

Behavioral Analyses of Retailers' Ordering Decisions

Sebastián Villa

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Ph.D in Management

Supervised by

Prof. Paulo Gonçalves, Ph.D.

Università della Svizzera italiana

Facoltà di scienze economiche, Institute of Management

Lugano

October 2016

To my wife, Gloria

Notes on Software and Documentation

The experimental settings in chapters 2 was programmed and run by the author using Powersim Studio. The experimental settings in chapters 3 and 4 were programmed and run by the author in z-Tree (Fischbacher, 2007). The author used Microsoft Office Excel to compile all the experimental results. The system dynamics model in chapter 3 was implemented in Vensim DSS. The econometric analyses were computed by the author using Stata version 12 and R.

Table of Contents

Chapter 1. Introduction	15
Chapter 2. Exploring Retailers' Ordering Decisions under Delays	19
2.1. Introduction	19
2.2. Model Description	22
2.2.1. Cost Objective	23
2.3. The Experiment	24
2.3.1. Experimental Treatments	25
2.3.2. Experimental Protocol	26
2.3.3. Optimal Simulated Trajectory	27
2.4. Results	27
2.4.1. Subjects' Order Decisions Behavior	27
2.4.2. Subjects' Cost Performance	29
2.5. Modeling Decision Rules	32
2.5.1. Least Squares Analysis	34
2.5.2. Panel Data Analysis	38
2.6. Conclusion	42
Chapter 3. Behavioral Analysis of the Effect of Duplication Orders in Single-Supplier Multi-Retailer Supply Chains	51
3.1. Introduction	51
3.2. Formal Model	54
3.2.1. Retailers' Model	55
3.2.2. Supplier's Model	56
3.3. The Experiment	58
3.3.1. Experimental Protocol	58
3.3.2. Study I: Stationary and known demand under retailer competition	59
3.3.3. Results Study I	61
3.3.4. Study II: Duplications and time to build supplier's capacity under retailer competition	63

3.3.5. Simulated trajectory benchmark	64
3.3.6. Results Study II	66
3.3.7. Study III: Proportional vs. Turn-and-earn allocation under retailer competition	69
3.3.8. Results Study III	70
3.4. General results	71
3.4.1. Average Costs	71
3.4.2. Retailers' Decision Rule	73
3.4.3. Evaluation of decision rule	74
3.5. Conclusions	78
Chapter 4. Transshipments in Supply Chains: Beyond the Analytical Models	83
4.1. Introduction	83
4.2. Literature Review	85
4.3. Analytical Background	88
4.3.1. Isolated Newsvendor	88
4.3.2. Newsvendor problem with transshipments	88
4.4. Base Case Experiment (BC): Human vs. Human	90
4.4.1. Design of the BC experiment	91
4.4.2. Analyses and results of the BC experiment	92
4.4.2.1. Analytical estimations	92
4.4.2.2. Behavioral analyses	93
4.4.2.3. Development of Behavioral Models	97
4.4.2.4. Estimation of the Behavioral Models	98
4.5. Practical improvements to the system	100
4.5.1. Communication Experiment (C): Human vs. Human	100
4.5.1.1. Design of the C experiment	100
4.5.1.2. Analyses and results of the C experiment	101
4.5.2. Best Response Experiments (BR): Human vs. Computer	104
4.5.2.1. Design of the BR experiment	106
4.5.2.1. Analyses and results of the BR experiment	107
4.6. Conclusions	113

Chapter 5. General Conclusions 123

5.1. Main Contributions..... 123

5.2. Future work 125

List of Figures

Figure 2.1. Supply Chain Structure	22
Figure 2.2. Overview of model structure.....	23
Figure 2.3. Typical experimental results (P_j indicates the subject ID with $j=1, \dots, 15$).	28
Figure 2.4. Final customer demand, optimal and average subjects' orders (AO) in each treatment. ...	29
Figure 2.5. Box plots with the coefficient estimations using least squares, expected values (triangle), and panel data estimations (circle).	40
Figure 2.6. Average subjects' orders (AO), simulation of the proposed heuristic (Heuristic), adjusted model with panel parameters (Panel).	41
Figure 3.1. Structure for a single-supplier two-retailer supply chain with order duplications.	55
Figure 4.1. Transshipments scenarios	89
Figure 4.2. Transshipments best response policy and Nash equilibrium order	90
Figure 4.3. Structure for a single-supplier multiple-Newsvendor supply chain with transshipments..	91
Figure 4.4. Optimal isolated Newsvendor orders, best response function and Nash equilibria for a) High Profit, and b) Low Profit conditions.....	93
Figure 4.5. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) for a) T1, and b) T2	95
Figure 4.6. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) for a) T3, and b) T4.....	103
Figure 4.7. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) T5-T10	112

List of Tables

Table 2.1. Experimental treatments	26
Table 2.2. Total cumulative, average, and optimal costs across treatment for the experiment	30
Table 2.3. Costs distribution given by Orders and Supply gap	32
Table 2.4. Coefficient estimates	34
Table 2.5. Coefficient estimates of decision rule for each individual for all treatment.....	35
Table 2.6. Coefficient estimates of decision rule for treatment as panel data	39
Table 3.1. Experimental treatments Study 1	60
Table 3.2. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas) and Customer Demand (solid line)	62
Table 3.3. Experimental treatments Study II	64
Table 3.4. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas), Benchmark trajectory (solid line).....	65
Table 3.5. Mean comparisons among treatment variables.....	67
Table 3.6. Standard deviation of subjects' orders by experimental treatment.....	68
Table 3.7. Experimental treatments Study II	69
Table 3.8. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas), Benchmark trajectory (solid line).....	70
Table 3.9. Average, Ordering, Inventory and Backlog cost by experimental variable.....	72
Table 3.10. Parameter estimations for the decision rule.....	74
Table 4.1. Performance measures for treatments T1 and T2	94
Table 4.2. Panel Data estimations for the BC treatments	96
Table 4.3. Structural estimation of Behavioral Models	99
Table 4.4. Experimental treatments for the C experiment.....	101
Table 4.5. Performance measures for treatments T3 and T4	102
Table 4.6. Panel Data estimations for the C treatments	104
Table 4.7. Experimental treatments for the BR experiment	106

Table 4.8. Performance measures for treatments T5, T7 and T9 – High Profit Treatments..... 108

Table 4.9. Performance measures for treatments T6, T8 and T10 - Low Profit Treatments 111

Table 4.10. Panel Data estimations for the BR treatments 112

List of Appendices

Appendix 2.1. Complete Model specifications	46
Appendix 2.2. Instructions for T3 – Agile-Supplier Treatment (Translated into English)	47
Appendix 2.3. Experiment Environment.....	49
Appendix 2.4. Interface of the experiment in Powersim (in Translated into English).....	49
Appendix 2.5. Linearization of Heuristic	50
Appendix 3.1. Interface of the experiment in Z-tree (in Spanish).....	81
Appendix 3.2. <i>p</i> -values for comparison of orders’ means for treatments in Study II	81
Appendix 3.3. <i>p</i> -values for comparison of orders’ standard deviations for treatments in Study II.....	82
Appendix 3.4. Overall experimental treatments.....	82
Appendix 4.1. Finesse analytical solution for a Newsvendor problem with transshipments	116
Appendix 4.2. Interface of the experiment in Z-tree (in Spanish).....	118
Appendix 4.3. Instructions Base Case – High Profit condition (in Spanish)	119
Appendix 4.4. <i>p</i> -values for the t-tests comparing the performance measures computed for treatments T1 and T5 to 10	120

Chapter 1. Introduction

Previous research on behavioral operations has focused on describing decision making biases and deriving heuristics that aim to explain those biases in a lineal supply chains (Croson et al., 2014; Sterman & Dogan, 2015; Sterman, 1989a) or under a simple newsvendor framework (Bolton & Katok, 2008; Schweitzer & Cachon, 2000). However, limited behavioral work has been done on analyzing interactions among multiple retailers and on understanding how to take advantage of subjects' behaviors to create policies that lead to better supply chain performance. Following this gap in the behavioral operations literature, the main objective I pursue in this thesis is to better understand how different factors may independently and in combination influence retailers ordering decisions under different supply chain structures (single agent and multi agent), different demand uncertainty (deterministic and stochastic), and different interaction among retailers (no interaction, competition and cooperation). I developed three different studies that allow me to better understand the main dynamics and biases around the ordering decisions in different supply chain structures.

One of the main topics that I discuss in the first two chapters of the thesis refers to order amplifications. Amplifications usually take place in supply chains with tight capacity. Under scarce supply, the supplier rations the allocation of available supply to satisfy retailers' orders, while retailers receiving only a fraction of previous orders, amplify future ones in an attempt to secure more units (Lee et al., 1997a, 1997b). The amplification of retailers' orders creates problems such as excessive supplier capital investment, inventory gluts, low capacity utilization, and poor service, among others (Armony & Plambeck, 2005; Gonçalves, 2003; Lee et al., 1997a; Sterman, 2000). Cisco System's 2001 inventory write-off provides an instructive practical example.

Behavioral research in order amplification has focused mainly on understanding the biases and underperformances presented in a typical serial supply chain (e.g. The Beer Game) (Croson & Donohue, 2005; Sterman & Dogan, 2015; Sterman, 1989a). However, limited work has been done on analyzing (i) the effect of different types and magnitudes of delays, (ii) the interactions among competing retailers, and (iii) the effect of different allocation mechanism in subjects' behavior.

In each of the chapters of this thesis, I aim to understand how people make their ordering decisions when they assume the role of retailers. Therefore, I created different decision-making-laboratory experiments following standard experimental economics protocol (Friedman & Sunder, 1994, 2004; Katok, 2011). Initially, I developed a formal model incorporating standard operations management processes for (i) supplier's capacity investment, (ii) retailers' inventory management, and (iii) final customers' orders. Second, I ran decision making laboratory experiment based on the developed models to study how human subjects playing the role of retailers make ordering decisions. Based on subjects' decisions and on the system dynamics, I used econometric methods to analyze the results obtained and shed light on the decision patterns used by subjects in the retailer role. Finally, subjects' decisions were compared against some theoretical benchmarks to determine subjects' performance in each experiment.

In 0, I analyze order amplification in a single-supplier single-retailer supply chain. In this chapter, I used a behavioral experiment to test retailers' orders under different *ordering delays* and different *times to build supplier's capacity*. Results provide (i) a better understanding of the endogenous dynamics leading to retailers' ordering amplification, (ii) a description of subjects' biases and deviation from optimal trajectories, despite subjects have full information about the system structure, and (iii) some practical implications and recommendations that may lead to an increase on supply chain performance.

In Chapter 3, I analyze how the amplification of orders can also take place when there is fierce retailer competition and limited supplier capacity. For this study, I built on Armony and Plambeck (2005)'s analytical work on the impact of duplicate orders on upstream suppliers' demand estimation and capacity investment. I study how different factors (different time to build supplier capacity, different levels of competition among retailers, different magnitudes of supply shortage and different allocation mechanisms) may independently and in combination influence retailers' order in a system with *two* retailers under supply competition. Results show that (i) the bullwhip effect persists even when subjects do not have incentives to deviate and that the order amplification do not disappear over time, (ii) subjects amplify their orders in an attempt to build an unnecessary safety stock to respond to potential deviations from the other retailers, (iii) subjects' biases do not increase when subjects face systems with higher complexity, and (iv) retailers' underperformance varies with the allocation mechanism used by the supplier.

In 0 and Chapter 3, I consider a system where retailers need to control their inventory level when final customer demand is assumed to follow a known pattern and when retailers had the opportunity to store their inventory over time. In Chapter 4, I remove the ability to indefinitely store inventory, due to perishability or obsolescence of the product, and I include uncertainty in the final customer demand. This problem is commonly known as the newsvendor problem (Arrow et al., 1951). The newsvendor problem characterizes situations where a retailer needs to decide how many units to order to the supplier to satisfy an uncertain final customer demand. In this case, both leftovers and shortages at the end of the selling period are costly. Previous research in behavioral operations on the newsvendor problem has focused mainly on describing decision making biases and/or deriving heuristics that aim to explain those biases in a single actor problem (Bolton & Katok, 2008; Bostian et al., 2008; Croson & Ren, 2013; Schweitzer & Cachon, 2000). However, limited work has been done (i) on analyzing interactions among *multiple* subjects and (ii) on understanding subjects' behaviors as a way to create better interaction policies that could improve supply chain coordination. I contribute to this literature by experimentally exploring the effect of *transshipments* among retailers in a single-supplier multi-retailer supply chain. Specifically, I explore retailers' orders under different profit (Schweitzer & Cachon, 2000) and communication conditions (Ahn et al., 2011). Finally, I integrate analytical and behavioral models to improve supply chain performance. Results show that (i) the persistence of common biases in a newsvendor problem (pull-to-center, demand chasing, loss aversion, psychological disutility), (ii) communication could improve coordination and may reduce demand chasing behavior, (iii) supply chain performance increases with the use of behavioral strategies embedded within a traditional optimization model, and (iv) dynamic heuristics improve overall coordination, outperforming a simple Nash Equilibrium strategy.

Chapter 2. Exploring Retailers' Ordering Decisions under Delays

(with Paulo Gonçalves and Santiago Arango)

Abstract

When final customer demand exceeds available supply, retailers often hedge against shortages by inflating orders to their suppliers. While this amplification in orders is clearly described in the literature, there is little experimental research quantifying the factors influencing these amplifications. We use an experiment to test subjects' ordering decisions under different ordering and supplier's capacity acquisition delays. Subjects in the experiment display limited ability to process the impact of delays and feedback. The order trajectories follow a pattern of overshoot and subsequent undershoot until reaching an equilibrium. However, the initial overshoot is less intense and lasts longer than the optimal behavior, when subjects face longer delays. In addition, subjects inflate their orders when the supplier faces longer capacity acquisition delays and when orders take longer to be perceived by the supplier. Econometric estimates show that the proposed anchoring and adjustment heuristic is a possible heuristic for explaining subjects' ordering behavior.

Keywords: Order Amplification, Laboratory Experiment, Behavioral Operations, Supply Chain Management, Demand Bubbles, System Dynamics.

2.1. Introduction

One of the most common and costly problems in supply chains is caused by retailer orders' amplification (Armony & Plambeck, 2005). These amplifications have been captured in the literature as early as 1924, when Mitchell described the case of retailers inflating their orders to manufacturers when competing with other retailers for scarce supply. He argued "if [retailers] want 90 units of an article, they order 100, so as to be sure, each, of getting the 90 in the pro rata share delivered" (Mitchell, 1924, p. 645). When faced with limited capacity, suppliers typically allocate available supply among retailers. In turn, a retailer receiving only a fraction of previous orders, amplifies future ones in an attempt to secure more units (Lee et al., 1997a, 1997b). This phenomenon can propagate through the supply chain causing orders (and subsequently inventories) to chronically overshoot and undershoot desired levels. These fluctuations can lead retailers and suppliers alike to overreact, leading to problems such as excessive

supplier capital investment, inventory gluts, low capacity utilization, and poor service (Armony & Plambeck, 2005; E & Fine, 1999; Gonçalves, 2003; Lee et al., 1997a; Sterman, 2000).

Academic interest in the subject has its roots on real and frequent problems faced by businesses in diverse industries. For example, in the 1980's, the computer industry faced shortages of DRAM chips in several occasions: orders surged because of retailers anticipation (Li, 1992). Similarly, excessive reseller orders for Hewlett-Packard LaserJet printers led to excess inventory and unnecessary capacity (Lee et al., 1997a). In 2000, shortages of key components at Cisco caused customer orders amplification, leading to overestimated sales forecasts and a strong production capacity expansion through long-term contracts with OEMs. Once production capacity became available and delivery delays went back to normal, customers canceled duplicated orders, leaving Cisco with significant excess capacity, rigid long-term contracts and high amount of inventory (Byrme & Elgin, 2002).

Informed by these industry experiences, our research fits in the growing field of behavioral operations management, which analyzes the relationships between operations management and human behavior (Croson et al., 2014; Katok, 2011). Previous research in this stream of the literature using the Beer Game estimated individual decision rules for subjects' ordering decisions under complex system structure, resulting in costly oscillations and system instability (Sterman, 1989a; Van Ackere et al., 1993). These results are consistent with those of Croson et al. (2014) also using the Beer Game. Croson et al. (2014) also find oscillations and amplification in orders even when demand uncertainty is eliminated and subjects have access to a perfect demand forecast. In our study, we conducted some experiments to understand the impact that different delays may have on subjects' ordering patterns. Our results also lead to order amplification and oscillation, even though subjects had complete information on the structure of the system and final customer demand.

Our approach for modeling the dynamics of this single-supplier single-retailer supply chain is based on a system dynamics model adapted from Gonçalves (2003). Despite other models described in the literature (e.g. Sterman (1989b), Barlas and Özevin (2004), Armony and Plambeck, (2005)) could also explain the main dynamics of our system, Gonçalves (2003) offers a parsimonious model that could be used to represent supplier's capacity investment and performance. Although Gonçalves' model focuses mainly on the supplier perspective, his model is able to represent the main dynamics and the

positive feedbacks presented in industries affected by long delays and the effect of these delays on retailers' orders. Our experimental setting focuses on the dynamics around retailer's orders. Specifically, we analyze retailers' order amplification, when a retailer faces scarce supply and long delays, and we use a decision rule to explain retailers' decisions. The simple decision rule used in this chapter is based on the same anchor and adjustment concept (Tversky & Kahneman, 1974) previously adopted by Sterman (1989a, 1989b). Sterman's rule for ordering uses the demand forecast as the anchor and adjustments are made in response to the adequacy of the desired inventory and supply line levels. In our decision rule, however, the anchor term captures retailer's intention to place sufficient orders to meet their customers' orders and the adjustment term closes the gap between retailer's desired and actual backlog of orders. In addition, we assume that the supplier behavior (investment in capacity) follows a behavioral heuristic as the one identified by Gonçalves and Arango (2010).

Our research explores the impact that delays may have on subjects' ordering decisions. We hypothesize subjects' performance deteriorates with longer retailer ordering delays and supplier capacity acquisition delays. Both conditions are consistent with studies by Sterman (1989a, 1989b), Gonçalves (2003) or Barlas and Özevin (2004). Our results show that subjects' orders systematically deviate from an optimal order trajectory, experience longer capacity acquisition and ordering delays complicate the system, and when subjects' experience them together it leads to higher costs and lower performance. While subjects' ordering behavior is not optimal, it can be explained econometrically by a simple anchoring and adjustment decision rule. These results are also consistent with Yasarcan (2005; 2011), where he explains the consequences of ignoring delays as a way of ignoring the supply line and shows that the anchor and adjustment decision heuristic, which represents subjects' behavior, is not optimal and that significant delays undermine subjects' performance.

This chapter proceeds as follows. The next section describes and analyzes the proposed mathematical model. Then we detail a decision-making laboratory experiment based on the proposed mathematical model. The following section discusses our results and the impact of ordering and capacity acquisition delays on subjects' performance. Afterwards, we derive an econometric model to analyze subjects' decision rules. Finally, we discuss our main findings.

2.2. Model Description

We build upon a model proposed by Gonçalves (2003) capturing a supply chain with a single supplier offering a unique, non-substitutable product to retailers. The emphasis of our analysis is on the ordering behavior of a single retailer trying to match products received from its supplier with final customer demand. **Error! Reference source not found.** displays the structure of our supply chain structure.



Figure 2.1. Supply Chain Structure

Similarly, Figure 2.2 provides an overview of the supplier-retailer model that serves as the basis for a laboratory experiment. The hexagon in the middle box represents our variable of interest, retailer's orders, where subjects implement their ordering decisions during the experiments. To model the supplier system, we first define the supplier's backlog of orders (B) as a function of retailer's orders (R_D) and supplier shipments (S).

$$\dot{B} = R_D - S \quad (2.1)$$

Shipments (S) are typically given by the minimum between the desired shipments and the available capacity. However, since we are interested in situations characterized by supply shortages, we model shipments as always constrained by available capacity (K).

$$S = K \quad (2.2)$$

The supplier can change capacity (K) over time to adjust to retailer's demand. The change in supplier's capacity (\dot{K}) is given by a first order exponential smooth between desired shipments (S^*) and capacity (K), with an adjustment time given by the time to build capacity (τ_K). This formulation captures a naïve capacity adjustment process, where the supplier tries to maintain sufficient capacity to satisfy retailer demand with a target delivery delay. Finally, desired shipments (S^*), given by the ratio of Backlog (B) and the Target Delivery Delay (τ_D), capture the shipment rate required to maintain delivery delays at the target level for the existing level of backlog.

$$TC = \sum_{t=1}^T (C_{gap} + C_o) \quad (2.6)$$

Where,

$$C_{gap} = \theta \cdot (D_r - E_s)^2 \quad (2.7)$$

$$C_o = \gamma \cdot R_D^2 \quad (2.8)$$

In addition, we assume quadratic cost functions because of three reasons. First, they are reasonable approximations to the loss function in many stock management settings (Holt et al., 1960). Furthermore, quadratic cost functions allow us to penalize higher deviations. Finally, motivated by the work of Diehl and Sterman (1995), we calibrate the cost coefficients using pilot experiments and simulations, allowing us to balance the contribution of each of the two terms in the cost function. We find that coefficients for θ and γ , of $\theta=0.001$ and $\gamma=0.002$ (a 1:2 proportion) represent a balanced trade-off between the ordering and supply gap cost. The higher value of the γ coefficient reflects a higher sensitivity of the cost function to ordering costs, requiring that subjects be mindful about their ordering decisions. The proposed values for θ and γ allow participants to work with cost magnitudes that are manageable and understandable. Appendix 2.1 presents the general units of measure used for each variable or parameter of the model.

2.3. The Experiment

We use the model described above as a basis for a “management flight simulator” (Senge & Sterman, 1992; Sterman, 1989b). Subjects play the role of a single retailer, placing orders to a supplier and trying to minimize total costs. As in the Beer Game, the experiment starts in dynamic equilibrium, where the supplier has sufficient production capacity (100 units/week) to meet total retailer’s demand (100 units/week) according to the target delivery delay. After the third period (week), the retailer faces a sudden increase in final customer orders. This step in the final customer demand is also a common approach in the Beer Game. In addition, despite the fact that real world examples do not include complete information sharing, subjects in our experiments were informed that customer demand will increase in 20% and that the supplier faces a delay to build additional capacity. We give subjects

complete demand information because (1) it facilitates the estimation of the heuristic, (2) it avoids subject's need to forecast future demand, and (3) it has been shown (Croson & Donohue, 2003; Croson et al., 2014) not to eliminate subject's underperformance in supply chains. Subjects must decide how many units to order from the supplier each week during 35 simulated weeks. Subjects are asked to minimize the total accumulated cost (TC) throughout the simulated horizon. An experiment horizon of 35 simulated periods was selected to ensure sufficient time for the dynamics to unfold.

2.3.1. Experimental Treatments

Our experiment explores two characteristics previously identified by Gonçalves (2003) and Gonçalves and Arango (2010) affecting the performance of retailer's decisions. The first one is related to the retailer's ability to get their orders in place, either for internal process of the retailer or for possible delays with the supplier to process the orders that they receive: *retailer ordering delays* (Δ_o). The second characteristic is related to the ability of the supplier to adjust to the orders that they are receiving: *supplier capacity acquisition delays* (τ_k). We model the retailer ordering delay (Δ_o) as a pipeline delay and explore the impact of short ($\Delta_o = 2$) and long ($\Delta_o = 3$) delays on retailer ordering behavior. In addition, supplier capacity acquisition delays (τ_k) are captured as the time constant in the exponential smoothing equation. Analogously, we explore the impact of a short ($\tau_k = 1$) and long ($\tau_k = 3$) time to build capacity. We run a full experimental design, with four experimental treatments. The first treatment (T1) presents an agile system. This is the system with less dynamics in our experiments, where we account for the lowest value in our experimental variables. The fourth treatment (T4) is the most dynamically complex system (slow system), where our experimental variables take the highest possible value. Treatment 2 (T2) presents an agile retailer with a slow supplier, where we combine short retailer ordering delay and a long supplier capacity acquisition delay. Finally, Treatment 3 (T3) presents only an agile supplier with a slow retailer, where we combine short supplier capacity acquisition delay with a long retailer ordering delay. Table 2.1 characterizes each treatment conducted and the number of participants (n) in each treatment.

2.3.2. Experimental Protocol

We followed the standard experimental economics protocol (see Friedman & Sunder, 1994, 2004). Subjects were fourth and fifth year Industrial and Management Engineering students at the National University of Colombia, in the autumn of 2010. Subjects did not have previous experience in any related experiment.

Table 2.1. Experimental treatments

		Supplier's Capacity Investment Delay (τ_k)	
		1	3
Retailer's Order Decision Delay (Δ_o)	2	Agile System ($n=20$)	Agile Retailer ($n=20$)
	3	Agile Supplier ($n=20$)	Slow System ($n=20$)

Participants were told they would earn a show-up fee of Col\$10.000 (approximately US\$5) and a variable amount contingent on their performance, between Col\$0 and Col\$30.000 (US\$0 - US\$15) for an overall average payoff of Col\$24.000 (US\$12). The experiment ran for around one hour and students were informed about the duration of the experiment beforehand. The payoff was more than two times larger than the opportunity cost for an undergraduate student in a public university in Colombia. The students were also given a set of instructions describing the production system, the decisions and the goals of the game (shown in detail in Appendix 2.2).

We ran the experiment with 20 subjects per treatment. Upon arrival, subjects were seated behind computers and one of the four treatments was assigned randomly (see Appendix 2.3). Participants were allowed to ask questions and test out the computer interface (see Appendix 2.4). All the experiment parameters were common knowledge to all participants. We ran the experiment using the computer simulation software *Powersim-Constructor-2.51*®. The software ran automatically and kept record of all variables, including subjects' decisions. Subjects wrote their decisions on a sheet of paper, which served as a physical backup of the data.

2.3.3. Optimal Simulated Trajectory

To properly assess subjects' performance, we compare their ordering behavior with the optimal simulated order trajectory in each treatment. These optimal order trajectories are estimated using the Solver in Powersim Studio 8 and minimizing the total cost over all periods. Powersim Studio 8 uses an optimization method called evolutionary search. Inspired by Darwin's evolutionary theory, the method is a goal-seeking process where successive runs take place and where the best inputs from a run are used in the next run to generate new inputs to a simulation and try to find the optimum. Figure 2.4 shows the behavior of these optimal trajectories (thick continuous line) in each treatment. The optimal ordering trajectories are characterized by a large initial order at the moment the demand surges. The magnitude of this optimal initial order increases with the complexity (longer delays) of the system. Then, orders exponentially decrease with a damped oscillation until settling into equilibrium. The magnitude of the damped oscillation increases with system complexity. Finally, optimal orders settle at 120 units per week for the rest of the trajectory.

2.4. Results

In this section, we present the overall results of our experiments. Our experimental results are based on 61 subjects (15 in the agile-system treatment, 17 in agile-retailer treatment, 14 in the agile-supplier treatment and 15 in the slow-system treatment), chosen among all the subjects after excluding outliers. In order to identify outliers in our experiment, we use both qualitative and quantitative analyses. First, we identified the subjects that clearly did not understand the system; and then we conducted four different quantitative tests to remove the remaining extreme cases. In general, in our quantitative identification of outliers, we used different univariate methods as the ones presented by Ben-Gal (2005), Croson et. al. (2014) and Sterman and Dogan (2015).

2.4.1. Subjects' Order Decisions Behavior

Subjects received information on the system structure, delays and costs (see Appendix 2.2 for the description of the instruction) and then were asked to place orders that would minimize total simulated long-run costs. Figure 2.3 shows ordering behavior for four selected subjects (one in each treatment) capturing typical behavior of subjects. The results suggest a common pattern: subjects' orders initially

over-shoot, then under-shoot until settling around equilibrium close to 120 units (the final customer demand).

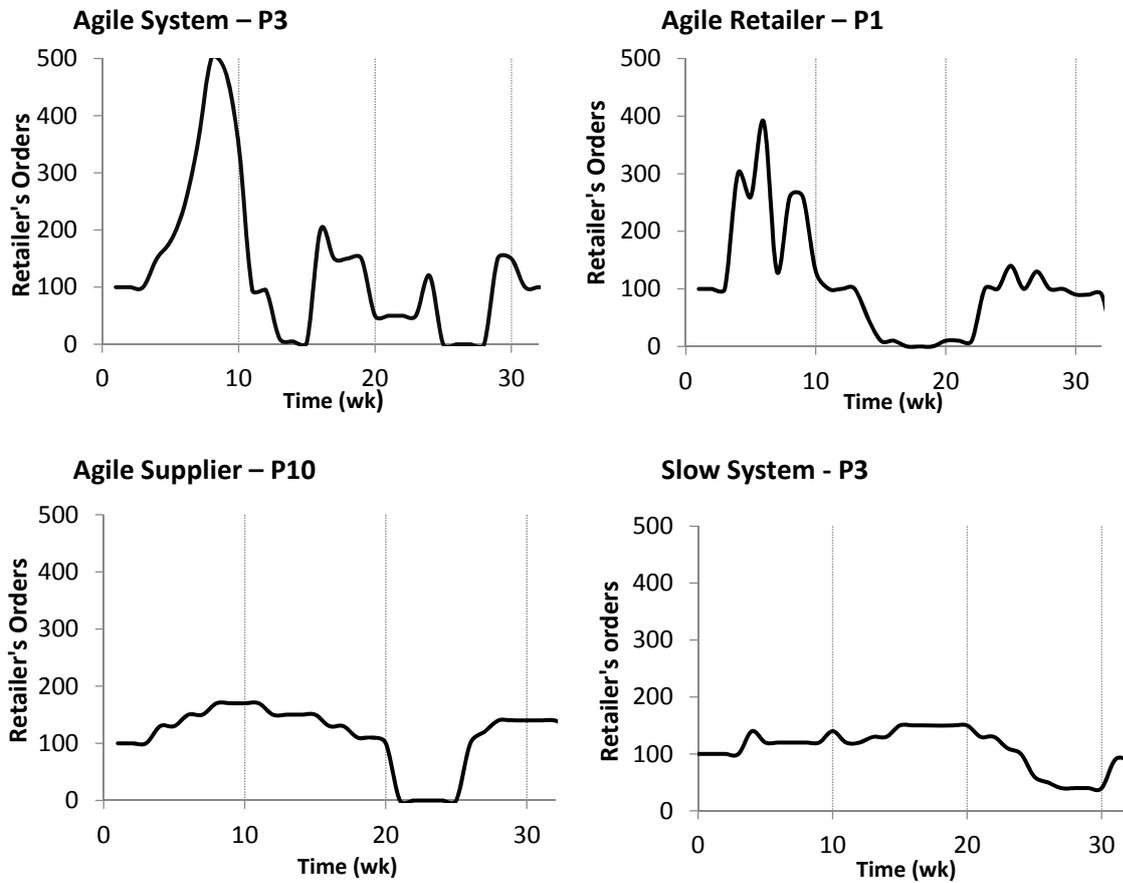


Figure 2.3. Typical experimental results (P_j indicates the subject ID with $j=1, \dots, 15$).

Figure 2.3 also shows that subjects in the agile-system and agile-retailer treatments (with shorter ordering delays) over-order for shorter periods of time (around 10 weeks) compared with subjects in the agile-supplier and slow-system treatments who over-order for longer periods, but displays less variability. In the agile-system and agile-retailer treatments, the shorter ordering delays allowed subjects to more quickly adjust their orders. To compare overall subject behavior in each treatment with the optimal ordering decisions, we compute the average retailer's orders (AO) for players in each treatment. Figure 2.4 suggests that subjects fail to place sufficiently large initial orders, and also fail to reduce them quickly toward the equilibrium value. Instead, subjects place orders with magnitudes averaging half of the desired initial value, but maintain high orders for a longer period than desired. When subjects finally reduce their orders, they do so more than the optimal values. As a result, subjects' orders fluctuate around the optimal trajectory in all treatments. While the pattern presents similarities across treatments,

it is also possible to identify differences. The high initial subjects' orders tend to remain high for a longer period in treatments with longer retailer ordering delays (agile-supplier and slow-system treatments). Subjects' decisions are less stable and take longer to settle in the treatment with higher delays (slow-system treatment).

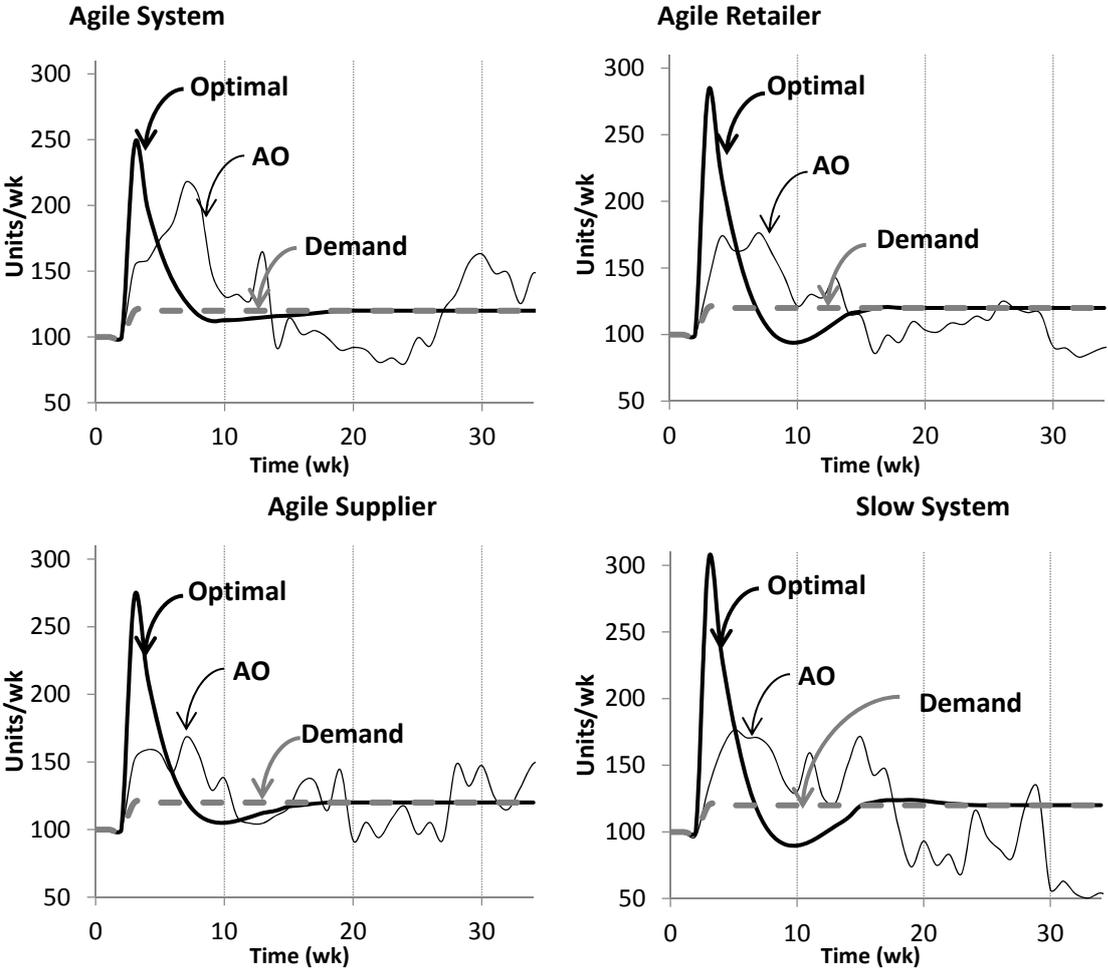


Figure 2.4. Final customer demand, optimal and average subjects' orders (AO) in each treatment.

2.4.2. Subjects' Cost Performance

The subjects' main objective in the experiment was to minimize cumulative costs. Table 2.2 presents total cumulative costs per subject and the average, the median, the minimum and the optimal for each treatment. A general observation is that most of the subjects perform far from optimal for all treatments. The lowest total cost achieved by a subject was 20% higher than the optimal of the treatment, which occurred for subject P12 in the agile-retailer treatment. The best performances observed in the other treatments were also above optimal costs: 32% above optimal in the agile-system treatment, 37% above optimal in the agile-supplier treatment and 95% above optimal in the slow-system treatment.

Table 2.2. Total cumulative, average, and optimal costs across treatment for the experiment

Subject	Agile		Agile	Slow
	System (\$)	Retailer (\$)	Supplier (\$)	System (\$)
P1	2,331.95	3,243.45	1,285.69	1,576.23
P2	16,186.75	10,474.18	6,349.57	1,976.48
P3	17,921.82	1,313.24	1,439.31	19,297.83
P4	3,995.25	6,806.07	2,441.01	12,619.97
P5	845.60	3,017.32	1,407.55	30,220.81
P6	3,834.15	878.90	2,086.54	4,258.17
P7	6,805.24	899.15	3,946.30	2,214.28
P8	25,358.16	14,624.58	2,410.65	1,403.54
P9	4,056.73	2,712.87	2,958.85	2,294.44
P10	1,664.46	854.73	1,202.67	2,649.15
P11	1,511.78	10,944.48	1,106.48	13,445.40
P12	1,193.47	781.92	885.04	35,200.88
P13	4,790.34	2,438.42	961.68	1,388.21
P14	805.86	1,002.29	1,719.34	5,960.91
P15	27,068.96	2,792.65		16,640.05
P16		7,144.60		
P17		5,362.77		
Average	7,891.37	4,428.92	2,157.19	10,076.42
(Standard Error)	(2344,72)	(1018,97)	(396,83)	(2845,09)
Median	3,995.25	2,792.65	1,579.33	4,258.17
Min	805.86	781.92	885.04	1,388.21
Optimal	610.94	654.99	646.86	712.50
Min/Optimal	1.32	1.20	1.37	1.95

Subjects' average performances vary from 333% to 1414% higher than the optimal. These results are conservative since we have excluded subjects with outlying ordering behavior, who present even higher costs (including outliers we get average values of \$23,312.73; \$6,792.49; \$4,724.34 and \$13,551.26 for treatments 1 to 4 respectively). The lowest optimal costs is observed in the agile-system treatment (\$610.94) and highest is in the slow-system treatment (\$712.50), these results highlight the increasing system difficulty when higher delays are introduced producing lower performances. In general, subjects' decisions present higher total cumulative cost in the less agile system (slow-system treatment). In addition, results from the agile-retailer and agile-supplier treatments present lower total cumulative costs than results from the slow-system, as expected. However, results in the agile system do not completely fit the pattern expected if we think shorter time delays will lead to lower total cumulative costs. In this case, both the average and median costs in the agile-system treatment are higher than the average and median cost of the agile-retailer and agile-supplier treatments. This could have a methodological explanation. For example, we could have improved the experimental design, emphasizing higher difference in delays among treatments. Probably, under the lack of large enough sample size, the current setup does not allow us to identify significant differences in costs among treatments and the potential unexpected results could be given just by a normal increase in the orders' variability or a potential sampling selection problem in one of the treatments (in this case the agile-system treatment).

Table 2.3 shows how cost components contribute to optimal and average subjects' total cost in each treatment. These results are robust to changes in cost parameters. The cost breakdown in the optimal trajectory suggests that most of the costs are given by the ordering component. Hence, the choice of parameters θ and γ induce optimal orders that minimize the Supply Gap and its associated cost. In contrast, the cost breakdown for the subjects' decisions shows that subjects have difficulties balancing supply and demand, placing orders that fail to minimize the Supply Gap. Thus, a disproportionally high fraction of the subjects' costs is due to the Supply Gap cost component. As expected, in the most dynamically complex treatment (slow-system treatment), subjects incur the highest proportion of costs due to the Supply Gap. These results are also conservative, because if we include outliers, the cost

percentage given by the Supply Gap will be higher (76.04% for the agile system, 69.38% for the agile retailer, 65.12% for the agile supplier and 82.49% the slow system). This is because the level of underperformance was higher in those subjects excluded from the analysis meaning that, their capacity to minimize the Supply Gap was even lower.

Table 2.3. Costs distribution given by Orders and Supply gap

	% Cost given by Orders	%Costs given by Supply Gap	% Cost given by Orders	%Costs given by Supply Gap
	Agile System		Agile Retailer	
Average	26.9%	73.1%	32.2%	67.7%
	(21.2%)	(21.2%)	(26.4%)	(26.4%)
Optimal	95.7%	4.2%	91.1%	8.9%
	Agile Supplier		Slow System	
Average	41.5%	58.4%	19.6%	80.4%
	(16.3%)	(16.3%)	(13.8%)	(13.8%)
Optimal	92.5%	7.5%	86.6%	13.4%

Standard Deviation in parenthesis

Given the qualitative similarity of the decision patterns and the results shown in this section, one might argue that subjects use a heuristic with common features in order to make their orders (Sterman, 1989a). In the next section, we discuss a specific decision rule and test the accuracy of the rule using econometrical analyses.

2.5. Modeling Decision Rules

For modeling the subjects' decision rules, we test the heuristic (equation (2.9)) proposed by Gonçalves (2003). Gonçalves modeled retailer's orders, R_D , using an *anchor and adjustment heuristic*, where the retailer anchors its orders on a demand forecast, and then adjust it up or down to maintain orders at a desired level. The anchor term captures retailer's intention to place sufficient orders to meet their customers' orders. The adjustment term closes the gap between retailer's desired and actual backlog of orders within a specific adjustment time. Gonçalves (2003) also assumes that each retailer adopts the same heuristic with the model capturing total values for customer demand forecast (d), actual backlog

of orders (B), desired backlog of orders (B^*), and adjustment time (τ_B). Finally, total retailer's orders are non-negative (no cancellations).

$$R_D = \text{Max} \left(0, d + \frac{B^* - B}{\tau_B} \right) \quad (2.9)$$

Where, retailer's desired backlog of orders (B^*) is given by the product of the demand forecast (d) and the expected delivery delay to receive orders from the supplier (ED).

$$B^* = d \cdot ED \quad (2.10)$$

Now, let's assume that the expected delivery delay is given by a linear function of the actual delivery delay (AD) with slope α . This function captures retailer's delivery delay adjustment, that is, when faced with long delivery delays, a retailer sets its expected delivery delay (ED) above the actual delivery delay (AD) quoted by the supplier. Longer expected delivery delays (ED) rather than actual (AD) leads to higher desired backlog of orders (B^*) and higher retailer's orders.

$$ED = \alpha AD, \text{ where } \alpha \geq 1 \quad (2.11)$$

Where, actual delivery delay (AD) is given by the ratio of the order backlog (B) to shipments (S).

Substituting equations (2.10) and (2.11) into (2.9), we obtain equation (2.12), which can be used as a heuristic to test if retailers' orders are well represented by an anchoring and adjustment heuristic.

$$R_D = \text{Max} \left(0, d + \frac{d \cdot \alpha \cdot \frac{B}{K} - B}{\tau_B} \right) \quad (2.12)$$

The system determined by equation (2.12) involves a nonlinearity associated with the ratio of the two states: order backlog (B) and capacity (K). As proposed in Gonçalves and Arango (2010), we linearize the system using a Taylor series approximation of the ratio of the two states (B/K) around the initial backlog (B_0) and capacity (K_0) and neglect higher order terms (details of the linearization process can be found in Appendix 2.5). We get a linear approximation of the *anchor and adjustment heuristic*, which can be tested econometrically:

$$R_D = \text{Max} \left(0, \left(d + \frac{\alpha d \tau_D}{\tau_B} \right) - \left(\frac{\alpha d \tau_D}{\tau_B K_0} \right) K + \left(\frac{\alpha d - K_0}{\tau_B K_0} \right) B \right) \quad (2.13)$$

Below, we analyze this linearized heuristic using two different methods. First, in order to test if this linearized heuristic represents each subject behavior, we estimate the model coefficients for each subject in each treatment using least squares. Second, to test if the linearized heuristic is able to explain the general behavior, we structure the data as a panel estimating a single model for all subjects and treatments. The panel data estimation increases the efficiency of our estimate and its representativeness.

2.5.1. Least Squares Analysis

To econometrically analyze the model stated in equation (2.13), we first omit the maximum operator due to the low incidence of zero decision occurrences (~10%) within each subject's decisions and then we make the following parameter substitutions:

$$\beta_0 = d + \frac{\alpha d \tau_D}{\tau_B}; \beta_1 = -\frac{\alpha d \tau_D}{\tau_B K_0} \text{ and } \beta_2 = \frac{\alpha d - K_0}{\tau_B K_0} \quad (2.14)$$

Then equation (2.13) can be re-written as:

$$R_{Dij} = \beta_{0ij} + \beta_{1ij} K_t + \beta_{2ij} B_t + \varepsilon_{ij} \quad (2.15)$$

Where β_{kij} represents coefficient k , for subject i and treatment j , where $k=0,1,2$, $i=0,1,\dots,15$ and $j=1,\dots,4$, and ε_{ij} is the error term. The parameter values in equation (2.15) suggest that we should expect coefficient β_0 to be positive ($\beta_0 > 0$), β_1 to be negative ($\beta_1 < 0$), and β_2 to be positive ($\beta_2 > 0$). A positive coefficient for β_0 is reasonable since this anchor is the sum of two positive terms. A negative coefficient for β_1 is also intuitive since a higher value of supplier capacity (K) induces lower orders by the retailer. Finally, a positive coefficient for β_2 suggests that faced with a large backlog (B) a retailer will order more in an attempt to receive what she needs. Substituting the parameter values used in the simulation model ($\tau_D=10$, $\tau_B=4$, $K_0=100$, $\alpha=1.1$ and $d=120$), we can obtain estimates for the:

β_0	β_1	β_2
450	-3.3	0.08

We used the R software to obtain the estimates for the model in equation (2.15). Table 2.5 provides the results.

Table 2.5. Coefficient estimates of decision rule for each individual for all treatment

Subject	Agile System				Agile Retailer			
	β_0	β_1	β_2	R^2	β_0	β_1	β_2	R^2
1	120.00†	0.00†	0.00†	0.49	130.07†	-2.35†	0.19†	0.16
2	51.82†	-2.72†	0.31†	0.18	222.44†	-3.32†	0.24†	0.35
3	180.92†	-2.48†	0.19	0.12	262.94†	-2.14†	0.11†	0.52
4	157.95†	-6.96†	0.66†	0.47	159.43	-1.45	0.12	0.04
5	362.94†	2.03†	-0.40†	0.75	411.27†	-2.94†	0.02	0.21
6	213.21†	-1.57	0.06	0.02	300.41†	-1.45†	0.01	0.72
7	53.03†	-7.56†	0.79†	0.50	325.37†	-1.16†	-0.04	0.59
8	200.65†	-1.00	0.04	0.02	148.17†	-1.05	0.08	0.02
9	149.31†	-5.85†	0.55†	0.31	132.54	-0.13	0.01	0.00
10	238.34†	0.11	-0.11	0.05	425.83†	-1.51†	-0.09	0.29
11	145.41†	-0.07	-0.02	0.02	269.08	1.59	0.35†	0.11
12	72.64†	-6.70†	0.71†	0.52	379.68†	-1.99†	-0.01	0.79
13	33.34	1.79	-0.07	0.47	127.95†	-3.45†	0.19†	0.32
14	209.31†	-0.48	-0.02	0.01	199.24†	-0.79†	0.02	0.25
15	349.85†	-2.46	0.13	0.12	158.76†	-2.26†	0.19†	0.13
16					-60.55	2.31	-0.08	0.09
17					175.32†	-1.67	0.09	0.07
Average*	178.96	-3.78	0.37	0.46	251.34	-2.12	0.21	0.30
Using								
Average	182.59†	-5.64†	0.52†	0.42	116.51†	-2.37†	0.05	0.24
Decisions								
Subject	Agile Supplier				Slow System			
	β_0	β_1	β_2	R^2	β_0	β_1	β_2	R^2
1	251.20†	-1.91	0.08	0.05	149.92†	-2.09†	0.19†	0.13
2	314.39†	2.17	-0.32†	0.19	321.35†	-3.77†	0.20†	0.38
3	266.48†	1.24	-0.22	0.06	22.72	-2.02†	0.25†	0.27
4	196.76†	1.66	-0.24	0.01	285.54†	-2.24†	0.14†	0.15
5	493.60†	-3.08†	-0.02	0.72	233.13†	-3.20†	0.23†	0.25
6	212.80†	-1.09	0.01	0.03	172.78†	-1.30	0.09	0.08
7	92.25	-3.69†	0.45†	0.13	266.71†	-0.60	-0.06	0.32
8	213.40†	-3.65†	0.29†	0.13	138.83	1.36	-0.21†	0.13
9	174.03†	-0.18†	-0.05	0.03	78.76	-0.16	0.05	0.04
10	235.47†	-3.69†	0.26	0.25	195.82†	-4.30†	0.38†	0.57
11	173.55†	1.27†	-0.17†	0.40	140.76†	-1.76†	0.14†	0.08
12	243.46†	-2.88†	0.19†	0.17	173.35†	-0.88	0.06	0.10
13	314.00†	-1.95	0.04	0.21	9.40	-0.19	0.08	0.01
14	87.71†	-3.60†	0.39†	0.20	244.36†	-2.04†	0.11	0.11
15					417.05†	-0.27	-0.13	0.11
Average*	244.37	-2.44	0.14	0.20	236.43	-2.68	0.17	0.26
Using								
Average	317.51†	-1.88†	0.03	0.38	145.89†	-2.80†	0.25†	0.31
Decisions								

† Significant at 10% level, * Average using significant values, Italics when OLS is used

Initially, we estimated the parameters for each subject using Ordinary Least Squares (OLS) and then, checked if the OLS assumptions were satisfied. In the cases where all the OLS assumptions were satisfied, those parameter estimations were kept. However, in some cases we found significant autocorrelation of the error term, which makes the OLS estimator inefficient. In the cases with significant autocorrelation of errors, the parameters were re-estimated using Generalized Least Squares with an Auto-Regressive model of order 1 (AR(1)) as model for the residuals. Table 2.5 also shows the method used in each regression. In addition, residual analyses do not show heteroskedasticity, so it is unlikely to bias estimation significantly.

In general, the estimations satisfied the OLS assumptions in 36% of the cases; in the remaining 64% of the cases, we had to use GLS. Results show that a high fraction of the estimated models is significant. For instance, we found significant values for all three parameters in 38% of all subjects and that more than 30% of R^2 values are larger than 0.30. Table 2.5 also computes the R^2 of the “average decision rule” obtained running the model using the average decisions for each treatment (ranging between 0.24 and 0.47), which in social sciences are considered as a moderate explanation. However, these results are not completely conclusive. While the proposed decision rule is consistent for some of our subjects, it is not able to explain subjects’ behavior in other cases. For instance, we have a R^2 lower than 0.2 in 60% of the cases. This suggests that individuals could be using different strategies to make their ordering decisions. Some subjects could be using a rule that combines forecasting and feedback structures as proposed by Paich and Sterman (1993) or they could also be following a non-linear expectation rule or any classical discrete inventory control rule as presented by Barlas and Özevin (2004). However, analyzing the accuracy of these alternative heuristics goes beyond the scope of this study.

In addition, considering the specific results for each coefficient, we observe that the constant β_0 is positive and significant for 83% of all subjects. Coefficient β_1 is also consistent with our expectations, with negative and significant values for most subjects. More specifically, we find significant values for β_1 in 53%, 65%, 57%, and 53% of subjects in the agile-system, agile-retailer, agile-supplier and slow-system treatments, respectively. In addition, most of the signs (83%) of β_1 are negative and 62% of them are significant. The estimates obtained for coefficient β_2 are also as expected with positive values for

most subjects (70%), and β_2 has significant values for 47%, 35%, 42%, and 53% of subjects in the agile-system, agile-retailer, agile-supplier and slow-system treatments, respectively. Most of the signs (70%) of β_2 are positive and 53% of them are significant. Table 2.5 also shows that, after using the least squares estimation with the average decisions, most of the parameters are significant (83%) and all of them have the right signs. (Our qualitative and quantitative analyses do not change significantly if we include outliers.) In addition, the average parameter estimation for each treatment also has the correct signs. Finally, running the model used in the experiment with the estimators obtained with each subject, does not return a single order for zero units. This result suggests that there is no violation of the non-negativity constraint (cancellations were not allowed during the experiment), which supports our estimation omitting the maximum operator in equation (2.13).

Comparing the parameter values for β_1 and β_2 (Table 2.5) with the expected values obtained using the linearized heuristic (Table 2.4), we see that the econometric estimations for β_1 and β_2 have the proper signs. In addition, the estimated value for β_1 is fairly close to the value derived by the linearized heuristic (-3.3); however, the estimated value found for β_2 is farther from the expected one. This result could be given by non-significant estimations, due to an expected β_2 value that is close to zero (0.08) and a limited time series (35 periods). However, if we consider the average of all β_2 coefficients (both significant and insignificant), we obtain a value of 0.11, which is close to the expected value of β_2 (0.08).

Finally, for a few subjects we find statistically significant estimates for parameters β_1 and β_2 with unexpected signs. This is the case for four players: P5 in the agile-system treatment, P2 and P11 in agile-supplier treatment, and P8 in the slow-system treatment. These switches in parameter signs could have occurred because those subjects were able to place high orders at the very beginning of the experiment, when the backlog was relatively low (changing the sign of β_2). Then, β_1 had to control (at least weakly) for the subjects' order increments presented in the remainder of the experiment. These results could also mean that subjects may be using a different decision rule or that they change their decision rule over time. Alternatively, subjects could be using a dynamic decision rule as the one presented by Sterman and Dogan (2015).

To more deeply investigate our understanding of human behavior and the effectiveness of our heuristic, we analyze the data using a panel data analysis. This allows us to get information about the collective subject behavior, controlling for subjects' individual effects.

2.5.2. Panel Data Analysis

We structure the data from the experiments as a panel to control for omitted variables that vary over time but are constant between subjects in each treatment (e.g., temperature, time of the day, day of the week, etc.) and to account for individual heterogeneity, controlling for variables that cannot be observed or measured (e.g., cultural factors). The panel increases the efficiency of the estimations of the linearized model and improves the potential representativeness of the decision rule.

Before making the panel data analysis, we had to decide whether to use random effects, fixed effects, or simple (pooling) least squares. The Breusch Pagan Lagrange multiplier (BP-LM) test helps us decide between random effects and a simple regression. After running this test for each treatment, we found significant difference across subjects (i.e. panel effect): Prob > Chi² is 0.00 all for treatments, which allows us to conclude that random effects are more appropriate.

Next, we run a Hausman test to decide between random or fixed effects. This test checks whether the unique errors are correlated with the regressors. If the effects are exogenous, random effect is efficient, and the fixed effect is just consistent; therefore, we should use random effects. However, if the effects are not exogenous, the fixed effect is efficient, and the random effect is biased, we should use fixed effects. After running the Hausman test for each treatment as a null hypothesis with the preferred model as random effects (and as an alternative to the fixed effects), we found significant differences in all treatments: Prob > Chi² is less than 0.1 (0.00 for the agile-system, agile-supplier and slow-system treatments, and 0.095 for the agile-retailer treatment). Hence, we can reject the null hypothesis and adopt fixed effects as the appropriate approach. To explain overall subjects' behavior, we also control for time-fixed effects. Table 2.6 provides the results of the panel data analysis using Stata.

The significance test of the model using the statistic F shows that all the p -values are small (p -values ~ 0.00) suggesting that the proposed model (of the linearized heuristic) is an acceptable way for explaining subjects' ordering decisions. Table 2.6 also shows the R^2 for each treatment. The proposed heuristic estimated using the panel data approximation, controlling for the individual effects, explains

on average the 25% of the variability in subjects' behavior. In addition, the heuristic performs better explaining subjects' decision when the retailer's order decision delay is low (agile-system and agile-retailer treatments).

Table 2.6. Coefficient estimates of decision rule for treatment as panel data

Regressors	Agile System	Agile Retailer	Agile Supplier	Slow System
β_0 (Intercept)	141.97†	186.30†	187.11†	183.13†
β_1 (Capacity)	-3.40†	-2.59†	-3.28†	-1.94†
β_2 (Backlog)	0.35†	0.21†	0.29†	0.15†
p-value	.00	.00	.00	.00
Corr (ν, B, K)	0.09	-0.06	-0.01	-0.09
R^2 (within)	0.25	0.32	0.15	0.26
N° Observations	495	561	462	495

† Significant at 1%

Furthermore, the three coefficients in all treatments are all highly significant and have the expected signs. The β_1 coefficients are negative and with the expected value (Table 2.4 & Table 2.6). The β_2 coefficients are positive for all treatments as expected, however, the estimated values overestimate the expected magnitude 2 to 4 times (Table 2.4 & Table 2.6). A possible explanation for the overestimation of β_2 may be due to the complexity of the task. In particular, subjects overestimate β_2 to compensate for the underestimation they make for β_0 (the anchor in the linearized heuristic).

Additional insight (that we could not do with the cost analyses - see Table 2.2) could be obtained now from the parameter estimates. Initially, the β_1 estimation in the agile-system treatment (-3.40) and the agile-retailer treatment (-2.59) are lower than the estimations in the agile-supplier treatment (-3.28) and the slow-system treatment (-1.94), respectively. This could mean that subjects take into account the supplier's capacity investment delay and do not need to inflate their orders when the supplier is able to quickly satisfy their orders ($\tau_K = 1$). Similarly, the β_1 estimation in the agile-system treatment (-3.40) and the agile-supplier treatment (-3.28) are lower than the estimations in the agile-retailer treatment (-2.59) and the slow-system treatment (-1.94), respectively. This means that subjects are accounting for

the effect of their ordering decision delay and they increase their orders when their orders take longer to be perceived by the supplier ($\Delta_0 = 3$). A similar analysis can be done for the effect of β_2 and the experimental variables in subjects' ordering decisions. Hence, despite the cost analyses do not show a clear consistency on some results (especially for the agile system), the panel estimation allows us to see some consistency about how accurately subjects could be making their decisions, taking into account the effect of our experimental variables.

Figure 2.5 summarizes the results obtained in the previous sections. First, we build box-plots using the estimated parameters, obtained by least squares, for all subjects in each treatment (Table 2.5). Second, we let triangles represent the expected values of the linearized heuristic (Table 2.4). Finally, we capture the panel data estimations with circles (Table 2.6). The box-plots show the general ranges and distribution for each parameter, indicating whether (or not) they are skewed. In comparing the results, it is important to note the different scales for parameters β_0 , β_1 and β_2 . Overall we note that estimates obtained by OLS or GLS, and those obtained by the panel data are similar for all coefficients. However, the panel data estimates for β_1 are closer to the expected values, and the OLS & GLS estimates for β_2 are closer to the expected values. We also note that subjects underestimate β_0 in all treatments, that is, econometrically estimated β_0 coefficients are lower than the expected values of the linearized heuristic. Subjects tend to underestimate β_1 (values closer to 0) in all treatments.

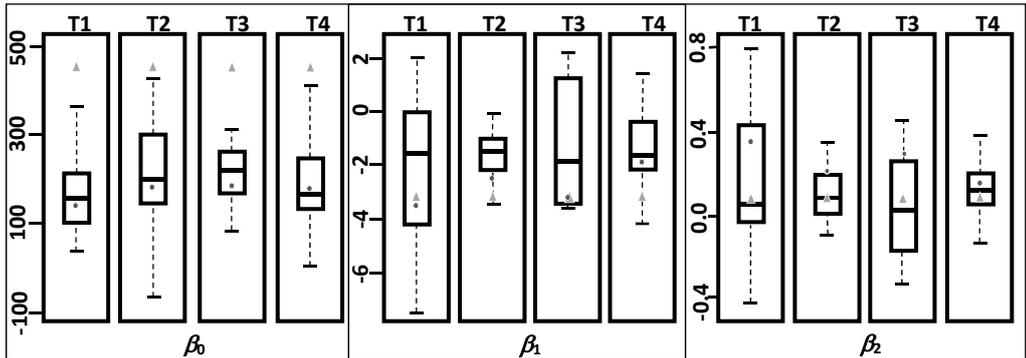


Figure 2.5. Box plots with the coefficient estimations using least squares, expected values (triangle), and panel data estimations (circle).

Figure 2.5 also shows that 63% of the expected values and panel data estimations are between the first and third quartile of the individual least square estimations (represented with the boxplots) in all treatments (T1: agile system, T2: agile retailer, T3: agile supplier and T4: slow system). This percentage

is not higher because the expected values of β_0 are higher than the experimental results, meaning that the heuristic is creating an overestimation of the independent parameter. However, the results for β_1 (Capacity coefficient) and β_2 (Backlog coefficient) show that the estimated values using the heuristic and the panel data analysis are generally with the right sign and also within the expected range.

Figure 2.6 shows that the AO estimations in the agile-system and agile-supplier treatments (those with lower supplier capacity investment delay) present higher variability, increasing the uncertainty in parameter estimation. This increase in variability means that when subjects (as retailers) are able to get a faster response from their supplier, they present more unpredictable behavior, which affects the supplier’s planning. Estimations in the agile-retailer and the slow-system treatments show that subjects are more consistent in their decisions, increasing predictability in their behavior.

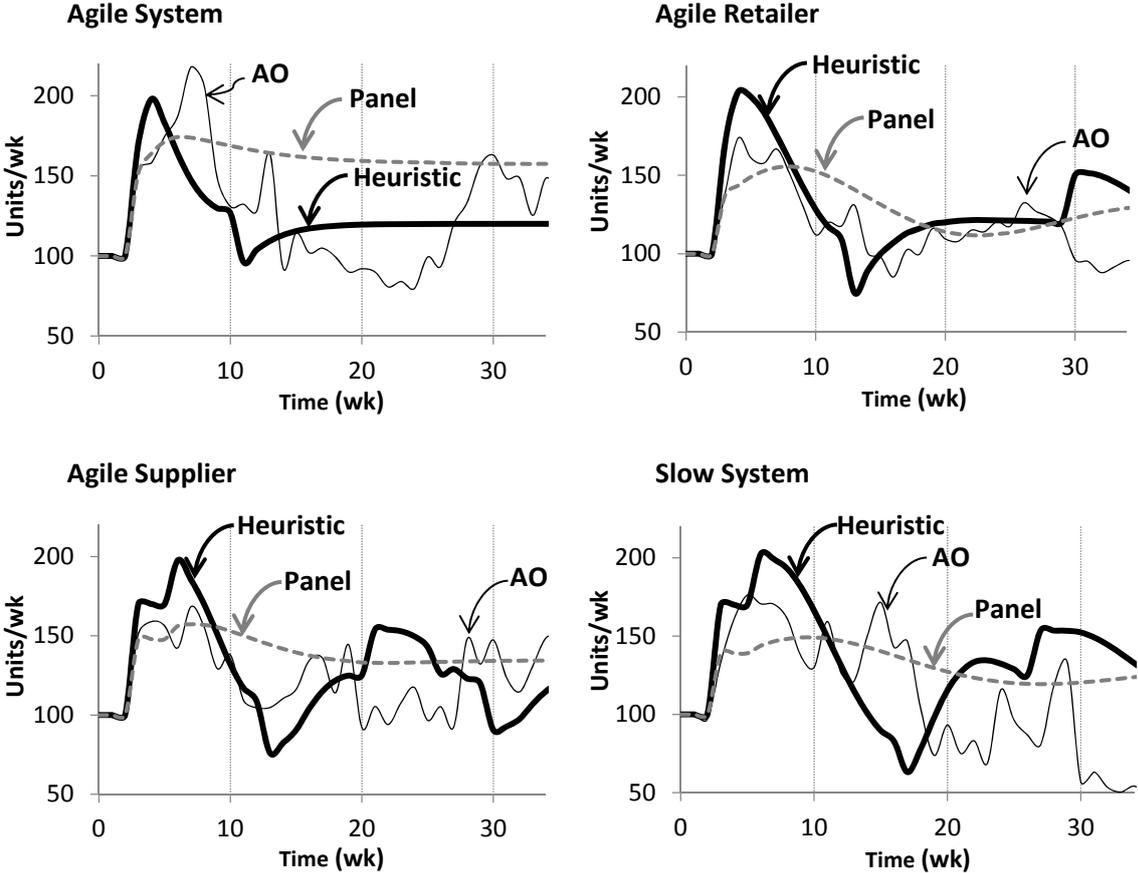


Figure 2.6. Average subjects’ orders (AO), simulation of the proposed heuristic (Heuristic), adjusted model with panel parameters (Panel).

Using the estimators obtained with the panel data analysis (Panel), we inserted and ran them into the same model of the experiment. Figure 2.6 shows the behavior of these runs over time. As it was

predicted by the goodness of fit estimators, these simulations reflect a moderate explanation of the data, seeming to smooth the pattern given by the actual subjects' behavior (AO). Finally, the heuristic proposed in equation (2.12) (Heuristic), with parameter values used in the simulation model ($\tau_D=10$, $\tau_B=4$, $K_{\theta}=100$, $\alpha=1.1$ and $d=120$), presents a better approximation to the actual subjects' decisions in all treatments. Hence, based on our analysis, the heuristic moderately characterizes the decision making rule of some subjects assuming the role of a retailer in a single-supplier single-retailer supply chain.

2.6. Conclusion

In this chapter, we developed a laboratory experiment to explore how subjects playing the role of a retailer place orders in response to a surge in final customer demand. Subjects must minimize total cumulative costs, given by the sum of two cost components: a Supply Gap Cost and an Ordering Cost. In the experiment, subjects face two types of delays: retailer ordering delays and supplier's capacity acquisition delays.

Theoretical Implications

To establish a normative performance benchmark, we estimated the optimal ordering trajectory for each experimental treatment. The optimal trajectory is characterized by a large initial order followed by an exponential decrease that undershoots below initial orders and a dampened oscillation into the final equilibrium of 120 units per week. The magnitude of the peak in the optimal ordering trajectory varies across treatments, increasing with longer system delays. Our selection of cost parameters results in optimal trajectories with total costs driven by the ordering component (e.g., 95.7% of the total costs in the agile-system treatment) and the retailer's ability to close any supply-demand gap.

Our experimental results show that subjects underperform when compared to the optimum, even when demand is known and constant and the system begins in equilibrium. This result is consistent with previous research (e.g., Croson et al., 2014). In addition, even with full access to system information, subjects have limited ability to process and interpret the impact of delays and feedback on the overall system behavior, commonly known as misperception of feedback (Sterman, 1989a). Compared to the optimal ordering trajectory, subjects fail to place sufficiently large initial orders and fail to reduce them quickly toward the equilibrium. Instead, subjects' orders are lower than the amounts initially required

but are kept high for longer than optimal. When subjects' orders finally are reduced, subjects do so excessively, under-ordering below the optimal level. As expected, subject performance differs in each treatment, and in particular, it decreases with higher dynamic complexity, consistently with previous findings by Diehl and Serman (1995). The experimental costs provide clues about the sources of subjects' underperformance. In general, cumulative costs are closer to optimal in treatments with shorter delays (agile-system, agile-retailer and agile-supplier treatments) and further from the optimal in treatments with longer delays (slow-system treatment). However, the cost analysis in the agile-system treatment does not completely fit the expected pattern. That is, shorter time delays do not necessarily lead to lower total cumulative costs. Comparing the costs associated with subjects' orders to the optimal ones, we observe higher costs, ranging from 333% to 1414%. The lowest subject cost still is 33% higher than the optimal. In addition, subjects fail to minimize the supply gap during the experiment, incurring high long-term costs. For instance, in the slow-system treatment, the Supply Gap cost accounts for 80% of the subjects' total costs.

Given their limited processing and cognitive capability, people make decisions translating complex information into simple models, either by capturing essential features from problems and not considering all the features, or by developing habits and routines (Lazaric, 2000; Simon, 1982). Our analysis suggests that the anchoring and adjustment heuristic (Tversky & Kahneman, 1974), which Serman (1989a) identified as representative of subjects' behavior in the Beer Game, is a simple rule that represents the ordering decision of some of our subjects.

We tested the anchoring and adjustment heuristic using a linearized econometric model as a function of supplier capacity and backlog. Results from least squares and panel data suggest that the model characterizes 25% of the subjects' ordering decisions. This also suggests that the proposed heuristic is only one possible heuristic explaining subjects' behavior. Other heuristics, such as those found in Paich and Serman (1993) and in Barlas and Özevin (2004) could potentially explain the decision making process followed by other subjects.

The resulting econometric models are significant. The coefficients for capacity and backlog in the individual regressions are also significant and have the expected signs (negative and positive, respectively). In addition, the panel estimation allows us to see some consistency on how subjects place

orders. The results suggest that they take into account the effect of supplier's capacity investment delay and ordering decision delay. For instance, subjects do not inflate their orders as much when the supplier is able to quickly satisfy their orders ($\tau_K = 1$) and subjects increase their orders when they take longer to be perceived by the supplier ($\Delta_0 = 3$).

Practical implications

Our research also provides insights relevant to decision makers interested in the importance of improving decision-making and implementing Business Process Redesign (Van Ackere et al., 1993). Design can play an important role in production-distribution systems by reducing total costs and improving system stability. However, such redesign requires breaking existing habits, understanding and carefully evaluating present processes (Van Ackere et al., 1993). Our results suggest that, when possible, retailers should try to decrease the delays inherent in their ordering processes. In doing so, retailers would reduce the complexity of the system, improve their ordering decisions as well as improve their ability to manage mismatches between supply and demand. Hence, shorter ordering delays lead to simple systems, which generally yield lower costs. These results are consistent with those of Sterman (1989b), Kaminsky and Simchi-Levi (1998), and Gupta et al. (2002).

In practice, managers should be careful when relying on rules-of-thumb. In our simulated experiment, the adopted heuristic performs substantially worse than the optimal, which suggests significant opportunity for improvement. Naturally, while heuristics are simple and useful, if they are not good enough, they could lead to consistent biases, limited search, and resistance to change (Lazaric, 2000; Leonard-Barton, 1992).

Limitations and future research opportunities

Despite the meaningful discussions presented here, this research has limitations that could be addressed in future research. First, the analysis focuses on a simple supply chain where a single subject is making ordering decisions and interacting directly with the computer simulation. However, in practice, there is not just one retailer; there is competition amongst multiple retailers. Therefore, it would be interesting to study multiple subjects (retailers) interacting amongst themselves and placing orders to the same supplier who will allocate the available supply in proportion to orders placed. Different allocation mechanisms and cost functions could be explored. It would be interesting to see if the performance

during the game using multiple subjects is improved by suggesting a coordination stock as a buffer against strategic uncertainty (Croson et al., 2014).

Future research could also explore possible ways to improve retailers' decisions either by prominently displaying important information, or by providing guidance regarding some heuristics that could be followed. For example, subjects could be informed about the improved decision-making heuristic introduced by Yasarcan (2011), which creates a stable and fast response in a dynamic behavior of the stock, accounting for delayed information. Building on this, such efforts could provide clues to the required training managers would need in order to improve their performance in supply chain systems. For instance, we could add information sharing as a control variable; hence, we would be able to conclude what would be its effect in the supply chain performance (Van Ackere et al., 1993).

In addition, other (dynamic) heuristics as the ones proposed by Paich and Sterman (1993), Barlas and Özevin (2004) or Sterman and Dogan (2015) could be used to test subjects' behavior under different supply chain setting such as retailer competition and to improve the understanding of the decision process followed by the outliers. Finally, this research does not directly analyze other relevant dynamics like learning effects, where other meaningful implications could be found.

Appendix 2.1. Complete Model specifications

Variable, Stock or Parameter	Symbol	Initial Value	Equation	Units
Backlog of orders	B	1000	$\sum_t (R_{Dt} - S)$	Units
Retailer's desired backlog of orders	B^*	1000	$d \cdot ED$	Units
Expected delivery delay	ED	10	αAD	wk
Cumulative Customer Orders	D_R	100	$\sum_t d$	Units
Cumulative Supplier Shipments	E_S	100	$\sum_t S$	Units
Retailer's orders	R_D	100	Decision Variable*	Units/wk
Delayed Retailer's orders	R_{Dt}	100	$R_D (t - \Delta_0)$	Units/wk
Supplier Shipments	S	100	K	Units/wk
Supplier Capacity	K	100	$\sum_t K$	Units/wk
Change in Capacity	\dot{K}	0	$\frac{B / \tau_D - K}{\tau_K}$	Units/wk/wk
Final Customer Orders	d	100	$100 + \text{step}(3, 20)$	Units/wk
Retailer's total costs	TC	10	$\sum_t (C_{\text{gap}} + C_o)$	\$
Supply Gap Costs	C_{gap}	0	$2 \cdot 10^{-3} \cdot (D_R - E_S)^2$	\$/wk
Order costs	C_o	10	$1 \cdot 10^{-3} \cdot R_D^2$	\$/wk
Supply Gap Cost Coefficient	θ	0.001	0.001	\$/Unit ²
Order Cost Coefficient	γ	0.002	0.002	\$/Unit ²
Target Delivery delay	τ_D	10	10	wk
Time to build Capacity	τ_K	1 or 3	1 or 3	wk
Time Adjust Backlog	τ_B	4	4	wk
Retailer's Order Decision Delay	Δ_0	2 or 3	2 or 3	wk
Linear Coefficient	α	1.1	1.1	Dimensionless

*The heuristic proposed for the Retailer's order is shown in equation (2.12).

Appendix 2.2. Instructions for T3 – Agile-Supplier Treatment (Translated into English)

INSTRUCTIONS PLEASE DO NOT TOUCH THE COMPUTER UNTIL YOU ARE TOLD

Welcome, from this moment you are part of an experiment about decision-making, in which you will assume the role of a wholesale manager. Your goal in the game is to **minimize cumulative costs** at the end of the game simulation (50 weeks). According to your performance, you will get cash payment as a reward. This money comes from a research project sponsored by the Universidad Nacional de Colombia, Sede Medellín.

Your weekly decision is to define how many units to **order** to your supplier, with the objective to cover all your customer demand (in the experiment, this decision is taken in the cell placed in front of “*Order decision*”). The decision you take will be received by your supplier three weeks after the order is made and it will be accumulated in the supplier backlog. The initial production capacity of your supplier is 1000 units by week. However, the supplier has the ability of changing their capacity according to the orders you make: more orders, more investment in capacity. The supplier’s capacity building time is one week. If your supplier does not have enough capacity to satisfy your orders, he is going to have delivery delays (greater than 10 weeks) and therefore, you will also be delayed with your customers.

You incur costs every week from two components:

1. Ordering Cost (C_o):

$$C_o = \frac{1 * (\text{Order Decision})^2}{1000}$$

With an initial capacity of 100 units, this cost in the first week is \$10.000.

2. Supply gap cost (C_{gap}):

$$C_{gap} = \frac{2 * (\text{Supply Gap})^2}{1000}$$

With a deficit of 0 units, the cost in the first week is \$0.

In this way, the total cumulative cost TC is the sum of these costs during the whole simulation:

$$TC = \sum_{t=1}^T (C_{gap,t} + C_{o,t})$$

Initially, you order 100 units per week, which allows your supplier to maintain a target delivery delay of 10 weeks as an initial condition. Recently, novel applications of your product created a surge in demand. You estimate the increase in demand to be permanent and in the order of 20 units per week. Because you were not attentive to these novel applications, the surge in demand caught you by surprise. You notice your deficit is increasing and therefore, you are losing customers and prestige.

You will begin during 3 weeks deciding 100 units as a learning period. Afterwards, your task is to manage the company during the simulation, deciding how much to order to your supplier while minimizing the total cumulative cost, TC.

PAYMENT: The payment will be in cash at the end of the experiment. It corresponds to a fixed amount for participation of COP\$10000 plus a variable amount between COP\$0 and COP\$30000, depending on the TC result. The lower the total cost TC, the greater the payment.

NOTE: Please do not share information about the experiment with your peers to avoid losing scientific validity of the experiment.

GLOSSARY

(ABOUT THE RESULTS THAT ARE OBSERVED IN “REPORTS”)

Operations Section: It gives information about the wholesaler’s system (you).

OPERATIONS

DEMAND [units/wk]	100	➔ 1
SHORTAGE [units]	0	➔ 2
UNITS RECEIVED [units]	100	➔ 3
DELIVERY DELAY [wk]	10.00	➔ 4
SUPPLIER BACKLOG [units]	1000	➔ 5

1. Orders that you receive from your final customers. This is the demand that you should cover each period.
2. Units that you need to deliver (if it is **negative**, it indicates **inventory**).
3. Units that arrive every period to the wholesaler (you) from the supplier. Those are the available units you have by period to satisfy the demand.
4. Average delivery-delay time of the orders, counted from the moment you make the order until the moment you receive them. The ideal delivery delay is 10 weeks.
5. It accumulates the difference between the orders made and received by the wholesaler (you) over time. Initially you have total accumulated orders of 1000 units, which will be received in batches of 100 units during 10 weeks.

Costs Section: It gives information about each cost component.

ORDERING COSTS [\$ /wk]	10	➔ 6
SUPPLY GAP COSTS [\$ /wk]	0	➔ 7
TOTAL WEEKLY COSTS [\$ /wk]	10	➔ 8
CUMULATIVE COSTS [\$]	0	➔ 9

6. Cost due to the order decision made each period.
7. Cost of having inventory or shortage for the final customer (\$/week).
8. Sum of the two cost components every week.
9. Total accumulated cost incurred during the simulation CT.

Appendix 2.3. Experiment Environment



Appendix 2.4. Interface of the experiment in Powersim (in Translated into English)

GAME CONTROL

ORDER DECISION [units/wk]

100

NEXT WEEK >>

Week: 1

REPORTS

OPERATIONS

DEMAND [units/wk]	100
SHORTAGE [units]	0
UNITS RECEIVED [units]	100
DELIVERY DELAY [wk]	10.00
SUPPLIER BACKLOG [units]	1000

COSTS

ORDERING COSTS [\$ /wk]	10
SUPPLY GAP COSTS [\$ /wk]	0
TOTAL WEEKLY COSTS [\$ /wk]	10
CUMULATIVE COSTS [\$]	0

ORDERS

DEMAND

COSTS

Appendix 2.5. Linearization of Heuristic

Given the system determined by equation (2.12) involves a nonlinearity associated with the ratio of the two states: order backlog (B) and capacity (K), we linearize the system. We use a Taylor series approximation of the ratio of the two states (B/K) around the initial backlog (B_0) and capacity (K_0) and neglect higher order terms.

$$AD = \frac{B}{K} = \frac{B_0}{K_0} + (B - B_0) \frac{1}{K} \Big|_{B_0, K_0} + (K - K_0) \frac{B}{K^2} \Big|_{B_0, K_0} = \frac{B}{K} = \frac{B_0}{K_0} + (B - B_0) \frac{1}{K_0} + (K - K_0) \frac{B_0}{K_0^2} \quad (\text{S4})$$

and since in equilibrium we have that the supplier's initial supplier capacity (K_0) is equal to $K_0 = B_0/\tau_D$, the linearized form for delivery delays is given by:

$$AD = \frac{B}{K} = \tau_D \left(1 + \left(\frac{B/\tau_D - K}{K_0} \right) \right) \quad (\text{S5})$$

Substituting S5 in 12, we get:

$$R_D = \text{Max} \left(0, d + \frac{d\alpha\tau_D \left(1 + \left(\frac{B/\tau_D - K}{K_0} \right) \right) - B}{\tau_B} \right) \quad (\text{S6})$$

Finally, grouping terms and taking the linear part of K and B , we get a linear approximation of the *anchor and adjustment heuristic*, which can be tested econometrically:

$$R_D = \text{Max} \left(0, \left(d + \frac{\alpha d \tau_D}{\tau_B} \right) - \left(\frac{\alpha d \tau_D}{\tau_B K_0} \right) K + \left(\frac{\alpha d - K_0}{\tau_B K_0} \right) B \right) \quad (\text{13})$$

Chapter 3. Behavioral Analysis of the Effect of Duplication Orders in Single-Supplier Multi-Retailer Supply Chains

(with Paulo Gonçalves)

Abstract

The bullwhip effect, the tendency of subjects to inflate their orders to suppliers to satisfy customers' demand, is a frequent and costly source of instabilities in supply chains. One typical operational cause for such problem occurs when there is a horizontal competition among retailers for scarce supply. While this operational cause is intuitive, there is little research quantifying the impact it causes in retailers' behavior. In this chapter, we model a one-supplier two-retailer supply chain and we run three different behavioral studies to determine how subjects react to different supply chain conditions. We evaluate the effect of different duplication probabilities, different supplier's capacity acquisition delay and different steps in final customer demand. Results show that the bullwhip effect persist in a system with two competing retailers even when subjects do not have incentives to inflate their orders. A cost analysis shows that systems with lower dynamic complexity lead to lower average costs. However, econometric results of our decision rule show that when subjects face situations with higher complexity, subject's rational system (System II) is more active, leading to a reduction in the observed biases. Finally, we analyze the effect of two different supplier allocation mechanism (proportional and turn-and-earn) on supply chain performance. Results show that under identical supply chain structure, neither allocation mechanism eliminates order amplifications, but the use of the turn-and-earn allocation mechanism leads to a reduction of the bullwhip effect.

Key words: Bullwhip effect, Allocation mechanism, Duplications, Behavioral analysis

3.1. Introduction

The bullwhip effect is a typical supply chain problem that takes place when orders' variability increases as we move up the supply chain. This order variability leads to inefficiencies and instabilities in the whole supply chain. This increase in variability means that the variability of final customers' orders is lower than the variability of retailers' orders, while the variability of retailers' orders is lower than suppliers' orders, and so on (Croson et al., 2014; Lee et al., 1997a). This pattern of behavior takes place because subjects distort the information they received from their immediate downstream customer to

the upstream supplier. The distortion of information from customers to suppliers leads to high operational inefficiencies such as high inventory levels and unnecessary capital investment. This phenomenon is frequently evident in many industries, even when the final consumer demand does not vary much (Cachon et al., 2007; Lee et al., 1997b; Sterman, 2000).

In an attempt to ameliorate the consequences of the bullwhip effect, managers have been improving their information systems and organizational relationships to identify and reduce of the main causes of the bullwhip effect (Ellram, 2010). In a similar effort to understand this phenomenon, previous studies of the bullwhip effect have shown that both operational and behavioral causes can provide explanations on its origin (Lee et al., 1997a, 1997b; Sterman, 1989a). The operational causes suggest that even if subjects were fully rational, supply chain instability may persist due to the actual structure of the system. Lee et al. (1997b) identified four typical operational causes of the bullwhip effect, which should be controlled and taken in to account for improving process design. The four operational causes are: (i) order batching, (ii) price fluctuations, (iii) rationing and shortage gaming, and (iv) demand forecast updating. However, the behavioral causes predict that, due to subjects' bounded rationality, the bullwhip effect will take place even if we control for the four operational causes. Therefore, behavioral causes emphasize how limitations in subject's rationality lead to behavior that diverges from the theoretical predictions (Diehl & Sterman, 1995; Schweitzer & Cachon, 2000; Sterman & Dogan, 2015).

Previous studies have mainly focused in understanding these behavioral deviations by analyzing subjects' behavior in serial supply chains (Croson et al., 2014; Sterman & Dogan, 2015; Sterman, 1989b) or in a single-actors systems (Villa et al., 2015; Yasarcan & Barlas, 2005). However, there is scant behavioral research on analyzing subjects' behavior in non-serial supply chains, where analytical research in operations management has claimed that the existence of shortages in a system with horizontal competition may lead subjects to over-order (Armony & Plambeck, 2005; Gonçalves, 2003; Lee et al., 1997a; Sterman, 2000). Similarly, empirical studies have claimed that when we account for non-serial and competitive supply chains, in addition to the inflated orders, customers may duplicate their orders with certain probability by placing additional orders to multiple retailers. Hence, duplications take place in a competitive environment where a customer finds that his retailer is out of stock, and therefore, the customer decides to place simultaneous orders to multiple retailers (Armony &

Plambeck, 2005). In addition, in an attempt to guarantee higher service levels, retailers' inventory levels may be worsened (increased) as the probability of duplications and retail competition level increases (Cachon & Olivares, 2009). Summarizing these points, there is an opportunity to make behavioral research considering non-serial supply chains with horizontal competition and duplicated orders from final customers, to understand the consequences in retailers' behavior.

A well-known example of the consequences of duplications in a real supply chain is the one experienced by Cisco Systems in 2001, when the demand for their products increased considerably leading to significant stock-outs. Cisco lost more than US\$ 2.2 billion in inventory write-offs due to the combination of different factors such as stock-outs, long periods to build production capacity and an excessive increase in the number of indirect resellers (Adelman, 2001). These factors led final customers to duplicate their orders by placing orders to multiple resellers. Therefore, due to the horizontal competition among resellers and in an attempt to receive higher allocation of units, resellers inflated their orders to the Cisco. This situation led to a misread of the final customer demand and a catastrophic drop in Cisco's net income, which reached a value of US\$ -2.7 billion by the third quarter of 2001 (Armony & Plambeck, 2005).

As claimed by (Clark & Scarf, 1960; Scarf, 1959), a supplier should control her inventory under demand uncertainty, trying to find out retailer's ordering policy while learning about final customer demand distribution, especially in systems with long delays and high level of retailer aggressiveness for scarce resources (Gonçalves & Arango, 2010). Horizontal competition poses an additional challenge to suppliers, where in case of stock-outs, they will have to define a specific allocation mechanism to distribute the scarce supply among the retailers (Cachon & Lariviere, 1999c). Allocation mechanisms may consider the amount of past orders, or the amount of past sales, or the difference among retailers' orders, among others. However, retailers may order differently depending on the allocation scheme chosen by the supplier. Therefore, some specific allocation mechanisms may lead retailers to make decisions that are closer to the theoretical predictions, while others may induce retailers to highly deviate from the theoretical predictions in an effort to gain a better allocation.

Motivated by the enormous Cisco's inventory write off and the analytical model developed by Armony and Plambeck (2005), we study how customers' duplicated orders and cancellations can lead

retailers to make faulty ordering decisions. We investigate the effect of order duplications in the decision process and performance of retailers in a decentralized supply chain. The system is composed by (i) one supplier, who has a durable production capacity that can be carried over from period to period, and (ii) two competing retailers who face a stable and known final customer demand. Initially, we extend Armony and Plambeck's (2005) work by developing a system dynamics model encompassing endogenous decision policies for both supplier's capacity investment and customers' ordering. This formal mathematical model represents the main structure and dynamics of a single-supplier two-retailer supply chain. Then, we run three different experimental studies aiming to describe the dynamics and biases around ordering decisions. We evaluate the effect of different duplication probabilities, different supplier's capacity acquisition delay, different steps in final customer demand and two different supplier allocation mechanisms on supply chain performance.

Finally, to analyze how subjects' decisions change based on the specific structure of the system, we build on the behavioral models proposed by Croson and Donohue (2005) and Oliva and Gonçalves (2005) to create a parsimonious model that can be estimated for subjects in each experiment. Previous studies have shown how similar rules are able to highlight the biases presented in subjects' decisions (Bolton & Katok, 2008; Bostian et al., 2008; Croson & Ren, 2013; Schweitzer & Cachon, 2000; Sterman, 1989a). However, they do not describe how the main behavioral parameters change by the fact that people are facing different operational challenges in a multi-agent system where there is competition for scarce supply.

In the next section, we present our mathematical model by incorporating one supplier and two retailers that compete for supply allocation. Then, we describe the main considerations in our experimental design. Afterwards, we describe our three experimental studies and discuss the main quantitative results of the behavioral rule. Finally, we conclude and provide some insights and recommendations.

3.2. Formal Model

To create a proper framework that allows us to systematically analyze the effect of retailers' orders in supply chain performance, we build upon the model proposed by Armony and Plambeck (2005) and upon the heuristics used by Sterman (Sterman, 1989a; Van Ackere et al., 1993). Our model considers a

supply chain of a durable good that can be stored over time. One supplier with two competing retailers compose the structure of the supply chain, where retailers face a stationary final customer demand. In addition, the final customers may duplicate or cancel orders according to retailer product availability. Figure 3.1 shows the general structure of the supply chain, where the dashed lines represent the flow of information (actors' orders) and the solid lines represent the flow of units (shipments).

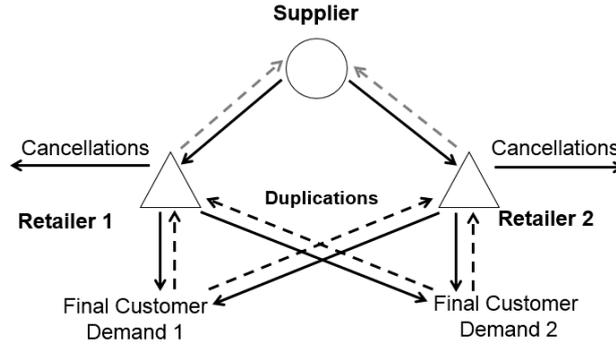


Figure 3.1. Structure for a single-supplier two-retailer supply chain with order duplications.

3.2.1. Retailers' Model

Our model includes two independent retailers that may compete for scarce supply to satisfy the final customer demand. For each retailer i ($i=1,2$), final customer demand ($o_{t,i}$) at any time t is initially exogenous. However, if a final customer finds out that his retailer is out of stock, he will then place a duplicated order with the other retailer ($d_{t,i-}$) with a given probability α . Therefore, retailers may receive orders both from his own customers and from customers of the other retailer. The level of on hand inventory for each specific retailer ($I_{t,i}$) increases with the number of units received from the supplier ($S_{t,i}$) and decreases with the number of units shipped (s_t^i)¹. The number of units shipped are a function of the units sent to his own final customers ($s_{t,i}^i$) and to the potential units shipped to the other retailer's customers ($s_{t,i-}^i$) (Equation (3.1)). The total number of units shipped to the final customers will be given by the minimum between the on hand inventory ($I_{t,i}$) and the unsatisfied demand (backlog) ($b_{t,i}$) (Equation (3.2)).

As soon as one retailer supplies an initially unsatisfied final customer with his original desired amount, the final customer cancels any duplicated order ($c_t^{d,i}$), which may include cancelations with his

¹ Variables in bold refer to summation of individual variables. Example: $\mathbf{s}_t^i = s_{t,i}^i + s_{t,i-}^i$.

own retailer ($c_{t,i}^{d,i}$) or with the other retailer ($c_{t,i-}^{d,i}$). Furthermore, customers may get upset for not being supplied after an average waiting time (τ_w), hence, they will cancel some outstanding orders ($c_t^{w,i}$) (Equation (3.3)). Therefore, every period the backlog increases with the incoming orders ($o_{t,i}$) and duplicated orders ($d_{t,i-}$), and decreases with the shipments (s_t^i) and cancellations (c_t^i) (Equation (3.4)).

$$\dot{l}_{t,i} = S_{t,i} - s_t^i \quad (3.1)$$

$$s_t^i = \min(l_{t,i}, b_{t,i}) \quad (3.2)$$

$$c_t^i = c_t^{w,i} + c_t^{d,i} \quad (3.3)$$

$$\dot{b}_{t,i} = o_{t,i} + d_{t,i-} - s_t^i - c_t^i \quad (3.4)$$

Notice that the backlog represents the number of outstanding units that each retailer should satisfy in the following periods. The orders that could not be supplied accumulate in the respective backlog, i.e. the backlog represents the customers' orders that are waiting to be delivered by the retailer. If a retailer has enough inventory, he will be able to supply the orders to both customers. Otherwise, retailers will ship the available product to customers in proportion to their share of the backlog and orders received from each one of them. Finally, we can define the effective inventory for each retailer ($RI_{t,i}$) as the difference between the on hand inventory and the backlog, so that a positive value for $RI_{t,i}$ represents the existence of on hand inventory; while a negative $RI_{t,i}$ represents the existence of unsatisfied customers. Note that the actual unsatisfied demand at time t is given by the total backlog in that period minus the total number of duplications. Therefore, if $RI_{t,i} > 0$, for any retailer i , then $d_{t,i} = 0$.

3.2.2. Supplier's Model

Supplier's production capacity determines her ability to deliver products to retailers. One of the objective of the supplier is to define how to adjust her production capacity (K_t) to satisfy retailers' orders (O_t). Supplier's shipments to retailers (S_t) over time is a function of the available production capacity and the outstanding orders to the retailers (B_t) (Equation (3.5)). Similarly, the number of outstanding orders increases with the actual number of retailers' orders and decreases with the number of shipments (Equation (3.6)).

$$S_t = \min(B_t, K_t) \quad (3.5)$$

$$\dot{B}_t = O_t - S_t \quad (3.6)$$

To model the way the supplier adjusts her capacity, we follow a simple anchoring and adjustment heuristic (Gonçalves, 2003; Sterman, 1989a, 1989b). Initially, no change in capacity is desired, therefore, the supplier anchors her change in capacity to zero. Then, the adjustment is given by the gap between the actual and desired capacity. The desired capacity is given by an initial forecast about future retailers' orders (\hat{O}_t) and an adjustment of the current outstanding orders ($AdjB_t$). Then, the supplier compares this desired capacity with her actual capacity to determine the capacity gap and the indicated increase/decrease in capacity. Finally, an increase or a decrease in capacity requires a time to build/reduce capacity (τ_K). Hence, the heuristic uses this time to divide the capacity gap and find the change in capacity over time (\dot{K}_t) (Equation (3.7)).

$$\dot{K}_t = \frac{(\hat{O}_t + AdjB_t) - K_t}{\tau_K} \quad (3.7)$$

The supplier's production capacity determines her ability to deliver products to retailers (and, therefore, to customers), and her estimation of retailers' orders serves as a basis for her investment on production capacity. Therefore, supplier's forecast about retailers' orders is updated over time based on the last period forecast and the actual retailers' orders from last period (Equation (3.8)). Finally, the $AdjB_t$ is computed based on the current outstanding orders and on a desired level of outstanding orders (\hat{B}_t) (Equation (3.9)). \hat{B}_t is defined as a function of the expected retailers' orders and on a desired coverage period τ_C ($\hat{B}_t = \tau_C \hat{O}_t$).

$$\dot{\hat{O}}_t = \frac{\hat{O}_{t-1} - O_{t-1}}{\tau_o} \quad (3.8)$$

$$AdjB_t = \frac{B_t - \hat{B}_t}{\tau_B} \quad (3.9)$$

In case of insufficient production capacity, the supplier may distribute her limited production capacity following one of the following allocation mechanisms: *Proportional* (Lee, 1997), or *Turn-and-Earn* (Cachon & Lariviere, 1999b; Lu & Lariviere, 2012). Under a Proportional allocation mechanism, retailers receive a number of units that is proportional to the outstanding orders to be received: $S_{t,i} = S_t * B_{t,i} / B_t$. Under a Turn-and-Earn allocation mechanism, the retailer with higher past sales gets a

more favorable allocation. First, the supplier divides her production capacity in two: reserved and unreserved capacity. Then, the reserved capacity is guaranteed for the last-period sales leader and it is equal to the difference between retailers' sales, and finally the unreserved capacity is allocated evenly among retailers.

3.3. The Experiment

We strive to advance knowledge about amplification of retailer's orders (O_t) in a single-supplier multi-retailer supply chain by running a decision-making laboratory experiment based on the model previously developed. In our experiment, each subject assumes the role of one of the two positions of the retailers, so that there are two retailers making decisions simultaneously in each system. Subjects place orders to an automatized supplier while trying to satisfy the final customer demand. Retailers make decisions for 35 periods. Every period, retailers' orders take place after retailers receive the shipment from the supplier and the final customer order is fulfilled with the available inventory. We ran three different studies with different treatments and we exposed each subject to only one of the treatments. The studies help us understand the relative strength of different factors in generating the inflationary ordering behavior. Our three studies explore different characteristics previously discussed by Armony and Plambeck (2005), Gonçalves (2003), and Oliva and Gonçalves (Oliva & Gonçalves, 2005) affecting supply chain stability: *probability* of customer order duplications, supplier *capacity acquisition delay*, *step* in final customer demand and supplier's *allocation mechanisms*.

3.3.1. Experimental Protocol

We followed the standard experimental economics protocol (Friedman & Sunder, 1994, 2004; Katok, 2011). We ran formal experiments with undergraduate students in management and industrial engineering. We ran each experimental treatment with an average of 28 participants. Upon arrival to the experiment, subjects were seated behind computers and a partner was assigned randomly. Subjects were given a set of instructions describing the production system, the decisions and the goals of the experiment. Participants were allowed to ask questions and test out the computer interface (See Appendix 3.1). Subjects had *full information* about the system structure, delays and main parameters. To guarantee the independence of experimental results from individual interests, experimental subjects

were rewarded according to their performance (Smith's (1976, 1982) Induced Value theory). Subjects knew before the experimentation that they would earn a show-up fee for participating and a variable amount contingent on their performance (this payoff was larger than the subjects' opportunity cost). Our mathematical model was inserted into a computer simulation software Z-tree (Fischbacher, 2007) where the computer automatically ran the decisions of the supplier and the final customers, while subjects make their ordering decisions as retailers. The software kept record of all variables, including subjects' decisions.

Subjects' objective was to minimize total costs (TC) during the 35 simulated periods. As in the Beer Distribution game (Croson & Donohue, 2005; Serman, 1989a), we accounted for three different cost components (Ordering cost (OC_t), Inventory cost (IC_t), and Backlog cost (BC_t)). We used quadratic costs to penalize higher deviations (Diehl and Serman, 1995) and we allocated a higher per-unit cost to backlogged units than to positive effective inventory to emphasize the cumulative nature and financial impact of backlogs (Oliva & Gonçalves, 2005). Finally, the objective function for each subject in the experiment can be expressed as:

$$\min TC = \sum_{t=1}^{35} (OC_t + IC_t + BC_t) \quad (3.10)$$

Where,

$$OC_t = 0.01 * O_t^2 \quad (3.11)$$

$$IC_t = 1 * \max(0, RI_t^2) \quad (3.12)$$

$$BC_t = 1.8 * \max(0, (-RI_t)^2) \quad (3.13)$$

In each treatment, the game starts in equilibrium with sufficient supply to meet retailers' orders. This is, the final customer demand starts in 50 units/period for each retailer and the supplier's capacity is equal to 100 units/period. Then, in period 4, retailers face a percentage increase in final customer demand (*step*), such that total expected demand exceeds available supplier capacity.

3.3.2. Study I: Stationary and known demand under retailer competition

In this study, we run three treatments that systematically evaluate the effect of the existence of probability of duplication and of the step in the final customer demand in subjects' ordering decisions. In the first treatment (T1), we do not allow neither for duplications ($\alpha=0$) nor for step in the final

customer demand (*step*= 0). In treatment 2 (T2), we use a step of 5 units in the final customer demand and we do not allow for duplications ($\alpha= 0$). Finally, treatment 3 (T3) includes both the probability of duplications ($\alpha= 0.1$) and the step in the final customer demand (*step*= 5). In these treatments, we fix the *supplier capacity acquisition delay* (τ_K) to 1 and we use *proportional allocation* as the allocation mechanism that the supplier will use in case of facing excess demand. Table 3.1 characterizes each treatment conducted in this study. To run our analyses, we excluded 2 observations from T1, 1 observation from T2 and 1 observation from T3 because subjects' behavior in these observations led to total cumulative costs that were separated more than three standard deviations from the average cumulative cost of the remaining observations.

Table 3.1. Experimental treatments Study 1

		$\alpha=0$	$\alpha=0.1$
$\tau_K=1$	Step= 0	T1	
	Step= 5	T2	T3

Notice that in T1, we mainly use the traditional strategies used in the Beer Game to eliminate the first three operational causes of the bullwhip effect (order batching, price fluctuations, shortage gaming), and we control for the fourth operational cause (demand forecast updating) by using a stationary customer demand and informing subjects about the pattern of the final customer demand before the game begins. Similarly, in T2 and T3, we control for three of these operational causes but we allow for shortage gaming in systems where customers may duplicate (T3) or not (T2) their orders. The main difference between this experimental study and prior behavioral studies analyzing ordering decisions is the systematic control of the main operational causes of the bullwhip effect in a competitive setting. In some previous studies, customer demand was either stochastic or non-stationary and it was unknown to participants (Steckel et al., 2004; Sterman, 1989a). In these cases, subjects need to forecast future demand, which may lead to the presence of the bullwhip effect due to forecasting errors (F. Chen et al., 2000; L. Chen & Lee, 2011). In other cases, customer demand was stationary and known by the participants, but there was no retailer competition. In these cases, subjects do not follow a simple base

stock policy due to the lack of trust on the other players in their supply chain (Croson & Donohue, 2005; Croson et al., 2014; Sterman & Dogan, 2015). This behavior leads to the presence of the bullwhip effect.

In T1 and T2, the demand is stable and publicly known and there is no probability of duplications, therefore, there is no need for safety stock and subjects will minimize their costs by ordering the same demand they perceive. Relevant literature in operations management has shown that under horizontal competition the presence of shortages can lead retailers to over-order (Armony & Plambeck, 2005; Lee et al., 1997b; Sterman, 2000). However, given that the system starts in equilibrium and there is full information about the stable final customer demand, there is no incentive to adjust the inventory level and the system should remain in equilibrium. Therefore, we do not expect order oscillation. In treatment 3, there is competition for an initial scarce supply and there is some probability of duplication of the final customer demand, which may pose some incentive to the retailer to over-order more than the other retailer to get a greater allocation.

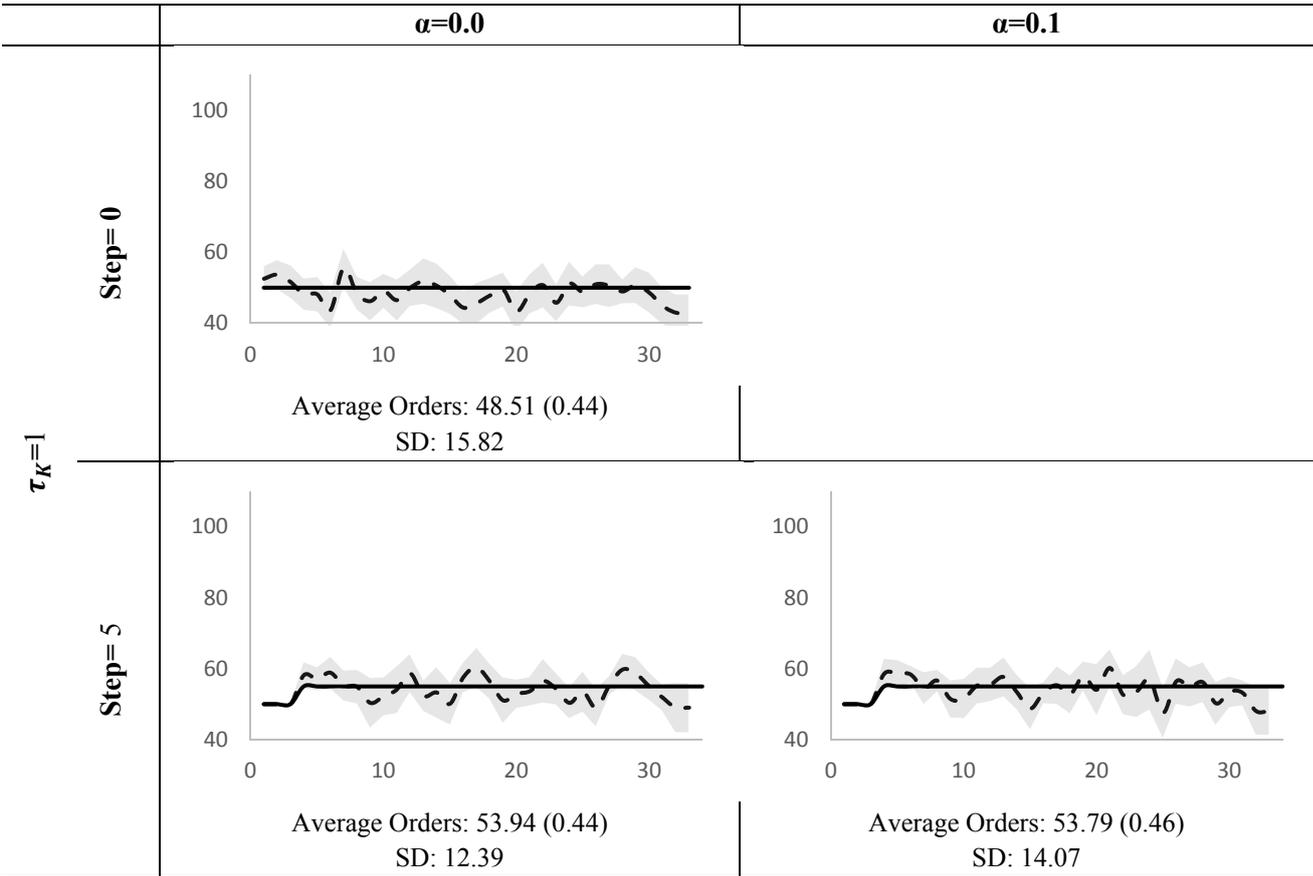
3.3.3. Results Study I

Table 3.2 shows the average order trajectories, the 95% confidence intervals, and the actual final customer demand for each treatment. As in previous behavioral studies (see Gonçalves & Villa, 2016), results show a clear underperformance in subjects' decisions. In this case, subjects initially over-order, trying to stock some extra units, then subjects under-order, trying to get rid of excess inventory. This behavior leads to oscillations around the final customer demand (50 or 55 units). In T1 and T2, these frequent oscillations are contrary to our expectations. Given that there is no probability of duplications and that there is full information about system structure, subjects do not need to forecast future final customer demand, and they do not have any incentive to deviate from a base-stock strategy. However, subjects over-order trying to anticipate a potential deviation from the other retailer, which may leave them in a disadvantageous position. In T3, however, we argue that when retailers not only compete for supply but also compete for customers, they have more incentives to carry more inventory by amplifying their orders (Anupindi & Bassok, 1999; Cachon & Olivares, 2009; Netessine et al., 2001); therefore, some oscillations are expected.

Now, by simple inspection, we can see the increase on orders' variance as one moves from the final customers (solid line) to retailers (dashed line). There is evidence of the existence of the bullwhip

effect. To rigorously check for the increase in variance of subjects' orders, we estimate the standard deviation (SD) of subjects' orders in each experimental treatment (see Table 3.2). These standard deviations are significantly higher than zero, which means that the bullwhip effect is statistically significant in all treatment conditions (p -values $<.01$, in all cases). In addition, the standard deviations of subjects' orders on T1 and T3 are significantly higher than in T2 (p -value $<.01$). This means that (i) the inclusion of a small step in the final customer demand could even reduce the bullwhip effect, because subjects' overreaction, that were unexpected in T1, can be ameliorated by the wish to properly respond to final customers; and (ii) even a small increment in the probability of duplicated orders make retailers over-react.

Table 3.2. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas) and Customer Demand (solid line)



Standard Errors in parentheses

To understand why participants do not follow the simple base-stock policy in T1, we collected participants' responses using a post-game questionnaire. Some typical responses to the strategies that subjects implemented during the experiment are:

- "I ordered always the demand to avoid inventory and backlog costs"

- “I was ordering always 50 units; however, in some cases I ordered more to see how the system reacted”
- “I tried to keep a small inventory and then I tried to balance retailers’ orders with my orders”
- “I cooperated when the other retailer did it, but if he did not, I over-ordered to get a higher fraction of units”
- “In some periods, I ordered above 50 to build some inventory, and then in some periods I ordered below 50”.

These responses suggest that the combination of different strategies may lead to system instability and to the existence of the bullwhip effect, even when the demand is known and stationary. One of the main reasons is the desire to have positive inventory. Participants prefer to have extra units to respond to potential deviations from the other retailer. This way, if the other retailer deviated from the coordinating strategy and they did not receive what they ordered, they would have enough inventory to satisfy their observed demand. Another cause of deviation from the base-stock policy is curiosity. A simple emotional factor may lead subjects to deviate and affect the general performance of the whole supply chain. Finally, another typical reason for amplifications is given by a tit-for-tat strategy. In this case, subjects try to cooperate, but in case the other retailer deviated from the coordinating strategy (due to curiosity or desire to build some inventory), they would over-order in the following periods, trying to penalize the deviation of the other retailer. These results complement previous findings by Croson et al. (2013, 2014), Sterman and Dogan (2015), where they argue that subjects deviate from the equilibrium because they did not trust their other partners in the supply chain.

3.3.4. Study II: Duplications and time to build supplier’s capacity under retailer competition

In this study, we ran a full experimental design with two different levels of three important experimental variables: *probability of duplication* ($\alpha= 0.1$ and $\alpha= 0.4$), *time to build supplier capacity* ($\tau_K=1$ and $\tau_K=4$) and *step of the final customer demand* (*step*= 5 and *step*= 20). Table 3.3 characterizes each treatment conducted in this study. As it was used in Study I, final customer demand was stationary and known by the participants. We excluded from the analysis the observations that led to total cumulative costs that were separated more than three standard deviations from the average cumulative cost of the remaining observations. Finally, notice that the results for T3 are the same that we get from Study I.

Table 3.3. Experimental treatments Study II

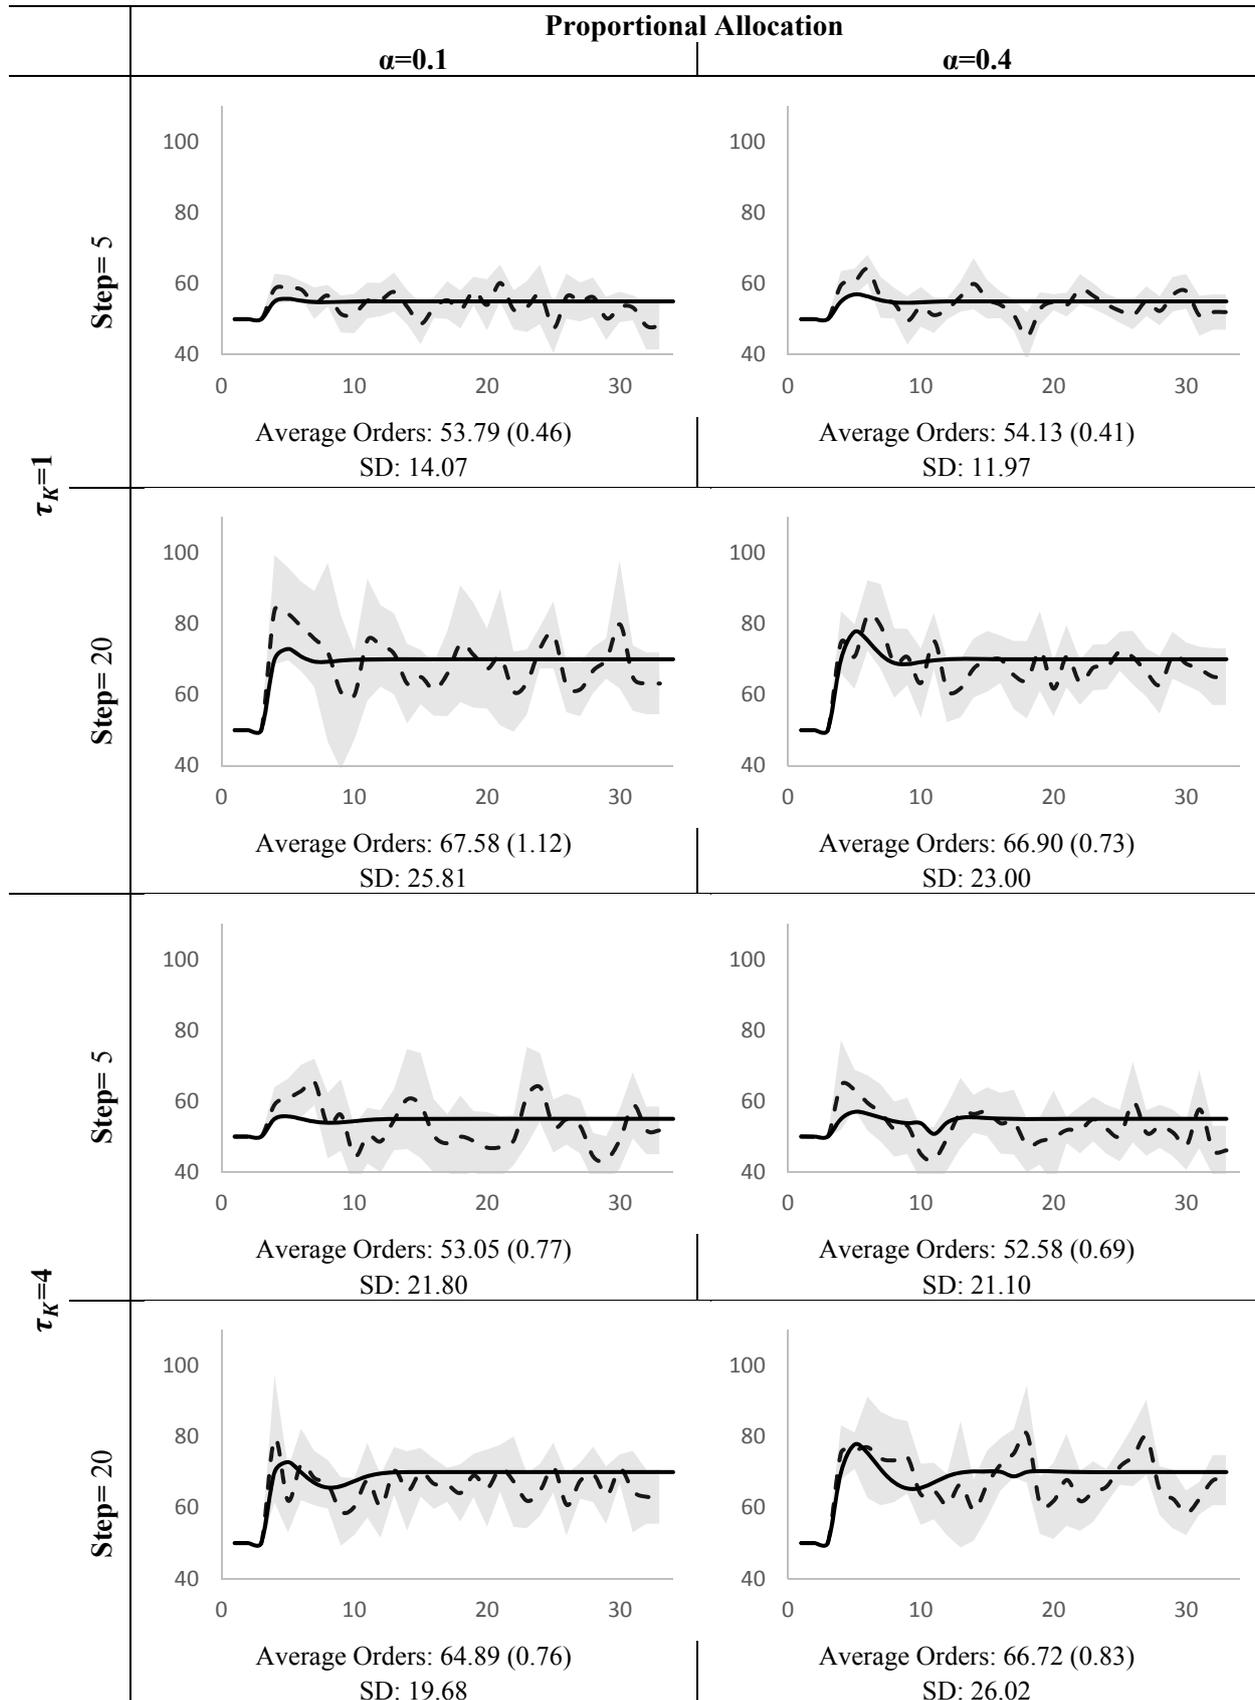
		$\alpha=0.1$	$\alpha=0.4$
$\tau_K=1$	Step= 5	T3	T7
	Step= 20	T4	T8
$\tau_K=4$	Step= 5	T5	T9
	Step= 20	T6	T10

3.3.5. Simulated trajectory benchmark

We inserted our mathematical model in Vensim DSS to determine a general strategy that can be used as benchmark to assess subject behavior in each treatment. Vensim DSS uses the Powell method as optimization method, which allows us to estimate a simulated ordering trajectory that minimizes the total cumulative costs of retailers over the 35 simulated periods. Table 3.4 shows the behavior of the optimal trajectories for each treatment in comparison with subjects' average order decisions. The trajectories of these benchmarks are characterized by an initial increase in orders, due to the surge in the final customer demand. Then, orders settle smoothly into equilibrium with a small oscillation. The magnitude and duration of the oscillation increases with the complexity of the system (higher duplication probabilities, longer delays, higher step in final customer demand).

For treatments with higher probability of duplications, we expect higher level of competition between symmetric retailers for customers; therefore, we expect a higher level of bullwhip effect (Anupindi & Bassok, 1999). Similarly, a higher increase in the final customer demand may lead to a higher over-reaction of retailers and to a higher order variability. Finally, longer time to build suppliers' capacity increases the complexity of the system; therefore, higher amplifications are expected (Gonçalves & Arango, 2010; Sterman, 2000). However, once subjects reach the equilibrium, there is no need for safety stock and subjects will minimize their costs by ordering the demand they perceive. Therefore, we do expect some order oscillation at the beginning of the experiment due to the delays and competition for an initial scarce supply, which may pose some incentive to the retailer to over-order above the other retailer to get a greater allocation. However, we also expect a reduction in this variability due to the stationarity of the final customer demand and the ability of the supplier to build capacity.

Table 3.4. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas),
Benchmark trajectory (solid line)



3.3.6. Results Study II

Table 3.4 shows the average order trajectories and the 95% confidence intervals for each treatment. We observe a higher level of orders and wider confidence intervals in treatments with higher step in the final customer demand and longer time to build supplier capacity. During the first three periods, the system is in equilibrium given that the supplier is able to satisfy retailers' orders and retailers are able to satisfy customers' demand on time. Once the final customer demand increases, retailers face a backlog very quickly, causing the probability of receiving duplicated orders to increase. The presence of backlog and inflated final customer demand lead to an increase in retailers' orders. With time, the supplier builds capacity to meet the increase in retailers' demand, so that the supplier capacity increases and surpasses retailers' orders. When the supplier capacity is large enough, she is able to satisfy retailers' demand, and with backlogs at the desired levels, final customers can cancel duplicated orders.

In addition, Table 3.4 shows the average retailers' orders and standard deviation of subjects' decisions and Appendix 3.2 shows the p -values that evaluate the significance difference of subject behavior between experimental treatments. Results show that despite the oscillations in subjects' decisions, the average orders in treatments with the same step in final customer demand are not significantly different. This is because subjects can backlog unsatisfied orders until enough capacity becomes available, which means that in the long term, retailers will be able to reduce their backlog and satisfy their customers' orders.

Subjects deviate from optimal trajectories in all treatments. Results show that subjects' ordering behavior fluctuates around the optimal trajectory in all treatments during the whole simulation horizon. Taking into account the different treatment variables analyzed in this study, we quantify the effect of each experimental variable on retailers' orders deviations (see Table 3.5). We compared the average deviation from the subjects' orders with respect to the optimal trajectories. Deviations are computed as the sum of absolute values of the difference between subjects' decisions and the optimal ordering trajectory. Initially, as shown in Table 3.4, Table 3.5 shows that there is a significant difference between subjects decisions and optimal ordering quantities for all experimental variables (all deviations are significantly different from zero).

Table 3.5. Mean comparisons among treatment variables

Variable	Variable Values	Deviation	<i>p</i> -value
			Difference of Deviations
Probability of duplication	$\alpha=0.1$	10.83 (0.31)	0.06
	$\alpha=0.4$	11.49 (0.29)	
Step in demand	Step=5	9.35 (0.27)	0.00
	Step=20	13.26 (0.34)	
Time to build capacity	$\tau_K=1$	9.90 (0.27)	0.00
	$\tau_K=4$	12.49 (0.32)	

Standard errors in parentheses

Results show that there is significant difference in the level of underperformance when we increase the probability of duplications from 0.1 to 0.4 ($\text{Diff}_{\alpha=0.1-\alpha=0.4}=-0.66$; $p\text{-value}=.06$). In addition, higher step in final customer demand and a higher time to build capacity also lead to a significant increase in the deviations from the optimal trajectory ($\text{Diff}_{\text{step}=5-\text{step}=20}=-3.91$; $p\text{-value}=.00$; $\text{Diff}_{\tau_K=1-\tau_K=4}=-2.59$; $p\text{-value}=.00$).

Now, in order to understand the effect of these experimental variables on the bullwhip effect, the differences in subjects' behavior should be analyzed in terms of ordering variance and the ability of the supplier to respond to retailers' orders. Therefore, to get a better understanding of the difference in subjects' performance, we need to compare the standard deviations of subjects' orders among treatments. To make a clean comparison among the different treatments, we analyze the standard deviation of the difference between subjects' orders and the optimal trajectories. In this way, we will discount the expected variance of the optimal solutions. Table 3.6 shows the standard deviations of these deviations from the optimal ordering trajectories in each treatment and Appendix 3.3 shows the p -values obtained by performing comparison tests under the hypothesis of equality of standard deviations between treatments.

Results show a switch in the deviations as we increase the time to build suppliers' capacity. For the cases where we have short time to build capacity ($\tau_K=1$), an increase in the probability of

duplications lead to less variation in subjects' orders (p -value $<.01$ for comparisons of T3 vs. T7 and T4 vs. T8). However, for the cases where we have long time to build capacity ($\tau_K=4$), there is not decrease in the variability of subjects' orders. For T5 and T9, there is not significant difference (p -value=0.82), and for T6 and T10, there is a significant increase in orders' variability (p -value $<.01$).

In addition, as we expected, an increase in the step of the final customer demand leads to a higher level of variation in subjects' orders in three out of four cases. Therefore, more aggressive change in final customer demand lead to more unstable retailers' orders and more instabilities for the whole supply chain.

Finally, due to the stationarity of the final customer demand, the ability of the supplier to build capacity and the full information subjects receive about the system, we were expecting a complete reduction of the bullwhip effect during the last periods of the experiment. Therefore, we extract the last 10 periods of the experiments (*after 24*) of our data and compared them with the first periods (*before 24*) to analyze the evolution of the bullwhip effect, i.e. the ordering variance, in these two time frames. We chose the last 10 periods, because we expect orders had reached the equilibrium at that moment.

Table 3.6. Standard deviation of subjects' orders by experimental treatment

			Proportional allocation	
			$\alpha=0.1$	$\alpha=0.4$
$\tau_K=1$	Step=5	All	T3 12.13	T7 10.14
		Before 24	14.72	9.98
		After 24	10.78	10.46
	Step=20	All	T4 20.94	T8 17.67
		Before 24	22.15	18.44
		After 24	17.06	15.61
$\tau_K=4$	Step=5	All	T5 17.52	T9 17.65
		Before 24	18.37	17.23
		After 24	15.39	18.59
	Step=20	All	T6 15.89	T10 21.33
		Before 24	16.30	22.65
		After 24	14.88	17.56

Table 3.6 shows the estimations of the standard deviations of the deviations of subjects' decisions before and after period 24. Results show a small reduction in the standard deviation in almost all treatments (p -values $<.01$), except in T7 (p -value $=.36$) and T9 (p -value $=.13$), where there is not even a significant reduction of the order variability.

3.3.7. Study III: Proportional vs. Turn-and-earn allocation under retailer competition

Finally, this study explores the effect of different supplier allocation mechanisms on retailers' orders. As previously argued by Cachon and Lariviere (1999c), retailers' orders depend on the allocation mechanism used by the supplier. Therefore, we are interested in analyzing the main behavioral changes that may occur when we analyze our system under two different allocation mechanisms: *proportional* (Lee et al., 1997a), and *turn-and-earn* (Cachon & Lariviere, 1999b). Table 3.7 characterizes the two treatments conducted in this study. As it was used in previous studies, final customer demand was stationary and known by the participants. In addition, we use T10 as our control group, which means that for this study we set $\alpha=0.4$, $\tau_K=4$ and Step= 20. Notice that the results for T10 are the same that we get from Study II. As in previous treatments, in this study, we excluded two observations from T11 because these observations lead to total cumulative costs that are separated more than three standard deviations from the average cumulative cost of the remaining observations. Proportional allocations is probably the most traditional allocation mechanism. In this case, a retailer receives a fraction of supplier's capacity equally to the fraction of his orders with respect to the total orders. However, a retailer expecting to get a higher allocation will have an additional incentive to order more units and therefore increase the oscillations in subjects' ordering decisions. On the other hand, turn-and-earn is a traditional allocation mechanism widely used in the automobile industry to incentivize sales. In this case, the number of units that a retailer receives from the supplier will be a function of the units sold last period. Under this sales-based allocation mechanism, retailers do not have an incentive to inflate their orders above the observed demand (actual orders plus duplicated orders). Therefore, we expect a reduction in subjects' biases with respect to the proportional allocation.

Table 3.7. Experimental treatments Study II

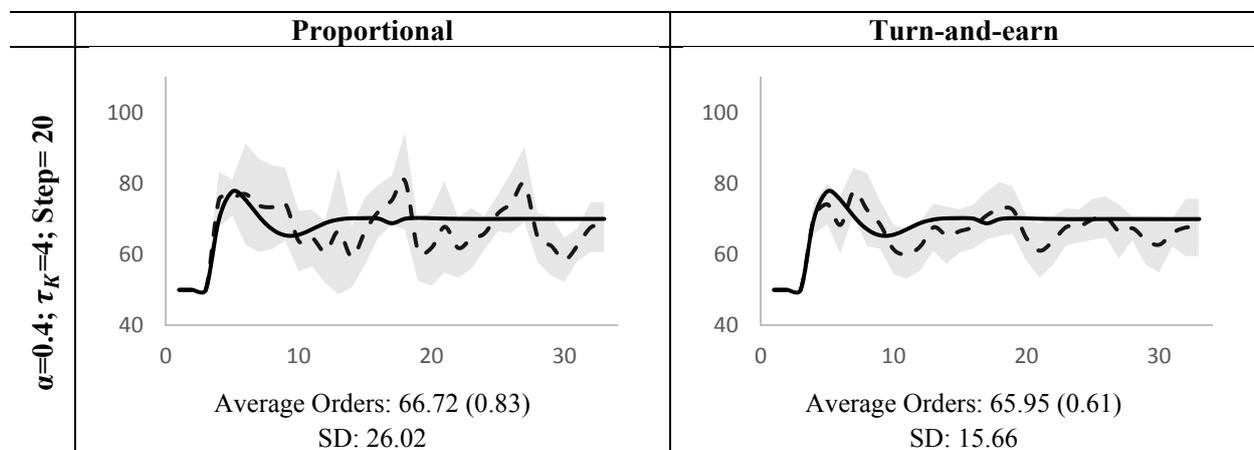
	Proportional	Turn-and-earn
--	---------------------	----------------------

$\alpha=0.4$	T10	T11
$\tau_K=4$		
Step= 20		

3.3.8. Results Study III

Table 3.8 shows the average order trajectories, the 95% confidence intervals and the optimal trajectory for each treatment in this study. We can see that despite the use of different allocation mechanisms, subjects still deviate from the optimal trajectories. However, the turn-and-earn allocation mechanism leads to a lower average deviation (8.83) from the optimal trajectory than the proportional allocation (13.95). A comparison test between the two average orders shows that there is not significant change in the average orders (p -value=.49). This is mainly because we ran the experiment for a long period of time, allowing retailers to have enough time to supply any initial unsatisfied customers. Furthermore, turn-and-earn leads to a decrease of the bullwhip effect. There is a reduction in the orders' standard deviation (p -value=.01) and, therefore, in the instabilities of the system. This means that even if both treatments have exactly the same system structure, having the supplier allocate her capacity using a reserved quantity based on previous sales (turn-and-earn), reduces subjects' incentives to over-order. However, this reduction seems to be unconscious because none of the participants' responses to the post-game questionnaire mentions the effect of the allocation mechanism in their ordering strategies.

Table 3.8. Figures of Average Orders (dashed line) with 95% Confidence Intervals (shaded areas), Benchmark trajectory (solid line)



Standard errors in parentheses

So far, we have described the main characteristics of subjects' ordering behavior and the consequences of their decisions in our three studies. Now, we will integrate the three studies to describe some general results and we will use an econometrical model to explain subjects' biases in each treatment.

3.4. General results

3.4.1. Average Costs

We analyze the effect of the different experimental variables on subjects' average cost. Table 3.9 shows the average costs separated by experimental variable, when we control for the other variables' effects. In addition, Table 3.9 shows the average ordering, inventory and backlog cost, and the corresponding percentage over the total costs for each condition of each experimental variable. Given that subjects were informed about the pattern of the final customer demand, we were expecting subjects to keep no inventory and no backlog and therefore, to allocate most of their cost to the ordering component. However, in none of the conditions the percentage of ordering cost exceeds the 4% of the total costs. This situation of not being able to make decisions that lead to the minimization of inventory and backlog costs is a clear evidence of subjects' bounded rationality and may be an explanation to the existence of the bullwhip effect. The first section of Table 3.9 shows the effect of duplications on retailers' costs. It shows how including duplications from customers significantly increases the average retailer's cost. There are not significant differences in costs when we compare $\alpha=0.1$ and $\alpha=0.4$. However, it seems that subjects prefer to make decisions that lead to increase their costs of inventory when they face a low probability of duplication, while in case of high probability, most of the costs are due to backlog. This result is expected because under the same level of unsatisfied customers, subjects in treatments with $\alpha=0.4$ will face a higher levels of backlog. Higher level of backlogs lead to higher costs and the situation is worsen when we account that we are using a quadratic cost function. The second and third sections of Table 3.9 show the effect of the step in the final customer demand and of the time to build supplier capacity on the average costs. We see that higher values of the step in the final customer demand or longer time to build supplier capacity lead to higher total average costs (p -value $<.01$ in all cases). Changing the system structure by having a longer time to build capacity makes the supplier poorly

respond to retailers' orders and therefore retailers place orders that generate significant and costly oscillations, which is a typical characteristic of the misperception of the feedback dynamics (Gonçalves & Villa, 2016; Sterman, 1989a). In addition, there is a small tendency (probably due to the cost structure) of retailers to make ordering decisions that slightly favor backlog costs when retailers face high increases in the final customer demand or when they face long time to build supplier's capacity. Finally, the last section of Table 3.9 shows the cost comparison between the two different allocation mechanisms. In this case, both systems have the same structure and results show no significant difference on average total costs. However, the turn-and-earn allocation leads to a significant reduction of costs due to the available inventory (p -value<.01). Probably, subjects realized that over-ordering and keeping inventory was not a good strategy because at the end they would receive units from the supplier based on last-period sales and not on the amount of orders placed.

Table 3.9. Average, Ordering, Inventory and Backlog cost by experimental variable

		Average Costs	Ordering Cost	Inventory Cost	Backlog Costs
Probability of duplication	$\alpha=0.0$	685.36 (49.80)	27.10 (0.36)	301.75(28.48)	356.5(42.15)
			4.0%	44.0%	52.0%
	$\alpha=0.1$	2,030.93 (102.42)	37.74 (0.60)	1,112.20 (86.16)	880.95 (61.18)
			1.9%	54.8%	43.4%
	$\alpha=0.4$	2,141.02 (79.92)	41.01 (0.54)	997.11 (62.61)	1,102.90 (54.46)
			1.9%	46.6%	51.5%
Step in demand	Step=0	921.57 (78.88)	25.39 (0.49)	386.36 (44.81)	509.81 (67.32)
			2.8%	41.9%	55.3%
	Step=5	1,638.25 (77.36)	30.62 (0.36)	1,028.2 (65.89)	579.39 (43.86)
			1.9%	62.8%	35.4%
	Step=20	2,242.39 (84.72)	47.85 (0.64)	877.98 (64.04)	1,316.60 (60.68)
			2.1%	39.2%	58.7%
Time to build capacity	$\tau_K=1$	1,164.60 (47.81)	34.41 (0.36)	632.93 (39.38)	497.26 (29.17)
			3.0%	54.3%	42.7%
	$\tau_K=4$	2,618.41 (98.34)	40.35 (0.59)	1210.40 (77.93)	1,367.70 (66.51)
			1.5%	46.2%	52.2%
Allocation Mechanism	Proportional	2,781.46 (216.37)	49.78 (1.63)	1,441.00 (198.67)	1290.70 (105.50)
			1.8%	51.8%	46.4%
	Turn & Earn	2,787.12 (226.34)	44.47 (0.82)	854.69 (117.72)	1,888.00(205.71)
			1.6%	30.7%	67.7%

Standard errors in parentheses

3.4.2. Retailers' Decision Rule

We now use a formal decision rule to model retailers' ordering behavior. We build our decision rule based on the model previously proposed by Croson and Donohue (2005). Our decision rule regresses the orders placed by subjects in period t ($O_{t,i}$), against the initial effective inventory level ($RI_{t-1,i}$), total incoming orders ($to_{t,i} = o_{t,i} + d_{t,i-}$), cancellations (c_t^i), shipments received from the supplier ($S_{t,i}$), and retailer's total outstanding orders ($B_{t,i}$). In addition, following Oliva and Gonçalves (2004), we complemented the decision rule by differentiating between positive effective inventory ($RI_{t-1,i}^+$) when the inventory ($I_{t,i}$) is higher than the backlog ($b_{t,i}$) and negative effective inventory ($RI_{t-1,i}^-$), when the opposite is true. Therefore, the decision rule for a given subject at period t can be expressed by equation (3.14). Finally, a *max* function is included in the decision rule because retailers are not allowed to cancel their previously placed orders to the supplier. Therefore, retailers' orders must be nonnegative.

$$O_{t,i} = \max(0, \beta_0 + \beta_I RI_{t-1,i}^+ + \beta_b RI_{t-1,i}^- + \beta_{to} to_{t,i} + \beta_c c_t^i + \beta_S S_{t,i} + \beta_B B_{t,i} + \varepsilon_{t,i}) \quad (3.14)$$

Notice that equation (3.14) is able to test subjects' reaction to the main operational variables involved in the system as in Sterman (1989a) and Croson et al. (2014). However, equation (3.14) uses a more general structure than the one used by Sterman's (1989a) and Croson's et al. (2014) model. In contrast to their work, our equation tests the independent effect of the main operational variables on subjects' behavior and does not make assumptions about a specific heuristic or anchor used by the subjects. Hence, equation (3.14) eliminates any assumption about the type of model subjects use to forecast (Croson & Donohue, 2005).

If subjects were rational and there were no evidence of the bullwhip effect at the retailer level, retailers would order exactly the same number of units as the total incoming orders. This means that a one-unit increase on total incoming orders from final customers should increase retailers' orders to suppliers by one unit, as well as a one-unit increase in the negative effective inventory (backlog). Therefore, we will expect a value of 1 to the parameters β_b and β_{to} in our model. Similarly, a one-unit increase in the positive effective inventory, in the cancellations, in shipments received from the supplier or in the outstanding orders (supply line) should decrease retailers' orders by one unit. Therefore, we would expect a value of -1 for parameters β_I , β_c , β_S and β_B .

3.4.3. Evaluation of decision rule

To improve our understanding about how subjects make ordering decision when facing different conditions, we used the data obtained from the experiments to estimate the unknown parameters of equation (3.14) for each treatment. To estimate these parameters, we relaxed the non-linearity condition of the decision rule and we structure the data from the experiments as a panel. This panel estimations allow us to control for individual heterogeneity. In addition, to get unbiased estimations, we run a Hausman test (H-test) to determine whether it would be better to use random or fixed effects in the estimation of our model. This test evaluates whether the effects are exogenous - in which case it would be better to use random effects -, or whether the errors are correlated with the regressors - in which case it would be better to use fixed effects -. After running the Hausman test for each treatment and assuming random effects as the preferred model, we did not find support to our hypothesis (all p -values < 0.01 for all H-tests). Therefore, we adopt fixed effects for the estimation of our model.

Table 3.10. Parameter estimations for the decision rule

Parameter	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11
β_I	-0.41 (0.03)	-0.38 (0.04)	-0.13 (0.02)	-0.37 (0.04)	-0.22 (0.02)	-0.33 (0.06)	-0.20 (0.02)	-0.33 (0.03)	-0.31 (0.03)	-0.30 (0.03)	-0.19 (0.03)
β_b	-0.01 (0.04)	0.62 (0.06)	-0.22 (0.13)	1.13 (0.25)	-0.13 (0.10)	-0.22 (0.09)	-0.07 (0.05)	0.17 (0.06)	-0.13 (0.05)	0.05 (0.06)	-0.18 (0.03)
β_{to}	-	1.40 (0.27)	1.13 (0.28)	1.39 (0.17)	1.01 (0.43)	1.17 (0.13)	1.12 (0.14)	1.02 (0.09)	0.91 (0.17)	1.05 (0.11)	0.87 (0.08)
β_c	-	-	1.29 (0.50)	-2.43 (0.92)	0.50 (0.36)	0.53 (0.32)	-0.06 (0.08)	0.06 (0.11)	0.16 (0.09)	-0.03 (0.12)	0.20 (0.06)
β_s	-0.29 (0.03)	-0.49 (0.04)	-0.12 (0.04)	-0.43 (0.04)	-0.21 (0.04)	-0.26 (0.05)	-0.07 (0.04)	-0.21 (0.04)	-0.28 (0.04)	-0.20 (0.04)	-0.01 (0.04)
β_B	-0.17 (0.04)	-0.01 (0.05)	0.07 (0.05)	-0.01 (0.06)	0.20 (0.04)	-0.09 (0.05)	0.32 (0.05)	-0.18 (0.05)	0.03 (0.04)	-0.03 (0.04)	0.17 (0.04)
β_0	66.67 (1.86)	4.91 (14.29)	-0.57 (15.07)	3.38 (10.97)	10.49 (23.10)	5.32 (8.37)	-2.20 (7.64)	11.97 (6.02)	20.47 (9.25)	10.57 (7.35)	8.04 (4.89)
Correlation	-0.53	-0.32	-0.60	-0.27	-0.41	-0.25	-0.42	-0.40	-0.56	-0.38	-0.43
F (Wald)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
H-test	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
F-Effects	0.06	0.12	0.09	0.27	0.17	0.39	0.00	0.27	0.32	0.12	0.13

Standard errors in parentheses

Finally, notice that our model does not include an independent variable t controlling for any potential trends as in the model used by Croson and Donohue (2005). Instead, we test our model for the need of time-fixed effects (F-Effects). We perform a simple Wald tests to test the likelihood-ratio of the model with and without time-fixed effects. We failed to reject the null hypothesis that all period coefficients are jointly equal to zero (p -values $>.05$ in 10 out of 11 treatments); therefore, there is no need to include time fixed-effects. We used Stata 12 to estimate the unknown parameters in each treatment. Table 3.10 shows the main parameter estimations and Appendix 3.4 shows a summary of the experimental variables used for each experimental treatment, as a reference to the reader to better follow the panel data analyses.

As expected, all the β_{t_0} coefficients are positive and close to 1. However, we identify some trends on how subjects make their decisions taking into account the orders of the final customers. First, when the delay is short ($\tau_K=1$), the coefficients are always higher than 1, signaling the tendency of subjects to overreact to final customer demand by inflating their orders up to a 40%, while longer delays allow subjects to adjust smoothly their orders to the suppliers according to the their final customer demand. Similarly, when the step in the final customer demand increases, retailers' orders increase. This shows a tendency of participants to inflate even more their orders when the demand increases. Also, it seems that subjects' overreact a bit more to the final customers' orders when they face a lower probability of duplications. Probably, the fact that there is high probability of duplications make subjects more aware of the existence of phantom orders and therefore they make decisions that are more consistent. Finally, as it was expected, the turn-and-earn mechanism reduces the tendency of retailers to over-order. However, in this case, retailers prefer to order less units than the amount of units ordered from their customers.

As expected, all the β_I coefficients are negative; indicating that for one additional unit of positive inventory, the retailers order will decrease. However, the coefficients are significantly higher than -1, meaning that the reduction of the orders does not reflect the same amount of the on hand inventory.

We observe that in general the β_I estimations are close to -0.3 in all treatments. However, there is a small tendency of the estimations to move toward -1 (the expected value) as the step in the final

customer demand moves from 5 to 20. This means that a higher peak in demand gives more incentives to subjects to reduce their orders based on their inventory levels.

We expected the effect of the backlog's parameter (β_b) to be positive, representing an increase in the orders according to the number of units owed to the final customers. However, we found that 6 out of 11 of our estimations have negative values. Although they are not significantly different from zero for T1, T3, T5 and T7. This means that subjects fail to properly consider their backlog levels when they make ordering decisions. However, there is a small tendency of the estimations to move toward the positive region for a short time to build suppliers' capacity. Finally, we see that the absolute value of the estimations of β_b are significantly different from the estimations of β_I , which shows that subjects react different to inventory than to backlog.

We expected the coefficients for the cancellations (β_c) to be negative, i.e. for each unit cancelled, we would expect the order to be reduced by one unit as well. Here, there is no effect of cancellations for T1 and T2, because there are no duplications. Results show that most of the coefficients are not significantly different from zero, especially when subjects face a high level of duplications or long time to build suppliers' capacity. Probably, a higher complexity in the system makes subjects to disregard the cancellations coming from their final customers. However, for the simplest treatments (T3 and T4), subjects show a mixed behavior. Finally, we see a significant and positive effect of cancellations on subjects' orders when the supplier allocates her constraint capacity using a turn-and-earn strategy. This result probably highlights the importance that selling an additional unit has to the retailer, so that facing cancellations is unacceptable and they prefer to over-order to avoid future cancellations and potential penalizations in the allocation process.

As expected, all signs for the coefficient of shipments received from the supplier (β_S) are negative. This means that retailers reduce their orders as they receive additional units from the supplier. However, all of the estimations are significantly higher than -1 (the expected parameter estimation). As a general trend, we observe that when the step increases (from 0 to 5, and from 5 to 20), the reduction in the orders due to order fulfillment is also bigger. Hence, it seems that the higher step allows for a better recognition of the orders being fulfilled, so that subjects reduce their future orders, to avoid having high inventories. In addition, the turn-and-earn mechanism gives lower incentives to reduce the orders based on suppliers'

order fulfillment. In fact, this effect is not different from zero. This may happen because turn-and-earn makes subjects focus on what they are actually selling, shifting the attention away from the orders received.

We expected the effect of the outstanding orders (β_B) to be negative, representing a decrease in the ordering decision according to the supply line, i.e. to the number of units already ordered to the supplier but not yet received. However, most of the coefficients (T2, T3, T4, T6, T9, T10) are not significantly different from zero, which means that subjects do not consider the outstanding orders they have with the supplier when they make ordering decisions. In addition, if subjects were properly accounting for the supply line, the estimations for β_B (outstanding orders) should be equal to β_I (inventory level). However, estimations show that $\beta_B > \beta_I$ (closer to zero) in all treatments. This is a clear evidence that subjects underweight the supply line (Serman, 1989). Finally, there is a weak tendency in the β_B parameters that shows that a higher step in the final customer makes subjects more aware of the supply line and demand. Therefore, we see lower values (closer to -1) for β_B in treatments with higher step in the final customer demand.

These results show a high level of subjects' deviation with respect to the expected behavior. However, parameter estimations show that these deviations (biases) are in general reduced when subjects face systems with higher complexity (higher step in final customer demand, longer time to build supplier's capacity and higher probability of duplication). These results are also aligned with previous physiological studies of dual process theory, where researchers have claimed that subjects' decisions are driven by two independent systems: System I or the automatic system, and System II or the rational system. In this case, subjects facing low effort tasks will make decisions mainly driven by System I; however, more complex tasks induce conscious judgments and therefore, lower biases (Kahneman, 2011; Stanovich & West, 2000). At the same time, results from Table 3.9 show that less complex systems lead to average costs that are significantly lower than the costs in more complex systems. This means that there is a higher penalization of subjects' underperformance when they are immersed in systems that are more complex; therefore, the current reduction of subjects' biases is not enough to have lower costs (or higher performance) if they need to deal with highly complex system dynamics. Consequently, in order for subjects dealing with high complexity to get (at least) the same level of

performance (costs) than less complex systems, they would need to have a higher reduction of the biases. Finally, when comparing systems with the same structure but different supplier allocation mechanism (T10 vs. T11), we find that the proportional allocation mechanism makes subjects significantly reduce their overall biases but the turn-and-earn allocation leads to a reduction of orders' variability and, therefore, of the bullwhip effect.

3.5. Conclusions

This chapter presents a behavioral study of the existence and consequences of the bullwhip effect in a non-serial supply chain. We use a simple supply chain composed of one supplier and two competing retailers who face a stable and known final customer demand. In an attempt to focus on the main behavioral causes leading to an increase in order variability, we built a clean experimental design to systematically control for the main operational variables leading to supply chain instability: order batching, price fluctuations, rationing and shortage gaming, and demand forecast updating (Cachon & Lariviere, 1999a; Lee et al., 1997a).

Initially, we develop a mathematical model to capture endogenous decision policies for both supplier's capacity investment and customers' ordering. We evaluate the effect of different duplication probabilities, different supplier's capacity acquisition delay, different strength in final customer demand and two different supplier allocation mechanisms on retailers' ordering decisions. We created three different and complementary studies that make significant contributions to previous work on the behavioral operations area.

Results from the first study evidence that the bullwhip effect persists in a system with two competing retailers even when subjects do not have incentives to inflate their orders. We show that in a system where there is neither positive probability of duplications nor changes in the final customer demand and where retailers have full information about the system structure, subjects still deviate from an equilibrium strategy. Retailers may deviate from the equilibrium for multiple reasons. They amplify their orders in an attempt to build an unnecessary safety stock to respond to potential deviations from the other retailer. Alternatively, the deviation could be driven by an emotional factor such as curiosity. The desire of subjects to know what would happen should they deviate from the equilibrium, leads to

interaction and penalization strategies (e.g. tit-for-tat) between the retailers, which activates unnecessary system dynamics that lead to order oscillations and higher overall costs.

In our second study, we create a full experimental design to evaluate the effect of three different experimental variables (duplication probabilities, different supplier's capacity acquisition delay, different strength in final customer demand) in subjects' performance. Results show that subjects deviate significantly from the optimal trajectories in all treatments. These deviations are characterized by a continuous fluctuation of subjects' orders around the optimal trajectory. A cost analysis shows that subjects fail to make decisions that allocate most of the cost to the ordering component. In addition, systems with lower dynamic complexity (lower duplication probability, shorter time to build capacity and smaller step in the final customer demand) lead to lower average costs, while subjects' low performance in systems with higher dynamic complexity is explained by the difficulty of subjects to reduce their backlog cost. However, estimations of the parameters of our decision rule show that subjects' biases are reduced when subjects face systems with higher complexity. In more complex treatments, for example, subjects are more aware of the supply line. These findings can be explained by the concepts of dual process theory. In this case, when subjects face situations with higher complexity, their mental System II (rational system) is activated. Therefore, subjects are expected to make more rational decisions. In contrast, when facing simpler systems, subjects' System I (automatic system) plays a higher role, leading subjects to make more reactive decisions (Kahneman, 2011; Stanovich & West, 2000). The results of this study show that the cost function places a higher penalization of subjects' underperformance when they are immersed in systems that are more complex and where higher deviations are expected. Given this higher penalization, despite the increased ability of subjects to make more rational decisions (System II) when immersed in high complex systems, it is not possible for them to achieve the same level of performance (costs) than if they were making decisions in a less complex system. To increase their performance (reduce costs), a higher reduction of their biases would be needed.

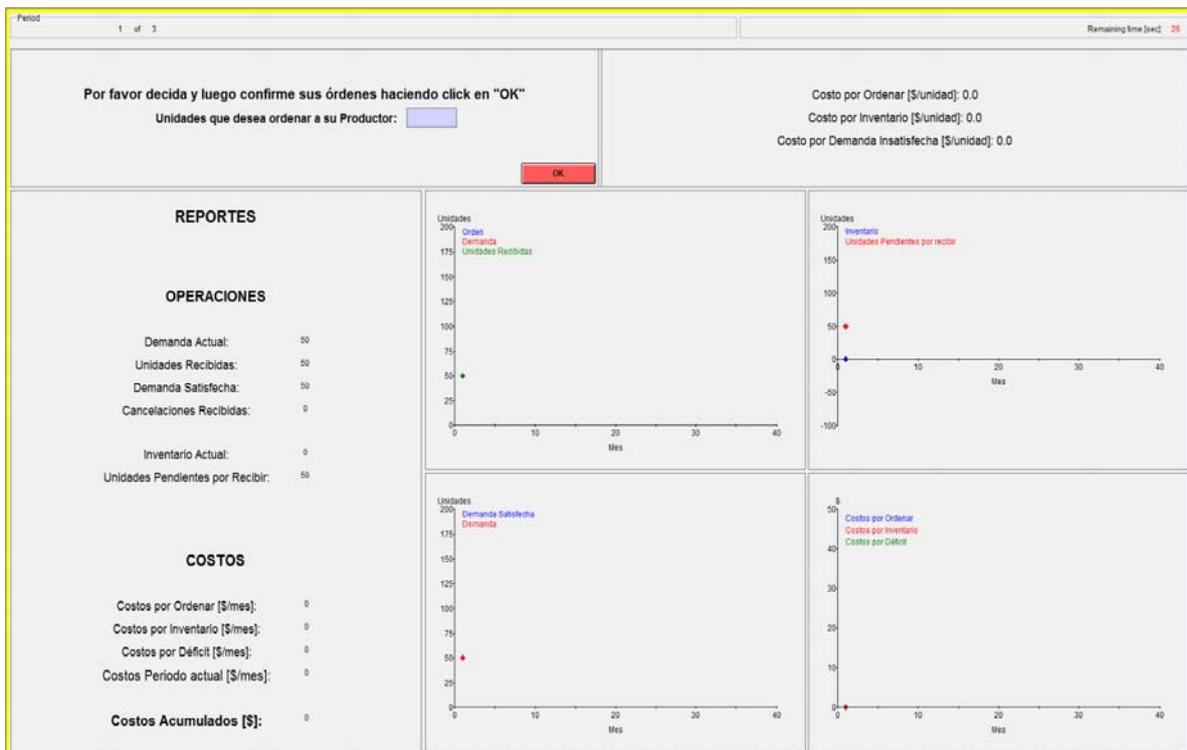
An additional finding from our second study is related with the poor ability of subjects to eliminate the bullwhip effect. Given that subjects were facing a system with a constant final customer demand, with a supplier able to build capacity and with full information about the whole supply chain structure, we were expecting a complete reduction of orders' variability after a long period of time had

passed. Results show a small but significant reduction in the order variability for decisions made during the last 10 periods of the experiments. However, the bullwhip effect is never completely eliminated. These results are also aligned with the findings we discuss in our first study.

Finally, the third study of this chapter presents a clear contribution of the effect of different allocation mechanisms on subjects' behavior. Results show that the biases leading to the amplification of orders are still existent in subjects' behavior regardless the allocation mechanism used. However, the level of subjects' underperformance may vary from mechanism to mechanism. In this chapter, we evaluate the effect of two different allocation mechanisms: one mechanism that encourages retailers to compete for supplier's constraint capacity by making an allocation proportional to the outstanding retailers' orders (proportional allocation) and the other mechanism reduces subjects' incentives to over-order by making an allocation based on last-period sales (turn-and-earn). Results show that, under identical system structure, the turn-and-earn allocation mechanism leads to a significant reduction of the bullwhip effect. This means that having the supplier allocate her capacity using a last previous sales (turn-and-earn) reduces subjects' biases, which is also aligned with previous analytical results (Cachon & Lariviere, 1999b; Lu & Lariviere, 2012). However, it is important to notice that the reduction in order variability seems not to be a conscious process in subjects mind. We arrived at this conclusion because none of the participants' claimed that their ordering decisions were affected by the allocation mechanism used by the supplier.

This work can be extended in many directions. For example, in this chapter we have considered a parsimonious model to study subjects' behavior; therefore, different decision rules can be used to explain these behaviors and try to find mechanisms that may reduce subjects' biases. Second, there are different allocation mechanisms that may be used and tested to determine subjects' reaction to the different kind of incentives offered by each one of them. In addition, it would be ideal to test how subject's behavior would change if we automated the decisions of one of the retailers, so that the automatized retailer uses an optimizing strategy. Finally, in our experiment the supplier was automatized by the computer. Therefore, it would be interesting to relax this condition and allow an additional subject to play the role of the supplier to test what is the preferred allocation strategy followed by the supplier and how retailers react to the introduction of a real supplier.

Appendix 3.1. Interface of the experiment in Z-tree (in Spanish)



Appendix 3.2. p -values for comparison of orders' means for treatments in Study II

	T3	T4	T5	T6	T7	T8	T9
T4	0.00						
T5		0.00					
T6	0.00	0.0419	0.00				
T7	0.58						
T8		0.59			0.00		
T9			0.65		0.06	0.00	
T10				0.12	0.00	0.87	0.00

Appendix 3.3. *p*-values for comparison of orders' standard deviations for treatments in Study II

	T3	T4	T5	T6	T7	T8	T9
T4	0.00						
T5	0.00						
T6		0.00	0.01				
T7	0.00						
T8		0.00			0.00		
T9			0.82		0.00	0.00	
T10				0.00			0.00

Appendix 3.4. Overall experimental treatments

		Proportional			Turn-and-earn
		$\alpha=0.0$	$\alpha=0.1$	$\alpha=0.4$	$\alpha=0.4$
$\tau_K=1$	Step= 0	T1			
	Step= 5	T2	T3	T7	
	Step= 20		T4	T8	
$\tau_K=4$	Step= 5		T5	T9	
	Step= 20		T6	T10	

Chapter 4. Transshipments in Supply Chains: Beyond the Analytical Models

Abstract

Behavioral operations studies in inventory management have focused on understanding the decision-making processes and heuristics that explain subjects' behaviors in a single actor problem. However, interactions among multiple agents has received low attention, despite its importance for the creation of better policies that lead to the improvement of real operations. I contribute to this literature by experimentally exploring the effect of transshipments on a multi-retailer problem. I consider a system composed of one supplier and two symmetric retailers (Newsvendors) at the same echelon level. Building on a formal mathematical model, I test whether subjects coordinate through any transshipment strategy. I ran different experimental treatments analyzing the effect of different (i) profit conditions, (ii) communication strategies and (iii) behavioral best response heuristics. Results show a general retailers' underperformance in all profit conditions (presence of the pull-to-center behavior). However, mechanisms like Face-to-Face communication and best response heuristics lead to an increase in supply chain coordination.

Key words: Transshipments, Communication, Newsvendor, Nash Equilibrium

4.1. Introduction

The Newsvendor problem is one of the most studied problems within the Operations Management field. This problem captures a single-period decision problem, where a manager makes a procurement order, a production order, or inventory plan before the realization of an uncertain demand. In this kind of problem, both leftovers and shortages are costly (Eeckhoudt et al., 1995; Schweitzer & Cachon, 2000): if a manager orders too much, she will have to salvage leftovers or dispose them at a loss; or if she orders too little, she will forgo additional profits (Arrow et al., 1951).

Previous research in Behavioral Operations on the Newsvendor problem has focused mainly on describing decision making biases (e.g. pull-to-center behavior, demand chasing, loss aversion) and on deriving behavioral models (e.g. anchoring and adjustment, loss aversion) that aim to explain those biases in a single-actor problem (Bolton & Katok, 2008; Bostian et al., 2008; Croson & Ren, 2013;

Schweitzer & Cachon, 2000). However, limited behavioral extensions have been done to understand its applications within a multiple-actor setting.

In addition, real world problems are characterized by the interaction of multiple actors within a modern information technology system that allows both a continuous communication among actors and a easy track of goods (Axsäter, 2003). Therefore, a better understanding of the behavioral interactions among different agents may improve coordination in a supply chain.

In this regard, I am interested in understanding the behavioral factors that influence subjects' ordering decision in a multi-retailer setting with transshipments. Transshipments are known as the monitored movement of material among multiple locations at the same echelon level (Herer et al., 2006). Transshipments are a common practice in many industries (both intra- and interfirm) as a mechanism to alleviate the problem of localized demand shocks, which encourages independent retailers to coordinate and share their inventories to achieve a better match between supply and demand (Dong & Rudi, 2004; Rudi et al., 2001, Sošić, 2006). Therefore, accurate transshipment decisions may improve stock polices, reduce costs and create better customer service by gaining a source of supply whose reaction time is shorter than the regular supply (Herer et al., 2006).

I present a mathematical model to characterize a system formed by two identical Newsvendors that place decentralized orders to a unique supplier. The supplier has enough capacity to meet Newsvendors' orders, and Newsvendors face the same cost structure and (uncertain) demand distribution. Additionally, I ran different behavioral experiments under different product margin conditions and different interactions among Humans (behavioral Newsvendor) and Computerized heuristics (knowledgeable Newsvendor) to evaluate the effect of Nash Equilibrium, Face-to-Face communication, static best response and dynamic best response policies on the supply chain performance.

The rest of the chapter is organized as follows. Sections 2 and 3 present the literature review involving a Newsvendor problem with transshipments. I formulate the research questions and provide theoretical benchmarks for traditional Newsvendor tasks. Section 4 describes the system of study, provides the experimental design, and describes the main results for the main experimental treatments. Sections 5 takes advantage of the results from section 4 and the estimation of behavioral models to

analyze different practical improvements for a supply chain system with transshipments. Finally, I summarize the main results and provide managerial implications and opportunities for future research.

4.2. Literature Review

Analytical models of the Newsvendor problem date back to the studies of Edgeworth (1888), who argued that the amount of money available in a bank should be properly managed to satisfy uncertain cash withdrawals from account holders. Later, Arrow et al. (1951) formalized the model incorporating inventory control policies under demand uncertainty. Since then, many scholars in the area of inventory management under demand uncertainty have analyzed the Newsvendor problem attempting to introduce extensions in various directions. An analytical extension of the Newsvendor problem allowing transshipments among multiple retailers seems to be that of Krishnan and Rao (1965), which assumes single-period order-up-to policy and equal costs at each retailer location. They show that when the locations are identical in their cost parameters, it is optimal for the retailers to stock at an equal fractile. Robinson (1990) extends Krishnan and Rao study to the multi-period case and present the optimality proofs. Sošić (2006) takes into account the effect of free communication among retailers and creates an analytical model assuming that retailers can freely join or leave alliances. Tagaras (1989) defines a set of assumptions that lead to "complete pooling". Complete pooling means that if one location has excess stock while another location is short, the number of units transshipped will be the minimum of the excess and the shortage (Paterson et al., 2011). Rudi et al. (2001) study transshipments between two independent Newsvendors and shows that, in general, maximizing the profit of each Newsvendor will not lead to the maximization of the whole system. However, Hu et al. (2007) formulate the conditions under which the system can be coordinated. These emerging interactions among retailers can be modeled as a game among independent actors (Newsvendors), and the game can be analyzed using Nash equilibria (Rudi et al., 2001).

However, the behavioral approaches to study the Newsvendor problem are recent, dating back to Schweitzer and Cachon (2000)'s seminal Newsvendor problem laboratory experiment. Experimental results show that, despite the Newsvendor problem's simple structure, individuals systematically deviate from the expected profit-maximizing quantity (e.g., Benzion et al., 2008; Bolton & Katok, 2008; Bostian et al., 2008). As described by Schweitzer and Cachon (2000), "subjects consistently ordered amounts

lower than the expected profit-maximizing quantity for high-profit products and higher than the expected profit-maximizing quantity for low-profit products” (p. 418). This systematic bias is known as *pull-to-center behavior* (Bostian et al., 2008).

Given the pervasiveness of pull-to-center behavior (e.g., De Véricourt et al., 2013; Kremer et al., 2010; Moritz et al., 2013), additional lab experiments have explored different de-biasing mechanisms. For example, motivated by the theory that feedback on foregone options can help learning process (Brown, 1951), multiple authors tested the effect of experience and feedback on subjects’ performance (Bolton & Katok, 2008). Results show that allowing subjects to make ordering decisions during multiple periods and providing individuals with an improved outcome feedback usually leads to ordering decisions that are closer to the optimal inventory ordering quantities (Benzion et al., 2008; Bolton & Katok, 2008; Bostian et al., 2008).

Another common bias presented in behavioral experiments of the Newsvendor problem is *demand chasing*. Demand chasing is the propensity of adjusting ordering decisions towards the prior observed demand. In their seminal paper, Schweitzer and Cachon (2000) claim that it is more likely that subjects adjust their order quantities toward the prior demand realization than not. Benzion et al. (Benzion et al., 2008) presented the same results even under different demand distributions. Under different profit conditions, Moritz et al. (2013) found that subjects with low cognitive reflection scores show stronger demand chasing than subjects with high cognitive reflection scores.

In order to explain subjects’ behavior under different Newsvendor settings, different structural models have been proposed and estimated. For example, Bostian et al. (Bostian et al., 2008) explored anchoring and adjustment models. Ho et al. (2010) explored a reference dependence model that includes asymmetric psychological costs of leftovers and shortages in a multi-location Newsvendor setting. Becker-Peth et al. (2013) explored a behavioral model that includes anchoring, loss aversion and mental accounting for designing better buyback contracts. These different models help better understand how different behavioral factors take place under different situations.

Previous research in behavioral operations on the Newsvendor problem has focused mainly on describing decision-making biases, testing different de-biasing mechanisms and deriving heuristics that aim to explain subjects’ decisions in a single actor problem (Croson & Ren, 2013). However, the

understanding of the interactions among multiple subjects have been barely studied. A couple of exceptions are the research done by Becker-Peth et al. (2013) where they analyze buyback contracts for behavioral Newsvendors; and the research done by Ovchinnikov et al. (2015), where they describe the behavioral aspects of two Newsvendors under competition.

Analyzing the interactions and collective action among multiple subjects and understanding subjects' behaviors requires also understanding the difficulty for people to reach common goals (Olson, 1965). Ostrom (1990, 2000) explains that individuals neither apply nor learn Nash equilibrium strategies when they face different kind of dilemmas, but they rather use other types of (behavioral) models, that usually lead to non-cooperative behaviors (Cardenas, 2000). However, biases of these behavioral models can be reduced by including Face-to-Face communication as a coordination mechanism that creates a trusty environment that reduces non-cooperative behaviors (Ahn et al., 2011; Castillo & Saysel, 2005; Ostrom, 1998).

I contribute to the literature of Behavioral Operations by experimentally exploring the effect of transshipments on multi-Newsvendor system. Whereas it is likely that the Newsvendor problem with transshipments will affect subjects' ordering decisions compared to single Newsvendor setting, the inclusion of subjects' interactions make it difficult to determine the magnitude and direction of the effects of each variable. Therefore, this chapter is exploratory in this regard.

In particular, the main questions this chapter aims to answer stem from the research gaps briefly described above, and can be summarized in: (i) Are Newsvendors, in a transshipment setting, prone to the common biases observed in Newsvendor setting? (ii) Are Newsvendors, in a transshipment setting, prone to other biased inventory ordering behaviors not previously observed in Newsvendor settings? (iii) How previously proposed behavioral models explain ordering behavior in a Newsvendor setting with transshipments? (iv) How do different levels of communication among Newsvendors (in a transshipment setting) affect inventory-ordering behaviors? (v) How should a Newsvendor strategically respond to another Newsvendor strategy?

4.3. Analytical Background

4.3.1. Isolated Newsvendor

The Newsvendor problem characterizes situations where a manager has to make an inventory decision before a realization of an unknown demand occurs, and where both leftovers and shortages are costly (Eeckhoudt et al., 1995; Schweitzer & Cachon, 2000). Therefore, the basic Newsvendor problem focuses on minimizing the expected cost function given by the following expression:

$$E[C] = C_o \int_0^Q (Q - D)f(D)dD + C_u \int_Q^\infty (D - Q)f(D)dD \quad (4.1)$$

where Q is the Newsvendor ordering decision, f is the probability density function of demand D , and C_o and C_u are the costs associated with over-ordering and under-ordering, respectively. The first-order condition of (1) leads to the optimal decision Q^* of a classic Newsvendor problem. This optimal solution is a base-stock policy that provides a direct relationship between the overage/underage cost ratio and the probability of overestimating D at the optimum (Cachon & Lariviere, 1999a). It is well known that the optimal inventory order quantity Q^* is given by the critical fractile solution:

$$F(Q^*) = \frac{C_u}{C_u + C_o} \quad (4.2)$$

where F is the cumulative distribution function of demand. This critical fractile is commonly used to classify products; a high-profit product is considered when $C_u/(C_u+C_o) \geq 0.5$, and a low-profit product is considered when $C_u/(C_u+C_o) \leq 0.5$ (Schweitzer & Cachon, 2000).

4.3.2. Newsvendor problem with transshipments

Of particular interest for this chapter are extensions to Newsvendor problem with transshipments. The Newsvendor model provides the basis of most existing literature on transshipment (Rudi et al., 2001). The sequence of events in a transshipment problem under complete pooling occur in the following order: (i) Newsvendors place orders, (ii) Supplier supplies each Newsvendor orders, (iii) Final customer demand takes place, (iv) Demand is satisfied, and (v) Potential transshipments among Newsvendors and additional final customer demand satisfaction take place. Figure 4.1 represents the four scenarios that both retailers could face during a selling season. The X- and Y- axes represent the potential final customer demand faced by each Newsvendor (D_1 and D_2) and q_1 and q_2 are the ordering decisions placed

by each Newsvendor before the beginning of the selling season. In sector I and II, both Newsvendors face surplus and shortage, respectively; hence, transshipments are not feasible. In sector III, some transshipment are feasible from one Newsvendor to (completely) satisfy the shortage faced by the other retailer. Finally, in sector IV, some transshipments are feasible from one Newsvendor to partially satisfy the shortage faced by the other Newsvendor (Krishnan & Rao, 1965; Rudi et al., 2001).

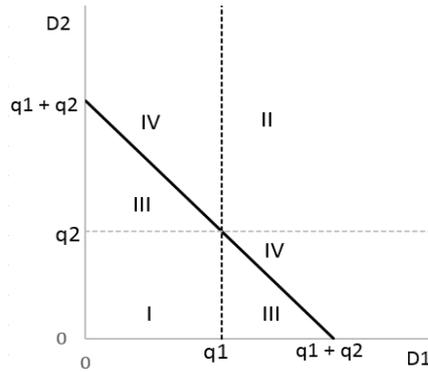


Figure 4.1. Transshipments scenarios

Now, taking advantage of the optimal solution of a typical Newsvendor problem and the potential transshipment scenarios explained in Figure 4.1, I can redefine the C_o and C_u for a transshipment problem as:

$$C_o = (c - b) * P(D_1 + D_2 < q_1 + q_2 | D_1 < q_1) + (t - r + c) * P(D_1 + D_2 > q_1 + q_2 | D_1 < q_1) \quad (4.3)$$

$$C_u = (r - c) * P(D_1 + D_2 > q_1 + q_2 | D_1 > q_1) + (t - c + b) * P(D_1 + D_2 < q_1 + q_2 | D_1 > q_1) \quad (4.4)$$

Where, c is the unit cost, b is the salvage value, t is the transshipment cost and r is the unit revenue. Integrating these terms into the critical ratio, simplifying and rearranging leads to a critical ratio for a Newsvendor problem with transshipments when the locations are identical in their cost parameters (Krishnan & Rao, 1965). This gives the best response order quantity of Newsvendor 1 (q_1^*) as a function of q_2 (See Appendix 4.1 for further details on this formulation).

$$q_1^* = F_1^{-1} \left(\frac{r - c + (t + b - r) * F_{12}(q_1^* + q_2)}{t} \right) \quad (4.5)$$

Where F_1 and F_{12} are D_1 's and $(D_1 + D_2)$'s cumulative density distribution. Equation 3 on Dong and Rudi (2004) presents an equivalent expression for equation (4.5) (although they made the mathematical deduction directly from a retailers' profit function). Robinson (1990) proves that this retailer problem is piece-wise linearly concave in the order quantities, therefore the problem can be

solved using the first order condition of the retailers profit function with respect to the order quantities. The solution of this first order condition lead to the critical fractile presented on equation (4.5). To assure complete coordination between these two independent and symmetrical Newsvendors, the first best orders (the optimal Newsvendor order quantities that maximize the overall system profit) can be obtained making the critical fractile of both Newsvendors equivalent (Dong & Rudi, 2004; Krishnan & Rao, 1965). Figure 4.2 presents an example of (i) the best response polices for each Newsvendor as a function of the decision of the other Newsvendor (e.g., $q_1^*(q_2)$), and (ii) the Nash Equilibrium (first best) solutions. The Nash equilibrium is located where the Newsvendor-1's best response ($q_1^*(q_2)$) crosses Newsvendor-2's best response ($q_2^*(q_1)$).

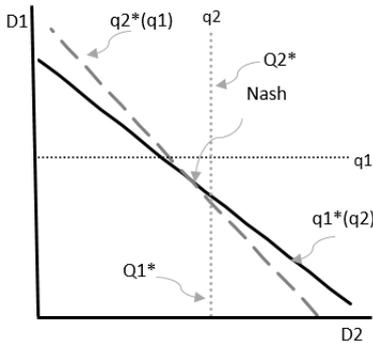


Figure 4.2. Transshipments best response policy and Nash equilibrium order

These benchmarks allow comparing the performance of subjects' ordering decisions in any experimental setting. Notice that the Nash equilibrium ($q_1^*(q_2^*)$) does not necessarily coincide with the optimal solution obtained for the isolated Newsvendor (Q^*).

4.4. Base Case Experiment (BC): Human vs. Human

This experiment allows investigating if the typical Newsvendor biases (pull-to-center, demand chasing and asymmetric reaction to over- and under-ordering) under two profit conditions: low profit (LP) and high profit (HP) are still present in a decentralized supply chain composed of one computerized supplier and two independent Newsvendors. In this system, the supplier has enough capacity to supply Newsvendor' orders before the final customer demand is known. Newsvendors face the same cost and stochastic demand structures and they sell an identical product under a complete pooling policy. Figure 4.3 shows the structure of the system, where the solid lines indicate product flow and the dashed lines indicate the information (orders) flow.

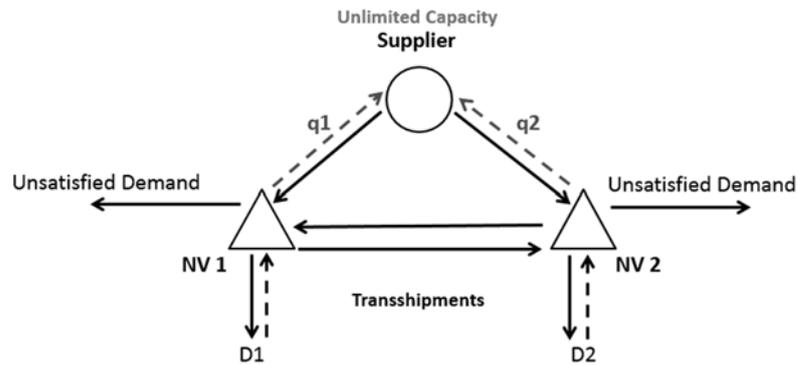


Figure 4.3. Structure for a single-supplier multiple-Newsvendor supply chain with transshipments

4.4.1. Design of the BC experiment

The experiment uses a set of parameters and a design used in previous experiments. Each Newsvendor faces a final demand D that follows an uniform distribution $D \sim U[1, 100]$ with integer values. Therefore, $(D1 + D2)$'s cumulative density distribution will be a triangular distribution, which will be easy to manage for estimating F_{12} in equation (4.5). The unit revenue r is 12 \$/item. Leftovers at the end of each round will be lost ($b=0$). The unit cost (c) determines the two treatment conditions analyzed: the first treatment (T1) considers a unit cost (c) of 2 \$ in High Profitability and the second treatment (T2) considers a unit cost of 10 \$/item in Low Profitability. Transshipment cost (t) is 1\$/item, and will be assumed by the retailer receiving the items. It is common to fix the transshipment cost using direct negotiation between the retailers (Rudi et al., 2001); however, I fixed the transshipment cost in our experiments, which allow me to focus in understanding the behavioral aspects of ordering decisions, isolating it from the subgame involved in the pricing strategies among the Newsvendors.

I follow the standard experimental economics protocol to guide the laboratory (Friedman & Cassar, 1995; Friedman & Sunder, 2004; Smith, 1982). I implemented the experimental treatments in a management flight simulator, using the computer simulation software Z-tree (see Appendix 4.2 for an example of the interface designed in this software) (Fischbacher, 2007). For each BC treatment, I recruited 32 undergraduate students in management and industrial engineering at the Universidad Nacional de Colombia. Participants were told that for their participation in the experiment, they would earn a show-up fee of COP\$10.000 (approximately US\$5) and a variable amount contingent on their performance, between COP\$0 and COP\$30.000 (US\$0 - US\$15) (Induced Value theory - Smith's

(1976)). Subjects were randomly assigned to treatments following a between-subjects design, i.e. a subject takes part in one treatment only, ruling out confounding factors due to order-of-presentation effects. Upon arrival, participants were given a set of instructions (see Appendix 4.3 for an example) describing the specific experimental treatment and they were allowed to ask questions and to test the computer interface, before the official experiment started. Subjects were paired with another subject in the room (Human vs. Human), such that both of them were making decisions at the same market and with the same (virtual) supplier. Then, a 15-minutes briefing to the participants took place. During the briefing section, a moderator read loudly the instructions and then made a small presentation about the main features of the system. Then, subjects were asked to answer 10 control questions about different scenarios that reflected potential outputs they could face during the official experimental session. After the control questions, subjects took few minutes to familiarize with the simulator and understand the experiment's flow.

After the warm up section, the official experimental session started. Subjects played the role of one of the Newsvendors placing orders to the supplier and trying to maximize their cumulative profits at the end of the experiment (30 independent periods). Each period, subjects had to choose an order quantity q of a general item, which arrived before the start of the next selling period. Transshipments among subjects were automatically generated by the system at the end of each period. Subjects had full information about the experimental parameters and were able to see a result screen showing full information from past periods. This full-information approach tries to simulate the role of access to point of sale (POS) information, commonly used in actual supply chains where transshipments are allowed. The use of POS data has been demonstrated to improve supply chain performance (Croson & Donohue, 2003; Lee et al., 1997a). The results screen showed to each subject the information about past orders, profits, final customer demands, units discarded, unsatisfied demand and units transshipped from both players.

4.4.2. Analyses and results of the BC experiment

4.4.2.1. Analytical estimations

Using the experimental parameters presented in the previous section, the optimal ordering for an isolated Newsvendor (Q_i^*) are 83 and 17 units for a high profit and low profit conditions, respectively. Similarly,

the standard theory of the Newsvendor problem facing transshipments provide estimations for Nash equilibria that let estimate the potential biases on subjects ordering decisions in each treatment. The Nash equilibria ordering decision (Nash) are 72 and 28 units for a high profit and low profit conditions, respectively. Figure 4.4 presents a summary of the main theoretical estimations. It shows that under both high and low profit conditions, the Nash equilibria are located somewhere between the optimal isolated Newsvendor Orders Q_i^* (83 and 17 units) and the mean final customer demand (50 units). This result is consistent with Proposition 1 in Dong and Rudi (2004), and the rationale behind this result is that transshipments allow a better match between supply and demand, which decreases the need of over- or under-ordering, which moves the Newsvendor orders towards the mean.

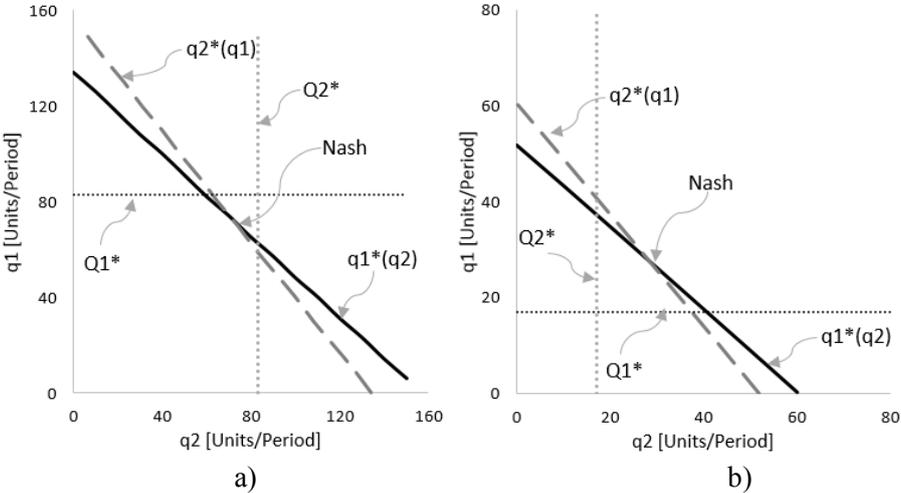


Figure 4.4. Optimal isolated Newsvendor orders, best response function and Nash equilibria for
a) High Profit, and b) Low Profit conditions

4.4.2.2. Behavioral analyses

Table 4.1 presents a summary of different performance measures used to evaluate subject’s decisions. First, as a means of comparison against the theoretical benchmarks (Nash equilibria), I compute the average orders during the 30 experimental periods for all subjects in each experimental treatment. Results in both the high (T1) and low (T2) profit conditions show that subjects make ordering decisions that are distant to the Nash equilibria. In the high profit condition, subjects order significantly less than estimated Nash equilibrium (p -value<.01), and in the low profit condition, subjects order more than the estimated Nash equilibrium (p -value<.01).

In addition, in order to define a unit of measure that defines the deviations of the decisions made by both Newsvendors in a system, I compute the Euclidean distance of subjects' decisions to the Nash equilibrium (Euclidean distance $= ED_t = \sqrt{(q_1^*(q_2^*) - q_{1t})^2 + (q_2^*(q_1^*) - q_{2t})^2}$). Table 4.1 shows that the average Euclidean distance is significantly higher than zero in both treatments, meaning not only that subjects made decisions that are far from optimum (on directions toward the mean final customer demand) but also that they do not coordinate among themselves trying to generate average orders that are closer to the Nash equilibria. These results evidence the presence of the pull-to-center behavior previously encountered in traditional Newsvendor experiments (Bolton & Katok, 2008; Bostian et al., 2008; Schweitzer & Cachon, 2000). Moreover, in this system – allowing transshipments among multiple Newsvendors –, the Euclidean distances do not show difference in performance between the high and low profit condition (p -value = 0.92, t = .1063). This means that there is a symmetric deviation towards the mean for high and low profit conditions.

Table 4.1. Performance measures for treatments T1 and T2

	T1	T2
	High Profit	Low Profit
Av. Orders	54.34 (0.56) ^a	46.02 (0.58)
Av. Euclidean Distance (ED)	31.21 (0.50)	31.13 (0.57)
Av. Transshipments (T)	0.62 (0.09)	0.47 (0.07)
Av. Profit – Subjects (P)	353.17 (6.77)	-50.90 (6.05)
Av. Units Discarded (UD)	15.80 (0.64)	11.87 (0.55)
Av. Customer Satisfaction (CS)	0.88 (0.01)	0.82 (0.01)

Standard Errors in parentheses.

Figure 4.5 shows the average decisions (blue area) for each subject and the Nash equilibria (point where the two best response lines cross each other) for the high and low profit conditions. The blue area shows that in general the average orders placed by the retailers are closer to the mean demand (red dot) than to the Nash equilibrium.

Table 4.1 also provides descriptive information about the average profits and number of transshipments in each specific treatment. In addition, measures of the average units discarded and customer satisfaction (measured as the fraction of final customers satisfied) are provided in Table 4.1.

As expected, the average units discarded and the average customer satisfaction are higher in a high profit condition (T1). This happens because subjects' average orders in T1 (high profit condition) are higher than the mean (expected) demand. Therefore, it is more likely that subjects would be able to satisfy more customers at end of each experimental period.

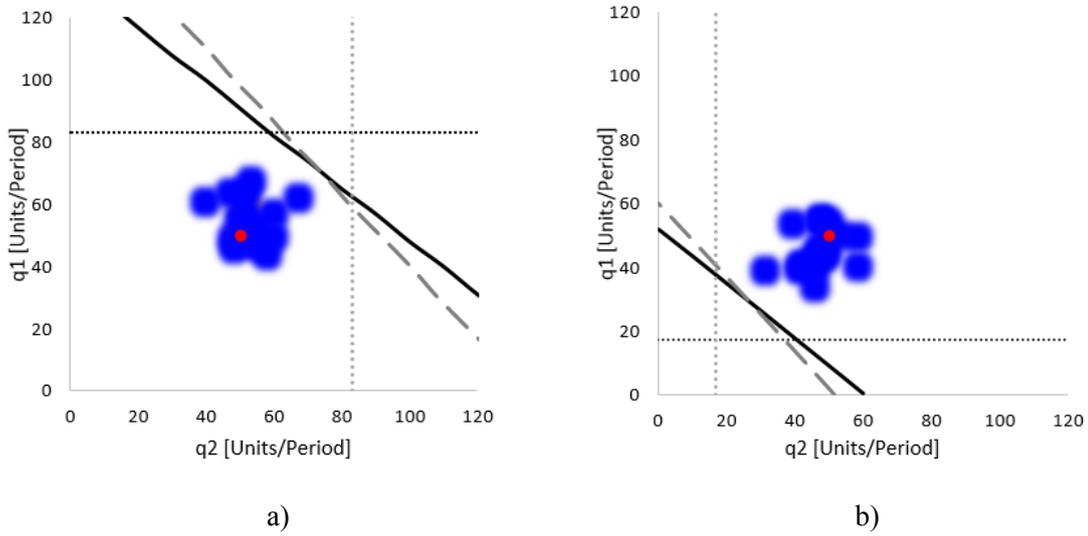


Figure 4.5. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) for a) T1, and b) T2

Finally, I conduct an econometric analysis using a panel data approach to evaluate how orders evolve over time. I test for the presence of demand chasing behavior and the effect of last period's overage and underage on subjects' decisions. I use a parsimonious model expressing the relation between subjects' decisions (q) and a set of predictors: lagged demand ($D_{i,t-1}$), lagged overage ($OV_{i,t-1}$) and lagged underage ($Un_{i,t-1}$) amounts (Ovchinnikov et al., 2015):

$$q_{it} = \beta_{0i} + \beta_1 \cdot D_{i,t-1} + \beta_2 \cdot OV_{i,t-1} + \beta_3 \cdot Un_{i,t-1} + \varepsilon_{it} \quad (4.6)$$

Where ε_{it} is the error term, and i and t are indexes for subjects and time, respectively. Analyses to the orders' time series show they are stationary (Phillips-Peron test) and there is no evidence of autoregressive processes (based on the analysis of ACF and PACF graphs), therefore no lagged dependent variable is required in the model. Moreover, given the random assignment used in experiments, there is no expectation to have time-invariant omitted variables between subjects and, thus, fixed effects are not necessary. Consequently, and to allow for variation between subjects, I use a random effects intercept for the model. Results are shown in Table 4.2. Estimations of parameter β_1 are

positive and significantly higher than zero in both T1 (p -value $<.01$, $t= 3.49$) and T2 (p -value $<.01$, $t= 3.07$), which reflects a clear evidence of demand chasing behavior. This means that subjects anchor their decisions in period t based on the observed demand in period $t-1$. In addition, the magnitude of β_2 and β_3 are similar and significant (p -value $<.05$ in all cases) but their direction is opposite. In this case, the positive sign on β_2 means that the more units were discarded in period $t-1$, the more units subjects will order in period t , probably expecting a higher demand. Similarly, the negative sign on β_3 means that the more customers were unsatisfied in period $t-1$, the less units subjects will order in period t .

The features of the final customer demand distribution could explain this behavior. When subjects perceive a low (high) final customer demand, leading to overage (underage) at the end of the period, subjects can –erroneously- assume that next period the demand will move on the opposite direction and they decide to increase (decrease) their orders. These results are consistent with previous behavioral studies on Newsvendor problems (Bolton & Katok, 2008; Bostian et al., 2008; Ovchinnikov et al., 2015).

Table 4.2. Panel Data estimations for the BC treatments

	T1 High Profit	T2 Low Profit
Fixed part		
β_0 (Intercept)	42.51† (2.43)	34.75† (2.18)
β_1 (D_{t-1})	0.24† (0.04)	0.26† (0.04)
β_2 (OV_{t-1})	0.11† (0.04)	0.10† (0.04)
β_3 (Un_{t-1})	-0.16† (0.05)	-0.18† (0.05)
Stochastic part		
Std. Dev. Intercept	5.26	5.11
Std. Dev. Error	15.85	16.50
AIC	8098	8171

Standard Errors in parentheses; † p -value $< .01$.

Based on the results found for these BC treatments, I argue that in the Human vs. Human setting, the systematic biases observed in a typical Newsvendor problem would persist in a Newsvendor problem with transshipments, both in the high- and low-profit conditions (T1 and T2).

4.4.2.3. Development of Behavioral Models

As in any typical Newsvendor setting, subjects in the experiments with transshipments do not make ordering decisions following the theoretical model. Therefore, the objective in this section is to use different behavioral models that have been used in recent behavioral operations studies to explain the deviation observed in subjects' decisions. I consider three parsimonious behavioral models that allow us to explain some features hidden in the data: (i) anchoring toward the mean (e.g., Benzion et al., 2008; Bostian et al., 2008), (ii) loss aversion from leftover inventory (e.g., Becker-Peth et al., 2013; Chen & Davis, 2014; Davis, 2015), and (iii) psychological disutility (Ho et al., 2010). Given the nature of the behavioral models, I assume participants are fully rational and make ordering decisions following their individual rather than economical preferences. To adapt each of the behavioral models to the particular system, I build on the formal theoretical solution presented in equation (4.5) and then apply algebra to determine the specific model structure.

Anchoring toward the mean: In this model, I set α as the degree of subjects' mean demand anchoring. The model allocates a weight of α ($0 \leq \alpha \leq 1$) to the expected final customer demand μ and a weight of $(1 - \alpha)$ to the theoretical solution. Therefore, the closer the value of α to zero, the lower the anchoring toward the mean.

$$q_1^* = (1 - \alpha) * F_1^{-1} \left(\frac{r - c + (t + b - r) * F_{12}(q_1^* + q_2^*)}{t} \right) + \alpha \mu \quad (4.7)$$

Loss aversion: In this model, β (≥ 1) is defined as the degree of subjects' loss aversion towards leftover inventory. Building on the behavioral model proposed by Becker-Peth et al. (2013), it is possible to redefine equation (4.2) in terms of the loss aversion parameter as:

$$F(Q^*) = \frac{C_u}{C_u + C_o + p * (\beta - 1)} \quad (4.8)$$

Then, using a similar process as the one explained in section 3.2, it is possible to arrive to a loss aversion model for a Newsvendor problem with transshipments. In this case, the closer the value of β to 1, the lower the loss aversion.

$$q_1^* = F_1^{-1} \left(\frac{r - c + (t + b - r) * F_{12}(q_1^* + q_2^*)}{t + p * (\beta - 1)} \right) \quad (4.9)$$

Physiological disutility: In this model, $\delta_o (\geq 0)$ and $\delta_u (\geq 0)$ represent the psychological per-unit cost of over- and under-ordering, respectively. The behavioral model for an isolated Newsvendor accounting for these psychological costs is defined as (Ho et al., 2010):

$$F(Q^*) = \frac{C_u + \delta_u}{C_u + C_o + \delta_o + \delta_u} \quad (4.10)$$

Equation (4.11) presents the transformation of a Newsvendor problem with transshipments to a behavioral model. If $\delta_o > \delta_u$ (or $\delta_u > \delta_o$), it would evidence the asymmetric effect of over- and under-ordering on subjects decisions under each profit condition.

$$q_1^* = F_1^{-1} \left(\frac{r - c + \delta_u + (t + b - r) * F_{12}(q_1^* + q_2^*)}{t + \delta_o + \delta_u} \right) \quad (4.11)$$

4.4.2.4. Estimation of the Behavioral Models

I used the data of both experimental treatments (T1 and T2) to estimate the parameters of each behavioral model because the behavioral parameters are specified to be common across profitability conditions (Ho et al., 2010). Then, I can structurally estimate the behavioral parameters following a maximum likelihood estimation (MLE) approach (Becker-Peth et al., 2013; Ho et al., 2010; Olivares et al., 2008). It is assumed that the order quantities placed by participants have a mean q_1^* (defined in each structural model) and subjects exhibit errors in their decision that are normally distributed with mean zero and standard deviations τ_{HP} and τ_{LP} for Treatment 1 (HP) and Treatment 2 (LP), respectively. Equation (4.12) shows a representation of the likelihood function used to estimate the behavioral models.

$$L(q_1^* | c_i, t, r, b, \mathbf{p}, \tau_{HP}, \tau_{LP}) = \prod_{i=1}^n g(q_i; c_i, t, r, b, \mathbf{p}, \tau_{HP}, \tau_{LP}) \quad (4.12)$$

Where $g(\cdot)$ denotes the probability density function for the order quantity q_i made by an individual i given the treatment parameters $c_i, t, r, b, \mathbf{p}, \tau_{HP}, \tau_{LP}$, where \mathbf{p} represents the set of behavioral parameters in each behavioral model: α, β or δ_o and δ_u (Davis, 2015).

Table 4.3 presents the estimations for the parameters of each behavioral model. In the first behavioral model, the parameter α is significantly different from zero (p -value=.000) and with an estimated value close to one. Despite the fact that the Nash equilibrium in both treatments is located in

a place between the optimal isolated Newsvendor order and the mean demand, subjects still display a strong tendency to anchor their decisions toward the mean demand and place their order far from the Nash equilibrium.

Table 4.3. Structural estimation of Behavioral Models

Estimators	Anchoring	Loss Aversion	Physiological disutility
α	.87† (0.01)		
β		1.026† (0.00)	
δ_o			3.08† (0.36)
δ_u			3.22† (.37)
τ_{HP}	17.29† (0.55)	17.32† (0.96)	17.26† (0.39)
τ_{LP}	17.74† (0.39)	38.97† (0.88)	17.73† (0.55)
-LL	8496.0	9277.7	8493.1

Standard Errors in parentheses; † p-value < .01.

For the loss aversion model, β parameter is significantly higher than 1 (p -value=.000), which is the expected value for the theoretical benchmark for a person with a neutral preferences towards losses. This means that subjects experience loss aversion from leftovers, even if the strength of the effect is low (β close to 1). This low strength of the loss aversion from leftovers can be attributed to the presence of transshipments in the system. This is, given that a subject could transship (if needed) some of the leftover units to the other Newsvendor at the end of each simulated period, subjects become less afraid of leftovers and they order more than what they would have ordered in a system without transshipments.

For the psychological disutility model, results show that $\delta_u > \delta_o > 0$. In each period, every unit that subjects could have sold but were not available to supply to the final customer leads to a psychological cost of 3.22 points (p -value=.000). Similarly, every unit bought and not sold in each period brings a psychological cost of 3.08 (p -value=.000). This means that subjects on average prefer to have extra units at the end of each period than having unsatisfied final customer demand. This result shows a switch in the direction of the psychological disutility estimations compared with those obtained in previous studies where the psychological per-unit cost of over-ordering was lower than the under-

ordering cost (Ho et al., 2010). This finding about the switch in the importance of the psychological cost is consistent with the results explained in the loss aversion model. Subjects prefer to place higher orders in case they face a high final customer demand, but if the final customer demand is lower than their orders, decision makers still expect to transship some units to the other Newsvendor, incurring in zero cost for those transshipped units. Consequently, subjects allocate a higher psychological cost to under-ordering.

4.5. Practical improvements to the system

Adding to the explorative nature of this chapter, this section evaluates two practical strategies that could improve decision-makers' profits in a Newsvendor problem with transshipments. These strategies consist of: (i) face-to-face communication (Ahn et al., 2011), and (ii) different behavioral best response heuristics (Becker-Peth et al., 2013; Ovchinnikov et al., 2015). In the next subsections, I describe the experimental treatments and results of these two practical strategies.

4.5.1. Communication Experiment (C): Human vs. Human

Although access to POS data improves supply chain performance, this full information availability does not guarantee cooperation among agents. I build on the ideas of Cardenas (2000) and Ostrom (1998) and evaluate the effect of face-to-face communication as a mechanism that creates a trusty environment for reducing non-cooperative behaviors (Ahn et al., 2011; Castillo & Saysel, 2005). Therefore, I expect that the inclusion of face-to-face communication to the BC treatments would reduce the systematic biases discussed in the previous section.

4.5.1.1. Design of the C experiment

This experiment builds on the results obtained from the BC experiment. I consider a full experimental design, with four experimental treatments – two different profit conditions (High and Low) and two different kind of communication conditions (No Communication and Face-to-Face Communication). Table 4.4 specifies all treatments conducted in this C experiment and the number of subjects used in each treatment.

For the No Communication treatments, I use the data and analyses made in the BC experiments: T1 and T2. For the Face-to-Face communication treatments (T3 and T4), I ran a new set of experiments using

the same set of parameters used in the BC experiments, which guarantees that the estimation of the first best strategies for the high and low profit conditions remain unchanged.

Table 4.4. Experimental treatments for the C experiment

	Profit Conditions	
	High	Low
No Communication	T1 (<i>n</i> =32)	T2 (<i>n</i> =32)
Face-to-Face Communication	T3 (<i>n</i> =28)	T4 (<i>n</i> =26)

In addition, I followed an identical experimental protocol and briefing process as the used in the BC experiment. However, to facilitate the communication between participants in the same experimental market, a small modification in the experimental process was implemented. Upon arrival to the experimental session, subjects participating in the same market were sat next to each other to allow face-to-face communication. During the experiment, before registering the ordering decisions in the computer, subjects were allowed to talk with the other Newsvendor in the market for a minute. This minute of communication provides an opportunity for discussing or creating transshipment strategies that could work for the benefit of the whole system. Again, the experiment ran during 30 independent periods and subjects had full information about the system outputs (demand, orders placed, profits, demand unsatisfied, units discarded, etc.) at the end of each period.

4.5.1.2. Analyses and results of the C experiment

Table 4.5 presents a summary of the different performance metrics used to evaluate subjects' biases in treatments T3 and T4. Despite subjects in the same experimental market were allowed to communicate with each other during the experiment, results in both T3 and T4 show that subjects still make ordering decisions that are distant to the Nash equilibria. Both the average orders and the Euclidean distance show that the pull-to-center behavior remains. Now, focusing on the effect of communication on subjects' performance, for the high profit treatments (T1 vs. T3) results show that there are not significant differences in the average orders ($\mu_{q,T1} - \mu_{q,T3} = .06, t = -0.711, p = .48$) and average

Euclidean distance ($\mu_{ED,T1} - \mu_{ED,T3} = 3.44, t = -.887, p = .38$). However, there is an increase in the units transshipped ($\mu_{T,T1} - \mu_{T,T3} = -.26, t = -1.75, p = .07$) keeping similar customer satisfaction ($\mu_{CS,T1} - \mu_{CS,T3} = .00, t = 0.11, p = .91$), units discarded ($\mu_{UD,T1} - \mu_{UD,T3} = -.77, t = -.79, p = .43$) and average profits ($\mu_{P,T1} - \mu_{P,T3} = 3.44, t = .35, p = .73$). Therefore, it seems that during the communication part of the experiment, subjects were able to agree on a strategy where one of them places a higher order than the other (higher order variability), so that they can increase the number of transshipments but keeping similar average orders.

Table 4.5. Performance measures for treatments T3 and T4

	T3	T4
	Communication - HP	Communication - LP
Av. Orders	54.94 (0.63)	41.83 (0.62)
Av. Euclidean Distance (ED)	31.86 (0.53)	21.79 (0.53)
Av. Transshipments (T)	0.88 (0.12)	0.69 (0.11)
Av. Profit – Subjects (P)	349.73 (7.20)	-36.18 (6.19)
Av. Units Discarded (UD)	16.57 (0.72)	9.93 (0.57)
Av. Customer Satisfaction (CS)	0.88 (0.01)	0.78 (0.01)

Standard Errors in parentheses.

For the low profit treatments (T2 vs. T4) results show a significant reduction both in the number of average orders ($\mu_{q,T2} - \mu_{q,T4} = 4.19, t = 4.93, p < .01$) placed by the subjects during the experiments and in the Euclidean distance ($\mu_{ED,T2} - \mu_{ED,T4} = 3.57, t = 4.91, p < .01$) to the Nash equilibrium. Subjects in the low profit condition also used communication as a way to improve coordination and increase the number of units transshipped ($\mu_{T,T2} - \mu_{T,T4} = .22, t = -1.69, p = .09$). Contrary to the high profit condition, subjects in T4 were able to increase their average profit ($\mu_{P,T2} - \mu_{P,T4} = -14.72, t = -1.70, p = .09$) even if they had to forgo a small share of customer satisfaction

$(\mu_{CS,T2} - \mu_{CS,T4} = .04, t = 2.75, p = .01)$. Thus, for the low profit condition, the use of communication has a positive effect on subjects' performance, reducing the pull-to-center behavior.

Figure 4.6 shows the distribution of the average orders for each pair of subjects and the Nash equilibria for the high and low profit conditions allowing Face-to-Face communication. The blue area shows that in general the average orders start to move closer the Nash equilibria compared to the results observed in Figure 4.5. In fact, for the low profit condition, Figure 4.6 shows that in some experimental markets, subjects were placing average decisions that are really close to the Nash equilibrium.

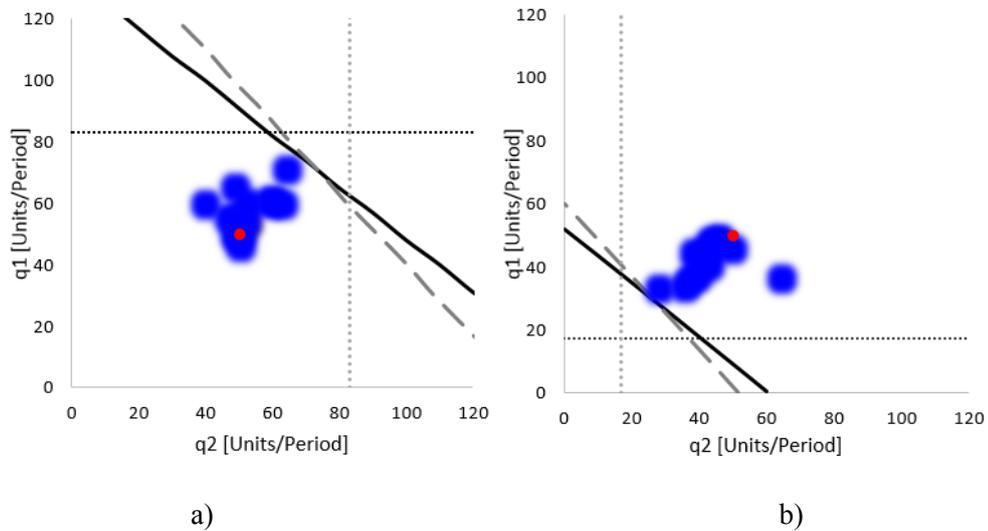


Figure 4.6. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) for a) T3, and b) T4

Finally, I test for the presence of demand chasing behavior and the effect of last period's overage and underage on subjects' decisions when they were allowed to communicate. I use the same parsimonious model presented in equation (4.6) and Table 4.6 presents the estimations. For the high profit treatment (T3), neither of the β parameter of interest is significant (p -value $> .1$ in all cases). Therefore, the use of communication allowed subjects to make orders in period t that are not anchored towards the outputs observed in period $t-1$. For the low profit treatment (T4), communication among subjects allowed a significant reduction on the anchoring towards the last period outputs ($\beta_{2,T2} - \beta_{2,T4} = .11, t = 1.81, p = .04$); however, in this case, subjects still present (weakly) the biases observed in a traditional Newsvendor setting. These results are also consistent with previous studies, where there is an asymmetric demand chasing behavior (Bolton & Katok, 2008).

Based on the results shown for these C treatments, I argue that including communication in a Newsvendor setting with transshipments reduces the systematic biases observed in a typical Newsvendor problem, especially for a low profit condition. Communication creates a trusty environment that increases the cooperative behavior of subjects, which is aligned with previous collective action and communication studies (Ahn et al., 2011; Cardenas, 2000; Castillo & Saysel, 2005; Ostrom, 1998).

Table 4.6. Panel Data estimations for the C treatments

	T3 High Profit	T4 Low Profit
Fixed part		
β_0 (Intercept)	55.56 [†] (2.85)	36.00 [†] (2.44)
β_1 (D_{t-1})	-0.01 (0.05)	0.15 [†] (0.05)
β_2 (Ov_{t-1})	-0.05 (0.05)	0.05 (0.05)
β_3 (Un_{t-1})	0.07 (0.05)	-0.12 [†] (0.05)
Stochastic part		
Std. Dev. Intercept	6.49	5.11
Std. Dev. Error	17.18	16.50
AIC	7227	8171

Standard Errors in parentheses; † p-value < .01.

4.5.2. Best Response Experiments (BR): Human vs. Computer

In the previous experimental conditions, I was analyzing Human vs. Human experiments, trying to understand interactions and behavioral features of subjects' decisions. In this part of the chapter, I build on the work done by Ovchinnikov et al. (2015), the parsimonious model proposed in equation (4.6) and the analytical model presented in equation (4.5) to create decision rules that can be used by a knowledgeable Newsvendor (in this case the computer) to estimate a best response order quantity. Hence, in this case I am proposing a Human (behavioral Newsvendor) vs. Computer (knowledgeable Newsvendor) experiments. I design different experimental treatments to analyze different strategies followed by the knowledgeable Newsvendor. These strategies are:

Nash-Equilibrium Response: In this case, independently of the behavioral Newsvendor decisions, the knowledgeable Newsvendor will always order his individual first best order quantity.

Static Best Response: In this case, the experiment will be divided in two parts. In the first part (periods 1 to 30), the knowledgeable Newsvendor will use average order quantities obtained in the BC experiment together with the analytical model (equation (4.5)) to determine a global but constant ordering strategy (General Static Best Response - GSBR). Regardless of the behavioral Newsvendor decisions, during this first part of the experiment the knowledgeable Newsvendor will always place this GSBR order. In the second part (periods 31 to 60), the knowledgeable Newsvendor will compute the average orders placed by the behavioral Newsvendor during first part of the experiment (Human vs. GSBR) and, using equation (4.5), will determine a particular strategy for each behavioral Newsvendor (Particular Static Best Response - PSBR). Independently of the behavioral Newsvendor decisions, the knowledgeable Newsvendor will always place this new order quantity in the second part of the experiment. The computer automatically runs the transition from the first part to the second part; therefore, subjects do not perceive changes or delays during the experiment.

Dynamic Best Response: As in the Static Best Response case, the experiment is divided in two parts. In the first part (period 1 to 30), the knowledgeable Newsvendor builds on the results obtained from the BC experiments and takes advantage of the estimations obtained for the dynamic model (equation (4.6), and Table 4.2) to predict a generic dynamic Newsvendor behavior. After determining the behavioral Newsvendor behavior, the knowledgeable Newsvendor uses the analytical model (equation (4.5)) to determine a global but dynamic best response strategy (General Dynamic Best Response - GDBR) that can be used to estimate its decisions each period. These dynamic decisions are updated period by period based on the evolution of the main variables of the system (including the behavioral Newsvendor's decisions). In the second part (periods 31 to 60), the knowledgeable Newsvendor uses the information from the first part of the experiment (Human vs. GDBR) to estimate the parameters of the dynamic model (equation (4.6)) for each behavioral Newsvendor. Then, using the analytical model (equation (4.5)) the knowledgeable Newsvendor determines a particular response strategy for each behavioral Newsvendor (Particular Dynamic Best Response - PDBR). This particular

strategy is computed period by period based on the evolution of the main variables in the dynamic model.

4.5.2.1. Design of the BR experiment

As in the C experiment, this design builds on the results obtained from the BC experiment. I consider a full experimental design, with eight experimental treatments – two different profit conditions (High and Low), the Human vs. Human interaction explained in the BC experiment and three different knowledgeable Newsvendor strategies (Nash-Equilibrium, Static Best Response and Dynamic Best response). Table 4.7 specifies all treatments conducted in this BR experiment and the number of subjects who participated in each treatment.

Table 4.7. Experimental treatments for the BR experiment

	Profit Conditions	
	High	Low
BC – Human vs. Human	T1 (<i>n</i> =32)	T2 (<i>n</i> =32)
Nash-Equilibrium	T5 (<i>n</i> =30)	T6 (<i>n</i> =30)
Static Best Response	T7 (<i>n</i> =24)	T8 (<i>n</i> =16)
Dynamic Best Response	T9 (<i>n</i> =16)	T10 (<i>n</i> =16)

For the experiments including any strategy of the knowledgeable Newsvendor, I use the same set of parameters, experimental protocol and briefing process as the one used in the BC experiment. However, in this case subjects in the experiments were not paired with another subject in the room, but they played individually in a market where the computer makes the decisions for the second subject (knowledgeable Newsvendor). In each treatment, the knowledgeable Newsvendor used a management science strategy to predict systematic regularities in the orders of the behavioral partner and built the corresponding decision model strategy (based on each treatment condition) to exploit the predictable irrationality of the subjects. The subjects in the experiment were thoroughly informed about the experimental setting. As in the BC experiment, subjects had full information about the system outputs (demand, orders placed, profits, demand unsatisfied, units discarded, etc.) at the end of each period.

4.5.2.1. Analyses and results of the BR experiment

Table 4.8 and Table 4.9 present a summary of the different performance metrics used to evaluate subjects' biases in treatments T5 to T10. In treatments T7 to T10, I separate the estimation in two parts (period 1 to 30 and period 31 to 60) to disentangle the effect of the General and Particular Best responses. Appendix 4.4 provides all the p -values from the t-tests comparing the performance measures among the different experimental treatments.

High profit treatments: Table 4.8 shows that subjects' average orders are significantly higher (closer to the Nash equilibrium) in T5, T7 and T9 than in T1 ($\mu_{q,T1} - \mu_{q,(T5,T7,T9)} < 0$, p -value $< .05$ in all cases). This means that assuming any of the proposed strategies for the knowledgeable Newsvendor makes the behavioral Newsvendor move his orders toward the Nash equilibrium. This result is also supported by the average Euclidean distance, where the distances are significantly closer to the Nash equilibrium in T5, T7 and T9 than in T1 ($\mu_{ED,T1} - \mu_{ED,(T5,T7,T9)} > 0$, p -value $< .05$ in all cases), especially on the best response treatments (T7 and T9) where the knowledgeable Newsvendor assumes an active position in the experiment. Although subjects in the Particular Best Response treatments (PSBR and PDBR) seem to place higher average orders than subjects in the General Best Response treatments (GSBR and GDBR), these differences are not significant ($\mu_{q,T7-1} - \mu_{q,T7-2} = -1.74$, $t = -1.14$, $p = .254$; $\mu_{q,T9-1} - \mu_{q,T9-2} = -.62$, $t = -.79$, $p = .43$).

For these treatments, where the knowledgeable Newsvendor plays a role in the experiment, I make distinctions between the two types of actors (subjects and computer) in some of the performance measures. The average Transshipments In (Out) refers to the average number of transshipped units that the behavioral Newsvendor received (sent) from (to) the knowledgeable Newsvendor during the experiment. Results show that the number of transshipped units received by the behavioral Newsvendor (Transshipments In - T_i) are significantly higher in T5, T7 and T9 than in T1 ($\mu_{T,T1} - \mu_{T_i,(T5,T7,T9)} < 0$, p -value $< .05$ in all cases) and the units sent to the knowledgeable Newsvendor (Transshipments Out - T_o) are significantly lower ($\mu_{T,T1} - \mu_{T_o,(T5,T7,T9)} > 0$, p -value $< .05$ in all cases). These results reflect in part the strategies followed by the knowledgeable Newsvendor, who usually takes advantage of the pull-to-center behavior presented by subjects and place higher ordering decisions. Placing higher orders decreases the probability of receiving transshipped units, while raising the probability of having extra

units at the end of the period, which can be sent to the behavioral Newsvendor to increase the final customer demand satisfied.

In addition, using any of the proposed strategies leads to higher profits both for the behavioral Newsvendor, who is able to satisfy more of his demand due to the units received from the knowledgeable Newsvendor, and for the knowledgeable Newsvendor, who takes advantage of subjects' biases for increasing his own profits selling more units. Results indicate that in all cases the knowledgeable Newsvendor gets a higher profit than the behavioral Newsvendor ($\mu_{P,(T5,T7,T9)} - \mu_{Pc,(T5,T7,T9)} < 0$, p -value $< .05$ in all cases) and that the average subjects' profits in T7-2 and T9 are higher than that obtained in the Human vs. Human interaction in T1 ($\mu_{P,T1} - \mu_{P,(T7-2,T9)} < 0$, p -value $< .05$ in all cases). Finally, given the increase in the overall orders, the units discarded and the customer satisfaction also increase.

Table 4.8. Performance measures for treatments T5, T7 and T9 – High Profit Treatments

	T5	T7-1	T7-2	T9-1	T9-2
	Nash	GSBR	PSBR	GDBR	PDBR
Av. Orders - Subjects	59.42 (0.55)	57.50 (0.82)	59.24 (0.88)	57.4 (0.87)	58.02 (0.89)
Av. Euclidean Distance (ED)	16.91 (0.50)	28.19 (0.45)	27.02 (0.54)	26.12 (0.51)	25.56 (0.51)
Av. Transshipments In (Ti)	1.24 (0.11)	2.47 (0.22)	4.15 (0.29)	2.66 (0.26)	3.44 (0.29)
Av. Transshipments Out (To)	0.04 (0.02)	0.02 (0.01)	0.11 (0.04)	0.08 (0.04)	0.01 (0.01)
Av. Profit – Subjects (P)	363.73 (07.25)	369.60 (9.47)	377.92 (8.97)	385.60 (11.56)	386.89 (10.67)
Av. Profit – Computer (Pc)	399.64 (8.56)	424.56 (12.96)	398.39 (10.97)	435.33 (15.41)	408.48 (13.88)
Av. Units Discarded (UD)	20.11 (0.70)	18.96 (0.89)	20.89 (0.89)	17.62 (1.00)	18.68 (0.99)
Av. Customer Satisfaction (CS)	0.92 (0.01)	0.89 (.01)	0.94 (.01)	0.90 (.01)	0.94 (0.01)

Standard Errors in parentheses.

Comparing the effect of each of the different knowledgeable Newsvendor strategies on the system performance, the Nash equilibrium strategy is the one that provides the lowest average total profit

$(\mu_{P,T5} + \mu_{PC,T5} = 763.37)$. It seems that sticking to the first best strategy under a high profit condition is a passive strategy that does not exploit the behavioral aspects of the subject to improve the overall supply chain profit. On the other hand, the strategies that provide higher profits for the supply chain are the GDBR $(\mu_{P,9-1} + \mu_{PC,T9-1} = 820.93)$ and PDBR $(\mu_{P,9-1} + \mu_{PC,T9-1} = 795.37)$. In these two treatments, the knowledgeable Newsvendor has a more active role in the market and updates frequently his strategies based on the information received from the previous period. However, it seems that creating a general dynamic best response policy (GDBR) would be the best option for increasing the average individual and overall profits, while dealing with the variability from the behavioral Newsvendor orders and the uncertain final customer demand.

Low profit treatments: Table 4.9 shows the performance measures for treatments T6, T8 and T10. In contrast to the high profit treatments, when the knowledgeable Newsvendor uses the Nash-equilibrium or the GDBS strategies, subjects' average orders do not have a significant change in magnitude compared with the subjects' orders in T2 $(\mu_{q,T2} - \mu_{q,T6,T8-1}) > 0, p\text{-value} > .1$ in both cases). On the other hand, subjects' average orders are significantly lower (closer to the Nash equilibrium) in T8-2 and T10 than in T2 $(\mu_{q,T2} - \mu_{q,(T8-2,T10)}) > 0, p\text{-value} < .05$ in all cases). However, the average Euclidean distance is closer to the Nash equilibrium in T6, T8 and T10 than in T1 $(\mu_{ED,T2} - \mu_{ED,(T6,T8,T10)}) > 0, p\text{-value} < .05$ in all cases), this is in part due to the strategic decisions made by knowledgeable Newsvendor in the experiment. In addition, subjects in the Particular Best Response treatments (PSBR and PDBR) place lower average orders than subjects in the General Best Response treatments (GSBR and GDBR) $(\mu_{q,T8-1} - \mu_{q,T8-2} = 4.39, t = 4.15, p\text{-value} < .01; \mu_{q,T10-1} - \mu_{q,T10-2} = 5.12, t = 4.54, p\text{-value} < .01)$. This can be explained as a simple learning process from one stage of the experiment to the other (Benzion et al., 2008; Bostian et al., 2008).

The analysis of the units transshipped is opposite to the high profit treatments. The number of transshipped units received by the behavioral Newsvendor (Transshipments In - Ti) are significantly lower in T6, T8 and T10 than in T2 $(\mu_{T,T2} - \mu_{Ti,(T6,T8,T10)}) > 0, p\text{-value} < .05$ in all cases) and the units sent to the knowledgeable Newsvendor (Transshipments Out - To) are significantly higher $(\mu_{T,T2} - \mu_{To,(T6,T8,T10)}) < 0, p\text{-value} < .05$ in all cases). In this low profit condition, the knowledgeable Newsvendor

takes advantage of the pull-to-center behavior presented in the subjects and place lower ordering decisions. Placing lower orders reduces the probability of having economical loses at the end of each period, and given that the behavioral Newsvendor will place higher orders, the likelihood of receiving transshipped units – in the case that at the end of the period part of the demand remains unsatisfied – will increase. Similarly, placing lower orders decreases the probability of sending transshipped units to the other Newsvendor.

In addition, when the knowledgeable Newsvendor uses any of the proposed dynamics or static strategies, he gets higher average profits than subjects do in treatment T2 ($\mu_{P,T2} - \mu_{Pc,(T8,T10)} < 0$, p -value $< .05$ in all cases). In these strategies, the knowledgeable Newsvendor takes advantage of subjects' over-ordering for decreasing his orders and reducing the probability of facing economical loses, while increasing his profits at the end of each period by selling the units received from the leftover units ordered by the behavioral Newsvendor. Moreover, the behavioral Newsvendor increases his profits in treatments where the knowledgeable Newsvendor uses strategies for predicting subjects' behavior (PSBR, GSBR, GDBR and PDBR). In these cases, given that the knowledgeable Newsvendor orders fewer units, the behavioral Newsvendor is able to get rid of leftover units that would have led to economic losses otherwise. In addition, there is a significant improvement in the behavioral Newsvendor profit from general to particular best response treatments. This improvement is due to both a learning process and to the improvements given by the particular response provided by the knowledgeable Newsvendor. Finally, given the decrease in overall orders, the units discarded and the customer satisfaction also decrease.

As in the high profit condition, the Nash equilibrium strategy is a passive strategy that provides low benefits to the total profit for the supply chain ($\mu_{P,T6} + \mu_{Pc,T6} = -21.95$). However, the particular strategies (PSBR and PDBR) are the ones providing higher profits for the supply chain ($\mu_{P,T8-2} + \mu_{Pc,T8-2} = 20.97$; $\mu_{P,T10-2} + \mu_{Pc,T10-2} = 35.43$) and especially for the knowledgeable Newsvendor ($\mu_{Pc,T8-2} = 22.98$; $\mu_{Pc,T10-2} = 24.58$). In these two treatments, the knowledgeable Newsvendor takes advantage of the understanding of each subject's behavior to create proper strategies that maximize the individual and global profit. Therefore, having a more active role in the market and updating every

period the strategies to respond to each Newsvendor is the best option for increasing performance under low profit conditions.

Table 4.9. Performance measures for treatments T6, T8 and T10 - Low Profit Treatments

	T6	T8-1	T8-2	T10-1	T10-2
	Nash	GSBR	PSBR	GDBR	PDBR
Av. Orders - Subjects	45.59 (0.72)	43.08 (0.82)	38.69 (0.81)	45.11 (0.86)	39.99 (0.75)
Av. Euclidean Distance (ED)	20.24 (0.63)	25.82 (0.51)	22.01 (0.53)	26.34 (0.64)	21.55 (0.62)
Av. Transshipments In (Ti)	0.04 (0.02)	0.003 (0.00)	0.05 (0.02)	0.02 (0.01)	0.03 (0.01)
Av. Transshipments Out (To)	1.84 (0.17)	3.23 (0.30)	2.50 (0.24)	3.81 (0.37)	3.64 (0.32)
Av. Profit – Subjects (P)	-44.64 (7.05)	-20.53 (8.64)	-2.01 (6.53)	-14.95 (8.24)	10.85 (5.65)
Av. Profit – Computer (Pc)	22.69 (2.26)	24.02 (0.74)	22.98 (2.14)	20.70 (1.09)	24.58 (1.42)
Av. Units Discarded (UD)	11.01 (0.65)	8.35 (0.77)	6.20 (0.58)	8.13 (0.72)	5.16 (0.49)
Av. Customer Satisfaction (CS)	0.79 (0.01)	0.75 (0.01)	0.74 (0.01)	0.76 (0.01)	0.75 (0.01)

Standard Errors in parentheses.

Figure 4.7 presents the distribution of the average orders for each pair of subjects and the Nash equilibria for the high and low profit conditions for treatments T5 to T10. The blue area shows that in general the average orders remain close to the Nash equilibria. For the Nash-equilibrium treatments, the blue region approaches the Nash equilibria; however, average orders still evidence the presence of the pull-to-center behavior. Given the existence of the pull-to-center behavior, the knowledgeable Newsvendor realizes that it would be beneficial for the whole system if he places orders that are higher (lower) than the Nash equilibrium for the high (low) profit conditions. Notice that in the treatments where the knowledgeable Newsvendor use a Particular strategy for each subject (PSBR and PDBR treatments), average orders fluctuate around the best response curves (close to the Nash equilibria). In fact, for the low profit condition, some systems present average orders that are really close to the Nash equilibrium, which leads to improvements in the supply chain overall profits.

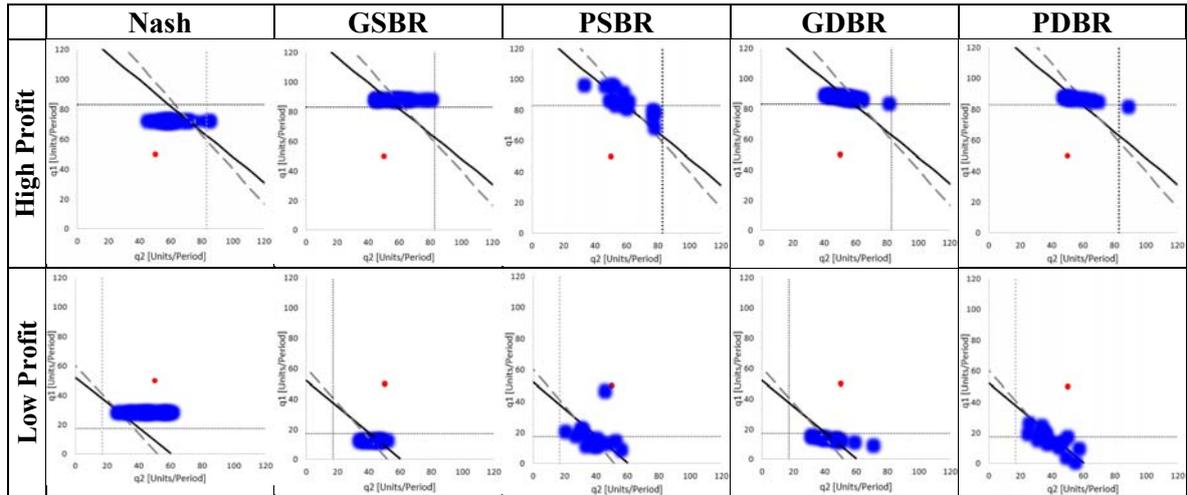


Figure 4.7. Nash equilibrium, mean demand (red dot) and average retailers' orders (blue area) T5-T10

Table 4.10. Panel Data estimations for the BR treatments

	T5	T7-1	T7-2	T9-1	T9-2	
	Nash	GSBR	PSBR	GDBR	PDBR	
Fixed part						
High Profit	β_0 (Intercept)	48.90 [†] (2.85)	35.57 [†] (3.17)	44.88 [†] (3.78)	38.28 [†] (3.76)	38.85 [†] (3.96)
	β_1 (D_{t-1})	0.15 [†] (0.04)	0.39 [†] (0.05)	0.29 [†] (.05)	0.33 [†] (0.05)	0.33 [†] (0.05)
	β_2 (Ov_{t-1})	0.13 [†] (0.04)	0.21 [†] (0.05)	0.01 (0.05)	0.18 [†] (0.06)	0.26 [†] (0.06)
	β_3 (Un_{t-1})	0.06 (0.04)	-0.17 [†] (0.05)	0.03 (0.07)	-0.15 [†] (0.06)	-0.21 [†] (0.08)
Stochastic part						
	Std. Dev. Intercept	7.71	7.31	11.67	6.94	8.12
	Std. Dev. Error	14.73	18.68	19.17	16.53	15.82
	AIC	7908	6109	6373	4102	3930
	T6	T8-1	T8-2	T10-1	T10-2	
	Nash	GSBR	PSBR	GDBR	PDBR	
Fixed part						
Low Profit	β_0 (Intercept)	33.02 [†] (2.60)	23.18 [†] (2.72)	27.05 [†] (3.04)	22.72 [†] (3.12)	25.47 [†] (2.85)
	β_1 (D_{t-1})	0.25 [†] (0.05)	.46 [†] (0.06)	0.29 [†] (0.05)	0.53 [†] (0.05)	0.40 [†] (0.05)
	β_2 (Ov_{t-1})	0.20 [†] (0.04)	0.34 [†] (0.06)	0.10 (0.06)	0.27 [†] (0.06)	0.15 [†] (0.06)
	β_3 (Un_{t-1})	-0.12 [†] (0.05)	-0.35 [†] (0.06)	-0.16 [†] (0.06)	-0.41 [†] (0.06)	-0.32 [†] (0.06)
Stochastic part						
	Std. Dev. Intercept	7.98	1.98	7.75	6.72	7.16
	Std. Dev. Error	19.14	16.37	14.93	14.91	12.68
	AIC	7941	3935	4010	4006	3729

Standard Errors in parentheses; [†] p-value < .01.

Finally, I test for the presence of demand chasing behavior and the effect of last period's overage and underage on subjects' decisions in treatments T5 to T10 using the same parsimonious model presented in equation (4.6). As in the performance measures, in treatments T7 to T10 I disentangle the effect of the General and Particular Best responses by separating the estimation in two parts. Table 4.2 provides the estimates for the base case experiment and Table 4.10 presents the estimates for treatments T5 to T10.

Overall results show no clear effect of the different strategies followed by the knowledgeable Newsvendor on the last-period anchoring biases of the behavioral Newsvendors. In all treatments (T6 to T10), results show that subjects anchor their decisions in period t based on the observed demands, overage and underage units of period $t-1$. Estimations of parameters β_1 , β_2 and β_3 are in the same direction and with similar magnitude (p -value $> .05$ in -almost- all cases) than the estimations obtained in the BC experiment.

4.6. Conclusions

This chapter mainly contributes to the field of Operations Management (OM). Within the OM field, the Newsvendor problem presents the basics for models of inventory management under demand uncertainty (Bolton & Katok, 2008). This chapter responds to recent calls for further experimental analyses considering the interactions among multiple subjects in inventory management problems (Becker-Peth et al., 2013; Ovchinnikov et al., 2015). I created a multi-agent experiment to evaluate different factors and strategies that lead to a reduction of the persistent behavioral biases presented in a typical Newsvendor problem with transshipments.

Initially, I show how Newsvendors, in a transshipment setting, are prone to the common biases (pull-to-center behavior, demand chasing) observed in a typical Newsvendor setting. Additionally, I built up on the basics of previous studies to derive and evaluate different behavioral models (Anchoring toward the mean, Loss aversion and Physiological disutility) that can explain subjects' biases. Estimations show that (i) subjects display a strong tendency to anchor their decisions toward the mean demand and place their order far from the estimated Nash equilibrium, (ii) subjects show loss aversion from leftovers, and (iii) subjects on average prefer to have extra units at the end of each period than having unsatisfied final customer demand.

This improved understanding of the Newsvendor behavior allows me to formulate and evaluate different strategies that may provide managers with useful information to redesign and improve real operations. As a first strategy, I evaluate the effect of Face-to-Face communication as a mechanism that creates a trusty environment for reducing non-cooperative behaviors (Ahn et al., 2011; Cardenas, 2000; Castillo & Saisel, 2005; Ostrom, 1998), which exist even in systems where agents have access to POS data. Inclusion of direct communication, in a Newsvendor problem with transshipments, reduces the systematic biases observed in a typical Newsvendor problem. There are significant reductions both in the pull-to-center behavior and in the anchoring toward outputs from previous periods.

In the last part of the chapter, I used a set of Human vs. Computer experiments. This design takes advantage of the mathematical deduction for the Newsvendor model with transshipments to formulate three different strategies that can be used by a knowledgeable Newsvendor (the computer) to offset the behavioral Newsvendor decisions (subjects in the experiment). The three strategies are: Nash-equilibrium response, Static Best Response and Dynamic Best Response, and both the Static Best Response and Dynamic Best Response strategies can be defined for the overall population (General) or for each single subject (Particular). Findings show, on one hand, that the strategies followed by the knowledgeable Newsvendor take advantage of the behavioral biases existing in subjects' decisions to place orders that increase the average profits for each actor in the supply chain. On the other hand, the behavioral Newsvendor usually responds to the knowledgeable Newsvendor in a way that moves the systems close to the Nash Equilibrium.

In addition, results from these managerial strategies suggest that for a knowledgeable Newsvendor (or any manager in a real supply chain) placing orders equal to the Nash equilibrium as a response to a behavioral Newsvendor (other symmetric manager in a real supply chain) does not bring many benefits for the whole supply chain. In contrast, to achieve higher profits while dealing with the variability of the behavioral Newsvendor orders and the uncertain final customer demand, a knowledgeable Newsvendor should have a more active role in the market. They should create and update every period a best response strategy, that allows them to understand the market and make better decisions. In particular, for a high profit condition, a knowledgeable Newsvendor could use a general

rule (GDBR) to respond to the other behavioral Newsvendors, while for a low profit condition, it would be better to generate a specific response rule (PDBR) for each behavioral Newsvendor.

The main results of this chapter offer practical guidance on how to exploit behaviorally biased Newsvendor orders to improve the overall performance in a single-supplier multi-retailer (Newsvendors) supply chain, where transshipments among Newsvendor are allowed and the final customer demand is uncertain. Therefore, this chapter becomes a building block for behavioral studies in inventory management for a system of the described characteristics.

This work can be extended in many directions. First, in this chapter, I have considered a parsimonious dynamic model to predict subjects' behavior; therefore, improvements in the response can be obtained by using more accurate models, although the main behavioral conclusion would likely remain unchanged. Second, there are different effects of learning that may be useful to understand. For example, how does previous experience in an isolated Newsvendor problem matter to the performance in a Newsvendor problem with transshipments? Third, I have considered in our experiments a system composed by two symmetrical Newsvendors and a supplier with unlimited capacity, therefore, I may ask: What would be the effect on subjects' orders when the system is asymmetric? How would the subjects' decisions change knowing that they may compete for limited capacity? How different types of contracts between supplier and Newsvendor may increase overall performance? Fourth, in our experiments, there are no backorders, transshipment cost is fixed and complete pooling is assumed; then, how the inclusion of backorders and the relaxation of the complete pooling assumption would change subjects' behavior? What strategies would subjects use to set transshipment costs? How would subjects coordinate? What would be the effect of Face-to-Face communication? Finally, our experiments consider a system with two independent Newsvendors making decentralized decisions; then, what would be the effect of centralizing the decision process?

Appendix 4.1. Finesse analytical solution for a Newsvendor problem with transshipments

Taking advantage of the optimal solution of a typical Newsvendor problem (equation (4.2)) and the potential transshipment scenarios explained in Figure 1, I use a marginal approach to redefine the C_o as follows:

$$C_o = (c - b) * P(D_1 + D_2 < q_1 + q_2 | D_1 < q_1) + (t - r + c) * P(D_1 + D_2 > q_1 + q_2 | D_1 < q_1) \quad (3)$$

$$C_o = (c - b) * (1 - P(D_1 + D_2 > q_1 + q_2 | D_1 < q_1)) + (t - r + c) * P(D_1 + D_2 > q_1 + q_2 | D_1 < q_1) \quad (A1)$$

$$C_o = (c - b) + (b + t - r) * P(D_1 + D_2 > q_1 + q_2 | D_1 < q_1) \quad (A2)$$

And C_u as follows:

$$C_u = (r - c) * P(D_1 + D_2 > q_1 + q_2 | D_1 > q_1) + (t - c + b) * P(D_1 + D_2 < q_1 + q_2 | D_1 > q_1) \quad (4)$$

$$C_u = (r - c) * (1 - P(D_1 + D_2 < q_1 + q_2 | D_1 > q_1)) + (t - c + b) * P(D_1 + D_2 < q_1 + q_2 | D_1 > q_1) \quad (A3)$$

$$C_u = (r - c) + (t + b - r) * P(D_1 + D_2 < q_1 + q_2 | D_1 > q_1) \quad (A4)$$

Therefore, equation (4.2) can be redefined for a Newsvendor problem with transshipment as follows:

$$F(q_1^*) = \frac{C_u}{C_u + C_o} \quad (A5)$$

$$(C_u + C_o)F(q_1^*) = C_u \quad (A6)$$

$$C_o F(q_1^*) = C_u(1 - F(q_1^*)) \quad (A7)$$

Incorporating C_o and C_u :

$$\begin{aligned} & \left((c - b) + (b + t - r) * P(D_1 + D_2 > q_1^* + q_2^* | D_1 < q_1^*) \right) F(q_1^*) \\ & = \left((r - c) + (t + b - r) * P(D_1 + D_2 < q_1^* + q_2^* | D_1 > q_1^*) \right) (1 - F(q_1^*)) \end{aligned} \quad (A8)$$

$$\begin{aligned} & \left((c - b) + (b + t - r) * \frac{P(D_1 + D_2 > q_1^* + q_2^* \& D_1 < q_1^*)}{P(D_1 < q_1^*)} \right) F(q_1^*) \\ & = \left((r - c) + (t + b - r) * \frac{P(D_1 + D_2 < q_1^* + q_2^* \& D_1 > q_1^*)}{P(D_1 > q_1^*)} \right) (1 - F(q_1^*)) \end{aligned} \quad (A9)$$

$$\begin{aligned} & \left((c - b)F(q_1^*) + (b + t - r)P(D_1 + D_2 > q_1^* + q_2^* \& D_1 < q_1^*) \right) \\ & - \left((r - c)(1 - F(q_1^*)) + (t + b - r)P(D_1 + D_2 < q_1^* + q_2^* \& D_1 > q_1^*) \right) = 0 \end{aligned} \quad (A10)$$

$$\begin{aligned} & \left((c-b)F(q_1^*) + (b+t-r) \left(F(q_1^*) - P(D_1 + D_2 < q_1^* + q_2^* \& D_1 < q_1^*) \right) \right) \\ & - \left((r-c) \left(1 - F(q_1^*) \right) + (t+b-r)P(D_1 + D_2 < q_1^* + q_2^* \& D_1 > q_1^*) \right) = 0 \end{aligned} \quad (\text{A11})$$

$$(c-b)F(q_1^*) - (r-c) \left(1 - F(q_1^*) \right) + (t+b-r) \left(F(q_1^*) - F_{12}(q_1^* + q_2^*) \right) = 0 \quad (\text{A12})$$

$$-(t+b-r)F(q_1^*) - (r-c) + tF(q_1^*) + (t+b-r) \left(F(q_1^*) - F_{12}(q_1^* + q_2^*) \right) = 0 \quad (\text{A13})$$

$$-(r-c) + tF(q_1^*) - (t+b-r)F_{12}(q_1^* + q_2^*) = 0 \quad (\text{A14})$$

$$tF(q_1^*) = (r-c) + (t+b-r)F_{12}(q_1^* + q_2^*) \quad (\text{A15})$$

$$F(q_1^*) = \frac{(r-c) + (t+b-r)F_{12}(q_1^* + q_2^*)}{t} \quad (5)$$

Appendix 4.2. Interface of the experiment in Z-tree (in Spanish)

Period Trial 1 of 1 Remaining time [sec] 45

Por favor decida y luego confirme sus órdenes haciendo click en "OK"

Costo por unidad: 3
Precio de venta por unidad: 12
Costo de traspaso por unidad: 1

Unidades que desea ordenar a su distribuidor:

OK

The interface contains four empty coordinate graphs arranged in a 2x2 grid. Each graph has a vertical axis labeled 'Unidades' (Units) ranging from 0 to 100 and a horizontal axis labeled 'Periodo' (Period) ranging from 0 to 30. The top-right graph also has a vertical axis labeled '€\$' ranging from -300 to 900. All graphs are currently blank.

Appendix 4.3. Instructions Base Case – High Profit condition (in Spanish)

DECISIÓN GERENCIAL

Tiempo disponible: 60 minutos

En esta parte del experimento usted asumirá el papel de gerente de una empresa mayorista que vende un artículo específico. Su responsabilidad es decidir cuántas unidades **ordenar** a su proveedor antes de saber cuánto será el verdadero valor de la demanda.

Sin embargo, usted sabe que la demanda por los artículos durante la temporada de ventas está **uniformemente distribuida entre 1 y 100 unidades**. Es decir, la demanda durante cada temporada de ventas puede tomar cualquier valor entre 1 y 100 con una probabilidad de 1/100 (1%) para cada valor. Además no existe ninguna relación entre la demanda actual y la demanda pasada.

Usted compra los artículos a su distribuidor a un costo de **\$2** cada uno, y los vende a sus consumidores durante la temporada de ventas a un precio de **\$12** cada uno.

Por otra parte, usted no es el único mayorista en el mercado. En esta sala hay una persona más que está trabajando en su mismo mercado y la cual enfrenta una demanda igual a la suya. De ser posible, usted colabora con esta otra persona según las siguientes reglas:

- Los artículos que usted compra antes de que comience la temporada y que no logra vender durante ésta (exceso de oferta), usted los traspasa al otro mayorista en caso de que él necesite más unidades. En este caso el otro mayorista le pagará a usted los costos iniciales de compra (**\$2/unidad**) y él asumirá los costos de traspaso de las unidades traspasadas.
- Usted deja de ganar **\$10** por cada artículo que usted **no** compre antes de que comience la temporada y que podría haber vendido durante ésta (exceso de demanda). Sin embargo, en caso de que el otro distribuidor cuente con exceso de oferta, usted podrá recibir unidades de él. En este caso usted asumirá el costo de compra del artículo (**\$2/unidad**) más un costo de traspaso de **\$1/unidad**.
- Traspasos no serán posibles en el caso que ambos mayoristas cuenten con excesos de oferta o exceso de demanda.
- Todas las unidades no vendidas o no traspasadas al final de cada temporada no implican ninguna ganancia para usted, ya que estas unidades simplemente se desechan.

Su objetivo es decidir cuántas unidades **ordenar** a su proveedor cada período con el fin de **maximizar las ganancias acumuladas** del sistema al final de la simulación (**30 temporadas de venta**), y de acuerdo a su desempeño individual obtendrá un pago en dinero efectivo.

PAGO: El pago de esta parte del experimento será una suma variable entre \$0 y \$15000 en función de las ganancias acumuladas a lo largo de los 35 períodos; a mayor ganancia acumulada mayor pago.

Appendix 4.4. *p*-values for the t-tests comparing the performance measures computed for treatments T1 and T5 to 10

High Profit Treatments

	Av. Orders					Av. Euclidean Distance					Av. Transs/ In					Av. Transs/ Out				
	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1
T5	.00					.00					.00					.00				
T7-1	.00	.05				.00	.00				.00	.00				.00	.29			
T7-2	.00	.86	.25			.00	.00	.09			.00	.00	.00			.00	.09	.01		
T9-1	.00	.05	.93	.14		.00	.00	.00	.23		.00	.00	.59	.00		.00	.35	.15	.65	
T9-2	.00	.18	.67	.33	.43	.00	.00	.00	.05	.39	.00	.00	.01	.08	.02	.00	.08	.42	.01	.09
	Av. Profit - Subjects					Av. Profit - Computer					Av. Units Discarded					Av. Customer Satisfaction				
	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1	T1	T5	T7-1	T7-2	T9-1
T5	.29					.00					.00					.00				
T7-1	.16	.62				.00	.11				.00	.31				.46	.01			
T7-2	.03	.22	.43			.00	.93	.15			.00	.49	.20			.00	.00	.00		
T9-1	.02	.11	.28	.60		.00	.04	.59	.05		.12	.04	.32	.02		.32	.06	.77	.00	
T9-2	.01	.07	.23	.52	.51	.00	.59	.40	.57	.53	.01	.24	.83	.10	.71	.00	.02	.00	.55	.00

Low Profit Treatments

	Av. Orders					Av. Euclidean Distance					Av. Transs/ In					Av. Transs/ Out				
	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1
T6	.65					.00					.00				.00					
T8-1	.00	.02				.00	.00				.00	.02				.00	.00			
T8-2	.00	.00	.00			.00	.03	.00			.00	.88	.08		.00	.03	.02			
T10-1	.38	.67	.09	.00		.00	.00	.52	.00		.00	.25	.23	.29	.00	.00	.22	.00		
T10-2	.00	.00	.01	.24	.00	.00	.14	.00	.58	.00	.00	.63	.04	.59	.50	.00	.00	.35	.00	.56
	Av. Profit - Subjects					Av. Profit - Computer					Av. Units Discarded					Av. Customer Satisfaction				
	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1	T2	T6	T8-1	T9-2	T10-1
T6	.50					.00					.31					.01				
T8-1	.00	.03				.00	.58				.00	.01				.00	.02			
T8-2	.00	.00	.08			.00	.93	.45			.00	.00	.04			.00	.00	.31		
T10-1	.00	.00	.64	.22		.00	.43	.01	.34		.00	.00	.83	.04		.00	.08	.57	.17	
T10-2	.00	.00	.00	.14	.00	.00	.48	.73	.54	.04	.00	.00	.00	.17	.00	.00	.02	.91	.36	.44

Chapter 5. General Conclusions

This thesis contributes to the literature of Behavioral Operations Management by providing a wide range of analysis to understand retailers' ordering decisions and their implications in the supply chain performance. The main objective of this thesis is to better understand how different operational factors may independently and in combination influence retailers ordering decisions under different supply chain structures (single agent and multi agent), different demand uncertainty (deterministic and stochastic), and different interaction among retailers (no interaction, competition and cooperation). I developed three different studies that allow me to better understand the main dynamics and biases around retailers' ordering decisions. A summary of each of these three studies is presented below, followed by a discussion about future research opportunities.

5.1. Main Contributions

One of the main topics that I discuss in Chapter 2 refers to order amplifications. Amplifications usually take place in supply chains with tight capacity and long acquisition delays. Under scarce supply, the supplier needs to ration the allocation of available supply to satisfy retailers' orders, while retailers receiving only a fraction of previous orders, amplify future ones in an attempt to secure more units (Lee et al., 1997a, 1997b). Providing some initial constraints in supplier's capacity, this study analyzes the ordering behavior of a single retailer trying to match products received from its supplier with a deterministic and known final customer demand. In this system, retailers had the opportunity to store inventory from one selling period to the other.

I use an experiment to test subjects' ordering decisions under different ordering and supplier's capacity acquisition delays. Main results from this experiments show that subjects display limited ability to process the impact of delays and feedback, even when demand is known and constant and the system begins in equilibrium. The order trajectories follow a pattern of overshoot and subsequent undershoot until reaching an equilibrium. However, the initial overshoot is less intense and lasts longer than the optimal behavior, when subjects face longer delays. In addition, subjects inflate their orders when the supplier faces longer capacity acquisition delays and when orders take longer to be perceived by the supplier. Similarly, the analysis suggests that a simple anchoring and adjustment heuristic is able to

represent the ordering decision process. Finally, this chapter provides insights relevant to decision makers interested in the importance of improving decision-making and implementing Business Process Redesign (Van Ackere et al., 1993). However, such redesign requires breaking existing habits, understanding and carefully evaluating present processes (Van Ackere et al., 1993). Results suggest that retailers should try to decrease the delays inherent in their ordering processes. In doing so, retailers would reduce the complexity of the system, improve their ordering decisions as well as improve their ability to manage mismatches between supply and demand. In addition, managers should be careful when relying on rules-of-thumb because these heuristics perform substantially worse than optimal, which suggests significant opportunity for improvement. Naturally, while heuristics are simple and useful, if they are not good enough, they could lead to consistent biases, limited search, and resistance to change (Lazaric, 2000; Leonard-Barton, 1992).

Chapter 3 presents an extension to Chapter 2 by analyzing the effect of duplicated orders and retailers' competition on the bullwhip effect. I used a system composed by one supplier and two retailers, where retailers face a stable and known final customer demand, and the supplier may use two different allocation mechanisms. Results show that the bullwhip effect persists even when subjects have no incentives to inflate their orders. In addition, this chapter provides three important practical contributions: (i) a systems with lower dynamic complexity lead to lower average costs, (ii) when subjects face situations with higher complexity, subjects become more rational (activation of System II), and (iii) under identical supply chain structure, the use of the turn-and-earn allocation leads to a less intense bullwhip effect than a proportional allocation mechanism.

Finally, the Chapter 4 analyzes retailers' ordering decisions in situations where retailers need to make their decisions under demand uncertainty. I design a study to understand subjects' behavior in a single-supplier multi-retailer supply chain where I allow for transshipments among retailers. This chapter experimentally explores the effect of different profit and communication conditions as a way to create better interaction policies that may improve supply chain coordination. Results show a general retailers' underperformance in all profit conditions (presence of the pull-to-center behavior). These results are similar to the behavioral results observed in a typical newsvendor problem. However, some

practical improvements like Face-to-Face communication and the integration of analytical and behavioral models lead to a reduction of subjects' biases and to an increase in supply chain coordination.

5.2. Future work

Despite the meaningful discussions presented here, this thesis opens the doors for different extensions that could be addressed in future research.

As a result from Chapter 2, future research could explore possible ways to improve retailers' decisions either by prominently displaying important information, or by providing guidance regarding some heuristics that subjects should follow. This process provides clues to the required training that managers would need in to improve their decision-making processes. Similarly, other (dynamic) heuristics could be used to test subjects' behavior under different supply chain setting to improve the understanding of the decision process followed by the retailers. Finally, it would be desired to directly analyze other relevant dynamics (like learning effects) where other meaningful implications could be found.

Chapter 3 also provides some guidance about future research directions. For example, different decision rules can be used to explain subjects' behaviors and try to find mechanisms that may reduce subjects' biases. Also, given that I only evaluate the effect of two different allocation mechanisms on subjects behavior, we could potential evaluate subjects' reaction to the different kind of incentives offered by other allocation mechanism. Similarly, it would be ideal to test how would subject behavior change if we automate the decisions of one of the retailers and we assume that automatized retailer uses an optimizing strategy. Finally, in this experiment the supplier was automatized by the computer, it would be interested to relax this condition and allow and additional subject to play the role of a supplier and then test what is the preferred allocation strategy followed by the supplier and how retailers react to the introduction of a real supplier.

Finally, Chapter 4 can be extended in many different ways. First, we observed that there are different effects of learning that may require a deeper analysis. For example, how does previous experience in an isolated Newsvendor problem matter to the performance in a Newsvendor problem with transshipments? In addition, I have considered in these experiments a system composed by two symmetrical Newsvendors. Therefore, I can evaluate the effect on subjects' orders in an assymetrical

system where one of the retailers has higher market power. Similarly, these experiments consider a system with two independent Newsvendors making decentralized decision; then, what would be the effect of centralizing the decision process? Finally, I fixed the transshipment price. However, in practice, transshipment prices are usually determined by negotiation between the retailers involved in the transaction; they are not set centrally. Therefore, it would be interesting to use laboratory experiments to understand how retailers set transshipment prices and how their decisions deviate (or not) from the theoretical channel-coordinating benchmarks. In this case, we could even evaluate how the bargaining power or the type of negotiation protocol used affects the observed biases.

Bibliography

- Adelman, D. (2001). First Loss Sets Gloomy Tone at CISCO. *Wall Street Journal*, p. B1.
- Ahn, T. K., Ostrom, E., & Walker, J. (2011). Reprint of: A common-pool resource experiment with postgraduate subjects from 41 countries. *Ecological Economics*, 70(9), 1580–1589.
<http://doi.org/10.1016/j.ecolecon.2010.12.002>
- Anupindi, R., & Bassok, Y. (1999). Centralization of Stocks: Retailers vs. Manufacturers. *Management Science*, 45(2), 178–191. <http://doi.org/10.1287/mnsc.45.2.178>
- Armony, M., & Plambeck, E. L. (2005). The Impact of Duplicate Orders on Demand Estimation and Capacity Investment. *Management Science*, 51(10), 1505–1518.
<http://doi.org/10.1287/mnsc.1050.0371>
- Arrow, J. K., Harris, T., & Marschak, J. (1951). Optimal Inventory Policy Author. *Econometrica*, 19(3), 250–272.
- Axsäter, S. (2003). A New Decision Rule for Lateral Transshipments in Inventory Systems. *Management Science*, 49(9), 1168–1179. <http://doi.org/10.1287/mnsc.49.9.1168.16568>
- Barlas, Y., & Özevin, M. (2004). Analysis of stock management gaming experiments and alternative ordering formulations. *System Research and Behavioral Science*, 21(4), 439–470.
- Becker-Peth, M., Katok, E., & Thonemann, U. W. (2013). Designing buyback contracts for irrational but predictable newsvendors. *Management Science*, 59(8), 1800–1816.
<http://doi.org/10.1287/mnsc.1120.1662>
- Ben-Gal, I. (2005). Outlier detection. In O. Maimon & L. Rockach (Eds.), *Data Mining and Knowledge Discovery Handbook: A Complete Guide for Practitioners and Researchers* (pp. 1–16). Boston, MA: Kluwer Academic.
- Benzion, U., Cohen, Y., Peled, R., & Shavit, T. (2008). Decision-making and the newsvendor problem: an experimental study. *Journal of the Operational Research Society*, 59(9), 1281–1287.
<http://doi.org/10.1057/palgrave.jors.2602470>
- Bolton, G. E., & Katok, E. (2008). Learning by Doing in the Newsvendor Problem: A Laboratory Investigation of the Role of Experience and Feedback. *Manufacturing & Service Operations Management*, 10(3), 519–538. <http://doi.org/10.1287/msom.1060.0190>

- Bostian, a. 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. <http://doi.org/10.1287/msom.1080.0228>
- Byrme, J. A., & Elgin, B. (2002). Cisco behind the hype: CEO John Chambers still thinks his company can growth 30% a year. But critics question its aggressive accounting. *Business Week*, p. 54.
- Cachon, G. P., & Lariviere, M. (1999a). Capacity Choice and Allocation: Strategic Behavior and Supply Chain Performance. *Management Science*, 45(8), 1091–1108. <http://doi.org/10.1287/mnsc.45.8.1091>
- Cachon, G. P., & Lariviere, M. a. (1999b). Capacity Allocation Using Past Sales: When to Turn-and-Earn. *Management Science*, 45(5), 685–703. <http://doi.org/10.1287/mnsc.45.5.685>
- Cachon, G. P., & Lariviere, M. A. (1999c). An Equilibrium Analysis of Linear, proportional and uniform allocation of scarce capacity. *IIE Transactions*, 31, 835–849.
- Cachon, G. P., & Olivares, M. (2009). Competing Retailers and Inventory : An Empirical Investigation of General Motors’ Dealerships in Isolated U.S. Markets. *Management Science*, 55(9), 1586–1604. <http://doi.org/10.1287/mnsc.1090.1050>
- Cachon, G. P., Randall, T., & Schmidt, G. M. (2007). In Search of the Bullwhip Effect. *Manufacturing & Service Operations Management*, 9(4), 457–479. <http://doi.org/10.1287/msom.1060.0149>
- Cardenas, J.-C. (2000). How Do Groups Solve Local Commons Dilemmas? Lessons from Experimental Economics in the Field. *Environment, Development and Sustainability*, 2(3-4), 305–322. <http://doi.org/10.1023/A:1011422313042>
- Castillo, D., & Saysel, A. K. (2005). Simulation of common pool resource field experiments: A behavioral model of collective action. *Ecological Economics*, 55(3), 420–436. <http://doi.org/10.1016/j.ecolecon.2004.12.014>
- Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000). Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information. *Management Science*, 46(3), 436–443. <http://doi.org/10.1287/mnsc.46.3.436.12069>
- Chen, L., & Davis, A. M. (2014). *Validating Models for Bounded Rationality and Decision Biases in a*

Newsvendor Experiment.

- Chen, L., & Lee, H. L. (2011). Bullwhip effect measurement and its implications. *Operations Research*, 60(4), 771–778.
- Clark, A., & Scarf, H. (1960). Optimal policies for a multi-echelon inventory problem. *Management Science*, 6(4), 475–490.
- Crosron, R., & Donohue, K. (2005). Upstream versus downstream information and its impact on the bullwhip effect. *System Dynamics Review*, 21(3), 249–260. <http://doi.org/10.1002/sdr.320>
- Crosron, R., Donohue, K., Katok, E., & Sterman, J. D. (2014). Order stability in supply chains: Coordination risk and the role of coordination stock. *Production and Operations Management*, 23(2), 176–196. <http://doi.org/10.1111/j.1937-5956.2012.01422.x>
- Crosron, R., & Donohue, K. L. (2003). Impact of Pos Data Sharing on Supply Chain Management : an Experimental Study. *Production and Operations Management*, 12(1), 1–11. <http://doi.org/10.1111/j.1937-5956.2003.tb00194.x>
- Crosron, R., & Ren, Y. (2013). Overconfidence in Newsvendor Orders: An Experimental Study. *Management Science*, 59(11), 2502–2517. <http://doi.org/doi:10.1287/mnsc.2013.1715>
- Crosron, R., Schultz, K., Siemsen, E., & Yeo, M. L. (2013). Behavioral operations: The state of the field. *Journal of Operations Management*, 31(1-2), 1–5. <http://doi.org/10.1016/j.jom.2012.12.001>
- Davis, A. M. (2015). An experimental investigation of pull contracts in supply chains. *Production and Operations Management*, 24(2), 325–340. <http://doi.org/10.1111/poms.12224>
- De Véricourt, F., Jain, K., Bearden, J. N., & Filipowicz, A. (2013). Sex, risk and the newsvendor. *Journal of Operations Management*, 31(1-2), 86–92. <http://doi.org/10.1016/j.jom.2012.11.001>
- Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes*, 62(2), 198–215. <http://doi.org/10.1006/obhd.1995.1043>
- Dong, L., & Rudi, N. (2004). Who Benefits from Transshipment? Exogenous vs. Endogenous Wholesale Prices. *Management Science*, 50(5), 645–657. <http://doi.org/10.1287/mnsc.1040.0203>
- E, A., & Fine, C. (1999). Business cycles and productivity in capital equipment supply chains. In S.

- Tayur, R. Ganeshan, & M. Magazine (Eds.), *Quantitative Models for Supply Chain Management* (pp. 381–415). Norwell, MA: Kluwer Academic.
- Edgeworth, F. Y. (1888). The Mathematical Theory of Banking. *Journal of the Royal Statistical Society*, *51*(1), 113–127.
- Eeckhoudt, L., Gollier, C., & Schlesinger, H. (1995). The Risk-Averse (and Prudent) Newsboy. *Management Science*, *41*(5), 786–794. <http://doi.org/10.1287/mnsc.41.5.786>
- Ellram, L. M. (2010). Introduction to the forum on the bullwhip effect in the current economic climate. *Journal of Supply Chain Management*, *46*(1), 3–4. <http://doi.org/10.1111/j.1745-493X.2009.03178.x>
- Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178. <http://doi.org/10.1007/s10683-006-9159-4>
- Friedman, D., & Cassar, A. (1995). Dialogues with the data.
- Friedman, D., & Sunder, S. (1994). *Experimental Methods: A Primer for Economists*. Cambridge, Massachusetts: Cambridge University Press.
- Friedman, D., & Sunder, S. (2004). *Economics Lab: An Intensive Course in Experimental Economics*. London, UK.
- Gonçalves. (2003). *Demand bubbles and phantom orders in supply chains*. Massachusetts Institute of Technology. Retrieved from <http://moya.bus.miami.edu/~pgoncalves/PhD Thesis Goncalves.PDF>
- Gonçalves, P., & Arango, S. (2010). Supplier Capacity Decisions Under Retailer Competition and Delays: Theoretical and Experimental Results. In *System Dynamics Conference* (p. 28).
- Gonçalves, P., & Villa, S. (2016). Implications of Misperception of Feedback on Behavior in Operations. In M. Kunc, J. Malpass, & L. White (Eds.), *Behavioural Operational Research: Theory, Methodology and Practice*. London, UK: Palgrave Publishers.
- Gupta, S., Steckel, J. H., & Banerji, A. (2002). Dynamic Decision Making in Marketing Channels: An Experimental Study of Cycle Time, Shared Information and Customer Demand Patterns. *Experimental Business Research*.
- Herer, Y. T., Tzur, M., & Yücesan, E. (2006). The Multi-Location Transshipment Problem. *IIE*

- Transactions*, 38(3), 1–33.
- Ho, T.-H., Lim, N., & Cui, T. H. (2010). Reference Dependence in Multilocation Newsvendor Models: A Structural Analysis. *Management Science*, 56(11), 1891–1910.
<http://doi.org/10.1287/mnsc.1100.1225>
- Hu, X., Duenyas, I., & Kapuscinski, R. (2007). Existence of Coordinating Transshipment Prices in a Two-Location Inventory Model. *Management Science*, 53(8), 1289–1302.
<http://doi.org/10.1287/mnsc.1060.0694>
- Kahneman, D. (2011). *Thinking, Fast and Slow* (First Edit). New York, NY.
- Kaminsky, P., & Simchi-Levi, D. (1998). A new computerized beer game: A tool for teaching the value of integrated supply chain management. In H. L. Lee & S. M. Ng (Eds.), *Supply Chain and Technology Management* (Vol. 1, pp. 216–225). The Production and Operations Management Society. Retrieved from http://www.ieor.berkeley.edu/~kaminsky/Reprints/PK_DSL_98b.pdf
- Katok, E. (2011). Using Laboratory Experiments to Build Better Operations Management Models. *Foundations and Trends in Technology, Information and Operations Management*, 5(1), 1–86.
<http://doi.org/10.1561/02000000022>
- Kremer, M., Minner, & van Wassenhove, L. N. (2010). Do Random Errors Explain Newsvendor Behavior? *Manufacturing & Service Operations Management*, 12(4), 673–681.
<http://doi.org/10.1287/xxxx.0000.0000>
- Krishnan, K. S., & Rao, V. R. K. (1965). Inventory control in N warehouses. *Journal of Industrial Engineering*.
- Lazaric, N. (2000). The role of routines, rules and habits in collective learning: Some epistemological and ontological considerations. *European Journal of Economic and Social Systems*, 14, 157–171. <http://doi.org/10.1051/ejess:2000115>
- Lee, H. L., Padmanabhan, V., & Whang, S. (1997a). Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science*, 43(4), 546–558. <http://doi.org/10.1287/mnsc.43.4.546>
- Lee, H. L., Padmanabhan, V., & Whang, S. (1997b). The Bullwhip Effect In Supply Chains 1. *Review*, 38(3), 93–102. <http://doi.org/10.1287/mnsc.43.4.546>
- Leonard-Barton, D. (1992). Strategy process: managing corporate self-renewal. *Strategic Management*

- Journal*, 1(34), 111–125.
- Li, L. (1992). The role of inventory in delivery time competition. *Management Science*, 38(2), 182–197.
- Lu, L. X., & Lariviere, M. a. (2012). Capacity Allocation over a Long Horizon: The Return on Turn-and-Earn. *Manufacturing & Service Operations Management*, 14(1), 24–41.
<http://doi.org/10.1287/msom.1110.0346>
- Mitchell, T. W. (1924). Competitive Illusion as a Cause of Business Cycles. *The Quarterly Journal of Economics*, 38(4), 631–652.
- Moritz, B., Hill, A. V, & Donohue, K. L. (2013). Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management*, 31, 72–85.
- Netessine, S., Rudi, N., & Wang, Y. (2001). *Dynamic inventory competition and customer retention*. Philadelphia, PA.
- Oliva, R., & Gonçalves, P. (2005). Behavioral Causes of the Bullwhip Effect: “Satisficing” Policies with Limited Information Cues. *Proceedings of the 23rd International Conference of the System Dynamics Society*, (118-119).
- Olivares, M., Terwiesch, C., & Cassorla, L. (2008). Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Science*, 54(1), 41–55.
<http://doi.org/10.1287/mnsc.1070.0756>
- Olson, M. (1965). The Logic of Collective Action: Public Goods and the Theory of Groups. *Social Forces*. <http://doi.org/10.2307/2576430>
- Ostrom, E. (1990). Governing the commons: the Evolution of Institutions for Collective Action. *Cambridge University Press*, 271. <http://doi.org/10.2307/1243016>
- Ostrom, E. (1998). Theory of Collective Action A Behavioral Approach to the Rational Choice Presidential Address , American Political Science Association , *The American Political Science Review*, 92(1), 1–22.
- Ostrom, E. (2000). Collective Action and the Evolution of Social Norms. *Journal of Economic Perspectives*, 14(3), 137–158. <http://doi.org/10.1080/19390459.2014.935173>
- Ovchinnikov, A., Moritz, B., & Quiroga, B. F. (2015). How to Compete against a Behavioral

- Newsvendor. *Production and Operations Management*, 24(11), 1783–1793.
<http://doi.org/10.1111/poms.12361>
- Paich, M., & Sterman, J. D. (1993). Boom, bust, and learn in experimental markets. *Management Science*, 39(12), 1439–1458.
- Paterson, C., Kiesmueller, G., Teunter, R. H., & Glazebrook, K. D. (2011). Inventory models with lateral transshipments: a review. *European Journal of Operational Research*, 44(1524), 1–19.
<http://doi.org/10.1016/j.ejor.2010.05.048>
- Robinson, L. (1990). Optimal and Approximate Policies in Multiperiod, Multilocation Inventory Models with Transshipments. *Operations Research*, 38(2), 278–295.
<http://doi.org/10.1287/opre.38.2.278>
- Rudi, N., Kapur, S., & Pyke, D. F. (2001). A Two-Location Inventory with Transshipment and a Local Decision Making. *Management*, 47(12), 1668–1680.
- Scarf, H. (1959). Bayes solution of the statistical inventory problem. *Annals of Mathematical Statistics*, 30, 346–363.
- Schweitzer, M. E., & Cachon, G. P. (2000). Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. *Management Science*, 46(3), 404–420.
- Senge, P., & Sterman, J. D. (1992). Systems thinking and organizational learning: Acting locally and thinking globally in the organization of the future. In T. Kochan & M. Useem (Eds.), *Transforming Organizations* (pp. 353–371). Oxford: Oxford University Press.
- Simon, H. A. (1982). Models of Bounded Rationality: Empirically Grounded Economic Reason. *MIT Press*.
- Smith, V. L. (1976). Economics : Induced Experimental Value Theory. *American Economic Review*, 66(2), 274–279. <http://doi.org/10.2307/1817233>
- Smith, V. L. (1982). Microeconomic Systems as an Experimental Science. *The American Economic Review*, 72(5), 923–955.
- Sošić, G. (2006). Transshipment of inventories among retailers: Myopic versus farsighted stability. *Management Science*, 52(10), 1493–1508.
- Stanovich, K. E., & West, R. F. (2000). Individual difference in reasoning: implications for the

- rationality debate? *Behavioural and Brain Sciences*, 23(1), 645–726.
- 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. <http://doi.org/10.1287/mnsc.1030.0169>
- Sterman, J. D. (1989a). Misperceptions of feedback in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 43(3), 301–335. [http://doi.org/10.1016/0749-5978\(89\)90041-1](http://doi.org/10.1016/0749-5978(89)90041-1)
- Sterman, J. D. (1989b). Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*, 35(3), 321–339. <http://doi.org/10.1287/mnsc.35.3.321>
- Sterman, J. D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Chicago, Illinois: Irwin-McGraw Hill.
- Sterman, J. D., & Dogan, G. (2015). “I’m not hoarding, I’m just stocking up before the hoarders get here.” Behavioral causes of phantom orderin in supply chains. *Journal of Operations Management*, 39-40(C), 6–22. <http://doi.org/10.1016/j.jom.2015.07.002>
- Tagaras, G. (1989). Effects of pooling on the optimization and service levels of two-location inventory systems. *IIE Transactions*, 21(3), 250–257.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science (New York, N.Y.)*, 185(4157), 1124–1131. <http://doi.org/10.1126/science.185.4157.1124>
- Van Ackere, A., Larsen, E. R., & Morecroft, J. D. W. (1993). Systems thinking and business process redesign: An application to the beer game. *European Management Journal*, 11(4), 412–423.
- Villa, S., Gonçalves, P., & Arango, S. (2015). Exploring retailers’ ordering decisions under delays. *System Dynamics Review*, 31(1), 1–27. <http://doi.org/10.1002/sdr>
- Yasarcan, H. (2011). Stock management in the presence of significant measurement delays. *System Dynamics Review*, 27(1), 91–109.
- Yasarcan, H., & Barlas, Y. (2005). A generalized stock control formulation for stock management problems involving composite delays and secondary stocks. *System Dynamics Review*, 21(1), 33–68. <http://doi.org/10.1002/sdr.309>