of magnitude greater than achieved from even the larger of previous experimental studies, such as Katok and Wu (2009), who generate a total of 4,000 decisions from their participants. To collect millions of decisions with human participants would clearly be impracticable.

Previous studies seem to have only been able to collect data from unique pairings of participants. However, once the decision models have been created and embodied in the ABS it is possible to pair every retailer participant with every supplier participant. As such, the limitation on the number of pairings that can interact through the 'game' is not driven by the availability of the human participants.

So, the ABS enables as many participant decisions to be generated as required and all combinations of participants to be paired together. As a result, many more participant decisions and many more participant pairings can be observed than would be possible with human participants. The sample sizes generated by the ABS for each pairing of participants means that results with sufficient statistical power can be generated. Indeed, our results, generated from only seven human participants, accord with the findings from previous studies which involved many more participants.

Our approach has some other advantages. The requirement for only a small number of participants means that experimental research does not need to be restricted to the convenience of using students. It makes it possible to work with a small sample of real decision-makers. Smaller sample sizes also have the benefit of reducing the time involved in running experiments with human participants and the total cost of incentivising participation.

Of course, our experimental approach and use of ABS is not without its limitations. A key issue is the extent to which a small sample of participants provides reasonable coverage of all possible decision-making strategies. Figure 6 illustrates this problem. By using seven decision-makers, we have obtained results for the full set of 12 possible pairings of real decision-makers (4 retailers x 3 suppliers). These are discrete points in a much larger solution space in which there are many other possible outcomes, at least some of which are almost certainly better. We have no knowledge of where on the solution space the pairings lie. Nor do we know the shape of the surface; where better and worse ‘solutions’ might lie. The fact that our results accord with the findings of a number of studies with many more participants suggests that we have achieved a reasonable coverage of the solution space. This, however, is purely fortuitous given that we made no attempt to select a range of decision-maker types in our sample.

Figure 6 Schematic of decision-maker pairings and full solution space of all possible decision-making strategies and pairings

Whether this is a concern depends on the objective of the study. If our aim is to understand and predict the outcome of current decision-making strategies, then the nature of the solution space is of less importance. We may simply want to know why some pairings perform better in order to help the other decision-makers perform to the same level. However, if we wish to search the solution space, particularly with a view to ‘optimising’ decision-making, then the approach is inadequate since we have learnt little about the solution space. Of course, a large sample of participants would help to address this by giving a better understanding of the solution space. This returns us to the first sample size problem: the number of participants. We envisage two ways in which this could be mitigated. The first would involve using personality tests such as Myers-Briggs (Myers and Briggs, 1995) to select participants; a form of stratified sampling. The difficulty here is choosing an appropriate test for the decision situation. The second approach would not require additional participants, but rather the creation of artificial participants generated by adjusting the decision-making models. In our case this would mean adjusting the regression coefficients to create alternative decision-making strategies. This would require careful consideration of the range within which coefficients could be adjusted and of the practical meaning of the different decision-making models.

A second issue is the extent to which human participants strategies in a gaming environment are similar to those they would employ if the decisions were real. This issue not only affects the ABS approach, but any laboratory experiment. The use of financial incentives, frequently employed in

experimental research, is one way to try and mitigate this effect. The impact of financial incentives on participant effort and performance is a complex phenomenon. Both Camerer and Hogarth (1999) and Bonner and Sprinkle (2002) perform an extensive review of experimental studies and the use of financial incentives. They find a mixture of effects, complicated by the difficulty of the task and the nature of the participants.

The regression models for representing the participants' decision-making strategies all provided a reasonable good fit to the data ($0.724 \leq \text{adj. } R^2 \leq 0.958$). There remains, however, up to 28% of the variation in decisions that is not explained by the regression models. This will impact on the validity of the ABS model. To address this, alternative approaches for representing the decision-making strategies could be employed such as artificial neural networks or rule-based expert systems. We would, however, always expect some level of error (residuals for regression models) in the decision functions due to variability in human decision-making; given the same scenario an individual is unlikely to make exactly the same decision on every occasion that the scenario is presented. Approaches for representing this variability in agent behaviour could be explored.

A fourth limitation arises from the participants interacting with an automated decision-maker. This mitigates the effects arising from social preference and reputation, but it means that the ABS model only represents a single decision-making strategy for each participant. In practice, participants are likely to adjust their decision-making strategy according to whom they are interacting with. This may limit the benefit of learning a single participant strategy and then assuming that same strategy would be used when interacting with all other decision-makers.

Future research could aim to explore and address these limitations. Beyond this, it would be useful to repeat this study with different newsvendor problem parameter settings in order to understand the impact on individual's decision-making and on the ABS approach. It would also be of interest to adopt the ABS approach in a more complex setting with more than two decision-makers interacting with one another. An obvious candidate would be to perform similar experiments with the beer distribution game (Sterman, 2000).

7. Conclusion

As work in behavioural OR develops, laboratory based experimental research is set become of increasing importance to OR research. Our work demonstrates the potential for ABS to enhance the scope of such experimental research. We encourage further developments in the use of ABS to support experimental work and research in behavioural OR.

References

- Bakken, B., J. Gould and D. Kim. 1994. Experimentation in learning organisations: A management flight simulator approach. *Modelling for Learning Organisations*, edited by Morecroft, J.D.W. and J.D. Sterman. Productivity Press, Portland, OR: 243-266.
- Bearden, J.N. and A. Rapoport. 2005. Operations research in experimental psychology. *Tutorials in Operations Research*. INFORMS, New Orleans.
- Becker-Peth, M., E. Katok and U.W. Thomemann. 2013. Designing contracts for irrational but predictable newsvendors. *Management Science*. In press.
- 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**(9) 1281-1287.
- Bolton, G.E. and E. Katok. 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.E., A. Ockenfels and U.W. Thonemann. 2008. Managers and students as newsvendors: How out-of-task experience matters. Working Paper, University of Cologne.
- Bonner, S.E. and G.B. Sprinkle. 2002. The effects of monetary incentives on effort and task performance: theories, evidence, and a framework for research. *Accounting, Organizations and Society* **27**(4) 303-345.
- Bostian, A.A., C.A. Holt and A.M. Smith. 2008. The newsvendor pull-to-center effect: Adaptive learning in a laboratory experiment. *Manufacturing and Service Operations Management* **10**(4) 590-608.
- Box, G., G.M. Jenkins and G. Reinsel. 1994. *Time series analysis: Forecasting and control 3rd ed.* Prentice Hall, Upper Saddle River, NJ.
- Breusch, T.S. 1978. Testing for autocorrelation in dynamic linear models. *Australian Economic Papers* **17**(31) 334-356.
- Cachon, G.P. 2003. Supply chain coordination with contracts. *Supply Chain Management: Design, Coordination and Operation*, edited by Kok, A. and S. Graves. Elsevier, Amsterdam: 229-341.
- Camerer, C. 1995. Individual decision making. *Handbook of Experimental Economics* edited by Kagel, J. and A. Roth. Princeton University Press, Princeton, NJ: 313-327.
- Camerer, C.F. and R.M. Hogarth. 1999. The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty* **19**(1-3) 7-42.
- Croson, R. and K. Donohue. 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science* **52**(3) 323-336.

- de Véricourt, F., K. Jain, N. Bearden and A. Filipowicz (2013). Sex, risk and the newsvendor. *Journal of Operations Management* **31**(1/2) 86-92.
- Elahi, E., N. Lamba and C. Ramaswamy. 2013. How can we improve the performance of supply chain contracts? An experimental study. *International Journal of Production Economics* **142** 146-157.
- Frederick, S. 2005. Cognitive reflection and decision making. *Journal of Economic Perspectives* **19**(4) 25-42.
- Godfrey, L.G. 1978. Testing against general autoregressive and moving average error models when the regressors include lagged dependent variables. *Econometrica* **46**(6) 1293-1301.
- Greene, W.H. 2003. *Econometric Analysis*, 5th ed. Prentice Hall, Upper Saddle River, NJ.
- Hair, J.F., B. Black, B. Babin, R.E. Anderson and R.L. Tatham. 2006. *Multi-variate Data Analysis* 6th ed. Prentice Hall, Upper Saddle River, NJ.
- Halkos, G. and I. Kevork. 2011. Non-negative demand in newsvendor models: The case of singly truncated normal samples. Munich Personal RePEc Archive Paper No. 31842, <http://mpra.ub.uni-muenchen.de/31842> accessed October 2014.
- Hämäläinen, R. P., Luoma, J., & Saarinen, E. (2013). On the importance of behavioral operational research: The case of understanding and communicating about dynamic systems. *European Journal of Operational Research*, **228**(3) 623-624
- 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. 2010. Automating warm-up length estimation. *Journal of the Operational Research Society* **61**(9) 1389-1403.
- Kalkanci, B., K.-Y. Chen and F. Erhun. 2011. Contract complexity and performance under asymmetric demand information: An experimental evaluation. *Management Science* **57**(4) 689-704.
- Katok, E. and D.Y. Wu. 2009. Contracting in supply chains: A laboratory investigation. *Management Science* **55**(2) 1953-1968.
- Keser, C. and G.A. Paleologo. 2004. Experimental investigation of supplier-retailer contracts: The wholesale price contract. Working Paper, IBM T.J. Watson Research Center, Yorktown Heights, NY.
- Khouja, M. 1999. The single-period (news-vendor) problem: Literature review and suggestions for future research. *Omega International Journal of Management Science* **27** 537-553.
- Lariviere, M.A. 1999. Supply chain contracting and coordination with stochastic demand. *Quantitative Models for Supply Chain Management*, edited by Tayur, S., R. Ganeshan and M. Magazine. Kluwer Academic Publishers, Boston, MA: 233-268.
- Lariviere, M.A. and E.L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing and Service Operations Management* **3**(4) 293-305.

- Loch, C.H. and D.Y. Wu. 2007. *Behavioral Operations Management: Foundations and Trends in Technology, Information and Operations Management*. Now Publishers, Hanover, NH.
- Loch, C.H. and D.Y. Wu. 2008. Social preferences and supply chain performance: An experimental study. *Management Science* **54**(11) 1835-1849.
- Macal, C.M. and M.J. North. 2010. Tutorial on agent-based modelling and simulation. *Journal of Simulation* **4**(3) 151–162.
- 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.
- Mills T.C. 1990. *Time Series Techniques for Economists*. Cambridge University Press, Cambridge, UK.
- Moritz, B.B., A.V. Hill and K.L. Donohue. 2013. Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management* **31**(1) 72-85.
- Myers, I.B. and Myers, P.B. (1995). *Gifts Differing: Understanding Personality Type*, 2nd ed. Davies-Black Publishing, Mountain View, CA.
- Rao, P. and Griliches, Z. 1969. Small sample properties of several two-stage regression methods in the context of auto-correlated errors. *Journal of the American Statistical Association* **64**(325) 253-272.
- Robinson, S., T. Alifantis, J.S. Edwards, J. Ladbroke 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**(8) 912-921.
- Robinson, S., Lee, E.P.K. and Edwards, J.S. (2012). Simulation Based Knowledge Elicitation: Effect of Visual Representation and Model Parameters. *Expert Systems with Applications*, 39, pp. 8479–8489.
- Schiffels, S., A. Fügner, R. Kolisch, J.O. Brunner. 2014. On the assessment of costs in a newsvendor environment: insights from an experimental study. *Omega – The International Journal of Management Science* **43** 1-8.
- Schweitzer, M.E. and G.P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science* **46**(3) 404-420.
- Senge, P.M. 1990. *The Fifth Discipline: The Art and Practice of the Learning Organization*. Random House, London.
- Steckel, J.H., S. Gupta and A. Banerji. 2004. Supply chain decision making: Will shorter cycle times and shared point-of-sale information necessarily help. *Management Science* **50**(4) 458-464.
- Sterman, J.D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill, Boston, MA.
- Weisberg, S. 2005. *Applied Linear Regression*, 3rd ed. John Wiley and Sons, NJ.

- White, K.P. 1997. An effective truncation heuristic for bias reduction in simulation output. *Simulation* **69**(6) 323-334.
- White, K.P., Cobb, M.J. and S.C. Spratt. 2000. A comparison of five steady-state truncation heuristics for simulation. *Proceedings of the 2000 Winter Simulation Conference*, edited by Joines, J., R. Barton, K. Kang and P. Fishwick. IEEE, Piscataway, NJ: 755-760.
- Whitin, T.M. 1955. Inventory control and price theory. *Management Science* **2**(1) 61-68.
- Wu, D.Y. 2013. The impact of repeated interactions on supply chain contract: a laboratory study. *International Journal of Production Economics* **142**(1) 3-15.
- Wu, X. and J.A. Niederhoff. 2014. Fairness in selling to the newsvendor. *Production and Operations Management*. Published on-line.
- Wynder, M. 2004. Facilitating creativity in management accounting: a computerized business simulation. *Accounting Education* **13**(2) 231-250.