Int. J. Production Economics 142 (2013) 61–73

Contents lists available at SciVerse ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

A human experiment on inventory decisions under supply uncertainty

A. Ancarani [a,*], C. Di Mauro [1,a], D. D'Urso [2,b]

[a] Dipartimento di Ingegneria Civile e Ambientale, Università di Catania—Viale A. Doria, 6—95125 Catania, Italy
[b] Dipartimento di Ingegneria Industriale e Meccanica, Università di Catania—Viale A. Doria, 6—95125 Catania, Italy

ARTICLE INFO

Article history:
Received 31 March 2011
Accepted 3 September 2012
Available online 10 September 2012

Keywords:
Supply risk
Ambiguity
Stochastic lead times
Human experiments
Beer game
Inventory management

ABSTRACT

Controlled human experiments are adopted in this paper to investigate the impact of supply uncertainty on buyers' inventory management. The experiments aim at assessing the impact of one specific source of supply chain uncertainty, namely stochastic lead times, on inventory holdings and the extent of the bullwhip effect. Three experimental treatments are run within the framework of the beer game manipulating variability in demand and in lead times. Results confirm that the bullwhip effect arises in all experimental treatments and that the variance of orders is higher under stochastic lead times. Analysis of players' behaviour in the course of the game suggests that players react to higher uncertainty by holding fewer inventories, a behaviour consistent with the predictions of some psychological models of choice under ambiguity.

1. Introduction

An increasing degree of complexity characterises the supply chains of many sectors (Esposito and Passaro, 2009). Among the causes of such a trend are outsourcing, the enlargement of supplier networks, increased dependence on supplier capabilities, shorter product life-cycles, and international market and production expansion (Wagner and Neshat, 2010). Further, as firms try to reduce costs through the rationalisation and reduction of the supply base, the aim to secure a stable flow of materials has become more difficult to achieve (Harland et al., 2003). As a consequence of both higher complexity and leaner supply chains, the instability of supply chains and supply uncertainty have increased (Geary et al., 2006).

A paradigmatic representation of supply chain instability is the Bullwhip Effect (BWE). The BWE is generally triggered by demand uncertainty (Forrester, 1958), and it entails that, as external demand passes through the SC from the downstream to the upper levels of the chain, the variance of orders is amplified. This behaviour can imply substantial costs in terms of stock-out as well as inventory holding and obsolescence costs, thus worsening the performance of the SC.

While the impact of demand variability on SC instability and performance has been explored in several studies (Croson and Donohue, 2006; Steckel et al., 2004; Gupta et al., 2002; Sterman, 1989), the impact deriving from supply-side sources of uncertainty has received less attention, in spite of the fact that some authors have posited that a reduction in SC instability is best enabled via implementation of the principles of smooth material flow, and by decreasing actual or perceived shortage risk (Geary et al., 2002, 2006).

A small number of numerical simulations (Chatfield et al., 2004; Truong et al., 2008) has investigated the effects of supply uncertainty on the extent and consequences of SC instability by making supply uncertainty operational through stochastic lead times, one of the most relevant supply-side sources of uncertainty. These studies have shown that, generally, stochastic lead times contribute to worsen SC instability.

While numerical simulations can throw light on how rational and optimising agents can react to SC uncertainty, they cannot fully account for deviations from rationality or limited cognitive abilities of SC managers. In this direction, experimental research on human subjects in neighbouring disciplines to Operations Management has shown that decision makers apply heuristics in processing tasks characterised by uncertainty (Kahneman and Tversky, 1974), and that they may use these heuristics as a way "to live with risk" (Gigerenzer, 2002). Further, decision makers exhibit biases in processing probabilistic information, since they distort probabilities of outcomes even when they are objectively known (Kahneman and Tversky, 1979) and, according to the domain of outcomes (costs vs. revenues), they might dislike/prefer uncertainty (Ellsberg, 1961; Wakker, 2010).

Controlled human experiments have gained importance as a methodology for the study of SC instability since Sterman's (1989) finding that the BWE is a problem arising as a consequence of human decision making and stemming from the amplification of unanticipated changes in demand, and from a biased perception of the flows in transit through the SC pipeline.

* Corresponding author. Tel.: +39 095 738 2715; fax: +39 095 738 2748.
E-mail addresses: aancaran@dica.unict.it (A. Ancarani), cdimauro@unict.it (C. Di Mauro), ddurso@diim.unict.it (D. D'Urso).
[1] Tel.:+39 09570305216.
[2] Tel.:+39 0957382412.

0925-5273/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ijpe.2012.09.001

In the face of varying types and degrees of SC uncertainty, managers' risk mitigation actions depend on their individual attitudes, on their perceptions of the likelihood of supply disruptions, and on the degree of confidence they assign to available risk information (Ellis et al., 2010). Thus, there are grounds for positing that under supply uncertainty heuristics and biases used in assessing probabilities of outcomes affect either the size of the BWE, or inventory holdings and inventory policies, or both, in ways which might be at odds with the predictions of numerical simulations. Further, since heuristics may develop through time in repeated tasks, it is relevant to provide insight into the way individuals adapt their decisions as they experience a highly variable environment. This analysis can hardly be carried out in the field because of lack of sufficient control, thus suggesting that human experimentation may prove a useful tool to explore the effects of supply uncertainty on the BWE formation and SC costs (Bendoly et al., 2006; Ancarani and Di Mauro, 2011).

To the best of our knowledge, a test of the impact of supply uncertainty on SC performance using human subjects is still lacking. In this paper, we apply human experiments with the aim to study the behaviour of members of a SC in the face of lead time uncertainty, and contrast the performance of a SC with stochastic lead times with that of a SC with deterministic lead times. We carry out the study within the framework of the beer distribution game, which reproduces a serial supply chain with four echelons (retailer, wholesaler, distributor, and factory).

Two research questions are investigated:

1. What is the impact of supply uncertainty on supply chain instability, as measured by the BWE, and on supply chain performance, as measured by SC costs?
2. How do SC managers react to supply uncertainty in terms of ordering decisions and inventory holdings?

Our results show that supply uncertainty, in terms of stochastic lead times, gives rise to a higher variance of orders at every echelon of the supply chain. More intriguingly, we find that, when the SC is characterised by both demand uncertainty and stochastic lead-times, buyers hold fewer inventories, a behaviour that we attribute to an uncertainty loving attitude.

The paper is organised as follows: Section 2 reviews the relevant branches of the literature underpinning the hypotheses tested through the experiment, Section 3 presents the experimental design, while Section 4 develops benchmarks for behaviour in the experiment by means of numerical simulation. Section 5 presents the results of the human experiment. Section 6 discusses the main findings, and highlights implications for SC management and future research. Section 7 concludes the paper.

2. Factors investigated and hypotheses tested

2.1. The beer game and the bullwhip effect with stochastic lead times

The BWE has been widely studied in the context of the so-called "beer game". In the classic beer distribution game (Forrester, 1958) the supply chain consists of four echelons (retailer–wholesaler–distributor–factory). Inventory is managed according to the periodic review inventory model (order-up-to). During the game each i-participant, i [1,…,4], at each t-period, t [1,..,T], places orders, $O_{i,t}$, to the immediate upstream supplier and fills downstream customer's order, $D_{i,t}$. Typically, at each echelon, when a buyer places an order a delay of 1 week (Lead Time of Information—LTI) occurs before this latter is known to the upstream supplier, $D_{i,t}=O_{i-1,t-1}$; a 2 weeks lead time (Lead Time of Distribution—LTD) is requested to ship orders to the downstream echelon and the same happens to the factory when beer is produced (Lead Time of Production—LTP). At each echelon, goods received at time t, $R_{i,t}$, correspond to those shipped by the upstream supplier 2 weeks before, $S_{i+1,t-2}$. During the game the inventory balance is such that: $I_{i,t}=I_{i,t-1}+R_{i,t}-S_{i,t}$, where $I_{i,t}$ is the on hand quantity ($I_{i,t}\geq 0$); customer orders are filled completely if $I_{i,t-1}+R_{i,t}\geq D_{i,t}$ otherwise $S_{i,t}<D_{i,t}$ and backorders occur, $B_{i,t}=B_{i,t-1}+D_{i,t}-S_{i,t}$.

If the external demand distribution is non-stationary and/or unknown, larger oscillations of orders occur moving upstream the SC, giving rise to the BWE. When the demand distribution is both stationary and known, then no BWE arises if a simple base-stock inventory policy is used, whereby individuals place orders equal to the orders they receive (Chen, 1999).

Most of the extant literature on the BWE has focused on the effect of uncertainty in customer demand, while ignoring the potential impact of supply-side sources of uncertainty on the BWE. More recently, a few papers have addressed the problem of the BWE and of SC performance relaxing the assumption of deterministic lead times. Stochastic lead times complicate the picture of the standard beer game, because in every period the volume of shipments received by the buyer is probabilistic. Further, if shipments in every period are independent draws, order cross-over may occur.

Chatfield et al. (2004) use simulation to investigate the effects of stochastic lead times in a k-node SC. In particular, through a factorial design, the effect of various degrees of lead time variability is crossed with that of four different levels of information quality, and with the absence/presence of information sharing. A normally distributed customer demand is assumed. The first information level refers to the situation in which there is no updating of policy parameters during the game, whereas the other three levels involve some form of updating based on demand and/or lead time historic information. Generally, updating of policy parameters occurs if the distributions of demand and lead times are unknown. When demand and lead time distributions are known, the optimal order-up-to quantity can be chosen in advance. In this instance, stochastic lead times do not generate BWE. However, there is evidence that buyers may use historical information also with known lead times (Chatfield et al., 2004). Updating of policy parameters worsens the BWE, as the variance of lead time increases. In particular, the amplification of order variances is highest when historical information on both demand and lead times variances are used to update inventory parameters, while it is lowest when lead time information is not used, because of misperception of the variability or the belief that lead time variation is not important. Comparable results are confirmed by Truong et al. (2008) assuming either an AR(1) or an ARMA(1,1) model for customer demand.

Kim et al. (2006) present a model with stochastic lead time in which the case of information sharing (customer demand is common knowledge for all echelons of the chain) is contrasted with that of no sharing. Results show that the variance of orders increases nearly linearly in echelon stage with information sharing, and exponentially without information sharing. One managerial implication of this result is that the sharing of information on customer demand by all echelons is an effective way to reduce BWE also under conditions of supply uncertainty. However, information sharing per se does not eliminate BWE.

Heydari et al. (2009) isolate the impact of lead time uncertainty from that of demand uncertainty by simulating a four-stage SC in which customer demand is constant. Results show that uncertainty in lead time increases the variance of orders at each echelon but does not worsen BWE. Further, order variance is positively correlated to the variance of inventory levels and the size of stock-outs. Results of this study are not directly comparable with those already discussed, since Heydari et al. (2009) assume that if the supplier holds insufficient inventory to satisfy an order, the unavailable quantity will be lost. Similarly, all delayed orders will be lost. Further, the uppermost level of the chain does not face stochastic lead times.

To conclude, under the assumption of fully rational agents, most models reviewed conclude that without ex ante knowledge of the lead time distribution or if the buyer makes use of historic information on lead times and demand, stochastic lead times enhance the variance of orders at each echelon. However, to the best of our knowledge, these results have not been put to empirical test through human experimentation. The above discussion suggests the following hypothesis:

Hypothesis 1. In the presence of both customer demand uncertainty and stochastic lead time, if the buyer updates the inventory policy parameters on the basis of demand and/or lead time variance information, then the BWE is higher than in the case when only demand uncertainty is considered and the lead time is deterministic.

2.2. Supply uncertainty and inventory decisions: a behavioural perspective

Most decisions within organisations are made in the presence of uncertainty about their outcomes. Uncertainty can take the form of risk, which is present when there are multiple possible outcomes that could occur with well-defined probabilities (Bernoulli, 1738). Uncertainty can also refer to ambiguity, i.e. a situation in which there are multiple possible outcomes whose probabilities are vague or unknown (Knight, 1921).

According to the normative theory of individual behaviour under uncertainty, Subjective Expected Utility (SEU) theory (Savage 1954), even if probabilities are not objectively well-defined (i.e. externally given probabilities are not available), decision makers assign their own subjective beliefs to probabilities, which thus they perceive as defined and known to them. Hence, according to SEU theory, the distinction between risk and ambiguity is meaningless, and decision makers should not display any specific reaction in the face of probabilistic vagueness, as individuals are always acting "as if" probabilities were known.

However, behavioural analyses have shown that decision makers behave at odds with the predictions of SEU theory and that they exhibit biases in processing probabilistic information. For instance, they distort probability of outcomes even if they are objectively known (Kahneman and Tversky, 1979), and process decisions characterised by ambiguity differently from those with known probabilities (Ellsberg, 1961; Camerer and Weber, 1992; Wakker, 2010). According to Ellsberg (1961) the perception of ambiguity depends "on the type, amount, reliability and unanimity of information", and gives rise to "one's degree of confidence in an estimate of relative likelihoods" (p. 657).

A common reaction in the face of both risk and ambiguity is "pessimism", i.e. people tend to overestimate the less favourable outcomes. Also, people often tend to prefer well-defined probabilities to vague probabilities, and thus they exhibit greater pessimism under ambiguity. The difference in pessimism between ambiguity and risk reflects an individual attitude called ambiguity aversion (Abdellaoui et al., 2010). Whereas risk aversion is measured by the willingness of the decision maker to accept a sure sum lower than the expected value of a bet, ambiguity aversion is reflected by the willingness to assign a greater value to a bet whose probability distribution is considered to be more reliable. As an illustration of the difference between risk and ambiguity aversion, a risk averse decision maker strictly prefers getting 50 euro for sure rather than betting 100 euro on a fair coin landing head, whereas an ambiguity averse decision maker prefers to bet on a fair coin landing head rather than on a thumbtack landing head.

The attitude toward ambiguity may be of relevance in several decision domains. For instance, entering new markets or making venture investments may be modelled as a decision problem under ambiguity, insofar the probability of success associated with these decisions is not objectively known. While entrepreneurs are typically characterised as agents with low risk aversion (and thus showing a preference for a less likely but larger return over a highly likely but smaller one, when these options offer the same expected value), they may still be ambiguity averse. This means that given the same expected value, they strictly prefer the investment option in which the likelihood of possible returns is known (Rigotti et al., 2011).

Evidence on the practical implications of distinguishing between risk and ambiguity attitudes can be found in the field of product selection. Muthukrishnan et al. (2009) show that ambiguity aversion leads consumers to systematically prefer established brands even when these products are dominated on all attributes, and that risk aversion does not play any role in explaining the choice of brand. In the same vein, Hazen et al. (2012) show that there is a direct relationship between the level of tolerance for ambiguity and the customer willingness to pay for remanufactured products. Thus, in order to get higher prices in the marketplace, producers of remanufactured products should not focus on risk reduction measures (for instance through the provision of a warranty on defective products) but should rather work to reduce the level of ambiguity associated with their processes.

To conclude, risk and ambiguity attitudes are different cognitive phenomena which are not necessarily correlated (Wakker, 2010), and thus require distinct analytical formalization (Schmeidler, 1989; Tversky and Wakker, 1995).

Although individuals are generally averse to ambiguity, there is evidence that, according to the decision context, individuals may be ambiguity seeking rather than ambiguity averse (Abdellaoui et al., 2010), i.e. they might exhibit "optimism" in the face of vagueness in probability. In particular, psychological models (Einhorn and Hogarth, 1985, 1986; Hogarth and Einhorn, 1990) posit that, when ambiguity concerns the occurrence of a loss, ambiguity seeking behaviour prevails when the expected probability of that loss is high, while ambiguity aversion should be observed when the expected likelihood is low. Conversely, when ambiguity concerns the probability of a gain, ambiguity aversion prevails at high expected probabilities, while ambiguity seeking prevails at low expected probabilities.

In experimental research, ambiguity seeking behaviour has been reported more frequently when the decision is framed as a loss (Cohen et al., 1985; Viscusi and Chesson, 1999; Di Mauro and Maffioletti, 2004). For instance, in the insurance industry (where insurance purchasers are called to incur a lower sure cost in order to reduce higher potential costs in the future), consumers have shown a lower willingness to pay to cover against losses if probabilities are perceived as ambiguous (Hogarth and Kunreuther, 1985, 1989). Therefore, understanding the character of attitudes toward ambiguity is important to better understand and predict what precautionary self-protecting measures individuals and organisations will adopt in response to uncertainty.

Focusing on inventory management, where decisions are often characterised by risk and ambiguity, analytical models have investigated the impact of the decision maker's risk aversion (Jammernegg and Kischka, 2009), while the role of ambiguity attitude has not been taken into consideration yet.

In the beer game, ambiguity arises if the likelihoods of one or more relevant variables (demand, shipment, stock-out) are not well-defined. Elaborating on Ellsberg (1961), the ambiguity perceived in the beer game increases as buyers' degree of confidence in an estimate of relatively likelihood of demand and/or shipment diminishes. The increase in ambiguity stems either from the fact that the probability distributions of these variables are unknown, or from the fact that the decision maker cannot calculate well-defined probabilities because of computational limitations. These limitations, for instance, may be at the root of the finding that decision makers consider situations characterised by known but compound probabilities more ambiguous than situations in which simple probabilities are involved (Halevy, 2007). Following this line of reasoning, consider for instance the case of two beer game models having the same demand distribution but differing in the fact that in the first lead times are constant while in the second lead times are stochastic. Assume further that the probability distribution of lead times in the second model is known and has a mean value equal to the constant lead time of the first model. We expect that the decision maker considers the latter model more ambiguous than the former. The first reason is that even if the probability distribution of lead times is known, the amount received at each period is difficult to calculate, especially for the downstream echelons of the SC. This effect is exacerbated if the realisation of lead times of each echelon is an independent draw, and order cross-over is possible. Next, the risk of a stock-out becomes ambiguous because it is cognitively cumbersome to compute, especially if buyer and supplier do not share inventory information and several inter-related layers of buyer–supplier relationships are involved.

Likewise, assume we compare two beer game models sharing the same stochastic distribution of lead times, but differing because one has constant external demand and the other exhibits stochastic external demand. We expect the latter to be perceived as more ambiguous than the former, especially for the uppermost echelons which have to form subjective beliefs about the distribution of the demand of the downstream echelon.

When confronting ambiguity, an ambiguity averse buyer will protect against the risk of running out of stock by holding more inventory with respect to a situation with no ambiguity or relatively less ambiguity. Intuitively, the ambiguity averse decision maker's beliefs are pessimistic, i.e. more weight is put on a late delivery of stock or an unexpected increase in demand. Conversely, an ambiguity seeking buyer has optimistic beliefs and overweighs the probability of an early delivery, thus holding fewer inventories.

Inventory holding is a risk mitigation strategy, thus akin to buying insurance coverage: increasing inventories entails a lower but *sure* cost now vis-a-vis a *potential* benefit from avoiding the higher cost of a stock-out. Therefore, we may conjecture that, under ambiguity, the same behaviour observed in insurance settings applies to inventory management. Specifically, ambiguity preference is likely to prevail, i.e. when comparing inventory holdings in two environments characterised by different degrees of ambiguity, inventories will be lower in the more ambiguous environment. Although counter-intuitive, this behaviour is consistent with psychological models (Einhorn and Hogarth, 1985) and with the experimental evidence on behaviour in insurance settings discussed above (Hogarth and Kunreuther, 1985, 1989). Thus, we formulate the following hypothesis:

Hypothesis 2. Decision makers in the beer game will hold fewer inventories in settings characterised by more ambiguity.

3. Experimental design

In order to explore how human players react to supply and to demand uncertainty, three experimental treatments were run using human participants. The treatments were obtained by manipulating stochastic vs. known lead times and constant vs. variable demand. All three treatments can be assumed to mimic a non-integrated supply chain in which each buyer has a single supplier; no information sharing about actual demand, inventories, backlogs, and own lead time, is allowed among SC participants.

The first treatment (SBG hereafter) reproduced a beer game with four echelons ($i=1,\ldots,4$), i.i.d. normally distributed external demand with parameters known to all echelons ($\mu=100$, $\sigma=20$), known and constant lead times equal to one period for information lead times (LTI$=1$) and to two periods for distribution lead times (LTD$=2$). This design differs from Sterman's experiments, in which the retail demand is completely unknown and non-stationary and is represented by a simple step-function whereby demand starts at 4 units and jumps to 8 units after the eighth game period. However, subsequent studies by Croson and Donohue (2006) have shown that even with stationary and known demand distributions, the BWE arises. Thus, we expect the BWE to arise also in our setting.

In the second treatment (SLT henceforth) players face both demand and supply uncertainty, since demand is normally distributed as in the first treatment, and the LTD of all suppliers in the chain (including the factory's brewery) is uniformly distributed in the interval (1, 2, 3) periods. Thus, although the mean lead time is the same in both SBG and SLT, in the latter treatment both delays and anticipated deliveries with respect to the mean lead time are possible. In the SLT treatment, because of stochastic lead times, the possibility of order cross-over arises. The LTI remains constant and equal to 1, as in the first treatment.

A third treatment (CD_SLT) was implemented with the aim to isolate the impact of lead time uncertainty from that of demand uncertainty. In CD_SLT external demand was constant and equal to 100 pieces per period, while transportation lead times were stochastic and equally distributed in the interval 1–3 periods, like in SLT.

Comparison among the alternative treatments is meaningful insofar differences in players' behaviour and performance can clearly be attributed to a specific factor. Thus, the comparison between SBG and SLT (change in lead times variability given the same distribution of demand) allows isolating the impact of lead time uncertainty in a context where demand uncertainty exists, while the comparison between SLT and CD_SLT (change in demand variability given the same distribution of lead times) allows isolating the impact of demand variability in a context characterised by supply uncertainty. A direct comparison between SBG and CD_SLT is not granted by the experimental design, since the uncertainty created by demand variability in SBG is not equivalent to that determined by supply uncertainty in CD_SLT.

In all treatments, an order placed with the supplier can be partially fulfilled with a continuous distribution, depending on the supplier's inventory availability. Each role incurs unit inventory holding costs of €0.50 and unit backlogging costs of €1.00 per period (Sterman, 1989).

The information set available to each player includes the histories of incoming demands, of past shipments and of past purchases. From this information, outstanding orders and the history of lead times can be worked out.

Behaviour was observed for a number of periods (T), from a minimum of 36 up to 50. Players were not informed of the final period of the game to avoid end-of-game behaviour that might trigger over- or under-ordering.

Each echelon began with an initial inventory level $I_i(t=1)=300$ which allowed to deplete inventories by satisfying mean demand within the mean lead time, outstanding orders

$O_i(t=0, -1)=100$ for the previous two periods, and an incoming shipment $S_i(t=2, 3)=100$ in the following two periods. All experiments also used the same random number seed to generate demand, i.e., D_{it}, $t=1, \ldots, T$ was identical across groups. This allowed us to isolate variations due to ordering behaviour from variations due to different demand streams.

The beer game is a complex game with many feedbacks and loops. Many experiments use first time players and therefore, it could be conjectured that the BWE observed is a mere consequence of lack of adequate training. To keep account of this, we build on Wu and Katok (2006) who find that both hand-on training and communication must be provided to beer game players in order to obtain improvements of SC performance. To this end, each player participated in two different repetitions of the same beer game treatment, the second taking place about one month after the first. In our view, the long lag between the two repetitions should allow for a more effective internalisation and exchange of knowledge, which can be considered an essential component of learning, especially in complex tasks (Holsapple and Singh, 2001). Although participants were not told that they would be asked to play the same game a second time, we expect that after the first session they engaged in reflection and thinking over their actions in the game, as well as in exchange of views with other participants. This was confirmed by verbal communication we had with the participants.

In the second repetition, each player kept the same role he/she had been assigned in the first game but was assigned to a different chain, in order to avoid that members of a chain during the first session agreed on a specific strategy for the next. This would have obviously changed the SC from a non-integrated chain to a coordinated one. Results of the first repetition were considered as a form of hand-on training and are not reported in this paper.

Participants in all treatments were graduate students who had attended at least one course in Operations Management. This was done in order to enhance the chances that the experimental subjects may be considered "tomorrow's inventory professionals" (Croson and Donohue, 2006). A total of 124 students participated in the second repetitions of experiments that are reported in this paper. Median age was 24 years old, participants were equally split between males and females. Participants were randomly assigned to one of the treatments upon signing up. Once seated, they were oriented to the rules and objectives of the game by means of a tutorial lasting about 30 min. During the game, communication was strictly forbidden.

Table 1 summarises the distribution of players across the different cells of the experiment.

The incentive used in the game was both monetary and in terms of coursework grades. Participants were instructed that the members of the supply chain team with the lowest total costs (inventory + backlog costs) shared a final prize of €77. In addition, the members of the winning chain received an extra course grade (out of a total of 30 grades).

Table 1
Participants in the different treatments.

	Constant external demand	Stochastic external demand
Constant lead time	–	SBG 11 chains
Stochastic lead time	CD_SLT 10 chains	SLT 10 chains

The version of the beer game here adopted was developed in a Googledocs® software application which enables an Excel® spreadsheet to be shared by different SC players.

4. A numerical simulation of the beer game scenarios

Before we analyse the data to examine the effects of our manipulations, we coupled the human experiments (SBG, SLT and CD_SLT) with a simulation model in order to compute performance benchmarks corresponding to our treatments. Given the complexity of feedbacks and interactions in the beer game, we thought of little use to compare human behaviour with a theoretical benchmark consisting of a fully optimising, perfect-foresight agent. Rather, we envisaged a bounded rational buyer (Simon, 1957) who uses all the information provided during the game but has limited computational ability. We further assume that such a player minimises total costs of inventory related management (inventory holding costs plus backlog costs) and that she does not exhibit any cognitive bias in her perception of probabilities of the events and in her valuation of the pay-offs of the game, thus this player can be defined as risk and ambiguity neutral. In this sense, this buyer can provide plausible benchmark values for the standard deviations of orders and ensuing inventory holding and backlog costs, which can be compared with those obtained by humans who may display cognitive biases, including ambiguity aversion/preference.

With reference to the modelling of material flow, in all simulated scenarios, orders are filled from stock in a FIFO manner, with backordering used when stock-outs occur. Partial replenishments are permitted when there is not enough stock to fill an order completely. With stochastic lead times, the simulation model accommodates order crossover.

With reference to the information flow, supply chain nodes possess only local information and are "blind" to what is going on outside their level. Each node's supply chain knowledge-base is derived from the incoming demand flow from the downstream partner ($D_{i,t}$) and the outgoing flow of orders being placed with the upstream partner ($O_{i,t}$). We assume that each inventory manager (i) takes an ordering decision observing three variables: end of period inventories ($I_{i,t}$), downstream demand ($D_{i,t}$), and upstream replenishment ($R_{i,t}$).

Consistently with the human experiment, virtual players in the simulation are endowed with a starting inventory equal to three times the expected period demand of the external customer. Thus, starting inventory is higher than the equilibrium one, and inventory management must be performed by using an unsteady state model requiring a continuous control of the current inventory, and aimed at reaching a target inventory availability level. This latter differs from a policy based on mean demand, and it is based on the period by period comparison between upstream replenishment (R_{it}) and downstream demand (D_{it}). This policy allows managing the transition from the starting inventory to the target availability. At the end of each period, each buyer places replenishment order $O_{i,t}$ to raise or lower the inventory position to a target safety stock ($SS_{i,t}$) level. This target safety stock is meant to guarantee that $R_{i,t} \geq D_{i,t}$ for a proportion of periods corresponding to the expected level of service (LS).

As a consequence of this approach, the buyer considers $R_{i,t} - D_{i,t}$ as a random variable whose average value and standard deviation are as follows: $E_{Ri,t-Di,t} \sim 0$; $\sigma^2_{Ri,t-Di,t} = \sigma^2_{Di,t} + \sigma^2_{Ri,t}$.

Independency of $R_{i,t}$ and $D_{i,t}$ is assumed, either because of the stochastic replenishment lead time and the allowed order crossover, or because the virtual buyer lacks the cognitive ability to estimate the covariance matrix.

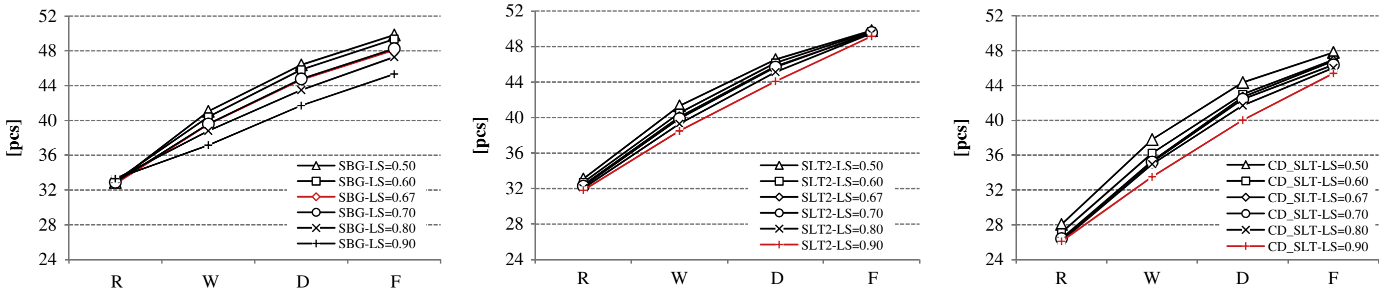

Fig. 1. Virtual player's standard deviations of orders (100 replications).

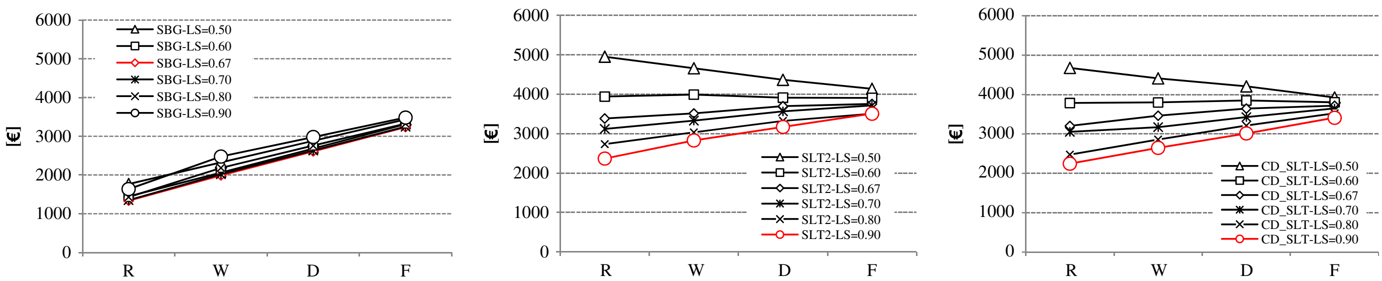

Fig. 2. Virtual player's total inventory costs (100 replications) (For interpretation of the references to colour in this figure, the reader is referred to the web version of this article.).

The inventory policy is an order-up-to level with review period, r, equal to 1, and target availability, TA_i, calculated as $TA_{i,t}=E_{Ri,t-Di,t}+SS_{i,t}$, where $SS_{i,t}=K(LS)\sigma_{Ri,t-Di,t}$, and $K(LS)$ denotes the generic factor which corresponds to the expected level of service. Further:

- $A_{i,t}=I_{i,t}+\sum_{k=1}^{t-1}O_{i,k}-\sum_{k=1}^{t-1}R_{i,k}$ is availability at the end of period t, before the replenishment order $O_{i,t}$ is placed; where $I_{i,t}=I_{i,t-1}+R_{i,t}-S_{i,t}$ is on-hand inventory which takes positive values; $\sum_{k=1}^{t-1}O_{i,k}$ is the sum of placed orders until the decision making period; $\sum_{k=1}^{t-1}R_{i,k}$is the sum of received replenishments from upstream level at the end of t period; S_{it} is the quantity shipped at each period:

$S_{i,t}=D_{i,t}+B_{i,t-1}$ if $I_{i,t-1}+R_{i,t}-D_{i,t}-B_{i,t-1}\geq 0$,
otherwise $S_{i,t}=I_{i,t-1}+R_{i,t}$
$B_{i,t}=B_{i,t-1}+D_{i,t}-S_{i,t}$ is the backlog quantity at time t;
$O_{i,t}=\max(0,TA_i-A_{i,t})$.

To estimate $E_{Ri,t-Di,t}$ and $\sigma^2_{Ri,t-Di,t}$ at each node, we assume that the virtual player uses all the historical information available (Chatfield et al., 2004). Finally, we assume that the decision maker optimises by choosing the level of service that minimises cumulated total costs. Such a level of service has been determined by means of a Monte Carlo simulation process by treatment (SBG, SLT, CD_SLT).

Fig. 1 report the effects of alternative levels of services (ranging from to 0.5 to 0.90) on the standard deviation of orders in each echelon (R=retailer, W=wholesaler, D=distributor, F=factory). We adopt the simplifying assumption that all echelons of the chains choose the same *LS*. Median standard deviation of orders resulting from 100 replications of 36 periods each are shown. Standard deviations per echelon decrease in the *LS*, but for the retailer.

Keeping the expected *LS* constant, standard deviations per echelon are only slightly higher in SLT. Since lead times in SLT are symmetrically distributed with mean equal to the distribution lead time in SBG, the standard deviations of orders by role end up being fairly similar in the two treatments. With constant demand and variable lead time (CD_SLT) standard deviations are lower than in the other two treatments for the three lower echelons of the chain, while they are of comparable magnitude to SBG for the Factory.

Fig. 2 displays the total costs in the three treatments, by *LS* and role. The curves corresponding to the cost minimising level of service are marked in red, and costs corresponding to this service level will be compared to experimental costs later in the paper.

In SBG costs are minimised for $LS=0.67$, implying that it is optimal to weight equally backlog and inventory holding costs over periods (following the assumption that unit backlog costs are double than unit inventory holding costs). With stochastic lead times (SLT and CD_SLT), cost minimisation requires an expected *LS* equal to 0.9, due to the player's imperfect ability to account for the complexity of the system created by stochastic lead times for all upstream echelons. In all treatments costs increase in echelon level, except in the case of $LS=0.5$ for SLT and CD_SLT. In fact, such a low *LS* brings about high backlog costs that are suffered especially by the downstream echelons of the chains. The minimised overall costs over a time horizon of 36 periods are fairly close in the three treatments for the Factory (3483 in SLT, 3410 in CD_SLT, and 3247 in SBG), due to the fact that the supplier of the Factory has always enough inventory availability to fulfil orders.

Fig. 3 reports the virtual player's cumulated orders by period and by role. For each treatment, the curves refer to the cost minimising levels of service. In all echelons the cumulated orders made by the virtual player in SBG are always lower than the corresponding cumulated orders in SLT, and the differences tend to increase moving upstream the SC. In CD_SLT the cumulated orders are always slightly lower than SLT. Thus, the numerical simulations imply that the cost minimising player should order more in the presence of stochastic lead times.

5. Results of human experiments

5.1. Evidence of bullwhip effect

We measure the bullwhip effect by the increase in the standard deviation of order quantities at each node (Chatfield

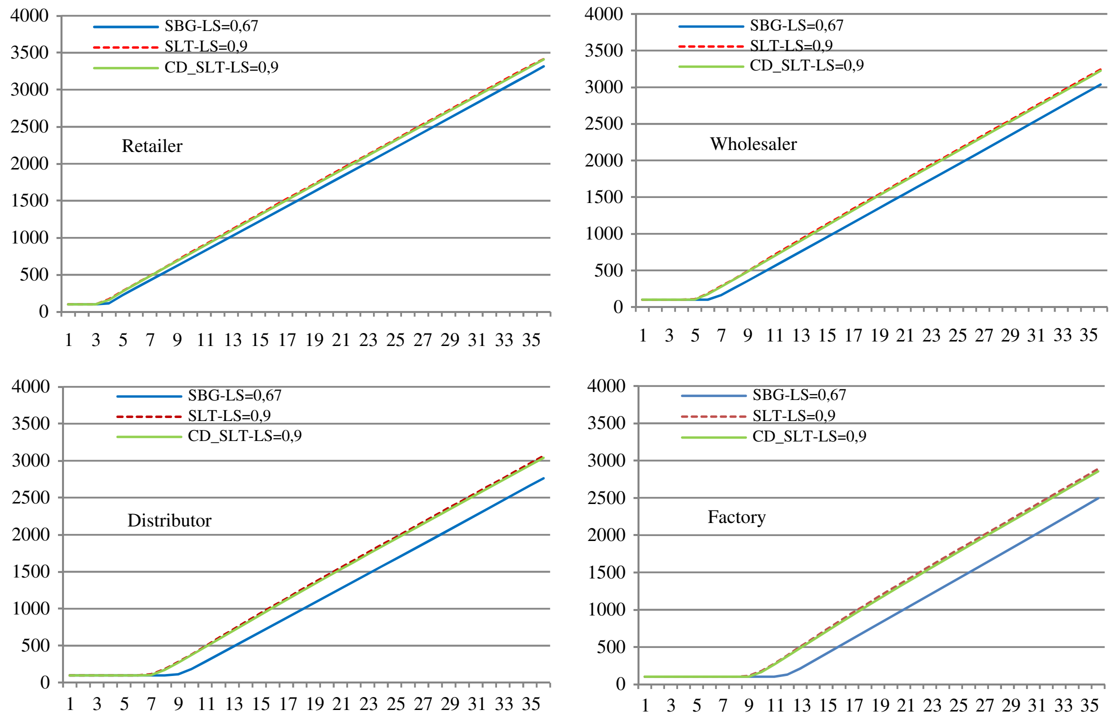

Fig. 3. Virtual player's cumulated orders by period and role (100 replications).

Table 2
Mean and median (in italics) standard deviation, by role and treatment.

	SBG	SLT	CD_SLT
R	19.8 *18.5*	41.1 *42.2*	34.3 *23.9*
W	40.3 *36.2*	53.4 *39.0*	51.9 *47.0*
D	47.7 *44.3*	90.6 *81.0*	59.5 *53.8*
F	62.2 *62.1*	108.0 *76.8*	68.5 51.8

Table 3
Wilcoxon test on amplification of standard deviations, and treatment.

	W vs. R	**D vs. W**	**F vs. D**	**Overall**
SBG	−2.395 **0.017**	−1.478 0.139	−2.497 **0.013**	−2.803 **0.005**
SLT	−1.600 0.110	−2.490 **0.013**	−0.978 0.328	−2.578 **0.01**
CD_SLT	−0.764 0.445	−**0.663** 0.508	−0.866 0.386	−1.886 0.059

et al., 2004). Table 2 summarises the results, by reporting mean and median standard deviation by role, in each cell of the experiment.

The BWE is clearly visible in all experimental treatments, including the one with constant demand (CD_SLT), in spite of the fact that players had acquired experience of the game in the warm-up session. In SBG mean and median standard deviations increase monotonically in echelon, with the ratio of the standard deviation of the factory to that of the retailer roughly equal to 3. With stochastic lead times the rate of increase of median values from one echelon to the next is smaller, with the median standard deviation of the factory being about twice that of the retailer (1.82 in SLT, 2.17 in CD_SLT). Table 3 reports the results of a Wilcoxon non parametric test of the statistical significance of variability amplification between contiguous echelons of the chains (R vs. W, W vs. D, and D vs. F) and the overall effect (R vs. F). The test shows that a significant overall effect of variance amplification is at play in SBG and SLT, although—as in other experiments, the difference between contiguous echelons is not always statistically significant (Croson and Donohue, 2006; Wu and Katok, 2006).

Comparison of standard deviations across treatments shows that the standard deviations of orders in SLT are statistically larger than in SBG for the retailer ($p<0.018$) and the distributor ($p<0.031$). Standard deviations in CD_SLT are larger than in SBG although the difference is not statistically significant. In summary, standard deviations in CD_SLT tend to be in between those in SBG and SLT.

These results are in line with the extant literature and with Hypothesis 1. Standard deviations are highest when variable demand is coupled with stochastic lead times, since players in this case find it more difficult to engage in the mental effort needed to compute outstanding orders and this gives rise to more chaotic behaviour. This conjecture is reinforced by the comparison of experimental standard deviations (Table 2) with those obtained from the numerical simulation (Fig. 1) which shows that with stochastic lead times human behaviour largely diverges from that of the virtual player, while the difference between predicted and observed standard deviations is smaller in SBG.

5.2. *Orders, inventory holdings and SC performance*

Several papers that study the supply chain through the beer game confine their analysis to the measurement of variance amplification of orders at each echelon, on the assumption that variance is directly correlated with SC costs. In actual facts, costs are also dependent on the dynamics of ordering behaviour throughout the game. For this reason, it is worthwhile to explore the pattern of orders through time in the different experimental

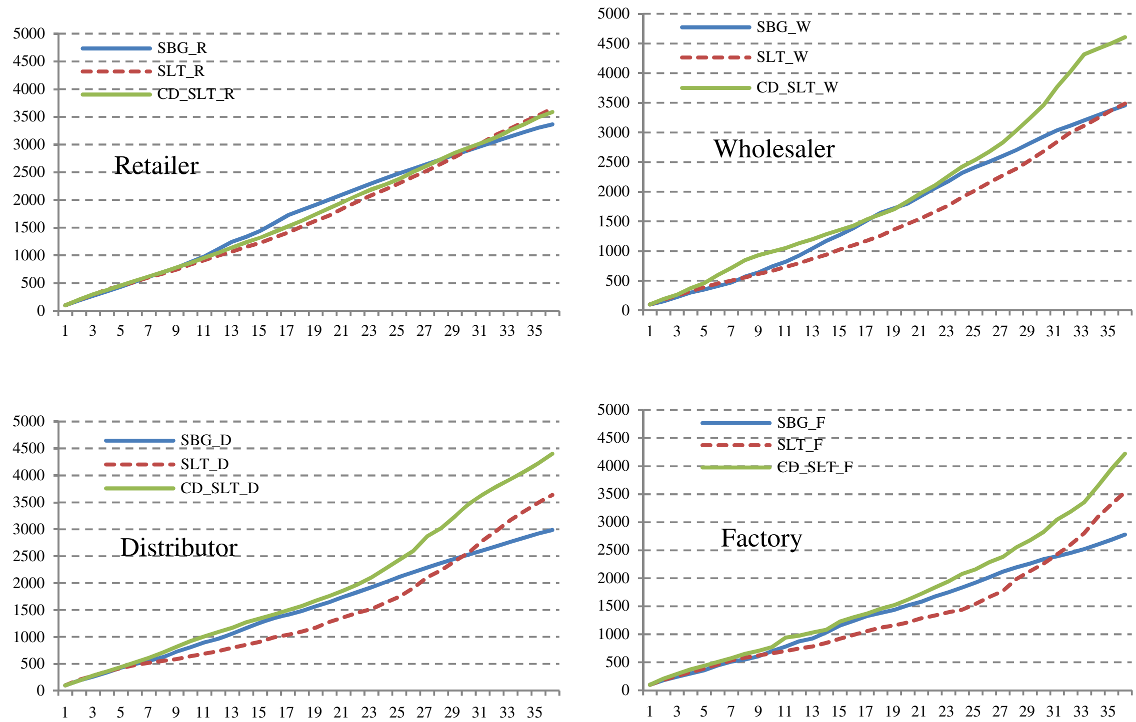

Fig. 4. Median experimental cumulated orders by role and by treatment.

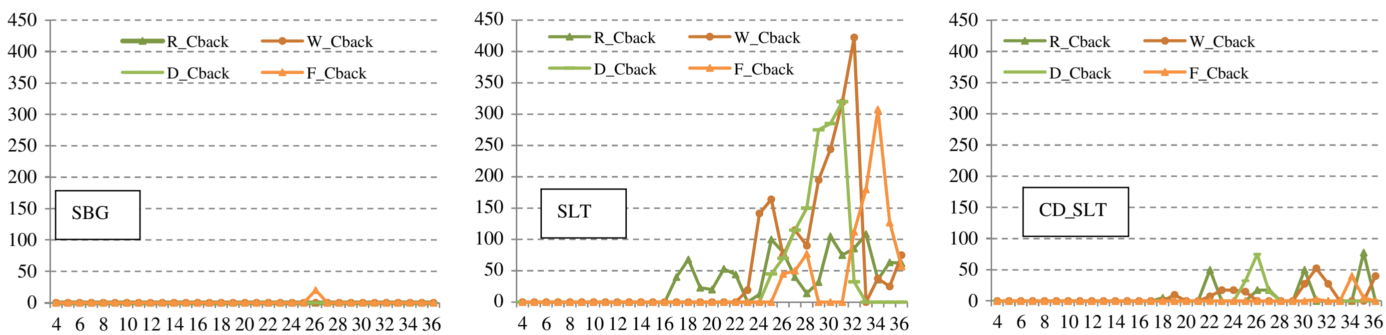

Fig. 5. Median experimental backlog costs (Cback) by role and by period.

treatments, and to compare ensuing inventory holding costs and backlog costs, in order to gain insight into how players manage inventory in the presence of different types of uncertainty (demand vs. supply).

Fig. 4 shows median cumulated orders up to period *t* in each of the three treatments and for each echelon. The figure allows assessing the impact of uncertainty, either in demand or in supply, on orders. To isolate the impact of demand variability, given stochastic lead times, we compare orders in CD_SLT with those in SLT. This comparison shows that orders are markedly higher in CD_SLT for all echelons except for the retailer, who observes the external demand and thus faces the lowest degree of demand uncertainty inside the chain. Since the only difference between the two treatments is demand variability, this result implies that uncertainty in demand leads to lower orders.

In order to identify the impact of stochastic lead times, we compare human behaviour in SLT and SBG. In SBG, orders are clearly higher than in SLT in most periods of the game. Total orders in SLT exceed those in SBG only in the very last leg of the game (after period 30) for the two uppermost layers of the chain. This change in behaviour can be attributed to the fact that, as backlogs build up, players react by ordering higher volumes. In fact, a plot of median backlog costs (across chains) by role in every period of the game (Fig. 5), shows that backlogs costs are steadily close to zero for all echelons of the chain in SBG, while they rise significantly in the second part of the game in SLT. The rise in backlogs is what induces players in SLT to increase orders in the last leg of the game.

Finally, Fig. 6 shows median inventory holding costs, backlogging costs, and total costs (backlog plus inventory holding) by role. Experimental values are shown by solid lines, whereas the two dashed lines represent minimised backlogging and inventory holding costs for a virtual player as obtained from the simulated models, and corresponding to the cost minimising service level of 0.67 in SBG and of 0.9 in SLT and CD_SLT. Higher total costs in SLT than in SBG are due to higher backlog costs in SLT. Total costs in SLT are also higher than in CD_SLT, but for the wholesaler.

Consistently with what already shown by the pattern of orders in Fig. 4, inventory holding costs in the three treatments suggest that an increase in uncertainty leads players to hold fewer inventories. In particular, when SLT and CD_SLT are compared, inventory holding costs are higher with constant demand, but for the Factory. Similarly, comparison between SLT and SBG shows that the addition of lead time uncertainty to a supply chain with variable demand leads the median player to hold fewer inventories. A Mann–Whitney non-parametric test comparing inventory holding costs in the various

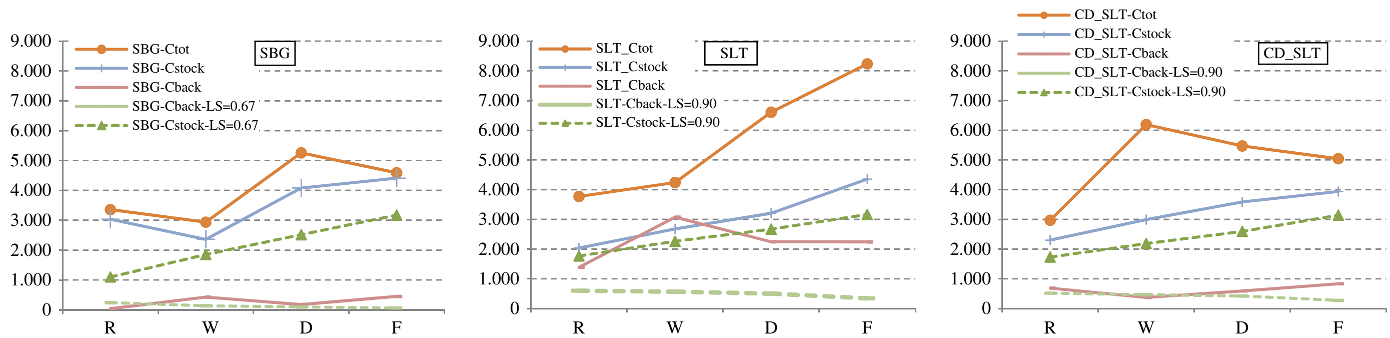

Fig. 6. Median inventory holding (Cstock), backlog (Cback), and total inventory costs (Ctot) by role.

treatments shows that inventory holding costs are statistically lower in SLT than in SBG for the retailer and the distributor ($R=-4.691$, $p<0.00$; $D=-3.798$, $p<0.00$). Differences between inventory holding costs in SLT and CD_SLT are never statistically significant.

5.3. Panel regression

In order to deepen our understanding of inventory management in the presence of stochastic lead times and to search for sound support for the result that players react to increases in ambiguity by holding fewer inventory, we built and estimated a dynamic panel model to track the determinants of orders in two of the experimental manipulations, namely SBG and SLT, which are directly comparable. Panel analysis allows capturing the richness of the longitudinal dimension of the data set (within player variations), while permitting to keep the heterogeneity among participants into account (between players variations). A different equation was estimated for each role of the chain.

In each equation, the dependent variable is current orders at time t, O_t, and the explanatory variables are current demand (D_t), current inventory (I_t), backlog pending at time t (B_t), past orders at times $t-1$, $t-2$, and t_3 (O_{t-1}, O_{t-2}, O_{t-3}), and two dummy variables, the first referring to periods 13–24 of the game (T2), and the second to orders taking place in the final leg of the game, i.e. periods 25–36 (T3).

We expect the coefficient of demand to be positive. In particular, if individuals place orders equal to the orders they receive, then the coefficient of the demand variable should equal 1 (Croson and Donohue, 2006). Inventory on hand and backlogs are introduced in the equation as two separate variables, although clearly non-zero inventory implies zero backlog and vice versa. The lagged dependent variables, O_{t-1}, O_{t-2}, and O_{t-3}, allow estimating the effect of orders still outstanding. Pending orders was not explicit information in both treatments, and players had to work it out for themselves using available data. Although outstanding orders do also depend on the backlogs of upstream suppliers, it is likely that players looked at past orders in order to estimate goods in transit. In SBG, O_{t-1} and O_{t-2} represent goods still in transit and thus the sign of their effect on current orders should be negative. If players underestimate the amounts in transit in the pipeline, as posited by Sterman (1989), we expect these coefficients to be not significantly different from zero. Thus, players may place orders in one period, but may not keep them into account when placing future orders. In SBG, O_{t-3} corresponds to the amount received at the beginning of the period and thus should lead to lower orders, whereas it may represent goods still in transit in SLT.

Finally, the dummy variables T2 and T3 capture any time trends present in the game and allow gathering further supportive evidence for the patterns observed in Fig. 4, i.e. a rise in orders in the end leg of the game.

Table 4a
Dynamic panel—Retailer.

	SBG_			SLT_		
	Coef.	**Std. err.**	**$P>z$**	**Coef.**	**Std. err.**	**$P>z$**
$O(T-1)$	0.102	0.061	0.097	0.166	0.061	0.007
$O(T-2)$	−0.059	0.054	0.281	−0.002	0.057	0.969
$O(T-3)$	−0.001	0.045	0.980	−0.008	0.055	0.889
$I(T)$	−0.195	0.024	0.000	−0.114	0.021	0.000
BKL(T)	−0.009	0.067	0.893	0.097	0.027	0.000
$D(T)$	0.182	0.042	0.000	0.262	0.091	0.004
T2	−15.946	2.973	0.000	4.735	5.844	0.418
T3	−26.619	3.765	0.000	8.063	7.651	0.292
Constant	114.787	14.377	0.000	62.682	13.88	0.000
Wald $\chi^2(8)$	308.69			136.98		
Prob. $>\chi^2$	0.0000			0.0000		

Table 4b
Dynamic panel—Wholesaler.

	SBG_			SLT_		
	Coef.	**Std. err.**	**$P>z$**	**Coef.**	**Std. err.**	**$P>z$**
$O(T-1)$	0.045	0.060	0.448	0.125	0.056	0.025
$O(T-2)$	−0.021	0.051	0.682	0.039	0.049	0.423
$O(T-3)$	−0.162	0.043	0.000	−0.003	0.046	0.936
$I(T)$	−0.287	0.033	0.000	−0.141	0.024	0.000
BKL(T)	0.392	0.118	0.001	0.092	0.016	0.000
$D(T)$	0.137	0.136	0.314	0.292	0.057	0.000
T2	−11.907	5.733	0.038	4.442	5.632	0.430
T3	−18.342	6.156	0.003	−5.601	7.609	0.462
Constant	133.559	18.811	0.000	59.527	10.721	0.000
Wald $\chi^2(8)$	218.80			264.93		
Prob. $>\chi^2$	0.0000			0.0000		

Table 4c
Dynamic panel—Distributor.

	SBG_			SLT_		
	Coef.	**Std. err.**	**$P>z$**	**Coef.**	**Std. err.**	**$P>z$**
$O(T-1)$	0.191	0.063	0.002	−0.112	0.056	0.046
$O(T-2)$	−0.199	0.052	0.000	0.001	0.053	0.995
$O(T-3)$	−0.093	0.041	0.023	−0.239	0.050	0.000
$I(T)$	−0.178	0.028	0.000	−0.083	0.054	0.127
BKL(T)	0.332	0.099	0.001	0.235	0.031	0.000
$D(T)$	0.2975	0.056	0.000	0.594	0.128	0.000
T2	−3.520	5.591	0.529	−11.503	14.796	0.437
T3	−24.552	7.079	0.001	23.501	18.095	0.194
Constant	124.530	19.919	0.000	62.885	23.473	0.007
Wald $\chi^2(8)$	425.11			177.03		
Prob. $>\chi^2$	0.0000			0.0000		

Table 4d
Dynamic panel—Factory.

	SBG_			SLT_		
	Coef.	**Std. err.**	$P>z$	**Coef.**	**Std. err.**	$P>z$
$O(T-1)$	−0.048	0.059	0.418	−0.020	0.054	0.709
$O(T-2)$	−0.400	0.055	0.000	−0.063	0.046	0.166
$O(T-3)$	0.097	0.037	0.009	−0.002	0.043	0.958
$I(T)$	−0.259	0.044	0.000	−0.242	0.031	0.000
BKL(T)	0.432	0.068	0.000	0.249	0.033	0.000
$D(T)$	0.47	0.071	0.000	0.337	0.062	0.000
T2	−14.840	6.634	0.025	−9.980	14.298	0.485
T3	−32.895	8.035	0.000	−10.037	18.142	0.580
Constant	131.279	20.719	0.000	117.484	16.829	0.000
Wald $\chi^2(8)$	541.34			390.19		
Prob. $>\chi^2$	0.0000			0.0000		

Estimation was achieved using the Arellano and Bond (1991) Generalised Method of Moments estimator for dynamic panels. Results are presented in Tables 4a–4d.

Estimates concerning retailers and wholesalers show that in SBG past orders are never statistically significant, confirming that, consistently with Sterman's findings, retailers underestimate goods in transit in the supply line. This is true, although to a lesser extent in SLT, in which orders lagged one period affect current purchases, consistent with a "recency" effect (Bostian et al., 2008). Inventory and demand are significant determinants of orders, but the coefficient of demand is well below 1, and slightly higher in the SLT game. Also, the time trend dummies are both negative and statistically significant in SBG but not in SLT, pointing to the fact that in SBG orders decrease in time. Finally, the constant term, which should indicate the optimal order-up-to level (Croson and Donohue, 2006), is higher in SBG than in SLT.

For distributors and factories in the majority of cases past orders have a significant impact on current orders (except for the factory in SLT). The sensitivity of orders to current demand also appears to be larger than that of downstream echelons. Again, there is evidence that the optimal order-up-to level is larger under deterministic than under stochastic lead times and that orders decrease in the last leg of the game in SBG.

Summing up, our analysis suggests that the following similarities exist between the two treatments: first, in both treatments there is some degree of underestimation of the supply line, which may be at the root of the observed BWE. This effect is on average less present in SLT, where, because of the variability in lead times, players tend to pay more attention to orders still outstanding. Second, in none of the equations we find evidence of individuals placing orders corresponding to current demand. Rather, it seems that they have an optimal order-up to level of inventory, which in the case of SBG is roughly equal to average demand plus one standard deviation, while it is sensibly smaller in the case of SLT (around 60 pieces, except for the factory). This latter result is in line with the descriptive analysis of orders and inventory holding costs discussed earlier in this section, thus confirming that players in SLT do not use inventory as a mitigation strategy for risk.

6. Discussion

The results of our study throw light on the way SC professionals may act in the face of supply uncertainty. Supply uncertainty (either due to stochastic lead times, or to uncertain supplier yield or to unexpected disruptions) is considered to be a challenge that many firms must address (Tomlin, 2009). Although inventory management is not the only response to supply uncertainty, it is still one of the most important risk mitigation strategies to counterbalance the possibility that a partial order is delivered or a delayed shipment occurs. While formal models provide guidance as to the normative response to such uncertainty, human experiments can provide evidence of whether those formal models are good predictors of observed behaviour, and whether learning eventually occurs.

In this direction, this study has brought together two strands of research, namely the investigation of the impact of supply uncertainty on inventory management and BWE, and the behavioural approach to operations management issues. In particular, the study has used human experimentation to identify whether behavioural biases exist in ordering and inventory decisions in a four echelons serial supply chain characterised by demand and/or supply uncertainty, and has tried to isolate the impact of human behaviour by comparing experimental results with the predictions of a simulated model. Here following, we discuss experimental results in detail, with the aim to provide explanations for their divergence from model predictions (if any), and to highlight the potential implications for theory and management.

6.1. Result 1—The Bullwhip Effect increases under supply uncertainty.

The numerical model predicts that the BWE should be only slightly higher when players face both supply and demand uncertainty (SLT) with respect to the case of demand uncertainty but reliable supply times (SBG), and that order variability should be lowest under supply uncertainty but constant demand (CD_SLT). The findings of the human experiments run contrary to this prediction. In fact, experimental order variability is significantly higher in SLT than in SBG. Also, stochastic lead times even with constant demand (CD_SLT) give rise to higher variability than in the SBG treatment.

According to the BWE literature (Chatfield et al., 2004), no order oscillations should arise when the distributions of demand and lead times are known, unless updating of policy parameters occurs on the basis of historic information. Thus, the fact that order amplification in SLT is higher than in SBG confirms Hypothesis 1 and also lends support to the fact that humans disregard the fact that they know the distribution of demand and lead times, and rely on past information in order to manage inventories.

Comparison between human behaviour and numerical simulation further shows that experimental order variability is higher than the predictions of the risk-neutral and ambiguity neutral model. In particular, while the difference between experimental and simulated standard deviations in SBG is fairly small, it becomes larger in CD_SLT and especially in SLT. Hence, while the model is a good predictor of order variability under low levels of uncertainty, when uncertainty increases, managers find it more difficult to identify the best inventory holding strategy, and end up behaving in a more chaotic fashion.

Since higher order variability is matched by higher overall inventory costs, the excess cost buyers suffer in SLT with respect to SBG provides a rough estimate of the expenditure firms would be willing to undertake in order to ensure the reliability of suppliers and the respect of shipment agreements. Likewise, the excess overall cost in SLT with respect to CD_SLT should indicate the willingness to spend to protect against demand oscillations.

6.2. Result 2—Human players in the beer game react to greater uncertainty by reducing orders

In a non-integrated SC, inventory buffers can potentially reduce the risk that the customer faces a stock-out as a consequence of supply uncertainty. However, firms face every day several causes of uncertainty in supply that often co-exist with demand volatility, leading to complex scenarios to manage and to higher costs of inventory management (Schmitt et al., 2010).

The combined analysis of the numerical simulation and of human behaviour gives insight into whether and how inventory management deviates from behaviour predicted by models, and into how supply chain members react to higher degrees of uncertainty.

First, the comparison between purchases predicted by the numerical model (Fig. 3) and actual experimental orders (Fig. 4) within each treatment shows that the human player buys more stock than the virtual player. Since the virtual player was assumed to be risk neutral, this entails that human players behave as risk averse, a feature that can be reconciled with the psychological and economic literature (Rabin and Thaler, 2001).

Second, in order to measure the impact on human behaviour of increases in the degree of uncertainty by separating at the same time the effect of the source of uncertainty, we performed a pairwise analysis of the three treatments. In particular, the comparison between human behaviour in SLT and in SBG allows isolating the impact of supply uncertainty, whereas the difference between SLT and CD_SLT can measure the impact of demand uncertainty. The numerical model suggests that when demand uncertainty and supply uncertainty co-exist (SLT), cost minimisation calls for higher orders than in the case in which only demand is stochastic (SBG). Also, orders are expected to be higher in SLT with respect to deterministic demand and stochastic lead times (CD_SLT). Our experimental results show that an increase in supply uncertainty (from SBG to SLT) reduces orders. Likewise, orders decrease when demand uncertainty increases (from CD_SLT to SLT). Median orders in the human experiment are lowest in SLT and this contrasts with the predictions of the numerical model, in which orders should be highest in SLT and lowest in SBG. As a consequence of human behaviour, backlog costs in SLT are higher than in the other two treatments not only because of the more pronounced BWE, but also due to the smaller inventory buffer built by players. Intuition may suggest that in SLT lower inventory holding costs are due to more frequent backlogs and not vice versa: once nervousness starts seeping in the system and the BWE gets "cascading" effects going, it becomes difficult to replenish inventories. However, this argument only partially explains such behaviour. In fact, not only cumulated purchases by role and by period show that in SLT players order less than in SBG and CD_SLT for most periods, indicating a deliberate choice to hold fewer inventories, but this finding is also reinforced by regression analysis, which shows that the order-up-to level is smaller in SLT for all echelons.

Taken together, these results confirm Hypothesis 2, and suggest that there is a general bias in players' decisions: individuals react to higher uncertainty by protecting relatively less, whatever the source of the increase in uncertainty.

From a behavioural perspective our results can be reconciled with an ambiguity loving attitude. Under either demand or supply uncertainty, a buyer perceives inventory as an effective way to protect against the risk of a stock-out, whereas in SLT her ability to protect becomes vaguer because of the combined effects of downstream and upstream uncertainty. The decision problem becomes ambiguous, and ambiguity makes it difficult for players to identify reliable defensive strategies. In this setting, inventory holding costs are perceived as certain costs with ambiguous benefits, and this ends up reducing the perceived marginal benefit of investing in inventory to protect against supply failures. As a consequence, ambiguity induces players to "bet" on the possibility that a late delivery of goods will not occur and to hold fewer inventories.

To recap, the net effect of human behaviour with respect to model predictions is to hold higher safety stock with respect to the benchmark, and to decrease stock in the face of higher uncertainty.

Third, comparison of experimental costs with simulated costs shows that while with deterministic lead times actual backlog costs follow expected backlogs closely, the size of the divergence is larger when lead times are stochastic, and especially when stochastic lead times are coupled with stochastic external demand. This is a direct consequence of the higher variability of orders and inventories across periods in SLT, i.e. the BWE. Further, in SLT the difference between experimental and predicted total inventory costs increases in echelon. The implication of such result is that the costs of the inability to handle demand and supply uncertainty are mostly suffered by the two uppermost echelons of the chain. Although the factory and the distributor are subject to less supply uncertainty than the downstream layers, they end up bearing the most of it. This is due to the model assumption that the stochastic lead time happens at all four echelons, so that for the factory it becomes relevant to keep into account not only the deliveries of the internal supplier but also of the downstream buyers. In fact, for the factory, the stochastic lead time in the downstream echelons of the chain translates in an increase in the variability of demand from the retailer upward.

Our results have relevant implications for SC management. If we take the SBG and CD_SLT to represent routine SC management problems and the SLT to stand for a non-routine or more turbulent situation, our results imply that purchasing managers neglect their chance to control for high turbulence, and that this effect is stronger for the uppermost layers of the supply chain. Since the end result of this behaviour is cost escalation, information sharing may be beneficial in these settings especially for the upper echelons, which appear to bear the greater share of cost increases.

Further, our results point out the importance of planning inventory holdings timely in order to avoid stock-outs. In fact, both the simulation and the human experiment show that when a stock-out occurs under variable demand and stochastic lead times (SLT), this is bound to persist for longer than in the other scenarios, because in a non-coordinated SC the variability of orders tends to increase and the system becomes chaotic. Thus, loss of control by managers represents a further implication of the ambiguity seeking attitude and optimistic bias it entails. This may suggest that managers should be trained to discount their expectations of success by removing the optimistic bias that leads to risk taking.

Although we did not test this finding rigorously, since the variability in demand and the variability in supply are not strictly equivalent, the comparison between CD_SLT and in SBG can potentially provide interesting insight into how players react to different types of uncertainty. The prediction of the numerical model that orders are higher with stochastic lead times (CD_SLT) than with stochastic demand (SBG) is confirmed by the experimental results. This may entail that decision makers perceive the need for higher protective investment against the risk of stock-outs under supply uncertainty than under demand uncertainty. However, as already discussed above, orders in SLT are fewer than in SBG. One possible interpretation of these findings is that when the degree of uncertainty is perceived to be manageable (being confined to either demand or supply uncertainty), the decision maker reacts by adopting standard defensive strategies (inventory holdings), but when uncertainty becomes "deep", decision makers prefer to run the risk of higher backlog costs, rather than protecting by incurring higher inventory holding costs.

7. Conclusions

Supply uncertainty is a major issue for inventory management in serial supply chains and lead-time uncertainty is one of the most prominent and common of its components. A potentially fruitful approach to this issue is to bring together the results and methodologies of two separate streams of literature. The first comprises behavioural studies that have long demonstrated that cognitive biases are pervasive in choice under uncertainty and have identified systematic deviations such as aversion or preference for ambiguity. The second comprises numerical simulation studies of the expected impact of supply uncertainty (and lead time uncertainty in particular).

This paper has attempted to apply this approach and has investigated the impact of uncertainty in lead times on the formation of the BWE, inventory holding and SC performance through a controlled human experiment, whose results have been compared with those of a numerical model simulating a risk neutral, ambiguity neutral player.

The paper positions itself among the contributions that aim at studying the relevance that SC vulnerability plays in SC management: in the experiment, lead time uncertainty in the presence of a single supplier can be considered a proxy of the vulnerability of the SC to stock-outs.

Although the experiment has not been carried out with professional managers, the fact that participants were graduate students with background in Operations Management makes it reasonable to assume that these results may be descriptive of the behaviour of actual purchasing managers.

The following limitations of the present study should be underlined: uncertainty in supply has been made operational through a specific distribution of lead-time, thus it is possible that different levels of uncertainty or different distributions give rise to different behaviour. In this vein, an interesting extension of the research may consider introducing stochastic lead times at only one echelon rather than at all four layers.

Next, the possibility to link inventory holding to the individual ambiguity attitude has been restricted by the choice to manipulate the presence/absence of stochastic lead times on a between subject basis. However, the indubitable gain of the between subject manipulation is that the observed differences between players' behaviour across treatments are free from potential framing effects created by the comparative setting.

As part of our future research agenda we plan to address these shortcomings and to gain sound support for the effect of various facets of supply risk on SC performance.

Acknowledgements

We thank two anonymous referees, Ram Narasimhan and Richard Pibernik for useful suggestions and criticism on earlier versions of this paper. Financial support by a M.I.U.R. grant is gratefully acknowledged.

References

Abdellaoui, M., Baillon, A., Placido, L., Wakker, P., 2010. The rich domain of uncertainty: source functions and their experimental implementation. American Economic Review 101 (2), 695–729.

Ancarani, A., Di Mauro, C., 2011. The human side of supply chains: a behavioural perspective on supply chain management, In: Evangelista, P., McKinnon, A., Sweeney, E., Esposito, E., (Eds). Supply Chain Innovation for Competing in Highly Dynamic Markets: Challenges and Solutions. IGI Global Publications, Hershey PA, pp. 290–314.

Arellano, M., Bond, S., 1991. Some tests of specification for panel data: Montecarlo evidence and an application to employment equations. Review of Economic Studies 58 (2), 277–297.

Bendoly, E., Donohue, E., Schultz, K.L., 2006. Behaviour in operations management: assessing recent findings and revisiting old assumptions. Journal of Operations Management 24 (6), 737–752.

Bernoulli, D., 1738. Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientarum Imperialis Petropolitanae 5, 175–192.

Bostian, A.J.A., Holt, C.A., Smith, A.M., 2008. Newsvendor "pull-to-center" effect: adaptive learning in a laboratory experiment. Manufacturing & Service Operations Management 10 (4), 590–608.

Camerer, C.F., Weber, M., 1992. Recent development in modeling preferences: uncertainty and ambiguity. Journal of Risk and Uncertainty 5 (4), 325–370.

Chatfield, D.C., Kim, J.G., Harrison, T.P., Hayya, J.C., 2004. The bullwhip effect—impact of stochastic lead time, information quality, and information sharing: a simulation study. Production and Operations Management 13 (4), 340–353.

Chen, F., 1999. Decentralized supply chains subject to information delays. Management Science 45, 1016–1090.

Cohen, M., Jaffray, J.Y., Said, T., 1985. Individual behaviour under risk and under uncertainty: an experimental study. Theory and Decision 18, 203–228.

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

Di Mauro, C., Maffioletti, A., 2004. Attitudes to risk and attitudes to uncertainty: experimental evidence. Applied Economics 36, 357–372.

Einhorn, H., Hogarth, R., 1985. Ambiguity uncertainty in probabilistic inference. Psychological Review 92, 433–439.

Einhorn, H., Hogarth, R., 1986. Decision making under ambiguity. Journal of Business 59, 225–250.

Ellis, S.C., Henry, R.M., Shockley, J., 2010. Buyer perception of supply disruption risk: a behavioural view and empirical assessment. Journal of Operations Management 28, 34–46.

Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. The Quarterly Journal of Economics 75 (4), 643–669.

Esposito, E., Passaro, R., 2009. The evolution of supply chain relationships: an interpretative framework based on the Italian inter-industry experience. Journal of Purchasing and Supply Management 15 (2), 114–126.

Forrester, J., 1958. Industrial dynamics: a major breakthrough for decision makers. Harvard Business Review 36, 37–66.

Geary, S., Childerhouse, P., Towill, D.R., 2002. Uncertainty and the seamless supply chain. Supply Chain Management Review July/August, 52–61.

Geary, S., Disney, S.M., Towill, D.R., 2006. On bullwhip in supply chains—historical review, present practice and expected future impact. International Journal of Production Economics 101, 2–18.

Gigerenzer, G., 2002. Reckoning With Risk: Learning to Live With Uncertainty. Allen Lane, The Penguin Press, New York.

Gupta, S., Steckel, J.H., Banerjii, A., 2002. Dynamic decision making in marketing channels. In: Rapoport, A. (Ed.), Experimental Business Research. Kluwer Academic Publishers, Boston, pp. 4–21.

Halevy, Y., 2007. Ellsberg revisited: an experimental study. Econometrica 75 (2), 503–536.

Harland, C., Brenchley, R., Walker, H., 2003. Risk in supply networks. Journal of Purchasing and Supply Management 9 (2), 51–62.

Hazen, B.T., Overstreet, R.E., Jones-Farmer, L.A., Field, H.S., 2012. The role of ambiguity tolerance in consumer perception of remanufactured products. International Journal of Production Economics 135 (2), 781–790.

Heydari, J., Kazemzadeh, R.B., Chaharsooghi, S.K., 2009. A study of lead time variation impact on supply chain performance. International Journal of Advanced Manufacturing Technology 40, 1206–1215.

Hogarth, R.M., Einhorn, H.J., 1990. Venture theory: a model of decision weights. Management Science 36 (7), 780–803.

Hogarth, R.M., Kunreuther, H., 1985. Ambiguity and insurance decisions. American Economic Review 75 (2), 386–390.

Hogarth, R., Kunreuther, H., 1989. Risk, ambiguity, and insurance. Journal of Risk and Uncertainty 2, 5–35.

Holsapple, C.W., Singh, M., 2001. The knowledge chain model: activities for competitiveness. Expert Systems with Applications 20 (1), 77–98.

Jammernegg, W., Kischka, P., 2009. Risk preferences and robust inventory decisions. International Journal of Production Economics 118 (1), 269–274.

Kahneman, D., Tversky, A., 1974. Judgement under uncertainty: heuristics and biases. *Science*, New Series 185 (4157), 1124–1131.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47 (2), 263–292.

Kim, J.G., Chatfield, D.C., Harrison, T.P., Hayya, J.C., 2006. Production, manufacturing and logistics quantifying the bullwhip effect in a supply chain with stochastic lead time. European Journal of Operational Research 173 (2), 617–636.

Knight, F.H., 1921. Risk, Uncertainty, and Profit. Houghton Mifflin, New York.

Muthukrishnan, A.V., Wathieu, L., Xu, A.J., 2009. Ambiguity aversion and the preference for established brands. Management Science 55 (12), 1933–1941.

Rabin, M., Thaler, R.H., 2001. Anomalies: risk aversion. The Journal of Economic Perspectives 15 (1), 219–232.

Rigotti, L., Rayan, M., Vaithianathan, R., 2011. Optimism and firm formation. Economic Theory 46 (1), 1–38.

Savage, L.J., 1954. Foundations of Statistics. Wiley, New York.

Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica 57, 571–588.

Schmitt, A.J., Snyder, L.V., Shen, Z.-J.M., 2010. Inventory systems with stochastic demand and supply: properties and approximations. European Journal of Operational Research 206 (2), 313–328.

Simon, H., 1957. Models of Man. Wiley, New York, USA.

Steckel, J.H., Gupta, S., Banerji, A., 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., 1989. Modeling managerial behaviour: misperceptions of feedback in a dynamic decision making experiment. Management Science 35 (3), 321–339.
Tomlin, B., 2009. Impact of supply learning when suppliers are unreliable. Manufacturing & Service Operations Management 11 (2), 192–209.
Truong, T.H.D., Luong, H.T., Kim, Y., 2008. A measure of the bullwhip effect in supply chains with stochastic lead time. International Journal of Advanced Manufacturing Technology 38, 1201–1212.
Tversky, A., Wakker, P., 1995. Risk attitudes and decision weights. Econometrica 63 (6), 1255–1280.
Viscusi, K., Chesson, H., 1999. Hopes and fears: the conflicting effects of risk ambiguity. Theory and Decision 47, 153–178.
Wakker, P., 2010. Prospect Theory: for Risk and Ambiguity. Cambridge University Press, Cambridge, UK.
Wagner, S.M., Neshat, N., 2010. Assessing the vulnerability of supply chains using graph theory. International Journal of Production Economics 126 (1), 121–129.
Wu, D., Katok, E., 2006. Learning, communication, and the bullwhip effect. Journal of Operations Management 24 (6), 839–850.