

ESSAYS ON BEHAVIORAL OPERATIONS:  
MANAGERS, ALGORITHMS, AND OPERATIONAL DECISIONS

A Dissertation

by

HAN KYUL OH

Submitted to the Graduate and Professional School of  
Texas A&M University  
in partial fulfillment of the requirements for the degree of  
DOCTOR OF PHILOSOPHY

Chair of Committee,	Rogelio Oliva
Committee Members,	James Abbey
	Gregory R. Heim
	Alina Sorescu
Head of Department,	Rich Metters

May 2022

Major Subject: Business Administration

Copyright 2022 Han Kyul Oh

## ABSTRACT

This dissertation presents three essays that investigate how managers leverage heuristics in a complex and dynamic business environment, focusing on retail as the primary business context. I provide insights on how managers use heuristics and interact with algorithms to make better operational decisions and how they make retail decisions involving consumer returns as well as order quantity and pricing decisions. The first essay examines, through econometric analysis of detailed transactional data, whether managers possess exogenous information (i.e., information not utilized by the algorithms) and the capability to transform that information into modifications that improve upon algorithmic decisions. I probe the boundary conditions of when algorithms work and how managers can contribute outside the boundaries. I find that managers can systematically and consistently improve restocking recommendations from algorithms, yielding an average savings of 2.5% of the cost of goods sold.

The second essay presents a behavioral analysis of return policy decision-making in a retail environment with aggregate demand and individual product valuation uncertainties. Leveraging a generalized newsvendor model, I conduct a randomized behavioral experiment to understand how individuals make order quantity, price, and refund amount decisions and the causal effect of salvage value on these decisions. I find that the salvage value plays an important role in affecting return policy decisions. Further, I reveal several time-dependent behavioral regularities in decision-making that I explain through a process theory, thus providing insights regarding how decision-makers use heuristics to dynamically make decisions and a new direction with testable hypotheses for future research.

The third essay investigates how decision-making environments affect managers' exploration of decision sets, which is fundamental in understanding the adaptive behavior in decision-making under uncertainty. I hypothesize the decision-maker's exploration behavior using two individual-level decision-making theories, prospect theory and bounded rationality. I analyze how individuals respond to environmental change and unpack their exploration behavior using the generalized

news vendor experiment introduced in the second essay. Further, using process theory as our theoretical lens, I test and explain the mechanism of how profit performance can mediate the effect of changing environment on explorations.

## DEDICATION

To my family with love.

## ACKNOWLEDGMENTS

A wise man once told me that “it is the people you meet along the way that makes Ph.D. worth it.” I am fortunate to have met many wonderful and intelligent people during my Ph.D., and without them completing my Ph.D. would have been impossible.

I want to thank Dr. Rogelio Oliva for all his guidance, wisdom, and patience in teaching me how to be a good scholar. Also, I want to thank the Mays OM family and my doctoral friends, Drs. Huseyn Abdulla, Seulchan Lee, and Mayukh Majumdar, who made this journey fun and enjoyable! Lastly, this dissertation was possible only because of the support from my family.

## CONTRIBUTORS AND FUNDING SOURCES

### **Contributors**

This work was supported by a dissertation committee consisting of Professor Rogelio Oliva [advisor], Professors James Abbey, Gregory R. Heim of the Department of Information and Operations Management, and Professor Alina Sorescu of the Department of Marketing.

The data analyzed for Chapter 2 was provided by Professor Rogelio Oliva. The work depicted in Chapter 3 were completed in collaboration with Huseyn Abdulla of the Department of Information and Operations Management.

All other work conducted for the dissertation was completed by the student under supervision of Professor Rogelio Oliva.

### **Funding Sources**

Graduate study was supported by a fellowship from Mays Business School, Texas A&M University. This work was made possible in part by grants from Mays Business School and Mays Innovation Research Center.

The data analyzed for Chapter 3 and 4 was collected with the approval from the Institutional Review Board of Texas A&M University (IRB2019-1124M).

Its contents are solely the responsibility of the author and do not necessarily represent the official views of the Mays Business School and Mays Innovation Research Center

## NOMENCLATURE

ASO	Automated Stock Ordering
BC 95% CI	Bias-Corrected 95% Confidence Intervals
COGS	Cost of Goods Sold
CRP	Consumer Return Policy
DE	Dual Entitlement
Fr	Experimental Franc
LMM	Linear Mixed-effect Model
OOS	Out Of Stock
SKU	Stock Keeping Unit

## TABLE OF CONTENTS

	Page
ABSTRACT .....	ii
DEDICATION .....	iv
ACKNOWLEDGMENTS .....	v
CONTRIBUTORS AND FUNDING SOURCES .....	vi
NOMENCLATURE .....	vii
TABLE OF CONTENTS .....	viii
LIST OF FIGURES .....	x
LIST OF TABLES.....	xi
1. INTRODUCTION.....	1
2. BETTER TOGETHER? HOW MANAGERS CAN COMPLEMENT ALGORITHMS ....	4
2.1 Introduction.....	4
2.2 Literature Review .....	7
2.3 Research Context.....	11
2.4 Research design and hypotheses .....	16
2.5 Data Sources and Variables .....	22
2.5.1 Variables .....	22
2.6 Empirical Models .....	27
2.7 Results .....	30
2.7.1 Robustness Test.....	36
2.8 Discussion and Conclusions .....	39
3. RETURN OF THE BEHAVIORAL NEWSVENDOR: AN EXPERIMENTAL ANALYSIS OF CONSUMER RETURN POLICY DECISIONS.....	45
3.1 Introduction.....	45
3.2 Normative Model and Hypotheses .....	47
3.3 Behavioral Experiment .....	51
3.3.1 Experimental Design .....	51
3.3.2 Procedure and Sample.....	52
3.4 Testing the Hypotheses .....	53



3.5	Building a Process Theory .....	58
3.5.1	Within- and Cross-Period Decision-Making .....	60
3.5.1.1	Order Quantity .....	61
3.5.1.2	Price .....	62
3.5.1.3	Refund Amount .....	64
3.5.2	Within-Round Decision-Making .....	66
3.5.3	Process Theory Summary .....	69
3.6	Conclusion .....	71
4.	SATISFICING AND EXPLORATION BEHAVIOR UNDER COMPLEX AND DYNAMIC DECISION-MAKING ENVIRONMENT .....	74
4.1	Introduction .....	74
4.2	Literature Review and Hypotheses .....	75
4.2.1	Bounded Rationality, Prospect Theory, and Exploration Behavior .....	76
4.3	Data, Variables, and Models .....	80
4.3.1	Variables .....	81
4.3.2	Empirical Models .....	81
4.3.3	Multilevel Moderated Mediation Model .....	83
4.4	Results .....	84
4.4.1	Multilevel Moderated Mediation Analysis .....	86
4.5	Discussion and Conclusion .....	87
5.	CONCLUSION .....	91
	REFERENCES .....	94
	APPENDIX A. BETTER TOGETHER? HOW MANAGERS CAN COMPLEMENT ALGORITHMS .....	109
	APPENDIX B. RETURN OF THE BEHAVIORAL NEWSVENDOR: AN EXPERIMENTAL ANALYSIS OF CONSUMER RETURN POLICY DECISIONS .....	111
B.1	Derivation of Nominal Price under Normally Distributed Valuation .....	111
B.2	Proof that Price is Monotonically Increasing in the Salvage Value .....	111
B.3	Experimental Task Interface .....	113
B.4	Pre-defined Exclusion Criteria .....	117
B.5	Correlation Table .....	119
B.6	Examining the Role of Individual Heterogeneity .....	120
B.6.1	Demographic Characteristics .....	120
B.6.2	Socio-Economic Preferences .....	123
B.6.3	Performance Implications of Individual Heterogeneity .....	123

## LIST OF FIGURES

FIGURE	Page
2.1 Histograms of 5th to 95th percentile of <i>Cost Savings</i> of Modified Proposals .....	15
3.1 Order Quantity Decisions by Treatment Blocks and Periods .....	61
3.2 Pricing Decisions by Treatment Blocks and Periods .....	63
3.3 Refund Decisions by Treatment Blocks.....	65
3.4 Preliminary Structural Causal Models .....	67
3.5 Final Structural Causal Model .....	69
4.1 Moderated Mediation Conceptual Diagram .....	80
4.2 Exploration Behavior by Treatment Blocks and Periods .....	86
4.3 Moderated Mediation Analysis—Profits by Treatment Blocks and Periods .....	87

## LIST OF TABLES

TABLE	Page
2.1 Distribution of <i>Cost Savings</i> of Modified Proposals .....	15
2.2 Average <i>Cost Savings Pct</i> of Modified Proposals .....	16
2.3 Frequency of Modification Decisions by Managers .....	25
2.4 Summary Statistics .....	26
2.5 Correlation Matrix .....	26
2.6 Estimation of Multinomial Logistic Regressions .....	32
2.7 Estimation of Heckman Sample Selection Model .....	34
2.8 Estimation of Heckman Sample Selection model with Normalized Cost Savings .....	37
3.1 Experimental Design and Parameterization .....	51
3.2 One-Sample Wilcoxon Sign Test Results .....	54
3.3 Results of Linear Mixed-Effect Model for Testing H2 .....	56
3.4 Effect of Treatment Sequence .....	59
3.5 Analysis of Structural Causal Models .....	68
3.6 Summary of the Proposed Process Theory .....	70
4.1 Effect of Treatment Sequence on Profit and Exploration .....	85
A1 Sample Selection Model with Varying Inventory Holding and Substitution Rates .....	109
A2 Estimation of Heckman Sample Selection Model with Varying Inventory Holding and Substitution Rates with Normalized Cost Savings .....	110
B1 Correlations between the Demographic and Socio-economic Measures .....	119
B2 Frequency Table of Demographic Characteristics .....	121
B3 Analysis of Individual Heterogeneity: Demographic Characteristics .....	122
B4 Analysis of Individual Heterogeneity: Socio-Economic Preferences .....	124

B5 Analysis of Demographic Characteristics: Deviations from Normative Decisions . . . . 126

B6 Analysis of Socio-Economic Characteristics: Deviations from Normative Decisions. 127

## 1. INTRODUCTION

As evidenced by the rise of studies involving data-driven decision-making (e.g., machine learning and artificial intelligence), the use of algorithms to automate managerial decisions have gained traction in both practice (Gil et al. 2020, Tarafdar et al. 2019) and academia (Dogru and Keskin 2020, Helo and Hao 2021, Mišić and Perakis 2020). These studies often highlight the benefits of algorithmic decision-making, such as efficiency and scalability, without fully considering the role of managers in the decision-making process (Athey 2017, Oh and Oliva 2021). Specifically, we are yet to fully understand how managers interact with algorithms and whether they can positively influence the algorithmic decision-making system in a real business environment.

To address this issue, the first essay (§2), “Better Together? How Managers can Complement Algorithms” examines – through econometric analysis of detailed transactional data – whether managers possess exogenous information (i.e., information not considered by the algorithms) and the capability to transform that information into modifications that improve upon algorithmic decisions in the context of retailers’ restocking decisions. In the retail industry, restocking decisions are important and challenging because inventory makes up a significant portion of most retailers’ statements of financial position (Gaur et al. 2005) and a vast number of SKUs present in each store complicates the decisions. To help the retailers make restocking decisions, inventory policies developed by operations management scholars have been implemented via an automated stock ordering (ASO) system that either makes restocking decisions (Chao et al. 2019) or recommends those decisions to managers (e.g., Van Donselaar et al. 2010). Given the important and challenging nature of restocking decisions, it is not surprising that ASO decisions have become standard and local retail managers are seldom given responsibility for restocking decisions (Begley et al. 2019, Felix et al. 2018).

In this essay, using data from a retailer that allows managers to modify ASO recommendations, I assess the performance of retail managers’ modifications by observing their impact in the real world. That is, I explore situations in a complex decision-making environment, characterized by

imperfect information and possible information asymmetry between human decision-makers and analytical models, in which the former, despite limited computational abilities, have the potential to outperform the latter (e.g., DeMiguel et al. 2009, Fildes et al. 2009). This research contributes to behavioral operations by empirically showing the benefits of allowing managers with local information to override algorithmic decisions, thereby challenging the research paradigm in behavioral operations that managers are a liability (due to behavioral bias) and not an asset (Becker-Peth and Thonemann 2019, Tversky and Kahneman 1974).

In my second and third essay, I use laboratory experiments to understand how human decision-makers leverage heuristics in a complex and dynamic decision-making environment, which may enlighten us on how human decision-makers can be incorporated into automated decision-making systems. Using retail decisions on consumer return policies as the research context, my second essay (§3), "Return of the Behavioral Newsvendor: An Experimental Analysis of Consumer Return Policy Decisions," investigates the behavioral aspects of return policy decisions and their interaction with order quantity and pricing decisions. Consumer return policy is an important managerial problem in today's retail environment because the annual value of consumer returns exceeds \$643 billion globally (Cheng 2015), which poses significant triple bottom-line challenges (i.e., focus on social and environmental concerns as well as profits) to retailers. Due to this importance retailers are experimenting with different return policies to reduce the burden of returns while trying to not compromise the value proposition and hurt sales performance. Cognizant of the increased importance of the problem for managers, academic research on consumer return policy (CRP) design has also grown significantly (see Abdulla et al. 2019, for a comprehensive review). Particularly, the literature on consumer return policy has focused on developing analytical models that prescribe optimal return policies (e.g., Akçay et al. 2013, Altug and Aydinliyim 2016, Su 2009). These models find that a partial refund is the optimal return policy and that the optimal refund amount depends on the salvage value of the returned product. I test the predictions of analytical models through a behavioral experiment. Through this research, I provide insights regarding 1) how managers set return policies, 2) how return policy decisions interact with pricing and ordering de-

cisions, and 3) how decision-makers use heuristics to dynamically make decisions under complex decision-making environments.

Further, I extend the research scope beyond the simple hypotheses testing and aim to explain the observed behavioral regularities, similar to the first behavioral study on the classical single-lever newsvendor model by Schweitzer and Cachon (2000). I employ an abduction process to provide explanations for the observed behaviors across three decision levers based on established judgment and decision-making theories. To this end, I first identify several time-dependent, dynamic behavioral regularities and dependencies in the decision process, and then articulate a *process theory* to explain these regularities, thus providing a new direction with testable hypotheses for future research.

In my follow up essay (§4), "Satisficing and Exploration Behavior Under Complex and Dynamic Decision-making Environment," I utilize the data from the second essay and extend the research by investigating the mechanism behind how the decision-making environment affects managers' explorations of decision sets. Managing exploration and exploitation tradeoff is an important behavior that forms the basis for decision-making under uncertainty and is fundamental in understanding the adaptive behavior in complex and dynamic decision-making environments (Laureiro-Martínez et al. 2010). However, despite this important little is known about the antecedents of exploratory behavior of individuals (Hardy III et al. 2014), and the mechanism of their exploration behaviors (Laureiro-Martínez et al. 2010). I contribute to both management and behavioral operations literature by using two individual-level decision-making theories – prospect theory (Kahneman and Tversky 1979) and bounded rationality (Simon 1957) – to hypothesize the exploration behavior in complex and dynamic decision-making environments. Specifically, I investigate how individuals respond to environmental change, and unpack their exploration behavior using a new experimental design, the generalized newsvendor experiment (Oh et al. 2021). Further, I test the mechanism of how profit performance can mediate the effect of changing environment on explorations. I do so using process theory (Mohr 1982) as our theoretical lens, thereby responding to calls for more process focused rationality (Levinthal 2011, Levinthal and March 1993).

## 2. BETTER TOGETHER? HOW MANAGERS CAN COMPLEMENT ALGORITHMS

### 2.1 Introduction

Inventory is a significant asset on most retailers' statements of financial position (Gaur et al. 2005). In 2011, inventory represented approximately 21% of total assets in an average U.S. public retailer (Gaur et al. 2014), and according to the U.S. Census Bureau, U.S. retailers carried inventories worth \$664 billion in May 2019 (seasonally adjusted), shadowing sales in the same period with an inventory to sales ratio of 1.46. To control inventory, operations management scholars have developed inventory policies that assume that inventory levels can be computed to minimize inventory costs, which include the cost of lost sales and inventory holding cost (Cachon and Terwiesch 2012, Porteus 2002, Silver et al. 1998). In the retail industry, these inventory policies are often implemented via an automated stock ordering (ASO) system that either makes restocking decisions (Chao et al. 2019) or recommends those decisions to managers (e.g., Van Donselaar et al. 2010). Given the vast number of SKUs in the retail space and complexity of inventory models that employ various forecasting methods, it is not surprising that ASO decisions have become standard and local retail managers are seldom given responsibility for restocking decisions (Begley et al. 2019, Felix et al. 2018).

Indeed, the evidence suggests that local managers are not good at managing inventory. Experiments reveal that managers faced with demand uncertainty consistently choose order quantities that deviate from optimal toward mean demand, referred to as pull-to-center bias (Schweitzer and Cachon 2000). In a supply chain context, the beer distribution game similarly shows managers at different tiers of the supply chain to choose restocking quantities that, typically ignoring the backlog of orders placed but not yet received, systematically deviate from optimal (Croson et al. 2014, Sterman 1989). These suboptimal decisions have traditionally been attributed to "behavioral biases" (Becker-Peth and Thonemann 2019) introduced by, or reflecting, cognitive limitations of managers (Feiler et al. 2013, Sterman 1989, Tong and Feiler 2016), giving rise to the argument



that managers are not as reliable as algorithms (e.g., Bolton and Katok 2008, Bolton et al. 2012, Bostian et al. 2008, Croson et al. 2014, Schweitzer and Cachon 2000).

These studies, typically ran in controlled lab experiments, compare performance between participants and an “optimal” inventory model by providing both with correct model assumptions and accurate parameter values. As such, the studies exclude the possibility of exogenous information that might render the model assumptions inoperative or alter parameter values (Silver 1981, Wagner 2002). Weather, for example, although not a factor commonly considered in inventory models, can significantly affect daily demand and, taken into account, yield more accurate estimates of consumer demand (Steinker et al. 2017). Nor is the hosting of sporting events commonly considered in inventory models, yet many studies have found a positive impact of sporting events on local economies (e.g., Gratton et al. 2006, Wilson 2006). Human decision makers, however, tend to be good at observing and extrapolating from exogenous information that algorithms typically ignore (Armstrong 1983, Brown 1996, Lawrence et al. 2006, Lim and O’Connor 1996). Studies in judgmental forecasting show experienced managers to be able to leverage exogenous information, such as weather, shipment delays, and interaction with customers, to improve the quality of demand forecasts in settings as varied as warehouses (Sanders and Ritzman 1992), manufacturing, and retail (Fildes et al. 2009), and Van Donselaar et al. (2010) found that managers in a supermarket who adjusted proposals from the ASO system to take into account in-store logistics costs achieved better balanced workloads and in-stock performance.

Note that the ‘exogenous information’ is not restricted to alternative data sources that managers may observe. Pearl and Mackenzie (2018) argue that humans can react differently compared to algorithms when exposed to the same information because humans are causal machines, whereas algorithms can only synthesize correlational information. This significant difference enables managers to react to potential causes to the observations rather than reacting to observations themselves.

In this study, using data from a retailer that allows managers to modify ASO recommendations, we assess managers’ ability to leverage exogenous information to improve an ASO system’s

recommendations. We test their ability to improve on the ASO system's recommendations by considering managers' computational capacity together with the decision-making environment, the two key factors in the theory of bounded rationality (Simon 1955). Acknowledging the impossibility of ascertaining precisely what information is available to managers and their priorities in making modification decisions, we assess managers' access to and effectiveness at converting information into cost-saving decisions by studying the ASO recommendations they choose to modify and the impact of their adjustments on store operations and sales. This strategy is similar to what Samuelson (1937, 1938, 1948) did to capture people's unobservable preference when developing revealed preference theory. If the system makes restocking recommendations based on adequate information, the resulting recommendation will be close to optimal and difficult to improve by an uninformed manager. We posit, however, that system performance can be improved by adjusting restocking proposals if managers have access to and the capacity and capability to process useful exogenous information.

We assess the performance of managers' modifications by observing their impact in the real world. That is, we explore situations in a complex decision-making environment, characterized by imperfect information and possible information asymmetry between human decision-makers and analytical models, in which the former, despite limited computational abilities, have the potential to outperform the latter (e.g., DeMiguel et al. 2009, Fildes et al. 2009).

Sample data from our research partner, a multinational retailer with more than 400 stores, consists of 26,993 inventory decisions made by local retail managers. Our finding that managers systematically and consistently improved restocking recommendations, yielding an average savings of 2.5% of cost of goods sold, suggests the presence of *blind spots* in system algorithms (Oliva and Watson 2009), whether through restricted access to information or constrained solution space, and the potential for managers to exploit exogenous information to inform decisions that reduce inventory costs. The industry practice of not allowing local modifications may thus deprive organizations of the opportunity to identify and leverage useful exogenous information. Our results show that afforded a complementary role, managers can introduce flexibility (e.g., adaptability) to effi-

cient, but inflexible ASO system decision-making algorithms. We further posit that establishing a *default choice* (i.e., the system’s proposed restocking decision) serves to ease managers’ cognitive load, enabling them to focus on important changes that need to be addressed. Our findings are in contrast to the majority of studies in the inventory decisions literature, which argue that managers are a source of behavioral bias and a liability.

The paper is structured as follows. We review the related literature in §2.2 and introduce our research context in §2.3. Our research hypotheses are articulated in §2.4 and data sources and variables are described in §2.5. We present our empirical models and results in §2.6 and §2.7. Implications of and conclusions from our findings are discussed in §2.8 together with limitations of our study and suggestions for future research.

## **2.2 Literature Review**

Our research, in exploring the impact and performance of managers in a complex, dynamic decision-making environment involving inventory restocking decisions builds on several streams of literature.

The possibility that the performance of inventory models may be improved by considering additional exogenous information was hypothesized in the inventory decisions literature by Silver (1981) and Wagner (2002). Silver (1981) suggests that the existence of exogenous information may significantly increase the discrepancy between simplified model assumptions (i.e., parameter specifications) and current business realities. Since the performance of the inventory model depends on how well its assumptions reflect the business realities (Rumyantsev and Netessine 2007), inventory decisions without considering the exogenous information may lead to suboptimal decisions. Wagner (2002), pointing out that data issues (e.g., missing observations, measurement errors, and discontinuous demand) might introduce significant errors in parameter estimates, suggests that to improve the quality of decisions, “on occasion human intervention is called for” (p. 224). Analytical studies by Eroglu et al. (2013) and Prak et al. (2016) suggest that most inventory models assume the true value of parameters to be known, and demonstrate that not knowing the parameters or using a wrong parameter estimate has a detrimental effect on performance. Em-

empirical evidence that managers have access to exogenous information, as suggested by Silver and Wagner, is lacking. While our study does not attempt to assess the accuracy of model parameter, it contributes to this stream of research by empirically testing the possibility that managers have access to exogenous information that might affect model assumptions and the effectiveness of a model's recommendations.

Behavioral research in the inventory decisions literature (see Becker-Peth and Thonemann (2019) for a summary), which focuses on the discovery and explanation of behavioral biases in inventory decisions, has consistently identified managerial biases in inventory decisions including pull-to-center bias (Schweitzer and Cachon 2000), overconfidence (Ren and Croson 2013), misperception of feedback (Sterman 1989), and censorship bias (Feiler et al. 2013). These studies often rely on experiments that compare performance between participants and an optimal inventory policy that has access to the same information as the decision makers. Under this setting, inventory decisions made by the participants that systematically deviate from the inventory policy are identified as behavioral biases. This research design follows the theory of “logical rationality. . . [in which human behaviors are evaluated] against the laws of logic or probability rather than success in the world” (Todd and Gigerenzer 2012, p.15), and suggest that managers are a liability (Tversky and Kahneman 1974). Our research expands the behavioral decision-making literature in three important ways.

First, our research departs from logical rationality, focusing instead on a broader model of decision-making to explain and assess managers' behavior. Simon (1955) argues that understanding human rationality requires analysis of the interactions between two key factors, the computational capabilities of the decision-maker and the structure of the decision-making environment. Drawing on Simon's work, Gigerenzer and Gaissmaier (2011) have shown there to be situations in a complex decision making environment in which human decision-makers, despite limited computational capabilities, can outperform analytical models. To date, studies in the field of behavioral operations have focused mainly on the computational capabilities of decision-makers; they have not evaluated these capabilities in a realistic decision-making environment (Katsikopoulos and

Gigerenzer 2013). Specifically, previous studies (e.g., Benzion et al. 2008, Bolton et al. 2012, Schweitzer and Cachon 2000) focused on limitations of computational capabilities in news vendor settings have kept environmental factors relatively constant by assuming an environment with perfect information. Our research contributes to the behavioral operations literature by considering the structure of the decision-making environment as well as computational limitations of the decision-maker.

Second, our research design takes into account, by observing managers' decisions in practice, information they may possess. Previous studies have shown that managerial understanding, whether accrued from repeated exposure to the decision-making task, that is, autonomous learning (Bolton and Katok 2008), or through explicit training (Bolton et al. 2012, Schweitzer and Cachon 2000), positively affects inventory decisions. Moreover, Bolton et al. (2012) and Moritz et al. (2013) assessed in laboratory experiments the expertise professional managers accrued from experience. Their studies found participants exposed to learning processes and experienced professional managers alike to exhibit behavioral (e.g., pull-to-center) bias. In practice, expert decision-making may depend on the information professional managers possess as well as on-the-job experience. Although contextual information available to managers is difficult to incorporate (or account for) in a laboratory setting, by evaluating the effect of managers' decisions on the performance of an ASO system in the real world, our research explores managerial ability to leverage information not available to the ASO system. Particularly, our approach considers the possibility that both behavioral biases and advantages of information not available to the ASO system in decision-making exist in practice and focuses on the combined effect of such behavior in the real world. This approach departs from the current paradigm in behavioral operations that focuses on isolating biases and attempting to explain managers' decisions by explaining small individual parts.

Third, our research contributes to the literature on behavioral decision theory in complex decision-making environments by assessing the impact of managerial decisions in the presence of a default choice. Previous research suggests that decision-makers are strongly inclined to stick with the default choice even in the presence of an alternative with higher (Samuelson and Zeck-

hauser 1988) and unambiguous (Roca et al. 2006) utility. This phenomenon, referred to as status quo bias (Samuelson and Zeckhauser 1988), is found in such varied decision-making contexts as retirement fund investments (Samuelson and Zeckhauser 1988) and organ donor decisions (Johnson and Goldstein 2003). The significant body of knowledge on status quo bias notwithstanding, prior studies of inventory decisions have not considered the impact of a default choice on managers' inventory decisions. Providing a default choice (i.e., ASO proposals) relieves managers of the need to modify every restocking proposal, enabling them to flexibly allocate their cognitive resources to a subset of inventory decisions, leveraging exogenous information to selectively modify proposals the quality of which may be improved. We contribute to the literature on status quo bias by analyzing the modification decisions of managers in practice. Specifically, we investigate whether managers systematically deviate from the default choice, and whether deviations yield improved performance.

One previous article has explored a similar context as our research paper. Van Donselaar et al. (2010) explore how managers use modifications to ASO proposals to balance their workload and improve in-stock performance. In their supermarket setting, managers' ability to balance the workload via order adjustments is dependent on short (i.e., less than one day) and reliable lead time. Through interviews with managers and empirical analysis of inventory orders, they find that managers consistently modified the ASO proposals to account for in-store logistics costs ignored by the ASO system. The authors propose an algorithm that accounts for these costs and show improved in-stock performance. Our work expands by considering a decision-making environment where managers have less control of inventory inflow due to variable and longer lead times, thus evaluating modifications under a more complex and prevalent context in the retail industry. Methodologically, we also expand by using a causal inference model to assess the impact of managerial modifications. In addition to these contextual and methodological differences, our research goals are also different. Rather than attempting to derive the heuristics that managers are using, we leverage the complexity and dynamism of our decision making environment to develop an understanding of the decision characteristics and the context in which managers operate and can successfully im-

prove the ASO proposals. The aim of our study is to lay the foundation for the design of heuristics that can improve decision-making depending on the available information and their operating environment. The prerequisites for this design phase of heuristics research are descriptive theories of heuristics and normative theories of how they work in different environments (Gigerenzer and Selten 2001, Gigerenzer et al. 1999).

### **2.3 Research Context**

Our research partner is a multinational retailer with 400+ stores and 40,000+ employees around the world. Each approximately 100,000 square foot store carries, on average, 33,000 SKUs. Replenishment decisions for thousands of stock keeping units (SKUs) are made every weekday by an internally-developed ASO system that implements a periodic review ( $R, s, nQ$ ) inventory policy (Hadley and Whitin 1963). Average review period ( $R$ ) per SKU is one week, and the reorder level ( $s$ ) is based on the greater quantity between the 98% service level and marketing related minimum display quantity. A restocking order is triggered when on-hand inventory falls below either the stock level representing the 98% service level or a marketing-determined minimum display quantity. Replenishment orders are placed in multiples ( $n$ ) of case pack size ( $Q$ ). The target (type 1, i.e., in-stock rate) service level of 98% across SKUs is appropriate as the average optimal service level, calculated by balancing the overage cost to the underage cost. For SKUs in our sample, the optimal service level is 99.3% with a standard deviation of 1.1%. The optimal service level is high in our sample (97.4% of items in our sample have optimal service level over 98%) because products are non-perishable and the review period is relatively short, thus making the cost of overage relatively low.

In each store, SKUs are grouped into fourteen categories, on average. For each category, a manager is responsible for inventory restocking decisions, assisting customers, displaying and organizing products on shelves, restocking products from the backroom, and counting inventory. Assisted by two to four associates, each category manager is responsible for, on average, approximately 2,300 distinct SKUs. Restocking recommendations are triggered by approximately 20% of a set of around 600 SKUs per category reviewed by the system every weekday. A printout of

reviewed SKUs and proposed restocking decisions is delivered to category managers in the morning for review and approval. Category managers may walk through the store to check the accuracy of system information (e.g., inventory levels) or collect other information pertinent to inventory decisions. For each ASO system-reviewed item, managers have three options: accept the order proposal (or lack thereof), or reduce (i.e., contract), or increase (i.e., expand) the proposed order quantity. Finalized inventory decisions and any corrections to inventory data managers might have made are processed in the afternoon and orders placed the following day. According to benchmarking studies by Accenture and AT Kearney, the company's ASO system performs as well as the best-in-class system used in the industry and, not considering managerial interventions, consistently delivers within 1% of the targeted service level.

We assess the effect of managers' modification decisions on inventory costs using two measures, *Cost Savings*, and cost savings as a percentage of cost of goods sold (*Cost Savings Pct*). We measure *Cost Savings* by first estimating inventory costs for the modified proposals based on the difference between the observed outcomes of the replenishment decisions and the counterfactual outcome had the decision not been modified. In our case, we can obtain an accurate estimate of the cost that would have been incurred had the proposal not been modified (the counterfactual outcome), as we have information of the original order recommendation and accurate account of product demand for the decision horizon. Because the inventory policy performs periodic reviews, the decision horizon is defined as the period during which an inventory decision, or its modification, affects inventory cost. This period starts on the day the order from the inventory decision is scheduled to be received and ends on either the day the order from the subsequent decision is delivered or scheduled to be delivered, whichever comes first.

Our cost function is defined by the objectives of the ASO system and the constraints under which it operates. Namely, maintaining a required service level while satisfying a minimum display quantity and re-stocking in multiples of the logistical case pack. Specifically, the *Cost Savings* measure is operationalized by considering two sources of inventory cost, the cost of excess inventory and cost of lost sales. The cost of excess inventory is the cost of carrying greater than the



required inventory level during the decision horizon. Specifically, since inventory is necessary for sales in retail (Balakrishnan et al. 2004, Krommyda et al. 2015, Stavroulaki 2011), it is not possible to unilaterally say that zero inventory is better. Thus, we consider as required inventory ( $RI$ ) the largest of the sales, marketing, or logistic requirements (i.e., 98% fill rate, minimum display quantity, and case pack size, respectively), and count as excess inventory ( $EI$ ) any inventory carried beyond the required inventory. The cost of excess inventory is the cost of holding that inventory over the decision horizon. The equations to calculate the cost of excess inventory for item  $s$  during period  $i$  are:

$$RI_{si} = \text{Max} \left( E[Sales]_{si}^{98th}, \text{Minimum order quantity}, \text{Merchandizing requirement} \right)$$

$$EI_{si} = \text{Max} (Avg(Inv)_{si} - RI_{si}, 0)$$

$$CEI_{si} = EI_{si} * t_{si} * c_s * \iota$$

where  $c$  represents the unit cost,  $t$  the duration, in days, of the decision horizon, and  $\iota$  the daily holding inventory rate. The firm's annual inventory holding rate of 24% for all SKUs is consistent with observations by Graves and Willems (2003). Note that the required inventory threshold creates a zero-cost region in the cost function because we do not penalize inventory up to the system-required inventory.

We estimate the cost of lost sales, incurred when an SKU is out-of-stock (OOS), as average daily demand for the product multiplied by the days it was OOS adjusted assuming a substitution rate ( $\sigma$ ) of 50%, a more conservative estimate than the substitution rate of 45% suggested by Corsten and Gruen (2003). We calculate the cost of lost sales using the following equation:

$$CLS_{si} = \text{Margin}_s * E[Sales]_s^d * \text{Days}_{si}^{(I=0)} * (1 - \sigma)$$

Total inventory cost is the sum of cost of excess inventory and cost of lost sales ( $CEI+CLS$ ). We estimate observed inventory costs based on the observed outcomes of inventory decisions, and

counterfactual inventory costs by calculating the inventory level for the duration of the decision horizon, assuming the original ASO recommendation to have been ordered without modification. For example, for a proposal modified by contracting order quantity by five, the counterfactual starting inventory is five units more than the observed starting inventory. Based on the counterfactual starting inventory level, we calculate the counterfactual days OOS and counterfactual excess inventory using the demand observed during the decision horizon.

*Cost Savings* is defined as counterfactual inventory cost minus observed inventory cost. *Cost Savings Pct* is calculated by dividing *Cost Savings* by the estimated cost of goods sold (COGS) for the SKU during the decision horizon. To develop a reliable measure of COGS and mitigate the problem of division by zero when no sales are observed, we estimate the COGS using average daily sales estimated from the demand history.

Analyzing the effect of managers' modifications of ASO proposals on inventory costs using the *Cost Savings* measure revealed that of the 37% (10,074 out of 26,993) of ASO recommendations that were modified, 54.3% generate savings and 33.2% incur higher total inventory cost (i.e., negative savings, refer to Table 2.1)<sup>1</sup>. One in eight of the modification decisions have no impact on *Cost Savings*, as they did not change inventory to the point that they incur excess inventory holding costs or lost sales costs. Figure 2.1 presents histograms of the two measures of cost savings, *Cost Savings* in Euros (left panel), and *Cost Savings Pct* (right panel). Most observations lie on the positive side (i.e., savings) and total positive savings are greater than total negative savings. When separating the modifications by their direction, we find that 86% of contractions result in cost savings, i.e., eliminating excess inventory holding cost, and only 3% of the contractions cut inventory so much that they result in lost sales. For expansions, the effect is the opposite. Most expansions (79.3%) are associated with cost increase, due to carrying cost of excess inventory, and only 5.8% result in cost savings where the extra inventory reduced lost sales.


Table 2.2 reports the average *Cost Savings Pct* of the decisions depending on the direction of modifications and cost savings (i.e., the quadrants in Table 2.2 match Table 2.1), and the last

---

<sup>1</sup>Full details of the data available and our sample are provided in §2.5

Table 2.1: Distribution of *Cost Savings* of Modified Proposals

Direction of Modifications	Cost Savings			Total
	+	-	No Change	
Contraction	5,236 86.0%	183 3.0%	668 11.0%	6,087 60.4%
Expansion	232 5.8%	3,163 79.3%	592 14.9%	3,987 39.6%
Total	5,468 54.3%	3,346 33.2%	1,260 12.5%	10,074 100.0%



Note: Modifications that had no effect on *Cost Savings* (i.e., no change) are omitted.

Figure 2.1: Histograms of 5th to 95th percentile of *Cost Savings* of Modified Proposals

column reports the weighted average (by the frequency) of these savings. Both contractions and expansions result in positive average savings — 1.2% and 4.6% of COGS, respectively — and overall modifications yield a savings of 2.5% of COGS in average. Closer inspection reveals that, on one hand, modifications that affect lost sales have high financial impact. Contracting too much, to the point of losing sales, results in an average loss (i.e., negative savings) of 33% of COGS, but expanding to avoid losing sales results in average savings of over 121%. On the other hand, modifications that alter the inventory holding cost translate in savings of 2.6% in the case of contractions, or negative savings of 3.2% when expanding beyond the required inventory. This asymmetric response to the two types of costs explains the positive overall outcomes of both types of decisions. Few expansions with large savings are enough to compensate for many expansions that result in

marginally higher cost, and few contractions with losses are compensated with a high number of contractions that result in cost savings. Despite the non-trivial fraction of modification that result in negative savings, totaled managerial modifications resulted in a 1.6% reduction of COGS of all the total sales in the sample, a significant improvement in an industry where average profit margin is 2% to 3% (Biery 2017, Evans and Mathur 2014, Reagan 2013) and COGS represents 70% of total cost.

Table 2.2: Average *Cost Savings Pct* of Modified Proposals

Direction of Modifications	Cost Savings Pct		Weighted Cost Saving Pct
	+	-	
Contraction	+2.6%	-33.5%	+1.2%
Expansion	+121.4%	-3.2%	+4.6%
Overall	+7.6%	-4.8%	+2.5%

## 2.4 Research design and hypotheses

The above evidence suggest that managers can systematically (consistently and reliably) improve on presumedly optimal restocking recommendations by the ASO system. This is surprising as ASO recommendations are optimized under the defined cost function and human modifications would be expected to deviate from the cost minimizing optimal. Furthermore, the evidence from behavioral decision-making research presented in the literature review suggests that humans are not very good at making these types of decisions.

Clearly, managers can decide to make modifications (a proactive action) if they have access to information arising from events that are not available for the ASO system – the so-called-blinds spots (Oliva and Watson 2009) - or if they are capable of extracting causal inferences that cannot be assessed by the ASO system (Pearl and Mackenzie 2018), or both. However, decision-making being particularly challenging in dynamic and complex environments, it is not possible for us to assess what information was being used by managers to make those modifications. Even if aware of exogenous events that could trigger demand changes, it would not be possible to assess what

information was being *actually* used by managers to make modifications to order recommendations outside a laboratory with controlled environmental stimuli. Instead, we leverage the access to the counterfactual outcomes of the modification decisions to assess the decision characteristics and the context in which managers operate and can successfully improve the ASO proposals. Our hypotheses capture the *attributes of the decision conditions* that would be more prone to be modified and are associated with cost savings. Specifically, our hypotheses, while not testing specific sources of information that managers might be using, test the conditions under which managers successfully complement algorithms — i.e., when is human input most valuable. Empirical identification of these conditions is necessary to spur future research on designing a system that can better incorporate manager’s judgmental decisions into a data-driven decision-making system. In elaborating our hypotheses, we draw on theories from diverse fields such as forecasting, signal processing, and behavioral decision making. Our goal is to explore where managers are capable of leveraging information exogenous to the ASO system — either information not coded into the system or causal inferences not visible to the algorithms — to modify recommendations that result in better inventory decisions.

### **Effect of Uncertainty**

In a high uncertainty environment, historical data are less effective at predicting the future. One explanation for this is that high uncertainty is driven by significant change in the underlying decision-making environment (e.g., shift in demand distribution), as argued in the literature on regime change detection (Barry and Pitz 1979, Kremer et al. 2011, Massey and Wu 2005). In a similar vein, Fildes et al. (2009) and Sanders and Ritzman (1992) argue in their work on judgmental forecasting that high uncertainty in demand is driven more by the effects of special events (i.e., non-random variations) than by statistical noise in the data. Thus, under high uncertainty, historical sales and inventory data on which ASO recommendations are based (Wagner 2002) will be less useful for generating restocking proposals. That is, high uncertainty in restocking proposals is driven by the effects of special events, of which category managers may possess some exogenous information. Hence, managers seeking to improve ASO recommendations are likely to target for

modification restocking proposals with greater uncertainty, which are likely to be associated with the existence of exogenous information.

We further hypothesize that managers with exogenous information will contract order quantity (i.e., reduce the ASO proposed quantity) for proposals under high uncertainty and reduce the inventory costs. In a periodic review inventory system, such as the  $(R, S, nQ)$  policy at the study site, calculation of target inventory quantity depends on uncertainties in model assumptions (Hadley and Whitin 1963). Specifically, high uncertainty leads to larger target inventory (Cachon and Terwiesch 2012). That is, inventory models translate high uncertainty into larger restocking quantities to satisfy the target service level. We expect managers who possess exogenous information under high uncertainty to decrease proposed order quantities because the injection of new information into the decision-making process decreases overall uncertainty with regard to the restocking proposal, alleviating the need to order a large quantity of inventory to cover the target service level. We thus hypothesize as follows.

**Hypothesis 2.1: Order proposals under high uncertainty are positively associated with contraction decisions.**

Moreover, if managers' modifications are based on useful exogenous information, we expect the modifications to reduce inventory costs. We, therefore, propose as a null hypothesis:

**Hypothesis 2.2: Order proposals under high uncertainty are positively associated with cost savings.**

The alternative hypothesis posits that managers' behavioral biases and cognitive limitations would increase inventory cost (i.e., negative savings) from the counterfactual inventory cost (as suggested by the literature presented in the introduction).

### **Effect of Quantized Decision Space**

The decision space in a  $(R, s, nQ)$  inventory policy being quantized into multiples of case pack size ( $Q$ ), both the ASO system and managers are forced to order at multiples of case pack size (see Gray and Neuhoff (1998) for a review of quantization). Given the degree to which quantization of the decision space varies with product, we hypothesize that for products with greater case pack

size, managers are more likely to contract proposed restocking quantity and reduce inventory costs. This hypothesis is suggested by two considerations.

First, greater case pack size correlates with larger quantization errors (i.e., rounding to the next largest quantum), the decision space of restocking orders not being continuous. Consider, for example, a case in which the calculated target inventory is 13 and the case pack size 10. Given the choice of ordering just shy of the calculated target inventory (10) or an excess quantity (20) in order to cover the target inventory level, the system, following the  $(R, S, nQ)$  inventory policy, will always order excess inventory rather than an amount less than the target inventory level. Managers who judge that the modification will not affect performance negatively will override the system and order slightly less than the target inventory level.

Second, products with high case pack size increase the use of temporary backroom storage while awaiting shelving, thereby increasing in-store logistics and operational complexity (Eroglu et al. 2013). Van Donselaar et al. (2010) suggest that due to the active use of backrooms, managers possess exogenous information about in-store logistics and shelf space availability for SKUs with high case pack size. Thus, local category managers possessing information about in-store logistics and shelf space availability, can avoid the need to inflate orders and improve performance by contracting restocking proposals for high case pack size SKUs. This suggests the following hypotheses.

**Hypothesis 2.3: Restocking case pack size is positively associated with contraction decisions.**

**Hypothesis 2.4: Restocking case pack size is positively associated with cost savings.**

#### **Effect of Preference for the Salient**

Perfectly rational managers, possessing “skill in computation that enables [them]... to reach the highest attainable point,” would search for the global optimum (Simon 1955, p. 99), leading to a prediction that they would modify every ASO decision for which they have relevant exogenous information, no matter how small the resulting improvement. Real world managers, however, exhibit bounded computational capacity and limited attention. Decision-makers compensate for limited attention under complex decisions with large choice sets by focusing their attention on

choices to which they are attracted. This phenomenon, termed preference for the salient (Dellavigna 2009, Tversky and Kahneman 1974), is found in individual investment (Barber and Odean 2008, Bordalo et al. 2012) and resource consumption (Tiefenbeck et al. 2018) decisions, among many other decision-making contexts.

A typical retail store having thousands of SKUs, it is difficult to imagine a manager keeping track of information related to each product. Managers are instead incentivized to focus attention on a subset of financially important products (DeHoratius and Raman 2008). Most retail stores use the Pareto principle (80/20 rule) or ABC classification to identify products that significantly affect financial performance (Teunter et al. 2010), and tie sales performance of these high impact products to managers' incentives (Van Donselaar et al. 2010). Inferring these products to be more salient because of their importance, we argue that restocking proposals for them are more salient and managers consequently likely to possess, or strive to acquire, better information about them. We therefore hypothesize managers' decisions about restocking proposals for more salient (i.e., high impact) products to be well-informed and resulting modifications to yield cost savings. The following hypotheses are thus suggested.

**Hypothesis 2.5: Saliency of proposals is positively associated with modifications.**

**Hypothesis 2.6: Saliency of proposals is positively associated with cost savings.**

#### **Effect of Common Factors**

In addition to information that affects managers' inventory decisions at the individual decision level we consider exogenous events that affect decision-making at the category or store level, or both. Siemsen et al. (2009) argue that confidence in knowledge leads to more knowledge sharing, even in organizational environments unfavorable to sharing knowledge (i.e., organization with low psychological safety (Edmondson 1999, Kahn 1990)). We thus expect managers with exogenous information that may affect store performance, especially if they are confident about the information, to share it with other store managers.

An unexpected natural disaster, for example, can affect a category of products across multiple stores, as when demand for snow blowers and de-icing materials skyrocketed in areas affected



by a blizzard that struck the Northeastern United States in December 2010 (Dodes 2010). A high rate of modifications of ASO proposals among category managers responsible for restocking snow gear would thus be expected, as the event significantly changed the underlying decision-making environment (e.g., shift in demand distribution). We thus hypothesize that more cost-saving modifications based on information about an event that affects multiple category managers would be observed on the day on which the modification rate of managers in charge of the same category (but in different stores) is high.

In the case of category managers within a given store responding to an event that affects the entire store (e.g., a large sporting event near the store (Gratton et al. 2006, Wilson 2006)), we would expect high modification rates resulting from managers within the store interacting and sharing information that may affect their inventory decisions (Argote and Ingram 2000). We hypothesize that more cost-saving modifications based on information about an event that affects an entire store would be observed on the day on which the modification rate of managers within the store is high. The following hypotheses are thus suggested.

**Hypothesis 2.7: The proposal modification fraction of managers in charge of the same category is positively associated with cost savings.**

**Hypothesis 2.8: The proposal modification fraction of inventory managers within a store is positively associated with cost savings.**

Hypotheses 2.1, 2.3, and 2.5, by hypothesizing the conditions for and direction of modifications, indirectly posit access to useful information exogenous to the system. Hypotheses 2.2, 2.4, 2.6, 2.7, and 2.8 assess managers' effectiveness at converting useful information into modification decisions that lead to inventory cost savings. The alternative hypotheses are that managers' judgmental decisions are prone to behavioral biases and their cognitive limitations increase inventory costs. Note, however, that we do not test for specific information sources, but rather use the empirical evidence to assess whether managers might have something useful to contribute to the decision-making process and under which conditions their contribution might be more useful. By assessing performance against the counterfactual outcomes, our empirical approach effectively assesses the

combined effects of managers' cognitive limitations and biases, *and* the potential benefits of including managers capable of detecting exogenous information useful to the decision-making task. This approach departs from the current paradigm in behavioral operations that focuses on isolating biases and attempting to explain managers' decisions by explaining small individual parts (Carter et al. 2007, Donohue et al. 2020, Gino and Pisano 2008, Tversky and Kahneman 1974).

## **2.5 Data Sources and Variables**

Our sample includes observations of ASO proposed order quantities and final inventory decisions made by category managers as well as information on order shipments from suppliers and sales data for each SKU in each store. We observe inventory decisions made by seven category managers in charge of replenishment decisions for the same category of products in seven different stores. We observe inventory decisions for 4,096 unique SKUs over a period of eight weeks, yielding an overall sample of 26,993 observations. On average, each store handles 1,605 SKUs in the relevant category and we have 3,856 decisions per manager. Our sample includes a total of 56 distinct suppliers across stores with an average of 44 suppliers per store. Structuring the sample as a repeated measure helps us investigate the effect of managers' decisions on performance (Curran and Bauer 2011). Our data do not include SKUs that were reviewed but did not trigger an order recommendation by the ASO system or manager.

### **2.5.1 Variables**

We employ five dependent variables. The first, *Modified*, is an indicator variable that classifies restocking proposals as modified by managers. We also identify with indicators (i.e., binary variables) whether modifications *Contracted* (the manager ordered less than) or *Expanded* (the manager ordered more than) the ASO recommendation. The remaining dependent variables, *Cost Savings* and *Cost Savings Pct*, which measure the effect of managers' modification decisions on inventory costs, were described in §2.3.

To measure uncertainties associated with order proposals, we use as independent variables the following four proxies, *Demand Volatility*, *Agreed Lead Time*, *Supplier Lateness*, and *Supplier Re-*

*liability*. The first two variables, *Demand Volatility* and *Agreed Lead Time*, measure the uncertainty associated with demand forecasts. *Demand Volatility* is represented by the coefficient of variation of the three most recent monthly sales figures (Fildes et al. 2009, Sanders and Ritzman 1992). The *Agreed Lead Time* promised by the supplier represents another source of uncertainty because the errors in demand forecasts increase as forecasts are made further away in the future (Prak et al. 2016). *Supplier Lateness* and *Supplier Reliability*, which measure uncertainty associated with suppliers, follow measures used in Shin et al. (2000). *Supplier Lateness* captures suppliers' average realized lead time relative to agreed lead times, and *Supplier Reliability* the standard deviation of the distribution of suppliers' realized lead times compared to agreed lead times. We calculate the latter variables by subtracting agreed lead time from realized lead time (a negative value represents early delivery) and dividing by agreed lead time for every restocking order observed in the sample. We calculate *Supplier Lateness* by averaging this value for each supplier-manager pair, and *Supplier Reliability* as the standard deviation of this value by each supplier-manager pair. We use the following equations, in which the indexes represent manager ( $i = 1 \dots I$ ), supplier ( $s = 1 \dots S$ ), and time ( $t = 1 \dots T$ ).

$$Supplier\ Lateness_{is} = Mean\left(\frac{Realized\ Lead\ Time_{ist} - Agreed\ Lead\ Time_{ist}}{Agreed\ Lead\ Time_{ist}}\right), \forall i, s$$

$$Supplier\ Reliability_{is} = StDev\left(\frac{Realized\ Lead\ Time_{ist} - Agreed\ Lead\ Time_{ist}}{Agreed\ Lead\ Time_{ist}}\right), \forall i, s$$

We measure the level of quantization of the decision space using the *Case Pack Size* (i.e. quantity of SKUs in a single pack) mandated by suppliers, following Van Donselaar et al. (2010) and Eroglu et al. (2013).

We use two variables as a measure of saliency (DeHoratius and Raman 2008), *Top 20% Margin* and *Top 20% Dollar Volume*. *Top 20% Margin* is an indicator variable for SKUs classified as being in the top 20th percentile of margin, and *Top 20% Dollar Volume* an indicator variable for SKUs classified as being in the top 20th percentile of sales dollar volume. We calculate sales dollar volume similar to DeHoratius and Raman (2008), by multiplying SKU sales price and sales over

the past three months. We classify restocking proposals in the top 20th percentile of margin or top 20th percentile of dollar volume as more salient for managers, following the Pareto principle that 20% of products drive 80% of performance. Corroborating interviews with company managers revealed that they are encouraged to pay special attention to the 20% of products that are the greatest drivers of performance.

We use two variables, *Category Modification Fraction* and *Store Modification Fraction*, as a measure of category and store level common factors, respectively. *Category Modification Fraction* assesses, for each of the managers in our sample, what the other six managers in the sample did — each with identical responsibilities and category but in different stores. Specifically, we compute the modification fraction of category managers by dividing the total number of proposals modified by category managers in charge of the same category across the other six stores, by the total number of restocking proposals observed by these managers. *Store Modification Fraction* assesses, for each of the managers in our sample, what other eight managers with identical responsibilities but in different categories in their corresponding stores did. We use only eight categories, as opposed to the fourteen total categories in the store, because the excluded six categories have structurally different interactions with the customers or suppliers. For example, lumber is not directly accessible to the customers in shelves but is ordered and delivered through store personnel. We calculate *Store Modification Fraction* by dividing the total number of proposals modified by the other category managers in a store, by the total number of restocking proposals observed by these managers.

We control for heterogeneity of proposals at the manager and SKU level and time period using the following four variables, *Number of Decisions*, *OOS*, *Subcategory*, and *Day of Week*. *Number of Decisions*, which controls for the effect of number of proposals (in a single day) on managers' modification decisions, is calculated by counting the number of restocking decisions proposed to a manager in a specific day. We control for the effect of stock outs on managers' modification decisions using the indicator variable *OOS* to identify proposals with zero stock at the time of the decision. The categorical variable *Subcategory* is used to control for heterogeneity in products. SKUs

in our category are further grouped into nine subcategories of products with similar attributes. The *Day of Week* variable is treated similarly. We account for heterogeneity of modification decisions among managers (see Table 2.3) by controlling for time-invariant manager characteristics with a manager fixed effect regressor. We do not include mean demand and inventory position as control variables because they are highly correlated with other independent variables (e.g., *Top 20% Dollar Volume*). Summary statistics and the correlation matrix for the variables are reported in Table 2.4 and Table 2.5.

Table 2.3: Frequency of Modification Decisions by Managers

	No Modification	Contraction	Expansion	Total
Manager 1	1,679 31.7%	2,371 44.7%	1,249 23.6%	5,299 100%
Manager 2	3,002 49.3%	1,761 28.9%	1,328 21.8%	6,091 100%
Manager 3	2,209 55.3%	952 23.8%	833 20.9%	3,994 100%
Manager 4	1,991 76.8%	426 16.4%	176 6.8%	2,593 100%
Manager 5	2,224 82.6%	318 11.8%	152 5.6%	2,694 100%
Manager 6	2,442 86.8%	217 7.7%	156 5.5%	2,815 100%
Manager 7	3,372 96.2%	42 1.2%	93 2.6%	3,507 100%
Overall	16,919 62.7%	6,087 22.5%	3,987 14.8%	26,993 100%

Table 2.4: Summary Statistics

Variable	# of Obs.	Mean	Std. Dev.	Min	Max
Modified	26,993	0.37	0.48	0.00	1.00
Cost Savings	10,074	0.22	2.39	-22.82	109.21
Demand Volatility	26,993	0.62	0.43	0.00	1.73
Agreed Leadtime	26,993	15.18	5.29	1.00	39.00
Supplier Lateness	26,993	-0.02	0.24	-2.32	0.65
Supplier Reliability	26,993	0.15	0.22	0.00	2.34
Case Pack Size	26,993	5.15	8.22	1.00	200.00
Top 20% Margin	26,993	0.19	0.39	0.00	1.00
Top 20% Dollar Volume	26,993	0.20	0.40	0.00	1.00
Category Modification Fraction	26,993	0.42	0.17	0.00	0.91
Store Modification Fraction	26,993	0.36	0.16	0.00	0.96
Number of Decisions	26,993	160.81	77.64	1.00	424.00
OOS	26,993	0.03	0.17	0.00	1.00

Table 2.5: Correlation Matrix

Variable	1	2	3	4	5	6	7	8	9	10	11	12
1 Cost Savings	1.000											
2 Demand Volatility	0.007	1.000										
3 Agreed Leadtime	0.006	-0.059***	1.000									
4 Supplier Lateness	0.005	0.023***	0.054***	1.000								
5 Supplier Reliability	0.011	0.013*	-0.507***	-0.385***	1.000							
6 Case Pack Size	0.008	-0.050***	0.006	-0.124***	0.003	1.000						
7 Top 20% Margin	0.065***	0.200***	0.010	-0.033***	0.058***	-0.202***	1.000					
8 Top 20% Dollar Volume	0.069***	-0.241***	-0.020**	-0.023***	0.076***	0.058***	0.118***	1.000				
9 Category Modification Fraction	-0.003	0.001	-0.052***	-0.057***	0.037***	-0.018**	-0.005	-0.021***	1.000			
10 Store Modification Fraction	-0.007	-0.035***	-0.089***	-0.041***	0.025***	0.012*	-0.015*	0.048***	0.304***	1.000		
11 Number of Decisions	-0.015	-0.077***	0.180***	-0.007	-0.071***	-0.109***	-0.081***	-0.079***	-0.155***	-0.190***	1.000	
12 OOS	0.048***	0.064***	-0.028***	0.011	0.004	-0.021***	0.056***	-0.010	0.010	-0.007	-0.034***	1.000

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ , significance level of correlation coefficients

## 2.6 Empirical Models

Managers are allowed to modify each ASO proposal by either contracting or expanding the proposed order quantity. Managers' inventory decisions are categorized into three mutually exclusive choices: Contraction, Expansion, and No Modification. Fitting a multinomial logit model using the three choices as the dependent variable is akin to fitting a logistic regression to all pairs of choices. For ease of testing our hypothesis, we set the choice of *No Modification* as the basis for comparison (i.e., the base case in the model).

In Equation 2.1,  $Y_{ijt}$  is a categorical variable that denotes the three mutually exclusive choices (contraction, expansion, or no modification) manager  $i$  ( $i = 1 \dots I$ ) can make for SKU  $j$  ( $j = 1 \dots J$ ) at time  $t$  ( $t = 1 \dots T$ ). The fraction on the left-hand side of Equations 2.1a and 2.1b denotes the two choices being compared. The choice on the denominator (i.e., No Modification) refers to the base case, the choice on the numerator to the choice (i.e., Contraction or Expansion) being compared to the base case. Note that in our case, comparison of the likelihood of a proposal being contracted to the likelihood of a proposal being expanded can be calculated by taking the difference between the sets of estimated coefficients in Equations 2.1a and 2.1b. Time  $t$  being a daily level index and each manager-SKU pair being reviewed for replenishment once a week (on average), our data structure is unbalanced panel data.

$$\begin{aligned}
 \frac{\Pr(Y_{ijt} = \textit{Contraction})}{\Pr(Y_{ijt} = \textit{No Modification})} &= \exp(\alpha_i + \alpha_1(\textit{Demand Volatility})_{ijt} \\
 &+ \alpha_2(\textit{Agreed Lead Time})_{ijt} + \alpha_3(\textit{Supplier Lateness})_{ij} + \alpha_4(\textit{Supplier Reliability})_{ij} \\
 &+ \alpha_5(\textit{Case Pack Size})_j + \alpha_6(\textit{Top 20\% Margin})_j + \alpha_7(\textit{Top 20\% Dollar Volume})_j \\
 &+ \alpha_8(\textit{Category Modification Fraction})_{it} + \alpha_9(\textit{Store Modification Fraction})_{it} \\
 &+ \alpha_{10}(\textit{OOS})_{ijt} + \alpha_{11}(\textit{Number of Decisions})_{it} + \alpha_{12}(\textit{Day of Week})_t \\
 &+ \alpha_{13}(\textit{Subcategory})_j)
 \end{aligned} \tag{2.1a}$$

$$\begin{aligned}
\frac{\Pr(Y_{ijt} = \textit{Expansion})}{\Pr(Y_{ijt} = \textit{No Modification})} &= \exp(\beta_i + \beta_1(\textit{Demand Volatility})_{ijt} \\
&+ \beta_2(\textit{Agreed Lead Time})_{ijt} + \beta_3(\textit{Supplier Lateness})_{ij} + \beta_4(\textit{Supplier Reliability})_{ij} \\
&+ \beta_5(\textit{Case Pack Size})_j + \beta_6(\textit{Top 20\% Margin})_j + \beta_7(\textit{Top 20\% Dollar Volume})_j \\
&+ \beta_8(\textit{Category Modification Fraction})_{it} + \beta_9(\textit{Store Modification Fraction})_{it} \\
&+ \beta_{10}(\textit{OOS})_{ijt} + \beta_{11}(\textit{Number of Decisions})_{it} + \beta_{12}(\textit{Day of Week})_t \\
&+ \beta_{13}(\textit{Subcategory})_j)
\end{aligned} \tag{2.1b}$$

To test hypotheses 2.1 and 2.3, we analyze the statistical significance of the estimated coefficients of Equation 2.1a, and hypothesis 2.5 by assessing the coefficients in Equations 2.1a and 2.1b.

Hypotheses 2.2, 2.4, 2.6, 2.7, and 2.8 require a statistical test for the association between our independent variables and cost savings from modifications. Because we observe the change in inventory cost caused by the modification only when managers choose to modify proposals, observations of cost savings from modifications are selected via a systematic process that may result in sample selection bias. We account for sample selection bias using the sample selection model developed by Heckman (1976, 1979), which serves three purposes. First, our empirical model tests hypotheses by comparing the observed outcomes to the counterfactuals. Our approach follows the potential outcome model or Rubin causal model (Holland 1986), thereby establishing the theoretical and statistical basis in our empirical model to test for causal inference. Note that comparing observations to counterfactuals is essentially what is achieved through randomized control trials (Angrist and Pischke 2009). Second, our sample selection model explicitly models exogenous information (i.e., private information in sample selection model terminology) and tests for it, which is the focus of our research. Specifically, testing for selection bias equates to the test of existence of exogenous information (although we cannot pin-point what information was used) because systematic selection, if significant, represents the “estimate of private information underlying-



ing" a choice and how it explains outcomes (Li and Prabhala 2007, pp. 44-45). Finally, the sample selection model accounts for omitted variable bias or endogeneity that arise from the manager's modifications. Specifically, since our independent variables capture the attributes of the decision conditions that may be associated with inventory cost improvements, these variables are correlated with exogenous information that drives the manager's modification – an example of endogeneity emanating from omitted variable bias. Our selection model is specified as follows:

$$\begin{aligned}
Modified_{ijt}^* &= \gamma_i + \gamma_1(Demand\ Volatility)_{ijt} \\
&+ \gamma_2(Agreed\ Lead\ Time)_{ijt} + \gamma_3(Supplier\ Lateness)_{ij} + \gamma_4(Supplier\ Reliability)_{ij} \\
&+ \gamma_5(Case\ Pack\ Size)_j + \gamma_6(Top\ 20\%\ Margin)_j + \gamma_7(Top\ 20\%\ Dollar\ Volume)_j \\
&+ \gamma_8(Category\ Modification\ Fraction)_{it} + \gamma_9(Store\ Modification\ Fraction)_{it} \\
&+ \gamma_{10}(Number\ of\ Decisions)_{it} + \gamma_{11}(OOS)_{ijt} + \gamma_{12}(Subcategory)_j \\
&+ \gamma_{13}(Day\ of\ Week)_t + \epsilon_{ijt}^1
\end{aligned} \tag{2.2a}$$

$$\begin{aligned}
Cost\ Savings_{ijt}^* &= \delta_i + \delta_1(Demand\ Volatility)_{ijt} \\
&+ \delta_2(Agreed\ Lead\ Time)_{ijt} + \delta_3(Supplier\ Lateness)_{ij} + \delta_4(Supplier\ Reliability)_{ij} \\
&+ \delta_5(Case\ Pack\ Size)_j + \delta_6(Top\ 20\%\ Margin)_j + \delta_7(Top\ 20\%\ Dollar\ Volume)_j \\
&+ \delta_8(Category\ Modification\ Fraction)_{it} + \delta_9(Store\ Modification\ Fraction)_{it} \\
&+ \delta_{10}(OOS)_{ijt} + \delta_{11}(Subcategory)_j \\
&+ \delta_{12}(Day\ of\ Week)_t + \epsilon_{ijt}^2
\end{aligned} \tag{2.2b}$$

$$Cost\ Savings_{ijt} = \begin{cases} Cost\ Savings_{ijt}^* & \text{if } Modified_{ijt}^* > 0 \\ 0 & \text{if } Modified_{ijt}^* \leq 0 \end{cases} \tag{2.2c}$$

where,  $\epsilon_{ijt}^1$  and  $\epsilon_{ijt}^2$  have a bivariate normal distribution and  $corr\epsilon_{ijt}^1, \epsilon_{ijt}^2 = \rho$

Formulation of the Heckman sample selection model in Equation 2.2 is based on a latent variable model formulation (Puhani 2000). We use two latent variables,  $Modified_{ijt}^*$  and  $Cost Savings_{ijt}^*$ , and one observed variable,  $Cost Savings_{ijt}$ , to formulate the sample selection process. Equation 2.2a is the selection equation that estimates the propensity for a proposal to be modified, Equation 2.2b the outcome equation that estimates the association between the independent variables and cost savings. Equation 2.2c expresses the fact that we observe cost savings from modifications only when restocking proposals are modified. We can test for hypotheses 2.2, 2.4, 2.6, 2.7, and 2.8 by analyzing the statistical significance of the estimated coefficients in Equations 2.2a and 2.2b.

Our analysis estimates the parameters in Equations 2.2a, 2.2b, and 2.2c together using the full-information maximum likelihood estimator (FIML) (refer to the likelihood function in Heckman (1976) and Puhani (2000)). We use the FIML estimator instead of Heckman's (1979) two stage estimator because it is more efficient (Heckman 1976, Puhani 2000). We use the number of restocking decisions proposed to a manager in a day as our exclusion restriction. The number of restocking decisions proposed in a day satisfies the requirement for an exclusion restriction by influencing the probability of modification but being uncorrelated with the cost savings generated by the modifications. Managers have limited processing capability, thus the probability of modification of an individual decision diminishes as the number of decisions increases. Also, managers intervene only in order proposals they can improve with the information available to them. The number of decisions proposed in a day is thus not correlated with cost savings.

We account for time-invariant manager characteristics by including manager fixed effects in both the multinomial logit model ( $\alpha_i$  and  $\beta_i$  in Equation 2.1) and Heckman sample selection model ( $\gamma_i$  and  $\delta_i$  in Equation 2.2), and account for potential heteroscedasticity by estimating standard errors using robust standard errors.

## 2.7 Results

Table 2.6 reports the estimation of the multinomial logistic regressions to test hypotheses 2.1, 2.3, and 2.5. Model 1 reports the estimation with only the control variables including fixed ef-

fects for *Subcategory*, *Day of Week*, and managers, Model 2 with the addition of variables related to hypotheses 2.1 and 2.3, and Model 3 the full empirical model as stated in Equation 2.1. Although we formulate no explicit hypotheses about their influence on the likelihood of contraction or expansion, we include the category and store modification fractions to be consistent with the full empirical model used in the Heckman sample selection model. Further, if hypotheses 2.7 and 2.8 are true, omitting these two variables, *Category Modification Fraction* and *Store Modification Fraction*, in the multinomial logistic regression would result in omitted variable bias.

The first column in Model 3, column (5) with the dependent variable *Contraction*, reports the estimated coefficients that explain managers' likelihood of contracting order proposals compared to the likelihood of making no modifications (Equation 2.1a), the second column in Model 3, column (6), the estimated coefficients that explain managers' likelihood of making expansion decisions compared to the likelihood of making no modifications (Equation 2.1b).

Analyzing the coefficients in Model 3, we found all four measures of uncertainty in order proposals, *Demand Volatility* ( $b=0.2557$ ,  $p<0.001$ ), *Agreed Lead Time* ( $b=0.0316$ ,  $p<0.001$ ), *Supplier Lateness* ( $b=0.4978$ ,  $p<0.001$ ), and *Supplier Reliability* ( $b=0.3433$ ,  $p=0.002$ ), to be positively and significantly correlated with the likelihood of a manager contracting an order proposal (hypothesis 2.1). The coefficient of *Case Pack Size* was positive and significant ( $b=0.0083$ ,  $p=0.001$ ), providing evidence that higher *Case Pack Size* is associated with a greater likelihood of contraction decisions (hypothesis 2.3). Hypothesis 2.5 predicted a positive relationship between saliency of proposals and modification decisions, and both measures of saliency, *Top 20% Margin* ( $b=0.1570$ ,  $p=0.006$ ) and *Top 20% Dollar Volume* ( $b=0.1283$ ,  $p=0.006$ ), were positively and significantly correlated with likelihood of contraction decisions.

In addition to the tests for hypotheses 2.1, 2.3, and 2.5, we found *Category Modification Fraction*, a variable for hypothesis 2.7, to be positively and significantly correlated with likelihood of expansion decisions ( $b=1.0379$ ,  $p<0.001$ ). *Store Modification Fraction*, a variable for hypothesis 2.8, was positively and significantly correlated with likelihood of contraction ( $b=3.2846$ ,  $p<0.001$ ) and likelihood of expansion ( $b=2.9187$ ,  $p<0.001$ ). *Category Modification Fraction* and *Store Mod-*

Table 2.6: Estimation of Multinomial Logistic Regressions

Dependent Variable	Model 1: Multinomial Logit		Model 2: Multinomial Logit		Model 3: Multinomial Logit	
	(1)	(2)	(3)	(4)	(5)	(6)
	Base case: No Modifications		Base case: No Modifications		Base case: No Modifications	
	Contraction	Expansion	Contraction	Expansion	Contraction	Expansion
Demand Volatility			0.2273*** (0.0403)	0.3572*** (0.0458)	0.2557*** (0.0427)	0.2993*** (0.0482)
Agreed Leadtime			0.0294*** (0.0038)	0.0025 (0.0048)	0.0316*** (0.0040)	0.0084 (0.0051)
Supplier Lateness			0.4162*** (0.0784)	-0.1792* (0.0901)	0.4978*** (0.0815)	-0.1177 (0.0934)
Supplier Reliability			0.2207* (0.1061)	-0.1682 (0.1240)	0.3433** (0.1101)	0.0025 (0.1266)
Case Pack Size			0.0059* (0.0025)	0.0087*** (0.0023)	0.0083** (0.0026)	0.0091*** (0.0024)
Top 20% Margin					0.1570** (0.0576)	0.1069 (0.0673)
Top 20% Dollar Volume					0.1283** (0.0463)	-0.2354*** (0.0549)
Category Modification Fraction					0.1604 (0.1379)	1.0379*** (0.1451)
Store Modification Fraction					3.2846*** (0.1562)	2.9187*** (0.1545)
Number of Decisions	-0.0031*** (0.0003)	-0.0027*** (0.0003)	-0.0035*** (0.0003)	-0.0027*** (0.0003)	-0.0020*** (0.0003)	-0.0011*** (0.0003)
OOS	-1.0592*** (0.1412)	0.4586*** (0.1032)	-1.0856*** (0.1422)	0.4325*** (0.1039)	-1.0889*** (0.1454)	0.4293*** (0.1059)
Subcategory FE	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes	Yes	Yes
Total observations	26,993		26,993		26,993	
Mcfadden's Pseudo R2	0.1542		0.1580		0.1770	
Log likelihood	-20803		-20709		-20242	

Robust standard errors in parentheses \* p<0.05 \*\* p<0.01 \*\*\* p<0.001

\*Note: The variables category and store modification fraction were included to be consistent with the full empirical model used in the Heckman sample selection model (i.e., Equation 2.2 and Model 6 in Table 2.7).

*ification Fraction* are thus associated with systematic modifications, which is partial evidence for hypotheses 2.7 and 2.8. We formally test hypotheses 2.7 and 2.8, whether systematic modifications lead to cost savings, using the Heckman sample selection model.

Hypotheses 2.2, 2.4, 2.7, and 2.8 predicted a positive association between our independent variables and inventory cost savings from modifications. We tested our hypotheses by estimating the Heckman sample selection model (Equation 2.2). Results are reported in Table 2.7. We first included the control variables (Model 4), then the main effects (Models 5 and 6). For Models

4, 5, and 6, the columns with the dependent variable *Modified* represent the coefficients from the selection equation (Equation 2.2a), the columns with the dependent variable *Cost Savings* the coefficients from the outcome equation (Equation 2.2b). Model 6, the main model of our research, represents the full empirical model as stated in equation 2.2. Model 7 reports the estimation of the truncated OLS regression, which is a regression of the outcome equation (Equation 2.2b) on the sample of modified proposals.

We evaluated whether use of the Heckman selection model was justified by testing for the existence of selection bias in our sample. The result from the Wald test rejected the null hypothesis of no selection bias ( $p < 0.0001$ ), suggesting that linear regression without accounting for selection bias would yield biased estimates. That is, unobserved exogenous information underlying a manager's decision to modify influences performance outcomes. Our analysis found that exogenous information used by managers has a positive effect (i.e., positive and statistically significant lambda with 95% confidence interval between 2.0121 and 3.4858) on the decision to modify as well as savings in inventory costs. We evaluated the direction of the effect of selection bias by comparing the coefficients of the main model (column (6) in Table 2.7) to the coefficients of the regression done on the sample of modified proposals (i.e., truncated OLS, Model 7 in Table 2.7). Our finding that all but two coefficients (i.e., *Top 20% Margin* and *Top 20% Dollar Volume*) in our main model have larger effect sizes than the coefficients in the truncated OLS regression suggests that failing to account for selection bias leads to underestimation of the coefficients, which might result in failure to detect a significant relationship (i.e., a type II error). Accounting for ASO proposals not modified by managers by estimating the propensity to modify thus reveals the true unbiased relationship between our independent variables and *Cost Savings*.

To assess the statistical fit of the main model (Model 6), we examined the strength of our exclusion restriction variable, *Number of Decisions*. This variable shows a statistically significant coefficient ( $b = -0.0003$   $p = 0.001$ ) in the selection equation, column (5) of Table 2.7 (Puhani 2000). Additionally, we found the McFadden's pseudo R<sup>2</sup> value of 0.23 using probit regression on the selection equation; correlations below 0.19 were found between all independent variables and the

Table 2.7: Estimation of Heckman Sample Selection Model

Dependent Variable	Model 4: Heckman Selection Model		Model 5: Heckman Selection Model		Model 6: Heckman Selection Model		Model 7: Truncated OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Modified (Yes/No)	Cost Savings	Modified (Yes/No)	Cost Savings	Modified (Yes/No)	Cost Savings	Cost Savings
Demand Volatility			0.1264*** (0.0200)	0.3159*** (0.0761)	0.1293*** (0.0203)	0.3109*** (0.0762)	0.0553 (0.0615)
Agreed Leadtime			0.0068*** (0.0020)	0.0177** (0.0061)	0.0107*** (0.0021)	0.0270*** (0.0066)	0.0002 (0.0059)
Supplier Lateness			0.0398 (0.0403)	0.2011 (0.1167)	0.0749 (0.0412)	0.2980* (0.1160)	0.0298 (0.1371)
Supplier Reliability			0.0214 (0.0536)	0.2337 (0.1529)	0.1015 (0.0549)	0.4187** (0.1526)	0.1250 (0.1690)
Case Pack Size			0.0043*** (0.0010)	0.0097*** (0.0028)	0.0043*** (0.0011)	0.0095*** (0.0028)	0.0000 (0.0027)
Top 20% Margin			0.1374*** (0.0286)	0.3074** (0.1000)	0.1416*** (0.0291)	0.3178** (0.1011)	0.3522*** (0.1070)
Top 20% Dollar Volume			0.0745** (0.0241)	0.1920* (0.0779)	0.0745** (0.0244)	0.1890* (0.0781)	0.3221*** (0.0866)
Category Modification Fraction					0.3175*** (0.0668)	0.7103** (0.2250)	0.1991 (0.2373)
Store Modification Fraction					1.6491*** (0.0803)	3.1139*** (0.4813)	-0.3903 (0.2436)
Number of Decisions	-0.0018*** (0.0001)		-0.0007*** (0.0001)		-0.0003** (0.0001)		
OOS	-0.1209* (0.0543)	0.7721 (0.5712)	0.0822 (0.0874)	-0.0199 (0.3621)	0.0900 (0.0883)	-0.0029 (0.3656)	0.7520 (0.5628)
Subcategory FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Rho	(95% C.I.)	(-0.0443, 0.0112)		(0.9104, 0.9670)		(0.9088, 0.9665)	
Sigma	(95% C.I.)	(1.8247, 3.0992)		(2.2928, 3.37332)		(2.2787, 3.7172)	
Lambda	(95% C.I.)	(-0.1073, 0.0284)		(2.0287, 3.5032)		(2.0121, 3.4858)	
Total observations		26,993		26,993		26,993	10,074
Log likelihood		-37321		-34412		-34040	
Wald test of independent equations		0.2414		0.0000		0.0000	
R2							0.0142

Robust standard errors in parentheses \* p<0.05 \*\* p<0.01 \*\*\* p<0.001

inverse Mills ratio. Compared to the guidelines derived by Certo et al. (2016), our estimates show desirable characteristics of a high pseudo R2 value and low correlation between the inverse Mills ratio and independent variables (Bushway et al. 2007, Leung and Yu 1996). Further, the standard error of coefficients in the outcome equation (Equation 2.2a) decreases as the strength of the exclusion restriction increases. That is, a weak exclusion restriction would still yield unbiased esti-

mates, albeit with larger standard errors (Certo et al. 2016). Our model thus provides conservative estimates of the true impact of managers' modifications on Cost Savings. With a better exclusion restriction variable, lower standard errors can be achieved, which would strengthen our findings. The foregoing tests give us confidence in the statistical fit of the model and interpretations of model estimates.

Hypothesis 2.2 predicted a positive relationship between uncertainty in order proposals and cost-saving modifications. We test this hypothesis by analyzing the coefficients of the four measures of uncertainty in order proposals reported in Model 6. Two of the four measures, *Demand Volatility* ( $\gamma_1 = 0.1293$ ,  $\delta_1 = 0.3109$ ) and *Agreed Lead Time* ( $\gamma_2 = 0.0107$ ,  $\delta_2 = 0.0270$ ), have positive and significant coefficients in both the selection and outcome equations ( $\gamma$  represents coefficients in column 5 or Equation 2.2a,  $\delta$  coefficients in column 6 or Equation 2.2b). All four coefficients ( $\gamma_1$ ,  $\gamma_2$ ,  $\delta_1$ ,  $\delta_2$ ) have p value less than 0.001 and fail to reject hypothesis 2.2. *Supplier Lateness* ( $\gamma_3 = 0.0749$ ,  $p = 0.069$ ;  $\delta_3 = 0.2980$ ,  $p = 0.010$ ) and *Supplier Reliability* ( $\gamma_4 = 0.1015$ ,  $p = 0.065$ ;  $\delta_4 = 0.4187$ ,  $p = 0.006$ ), the other two measures of uncertainty, have positive and significant coefficients in the outcome equation, but insignificant coefficients in the selection equation, rejecting hypothesis 2.2 under those measures of uncertainty in restocking proposals. Our evidence for hypothesis 2.2 is nuanced by the visibility of the information available to the managers. The first two measures of uncertainty (i.e., *Demand Volatility* and *Agreed Lead Time*) are easily observable by category managers, being directly reported to them for use in revising orders. *Supplier Lateness* and *Supplier Reliability* being subtler metrics of uncertainty, managers have only a generic impression of them based on experience.

The coefficients of *Case Pack Size* in Model 6 were positive and significant in both the selection equation ( $\gamma_5 = 0.0043$ ,  $p < 0.001$ ) and outcome equation ( $\delta_5 = 0.0095$ ,  $p < 0.001$ ). Hence, *Case Pack Size* is associated with more modifications that result in cost savings (hypothesis 2.4).

Hypotheses 2.7 and 2.8 predicted a positive relationship between category level common factors and cost-saving modifications. The coefficients of *Category Modification Fraction* were positive and significant in both the selection equation ( $\alpha_8 = 0.3175$ ,  $p < 0.001$ ) and outcome equation

( $\beta_8 = 0.7103, p = 0.002$ ) as were the coefficients of the *Store Modification Fraction* ( $\alpha_9 = 1.6491, p < 0.001$ ;  $\beta_9 = 3.1139, p < 0.001$ ). Our evidence for hypotheses 2.7 and 2.8 bolsters the evidence from the multinomial logistic regression (Model 3 in Table 2.6), which shows *Category Modification Fraction* and *Store Modification Fraction* to be associated with systematic modifications by demonstrating them to be associated with cost savings.

The coefficients of the measures of saliency, *Top 20% Margin* ( $\alpha_6 = 0.1416, p < 0.001$ ;  $\beta_6 = 0.3178, p = 0.002$ ) and *Top 20% Dollar Volume* ( $\alpha_7 = 0.0745, p = 0.002$ ;  $\beta_7 = 0.1890, p = 0.016$ ), were positive and significant in both the selection and outcome equations. SKUs classified as being in the top 20th percentile of margin or top 20th percentile of dollar volume are thus associated with modifications that result in cost savings. This result, however, is not a true test of hypothesis 2.6, it being possible that price and sales volume are positively correlated with cost savings.

As a test of hypothesis 2.6, we control for the effect of price and sales volume on cost savings by substituting the dependent variable in the main model with *Cost Savings Pct* (i.e., *Cost Savings* standardized by the cost of goods sold during the decision horizon), as reported in Model 8 of Table 2.8. The coefficients of *Top 20% Margin* from this model are positive and significant in both the selection equation ( $b = 0.0927, p = 0.001$ ) and outcome equation ( $b = 0.0344, p = 0.037$ ), the coefficients of *Top 20% Dollar Volume* insignificant in both the selection equation ( $b = -0.0307, p = 0.104$ ) and outcome equation ( $b = 0.0029, p = 0.717$ ). This finding, being consistent with the ABC classification criteria used by our research site, which focuses exclusively on high-margin items, suggests that category managers do indeed pay more attention to SKUs identified as financially important.

### **2.7.1 Robustness Test**

Note that the normalized measure of *Cost Savings* (i.e., *Cost Savings Pct*) represents a robustness test for hypotheses 2.2, 2.4, 2.7, and 2.8 in that it provides an alternative measure for the dependent variable. We found (Model 8 of Table 2.8) statistical significance for the variables representing hypotheses 2.2, 2.7, and 2.8, *Demand Volatility*, *Agreed Lead Time*, *Category Modi-*



Table 2.8: Estimation of Heckman Sample Selection model with Normalized Cost Savings

Dependent Variable	Model 8: Heckman Selection Model	
	(1)	(2)
	Modified (Yes/No)	Cost Savings Pct
Demand Volatility	0.2419*** (0.0216)	0.0680*** (0.0183)
Agreed Leadtime	0.0091*** (0.0021)	0.0038*** (0.0012)
Supplier Lateness	0.0884* (0.0362)	0.0099 (0.0159)
Supplier Reliability	0.1262* (0.0512)	0.0139 (0.0228)
Case Pack Size	0.0048*** (0.0012)	0.0008 (0.0005)
Top 20% Margin	0.0927** (0.0291)	0.0344* (0.0165)
Top 20% Dollar Volume	-0.0307 (0.0189)	0.0029 (0.0081)
Category Modification Fraction	0.3114*** (0.0635)	0.1017** (0.0340)
Store Modification Fraction	1.6021*** (0.0749)	0.5468*** (0.1036)
Number of Decisions	-0.0003*** (0.0001)	
OOS	0.0683 (0.0975)	0.0030 (0.0596)
Subcategory FE	Yes	Yes
Day of Week FE	Yes	Yes
Manager FE	Yes	Yes
Rho	(95% C.I.)	(0.9495, 0.9842)
Sigma	(95% C.I.)	(0.3447, 0.6636)
Lambda	(95% C.I.)	(0.3151, 0.6243)
Total observations		26,993
Log likelihood		-14939
Wald test of independent equations		0.0000

Robust standard errors in parentheses \* p<0.05 \*\* p<0.01 \*\*\* p<0.001

\*Note: The dependent variable in Column (2), *Cost Saving Pct*, is standardized by the cost of goods sold during the decision horizon. The coefficients reported in this table can be compared to those of our main model (Model 6 in Table 2.7).

*modification Fraction*, and *Store Modification Fraction*, similar to the result in the main model (Model 6 of Table 2.7). The *Cost Savings Pct* measure allows for an intuitive interpretation of the effect size of the estimated coefficients based on the percentage of savings relative to COGS. On average, with all other conditions remaining the same, one standard deviation increases in *Demand*

*Volatility*, *Agreed Lead Time*, *Category Modification Fraction*, and *Store Modification Fraction* are associated with respective savings in inventory cost of 2.92%, 2.01%, 1.73%, and 8.75% relative to COGS; a non-trivial amount in an industry with tight profit margins (Biery 2017, Evans and Mathur 2014, Reagan 2013).

As a robustness test for hypothesis 6, we investigated whether managers possessing exogenous information are attracted to ASO proposals that are improvable, that is, whether managers are able to detect ASO proposals that lead to high inventory cost if left unmodified. We used propensity score matching to estimate the differences in the counterfactual inventory cost by comparing modified proposals to unmodified proposals with similar characteristics (i.e., similar propensity score computed using the selection equation). We found modified proposals to have a higher counterfactual inventory cost. Specifically, we found a difference of 0.0192 ( $p < 0.001$ ), which means that managers, on average, choose to modify ASO proposals that, if left unmodified, would lead to 1.92% higher inventory cost as a percentage of COGS. This suggests that managers are able to search for savings opportunities and successfully choose proposals with high counterfactual inventory cost. Our finding provides evidence for hypothesis 6 and suggests that a subset of proposals that lead to high inventory cost has salience for managers.

Further, when comparing the observed inventory costs of modified proposals to unmodified proposals with similar characteristics we find a difference of 0.0049 ( $p = 0.013$ ). This result suggests that managers are good at finding proposals with high counterfactual inventory cost and, through modification, significantly lower inventory cost to a level that matches the average cost of ASO decisions on unmodified proposals.

Lastly, we analyzed the effect of inventory holding rate and substitution rates on the robustness of our results. We computed *Cost Savings* and *Cost Savings Pct* with a combination of varying inventory holding rate ( $\iota$ ) (12% and 36%) and substitution rate ( $\sigma$ ) (25% and 75%) and used them to replicate the main model (Equation 2.2). This sensitivity analysis revealed the statistical significance of the coefficients to be consistent, and robust to varying inventory holding rates and substitution rates (refer to Table A1 and A2 reported in Appendix A).

## 2.8 Discussion and Conclusions

Our research has found that managers can systematically and consistently improve the restocking recommendations made by sophisticated algorithms (e.g., ASO system). Results of selected modifications by managers represent an average savings of 2.5% of cost of goods sold relative to the baseline recommendation. These improvements are possible only because of limitations of the restocking system that generates the recommendations. The system's algorithms provide optimal recommendations based on the information available and constraints in place; that these recommendations can be improved is a signal that the algorithms do not have access to all relevant information, that some constraints limit the solution space, or that managers are capable of making causal inferences from the available information. We identify three areas in which managers consistently improve the focal system's recommendations.

First, the system's ability to address demand uncertainty or supplier unreliability is limited; given a targeted 98% service level, the system's response to uncertainty is to elevate desired inventory levels. Information that reduces that uncertainty or puts it in a broader context allows managers to make contraction decisions that outperform the system recommendations. We see evidence of this behavior when managers consistently show a preference for contracting restocking recommendations in the presence of demand variability and longer agreed lead time. These modifications result in inventory cost savings.

Second, restocking case pack size limits the precision of the system's restocking recommendations. Large case pack sizes result in recommendations that deviate from the optimal due not to a fault in the system's algorithms, but to quantization of the decision space. This lack of finesse in the system recommendations affords managers an opportunity to leverage their understanding of demand patterns and in-store operations to selectively reduce inventory levels for items that deviate from the desired performance. We see evidence of this behavior when managers focus contraction decisions on SKUs with large pack sizes. These modification decisions, too, result in inventory cost savings.

Third, having visibility only to historical demand patterns, the system assumes, like most in-

ventory management systems (Eroglu et al. 2013, Prak et al. 2016), that future demand can be appropriately characterized by the parameters describing the historical distribution (i.e., demand parameters are assumed to be known and persistent over time). Clearly, in a complex and dynamic world these conditions hardly ever hold, and demand patterns exhibit seemingly random deviations from historical demand distributions. Although it is impossible to instantaneously adjust to changes in demand characteristics, humans operating in the same environment can easily detect fundamental shifts to demand patterns and adjust more quickly than an algorithm responding to historical information. We see evidence of managers responding to exogenous information that affects demand distributions when we observe a high correlation in managers' modification patterns in a store and among same-category managers across different stores. That these modifications result in inventory cost savings is evidence that they are systematic responses to information unavailable to the restocking system, as first theorized by Silver (1981) and Wagner (2002).

The observation that an automated stock ordering system has limitations is not surprising (e.g., Silver 1981, Van Donselaar et al. 2010). Oliva and Watson (2009) describe blind spots in decision support systems (i.e., unintentional but systematic errors in the results emerging from the system), which they categorize as being informational, related to not having access to relevant information, and procedural, being shortcomings in the algorithms that support the decision-making process. The first two areas of improvement described above emerge from procedural blind spots, the third area of improvement from informational blind spots<sup>2</sup>. What is surprising is that we find evidence of managers systematically improving system performance based on those blind spots. This finding flies in the face of current practice in most retailers, which do not allow local managers to modify centralized ASO decisions (Felix et al. 2018), and is in line with the argument by Kremer et al. (2011) that algorithms work better in stable, and not so well in unstable environments. We suspect that the practice of not allowing modifications is based on the belief that managers are not capable of identifying and processing all relevant information in an unbiased way rather than a naïve belief

---

<sup>2</sup>Although blind spots are by definition unintentional, in practice, designers often opt not to include information or algorithms due to cost considerations or feasibility constraints. Once design choices are coded into the algorithms, however, system recommendations will suffer from, and users often not be aware of, such blinds spots (Oliva and Watson 2009).

that systems are exempt from blind spots. The existing practice of not allowing local modifications nevertheless deprives organizations of the opportunity to become aware of and address these blind spots.

How is it that managers can achieve consistent improvements in a complex, dynamic environment? Bounded rationality (Simon 1955) suggests that it should not be possible for category managers to accurately process and assess an average of 120 active (i.e., non-zero) restocking recommendations per day. The detail and dynamic complexity of these decisions would be overwhelming even to the most experienced and dedicated inventory manager, never mind a manager with customer service and merchandise display responsibilities. We find evidence of “satisficing” behavior (Simon 1955) in that managers in our sample modify only 37% of active reorder recommendations. We posit that by providing a default choice (i.e., the system’s proposed restocking decision), the system eases managers’ cognitive load, enabling them to focus on important changes that need to be addressed. Our observation that managers’ modifications consistently result in savings indicates that they are allocating their resources adequately (i.e., to recommendations from which they can extract savings), and their preference for high margin items indicates that their allocation of effort is in line with store requirements. By providing a reasonable anchor from which managers may decide to deviate (Epley and Gilovich 2001, Mussweiler and Englich 2005, Mussweiler and Strack 1999, Tversky and Kahneman 1974), the default choice mechanism leverages the advantages of a system generated proposal while eliminating the main driver of algorithm aversion (Dietvorst et al. 2018). The default choice thus functions to enable managers to partially address the detail and dynamic complexity they are facing. Default choice seems a promising alternative approach to existing approaches to integrating algorithms and human decision makers: decomposition of tasks (Lee and Siemsen 2017), and combining independent recommendations via a weighted average (Blattberg and Hoch 1990).

Our research setting enables us to assess the performance of managers responding in real time to a complex, dynamic environment. Although we cannot directly access what is it that managers know or to what they are responding at the time they are making a decision, we can evaluate their

performance based on decision outcomes. We find that managers are capable of selectively responding to external information in real time and subsequently translating it into knowledge useful for decision-making. Although it would be tempting to suggest capturing relevant exogenous information detected by managers as an input to the system, thus formally eliminating identified blind spots, to be able to do so would be contingent on managers' ability to articulate precisely what is it they are responding to and the organizational capability to code appropriate responses to those inputs. Moreover, this solution increases the number of parameters and risk of model overfitting, and there is no guarantee that a more complex algorithm will generalize well in unstable environments and changing circumstances (Chou et al. 2022, Chuang et al. 2021, Pitt et al. 2002). Our results suggest a way to integrate exogenous information available to local managers into the algorithm. Algorithms are efficient (i.e., can make thousands of decisions in seconds), but not flexible (e.g., adaptable) decision makers; managers are sensitive and adaptable to irregular signals, but not efficient when compared to an algorithm (Gigerenzer 2008). We argue that providing a *default choice* enables incorporation of a manager's judgmental decision-making into an ASO system, thereby creating an essentially hybrid system that exploits both the efficiency in decision-making provided by algorithms and flexibility (e.g., adaptability) provided by human decision makers. We believe this line of research and the eventual design of collaborative systems will have an increasing importance as we figure out how to interact the outputs of ever more sophisticated (and less transparent) machine learning and artificial intelligence algorithms (Chuang et al. 2021).

Like any empirical study, our research design has limitations. We believe that the operational constraints of the ASO system at our research site (e.g., periodic reviews, externally set service level, and reordering in multiples of pack size) are representative of industry practice, and insights from our hypotheses hence generalizable to other settings. Nevertheless, our sample is limited to a single firm and a specific system. Further evidence is needed from other sites and alternative algorithms, perhaps with fewer operational constraints. Also, as discussed above, our research relies on the outcome of the decision-making process (i.e., observations of final decisions), not on the thought process that led to those decisions. Ideally, we would like to observe both the

realized decision and specific information the manager used to arrive at it. Perhaps, following Flicker (2019), laboratory experiments that convey specific exogenous information to subjects and evaluate its effect on the outcomes of decisions might be a fruitful avenue for future research. Use of laboratory settings would also accommodate controlling for different levels of environmental dynamism and complexity, perhaps leading to a better understanding of the limits of potential managerial improvements.

Despite these limitations, the present study makes several important contributions. First, it is one of few studies of behavioral decision-making that assess the performance of managers' inventory decisions by observing their impact in the real world. This is in contrast to most research on behavioral aspects of inventory decisions, predominantly laboratory-based experiments that strictly control for environmental factors (Zhang and Siemsen 2019). We need to step away from controlled laboratory settings if we are to understand how human decision makers perform in realistic (e.g., complex and dynamic) decision environments. It is encouraging that despite the complexity of measuring effects in the real world—for example, our main measure of performance has low sensitivity to small modifications, and is affected by unaccounted environmental impacts—we still found significant evidence of managers' effects on decision making. Our work is among the first to push in this direction (see Craig et al. (2016) and Van Donselaar et al. (2010) for other studies of decision making in realistic contexts).

Second, we establish that managers can complement an ASO system and consistently, in real time, improve decision performance. This finding is contrary to most studies in the inventory decisions literature, which argue that managers are a source of bias and a liability (Becker-Peth and Thonemann 2019). Previous research, albeit limited, has shown flexibility and adaptability to be among the strengths of human decision makers (e.g, Gigerenzer 2008, Lawrence et al. 2006). We should explore ways to leverage these human capabilities, and the contexts in which they might be most useful. Our work illustrates a specific instance in which these capabilities are being effectively deployed.

Last, and more specifically, our study contributes to research on how to integrate decisions

from algorithms and human decision makers by identifying the *default choice* mechanism as a synergistic alternative to existing approaches (Blattberg and Hoch 1990, Lee and Siemsen 2017). Managers, although they cannot compete with the precision and essentially limitless data processing capability of a computer, are capable of detecting and identifying important signals in the environment and assessing how those signals might modify underlying assumptions of the algorithms. Providing a *default choice* is a way to enable human oversight of algorithms, thereby ensuring adaptability and flexibility to changing environmental conditions while still leveraging the computational intensity and detail traceability of algorithms.



### 3. RETURN OF THE BEHAVIORAL NEWSVENDOR: AN EXPERIMENTAL ANALYSIS OF CONSUMER RETURN POLICY DECISIONS

#### 3.1 Introduction

Two decades of research on behavioral inventory decisions since the seminal work by Schweitzer and Cachon (2000) has accumulated a large body of evidence for systematic deviations of actual order quantities from the normative order quantities that analytical models prescribe—known as the “pull-to-center” effect. Researchers have explored various theoretical accounts for the effect based on cognitive biases and heuristics (e.g., demand chasing and prospect theory), examined the role of different sources of individual heterogeneity in explaining the actual decisions, and tested various de-biasing techniques and decision support systems to improve decision-making efficacy (see Becker-Peth and Thonemann 2019, for a review). More recently, research has been extended to include pricing decisions, notably by Kocabıyıkoglu et al. (2016) and Ramachandran et al. (2018). Missing in this stream of research to date is the account of an important managerial problem in today’s retail environment: consumer returns.

Consumer returns have long been a quandary for the retail industry, with the annual value of returns exceeding \$643 billion globally (Cheng 2015) and \$428 billion in the U.S. alone (National Retail Federation 2021). The growing volume of consumer returns, exacerbated by overly-generous return policies, poses significant triple bottom-line challenges (i.e., focus on social and environmental concerns as well as profits) to retailers. The operational costs of handling returns, which amounts to about \$100 billion annually for U.S. retailers and manufacturers (Blanchard 2007), revenue losses due to refunds paid, and low salvage values for the returned products hurt firm profits. Similarly, excessive convenience, opportunistic, and fraudulent returns allowed by generous return policies negatively affect retail workers’ paychecks due to lost commissions (Abrams 2016). Further, consumer returns make significant environmental impact by generating 5 billion pounds of landfill in the U.S. annually, which is equivalent to trash produced by 5 million

Americans (Constable 2017), and, process of handling returns generate 15 million metric tons of carbon dioxide annually (Optoro 2018).

Retailers are experimenting with different return policies to reduce the burden of returns while trying to not compromise the value proposition and hurt sales performance. Cognizant of the increased importance of the problem for managers, academic research on consumer return policy (CRP) design has also grown significantly. To date, the literature on CRP design has predominantly focused on analytical models that prescribe optimal return policies and empirical investigations of consumer behavioral reactions to return policy leniency across different dimensions (see Abdulla et al. 2019, for a comprehensive review). However, we are yet to know much that is empirically established about the decision-making process behind and antecedents of CRP decisions in order to validate the assumptions and predictions of the analytical-theoretical models.

Vast majority of analytical literature views the salvage value of returned products to a retailer as a key determinant of the optimal return policy leniency in terms of refund amount paid to customers (e.g., Akçay et al. 2013, Altug and Aydinliyim 2016, Su 2009). Meanwhile, only two studies empirically examined the relationship between the salvaging capability of a retailer and its return policy leniency. The results of these studies are correlational and inconclusive. On one hand, Shang et al. (2017) found evidence, in an eBay transactional data set, for a positive association between offering a money-back guarantee and sellers' salvage capabilities (using store status and feedback volume as proxies). On the other hand, through a survey of 143 retailers, Davis et al. (1998) found no statistically significant association between return policy leniency and the extent of retailers' salvaging opportunities. Moreover, virtually no empirical research has been conducted to date regarding the relationship between CRP decisions and other key operational decisions, such as inventory and pricing decisions (Abdulla et al. 2019).

We close these gaps in the behavioral inventory decisions and consumer returns literature by investigating how decision-makers set return policy, together with pricing and ordering decisions. We consider a generalized newsvendor model in which a seller is to simultaneously decide order quantity, price, and refund amount for returned products (i.e., return policy leniency) in a retail

environment that involves both aggregate market demand and individual product valuation uncertainty (Su 2009). Through a behavioral experiment with a diverse sample of decision-makers, we examine whether and how actual order quantity, price, and refund amount decisions deviate from the normative decisions and investigate the causal effect of salvage value on these decisions. We find that individuals significantly deviate from the normative decisions across all decision levers. More interestingly, we find a significant causal effect of salvage value on actual order quantity, price, and refund decisions, in the directions prescribed by the normative model. In line with the recent behavioral operations literature focusing on the role of individual-specific factors in decision-making, we also examine the role of several key demographic characteristics and socio-economic preferences of the decision-makers in explaining the variability in the actual decisions. Overall, we find a lack of significant explanatory power of individual-level characteristics.

This being the first time that the behavioral return policy decisions are explored in a multi-lever decision-making context, we extend our research scope beyond the simple hypotheses testing and aim to explain the observed behavioral regularities, similar to the first behavioral study on the classical single-lever newsvendor model by Schweitzer and Cachon (2000). While Schweitzer and Cachon (2000) explore which of the behavioral models—with different preferences and utility functions—would explain the observed “pull to center” effect in order decisions, we employ an abduction process to provide explanations for the observed behaviors across three decision levers based on established judgement and decision-making theories. To this end, we first identify several time-dependent, dynamic behavioral regularities and dependencies in the decision process, and then we articulate a *process theory* to explain these regularities, thus providing a new direction with testable hypotheses for future research.

### **3.2 Normative Model and Hypotheses**

To derive the normative decisions for order quantity, price, and return policy, we considered a price-setting newsvendor framework with consumer returns, in a similar vein to Su (2009). In this framework, a retailer that sells a single product needs to decide order quantity  $Q$ , price  $p$ , and refund amount  $r$  for returns at the beginning of a single selling season. The retailer pays a procure-

ment cost  $c$  to order each unit of the product and charges price  $p$  for each unit he sells. The retailer issues a refund  $r$  for each unit of product returned by consumers. These three retail decisions (i.e.,  $Q, p, r$ ) are made in a decision-making environment with the following two uncertainties. First, the aggregate market demand for the product,  $X$ , is random with a known probability distribution having a cumulative distribution function  $F$ . Market demand is interpreted as a mass of consumers who consider buying the product. Second, each individual consumer can buy only a single unit and has uncertainty with respect to the value he or she receives from the product, denoted by the random variable  $V$  with a cumulative distribution function  $G$ . At the end of the selling season, the retailer salvages returned and unsold products at a value of  $s$  per unit. Without loss of generality, we assume that the salvage value is less than the procurement cost ( $s < c < E[V]$ ).

The specific sequence of events is as follows. The retailer simultaneously determines the order quantity  $Q$ , price  $p$ , and refund amount  $r$ . Then, the aggregate market demand is realized and the retailer sells the minimum between the order quantity and the market demand, i.e.,  $\min(X, Q)$ , if the price of the product does not exceed the expected consumer surplus (i.e.,  $p \leq E \max(V, r)$ , the participation constraint). If the price of the product exceeds the expected consumer surplus, i.e.,  $p \geq E \max(V, r)$ , no consumer will buy and retailer will not be able to sell any product. In case of understocking, the consumers who cannot buy the product leave the market empty-handed. The retailer cannot hold backlogs.

Consumers who have purchased the product realize their individual valuation of the product and make a keep versus return decision by comparing the utilities received from each option. Clearly, a given consumer will choose to keep the product if the value they receive from the product is greater than or equal to the value from returning ( $V \geq r$ ), otherwise the product will be returned for refund. Finally, the retailer salvages any leftover units and returned products at a salvage value  $s$ . Using the notation and the sequence of events, the retailer's profit maximizing objective function can be expressed as the following:

$$\begin{aligned}
\Pi(p, Q, r) = & \underbrace{p\bar{G}(r)E[\min(X, Q)]}_{\text{sold and kept}} + \underbrace{(p - r + s)G(r)E[\min(X, Q)]}_{\text{sold and returned}} \\
& + \underbrace{s(Q - E[\min(X, Q)])}_{\text{not sold}} - cQ
\end{aligned} \tag{3.1}$$

The profit-maximizing order quantity, price, and refund amount can be expressed as follows (see Su 2009, for the proofs):

$$\begin{aligned}
\bar{F}(Q^*) &= \frac{c - s}{p^* - s} = \frac{c - s}{\mathbf{E}[\max(V, s)] - s}, \\
p^* &= \mathbf{E}[\max(V, s)], \\
r^* &= s.
\end{aligned} \tag{3.2}$$

Notice that  $\bar{F}$  and  $\bar{G}$  represent the complementary cumulative distribution function of F and G, respectively.

We assume that the aggregate market demand is normally distributed with a known mean and standard deviation, i.e.,  $X \sim N(\mu_X, \sigma_X)$ , which is commonly used in behavioral operations literature involving the newsvendor framework (Zhang and Siemsen 2019). Similarly, we assume that the individual product valuation of consumers is drawn from independently and identically distributed random variables with a normal distribution, i.e.,  $V \sim N(\mu_V, \sigma_V)$ . Under these assumptions, the normative order quantity, price, and refund amount decisions can be expressed as follows (see Appendix B.1):

$$\begin{aligned}
Q^* &= \mu_X + \sigma_X \Phi^{-1} \left( \frac{p^* - c}{p^* - s} \right), \\
p^* &= \mu_V + (s - \mu_V) \Phi \left( \frac{s - \mu_V}{\sigma_V} \right) + \sigma_V \phi \left( \frac{s - \mu_V}{\sigma_V} \right), \\
r^* &= s.
\end{aligned} \tag{3.3}$$

where  $\phi$ ,  $\Phi$ , and  $\Phi^{-1}$  are the pdf, CDF, and inverse of a standard normal random variable, respectively. Notice that when  $s = 0$ , the model reduces to the classical newsvendor problem with no

consumer returns, i.e.,  $r^* = 0$ . Also note that when returns are not allowed, i.e.,  $r = 0$ , the retailer has to solve a price-setting newsvendor model. Thus, our model can be viewed as a generalized version of the classical newsvendor model. We know from the classical behavioral newsvendor literature that individuals deviate from the normative order quantity decisions (see Becker-Peth and Thonemann (2019), Zhang and Siemsen (2019) for reviews). As such, we expect the same to hold for our multi-lever behavioral newsvendor context, that is, individuals' order quantity, price, and refund decisions would deviate from the normative decisions. Thus, we hypothesize:

**Hypothesis 3.1: Decision-makers' actual order quantity, product price, and refund amount decisions deviate from the normative decisions in the set of equations (3).**

From the equations (3), we observe that all three managerial decisions are affected by the salvage value that the retailer extracts. In particular, as the salvage value increases (decreases), rational decision-makers should order more (less), increase (decrease) the price, and offer more (less) generous refunds (see Appendix B.2). Literature on managerial cognition discusses failures in problem sensing and taking appropriate actions in reactions to significant changes in the environment (Kiesler and Sproull 1982). Bounded rationality (Simon 1978) and failure in causal logics (Nadkarni and Barr 2008) are listed among the factors that drive such failures. In our context, we expect that even if decision-makers would not be able to find the optimal solution to the problem, they would recognize the causal link between salvage value and react to the change in salvage value in a boundedly rational matter, that is, in the directions prescribed by the normative model. Hence, we state the following hypothesis:

**Hypothesis 3.2: A higher (lower) salvage value of the product causes decision-makers to order in a higher (lower) quantity, charge a higher (lower) price, and choose a higher (lower) refund amount.**

Note that we are interested not only in testing the hypotheses per se, but also—this being the first time that return policy is and a multi-lever decision-making context are being assessed—in identifying and explaining systematic deviations from the normative decisions, reactions to change in the salvage value, and other behavioral regularities across decision levers.

### 3.3 Behavioral Experiment

#### 3.3.1 Experimental Design

We conducted an incentive-compatible randomized experiment with an AB/BA repeated measures crossover design (Wallenstein and Fisher 1977, Mitchell and Jolley 2010). The overview of the experimental design and parametrization is provided in Table 3.1. To elaborate, participants were randomly assigned to one of the two blocks that involved two treatments (i.e., two different salvage values), received in a reversed sequence. Each participant completed 15 rounds of decisions under each of the two treatments, hence the repeated measures feature of the design.

Table 3.1: Experimental Design and Parameterization

Fixed Task Parameters	Block	Period 1 (Rounds 1–15)				Period 2 (Rounds 16–30)			
		$s$	$Q^*$	$p^*$	$r^*$	$s$	$Q^*$	$p^*$	$r^*$
Market Demand ( $X \sim N(200, 50)$ )	1	0	171	50.04	0	30	230	51.67	30
Product Valuation ( $V \sim N(50, 20)$ )	2	30	230	51.67	30	0	171	50.04	0
Procurement Cost ( $c = 36$ )									

As noted earlier, when the salvage value is equal to zero, the analytical model prescribes a no-refund policy and the optimal order quantity becomes equal to the one prescribed by the classical newsvendor model. Therefore, we included salvage value of zero into the design as a treatment condition. The implied critical ratio to determine the normative order quantity, assuming zero salvage value and normative price, is 0.28. We operationalize the high salvage value with  $s = 30$ , which implies a critical ratio of 0.72. Overall,  $s = 0$  and  $s = 30$  treatments ensure normative order quantities that are equidistantly positioned relative to mean market size, aligned with most research designs in the behavioral newsvendor literature (Becker-Peth and Thonemann 2019). Note that the normative prices under these salvage values are close to the mean valuation of the product. Therefore, the current parametrization makes the low- and high-salvage value conditions analogous to high- and low-margin conditions per the classical behavioral newsvendor research. Given the

multi-lever, inter-related decision structure, the fact that the profit-maximizing price is close to the mean product valuation, which can be a reasonable anchor for decision-makers, is also favorable in terms of avoiding an overly-challenging task environment.

### **3.3.2 Procedure and Sample**

The experiment involved 136 participants recruited through the Prolific Academic crowdsourcing platform (Bhatia 2019, Quispe-Torreblanca et al. 2019). Prolific Academic is specifically designed for academic research purposes, as opposed to general purpose alternatives such as Amazon MTurk. Peer et al. (2017) shows that Prolific Academic provides comparable data quality to MTurk but provides a greater sample diversity. In line with Lee et al. (2018), we restricted the participants' geographic location to the United States and qualified only high-reputation workers (above 90% approval ratings) as participants.

The interface of the experimental task was designed using oTree open-source platform for behavioral research (Chen et al. 2016). After reading the instructions, participants answered four questions regarding the key aspects of the task and were not allowed to proceed to the task if they failed any of the questions twice. The experimental task interface alongside with the instructions and comprehension check questions are provided in Appendix B.3.

Participants who passed the quiz questions completed a 5-round warm-up session to explore the task interface and make inconsequential decisions. Then, participants were randomly assigned to one of the two experimental blocks and completed the 30 main rounds. Throughout the experiment, participants were provided with the complete history of their decisions and performance. All participants faced the same demand pattern, drawn from a normal distribution, during the task. In the task interface, the ordering of decision input boxes was randomized for each participant to control for order effects. Once assigned, the order of decision input boxes stayed the same for a given participant during the whole task. No upper bound was set for the decisions and participants were free to enter excessively large values, whereas no negative values were allowed. Once participants completed the 30 rounds of decisions, they completed the socio-economic preferences survey module based on Falk et al. (2016) and Falk et al. (2018) and answered several demographic-related



questions (see Appendix B.3).

Experimental Francs (Fr) were used as the currency in the task, which was converted to US Dollars based on an exchange rate of 3000 Fr = 1 USD. We paid the participants a flat \$3.00 compensation for participation and a performance based payment (mean = \$0.53, min = \$0, max = \$0.98). In total, the average payment to the participants was \$3.53 and the session lasted about 25 minutes. On an hourly basis, the compensation was significantly higher than the averages in online crowd-sourcing platforms to conduct behavioral experiments (\$6.00/hr) and the U.S. federal minimum wage (\$7.25 in 2020), ensuring significant salience for the economic incentives.

### **3.4 Testing the Hypotheses**

Before testing the hypotheses and as the result of applying a *pre-defined* data exclusion criteria, we dropped the data from 33 participants. The data exclusion criteria were based on various types of anomalous and hyper-irrational behaviors observed in a pilot study, such as unreasonably high decision value entries (i.e., outliers) and consistent pricing below the procurement cost. The details of the pre-defined criteria are listed in Appendix B.4. Overall, the rate of data loss in our experiment was comparable to data loss rates reported in online studies surveyed by Abbey and Meloy (2017).

A Shapiro-Wilk test of normality and a visual inspection of histograms and Q-Q plots revealed significant deviations from a normal distribution for the observed order quantity, price, and refund decisions. Thus, to test hypothesis 3.1, we used the non-parametric one-sample Wilcoxon signed-rank test that compared the median of actual decisions to the respective nominal decisions.

As a result, we found support for hypothesis 3.1. We found that participants showed significant deviations from the normative order quantity, price, and refund decisions. Specifically, we found that participants ordered significantly more than the expected profit-maximizing order quantity and charged less than the optimal price in the low salvage value condition. In the high salvage value condition, participants ordered significantly less than the expected profit-maximizing order quantity and charged less than the optimal price. These patterns are in alignment with the classical newsvendor behavior that has been extensively documented in behavioral operations literature

(Becker-Peth and Thonemann 2019). Meanwhile, the analysis of refund decisions revealed that under the low salvage value condition, participants set a significantly greater (more generous) refund than the optimal, whereas under high salvage values, participants set a lower (less generous) refund than the optimal. The results of the analysis are reported in Table 3.2.

Table 3.2: One-Sample Wilcoxon Sign Test Results

Treatment	s = 0	s = 30
$Q$	0.000 ( $M_a > 171$ )	0.011 ( $M_a < 230$ )
$p$	0.000 ( $M_a < 50.04$ )	0.000 ( $M_a < 51.67$ )
$r$	0.000 ( $M_a > 0$ )	0.000 ( $M_a < 30$ )

All p-values are for one-sided test against  $H_0$ : Median of Actual ( $M_a$ ) = Normative.  $H_A$  is given in parentheses.

To test hypothesis 3.2—the causal relationship between the salvage value and three operational decisions—we estimated a two-level hierarchical linear model with random intercepts and slopes (Raudenbush and Bryk 2002). In particular, our linear mixed-effect modeling (LMM) approach recognizes the nested structure of the data, that is, the fact that decisions across different rounds (level-1) are nested within individuals (level-2) and are not independent of each other (Oliva et al. 2022). Including random intercepts and slopes into the model (i.e., random effects) accounts for and allows an empirical assessment of participant-level deviations from the average round and period trends, as well as variations due to unobserved heterogeneities such as in task engagement, fatigue, and learning rates.

Using a multi-level specification, the derivation of the linear mixed-effect model is as follows (note that the models are identical for all three decisions):

$$Q_{ij}|p_{ij}|r_{ij} = \gamma_{0j} + \gamma_1 treatment_{ij} + \gamma_{2j} round1_{ij} + \gamma_{3j} round2_{ij} + \gamma_{4j} period_{ij} + \epsilon_{ij} \quad (3.4)$$

$$\gamma_{0j} = \beta_0 + u_{0j} \quad (3.5)$$

$$\gamma_1 = \beta_1 \quad (3.6)$$

$$\gamma_{2j} = \beta_2 + u_{1j} \quad (3.7)$$

$$\gamma_{3j} = \beta_3 + u_{2j} \quad (3.8)$$

$$\gamma_{4j} = \beta_4 + u_{3j} \quad (3.9)$$

where indices  $i$  and  $j$  capture decision round and decision-maker, respectively. In this two-level model, *treatment* is a binary variable that captures the salvage value condition a decision-maker faces in a round (with *treatment* = 1 for the high salvage value condition of  $s = 30$ ). The two round trend variables, *round1* and *round2*, designate the round number (1–15) during period 1 and period 2 of the task, respectively. Finally, *period* is a binary variable that indicates the two periods of the task.

In multi-level modeling nomenclature, equation (3.4) is referred to as level-1 (within subject, decision level) and equations (3.5)–(3.9) are known as level-2 (between subject, participant level) equations. Substituting equations (3.5)–(3.9) in the equation (3.4) results in a reduced form equation for linear mixed-effect model, equation (3.10).

$$Q_{ij}|p_{ij}|r_{ij} = \underbrace{\beta_0 + \beta_1 treatment_{ij} + \beta_2 round1_{ij} + \beta_3 round2_{ij} + \beta_4 period_{ij}}_{\text{Fixed Effects Portion}} + \underbrace{u_{1j} round1_{ij} + u_{2j} round2_{ij} + u_{3j} period_{ij} + u_{0j}}_{\text{Random Effects Portion}} + \epsilon_{ij}. \quad (3.10)$$

Note that, since we did not have an *a priori* hypothesis regarding the asymmetry in the effect of salvage value change on decisions (when salvage value increases versus decreases) and the observations are counterbalanced, we do not include a block-period interaction in the models in estimating the causal effect of salvage value. We will revisit this in §3.5. We estimated the models using maximum likelihood estimation with an unstructured covariance matrix that allows observations for each individuals to freely correlate across rounds, without any specific pattern assumed

(Littell et al. 2000). We estimated the models using Stata command *mixed*, and the results of the estimations are provided in Table 3.3. We find positive, statistically significant causal effects of salvage value on order quantity, price, and refund amount decisions (i.e.,  $\beta_1 > 0$  for  $Q, p, r$ ), supporting hypothesis 3.2. Thus, we find that, on average the reactions of decision-makers to changes in the salvage value follow the predictions of the normative model.

Table 3.3: Results of Linear Mixed-Effect Model for Testing H2

	(1) $Q$	(2) $p$	(3) $r$
Fixed Effects			
Intercept	181.730*** (5.439)	46.312*** (0.535)	19.520*** (1.201)
Treatment (s = 30)	16.803*** (2.078)	0.504* (0.226)	4.365*** (0.757)
Round1	-1.012** (0.355)	0.123*** (0.035)	-0.180*** (0.045)
Round2	-0.923*** (0.274)	-0.021 (0.026)	-0.259*** (0.072)
Period	4.186 (4.781)	1.706*** (0.493)	-4.708*** (0.990)
Random Effects			
Level 2 (between-subject)			
Intercept	2572.886*** (409.346)	24.576*** (3.911)	127.784*** (18.777)
Round1	8.650*** (1.811)	0.082*** (0.017)	0.147*** (0.029)
Round2	3.415*** (1.084)	0.025*** (0.009)	0.469*** (0.074)
Period	1640.566*** (328.951)	18.070*** (3.494)	90.520*** (14.112)
Level 1 (within-subject)			
Residual	1208.272*** (33.026)	11.788*** (0.322)	17.565*** (0.480)
$N$	3089	3089	3089
Standard errors in parentheses			
* $p < 0.05$ , ** $p < 0.01$ , *** $p < 0.001$			

Analytical CRP design literature consistently suggests that salvage value is a key determinant of the optimal refund amount—a dimension of return policy leniency—for a profit-maximizing retailer. Yet, the empirical relationship between the salvage value that the retailer can extract and its return policy leniency was not clear in light of the two studies to date that examined this

relationship (Davis et al. 1998, Shang et al. 2017). Our behavioral experiment finds support for a statistically significant positive causal effect of salvage value on refund amounts offered ( $\beta_1 = 4.365, p = 0.000$ ), in line with the predictions of the analytical-theoretical models.

We note that the positive causal effect of salvage value on return policy leniency can explain certain observations from the practice. For instance, in the apparel sector where a newsvendor environment is typical, many retailers such as Tommy Hilfiger, Michael Kors, Gap disallow final sale and clearance products to be returned (i.e., zero refund). Another case in point would be that of Nordstrom and Nordstrom Rack that are managed by the same firm. Here, the latter department store chain acts as a clearance or sale outlet for the items that do not sell during the season at the former. As such, Nordstrom Rack seems to provide a profitable salvaging opportunity for Nordstrom at the end of a season, whereas for Nordstrom Rack the salvage values would become much lower. Therefore, while Nordstrom is known for its no-questions asked, no-time-limit, full-refund return policy, Nordstrom Rack imposes a return limit of 45 days, and a 50% restocking fee for the products that are returned beyond the deadline.

Finally, in Table 3.3, the estimates of variances for the random effect terms ( $u_{0j}$ ) suggest that there is significant individual-specific variability in the decisions that is not captured by the common intercept and treatment, round, and period fixed effects. To further explore this variability, we examined the role of individual heterogeneity in explaining the actual decisions. Specifically, we studied two categories of individual-level heterogeneity: demographic characteristics and socio-economic preferences. Overall, we found that none of the socio-economic preferences could explain significant variability in order quantity, price, and refund decisions beyond what was already explained by salvage value treatment, round, and period effects. Among demographic characteristics, we found that a higher education level was associated with a smaller order quantity and a greater deviation from the normative order quantity. All other demographic characteristics could not explain significant variability in any of the three decisions. We refer readers to Appendix B.6 for the details from this analysis. In the next section, we focus on exploring the decision-making dynamics and providing a theoretical account of how actual decisions across the three levers are

made.

### 3.5 Building a Process Theory

Finding support for hypotheses 3.1 and 3.2 is hardly surprising. In this section, we continue our investigation by exploring some behavioral regularities in the observed decisions. Note that the experimental design that we employed—namely, AB/BA repeated measures crossover design—balances out order effects which ensures internal validity of the inferences regarding the testing of hypotheses 3.1 and 3.2 (Mitchell and Jolley 2010). Furthermore, our experimental design is subject to, and allows an examination of, period effects, sequence effects, and treatment carryover effects (Jones and Kenward 2015). These effects relate to the temporal changes in decision-making under a given treatment, differential reactions to the sequence of treatments received, and the differential influence of the initial treatment on decision-making under the subsequent treatment, respectively. Investigating whether and why such effects emerge in our decision-making context is of high practical and theoretical relevance, because managers in real life face changing operating conditions and apply managerial cognition developed in past periods under different operating conditions to the novel environments they face. For the ease of exposition, we will collectively refer to the above-mentioned effects as *time-dependent effects*.

To empirically assess the presence of time-dependent effects in our experiment, we estimate the following LMM:

$$\begin{aligned}
 Q_{ij}|p_{ij}|r_{ij} = & \underbrace{\beta_0 + \beta_1 block_{ij} + \beta_2 period_{ij} + \beta_3 block_{ij} \times period_{ij} + \beta_4 round1_{ij} + \beta_5 round2_{ij}}_{\text{Fixed Effects Portion}} \\
 & + \underbrace{u_{1j} round1_{ij} + u_{2j} round2_{ij} + u_{3j} period_{ij} + u_{0j}}_{\text{Random Effects Portion}} + \epsilon_{ij}.
 \end{aligned}
 \tag{3.11}$$

where *block* is a binary variable that identifies the decision-maker's randomly assigned block, and *period* is a binary variable that identifies the period in which the decision was made. We follow the growth model approach (Raudenbush and Bryk 2002) to control for potential within-period learning and fatigue effects using *round1* and *round2*. Note that by estimating block, period, and

the block  $\times$  period interaction effects—as opposed to the treatment, period, and their interaction effects—we focus on a sequence dependent temporal perspective that not only isolates the effects of the treatment, but also the time-dependent effects that earlier treatments have on subjects. While the models are mathematically equivalent, focusing on this formulation allows us to build time-dependent predictive margin plots (see Figures 3.1, 2 and 3). The results of the analyses for each decision lever are reported in Table 3.4.

Table 3.4: Effect of Treatment Sequence

	(1) $Q$	(2) $p$	(3) $r$
Fixed Effects			
Block	4.371 (7.829)	1.124 (0.726)	12.052*** (1.956)
Period	20.993*** (5.238)	2.249*** (0.547)	-0.476 (1.259)
Block $\times$ Period	-33.613*** (4.157)	-1.084* (0.460)	-8.471*** (1.516)
Round1	-1.012** (0.355)	0.123*** (0.035)	-0.180*** (0.045)
Round2	-0.923*** (0.274)	-0.021 (0.026)	-0.259*** (0.072)
Intercept	188.127*** (6.593)	45.993*** (0.643)	15.565*** (1.474)
Random Effects			
Level 2 (between-subject)			
Period	1640.588*** (328.935)	18.075*** (3.495)	90.349*** (14.067)
Round1	8.650*** (1.811)	0.082*** (0.017)	0.147*** (0.029)
Round2	3.415*** (1.084)	0.025*** (0.009)	0.469*** (0.074)
Intercept	2448.965*** (394.982)	24.667*** (3.930)	114.221*** (16.642)
Level 1 (within-subject)			
Residual	1208.272*** (33.026)	11.788*** (0.322)	17.565*** (0.480)
$N$	3089	3089	3089

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

We observe a significant block-period interaction effect on the three decision levers (i.e.,  $\beta_3 > 0$  for  $Q, p, r$ ), which indicates an asymmetry in the block and period main effects depending on the

sequence of salvage value treatment. Moreover, we observe differential impacts of the treatment sequences on decisions made across each lever; there are significant period as well as block-period interaction effects for both order quantity and price decisions, and significant block and block-period interaction effects for refund decisions. Collectively, our analysis suggests that, there are significant time-dependent effects, these effects are of different nature for each lever, and, therefore, the drivers of these effects may also be different.

In exploring these time-dependent effects we transition from testing of variance hypotheses predicated on the normative model to explaining time-dependent effects by articulating a process theory (Mohr 1982, Monge 1990). Process theories, in contrast to the variance theories that ignore time dependencies among contributing factors, provide causal explanations of how and why things happen and identify how entities participate in and are affected by the sequence of events, i.e., timing is critical to the outcomes in process theories (Mohr 1982). In theorizing, we follow “theory as a narrative” approach (Pentland 1999), which “requires that hypotheses detailing regularities in relations among variables be accompanied by plausible accounts of how the actions of real humans could produce the associations predicted and observed” (DiMaggio 1995, p. 391). See Serman et al. (2015) and Oliva (2019) for a discussion of process theories in OM. In §3.5.1 we focus on explaining within- and cross-period behavioral regularities, whereas in §3.5.2 we focus on explaining the within-round behavior.

### **3.5.1 Within- and Cross-Period Decision-Making**

In what follows, we employ abductive reasoning in the sense of making inference to plausible explanations for the observed within- and cross-period behavioral regularities (Peirce 1992). As such, while our explanations do account for the observed time-dependent effects (Rozeboom 1997) and draw upon existing theories from decision-making literature, they should be considered as conjectures or hypotheses when it comes to the specific causal mechanisms (Oliva et al. 2022). Therefore, the hypotheses that we propose in this section lend themselves to further empirical testing in future research (Donohue and Schultz 2019).



### 3.5.1.1 Order Quantity

Figure 3.1 reports estimated marginal means for order quantity decisions based on the block-period interaction model in column (1) of Table 3.4. We note that in period 1, both block 1 and 2 had similar average order quantity ( $\beta_1 = 4.371, p = 0.577$ ). In period 2, we see a change in the average order quantities between the two blocks that are asymmetrical in magnitude ( $\beta_3 = -33.613, p = 0.000$ ). Block 1, which saw an increase in salvage value, responded by increasing the average order quantity ( $\beta_2 = 20.993, p = 0.000$ ), while block 2, which saw a decrease in salvage value, responded by decreasing the average order quantity ( $\beta_2 + \beta_3 = -12.620, s.e. = 5.189, p = 0.015$ ).




Figure 3.1: Order Quantity Decisions by Treatment Blocks and Periods

Although a distinctive quantification of the carryover and sequence effects are not possible under the AB/BA crossover design (Jones and Kenward 2015), the fact that the first period average order quantities were not statistically different (i.e., no between subject treatment effect of salvage value) leads us to the conclusion that the first period order quantities across the two blocks were driven by a common mechanism and not by the respective salvage value treatment, which suggests no differential carryover effects from the first period going into the second period. Thus, we posit that a sequence effect is what the observed pattern in Figure 3.1 predominantly reflects.

In particular, we propose that the average first period order quantity decisions across both blocks are driven by a tendency to weight between the mean aggregate demand as an anchor and zero aggregate demand which was a highly salient event whose probability depended on price and refund decisions (Kahneman and Tversky 1972). As a result, decision-makers tend to order below the mean aggregate demand of 200, under both high and low salvage value treatments. Once salvage value changes (i.e., receiving the second treatment), decision-makers in both block 1 and block 2 react in a (boundedly) rational manner, increasing and decreasing order quantities, respectively. That is, decision-makers recognize the opportunity to make additional profit; they adjust their order quantities in reaction to the change in the salvage value and the direction of the adjustment is aligned with rationality. However, due to the similar under-ordering behavior in the first period—driven by probability weighting between mean and zero aggregate demand (i.e., empty market)—while the average order quantities of decision-makers in block 2 become fairly close to the normative order quantity, for decision-makers in block 1 the average order quantities still fall significantly below the normative order quantity. Overall, we conjecture that it is due to the zero market potential that our decision-making environment poses that the average orders remain below the mean aggregate demand under the high salvage value treatment, which deviates from the pull-to-center effect in the case of classical newsvendor behavior (Schweitzer and Cachon 2000).

### 3.5.1.2 Price

Figure 3.2 reports the estimated marginal means of price based on the analysis in column (2), Table 3.4. Overall, we observe that pricing decisions in both blocks do not show large variability across periods and blocks, floating between 45 and 50 Fr—the mean valuation for the product by individuals in the market. Our statistical analysis shows no statistical difference between block 1 and 2 in the first period ( $\beta_1 = 1.124, p = 0.122$ ). We find a similar significant increase in price in period 2 compared to period 1 for both block 1 ( $\beta_2 = 2.249, p = 0.000$ ) and 2 ( $\beta_2 + \beta_3 = 1.165, s.e. = 0.541, p = 0.031$ ). Further, we find a positive and significant learning effect in period 1 ( $\beta_4 = 0.123, p = 0.000$ ), which suggests that decision-makers, on average, priced the product slightly below the mean valuation (around 46 Fr) and slowly increased the price. In period

2, we do not observe a significant learning effect ( $\beta_5 = -0.021, p = 0.419$ ), which suggests that decision-makers were satisfied with the price decision by the end of period 1 and stayed with it. From the underlying model estimates, we abductively infer that the observed pattern reflects a significant period effect. That is, the main time-dependent effect we observe is driven by the temporal change rather than the effects from treatment sequence or the residual treatment effect from the prior period (note the lack of significance of the block effect, and marginal significance and small magnitude of the block  $\times$  period interaction effect).

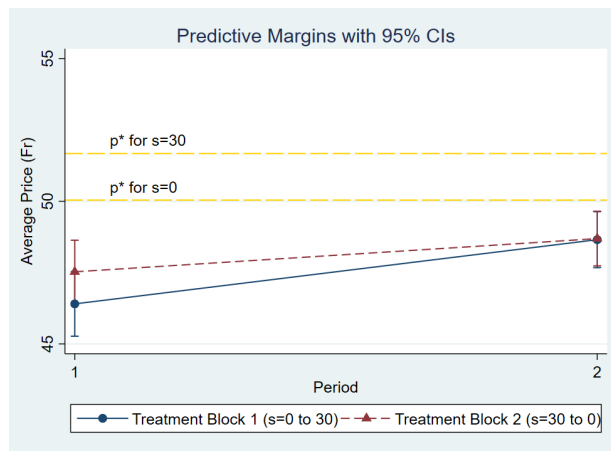


Figure 3.2: Pricing Decisions by Treatment Blocks and Periods

We propose that the mean valuation of the product was considered as an anchor by decision-makers. The observation that individuals under-priced relative to the expected profit maximizing price level can be explained by the dynamics of the task: individuals had to satisfy a market participation constraint and failing to do so (i.e., charging a very high price relative to the expected consumer surplus,  $p \geq E[\max(V, r)]$ ) results in an empty market and zero demand, causing significant losses. Hence, we posit that individuals used a satisficing heuristic to make price decisions. Specifically, pricing decisions were made based on the anchor on mean valuation and adjusted upwards until the decision-makers were satisfied with the price performance. Through this process, decision-makers also managed to effectively reduce the cognitive complexity of the task (Schwenk

1984). Interviews with the participants provide anecdotal evidence in support of our theoretical explanation. For example, one participant noted:

“[My tactic was to] start with the *price around the mean valuation*, and slowly increase the price in each round if all widgets were sold in the previous round...so even though market size and valuation were reset each round, *sticking with “successful” numbers* was generally a good tactic.”

### 3.5.1.3 Refund Amount

Figure 3.3 depicts the estimated marginal means of refund amounts based on the analysis in column (3), Table 3.4. We observe, on average, decision-makers in both treatment blocks offered refunds that are significantly greater than zero ( $\beta_0 = 15.565, p = 0.000$ ). Further, decision-makers who experienced the high salvage value treatment (i.e., block 2), on average, offered greater refund amount than those who experienced the low salvage value treatment ( $\beta_1 = 12.052, p = 0.000$ ), in period 1 (i.e., significant treatment effect of the salvage value). Moving to period 2, we find a significant asymmetry in terms of the magnitude of reactions to the change in salvage value across the two blocks ( $\beta_3 = -8.471, p = 0.000$ ). In particular, refund amounts offered by decision-makers in block 1 did not significantly increase when salvage value increased from 0 to 30 ( $\beta_2 = -0.476, p = 0.705$ ), whereas there was a significant decline in refund amounts offered by decision-makers in block 2 when salvage value decreased from 30 to 0 ( $\beta_2 + \beta_3 = -8.947, s.e. = 1.233, p = 0.000$ ). We infer that this asymmetrical reaction to the change in salvage value treatment is mainly due to the differential influence of the initial salvage value on refund amount decisions under the subsequent treatment (i.e., treatment carryover effect).

To explain this carryover effect, we employ the dual entitlement (DE) principle (Kahneman et al. 1986). The principle is often invoked in studies that examine fairness of various pricing tactics in a reaction to changes in costs (Vaidyanathan and Aggarwal 2003). An intriguing implication of the DE principle is that the seller’s profit entitlement takes precedence over the buyer’s price entitlement whenever both are threatened. The DE principle effectively implies that the seller is




Figure 3.3: Refund Decisions by Treatment Blocks

allowed to increase its profits even when there are costs reductions, once the consumer entitlement is met. In other words, “it is fair for prices and profits to only ever increase, because it is consistent with this norm of fairness for sellers to pass on costs increases and not cost decreases” (Kalapurakal et al. 1991, pg. 789).

As seen in Figure 3.2, the prices set by decision-makers remained relatively flat across different operating conditions, due to strong anchoring on the mean product valuation and a satisficing heuristic. In this case, we propose that decision-makers endogenously form a “fairness norm” for refund amounts in the first period, following the DE principle. In particular, decision-makers perceive a norm of what net profit margin they are entitled to versus what would be a fair compensation to consumers who experienced a low valuation of the product and returned it (i.e., consumer entitlement). Therefore, though decision-makers offer higher-than-optimal refunds (even if they experience low salvage values to meet the consumer’s entitlement) under the high salvage value treatment they are unwilling to offer a refund as high as the salvage value in an attempt to preserve the net margin they feel entitled to and instead offer less-than-optimal refunds. In both cases, however, decision-makers end up making sub-optimal refund decisions.

The DE principle also predicts an asymmetry in terms of passing gains (due to an increase in salvage value) to consumers by increasing the refund amount versus decreasing the refund amount

when the net margin is threatened. Aligned with this prediction, our analysis showed that refund amounts offered by decision-makers in block 1 did not significantly increase when salvage value increased from 0 to 30, whereas there was a significant decline in refund amounts offered by decision-makers in block 2 when salvage value decreased from 30 to 0. This finding supports the argument that the seller's profit entitlement takes precedence over the buyer's price entitlement whenever both are threatened, and sellers pass on cost increases to buyers but not cost decreases. Hence, the DE principle provides a parsimonious, plausible explanation for the carryover effect that we abductively inferred from Figure 3.3 and the underlying empirical analysis.

Our interviews with participants in the experiment provided supporting qualitative evidence that though decision-makers were aware of the negative implications of offering positive refunds when salvage value is zero, they nevertheless had to do so due to social norms and fairness (Bolton and Chen 2019). For example, one participant noted (emphasis added):

"I remember thinking that the refund amount was *unethically low*—I would never set a refund amount so low compared to how much it was sold for, but my goal was to maximize profit, so I had to set the refund amount low enough to keep a profit but *high enough* that it wasn't too risky."

### 3.5.2 Within-Round Decision-Making

Now, we turn our attention to within-round decision-making. In particular, we try to understand whether an independent or a conditional decision-making heuristic explains better the behavioral regularities that we observe in a typical round of decision-making. The independent decision-making heuristic implies that individuals make each decision largely independent of the remaining two and that each decision is subject to idiosyncratic sources of biases. In contrast, the conditional decision-making heuristic implies that individuals make decisions over a lever by conditioning on biased decisions over the other two levers. Note that the normative solution to the expected profit maximization problem, equations (3.3), also implies unbiased conditional decision-making, i.e., setting  $r^* = s$ , using  $r^*$  to calculate  $p^*$ , and using  $p^*$  to calculate  $Q^*$ .

By leveraging the exogenous variation in the salvage value, we can examine individuals' reliance on an independent versus conditional decision-making heuristic. To this end, we first consider three preliminary structural causal models a–c in Figure 3.4 that represent all simple direct and indirect (i.e., mediation) relationships among the three decision levers as caused by the salvage value. In doing so, our aim is to understand the most plausible representation of the interrelationships among the decision levers in a typical round of decision-making.




Figure 3.4: Preliminary Structural Causal Models

We estimate the models using STATA *sem* command with 5,000 bootstrapped samples for the first fifteen rounds (period 1) of data, since this part of the data is from a pure between-subject design and does not involve time-dependent effects. We use the bootstrapping technique for the estimation due to the non-normality of the indirect effects (Preacher and Hayes 2008), and use identical random seed to ensure stable comparison of standard errors across models. The individual path estimates and the indirect effects are reported in Table 3.5 columns (2) to (4). For reference, we also report (base model, column (1)) the coefficients of the direct effect of salvage value on each of the decisions, i.e., without indirect effects.

From the examination of the path coefficients and indirect effects across three models, we draw several inferences. First, model (a) estimates show that the effect of  $s$  on  $Q$  is mediated by  $p$ , with the  $p \rightarrow Q$  link also being highly significant. Model (b) estimates suggest that the effect of  $s$  on  $p$  is fully mediated by  $r$ . From model (c) estimates we observe the effect of  $s$  on  $r$  is only partially mediated by  $p$ . Overall, we find strong evidence for the indirect effect along the  $s \rightarrow r \rightarrow p$  path

Table 3.5: Analysis of Structural Causal Models

	(1) Base Model	(2) Model (a)	(3) Model (b)	(4) Model (c)	(5) Final Model
$s \rightarrow r$	12.360*** (0.581)	12.360*** (0.581)	12.360*** (0.581)	12.015*** (0.571)	12.360*** (0.581)
$s \rightarrow p$	0.929*** (0.279)	0.929*** (0.279)	-0.110 (0.378)	0.929*** (0.279)	-0.153 (0.379)
$s \rightarrow Q$	0.437 (2.778)	-3.455 (3.454)	0.437 (2.778)	0.437 (2.778)	-3.455 (3.454)
$r \rightarrow Q$		0.221 (0.145)			0.221 (0.145)
$p \rightarrow Q$		1.251*** (0.430)			1.251** (0.430)
$r \rightarrow p$			0.084*** (0.022)		0.088*** (0.022)
$Q \rightarrow p$			0.012*** (0.004)		
$Q \rightarrow r$				0.009 (0.006)	
$p \rightarrow r$				0.367*** (0.096)	
<b>Indirect Effects</b>					
$s \rightarrow r \rightarrow Q$		2.729 (1.807)			2.729 (1.807)
$s \rightarrow p \rightarrow Q$		1.162* (0.479)			-0.192 (0.490)
$s \rightarrow r \rightarrow p$			1.033*** (0.283)		1.082*** (0.284)
$s \rightarrow Q \rightarrow p$			0.005 (0.034)		
$s \rightarrow Q \rightarrow r$				0.004 (0.025)	
$s \rightarrow p \rightarrow r$				0.341* (0.132)	
$s \rightarrow r \rightarrow p \rightarrow Q$					1.354* (0.602)
Observations	1545	1545	1545	1545	1545

Bootstrapped standard errors in parentheses  
 \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

and marginally significant evidence from the paths  $s \rightarrow p \rightarrow Q$  and  $s \rightarrow p \rightarrow r$ . Thus, we proceed with a final structural causal model that includes  $s \rightarrow r \rightarrow p$  path into model (a) (see Figure 3.5), since this representation most comprehensively captures the statistically significant relationships among the variables. Note that, theoretically, the  $s \rightarrow r \rightarrow p \rightarrow Q$  path aligns with the sequential solution procedure implied by the normative model solution in §3.2.

The estimation of this model using the first period data (column (5) Table 3.5) suggests that  $r$






Figure 3.5: Final Structural Causal Model

and  $p$  decisions in series completely mediate the between-subject effect of salvage value on  $Q$  decisions. Thus, we find evidence that a conditional decision-making heuristic was employed during the first period. That is, individuals' decisions across each lever are not completely independent of each other and a sequential, conditional decision-making heuristic aligned with the solution procedure implied by the normative decisions explains the decisions made.

### 3.5.3 Process Theory Summary

The conditional decision making identified in the previous section can be combined with the proposed explanations for time-dependent effects to articulate the necessary conditions and causal mechanism required to specify a process theory (Mohr 1982, Oliva 2019). In a typical round, a conditional decision-making heuristic along the  $r \rightarrow p \rightarrow Q$  path is employed. Refund decisions are made based on the salvage value, social norms, and the dual entitlement principle (Kahneman et al. 1986, Bolton and Chen 2019). In particular, decision-makers establish entitlements based on social norms and the salvage value treatment they face in the first period, the influence of the first treatments carries over to the second period via the established entitlements and norms, and the adjustments to the refund amounts facing the second salvage value treatment are made based on the dual entitlement principle. Decision-makers set price decisions by anchoring (Tversky and Kahneman 1974) on the mean valuation of the product and adjust upward over time, while trying to satisfy market participation constraint to avoid empty market (Simon 1955). As rounds progress, decision-makers find a satisficing solution and by doing this they effectively reduce the cognitive complexity of the task. Within a round, price decisions are also impacted by sub-optimal refund amount decisions due to conditioning. In the first period, order quantities are predomi-

nantly determined by a probability weighting between mean demand as a salient anchor and zero aggregate demand as a salient event with dramatic consequences. Moving to the second period, the decision-makers adjust their decisions, in a boundedly rational manner, responding to the change in the salvage value treatments relative to their initial treatments. Due to conditioning, sub-optimal refund amount and price decisions also impact order quantity decisions in a typical round. Table 3.6 summarizes our process theory. Note that the constituents of Table 3.6 can be viewed as a set of new hypotheses that are subject to further testing.

	Within-Period Behavior	Proposed Explanation	Cross-Period Behavior	Proposed Explanation	Time-Dependent Effect Type
<i>r</i>	Decision-makers choose refund decisions based on social norms and the salvage value.	Social norms theory (Elster 1989)	Refund decisions are adjusted asymmetrically in response to the change in salvage value, conditional on the salvage value in the first period.	Dual-entitlement principle (Kahneman et al. 1986)	Carryover effect
<i>p</i>	Decision-makers anchor on the mean product valuation and make downward adjustment to avoid over-pricing and zero aggregate demand. Learning takes place as rounds progress.	Anchoring and adjustment (Tversky and Kahneman 1974)	Irrespective of salvage value change, prices are adjusted upward toward the mean valuation and stay relatively stable moving into the second period.	Learning and satisficing (Simon 1978)	Period effect
<i>Q</i>	Decision-makers order, on average, between the mean aggregate demand and zero aggregate demand—a highly salient event whose probability depends on price and refund decisions.	Anchoring and adjustment (Tversky and Kahneman 1974)	Moving into the second period, individuals insufficiently increase (decrease) orders in reaction to the increase (decrease) in the salvage value.	Bounded rationality (Simon 1955)	Sequence effect

Table 3.6: Summary of the Proposed Process Theory

The process theory developed in this section demonstrates the complexity of the decision-making process in a multi-lever operational environment and how decision-makers cope with this complexity. The main takeaway is that decision-makers are able to react, in a boundedly rational manner, to changes in operating conditions (i.e., salvage value) in the directions prescribed by the normative model. Further, decision-makers can employ a conditional decision-making heuristic that aligns with the procedure that can generate the normative decision structure. However, due to idiosyncratic sources of biases affecting decisions across each lever, individual decisions end

up far from optimality. Isolating the unique impacts of each of these biases and heuristics on the overall profit performance of the individuals is overly challenging given the complexity of the decision-making context and beyond the scope of the current work. Nevertheless, we think that future research in this direction would be fruitful.

### **3.6 Conclusion**

We examined, in a randomized behavioral experiment, joint decision-making across three operational decision levers in a retail environment with aggregate demand and individual product valuation uncertainties: order quantity, price, and refund amount. We make several contributions to the behavioral operations management and consumer returns literature streams.

To the best of our knowledge, our research is the first to examine operational decision-making via multiple (more than two) levers, stepping beyond Kocabıyıkoglu et al. (2016) and Ramachandran et al. (2018) who jointly examined price and order quantity decisions. We found systematic deviations from the normative decisions across all three decision levers. We also examined the role of various individual demographic and socio-economic characteristics in explaining the actual decisions and found that most of these characteristics cannot explain significant variability in the actual decisions. This led us to examine the time-dependent effects as allowed by our experimental design, which revealed interesting behavioral regularities. We developed a process theory by employing abductive reasoning to explain these time-dependent effects. This purposeful shift in the research mode allowed us to explain in a more nuanced way the effects that were embedded within the effects that we estimated while testing the hypothesis regarding the causal effect of salvage value on price, refund, and order quantity decisions.

Our investigation revealed that given an operating condition (i.e., a salvage value), individuals use a conditional decision-making heuristic, in which a biased decision over one lever influences the decisions over the other levers. When the operating condition changes, individuals react in a boundedly rational manner to the change to make order quantity decisions, follow the dual-entitlement principle to adjust refund amounts, and demonstrate learning through exploration and satisficing in setting prices. Overall, our research highlights the importance of examining tran-

sitions in decision-making heuristics based on the operating conditions (Oliva et al. 2022) and time-dependent effects. In this respect, we also view our investigation in this part as a contribution to the behavioral operations management literature as an application of statistical induction combined with abductive reasoning to provide plausible explanations to the observed phenomena, a research mode that Rozeboom (1997) refers as explanatory induction.

Second, we extended the research on behavioral inventory decision-making to also include consumer product returns, which is a critical managerial issue in today's retail environment. By doing so, we make two key contributions to the literature in consumer returns that respond to the recent calls by Abdulla et al. (2019). First, we provide the first behavioral examination of return policy decision-making. We revealed significant and systematic deviations from the optimal solutions of normative models that often prescribe a refund amount that is equal to the salvage value. We showed that decision-makers tend to choose and adjust refund amounts based on operating conditions (i.e., salvage value) and intrinsically established fairness norms that is in line with the dual entitlement principle. We also resolve an inconclusiveness, due to diverging correlational evidence, in the existing empirical literature regarding the relationship between salvage value and return policy leniency. In particular, we found a significant causal effect of salvage value on behavioral decisions with respect to return policy leniency. Further, we demonstrated that the magnitude of this effect is contingent on the direction that salvage value changes—in line with the DE principle. Finally, we show that demographic characteristics and socio-economic preferences of the individuals do not significantly predict refund amounts offered.

Our study opens many opportunities for future research. For instance, in this research we used the refund amount as the lever that characterized the retailer's return policy decision. As this is the first attempt to examine behaviorally return policy decision-making, we see opportunities to examine other return policy leniency levers available to the retailers, such as return time window. Further, as the first investigation in this area, we focused on a context where return policy leniency did not have a direct effect on the distribution of demand in subsequent periods, which is a common assumption in the analytical literature on consumer returns (Abdulla et al. 2019). Future research

can focus on a more complex environment where there is a demand stimulation potential of offering a lenient return policy in a given period.

Using the current decision-making context, future research can also examine the impacts of task decomposition and group decision-making on the efficacy of decisions across the levers and the overall performance. Another interesting avenue would be to examine the extent to which each of the levers would benefit most from implementation of a decision support system or nudge (Thaler and Sunstein 2009).

Finally, we see opportunities in exploring and explaining behavioral regularities in changing operational decision-making environments through process theorizing. Research in behavioral operations management has traditionally focused more on testing the variance hypotheses implied by normative models or simply the systematic deviations of actual decisions from normative decisions. While decision-making under different operating conditions, such as under high-margin versus low-margin conditions, has been explored in the existing behavioral newsvendor literature, understanding the interaction of this conditions with time in influencing decision-making and examining the underlying decision-making processes have been overlooked. For example, the ordering of high-margin and low-margin conditions is typically randomized in behavioral experiments for the purpose of controlling for potential order effects and certain adjustments are made to the pay-off structures to ensure same performance potential under both conditions. A deeper of examination of any time-dependent effect (i.e., period, sequence, carryover effects) that was observed has rarely been of a theoretical interest. However, it was the significant time-dependent effects that motivated us to engage in sense-making and use a process theory approach to explain the behavioral regularities observed in our data. We believe that our research provides an alternative perspective for behavioral operations research that can stimulate theoretically and practically relevant investigations. Beyond its practical relevance for retail operations, the behavioral task that we employ in this research can be used as an archetype of multi-lever, complex decision-making environment through which behavioral researchers can examine a broad variety of cognitive, affective, and conative processes.

## 4. SATISFICING AND EXPLORATION BEHAVIOR UNDER COMPLEX AND DYNAMIC DECISION-MAKING ENVIRONMENT

### 4.1 Introduction

The first two essays (§2 and 3) investigated the operational decisions made under a complex and dynamic decision-making environment in a retail context. The first essay found that the decision-making environment affects how retail managers interact with algorithmic decisions (i.e., ASO proposals). Specifically, we analyzed inventory decisions made by retail managers responsible for ordering and managing thousands of SKUs. In this complex environment, we found that managers can improve decisions proposed by algorithms by providing flexibility into an efficient but not flexible decision-making algorithm.

In the second essay, we analyzed the use of heuristics in a complex and dynamic decision-making environment through our controlled laboratory experiment. We found that as the environment changed over time through our manipulation of the salvage value, we observed different time-dependent effects on the three decision levers (i.e., order quantity, price, and refund amount). Particularly, decision-makers in the experiment faced identical cost parameters and distributions of aggregate demand and individual product valuations. Thus, the general insignificance of individual heterogeneity measures compared to the effect of treatment conditions (i.e., salvage value) suggests that the observed behaviors were predominantly driven by the environment that decision-makers faced in a given period.

The two essays, taken together, argue for the importance of accounting for the decision-making environment when analyzing decisions in a complex and dynamic environment—as theorized in Simon’s (1955) theory of bounded rationality. This essay further investigates the decision-making behavior by focusing on exploration behaviors, which consider the search processes for better outcomes through trial and error as used in the evolutionary algorithm (Adra and Fleming 2011, Črepinšek et al. 2013, McGinley et al. 2011) and machine learning literature (Kaelbling et al.

1996, Madhawa and Murata 2020, Thrun 1992). We make predictions—based on the theory of bounded rationality (Simon 1955) and prospect theory (Kahneman and Tversky 1979)—on how the environment affects decision-maker’s exploration in search of profit-maximizing choice. We test our prediction by leveraging the empirical observations made via controlled laboratory experiment described in §3.3.

The essay is structured as follows. We summarize the related literature and introduce our hypotheses in §4.2. Our data sources, variables, and empirical models are described in §4.3. We present our results in §4.4, and the implications of and conclusions from our findings are discussed in §4.5.

## **4.2 Literature Review and Hypotheses**

Managing exploration and exploitation tradeoff is an important behavior that forms the basis for decision-making under uncertainty and is fundamental in understanding the adaptive behavior in complex and dynamic decision-making environments (Laureiro-Martínez et al. 2010). Due to this importance, a number of scholars have conducted research on exploration and exploitation tradeoffs in decision-making, notably starting with the seminal work by March (1991). Particularly, management literature largely focuses on macro-level (firm-level) exploration behavior and studies how changes in factors such as decision-making environment (Audia et al. 2000, Posen and Levinthal 2012), incentive schemes (Ederer and Manso 2013), and uncertainty (Gershman 2019) drive exploration behaviors.

Despite the progress made in understanding the exploration behavior at the macro-level, relatively little is known about micro-level (individual decision-makers) exploration behavior (Eisenhardt et al. 2010, Laureiro-Martínez et al. 2015). Specifically, at the micro-level, little is known about the antecedents of exploratory behavior (Hardy III et al. 2014), and the mechanism of explorations (Laureiro-Martínez et al. 2010). Also, most studies focus on one-shot decisions; hence little is known about the “dynamic aspects of repeated search and decision-making” (Huang and Hutchinson 2013, p. 163).

Similar to the management literature, the topic of exploration behavior at the micro-level is

understudied in the field of operations management and, particularly, in the behavioral operations. This gap in the literature in behavioral operations is due to the focus on the uncovering of behavioral biases (Oh and Oliva 2021), based on the variance theory, thereby largely ignoring the time-dependent nature of exploration and exploitation tradeoff. We contribute to both management and behavioral operations literature by using two individual-level decision-making theories – prospect theory (Kahneman and Tversky 1979) and bounded rationality (Simon 1957) – to hypothesize the exploration behavior in complex and dynamic decision-making environments. Specifically, we investigate how individuals respond to environmental change, and we unpack their exploration behavior using a new experimental design, the generalized newsvendor experiment (Oh et al. 2021). Further, we test the mechanism of how profit performance can mediate the effect of changing environments on explorations. We do so using process theory (Mohr 1982) as our theoretical lens, thereby responding to calls for more process focused rationality (Levinthal 2011, Levinthal and March 1993).

Our study also contributes by utilizing an AB/BA repeated crossover research design to consider both the between-subject and within-subject effects. That is, we consider both the initial exploration effect under favorable and unfavorable environments, and then extending this into within-subject effect by also considering how the exploration decisions change when the environment changes. This is in contrast to the current paradigm in behavioral operations literature that focuses on average decisions (i.e., between-subject measures) and ignore the time-dependent effects (i.e., within-subject effects) by averaging the decisions made under the same treatment condition but at different time periods or by excluding it from the research design.

#### **4.2.1 Bounded Rationality, Prospect Theory, and Exploration Behavior**

Central to our theoretical account is the notion of bounded rationality (Simon 1955), which states that when individuals make decisions, their rationality is limited by the structures of the environment they face, the complexity of the decisions, and available cognitive resources. Thus, human decision-makers simplify the complex decision-making environment by stopping the search for the profit-maximizing choice once they find a satisfactory solution. This decision-making



strategy is known as satisficing heuristics (Simon 1955) and is in contrast to the rational decision-making strategy that searches for the optimal decision until it is found.

We consider a complex decision-making task that involves three interrelated decision levers with two sources of uncertainty. Thus, our experiment demands a significant amount of cognitive resources from the decision-makers to search for the profit-maximizing set of decisions. Also, in addition to the round level uncertainty, our experiment involves changes in the decision-making environment that regulates the expected profit, complicates the decisions, and influences the decision-making process. That is, the decision-makers are exposed to an exogenous shock to the environment via the treatment change in salvage value that may result in an asymmetrical response depending on the treatment sequence. Accordingly, our development of the hypotheses is grounded on the ideas of bounded rationality and process theory. We predict that the operating environment will affect how decision-makers explore different choice sets in search of better profit performance.

The decision-making environment impacts the decisions because it controls the expected profit of the decisions. Decision-making environments become favorable to the decision-makers when the expected profit increases. In our experiment, all things equal (*ceteris paribus*), higher salvage value leads to higher expected profit. Thus, decision-makers who stick with their decisions despite the environmental change will face either increased or decreased expected profit when the environment becomes favorable or unfavorable (i.e., increase or decrease in salvage value), respectively. Notice that experiencing favorable conditions (i.e., high salvage value) does not guarantee that decision-makers will gain superior profits. The realized profits observed at each round is dependent not only on the salvage value but also on the realized demand and decisions made at that particular round. Hence, the profits decision-makers face are highly dependent on their decisions and how they search for a better set of decisions.

Another theory central to our study is prospect theory, which reveals that people tend to be risk-taking when facing a loss and risk-averse when facing a gain (Kahneman and Tversky 1979). Following prospect theory, we propose that individuals become motivated to take risks when facing a

loss in an unfavorable operating environment (i.e., zero salvage value,  $s = 0$ ). The heightened risk-taking attitude in the loss domain leads individuals to explore with their decisions to recoup, resulting in greater variability in their decisions across subsequent rounds. Meanwhile, facing profit in a favorable operating environment (i.e.,  $s = 30$ ), individuals exploit initial decisions, which would result in smaller variability in subsequent decisions. This exploitation behavior is aligned with satisficing heuristics. Indeed, literature on exploration versus exploitation in decision-making links risk-taking to variance-generating exploration and satisficing to variance-reducing exploitation behavior (March 1991). Recent empirical evidence also suggests that individuals tend to explore immediately after incurring losses and exploit more immediately after gains (e.g., Sloman et al. 2019).

We postulate that operating environments that the individuals face, favorable or unfavorable environments, lead to different behavioral processes (i.e., exploration versus exploitation). Through this line of argument, we predict a greater degree of exploration across the rounds for the individuals in the unfavorable environment and less exploration (i.e., more exploitation) for those in the favorable environment during the first period of the task. In the second period, as the operating conditions change due to an exogenous shock (i.e., increase or decrease in salvage value), we posit that individuals set their first period performance as the reference point and those that see favorable environments in the second period will demonstrate more exploitation (i.e., less exploration) behavior compared to the unfavorable first period. We predict the opposite for the individuals who see unfavorable environments in the second period. Based on our theoretical argument, we state the following two hypotheses:

**Hypothesis 4.1: Decision-makers explore more (less) when they are in a decision-making environment with low (high) expected profit.**

**Hypothesis 4.2: Decision-makers increase (decrease) the degree of exploration when the expected profit of the decision-making environment decreases (increases).**

Further, based on the prospect theory, we hypothesize that the effect of the decision-making

environment on exploration decisions is mediated by the profit performances that decision-makers see. That is, the mechanism of how decision-makers decide to explore the choice sets or exploit the current choice depends on the decision-making environment, which affects the profit performance that ultimately affects the exploration versus exploitation decision.

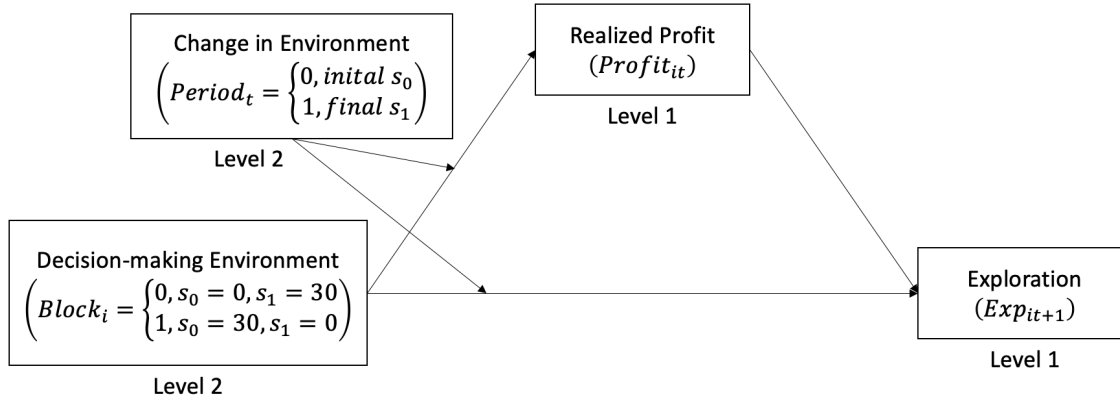
**Hypothesis 4.3: Realized profit negatively mediates the effect of decision-making environments on exploration behavior.**

We also predict that this mediation path is moderated by the treatment sequence that the decision-makers experience. That is, the different directional changes in the decision-making environment affect the profit in an opposite way, thus leading to divergence in exploration decisions. Particularly, we expect the decrease in exploration for decision-makers who go from favorable to unfavorable to be greater in magnitude than the increase in exploration for those who go from unfavorable to favorable environment. This is because bounded rationality states that decision-makers will stop searching for a better choice set once the performance reaches a satisficing level of utility. Thus, we predict that decision-makers in complex and dynamic decision-making environments will explore the choice sets more when they face losses compared to when they are facing gains. Specifically, decision-makers' satisficing level of utility is lower when going from unfavorable to favorable because they *gain profit* when they go from low to high profit, compared to the case when they are *losing profit* when they go from high to low profit.

We state the following moderated mediation hypothesis along with a conceptual diagram in Figure 4.1:

**Hypothesis 4.4: Environmental shift negatively moderates the mediation path.**

Hypotheses 4.1 and 4.3 are based on the variance theory where time-dependencies are not accounted for. Thus, we test such hypotheses using our first period measures, i.e., using between-subject estimates. Hypotheses 4.2 and 4.4 are based on the process theory where time-dependencies (i.e., historical change in the environment) are accounted for. We use both first and second period measures and focus on the within-subject estimates to test for time-dependent hypotheses.



\*Note that  $s_0$  and  $s_1$  represent salvage value at period 1 and 2, respectively. Level 1 and 2 signify the multilevel structure of research question and data.

Figure 4.1: Moderated Mediation Conceptual Diagram

### 4.3 Data, Variables, and Models

We leverage the observations from the experiment detailed in §3.3 to test our hypotheses. The experiment is useful for investigating exploration behaviors because of two reasons. First, the experimental setting creates a complex and dynamic decision-making environment. The three decision levers (i.e., order quantity, price, and refund amount) are decided simultaneously by the decision-makers under two different types of uncertainty (i.e., demand and individual product valuation uncertainty). This environment is ideal for testing our hypothesis regarding how decision-makers search for profit-maximizing solutions because the complex decision-making environment is repeated for multiple rounds. These repeated measures allow us to investigate the dynamic behavior and how decision-makers explore possible choice sets.

Second, the randomization process manipulates the decision-making environment via the change in salvage value. Specifically, decision-makers were randomly assigned to one of two salvage value conditions (i.e., treatment block), favorable ( $s=30$ ) or unfavorable ( $s=0$ ), in the first period. This statistically allows us to test hypotheses 4.1 and 4.3 in a causal manner (i.e., a test of hypotheses by leveraging the randomization processes). Further, the randomly assigned salvage value conditions change from favorable to unfavorable and vice versa, depending on the treatment block, in the second period—following AB/BA repeated measures crossover design (Wallenstein and Fisher

1977). This research design allows us to test for the time-dependent effect of treatment sequence, as conjectured in hypotheses 4.2 and 4.4.

### 4.3.1 Variables

In order to measure the degree of exploration behavior, the dependent variable, we modeled the response of decision-makers based on the information set available at the decision time (equation 4.1). Our measure follows the approach commonly used in the evolutionary algorithm literature that utilizes the Euclidean distance between consecutive decisions to measure exploration behavior (Adra and Fleming 2011, Črepinšek et al. 2013, McGinley et al. 2011), i.e.,  $[p_{i,t} - p_{i,t-1}, Q_{i,t} - Q_{i,t-1}, r_{i,t} - r_{i,t-1}]$ . These magnitudes of exploration at three decision levers have different impact on profit. Thus, to make the exploration along different decision levers comparable and in line with the goal of profit maximization, we normalized each lever according to the marginal effect they have on profit (i.e.,  $\nabla\Pi(p_{i,t-1}, Q_{i,t-1}, r_{i,t-1})$ ).

$$\text{Exp}_{it} = \|\nabla\Pi(p_{i,t-1}, Q_{i,t-1}, r_{i,t-1}) \cdot [p_{i,t} - p_{i,t-1}, Q_{i,t} - Q_{i,t-1}, r_{i,t} - r_{i,t-1}]\|_2 \quad (4.1a)$$

$$= \sqrt{\left(\frac{\partial\Pi}{\partial p} \times (p_{i,t} - p_{i,t-1})\right)^2 + \left(\frac{\partial\Pi}{\partial Q} \times (Q_{i,t} - Q_{i,t-1})\right)^2 + \left(\frac{\partial\Pi}{\partial r} \times (r_{i,t} - r_{i,t-1})\right)^2} \quad (4.1b)$$

We measure *profit*, the mediating variable, using the profit function in equation 3.1 and classify the decision-making environment, the independent variable, using the following two variables—*block* and *period*. *Block* represents one of the two treatment blocks that was randomly assigned to each decision-maker, and *period* represents the two periods that divided the 30 rounds into two equal parts. Specifically, the salvage value treatment assigned in the first period changed in the second period; thus, the interaction between *block* and *period* allow us to observe the time-dependent effect of environments on exploration behavior.

### 4.3.2 Empirical Models

We test our hypotheses using two empirical models based on linear mixed-effect modeling (LMM) approach and moderated mediation analysis. The first model, equation 4.2, tests for the

effect of decision-making environment on exploration behavior (hypotheses 4.1 and 4.2), while the second model, equations 4.3a and 4.3b, tests the mediating role of profit on the effect of environment on explorations (hypotheses 4.3 and 4.4).

Using LMM approach we estimate the following empirical model:

$$\begin{aligned}
 Exp_{it+1} = & \underbrace{\alpha_0 + \alpha_1 block_i + \alpha_2 period_t + \alpha_3 block_i \times period_t}_{\text{Fixed Effects Portion}} \\
 & + \underbrace{\alpha_4 round1_t + \alpha_5 round2_t}_{\text{Fixed Effects Portion}} + \underbrace{u_{1i} round1_t + u_{2i} round2_t + u_{3i} period_t + u_{0i}}_{\text{Random Effects Portion}} + \epsilon_{it}.
 \end{aligned} \tag{4.2}$$

where the indices  $i$  and  $t$  capture decision-maker and round, respectively. In this two-level model, *block* is a binary variable that captures the treatment sequence a decision-maker faced in the experiment (i.e.,  $block = 1$  for favorable to unfavorable environment and vice versa for  $block = 0$ ). The *period* is a binary variable that indicates the two periods of the experiment. The two round trend variables, *round1* and *round2*, designate the round number (1–15) during period 1 and 2 of the experiment, respectively. Note that by estimating block, period, and the block  $\times$  period interaction effects—as opposed to the treatment, period, and their interaction effects—we focus on a sequence dependent temporal perspective that not only isolates the effects of the treatment, but also the time-dependent effects that earlier treatments have on subjects. While the models are mathematically equivalent, focusing on this formulation allows us to build time dependent predictive margin plots (see Figures 4.2 and 4.3).

The random effects portion of the empirical model accounts for unobserved individual heterogeneity as well as potential learning and fatigue effects. Specifically, the random-intercept (i.e.,  $u_{0i}$ ) accounts for the time-invariant heterogeneity of decision-makers, and the random-coefficients (i.e.,  $u_{1i}$ ,  $u_{2i}$ , and  $u_{3i}$ ) control for potential within-period learning and fatigue effects, following the growth model approach. Finally, we drop the first round in each period because the exploration measure is based on the distance from the prior period.

### 4.3.3 Multilevel Moderated Mediation Model

The following multilevel moderated mediation model (equation 4.3) tests the mediating role of profit on the effect of the environment on explorations. The first equation (equation 4.3a) represents the mediating path between the decision-making environment (i.e., independent variable) and profit (i.e., mediator). Particularly, the decision-making environment is represented by the interaction between the treatment block and period, which captures the moderating effect of the treatment sequence on the mediation path. Further, we account for realized demand as individual realizations might impact profit.

$$\begin{aligned}
 Profit_{it} = & \underbrace{\beta_0 + \beta_1 block_i + \beta_2 period_t + \beta_3 block_i \times period_t}_{\text{Fixed Effects Portion}} \\
 & + \underbrace{\beta_4 round1_t + \beta_5 round2_t + \beta_6 realized\ demand_{i,t}}_{\text{Fixed Effects Portion}} \\
 & + \underbrace{u_{1i}^1 round1_t + u_{2i}^1 round2_t + u_{3i}^1 period_t + u_{0i}^1}_{\text{Random Effects Portion}} + \epsilon_{it}^1.
 \end{aligned} \tag{4.3a}$$

The second equation (equation 4.3b) in the moderated mediation analysis captures the mediating path between profit and exploration, as well as the direct path between decision-making environment and exploration.

$$\begin{aligned}
 Exp_{it+1} = & \underbrace{\gamma_0 + \gamma_1 Profit_{it} + \gamma_2 block_i + \gamma_3 period_t + \gamma_4 block_i \times period_t}_{\text{Fixed Effects Portion}} \\
 & + \underbrace{\gamma_5 round1_t + \gamma_6 round2_t}_{\text{Fixed Effects Portion}} + \underbrace{u_{1i}^2 round1_t + u_{2i}^2 round2_t + u_{3i}^2 period_t + u_{0i}^2}_{\text{Random Effects Portion}} + \epsilon_{it}^2.
 \end{aligned} \tag{4.3b}$$

Using the coefficients in both equations, we estimate the indirect effect of the decision-making environment on exploration behavior explained through profit performance via the following steps. First, we test whether the treatment sequence moderates the mediation paths (i.e., indirect effects) by estimating the index of moderated mediation (i.e.,  $\beta_3 \times \gamma_1$ ). If the index of moderated mediation is statistically significant ( $H_0: \beta_3 \times \gamma_1 = 0$ ) we reject the hypothesis that the mediation paths are not moderated (Hayes 2017).

Second, we estimate the following three indirect effects if we find evidence for moderated mediation. The first indirect effect is the between-subject treatment block effect (i.e.,  $\beta_1 \times \gamma_1$ ) that estimates the indirect effect based on the difference in the decision-making environment that the decision-makers were randomly assigned in the first period. The second and third indirect effects are the within-subject treatment sequence effect that estimates the indirect effect based on the change in the environment as it changes from unfavorable to favorable for block 1 (i.e.,  $\beta_2 \times \gamma_1$ ) and vice versa for block 2 (i.e.,  $(\beta_2 + \beta_3) \times \gamma_1$ ).

We preserve the multilevel structure when estimating the indirect effect (i.e., moderated mediation analysis) because without considering the clustered nature of our data, our estimates of the indirect effect will be biased (Krull and MacKinnon 2001). Specifically, in our experiment, the decision-makers were randomly assigned to a treatment block, and the decisions made by each decision-maker were dependent. Thus, we estimate the indirect effects by bootstrapping at the decision-maker level (i.e., level 2) because bootstrap resampling assumes independence of samples, which will be violated if we bootstrap at the individual decision level, i.e., level 1 (MacKinnon et al. 2007, Pituch et al. 2006). Further, we use bias-corrected bootstrap estimates, which corrects for potential bias and skewness in the distribution of bootstrap estimates that may bias the estimates when using percentile-based bootstrap (Hayes 2017, MacKinnon et al. 2004).

#### 4.4 Results

Table 4.1 reports the estimations of equations 4.2, 4.3a, and 4.3b in columns 1 to 3, respectively. Focusing on column 1, we find that decision-makers in period 1 that were randomly assigned to a favorable condition (i.e., block 2) explored less than those that were exposed to an unfavorable condition, i.e., block 1 ( $\alpha_1 = -462.107, p = 0.000$ ). Thus, failing to reject hypothesis 4.1.

In the second period, we observe a change in the environment, and we find a significant block-period interaction effect ( $\alpha_3 = 743.414, p = 0.000$ ), which indicate an asymmetry in the exploration behavior depending on the treatment sequence (i.e., favorable to unfavorable environments or vice versa). Figure 4.2 pictures the estimated marginal means for exploration behavior based on the block-period interaction model estimated in column 1 or equation 4.2. Block 1, which switched



Table 4.1: Effect of Treatment Sequence on Profit and Exploration

	Moderated Mediation Analysis		
	(1) Exploration	(2) Profit	(3) Exploration
Fixed Effects			
Profit			-0.229*** (0.011)
Block	-462.107*** (105.381)	1305.093*** (128.783)	-151.099 (99.018)
Period	-659.416*** (115.608)	2368.570*** (160.408)	-64.641 (111.004)
Block × Period	743.414*** (114.906)	-2574.583*** (137.333)	157.296 (112.335)
Round1	-3.077 (8.403)	72.103*** (10.599)	0.681 (7.878)
Round2	16.577* (8.137)	1.642 (10.889)	-3.279 (7.690)
Realized Demand		16.203*** (0.525)	
Intercept	1012.855*** (94.294)	-3741.271*** (171.442)	951.176*** (87.111)
Random Effects			
Level 2 (between-subject)			
Intercept	105046.183*** (27648.251)	241708.308*** (49759.785)	84879.512*** (23424.587)
Round1	1533.070*** (455.005)	0.000*** (0.000)	1483.744*** (411.474)
Round2	1079.894*** (390.155)	0.001 (0.138)	1099.639*** (354.381)
Period	0.000*** (0.000)	115113.157*** (57057.390)	0.000*** (0.000)
Level 1 (within-subject)			
Residual	1305612*** (36091.005)	2589932*** (70875.331)	1115968*** (30822.657)
<i>N</i>	2883	2883	2883

Standard errors in parentheses  
 \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

from unfavorable to favorable environment, responded by decreasing the degree of exploration ( $\alpha_2 = -659.416, p = 0.000$ ), while block 2 that shifted from favorable to unfavorable environment did not significantly alter their degree of exploration ( $\alpha_2 + \alpha_3 = 83.998, se = 113.900, p = 0.461$ ). Thus, failing to reject hypothesis 4.2 when decision-makers shift from unfavorable to favorable conditions but rejecting the hypothesis in the opposite direction.




Figure 4.2: Exploration Behavior by Treatment Blocks and Periods

#### 4.4.1 Multilevel Moderated Mediation Analysis

To estimate the indirect effect of decision-making environment on exploration as explained through profit, we estimated the two equations (equation 4.3a and 4.3b) and reported them in column 2 and 3 in table 4.1. Specifically, we estimated the indirect effects using 5,000 bootstrapped samples to compute the bias-corrected 95% confidence intervals (BC 95% CI). We find that the index of moderation ( $\beta_3 \times \gamma_1 = 590.721, se = 85.400, BC\ 95\% CI = [455.828, 811.289]$ ) is statistically significant, thereby rejecting the null hypothesis that the mediation paths are not moderated.

To visualize the indirect effect, we plotted predictive margins figures for column 2. Figure 4.3 represents the average profit based on the treatment blocks across periods, and we find evidence of an asymmetrical effect of environment on profit performance that suggests a moderating role of treatment sequence on the mediation paths. Further, we find that the direct effect of environment on exploration (i.e.,  $\gamma_2, \gamma_3,$  and  $\gamma_4$  in column 3 of Table 4.1) is statistically insignificant compared to figure 4.2 or coefficients  $\alpha_1, \alpha_2,$  and  $\alpha_3$  in column 1 in Table 4.1, thus implying the statistical significance of indirect effects.

The estimation of between-subject treatment block effect ( $\beta_1 \times \gamma_1 = -299.460, se = 53.353, BC\ 95\% CI = [-420.167, -211.263]$ ) suggests that decision-makers in a favorable envi-

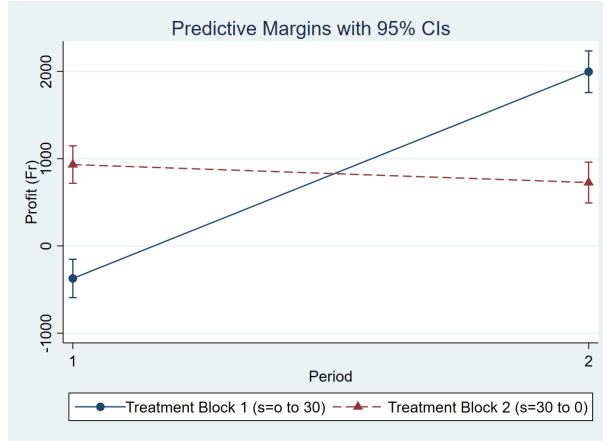


Figure 4.3: Moderated Mediation Analysis—Profits by Treatment Blocks and Periods

ronment explore less and this effect is negatively mediated by profit performance. Thus, failing to reject hypothesis 4.3.

The estimation of within-subject treatment sequence effect for block 1 ( $\beta_2 \times \gamma_1 = -543.480, se = 79.793, BC\ 95\% CI = [-739.320, -415.687]$ ) suggests that as decision-makers transition from unfavorable to favorable environment, their profit performance increase, which in turn decrease their exploration. Contrary to block 1, the decision-makers that transitioned from favorable to unfavorable environment saw statistically insignificant change in profit performance, which in turn kept the exploration behavior statistically unchanged ( $(\beta_2 + \beta_3) \times \gamma_1 = 47.271, se = 36.104, BC\ 95\% CI = [-12.658, 130.694]$ ). Considering the above two estimations, we find that the direction of environmental change significantly influences the exploration behavior as explained through profit and that this change is asymmetrical. Further, we conclude that our analysis failed to reject hypothesis 4.4 when the decision-making environment shifted from unfavorable to favorable (i.e., block 1) but rejected the hypothesis under the opposite case (i.e., block 2).

#### 4.5 Discussion and Conclusion

Through our randomized controlled experiment, we found that decision-makers are more active in exploring for better decision sets when faced with an unfavorable environment, which was defined as an environment with low expected profit. Further, we investigated the time-dependent

effect of the decision-making environment on exploration behavior by observing the change in explorations when the environment shifts from favorable to unfavorable and vice versa. As predicted, we found that exploration behavior is significantly reduced when the decision-making environment shifts from unfavorable to favorable. On the contrary, when the decision-making environment shifts from favorable to unfavorable, we find no significant change in the exploration behavior, thus rejecting our prediction based on the prospect theory.

Moreover, we tested the mechanism of how decision-making environments affect exploration by investigating how observed profit performance mediates the effect. We find that profit performance negatively mediates the effect of the decision-making environment on exploration behavior. Particularly, using between-subject analysis, we find evidence that decision-makers observe higher (lower) profit performance when exposed to a favorable (unfavorable) environment, which in turn leads to a lower (higher) exploration behavior.

However, we found that this mediation effect becomes more nuanced when we investigate the sequence of environmental shifts (i.e., a within-subject analysis focusing on the shift from favorable to unfavorable and vice versa). First, we find evidence that when decision-makers are initially exposed to an unfavorable environment, they observe (on average) a negative profit that leads them to explore more to recoup their loss. When the environment shifts favorably, in the second period, their profit performance increases, leading to a significant reduction in exploration for better decision sets. That is, the large increase in profit performance under a favorable environment easily satisfies the decision-makers leading to less explorations, and as decision fatigue sets in, their level of satisficing utility reduces, which also leads to less explorations. Our findings are inline with our predictions based on prospect theory and the theory of bounded rationality.

Second, we found that decision-makers exposed to the favorable environment first observed, on average, a higher positive profit than those exposed to the unfavorable environment, which led to lower exploration behavior. As the decision-making environment shifted to an unfavorable environment, the profit performance did not decrease and was sustained, and their level of exploration also stayed the same instead of increasing, contrary to our prediction. Perhaps this observation

can be explained by the fact that the decision-makers set profit performances observed in the first period as their reference point, and when they were facing a loss in the second period, they tried to recover from the pain coming from the loss by sustaining their prior level of exploration despite the onset of decision fatigue. Particularly, the fact that decision-makers achieve similar profit performance in the second period implies that they are using their positive performance in the first round as a reference point. Thus, they sustain their level of exploration only to the degree where they can achieve the same level of profit performance under an unfavorable environment.

We can further interpret the results in two ways. First, we find that decision-makers' exploration behavior in search of a better choice set and their satisficing level of utility that regulates the different exploration behavior is explained using prospect theory and bounded rationality. Particularly we find that change in the decision-making environment has a different impact on the exploration behavior, thus without considering the direction of the environmental change, we cannot fully explain the decision-maker's decisions. Our empirical observation leads us to the idea that rather than focusing on the characteristics of individual decision-makers to explain observed behavioral regularities, a fruitful line of inquiry would be to theorize on how the operating environment and the change in the environment may drive behavioral decision-making in a complex and dynamic environment with multiple decision-levers. This would require the use of both between-subject and within-subject research design (e.g., AB/BA repeated measures). Moreover, in essay 2, we found most of the demographic and socio-economics measures insignificant in explaining the exploration of decision sets beyond what was captured by the random intercept portion of the empirical model. This suggests that our experimental design and execution was successful in inducing our behavior of interest (i.e., exploration behavior of decision-makers) and kept confounds from "innate characteristics largely irrelevant" (Friedman and Sunder 1994, p. 13).

Second, our results are inline with the three main assumptions (i.e., dominance, salience, and monotonicity of incentives) of induced value theory (Smith 1976, 1982) that forms the basic structure in which all incentivized experiments are established upon. Particularly, induced value theory states that subjects' effort in the experimental task must be correlated with the incentives in or-

der to measure the subjects' behavior of interest without confounding the measure with subjects' preferences (Smith 1976). In our experiment, we were interested in the exploration behavior of decision-makers, and to measure such behavior, we tied the profit performance with the incentive pay. From our results, we found that the manipulation of the treatment (i.e., change in salvage value and the environment) influenced realized profit performance, which was tied to incentive pay that affected exploration decisions. Thus, profit was salient to the decision-makers, and they explored the decision sets in order to find better profits (i.e., dominance assumption). Further, our finding that profit performance, on average, only ever increased for both decision-makers in blocks 1 and 2 suggests that subjects preferred higher profit and actively searched for it (i.e., monotonicity assumption).

Our study is not without limitations. In this study, we do not investigate nor test the specific exploration strategies used by the decision-makers. Instead, we focus on testing the existence of general exploration and exploitation behavior and how that can be explained using two individual-level decision-making theories—prospect theory and bounded rationality. For future research, we will expand our research to include specific exploration strategies and how decision-makers utilize them. Also, we will measure an individual's satisficing level of utility, similar to Caplin et al. (2011), to better understand at the decision-level how satisficing affects explorations.

## 5. CONCLUSION

The three essays in this dissertation investigated how managers leverage heuristics in a complex and dynamic business environment, focusing on retail as the primary business context. The first essay found that managers can complement algorithms (e.g., ASO system) by systematically and consistently improving restocking recommendations, yielding an average savings of 2.5% of the cost of goods sold. Our finding that system recommendations can be improved signals the presence of blind spots in ASO system algorithms, consequence either of restricted access to information or constrained solution space, and that exogenous information to which managers have access can inform decisions that reduce inventory costs. Our results suggest that establishing default choices for an ASO system enables managers to participate in a complementary manner that imparts flexibility (e.g., adaptability) to an efficient but inflexible decision-making algorithm. Our findings are in contrast to the majority of the studies in the inventory decisions literature that argue that managers are a source of behavioral bias and a liability (Becker-Peth and Thonemann 2019, Tversky and Kahneman 1974).

This essay makes several important contributions. First, it is one of few studies of behavioral decision-making that assess the performance of managers' inventory decisions by observing their impact in the real world. This is in contrast to most research on behavioral aspects of inventory decisions, predominantly laboratory-based experiments that strictly control for environmental factors (Zhang and Siemsen 2019). Second, we establish that managers can complement an ASO system and consistently, in real time, improve decision performance. This finding is contrary to most studies in the inventory decisions literature, which argue that managers are a source of bias and a liability (Becker-Peth and Thonemann 2019). Previous research, albeit limited, has shown flexibility and adaptability to be among the strengths of human decision makers (e.g, Gigerenzer 2008, Lawrence et al. 2006). We should explore ways to leverage these human capabilities, and the contexts in which they might be most useful. This essay illustrates a specific instance in which these capabilities are being effectively deployed. Last, and more specifically, this study contributes

to research on how to integrate decisions from algorithms and human decision makers by identifying the *default choice* mechanism as a synergistic alternative to existing approaches (Blattberg and Hoch 1990, Lee and Siemsen 2017).

The second essay found that the salvaging capability causally affects return policy decisions. Further, the process theory developed in this essay demonstrates the complexity of the decision-making process in a multi-lever operational environment and how decision-makers cope with this complexity. Our investigation revealed that given an operating condition (i.e., a salvage value), individuals use a conditional decision-making heuristic, in which a biased decision over one lever influences the decisions over the other levers. When the operating condition changes, individuals react in a boundedly rational manner to the change to make order quantity decisions, follow the dual-entitlement principle to adjust refund amounts, and demonstrate learning through exploration and satisficing in setting prices.

This essay makes several important contributions to the behavioral operations management and consumer returns literature streams. First, to the best of my knowledge, this research is the first to examine operational decision-making via multiple (more than two) levers, stepping beyond Kocabıyıkoglu et al. (2016) and Ramachandran et al. (2018) who jointly examined price and order quantity decisions. Second, our research highlights the importance of examining transitions in decision-making heuristics based on the operating conditions (Oliva et al. 2022) and time-dependent effects. In this respect, we view our investigation in this essay as a contribution to the behavioral operations management literature as an application of statistical induction combined with abductive reasoning to provide plausible explanations to the observed phenomena, a research mode that Rozeboom (1997) refers as explanatory induction. Third, we extended the research on behavioral inventory decision-making to also include consumer product returns, which is a critical managerial issue in today's retail environment.

The third essay found that decision-makers are more active in exploring better decision sets when faced with an unfavorable environment. Also, investigating the time-dependent effects, we found that exploration behavior is significantly reduced when the decision-making environment



shifts from unfavorable to favorable. On the contrary, when the decision-making environment shifts from favorable to unfavorable, we found no significant change in the exploration behavior, thus rejecting our prediction based on the prospect theory. Moreover, we found that profit performance negatively mediates the effect of the decision-making environment on exploration behavior. However, we found that this mediation effect becomes more nuanced when we investigate the sequence of environmental shifts (i.e., a within-subject analysis focusing on the shift from favorable to unfavorable and vice versa). Our empirical observations leads us to the idea that rather than focusing on the characteristics of individual decision-makers to explain observed behavioral regularities, a fruitful line of inquiry would be to theorize on how the operating environment and the change in the environment may drive behavioral decision-making in a complex and dynamic environment.

Overall, this dissertation stresses the importance of including managers in the decision-making system and understanding human heuristics even in the era of data-driven decision-making. This dissertation opens many opportunities for future research. First, we can further explore how machine learning and data analytics can be better used to facilitate the synergies between managerial heuristics and the scalability of algorithmic decision-making. I believe this research has much to contribute because of the lack of understanding of managerial heuristics in real business environments and how they can be better utilized in operational decisions. Second, we can continue to expand the boundary of consumer returns and behavioral operations by focusing on decision-making in retail operations. Most behavioral studies on retail operations focus on how managers make decisions under uncertainty using a single managerial lever (i.e., order quantity decisions). However, retail managers have many operational decisions they need to make simultaneously in practice. We can contribute to behavioral research in retail operations by expanding the findings from single decision-lever studies and investigating how additional decision-levers provide insight on how managerial decisions are made in practice. Moreover, our research on consumer return policy focused on decision-making with firm external factors largely controlled for or excluded from the research scope (i.e., assumption of monopolistic market). We can include these factors in future studies.

## REFERENCES

- Abbey JD, Meloy MG (2017) Attention by design: Using attention checks to detect inattentive respondents and improve data quality. *Journal of Operations Management* 53:63–70.
- Abdulla H, Ketzenberg M, Abbey JD (2019) Taking stock of consumer returns: A review and classification of the literature. *Journal of Operations Management* 65(6):560–605.
- Abrams R (2016) Liberal return policies for consumers can reduce retail workers’ pay. *The New York Times* URL <https://www.nytimes.com/2016/06/14/business/liberal-returns-policies-for-consumers-can-reduce-retail-workers-pay.html>, accessed on February 28, 2022.
- Adra SF, Fleming PJ (2011) Diversity management in evolutionary many-objective optimization. *IEEE Transactions on Evolutionary Computation* 15(2):183–195.
- Akçay Y, Boyacı T, Zhang D (2013) Selling with money-back guarantees: The impact on prices, quantities, and retail profitability. *Production and Operations Management* 22(4):777–791.
- Altug MS, Aydinliyim T (2016) Counteracting strategic purchase deferrals: The impact of online retailers’ return policy decisions. *Manufacturing & Service Operations Management* 18(3):376–392.
- Angrist JD, Pischke JS (2009) *Mostly harmless econometrics: An empiricist’s companion* (Princeton, NJ: Princeton University Press).
- Argote L, Ingram P (2000) Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes* 82(1):150–169.
- Armstrong JS (1983) Relative accuracy of judgemental and extrapolative methods in forecasting annual earnings. *Journal of Forecasting* 2(4):437–447.
- Athey S (2017) Beyond prediction: Using big data for policy problems. *Science* 355(6324):483–485.
- Audia PG, Locke EA, Smith KG (2000) The paradox of success: An archival and a laboratory study of strategic persistence following radical environmental change. *Academy of Management journal* 43(5):837–853.
- Balakrishnan A, Pangburn MS, Stavroulakis E (2004) “Stack them high, let’em fly”: Lot-sizing policies when inventories stimulate demand. *Management Science* 50(5):630–644.

- Barber BM, Odean T (2008) All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies* 21(2):785–818.
- Barry DM, Pitz GF (1979) Detection of change in nonstationary, random sequences. *Organizational Behavior and Human Performance* 24(1):111–125.
- Becker-Peth M, Thonemann UW (2019) Behavioral inventory decisions: The newsvendor and other inventory settings. Donohue K, Katok E, Leider S, eds., *The Handbook of Behavioral Operations*, 393–432 (Wiley, Hoboken, NJ).
- Becker-Peth M, Thonemann UW, Gully T (2018) A note on the risk aversion of informed newsvendors. *Journal of the Operational Research Society* 69(7):1135–1145.
- Begley S, Hancock B, Kilroy T, Kohli S (2019) Automation in retail: An executive overview for getting ready. URL <https://www.mckinsey.com/industries/retail/our-insights/automation-in-retail-an-executive-overview-for-getting-ready>, accessed on February 28, 2022.
- 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.
- Bhatia S (2019) Predicting risk perception: New insights from data science. *Management Science* 65(8):3800–3823.
- Biery ME (2017) These industries generate the lowest profit margins. URL <https://www.forbes.com/sites/sageworks/2017/09/24/these-industries-generate-the-lowest-profit-margins/#6e2a3b22f49d>, accessed on February 28, 2022.
- Blanchard D (2007) Supply chains also work in reverse. *IndustryWeek* URL <https://www.industryweek.com/supply-chain/planning-forecasting/article/21954433/supply-chains-also-work-in-reverse>, accessed on February 28, 2022.
- Blattberg RC, Hoch SJ (1990) Database models and managerial intuition: 50% models + 50% manager. *Management Science* 36(8):887–899.
- Bolton GE, Chen Y (2019) Other-regarding behavior: Fairness, reciprocity, and trust. Donohue K, Katok E, Leider S, eds., *The Handbook of Behavioral Operations*, chapter 6, 199–229 (John Wiley & Sons, Hoboken, NJ).
- Bolton GE, 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, ISSN 1523-4614.
- Bolton GE, Ockenfels A, Thonemann UW (2012) Managers and students as newsvendors. *Management Science* 58(12):2225–2233, ISSN 0025-1909.
- Bordalo P, Gennaioli N, Shleifer A (2012) Saliency theory of choice under risk. *Quarterly Journal of Economics* 127(3):1243–1285.
- Bostian AA, Holt CA, Smith AM (2008) Newsvendor “pull-to-center” effect: Adaptive learning in a laboratory experiment. *Manufacturing Service Operations Management* 10(4):590–608.
- Brown LD (1996) Analyst forecasting errors and their implications for security analysis: An alternative perspective. *Financial Analysts Journal* 52(1):40–47.
- Bushway S, Johnson BD, Slocum LA (2007) Is the magic still there? the use of the heckman two-step correction for selection bias in criminology. *Journal of Quantitative Criminology* 23(2):151–178.
- Cachon GP, Terwiesch C (2012) *Matching Supply with Demand: An Introduction to Operations Management* (McGraw-Hill Education), 3rd ed edition.
- Caplin A, Dean M, Martin D (2011) Search and satisficing. *American Economic Review* 101(7):2899–2922.
- Carter CR, Kaufmann L, Michel A (2007) Behavioral supply management: A taxonomy of judgment and decision-making biases. *International Journal of Physical Distribution Logistics Management* 37(8):631–669.
- Certo ST, Busenbark JR, Woo HS, Semadeni M (2016) Sample selection bias and heckman models in strategic management research. *Strategic Management Journal* 37(13):2639–2657.
- Chao G, Cheung J, Haller K, Lee J, Tyson L, Wong CK (2019) The coming ai revolution in retail and consumer products. Report, IBM Corporation, URL <https://www.ibm.com/thought-leadership/institute-business-value/report/ai-retail-consumer-products#>, accessed on February 28, 2022.
- Chen DL, Schonger M, Wickens C (2016) otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9:88–97.
- Cheng A (2015) Consumers return \$642.6 billion in goods each year URL <https://www.marketwatch.com/story/consumers-return-6426-billion-in-goods-each-year-2015-06-18>, accessed on February 28, 2022.

- Chou P, Chuang HHC, Chou YC, Liang TP (2022) Predictive analytics for customer repurchase: Interdisciplinary integration of buy till you die modeling and machine learning. *European Journal of Operational Research* 296(2):635–651.
- Chuang HH, Chou Y, Oliva R (2021) Cross-item learning for volatile demand forecasting: An intervention with predictive analytics. *Journal of Operations Management* Forthcoming.
- Constable H (2017) Your brand new returns end up in landfill. *BBC* URL <https://www.bbcearth.com/blog/?article=your-brand-new-returns-end-up-in-landfill>, accessed on April 26, 2021.
- Corsten D, Gruen T (2003) Desperately seeking shelf availability: An examination of the extent, the causes, and the efforts to address retail out-of-stocks. *International Journal of Retail Distribution Management* 31(12):605–617.
- Craig N, DeHoratius N, Raman A (2016) The impact of supplier inventory service level on retailer demand. *Manufacturing Service Operations Management* 18(4):461–474.
- Črepinšek M, Liu SH, Mernik M (2013) Exploration and exploitation in evolutionary algorithms: A survey. *ACM computing surveys (CSUR)* 45(3):1–33.
- Crosan R, Donohue K, Katok E, Sterman J (2014) Order stability in supply chains: Coordination risk and the role of coordination stock. *Production and Operations Management* 23(2):176–196.
- Crosan R, Gneezy U (2009) Gender differences in preferences. *Journal of Economic literature* 47(2):448–74.
- Curran PJ, Bauer DJ (2011) The disaggregation of within-person and between-person effects in longitudinal models of change. *Annual Review of Psychology* 62:583–619.
- Davis S, Hagerty M, Gerstner E (1998) Return policies and the optimal level of “hassle”. *Journal of Economics and Business* 50(5):445–460.
- De Vericourt F, Jain K, Bearden JN, Filipowicz A (2013) Sex, risk and the newsvendor. *Journal of Operations Management* 31(1-2):86–92.
- DeHoratius N, Raman A (2008) Inventory record inaccuracy: An empirical analysis. *Management Science* 54(4):627–641.
- Dellavigna S (2009) Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2):315–372.

- DeMiguel V, Garlappi L, Uppal R (2009) Optimal versus naive diversification : How inefficient is the 1/n portfolio strategy? *The Review of Financial Studies* 22(5):1915–1953.
- Dietvorst BJ, Simmons JP, Massey C (2018) Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science* 64(3):1155–1170.
- DiMaggio PJ (1995) Comments on “What theory is not”. *Administrative Science Quarterly* 40(3):391–397.
- Dodes R (2010) Blizzard heats up sales for snow gear. URL <https://www.wsj.com/articles/SB10001424052970204527804576044053293555380>, accessed on February 28, 2022.
- Dogru AK, Keskin BB (2020) Ai in operations management: applications, challenges and opportunities. *Journal of Data, Information and Management* 2(2):67–74.
- Donohue K, Schultz K (2019) The future is bright: Recent trends and emerging topics in behavioral operations. Donohue K, Katok E, Leider S, eds., *The Handbook of Behavioral Operations*, chapter 18, 619–651 (John Wiley & Sons, Inc., Hoboken, NJ).
- Donohue K, Özer Zheng Y (2020) Behavioral operations: Past, present, and future. *Manufacturing Service Operations Management* 22(1):191–202.
- Ederer F, Manso G (2013) Is pay for performance detrimental to innovation? *Management Science* 59(7):1496–1513.
- Edmondson A (1999) Psychological safety and learning behavior in work teams. *Administrative Science Quarterly* 44(2):350–383.
- Eisenhardt KM, Furr NR, Bingham CB (2010) Crossroads—microfoundations of performance: Balancing efficiency and flexibility in dynamic environments. *Organization science* 21(6):1263–1273.
- Elster J (1989) Social norms and economic theory. *Journal of Economic Perspectives* 3(4):99–117.
- Epley N, Gilovich T (2001) Putting adjustment back in the anchoring and adjustment heuristic: Differential processing of self-generated and experimenter-provided anchors. *Psychological Science* 12(5):391–396.
- Eroglu C, Williams BD, Waller MA (2013) The backroom effect in retail operations. *Production and Operations Management* 22(4):915–923.
- Evans JR, Mathur A (2014) Retailing and the period leading up to the great recession: A model and a 25-year financial ratio analysis of us retailing. *The International Review of Retail, Distribution and Consumer Research* 24(1):30–58.

- Falk A, Becker A, Dohmen T, Enke B, Huffman D, Sunde U (2018) Global evidence on economic preferences. *The Quarterly Journal of Economics* 133(4):1645–1692.
- Falk A, Becker A, Dohmen TJ, Huffman D, Sunde U (2016) The preference survey module: A validated instrument for measuring risk, time, and social preferences. *IZA Discussion Paper 9674, IZA Institute of Labor Economics* .
- Feiler DC, Tong JD, Larrick RP (2013) Biased judgment in censored environments. *Management Science* 59(3):573–591.
- Felix I, Kuntze C, Silva L, Benoliel ET (2018) The route to no-touch planning: Taking the human error out of supply-chain planning. URL <https://www.mckinsey.com/business-functions/operations/our-insights/the-route-to-no-touch-planning>, accessed on February 28, 2022.
- Fildes R, Goodwin P, Lawrence M, Nikolopoulos K (2009) Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting* 25(1):3–23.
- Flicker B (2019) Managerial insight and “optimal” algorithms. Working paper, University of South Carolina, Columbia, South Carolina .
- Fréchette GR (2015) Laboratory experiments: Professionals versus students. *Handbook of Experimental Economic Methodology*, chapter 17 (Oxford University Press, New York).
- Friedman D, Sunder S (1994) *Experimental Methods: A Primer for Economists* (Cambridge University Press, New York).
- Garbarino E, Slonim R (2009) The robustness of trust and reciprocity across a heterogeneous us population. *Journal of Economic Behavior & Organization* 69(3):226–240.
- Gaur V, Fisher ML, Raman A (2005) An econometric analysis of inventory turnover performance in retail services. *Management Science* 51(2):181–194.
- Gaur V, Kesavan S, Raman A (2014) Retail inventory: Managing the canary in the coal mine. *California Management Review* 56(2):55–76.
- Gershman SJ (2019) Uncertainty and exploration. *Decision* 6(3):277.
- Gigerenzer G (2008) Why heuristics work. *Perspectives on Psychological Science* 3(1):20–29.
- Gigerenzer G, Gaissmaier W (2011) Heuristic decision making. *Annual Review of Psychology* 62:451–482.

- Gigerenzer G, Selten R (2001) *Bounded Rationality: The Adaptive Toolbox* (Cambridge, Massachusetts and London: The MIT Press).
- Gigerenzer G, Todd PM, ABC Research Group (1999) *Simple Heuristics That Make Us Smart* (Oxford University Press, New York and Oxford, UK).
- Gil D, Hobson S, Mojsilović A, Puri R, Smith JR (2020) AI for management: An overview. *The Future of Management in an AI World* 3–19.
- Gino F, Pisano G (2008) Toward a theory of behavioral operations. *Manufacturing Service Operations Management* 10(4):676–691.
- Gratton C, Shibli S, Coleman R (2006) The economic impact of major sports events: A review of ten events in the uk. *The Sociological Review* 52(2):41–58.
- Graves SC, Willems SP (2003) Supply chain design: Safety stock placement and supply chain configuration. de Kok AG, Graves SC, eds., *Handbooks in Operations Research and Management Science*, volume 11, 95–132 (Elsevier, Amsterdam).
- Gray RM, Neuhoff DL (1998) Quantization. *IEEE Transactions on Information Theory* 44(6):2325–2383.
- Hadley G, Whitin TM (1963) *Analysis of inventory systems* (Prentice Hall, Englewood Cliffs, NJ).
- Hardy III JH, Day EA, Hughes MG, Wang X, Schuelke MJ (2014) Exploratory behavior in active learning: A between-and within-person examination. *Organizational Behavior and Human Decision Processes* 125(2):98–112.
- Hayes AF (2017) *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (The Guilford Press), 2nd ed edition.
- Heckman JJ (1976) The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement* 5(4):475–492, ISSN 0691003637.
- Heckman JJ (1979) Sample selection bias as a specification error. *Econometrica* 47(1):153–161.
- Helo P, Hao Y (2021) Artificial intelligence in operations management and supply chain management: an exploratory case study. *Production Planning & Control* 1–18.
- Holland PW (1986) Statistics and causal inference. *Journal of the American Statistical Association* 81(396):945–960.
- Huang Y, Hutchinson JW (2013) The roles of planning, learning, and mental models in repeated dynamic



- decision making. *Organizational Behavior and Human Decision Processes* 122(2):163–176.
- Johnson EJ, Goldstein D (2003) Do defaults save lives? *Science* 302(5649):1338–1339.
- Jones B, Kenward MG (2015) *Design and analysis of cross-over trials* (CRC Press, Boca Raton, FL).
- Kaelbling LP, Littman ML, Moore AW (1996) Reinforcement learning: A survey. *Journal of artificial intelligence research* 4:237–285.
- Kahn WA (1990) Psychological conditions of personal engagement and disengagement at work. *Academy of Management Journal* 33(4):692–724.
- Kahneman D, Knetsch JL, Thaler RH (1986) Fairness and the assumptions of economics. *Journal of business* S285–S300.
- Kahneman D, Tversky A (1972) Subjective probability: A judgment of representativeness. *Cognitive psychology* 3(3):430–454.
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* 47(2):263–291.
- Kalapurakal R, Dickson PR, Urbany JE (1991) Perceived price fairness and dual entitlement. *NA - Advances in Consumer Research Volume 18*, 788–793 (Association for Consumer Research, Provo, UT).
- Katsikopoulos KV, Gigerenzer G (2013) Behavioral operations management: A blind spot and a research program. *Journal of Supply Chain Management* 49(1):3–7.
- Kiesler S, Sproull L (1982) Managerial response to changing environments: Perspectives on problem sensing from social cognition. *Administrative Science Quarterly* 27(4):548–570.
- Kocabıyıkoglu A, Göğüş CI, Gönül MS (2016) Decision making and the price setting newsvendor: experimental evidence. *Decision Sciences* 47(1):157–186.
- Kotz S, Balakrishnan N, Johnson NL (2004) *Continuous multivariate distributions, Volume 1: Models and applications*, volume 1 (John Wiley & Sons Inc., New York).
- Kremer M, Moritz B, Siemsen E (2011) Demand forecasting behavior: System neglect and change detection. *Management Science* 57(10):1827–1843.
- Krommyda IP, Skouri K, Konstantaras I (2015) Optimal ordering quantities for substitutable products with stock-dependent demand. *Applied Mathematical Modelling* 39(1):147–164.
- Krull JL, MacKinnon DP (2001) Multilevel modeling of individual and group level mediated effects. *Multivariate behavioral research* 36(2):249–277.

- Laureiro-Martínez D, Brusoni S, Canessa N, Zollo M (2015) Understanding the exploration–exploitation dilemma: An fmri study of attention control and decision-making performance. *Strategic management journal* 36(3):319–338.
- Laureiro-Martínez D, Brusoni S, Zollo M (2010) The neuroscientific foundations of the exploration- exploitation dilemma. *Journal of Neuroscience, Psychology, and Economics* 3(2):95.
- Lawrence M, Goodwin P, O’Connor M, Önkál D (2006) Judgmental forecasting: A review of progress over the last 25 years. *International Journal of Forecasting* 22(3):493–518.
- Lee YS, Seo YW, Siemsen E (2018) Running behavioral operations experiments using amazon’s mechanical turk. *Production and Operations Management* 27(5):973–989.
- Lee YS, Siemsen E (2017) Task decomposition and newsvendor decision making. *Management Science* 63(10):3226–3245.
- Leung SF, Yu S (1996) On the choice between sample selection and two-part models. *Journal of Econometrics* 72(1):197–229.
- Levinthal DA (2011) A behavioral approach to strategy—what’s the alternative? *Strategic Management Journal* 32(13):1517–1523.
- Levinthal DA, March JG (1993) The myopia of learning. *Strategic management journal* 14(S2):95–112.
- Li K, Prabhala NR (2007) *Self-selection models in corporate finance*, volume 1, 37–86 (Elsevier, Oxford, UK).
- Lim JS, O’Connor M (1996) Judgmental forecasting with time series and causal information. *International Journal of Forecasting* 12(1):139–153.
- Littell RC, Pendergast J, Natarajan R (2000) Modelling covariance structure in the analysis of repeated measures data. *Statistics in Medicine* 19(13):1793–1819.
- MacKinnon DP, Fairchild AJ, Fritz MS (2007) Mediation analysis. *Annu. Rev. Psychol.* 58:593–614.
- MacKinnon DP, Lockwood CM, Williams J (2004) Confidence limits for the indirect effect: Distribution of the product and resampling methods. *Multivariate behavioral research* 39(1):99–128.
- Madhawa K, Murata T (2020) Active learning for node classification: an evaluation. *Entropy* 22(10):1164.
- March JG (1991) Exploration and exploitation in organizational learning. *Organization Science* 2(1):71–87.
- Massey C, Wu G (2005) Detecting regime shifts: The causes of under- and overreaction. *Management Science* 51(6):932–947.

- Matsumoto Y, Yamagishi T, Li Y, Kiyonari T (2016) Prosocial behavior increases with age across five economic games. *PloS one* 11(7):e0158671.
- McGinley B, Maher J, O’Riordan C, Morgan F (2011) Maintaining healthy population diversity using adaptive crossover, mutation, and selection. *IEEE Transactions on Evolutionary Computation* 15(5):692–714.
- Mišić VV, Perakis G (2020) Data analytics in operations management: A review. *Manufacturing & Service Operations Management* 22(1):158–169.
- Mitchell ML, Jolley JM (2010) *Research design explained* (Wadsworth, Belmont, CA).
- Mohr LB (1982) *Explaining organizational behavior* (Jossey-Bass, San Francisco).
- Monge PR (1990) Theoretical and analytical issues in studying organizational processes. *Organization Science* 1(4):406–430.
- Moon A, Nelson LD (2020) The uncertain value of uncertainty: When consumers are unwilling to pay for what they like. *Management Science* 66(10):4686–4702.
- Moritz B, Siemsen E, Kremer M (2014) Judgmental forecasting: Cognitive reflection and decision speed. *Production and Operations Management* 23(7):1146–1160.
- Moritz BB, Hill AV, Donohue KL (2013) Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management* 31(1-2):72–85.
- Mussweiler T, Englich B (2005) Subliminal anchoring: Judgmental consequences and underlying mechanisms. *Organizational Behavior and Human Decision Processes* 98(2):133–143.
- Mussweiler T, Strack F (1999) Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model. *Journal of Experimental Social Psychology* 35(2):136–164.
- Nadkarni S, Barr PS (2008) Environmental context, managerial cognition, and strategic action: An integrated view. *Strategic Management Journal* 29(13):1395–1427.
- National Retail Federation (2021) \$428 billion in merchandise returned in 2020 URL <https://nrf.com/media-center/press-releases/428-billion-merchandise-returned-2020>, accessed on February 28, 2022.
- Oh HK, Abdulla H, Oliva R (2021) Better together? How managers can complement algorithms. Working paper, Texas A&M University, College Station, Texas.
- Oh HK, Oliva R (2021) Return of the behavioral newsvendor: An experimental analysis of consumer return

- policy decisions, <https://ssrn.com/abstract=4018719>.
- Oliva R (2019) Intervention as a research strategy. *Journal of Operations Management* 65(7):710–724.
- Oliva R, Abdulla H, Goncalves P (2022) Do managers overreact when in backlog? Evidence of scope neglect from a supply chain experiment. *Manufacturing & Service Operations Management*. Forthcoming.
- Oliva R, Watson N (2009) Managing functional biases in organizational forecasts: A case study of consensus forecasting in supply chain planning. *Production and Operations Management* 18(2):138–151.
- Optoro (2018) 2018 optoro impact report URL <https://www.optoro.com/2018-impact-report/>, accessed on February 28, 2022.
- Pearl J, Mackenzie D (2018) *The Book of Why: The New Science of Cause and Effect* (Basic Books, New York).
- Peer E, Brandimarte L, Samat S, Acquisti A (2017) Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70:153–163.
- Peirce CS (1992) Ketner KL, ed., *Reasoning and the Logic of Things: The Cambridge Conferences Lectures of 1898* (Harvard University Press, Cambridge, MA).
- Pentland BT (1999) Building process theory with narrative: From description to explanation. *Academy of Management Review* 24(4):711–724.
- Pitt MA, Myung IJ, Zhang S (2002) Toward a method of selecting among computational models of cognition. *Psychological Review* 109(3):472–491.
- Pituch KA, Stapleton LM, Kang JY (2006) A comparison of single sample and bootstrap methods to assess mediation in cluster randomized trials. *Multivariate Behavioral Research* 41(3):367–400.
- Porteus EL (2002) *Foundations of Stochastic Inventory Theory* (Stanford University Press, Stanford, CA).
- Posen HE, Levinthal DA (2012) Chasing a moving target: Exploitation and exploration in dynamic environments. *Management science* 58(3):587–601.
- Prak D, Teunter R, Syntetos A (2016) On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research* 256(2):454–461.
- Preacher KJ, Hayes AF (2008) Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods* 40(3):879–891.
- Puhani PA (2000) The heckman correction for sample selection and its critique. *Journal of Economic Surveys* 14(1):53–68.

- Quispe-Torreblanca EG, Stewart N, Gathergood J, Loewenstein G (2019) The red, the black, and the plastic: Paying down credit card debt for hotels, not sofas. *Management Science* 65(11):5392–5410.
- Ramachandran K, Tereyağoğlu N, Xia Y (2018) Multidimensional decision making in operations: An experimental investigation of joint pricing and quantity decisions. *Management Science* 64(12):5544–5558.
- Raudenbush SW, Bryk AS (2002) *Hierarchical linear models: Applications and data analysis methods* (SAGE Publications Inc., Thousand Oaks, CA).
- Reagan C (2013) What's behind the rush into the low-margin grocery business. URL <http://www.cnn.com/id/100794988>, accessed on February 28, 2022.
- Ren Y, Croson R (2013) Overconfidence in newsvendor orders: An experimental study. *Management Science* 59(11):2502–2517.
- Roca M, Hogarth RM, Maule AJ (2006) Ambiguity seeking as a result of the status quo bias. *Journal of Risk and Uncertainty* 32(3):175–194.
- Rozeboom WW (1997) Good science is abductive, not hypothetico-deductive. Harlow LL, Mulaik SA, Steiger JH, eds., *What if there were no significance tests?*, chapter 13, 335–392 (Psychology Press, New York).
- Rumyantsev S, Netessine S (2007) What can be learned from classical inventory models? a cross-industry exploratory investigation. *Manufacturing Service Operations Management* 9(4):409–429.
- Samuelson PA (1937) A note on measurement of utility. *The Review of Economic Studies* 4(2):155–161.
- Samuelson PA (1938) A note on the pure theory of consumer's behaviour. *Economica* 5(17):61–71.
- Samuelson PA (1948) Consumption theory in terms of revealed preference. *Economica* 15(60):243–253.
- Samuelson W, Zeckhauser R (1988) Status quo bias in decision making. *Journal of Risk and Uncertainty* 1(1):7–59.
- Sanders NR, Ritzman LP (1992) The need for contextual and technical knowledge in judgmental forecasting. *Journal of Behavioral Decision Making* 5(1):39–52.
- Schweitzer ME, Cachon GP (2000) Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science* 46(3):404–420.
- Schwenk CR (1984) Cognitive simplification processes in strategic decision-making. *Strategic management journal* 5(2):111–128.
- Shang G, Pekgün P, Ferguson M, Galbreth M (2017) How much do online consumers really value free

- product returns? Evidence from ebay. *Journal of Operations Management* 53:45–62.
- Shaw KL (1996) An empirical analysis of risk aversion and income growth. *Journal of Labor Economics* 14(4):626–653.
- Shin H, Collier DA, Wilson DD (2000) Supply management orientation and supplier/buyer performance. *Journal of Operations Management* 18(3):317–333, ISSN 1898805105.
- Siemsen E, Roth AV, Balasubramanian S, Anand G (2009) The influence of psychological safety and confidence in knowledge on employee knowledge sharing. *Manufacturing and Service Operations Management* 11(3):429–447.
- Silver EA (1981) Operations research in inventory management: A review and critique. *Operations Research* 29(4):628–645.
- Silver EA, Pyke DF, Peterson R (1998) *Inventory Management and Production Planning and Scheduling* (Wiley, New York), 3rd ed edition.
- Simon HA (1955) A behavioral model of rational choice. *The Quarterly Journal of Economics* 69(1):99–118.
- Simon HA (1957) *Administrative behavior: A study of decision-making processes in administrative organization* (Macmillan, New York), 2 edition.
- Simon HA (1978) Rationality as process and as product of thought. *The American Economic Review* 68(2):1–16.
- Sloman S, Goldstone RL, Gonzalez C (2019) Complex exploration dynamics from simple heuristics in a collective learning environment. *CogSci*, 2818–2824.
- Smith VL (1976) Experimental economics: Induced value theory. *The American Economic Review* 66(2):274–279.
- Smith VL (1982) Microeconomic systems as an experimental science. *The American Economic Review* 72(5):923–955.
- Stavrulaki E (2011) Inventory decisions for substitutable products with stock-dependent demand. *International Journal of Production Economics* 129(1):65–78.
- Steinker S, Hoberg K, Thonemann UW (2017) The value of weather information for e-commerce operations. *Production and Operations Management* 26(10):1854–1874.
- Sterman J, Oliva R, Linderman KW, Bendoly E (2015) System dynamics perspectives and modeling oppor-

- tunities for research in operations management. *Journal of Operations Management* 39–40:1–5.
- Sterman JD (1989) Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science* 35(3):321–339.
- Su X (2009) Consumer returns policies and supply chain performance. *Manufacturing & Service Operations Management* 11(4):595–612.
- Tarafdar M, Beath CM, Ross JW (2019) Using ai to enhance business operations. *MIT Sloan Management Review* 60(4):37–44.
- Teunter RH, Babai MZ, Syntetos AA (2010) Abc classification: Service levels and inventory costs. *Production and Operations Management* 19(3):343–354.
- Thaler RH, Sunstein CR (2009) *Nudge: Improving decisions about health, wealth, and happiness* (Penguin Books, London).
- Thrun SB (1992) The role of exploration in learning control. White DA, Sofge DA, eds., *Handbook of Intelligent Control: Neural, Fuzzy and Adaptive Approaches*, 1–27 (Van Nostrand Reinhold, Florence, Kentucky).
- Tiefenbeck V, Goette L, Degen K, Tasic V, Fleisch E, Lalive R, Staake T (2018) Overcoming salience bias: How real-time feedback fosters resource conservation. *Management Science* 64(3):1458–1476.
- Todd P, Gigerenzer G (2012) *Ecological rationality: Intelligence in the world* (Oxford University Press, New York).
- Tong J, Feiler D (2016) A behavioral model of forecasting: Naive statistics on mental samples. *Management Science* 63(11):3531–3997.
- Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.
- Vaidyanathan R, Aggarwal P (2003) Who is the fairest of them all? an attributional approach to price fairness perceptions. *Journal of Business Research* 56(6):453–463.
- Van Donselaar KH, Gaur V, Van Woensel T, Broekmeulen RACM, Fransoo JC (2010) Ordering behavior in retail stores and implications for automated replenishment. *Management Science* 56(5):766–784.
- Wagner HM (2002) And then there were none. *Operations Research* 50(1):217–226.
- Wallenstein S, Fisher AC (1977) The analysis of the two-period repeated measurements crossover design with application to clinical trials. *Biometrics* 33:261–269.

- Wilson R (2006) The economic impact of local sport events: significant, limited or otherwise? a case study of four swimming events. *Managing Leisure* 11(1):57–70.
- Zhang Y, Siemsen E (2019) A meta-analysis of newsvendor experiments: Revisiting the pull-to-center asymmetry. *Production and Operations Management* 28(1):140–156.



APPENDIX A

BETTER TOGETHER? HOW MANAGERS CAN COMPLEMENT ALGORITHMS  
 Table A1: Sample Selection Model with Varying Inventory Holding and Substitution Rates

	(1)	(2)	(3)	(4)
	Cost Savings	Cost Savings	Cost Savings	Cost Savings
Inventory Holding Rate	12%	12%	36%	36%
Substitution Rate	25%	75%	25%	75%
Outcome Equation				
Demand Volatility	0.4756*** (0.1142)	0.1742*** (0.0445)	0.4664*** (0.1142)	0.1546*** (0.0398)
Agreed Leadtime	0.0295** (0.0098)	0.0133*** (0.0040)	0.0405*** (0.0098)	0.0253*** (0.0037)
Supplier Lateness	0.1180 (0.1693)	0.1358* (0.0640)	0.4471* (0.1740)	0.4924*** (0.0723)
Supplier Reliability	0.4381 (0.2244)	0.2174** (0.0829)	0.6281** (0.2289)	0.4039*** (0.0912)
Case Pack Size	0.0126** (0.0042)	0.0053** (0.0016)	0.0143*** (0.0042)	0.0065*** (0.0017)
Top 20% Margin	0.4803** (0.1516)	0.1442* (0.0572)	0.4768** (0.1517)	0.1573** (0.0543)
Top 20% Dollar Volume	0.2170 (0.1163)	0.0799 (0.0459)	0.2835* (0.1171)	0.1513*** (0.0439)
Category Modification Fraction	1.1074*** (0.3336)	0.3720** (0.1352)	1.0655** (0.3375)	0.3383** (0.1255)
Store modification Fraction	4.5385*** (0.7230)	1.3221*** (0.2951)	4.6709*** (0.7219)	1.8145*** (0.2434)
OOS	0.0864 (0.5562)	-0.2282 (0.2305)	-0.0044 (0.5484)	-0.0839 (0.1809)
Selection Equation				
Demand Volatility	0.1271*** (0.0203)	0.1424*** (0.0212)	0.1293*** (0.0203)	0.1388*** (0.0204)
Agreed Leadtime	0.0102*** (0.0020)	0.0120*** (0.0021)	0.0107*** (0.0021)	0.0123*** (0.0021)
Supplier Lateness	0.0708 (0.0405)	0.1052* (0.0421)	0.0749 (0.0412)	0.0856 (0.0441)
Supplier Reliability	0.0981 (0.0544)	0.1080 (0.0561)	0.1015 (0.0549)	0.1192* (0.0569)
Case Pack Size	0.0044*** (0.0011)	0.0045*** (0.0012)	0.0043*** (0.0011)	0.0045*** (0.0011)
Top 20% Margin	0.1302*** (0.0290)	0.1223*** (0.0301)	0.1416*** (0.0291)	0.1692*** (0.0297)
Top 20% Dollar Volume	0.0642** (0.0244)	0.0477 (0.0244)	0.0745** (0.0244)	0.1044*** (0.0242)
Category Modification Fraction	0.3099*** (0.0666)	0.3101*** (0.0671)	0.3175*** (0.0668)	0.3396*** (0.0671)
Store modification Fraction	1.6446*** (0.0802)	1.7921*** (0.0820)	1.6491*** (0.0803)	1.6758*** (0.0794)
OOS	0.1048 (0.0896)	0.0430 (0.0841)	0.0899 (0.0883)	0.0247 (0.0814)
Number of decisions	-0.0003*** (0.0001)	-0.0006*** (0.0001)	-0.0003** (0.0001)	-0.0003** (0.0001)
Subcategory FE	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Observations	26993	26993	26993	26993
AIC	76039	59662	76367	57183
Log likelihood	-37960.7	-29772.1	-38124.7	-28532.6
Wald test of independent equations	0.0000	0.0000	0.0000	0.0000

Standard errors in parentheses: \* p<0.05, \*\*p<0.01, \*\*\*p<0.001

Table A2: Estimation of Heckman Sample Selection Model with Varying Inventory Holding and Substitution Rates with Normalized Cost Savings

	(1)	(2)	(3)	(4)
	Cost Savings Pct	Cost Savings Pct	Cost Savings Pct	Cost Savings Pct
Inventory Holding Rate	12%	12%	36%	36%
Substitution Rate	25%	75%	25%	75%
Outcome Equation				
Demand Volatility	0.1161*** (0.0284)	0.0376*** (0.0097)	0.1020*** (0.0274)	0.0263** (0.0089)
Agreed Leadtime	0.0047** (0.0017)	0.0018** (0.0006)	0.0058*** (0.0017)	0.0028*** (0.0007)
Supplier Lateness	0.0096 (0.0235)	0.0014 (0.0086)	0.0149 (0.0239)	0.0123 (0.0091)
Supplier Reliability	0.0231 (0.0335)	0.0044 (0.0123)	0.0209 (0.0342)	0.0087 (0.0133)
Case Pack Size	0.0013 (0.0007)	0.0003 (0.0003)	0.0013 (0.0007)	0.0005 (0.0003)
Top 20% Margin	0.0519* (0.0244)	0.0157 (0.0088)	0.0515* (0.0247)	0.0183* (0.0093)
Top 20% Dollar Volume	0.0025 (0.0119)	0.0039 (0.0042)	0.0044 (0.0121)	0.0027 (0.0045)
Category Modification Fraction	0.1644** (0.0500)	0.0518** (0.0187)	0.1526** (0.0509)	0.0440* (0.0195)
Store modification Fraction	0.8083*** (0.1551)	0.2570*** (0.0557)	0.8200*** (0.1553)	0.2927*** (0.0545)
OOS	0.0189 (0.0891)	-0.0225 (0.0322)	0.0046 (0.0894)	-0.0109 (0.0319)
Selection Equation				
Demand Volatility	0.2094*** (0.0214)	0.2190*** (0.0218)	0.2419*** (0.0216)	0.2940*** (0.0249)
Agreed Leadtime	0.0092*** (0.0021)	0.0107*** (0.0021)	0.0091*** (0.0021)	0.0104*** (0.0021)
Supplier Lateness	0.0770* (0.0358)	0.1045** (0.0382)	0.0885* (0.0362)	0.1125** (0.0379)
Supplier Reliability	0.1090* (0.0500)	0.1256* (0.0529)	0.1262* (0.0512)	0.1584** (0.0553)
Case Pack Size	0.0045*** (0.0012)	0.0052*** (0.0013)	0.0048*** (0.0012)	0.0056*** (0.0012)
Top 20% Margin	0.0924** (0.0286)	0.0938** (0.0292)	0.0928** (0.0291)	0.0865** (0.0302)
Top 20% Dollar Volume	-0.0291 (0.0187)	-0.0299 (0.0204)	-0.0307 (0.0189)	-0.0294 (0.0203)
Category Modification Fraction	0.3060*** (0.0634)	0.3137*** (0.0654)	0.3112*** (0.0635)	0.3218*** (0.0641)
Store modification Fraction	1.5895*** (0.0745)	1.7158*** (0.0784)	1.6022*** (0.0749)	1.6689*** (0.0787)
OOS	0.0846 (0.0981)	0.0583 (0.0930)	0.0683 (0.0976)	0.0104 (0.0916)
Number of decisions	-0.0003** (0.0001)	-0.0005*** (0.0001)	-0.0003*** (0.0001)	-0.0005*** (0.0001)
Subcategory FE	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Observations	26993	26993	26993	26993
AIC	37479	19408	38162	20274
Log likelihood	-18680.6	-9645.01	-19022.1	-10078.2
Wald test of independent equations	0.0000	0.0000	0.0000	0.0000

Standard errors in parentheses: \* p<0.05, \*\*p<0.01, \*\*\*p<0.001

## APPENDIX B

### RETURN OF THE BEHAVIORAL NEWSVENDOR: AN EXPERIMENTAL ANALYSIS OF CONSUMER RETURN POLICY DECISIONS

#### B.1 Derivation of Nominal Price under Normally Distributed Valuation

The nominal price under normally distributed product valuation can be derived as follows:

$$\begin{aligned}
 p^* &= E[\max(V, s)] = E[V|V > s]\Pr[V > s] + E[s|V \leq s]\Pr[V \leq s] = \\
 &\left( \mu_V + \sigma_V \frac{\phi\left(\frac{s-\mu_V}{\sigma_V}\right)}{1 - \Phi\left(\frac{s-\mu_V}{\sigma_V}\right)} \right) \left( 1 - \Phi\left(\frac{s-\mu_V}{\sigma_V}\right) \right) + s\Phi\left(\frac{s-\mu_V}{\sigma_V}\right) = \\
 &\mu_V \left( 1 - \Phi\left(\frac{s-\mu_V}{\sigma_V}\right) \right) + \sigma_V \phi\left(\frac{s-\mu_V}{\sigma_V}\right) + s\Phi\left(\frac{s-\mu_V}{\sigma_V}\right) = \\
 &\mu_V + (s - \mu_V)\Phi\left(\frac{s-\mu_V}{\sigma_V}\right) + \sigma_V \phi\left(\frac{s-\mu_V}{\sigma_V}\right)
 \end{aligned} \tag{B.1}$$

where  $\phi$ ,  $\Phi$ , and  $\Phi^{-1}$  are the pdf, CDF, and inverse of a standard normal random variable, respectively. Note that  $E[V|V > s]$  is the mean of a left-truncated normal random variable (Kotz et al. 2004). The expressions for  $Q^*$  and  $r^*$  follow straightforwardly from equations (2).

#### B.2 Proof that Price is Monotonically Increasing in the Salvage Value

Here, we prove that the nominal price is a monotonically increasing function of the salvage value. Taking the first derivative of the nominal price expression with respect to the salvage value,

$$\begin{aligned}
 \frac{dp^*}{ds} &= \Phi\left(\frac{s-\mu_V}{\sigma_V}\right) + (s - \mu_V)\Phi'\left(\frac{s-\mu_V}{\sigma_V}\right) + \sigma_V \phi'\left(\frac{s-\mu_V}{\sigma_V}\right) = \\
 &\Phi\left(\frac{s-\mu_V}{\sigma_V}\right) + (s - \mu_V)\frac{1}{\sigma_V}\phi\left(\frac{s-\mu_V}{\sigma_V}\right) + \sigma_V \frac{1}{\sigma_V} \left(-\frac{s-\mu_V}{\sigma_V}\right)\phi\left(\frac{s-\mu_V}{\sigma_V}\right) \\
 &= \Phi\left(\frac{s-\mu_V}{\sigma_V}\right).
 \end{aligned} \tag{B.2}$$

where we use the facts that the first order derivative of a CDF gives the pdf and for a standard normal random variable  $\phi'(x) = -x\phi(x)$ . Since  $\Phi$  is a positive function, we conclude that the first order derivative of the nominal price with respect to the salvage value is positive. Hence, the nominal price ( $p^*$ ) increases in the salvage value ( $s$ ).

## B.3 Experimental Task Interface

### Study Description (1 of 3)

**Please read carefully.**

General Overview:

Your earnings will be determined by the decisions you make in this task. You will make decisions for 35 rounds. Once you complete the task, we will randomly pick one of your rounds earning in Francs (Fr), convert it to U.S. dollars with 3,000 Fr = \$1 rate, and credit the converted amount to your account. The first 5 rounds will be for warm-up and will not be considered in the random draw to determine your payment. Earning negative profit in the task is possible, but you will not be charged anything if a round in which you earned a negative profit is picked in the random draw. In other words, the minimum amount you can earn in this task based on your decisions is \$0. In addition to the payment that is based on your decisions in the task, you will also receive \$3 payment for your participation and completion of the task.

**Note that you must complete the task to receive any payment.**

In this task, you will manage a retailer that sells widgets. Your goal is to maximize your profits. In each round, you need to make the following 3 decisions:

**1)Order Quantity:** how many widgets you want to order from the supplier,

**2)Price:** how much you want to charge your customers for each widget,

**3)Refund Amount:** how much you want to pay back to your customers as a refund if they return the widget after purchasing.

Each widget you order will incur a purchase cost, and the number of widgets you sell each round will depend on the market size and your decisions. At the end of each round, you will be reimbursed for any unsold widgets and widgets that are returned by customers at a predetermined salvage value per widget. Widgets cannot be carried over to the following rounds.

Next

## Study Description (2 of 3)

**Please read carefully.**

### Market Size:

In each round, the computer will randomly select a number from a Normal distribution with mean 200 and standard deviation 50. This number will be the market size. Market size represents the number of consumers in the market for widgets. The market size in any given round is **independent** of the market size in previous or later rounds. **Your decisions in a round will not affect the market size in the following rounds.**

### Consumer Valuation of Widgets and Consumer Demand:

Consumers can realize their exact valuation of the widget only after purchasing and experiencing the widget. However, before purchasing, each consumer knows that his or her valuation of the widget will come from a Normal distribution with mean of 50 Fr and standard deviation of 20 Fr. That is, before purchasing every consumer in the market has the same expectation of their individual valuation.

Note that all consumers in the market are rational. That is, if you set the price of widget very high relative to the value that consumers may receive from the widget, all consumers in the market will choose not to buy the widget. Therefore, you will not be able to sell any of the widgets that you ordered. In making purchase decisions, consumers also consider the refund amount that you offer for returned widgets. Each consumer who decides to buy will buy one unit of widget. Thus, the number of consumers in the market that decide to buy a widget given the price and refund amount that you choose will determine the consumer demand for widgets. **Depending on your price and refund decisions, the consumer demand for widgets can either be zero or equal to the market size.**

### Consumer Returns:

After purchasing the widget, each consumer will realize their individual valuation of the widget. If the value of the widget to the consumer turns out to be smaller than the refund amount you choose, the consumer will return the widget and receive the refund amount. Otherwise, the consumer will keep the widget.

### Profit:

If the consumer demand turns out to be greater than or equal to the amount you ordered, then you can at most sell the number of widgets that you ordered, at the price of your choice. If the consumer demand turns to be smaller than the amount you ordered, then you will sell the number of widgets demanded and salvage the unsold widgets. The widgets that are returned to you each round will also be salvaged.

For each round, your profit will be determined by the following equation:

$$\text{Your profit} = \text{Widgets sold} \times \text{Your price} - \text{Widgets returned} \times \text{Your refund amount} \\ + (\text{Widgets returned} + \text{Widgets unsold}) \times \text{Salvage Value} - \text{Purchase Cost} \times \text{Your order quantity}$$


Next

## Study Description (3 of 3)


Please read carefully.

Normally Distributed Market Size and Consumer Valuation of Widget:

Below is the normal distribution with mean 200 and standard deviation of 50 from which the market size will be drawn in each round.



Below is the normal distribution with mean 50 Fr and standard deviation of 20 Fr from which the individual consumer valuations of the widget will be drawn in each round.



Next

## Quiz to Proceed

1. Based on the instructions you just read, please indicate whether the following statement is TRUE or FALSE:  
Your decisions in each round will not affect the market size in the following rounds:

True  False

2. Suppose the market size in a round turns out to be 250. Based on your decisions, which of the following values can be the consumer demand for widgets in that round? Please select all that apply.

- 0  
 150  
 250  
 300

3. Suppose that in a round you ordered 100 widgets and consumer demand turns out to be 150 widgets. How many widgets will you sell in this round?

50  100  150  300

4. Suppose that you set the price of the widget as 34 Fr, and the purchase cost of each widget was 36 Fr. You ordered and sold 10 widgets, and none were returned. How much profit did you make?

40 Fr  20 Fr  0 Fr  -20 Fr  -40 Fr

Next

## [Round 1 of 30] Make Your Decision

- |   |  |
|---|--|
| <ul style="list-style-type: none"> <li>• Purchase cost: <b>36 Fr</b></li> <li>• Salvage value for unsold and returned widgets: <b>0 Fr</b></li> </ul> | <ul style="list-style-type: none"> <li>• Market size is distributed normally with <b>mean 200</b> and <b>standard deviation 50</b></li> <li>• Customer's valuation of a widget is normally distributed with <b>mean 50 Fr</b> and <b>standard deviation 20 Fr</b></li> </ul> |
|---|--|

Please enter your decisions and press "Next" to continue:

Refund Amount: <input type="text"/>	Price: <input type="text"/>	Order Quantity: <input type="text"/>	Next
--	--------------------------------	---	------

Your decision history:

Round #	Refund Amount	Price	Order Quantity	Realized Market Size	Widgets Sold	Widgets Returned	Widgets Unsold	Your Profit
1 (Current)								



1. Please indicate your gender:  
 Male  Female

2. Please indicate your age:

3. Have you ever worked in the retail industry?  
 Never  1-6 months  7-12 months  1-2 years  2-3 years  3-5 years  5-10 years  10+ years

3.1 If you have worked in the retail industry, what was your job title? If possible, please provide us with a short job description, too.

4. In general, how many years of full-time working experience do you have?  
 Never  1-6 months  7-12 months  1-2 years  2-3 years  3-5 years  5-10 years  10+ years

5. Please indicate your total annual household income:  
 \$0-\$35,000  \$35,000-\$69,999  \$70,000-\$99,999  \$100,000-\$149,999  \$150,000-\$199,999  \$200,000+

6. What is the highest degree of level of education you have completed?  
 No formal education  Some high school  High school  Bachelor's degree  Master's degree  Doctoral degree

## B.4 Pre-defined Exclusion Criteria

Our pre-defined exclusion criteria follow a three-step process. First, we identify anomalous decisions characterized as pricing, quantity, or refund decisions that are either abnormally high or low, which may have been driven by the extreme variability of open-ended decision levers (i.e., ability to set unlimited price, order quantity, and refund amounts, similar to Moon and Nelson 2020). We define abnormal pricing and order quantity decisions as decisions with values that are in the top or bottom 1 percentile or over 3 standard deviations away from the mean, where the mean and standard deviation are calculated using observations lying between the 1st to 99th percentiles (following winsorization approach to identify extreme values (Moon and Nelson 2020) and excluding observations of extreme values (Moritz et al. 2014)). Specific to pricing decisions, we also consider the price to be abnormally high if it exceeds the market participation constraint (i.e.,  $p \geq E \max(V, r)$ , which leads to zero sales) and abnormally low if the price is lower than the procurement cost ( $p < 36$ ). For refund amount, we flag overly generous refunds ( $r > p$ ) as anomalies. Second, we evaluate whether these anomalous decisions are driven by exploration

behaviors and drop only those participants who make three or more consecutive anomalous decisions (suggesting that their behavior is not driven by exploration behaviors). Using this criteria we dropped 24 subjects. Finally, we manually analyze the decisions from subjects with three or more anomalies to identify and drop those that make abnormal decisions often and with extreme values but never three consecutive times. Also, we check for clear typing error and for this case we do not drop the subject but just the decision. Using this criteria, we additionally dropped 9 subjects and one decision with clear typing error.

## B.5 Correlation Table

Table B1: Correlations between the Demographic and Socio-economic Measures

Variables	Age	Gender	Income	Education	Work Exp.	Retail Exp.	Risk Taking	Altruism	Trust	Pos. Rec.	Neg. Rec.	Patience
Age	1.000											
Gender	0.021 (0.173)	1.000										
Income	0.054 (0.001)	0.058 (0.000)	1.000									
Education	0.345 (0.000)	0.060 (0.000)	0.341 (0.000)	1.000								
Work Exp.	0.612 (0.000)	0.121 (0.000)	-0.011 (0.470)	0.361 (0.000)	1.000							
Retail Exp.	0.350 (0.000)	-0.116 (0.000)	-0.001 (0.962)	-0.015 (0.342)	0.327 (0.000)	1.000						
Risk Taking	0.077 (0.000)	0.080 (0.000)	0.169 (0.000)	0.211 (0.000)	0.119 (0.000)	-0.016 (0.318)	1.000					
Altruism	-0.057 (0.000)	-0.050 (0.002)	0.066 (0.000)	0.109 (0.000)	0.113 (0.000)	0.046 (0.003)	0.260 (0.000)	1.000				
Trust	0.109 (0.000)	0.068 (0.000)	0.101 (0.000)	0.219 (0.000)	0.142 (0.000)	-0.056 (0.000)	0.269 (0.000)	0.397 (0.000)	1.000			
Pos. Rec.	-0.013 (0.420)	-0.007 (0.655)	0.006 (0.726)	0.017 (0.278)	0.018 (0.264)	-0.037 (0.020)	0.069 (0.000)	0.356 (0.000)	0.314 (0.000)	1.000		
Neg. Rec.	0.029 (0.066)	0.159 (0.000)	0.063 (0.000)	-0.002 (0.896)	-0.084 (0.000)	0.033 (0.038)	0.100 (0.000)	-0.204 (0.000)	-0.256 (0.000)	-0.058 (0.000)	1.000	
Patience	-0.099 (0.000)	0.022 (0.161)	0.101 (0.000)	-0.047 (0.003)	-0.074 (0.000)	-0.067 (0.000)	-0.174 (0.000)	-0.086 (0.000)	0.074 (0.000)	0.172 (0.000)	-0.148 (0.000)	1.000

**Notes:** Gender is coded as Male = 1, Female = 0. Education is highest education level completed. *p*-values in parentheses.

## **B.6 Examining the Role of Individual Heterogeneity**

Existing literature suggests that different sources of individual-level heterogeneity, such as risk attitude, gender, work experience, may be related to the behavioral newsvendor decisions (Becker-Peth and Thonemann 2019, pp. 419–421). To highlight notable examples, De Vericourt et al. (2013) showed, using a sample from Amazon Mechanical Turk, that female decision-makers ordered less than male decision-makers in the high-margin condition but no differently in the low-margin condition and that the gender differences were mediated by the differences in risk attitudes. However, the study was limited to the demographic characteristics of gender and age and the socio-economic preferences of risk attitude, and the authors called for future research in this direction. Meanwhile, Becker-Peth et al. (2018) found a gender effect on order quantities, aligned with De Vericourt et al. (2013), whereas they also find that risk attitudes do not mediate this difference. In lab environments, Bolton et al. (2012) and Moritz et al. (2013) studied the performance differences among student decision-makers and managers in a classical newsvendor task and conclude that both groups showed comparable levels of pull-to-center bias.

We study two categories of individual-level heterogeneity: demographic characteristics and socio-economic preferences. Unlike the existing studies, we examine the associations among a more extensive set of individual-level factors under both categories and order quantity, price, and refund decisions.

### **B.6.1 Demographic Characteristics**

In the post-experimental survey, we collected data on participant’s age, gender, annual income, highest level of education completed, work experience, retail industry experience, and job role in retail, if any (coded in terms of whether the role was managerial or not). Table B2 reports frequency statistics of each variable.

To examine the role of the demographic characteristics in explaining the decisions of participants, we estimated a random-slope, random-intercept LMM as in equation 10 with the addition of fixed effect variables for the demographic characteristics. The result of the analysis is provided

Table B2: Frequency Table of Demographic Characteristics

Characteristic	Frequency	Percent	Cumulative
<i>Age</i>			
18–25	33	32.04	32.04
26–35	44	42.72	74.76
36–45	20	19.42	94.17
46–55	3	2.91	97.09
56+	3	2.91	100.00
<i>Gender</i>			
Male	47	45.63	45.63
Female	56	54.37	100.00
<i>Income</i>			
< \$35,000	21	20.39	20.39
\$35,000–\$69,999	33	32.04	52.43
\$70,000–\$99,999	23	22.33	74.76
\$100,000–\$149,999	15	14.56	89.32
\$150,000–\$199,999	5	4.85	94.17
≥ \$200,000	6	5.83	100.00
<i>Highest Education Level Completed</i>			
Some high school	1	0.97	0.97
High school diploma	33	32.04	33.01
Bachelor’s degree	40	38.83	71.84
Master’s degree	25	24.27	96.12
Doctoral degree	4	3.88	100.00
<i>Work Experience</i>			
None	16	15.53	15.53
1–6 months	5	4.85	20.39
7–12 months	2	1.94	22.33
1–2 years	7	6.8	29.13
2–3 years	9	8.74	37.86
3–5 years	12	11.65	49.51
5–10 years	20	19.42	68.93
10+ years	32	31.07	100.00
<i>Retail Experience</i>			
None	64	62.14	62.14
1–6 months	4	3.88	66.02
7–12 months	8	7.77	73.79
1–2 years	8	7.77	81.55
2–3 years	5	4.85	86.41
3–5 years	5	4.85	91.26
5–10 years	6	5.83	97.09
10+ years	3	2.91	100.00
<i>Retail Manager Experience</i>			
Yes	17	16.50	16.50
No	86	83.50	100.00
Total	103	100	100

in Table B3. We found that a decision-maker’s education level was negatively associated with the quantity ordered, all else equal. Other demographic characteristics such as age, gender, income, work experience, retail experience, and manager experience failed to significantly explain order quantity, price, and refund decisions beyond what was explained by the salvage value treatment,

round, and period effects.

Table B3: Analysis of Individual Heterogeneity: Demographic Characteristics

	(1) <i>Q</i>	(2) <i>p</i>	(3) <i>r</i>
Fixed Effects			
Intercept	233.245*** (20.023)	48.412*** (1.754)	26.873*** (5.346)
Treatment (s = 30)	16.822*** (2.078)	0.494* (0.226)	4.173*** (0.758)
Round1	-1.012** (0.355)	0.123*** (0.035)	-0.180*** (0.045)
Round2	-0.923*** (0.274)	-0.021 (0.026)	-0.259*** (0.072)
Period	4.187 (4.781)	1.706*** (0.493)	-4.714*** (0.989)
Age	0.056 (0.496)	-0.009 (0.043)	-0.114 (0.134)
Gender	0.344 (7.564)	1.050 (0.655)	-1.214 (2.039)
Income	-3.855 (2.799)	-0.148 (0.242)	0.089 (0.754)
Education Level	-12.341* (4.849)	-0.349 (0.420)	-1.490 (1.307)
Work Experience	1.727 (1.994)	-0.135 (0.173)	0.342 (0.537)
Retail Experience	0.573 (2.824)	0.086 (0.245)	0.369 (0.761)
Retail Manager Experience	-15.861 (15.731)	-0.017 (1.362)	2.929 (4.239)
Random Effects			
Level 2 (between-subject)			
Intercept	2308.355*** (393.308)	24.003*** (3.857)	125.764*** (19.183)
Round1	8.650*** (1.811)	0.082*** (0.017)	0.147*** (0.029)
Round2	3.415*** (1.084)	0.025*** (0.009)	0.469*** (0.074)
Period	1640.652*** (328.944)	18.070*** (3.494)	90.278*** (14.055)
Level 1 (within-subject)			
Residual	1208.272*** (33.026)	11.788*** (0.322)	17.565*** (0.480)
<i>N</i>	3089	3089	3089

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## **B.6.2 Socio-Economic Preferences**

In the post-experimental survey, we also measured a number of key socio-economic preferences of participants, namely, risk taking, altruism, trust, positive reciprocity, negative reciprocity, and patience (i.e., time discounting). To do so, we employed the socio-economic preferences survey module by Falk et al. (2018). The measures consisted of a series of direct survey questions and hypothetical choice tasks (see Appendix D in Falk et al. (2016) for the complete list). The data collected with these items were weighted according to weights estimated in Falk et al. (2016) as the result of cross-validation with financially incentivized experimental measures. We report correlations between the socio-economic preference and demographic variables in Table B1, Appendix B.5. It is worthwhile to note that the socio-economic preferences measures show significant correlations both in-between and with demographic characteristics in directions that are aligned with findings from numerous field- and lab-based behavioral economics studies. For example, the correlation table suggests that males are more risk taking compared to females, less altruistic (Croson and Gneezy 2009), more trusting and show a greater negative reciprocity (Garbarino and Slonim 2009), that older individuals are more trusting (Matsumoto et al. 2016), and that higher income is associated with greater risk taking (Shaw 1996).

Again, we estimated a random-slope, random-intercept LMM as in equation 10, adding socio-economic preferences. The result of the analysis is provided in Table B4. We found that none of the socio-economic preferences significantly explained the variability in order quantity, price, and refund decisions beyond what was already explained by salvage value treatment, round, and period effects.

## **B.6.3 Performance Implications of Individual Heterogeneity**

After examining the role of demographic characteristics and socio-economic preferences in explaining the actual order quantity, price, and refund decisions, we also analyzed how these factors relate to the decision-making performance measured in terms of absolute deviations from the expected profit maximizing decisions. We estimated six LMMs with absolute deviation from the

Table B4: Analysis of Individual Heterogeneity: Socio-Economic Preferences

	(1)	(2)	(3)
	$Q$	$p$	$r$
Fixed Effects			
Intercept	197.686*** (27.606)	44.195*** (2.301)	19.903** (7.243)
Treatment (s = 30)	16.763*** (2.078)	0.509* (0.227)	4.382*** (0.757)
Round1	-1.012** (0.355)	0.123*** (0.035)	-0.180*** (0.045)
Round2	-0.923*** (0.274)	-0.021 (0.026)	-0.259*** (0.072)
Period	4.186 (4.781)	1.706*** (0.493)	-4.708*** (0.990)
Risk Taking	4.038 (6.351)	-0.871 (0.525)	-1.151 (1.675)
Altruism	-1.331 (2.001)	0.024 (0.165)	0.110 (0.528)
Trust	-0.605 (1.498)	-0.114 (0.124)	-0.204 (0.395)
Pos. Reciprocity	-1.258 (3.185)	0.208 (0.263)	-0.021 (0.840)
Neg. Reciprocity	-0.328 (4.902)	0.025 (0.405)	0.731 (1.293)
Patience	0.667 (0.362)	0.042 (0.030)	-0.014 (0.096)
Random Effects			
Level 2 (between-subject)			
Intercept	2388.611*** (392.011)	23.771*** (3.843)	125.381*** (18.755)
Round1	8.650*** (1.811)	0.082*** (0.017)	0.147*** (0.029)
Round2	3.415*** (1.084)	0.025*** (0.009)	0.469*** (0.074)
Period	1640.348*** (328.885)	18.071*** (3.494)	90.544*** (14.119)
Level 1 (within-subject)			
Residual	1208.272*** (33.026)	11.788*** (0.322)	17.565*** (0.480)
$N$	3089	3089	3089

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

normative decision for a given decision, controlling for the other two decisions made, salvage value, round, period, demographic characteristics in three of the models, and socio-economic preferences in the remaining three models. We make the following observations. Decision-makers that are more patient deviated less from the optimal order quantity, everything else equal. Meanwhile, decision-makers with a higher education level deviated more from the optimal order quantity, ev-



everything else equal. The extent to which participants deviated from the optimal price and optimal refund amount were insignificantly predicted by both the demographic and socio-economic characteristics. We relegate the detailed results of these analyses to Table B5 and B6.

An interesting observation we make is that risk taking attitude did not explain individuals' decisions. Existing research has largely focused on exploring the role of risk in explaining order quantity decisions in a classical behavioral newsvendor framework (e.g., De Vericourt et al. 2013, Becker-Peth et al. 2018). As discussed earlier, this line of research had inconclusive findings regarding the explanatory role of risk taking attitude. By including other important socio-economic preferences into our analysis, we provide a broader picture of how these preferences relate to decision and performance variability in a multi-lever decision environment. For example, we found that when a number of key socio-economic preferences are controlled for, risk taking attitude does not significantly explain variability in order quantity decisions. Moreover, decision-makers with a longer overall work experience, retail experience, retail manager experience, higher income, and higher level of educational attainment did not perform significantly better. This is aligned with studies that have showed similar performances in a classical behavioral newsvendor task as well as other incentivized behavioral tasks by professionals and MBA students relative to undergraduate students (Moritz et al. 2013, Bolton et al. 2012, Fr chet 2015).

Overall, using a more diverse sample and in a multi-lever decision-making task, we come to a conclusion that most ex-ante characteristics that would render a decision-maker better suited for the task are not necessarily associated with significant decision-making performance advantages. That is, under experimental environments with perfect information (i.e., the parameters of statistical distributions are known), the characteristics of decision-makers do not seem to matter much.

Table B5: Analysis of Demographic Characteristics: Deviations from Normative Decisions

	(1) $ Q - Q^* $	(2) $ p - p^* $	(3) $ r - r^* $
<b>Fixed Effects</b>			
Intercept	-16.234 (15.721)	4.436** (1.624)	-3.505 (2.986)
Treatment (s = 30)	13.648*** (2.775)	1.358*** (0.177)	-4.133** (1.568)
Order Quantity		-0.010*** (0.001)	0.007*** (0.002)
Price	0.241 (0.131)		0.255*** (0.020)
Refund Amount	-0.134 (0.092)	0.039*** (0.010)	
Round1	-0.085 (0.212)	-0.147*** (0.027)	-0.078 (0.047)
Round2	0.218 (0.197)	-0.025 (0.023)	-0.090 (0.068)
Period	-9.937* (3.883)	-1.588*** (0.346)	6.012*** (1.632)
Age	0.068 (0.357)	0.009 (0.039)	0.069 (0.064)
Gender	4.863 (5.449)	-0.844 (0.590)	-0.455 (0.970)
Income	2.864 (2.015)	0.050 (0.218)	-0.309 (0.359)
Education Level	9.826** (3.491)	0.235 (0.378)	1.165 (0.622)
Work Experience	-0.525 (1.436)	0.138 (0.156)	-0.452 (0.256)
Retail Experience	0.728 (2.033)	-0.118 (0.220)	0.040 (0.362)
Retail Manager Experience	0.367 (11.322)	-0.350 (1.226)	1.906 (2.018)
<b>Random Effects</b>			
Level 2 (between-subject)			
Intercept	1247.983*** (205.671)	16.735*** (2.707)	89.764*** (13.657)
Round1	2.499*** (0.637)	0.046*** (0.011)	0.174*** (0.031)
Round2	1.893* (0.554)	0.023*** (0.007)	0.428*** (0.066)
Period	1183.586*** (214.892)	7.245*** (1.693)	265.893*** (38.395)
Level 1 (within-subject)			
Residual	575.453*** (15.735)	8.240*** (0.225)	13.692*** (0.374)
N	3089	3089	3089

Standard errors in parentheses  
 \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table B6: Analysis of Socio-Economic Characteristics: Deviations from Normative Decisions

	(1) $ Q - Q^* $	(2) $ p - p^* $	(3) $ r - r^* $
<b>Fixed Effects</b>			
Intercept	31.756 (20.120)	7.093*** (2.097)	4.128 (3.770)
Treatment (s = 30)	13.548*** (2.745)	1.339*** (0.179)	-4.267** (1.565)
Order Quantity		-0.010*** (0.001)	0.006** (0.002)
Price	0.255 (0.131)		0.255*** (0.020)
Refund Amount	-0.145 (0.091)	0.039*** (0.010)	
Round1	-0.089 (0.212)	-0.147*** (0.027)	-0.078 (0.047)
Round2	0.215 (0.197)	-0.025 (0.023)	-0.090 (0.068)
Period	-10.017** (3.883)	-1.592*** (0.346)	6.008*** (1.631)
Risk Taking	6.250 (4.409)	0.774 (0.474)	0.639 (0.804)
Altruism	1.581 (1.387)	-0.007 (0.149)	-0.250 (0.253)
Trust	1.004 (1.039)	0.038 (0.112)	0.042 (0.190)
Pos. Reciprocity	-0.796 (2.209)	-0.095 (0.237)	-0.142 (0.403)
Neg. Reciprocity	2.477 (3.398)	0.067 (0.365)	-0.175 (0.621)
Patience	-0.584* (0.251)	-0.035 (0.027)	-0.038 (0.046)
<b>Random Effects</b>			
Level 2 (between-subject)			
Intercept	1321.531*** (214.025)	16.082*** (2.626)	87.985*** (13.399)
Round1	2.497*** (0.637)	0.046*** (0.011)	0.174*** (0.031)
Round2	1.898* (0.555)	0.023*** (0.007)	0.428*** (0.066)
Period	1183.515*** (214.744)	7.244*** (1.693)	265.606*** (38.308)
Level 1 (within-subject)			
Residual	575.394*** (15.732)	8.240*** (0.225)	13.692*** (0.374)
<i>N</i>	3089	3089	3089

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$