# ESSAYS ON INDIVIDUAL CHOICE MODELING AND ANALYSIS

by

Yulia Vorotyntseva

APPROVED BY SUPERVISORY COMMITTEE:

___________________________________

Dorothée Honhon, Co-Chair

___________________________________

Elena Katok, Co-Chair

___________________________________

Ernan Haruvy

___________________________________

Kyle Hyndman

# ESSAYS ON INDIVIDUAL CHOICE MODELING AND ANALYSIS

by

YULIA VOROTYNTSEVA, BS, MS

DISSERTATION

Presented to the Faculty of

The University of Texas at Dallas

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY IN

MANAGEMENT SCIENCE

THE UNIVERSITY OF TEXAS AT DALLAS

August 2018

ACKNOWLEDGMENTS

My time at UT Dallas was a time of growth, both professional and personal, and I am deeply grateful to my advisors who provided the resources for my development and helped me to handle the pain of it. I want to thank Professor Dorothée Honhon for her support, advice, championship and especially the much needed inspiration she was giving me for the last three years of my Ph.D. program. I would like to express my deepest gratitude to Professor Elena Katok for her effective management of my dissertation progress and sharing her professional insights with me. I am incredibly fortunate to receive guidance and training in behavioral operations from one of those whose work has shaped this field.

I am greatly indebted to Professor Özalp Özer for his guidance, advice and feedback that helped me in writing the first essay of this dissertation. I want to thank Professor Canan Ulu for her advice, reviews of my numerous drafts and long hours of discussions that shaped the third essay in this dissertation. My work has also greatly benefited from comments and suggestions provided by Professors Kay-Yut Chen, Ernan Haruvy, Kyle Hyndman and Gary Bolton.

Many thanks to Professors Kyle Hyndman and Ernan Haruvy for serving on my dissertation committee. I am also thankful to Amanda and Joan Besch, Jennifer De Leon, Jamie Selander, Ashley Desouza for providing both administrative assistance and a sympathetic ear.

I thank all my fellow Ph.D. students: Xiao Zhang, Duygu Dagli, Bharadwaj Kadyala, Blair Flicker, Emre Ertan, Bahriye Cesaret, Yang Bo, Sina Shokoohyar and Harish Guda and many others, – for brightening up my Ph.D. experience with their friendship. Last but not least, I would like to thank my husband for his unconditional support that made the completion of this work possible.

June 2018

ESSAYS ON INDIVIDUAL CHOICE MODELING AND ANALYSIS

Yulia Vorotyntseva, PhD
The University of Texas at Dallas, 2018

Supervising Professors: Dorothée Honhon, Co-Chair
Elena Katok, Co-Chair

This dissertation includes three essays. The first two essays are concerned with empirical investigation of individual choice by means of laboratory experiment, and the third one is concerned with integrating one particular behavior pattern – consumer learning – into assortment decisions. The summaries of these essays are provided below.

In the first essay, we study how customers are choosing prices in a Pay-What-You-Want (PWYW) business model. Under PWYW the price for a product is fully determined by a buyer: the seller cannot reject any offer. We study two factors that can potentially affect PWYW prices: the seller's production costs and the buyer's private valuation of the product. We hypothesize that buyers may anticipate the seller's loss-aversion, so they will be reluctant to choose prices below the costs, and that the prices increase as the buyer's valuation increases. We suggest a model that incorporates this hypothetical behavior and estimate it using the dataset from a previously published experimental study. Based on the preliminary insights we obtain, we design and conduct our own controlled laboratory experiment. Our findings suggest that PWYW prices are indeed increasing on average in the buyer's valuation and seller's costs.

In the second essay we provide the results of empirical investigation of the behavior of human assortment planners. Assortment planning, that is, the selection of products to offer in a

store or the design of a product line for a manufacturer, is often performed by managers without any support from computerized optimization algorithms. The goals of our study are (1) to find if human decision makers deviate from the expected profit-maximizing solution in some systematic way and (2) to investigate the effect of decision support tools on the efficacy of assortment planners. To do this, we develop and conduct a behavioral experiment, where the subjects are repeatedly picking assortments in a simplified computerized market environment. We find that the subjects perform better when the profit maximizing assortment consists of fewer products and that subjects improve their decisions over time. The effect of decision support is somewhat surprising: under certain conditions providing subjects with more information resulted in worse performance.

In the third essay, we develop a model that explicitly incorporates consumer learning into a firm's assortment problem. Consumer's choice of a product from a particular category is influenced by her beliefs about how well the product fits her needs. When a consumer purchases a product repeatedly, her experience with it affects her beliefs and, consequently, her future choices. By providing the consumer with an assortment to choose from, the firm increases the chances that the consumer will find a product that is suitable for her and keep purchasing from the firm in the long term. However, by doing so the firm faces a risk that a less profitable product gets substituted for a more profitable one. The model we develop allows to investigate this tradeoff analytically.

TABLE OF CONTENTS

ACKNOWLEDGMENTS . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . iv

ABSTRACT . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . v

LIST OF FIGURES . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . ix

LIST OF TABLES . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . x

CHAPTER 1 INTRODUCTION . . . . . . . . . . . . . . . . . . . . . . . . 1

CHAPTER 2 A BEHAVIORAL STUDY OF PRICE SELECTION IN PAY-WHAT-YOU-WANT BUSINESS MODEL . . . . . . . . . . . . . . . . . . . . . . 6

2.1 Introduction . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 6

2.2 Literature Review . . . . . . . . . . . . . . . . . . . . . . . . . . . 11

2.3 Model Development . . . . . . . . . . . . . . . . . . . . . . . . . . 16

2.3.1 Problem Specification . . . . . . . . . . . . . . . . . . . . . . 16

2.3.2 The Fairness Model . . . . . . . . . . . . . . . . . . . . . . . 18

2.4 Study 1. Data from Schmidt et al. (2015) . . . . . . . . . . . . . . . 20

2.5 Study 2. Pay-What-You-Want Experiment . . . . . . . . . . . . . . . 25

2.5.1 Experimental Design and Procedure . . . . . . . . . . . . . . . 25

2.5.2 Results . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 27

2.6 Summary and Further Work . . . . . . . . . . . . . . . . . . . . . . 30

CHAPTER 3 A BEHAVIORAL STUDY OF RETAIL ASSORTMENT PLANNING 31

3.1 Introduction . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 31

3.2 The Assortment Planning Problem . . . . . . . . . . . . . . . . . . 32

3.2.1 Market Demand Model . . . . . . . . . . . . . . . . . . . . . . 33

3.2.2 Retailer's Profit Function . . . . . . . . . . . . . . . . . . . . 34

3.3 Research Question and Hypotheses . . . . . . . . . . . . . . . . . . 35

3.4 Experimental Design and Procedure . . . . . . . . . . . . . . . . . . 36

3.4.1 Design of Experiment . . . . . . . . . . . . . . . . . . . . . . 36

3.4.2 Experiment Implementation . . . . . . . . . . . . . . . . . . . 40

3.5 Results . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 41

3.5.1 Overall Performance Analysis . . . . . . . . . . . . . . . . . . 41

3.5.2 “Pull-to-Center” Effect . . . 44
3.5.3 Effect of Decision Support . . . 49
3.5.4 Analysis of the Clickstream Data . . . 50
3.6 Summary and Conclusions . . . 60
CHAPTER 4 A MODEL FOR PERSONALIZED ASSORTMENT PLANNING IN PRESENCE OF CONSUMER LEARNING . . . 62
4.1 Introduction . . . 62
4.2 Consumer Learning Model . . . 63
4.2.1 Problem Formulation . . . 63
4.2.2 Consumer Beliefs Evolution . . . 65
4.3 One-Product Case . . . 68
4.3.1 Consumer’s Abandonment Time . . . 68
4.3.2 Firm’s Profit . . . 75
4.4 Two-Product Case . . . 77
4.5 Summary and Conclusions . . . 86
APPENDIX A APPENDIX FOR CHAPTER 2 . . . 88
APPENDIX B APPENDIX FOR CHAPTER 3 . . . 89
B.1 Proof of Greedy Algorithm Optimality for Assortment Planning Model . . . 89
B.2 Subject Instructions . . . 90
B.3 Consent Form . . . 92
B.4 Additional Regression Models Output . . . 93
B.5 Additional Graphs . . . 95
APPENDIX C APPENDIX FOR CHAPTER 4 . . . 98
C.1 Proof or Theorem 1 . . . 98
REFERENCES . . . 100
BIOGRAPHICAL SKETCH . . . 106
CURRICULUM VITAE

LIST OF FIGURES


2.1 Sequence of events in a PWYW interaction . . . . . . . . . . . . . . . . . . . . 17

2.2 Buyer-determined price frequencies in the study by (Schmidt et al., 2015), grouped by the production cost value. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 22

2.3 Pooled results for rounds with positive seller's cost and positive surplus . . . . . 29

3.1 Decision stage screens for different levels of decision support information *Note: parameters shown are for low optimal variety (LOV) condition* . . . . . . . . . 36

3.2 Results stage screen . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 37

3.3 Expected profits for popular assortments . . . . . . . . . . . . . . . . . . . . . 38

3.4 Performance metrics summary . . . . . . . . . . . . . . . . . . . . . . . . . . . 43

3.5 Normalized expected profits . . . . . . . . . . . . . . . . . . . . . . . . . . . . 44

3.6 Assortment types over time . . . . . . . . . . . . . . . . . . . . . . . . . . . . 45

3.7 Average number of solutions of each type with 95% confidence intervals . . . . . 45

3.8 A box plot of number of solutions of each type per subject . . . . . . . . . . . . 47

3.9 Average number of different assortments seen per period . . . . . . . . . . . . . 51

3.10 Percentage of subjects who picked the best assortment among the ones seen . . 53

3.11 Decision time (in seconds – logarithm scale) . . . . . . . . . . . . . . . . . . . 55

3.12 Exploration in the first period and normalized expected profit : assortments seen 58

3.13 Exploration in the first period and normalized expected profit in the experiment: decision time . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 58

4.1 Examples of problem setup parameters that satisfy the assumption $q_B = 1 - q_G$ 67

4.2 Evolution of consumer's beliefs as a Markov Chain $\{Y_t\}$, conditional on the product state $S = s$ . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 72

4.3 Abandonment time distribution illustration for Example 2 . . . . . . . . . . . . 75

4.4 Sample path of consumer beliefs evolution for two-product case . . . . . . . . . 82

A.1 Subject interface for PWYW experiment, instructional stage . . . . . . . . . . . 88

B.1 Normalized expected profits per subject with quartiles . . . . . . . . . . . . . . 95

B.2 Assortment sizes over time. Both popular and nonpopular assortments are included 95

B.3 Time between the first and the last box clicks, seconds . . . . . . . . . . . . . . 96

B.4 Solution type conditional on the optimal assortment being selected at some point during the period . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 97

LIST OF TABLES

2.1 Confidence intervals: Robust estimators . . . . . . . . . . . . . . . . . . . . . . . . 23

2.2 Confidence intervals: Robust estimators . . . . . . . . . . . . . . . . . . . . . . . . 24

2.3 Subject demographics in PWYW experiment . . . . . . . . . . . . . . . . . . . . . . 28

2.4 Censored panel regression on PWYW prices for positive surplus values . . . . . 29

3.1 Popularity indices parameters . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 38

3.2 Problem parameters for each optimal variety condition . . . . . . . . . . . . . . . 38

3.3 Number of subjects per treatment . . . . . . . . . . . . . . . . . . . . . . . . . . 42

3.4 Panel regressions with factor interactions . . . . . . . . . . . . . . . . . . . . . . . 46

3.5 Censored regressions on the number of *too small* and *too large* popular assortments per subject . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 48

3.6 Number of different assortments seen and chosen by subjects over the entire experiment, mean and median . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 52

3.7 Assortments seen and choices in the first period . . . . . . . . . . . . . . . . . . . 54

3.8 Decision time, mean and median . . . . . . . . . . . . . . . . . . . . . . . . . . . 56

3.9 First period exploration and performance in the whole experiment: censored regression . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 59

B.1 Subject totals: censored regressions with full interactions . . . . . . . . . . . . . . 93

B.2 Cross sectional regression: last period . . . . . . . . . . . . . . . . . . . . . . . . . 93

B.3 Panel regressions with full interactions . . . . . . . . . . . . . . . . . . . . . . . . 94

# CHAPTER 1

# INTRODUCTION

With the rising skepticism regarding validity of simplistic economic models which are built on the assumption that firms and individuals maximize their expected pecuniary payoff, there is a need for both empirical investigation of economic decision making and the analytical tools that can accommodate this new knowledge for the purposes of business decision making. Individual choice models are gaining more importance in operations management as a foundation for both aggregate consumer demand models (e.g., Alptekinoğlu and Semple (2016), Jagabathula and Vulcano (2017)) and supply chain analysis (e.g., Croson and Donohue (2006), Katok and Pavlov (2013), etc.).

This dissertation includes three essays that focus on behavioral models in different contexts. The first two essays are concerned with empirical investigation of individual choice by means of laboratory experiment, and the third one is concerned with integrating one particular behavior pattern – consumer learning – into assortment decisions. In this introduction I will briefly describe the problems analyzed in each of these essays and the findings we have obtained.

In the first essay (Chapter 2), we study how customers are choosing prices in a Pay-What-You-Want business model.[1] Under Pay What You Want (PWYW) the price for a product is fully determined by a buyer: the seller cannot reject any offer. Recent studies have documented that even when buyers are allowed to pay zero, average payments are consistently higher. The distributions of buyer-determined prices vary significantly across different environments: payments around zero seem to be uncommon in PWYW restaurants, but they are prevalent in digital distribution. It is unclear whether such differences should be attributed exclusively to established social norms or they can be driven by some systematic

[1] The work reported in Chapter 2 of this dissertation constitutes a basis for a working paper (Özer and Vorotyntseva (2015)).

factors. We examine one of the possible factors: the seller's production cost. We hypothesize that buyers may anticipate the seller's loss-aversion, so they will be reluctant to choose prices below the costs. We suggest a model that incorporates this hypothetical behavior and estimate it using the dataset from the experiment of Schmidt et al. (2015). We also develop our own experiment and collect the data.

In the second essay (Chapter 3), we use a laboratory experiment to study the behavior of human assortment planners.[2] Assortment planning, that is, the selection of products to offer in a store or the design of a product line for a manufacturer, is often performed by managers without any support from computerized optimization algorithms. Managers may rely on sales data, demand forecasts and their hunches about interaction effects between products, but it is unclear how they incorporate this information into their decisions. To the best of our knowledge, there have been no empirical studies published that address the behavior of human decision makers in a context of assortment planning. In this study, we use a laboratory experiment to find if human decision makers deviate from the expected profit-maximizing solution in some systematic way and to investigate the effect of decision support tools on the efficacy of assortment planners.

Our subjects are asked to assume the role of a retailer on a competitive market and decide which products to include into their assortment. For each product included, they incur a fixed operational cost. To simulate the market demand, we use the seminal model of assortment planning from van Ryzin and Mahajan (1999), which uses the Multinomial Logit (MNL) model of consumer choice. In this setup, the optimal assortment has a simple structure: it must be a subset of products with the highest expected utility values. We are interested whether the subjects arrive to this conclusion, and if so, whether they include the optimal number of products in their assortment. In our experiment, the subjects are

[2] The work reported in Chapter 3 of this dissertation constitutes a basis for a working paper (Honhon et al. (2018)).

choosing from seven products, and the optimal assortment size is either two or six, depending on which treatment the subject is assigned to. In addition to that, we attempt to investigate the effect of decision support tools on the efficacy of assortment planners. To do so, we present subjects with intermediate metrics derived from the initial problem inputs, which are supposed to simplify the task of finding the optimal assortment. We implemented our experiment in a web application and recruited the subjects through the Amazon Mechanical Turk (mTurk). In total, 88 human subjects participated in our experiment.

Our results show that subjects who are facing the combinatorial problem of assortment planning make intelligent decisions. Most of them choose popular assortments, as the theory would suggest. However, subjects perform significantly worse when the optimal assortment consists of six products than when it consists in two products. Since in practice assortment planners have to deal with assortments of dozens, or even hundreds of products, the profits losses due to the flaws in the human assortment planners decision making might be much higher than our experiment seems to suggest. However, from the theoretical point of view our results are very encouraging: since human subjects decisions are overall consistent with implications of an analytical assortment planning model, even on a relatively small scale, the prospectives of further investigation of their decision making patterns are open, and our results can serve as a foundation for future research directions.

A somewhat surprising result is the direction of the decision support effect: in some treatment providing subjects with *more information* resulted in worse performance. This apparent adverse effect implies that the additional information can confuses and/or misguides the subjects, possibly by making the cannibalization effect (i.e., the fact that adding a product "steals" demand from previously included products) more salient, which causes the subjects to assign excessive weight to it.

In the third essay (Chapter 4), we develop a model that explicitly incorporates consumer's learning considerations into the firm's assortment problem.[3] Consumer's choice of a product from a particular category is influenced by her beliefs about how well the product fits her needs. When a consumer purchases a product repeatedly, her experience with it affects her beliefs and, consequently, her future choices. A particularly interesting case for studying the effects of consumer learning are the recently emerged subscription services, where at the fixed intervals of time a consumer receives one or more products from a certain category. This can be a pre-specified product (e.g., razor blades from the Dollar Shave Club), varying bundles picked by the seller (e.g., meal ingredients from Blue Apron), or a collection of items, from which the consumer can pick some and return the rest (e.g., clothes from Stitch Fix).

Since these services are novel, consumer learning dynamics has a very strong impact on the demand, while high variability of experience from consumption implies that a consumer may need multiple trials to understand whether or not the service is suitable for her. Due to random variations in experience even a potentially suitable product may produce in a series of unsatisfactory experiences that will result in the consumer abandoning the product.

Providing more than one product to choose from gives the firm an "second chance": it is more likely that the consumer will find a product that she deems suitable for her and stay with the firm indefinitely. However, larger assortment presents a challenge of cannibalization: a consumer may substitute a less profitable product for an equally suitable and more profitable one.

To do so, we construct a model where the consumer repeatedly purchases from a firm's assortment, each time picking the product that maximizes her current expected utility. After each purchase, the consumer gets a noisy signal about the product's utility and updates her beliefs on the product in a Bayesian fashion. If the expected utility of a product becomes lower than the utility of the outside option, the consumer stops buying.

[3] The work reported in Chapter 4 of this dissertation constitutes a basis for a working paper (Honhon et al. (2018)).

The firm is aware of the consumer's learning process but unable to view the signals she receives. We link this problem setup to a specific type of random walk known as a gambler's ruin problem and show how to represent it as a conditional renewal process. The derived properties of this model can potentially simplify the comparisons between different assortment policies in presence of learning consumers and make these comparisons more tractable.

# CHAPTER 2

# A BEHAVIORAL STUDY OF PRICE SELECTION IN PAY-WHAT-YOU-WANT BUSINESS MODEL

## 2.1 Introduction

Pay-What-You-Want pricing (PWYW) assumes a transaction of a good from a seller to a buyer, in which the buyer chooses and pays the price for the good. That is, unlike public goods funding, donations to charities and "pay it forward" contributions of open-source software developers, PWYW is a form of pricing. Classic economic theory implies that whenever buyers are allowed to choose a price for a commodity, they will always choose the lowest price possible, but there are exceptions to this law. In US restaurants customers decide how much to pay for the waiter's service, and in sandwich shop Panera Cares guests even choose how much to pay for their meal; Metropolitan Museum gives an access to its exhibitions for whatever fee visitors wants to pay, and Humble Bundle allows gamers to download Android apps for whatever price they choose. The recent researchers' attention to PWYW was brought by the musical band Radiohead who allowed the customers to download their album "In the Rainbows" and choose any price for the music.[1] This iconoclastic pricing strategy has attracted a lot of media attention. CNN Money has even included it into their rating: "101 Dumbest Moments In Business" with a sarcastic comment: "Can't wait for the follow-up album, 'In Debt' " and added: "Sixty-two percent, according to comScore, decide to pay nothing, while the other 38% voluntarily fork over an average of six buck."[2] Nevertheless, in various interviews the band members expressed deep satisfaction with the total revenue. Indeed, 38% of listeners paying for the music seems to be not bad in the age of

[1] A $0.90 required fee was charged for the download service.

[2] Editors of Fortune (2008, Jan 16). 101 Dumbest Moments in Business. Retrieved from money.cnn.com. Last Accessed on April 14, 2014

digital piracy. And according to the data of (Elberse and Bergsman, 2008), though a digital album sold on iTunes was typically priced at $9.99, an artist got only about $1.40 out of it. By distributing "In the Rainbows" through their own website, Radiohead band captured all the money that customers paid for their music.

Radiohead example illustrates an important potential benefit of PWYW pricing: a decrease in production and distribution costs. Since buyers choose prices themselves, a PWYW product does not need piracy protection, which can significantly reduce the production costs and contribute to customer satisfaction.[3] Summarizing the above, in the digital distribution area PWYW pricing can reduce institutional costs and lead to efficiency increase. Switching from list pricing to PWYW can benefit both a buyer and a seller, as long as the prices chosen by buyers are high enough. However, in practice customers seem to be less generous when it comes to digital distribution, compared to tipping or paying for meals settings. Riener and Traxler (2012) report the distributions of payments in PWYW buffet Wiener Deewan during their study, where the most common price for a meal was about 5 EUR, which is about the average price tag in similar restaurants nearby, and many customers paid more while virtually no one paid below 2 EUR. In contrast, the most common price for the the videogame World of Goo when it was offered under PWYW was 1 cent, which was the minimum allowed price, and very few people chose $20, which was the regular price for this game, about the same as for other games in its class.[4] In this paper, we attempt to answer the question: what can be the reason for such differences in customers' behavior?

Behavioral economics literature has repeatedly shown that even in the absence of external enforcement some people consistently sacrifice some part of their own monetary payoff in order to increase welfare of other people (see (Roth, 1995b) for a comprehensive review).

[3]For example, in its messages to customers a PWYW video game distributor Humble Bundle accents that their products are DRM-free. DRM stands for Digital Rights Management system that protects a product from illegal copying, but sometimes can cause problems with legal usage of the protected product as well.

[4]Retrieved from 2dboy.com. Last Accessed on May 15, 2014

Other-regarding behavior can be motivated by altruistic, fairness and reciprocity considerations. To measure fairness and altruism in environments where the decision-maker's behavior is not influenced externally, researchers often use decision-making experiments based on a so-called Dictator game. This is a trivial game: one person splits some amount of money between self and the other player; the latter one cannot make any action. A person who cares only about the own monetary payoff must allocate all the amount to self and leave nothing to the partner, but numerous laboratory experiments have found that most people share. A meta study by (Engel, 2011) finds that across 129 Dictator game studies published between 1992 and 2009, on average dictators give 28.35% of the total amount. In this paper we link an interaction between a buyer and a seller in a PWYW market to the Dictator game, which allows us to utilize existing findings of behavioral economics for investigating the buyer's choice of a price.

Understanding the factors that influence buyers' behavior is a necessary condition for building analytical models of PWYW pricing and forecasting its profitability. As of today, the research in this area is scarce, and the results of existing studies are sometimes very surprising. For example, though intuition suggests that buyers are likely to pay more when they are observed, in the famous series of PWYW field experiments reported in (Gneezy et al., 2010) and (Gneezy et al., 2012) varying the degree of payment anonymity had no effect on the average buyer-determined prices. They also observe a significant decrease in the quantity demanded under PWYW compared with a fixed low list price (see the Literature Review section for more details). Some anecdotal evidence also reflects what can be called a shrinking demand phenomenon. For instance, the owner of *Santorini Grill* restaurant in New York, switched from list pricing to PWYW and was pleased to see that the average payments for entrees did not go down.[5] However, *Santorini Grill* was closed just four months later. In

[5]Reddy, S. (2011, Nov 11). To pay or not to pay. Wall Street Journal (Online). Retrieved from proquest.com

the interview to Gothamist the owner explained that customers did not freeload, but they "just stopped coming".[6] This looks very surprising indeed, because unlike an economist would expect, under PWYW people can decrease their consumption compared to a list-price strategy.

In this paper we focus on two factors that might affect the price under PWYW: the seller's cost and the buyer's private valuation. The cost factor is very important for many PWYW sellers: in various media interviews owners of such businesses explain that their main goal is customers' satisfaction or serving as many people as possible, as long as revenues cover costs. Costs are an important concern even in a digital distribution: marginal costs per copy sold are typically perceived as zero by outside observes, but in reality they exist. Bandwidth and customer support expenditures are incurred with each sold copy, and can add up to very considerable amounts. PWYW businesses claim to be socially oriented, but very few of them can stay in business if every product sold adds to their losses. It is well known that people are more sensitive to losses than to gains, so it is possible that they apply this reasoning when thinking about other's payoff, because they "put themselves in the shoes of other person" or, in other words, due to empathy. (Batson et al., 2009) argue that empathy can induce altruistic behavior, and (Stahl and Haruvy, 2006) conclude that empathy concerns increase sharing in a Dictator game. If a buyer does indeed empathize with the seller's loss aversion, some of the buyers who would otherwise pay zero, may raise their prices if they learn about the seller's costs, ensuring that the seller does not incur losses. If such considerations take place, they can explain some differences in buyer-determined prices that are observed in different types of business. For instance, in the example of Wiener Deewan and World of Goo, described in the first paragraph, customers may perceive that downloading a video game for (almost) free does not do any harm to the developers, but restaurant visitors may

[6] Johnston, G. (2012, Mar 15). New York's Only "Pay What You Feel" Restaurant Closes. Gothamist. Retrieved from gothamist.com. Last Accessed on April, 13 2014

be concerned of the meal's productions costs and try to cover them. In this paper we verify if buyers on a PWYW market will indeed feel more concern about causing seller's losses than about just splitting the surplus in an inequitable way.

Understanding the effect of buyer's private valuation is necessary for establishing if PWYW can be profitable business model. Some researchers have argued that PWYW can elicit private willingness to pay of the customer and induce self-selection (e.g., Mak et al. (2015), Raju and Zhang (2010)). An example may be as follows: the listeners who are "fans" of a musical band will select higher prices for its song than "casual listeners". Therefore, PWYW may increase profits by exploiting the heterogeneity in consumers? willingness to pay. However, some evidence from behavioral research casts doubts on this conjecture. It has been repeatedly shown that the consumers perceive exploitation of higher willingness to pay as unfair (see Özer and Zheng (2012) or Roth (1995a) for comprehensive reviews). In addition to that, the theory of cognitive dissonance by Konow (2000) suggests that people tend to distort their fairness beliefs in their favor. Therefore, one might expect that buyers will not be willing to pay more if they valuation is high.

To study the effects of the seller's costs and the buyer's valuation, we build on the fairness theories that exist in the economics literature to model a buyer-seller interaction and derive the predictions for two possible cases: if a buyer empathizes on the seller's loss-aversion and if he or she does not. In addition to that, we use secondary data from an experiment from (Schmidt et al., 2015) to estimate our model, and this data appears to support our hypothesis. In addition to that, we design a laboratory experiment to test the theoretical predictions and verify if buyers indeed perceive the seller's gains and losses differently.

Our study also contributes to the literature on fairness theory. As we point out in the literature review, among the numerous studies on altruism, very few have studied altruism considerations in the situations where parties can incur losses. Our analytical framework illustrates how possible loss considerations can be incorporated into the existing fairness models.

## 2.2 Literature Review

In this section we provide a brief review of the academic studies that are relevant to PWYW pricing and to the methodology that we apply to investigate it. First, we address the empirical studies that shed some light on the factors which influence the buyers' behavior on a PWYW market. Both intuition and research suggest that buyer-determined prices depend on the buyer's fair-mindedness and altruistic concerns. Because of this, we then address economic literature on fairness and altruism, particularly, on distribution games – an extensively used tool for research of other-regarding behavior. Finally, we describe analytical models of PWYW pricing, suggested in the academic literature. We show how our findings can provide evidence that will help to improve PWYW models.

Academic literature on PWYW pricing consists primarily of empirical investigations: they describe the buyers' behavior in PWYW situations and attempt to elicit factors that influence the buyers' generosity. Perhaps the earliest study of PWYW pricing was conducted by Lynn (1990), who observed payments in a restaurant where customers could choose from four price tags for some entries (for example, for a Mexican Pizza was sold for \$4.70, \$5.45, \$6 or \$6.50). During the observations almost half of the guests selected prices that were higher than the allowed minimum. The paper argues that such behavior cannot not be explained by literature on customer behavior (existing by the time of publication) and calls for more studies. This call has been left unanswered for a while, until the successful PWYW-release of "In the Rainbows" album has inspired a new wave or research. Kim et al. (2009) reports on three field experiments where different goods (cinema tickets, hot beverages and buffet lunches) have been suggested under PWYW pricing. Though a zero price has been available in all treatments, all buyer-determined prices has been higher. In these experiments buyer interacts with a seller's representative face-to-face, and later research investigates if it is a necessary condition for obtaining positive payments. Gneezy et al (2012) conduct a field experiment in a PWYW buffet where they ensure anonymity of the visitor's payments.

Surprisingly, they find that in the absence of face-to-face interactions average payment do not decrease compared to the control group. This anonymity experiment is a part of notable PWYW experiment series summarized in Gneezy et al. (2010) and Gneezy et al. (2012). Other intriguing result they report is demand decrease that was observed in some PWYW treatments: when a souvenir photo from a boat tour have been suggested under PWYW pricing, less people have decided to purchase it compared to the treatment where a photo has been sold for $5. Based on these results, Gneezy et al. conjecture, that when choosing the price, buyers may want to look good in their own eyes, not in the eyes of other people. Among other factors that influence the buyer's choice of a price, different experimenters document buyers' mood Riener and Traxler (2012), preference for round prices (Lynn et al., 2013), seller's suggested price (Johnson and Cui, 2013), prices of similar products, and fairness considerations (Kim et al., 2009). We contribute to empirical studies of PWYW pricing by adding one more factor to the picture of voluntary payment determinants: the seller's production cost.

Fairness and altruism considerations are generally believed to play a key role in the buyer's behavior; as Mak et al. (2015) say:

> The prevailing wisdom on "pay what you want" (PWYW) pricing is that consumers' payments depend largely on their sense of altruism or of fairness.

Behavioral economics literature has a stream of studies on fairness and altruism, which can help our understanding of PWYW pricing. Fairness literature has made a long path from laboratory experiments on the Prisoners Dilemma in 1950-s to predictive analytical models emerged in 1990-s. Social preferences are typically measured by simple laboratory games: Prisoner's Dilemma, Public Goods game, Third-party Punishment game, Gift Exchange game, Ultimatum game, Trust game and Dictator game (see Camerer and Fehr (2004) for a detailed review). All of these games, except for perhaps an Ultimatum game, can be used for conceptualization of a PWYW framework.

Dictator game studies are very similar to our conceptual framework of PWYW pricing and to our experimental design, so we provide a brief review of this stream of literature as well. Dictator game has been designed to isolate altruistic considerations from other concerns in social preferences games. Its most recognized framework is first introduced in Forsythe et al. (1994): subjects are matched in pairs, and one person (dictator) in each pair divides some amount of money ("pie") between self and a partner (recipient). Forsythe et al. (1994) show, that though dictators transfers to recipients on average are less then transfers in games where participants can punish each other for "unfair" actions, most dictators still allocate positive amounts of money to their recipients, some give even 50% of a pie, and this behavior is independent of the pie size (they use $5 and $10 in their treatments). By now there exist many experiments which study factors that affect dictator giving. Hoffman et al. (1994) argue that most of dictator giving is explained by lack of subject-experimenter anonymity. They conduct a double-blind study and show that dictators' generosity is significantly decreased compared to a standard design, where an experimenter can observe dictators' actions. However, Bolton et al. (1998) cast doubt on this finding: in their experiment they do not find statistical difference between dictator giving in standard and double-blind procedures. They also find that when the dictator choice is restricted, and dictator's first-best choice is unavailable, dictator picks an option which is gives him or her higher monetary payoff than the first-best choice ("I'm-no-saint hypothesis"). Ruffle (1998) and Oxoby and Spraggon (2008) show that dictators are much more generous when the "pie" is "earned" by recipients (the pie size for each pair in their experiments is dependent on the recipient's performance on some test). Dana et al. (2006) finds that when potential dictators are given a choice between playing a Dictator game with $10 pie and receiving $9 (in the latter case a recipient does not learn about the game and gets nothing from it), many subjects decide to avoid the Dictator game. The researchers conclude, that dictator giving must be driven by recipients' expectations that are absent in an opt-out ($9) option. Stahl and Haruvy (2006) find, that in treatments where both subjects in a pair make allocation decisions, and then random chance determines which of two allocations is realized, dictators are significantly more generous than in treatments where the dictator is determined *ex ante*. They attribute this effect to empathy: empathy considerations are more likely to emerge when a dictator is forced to imagine being a recipient. In contrast, actual switching roles may have an opposite effect: Ben-Ner et al. (2004) show it in a two-part dictator game. In their design, after a standard dictator game a second part is announced, where the roles are swapped and subjects are rematched. Second-round dictator decision turns out to be strongly and positively correlated with what he or she has received in the first round, despite the rematching. Panchanathan et al. (2013) use Dictator game to investigate the "bystander effect".[7] In their experimental design there are several dictators per one recipient. Though the pie size to group size ratio are the same across all treatments, on average a recipient matched with a single dictator gets significantly more money than a recipient matched with two or three dictators. We take the previous findings into account when developing our theory of buyer's price selection and when choosing our experimental design and procedures.

Since a PWYW seller who incurs production costs, can incur losses, of particular interest to us are Dictator game studies where payoffs can be negative, but at this moment there are very few of them. Baquero et al. (2013) investigate the effect on the pie domain (positive or negative) on the subjects' behavior in Dictator and Ultimatum games. In Ultimatum game a recipient (in this context referred to as a responder) can "punish" the Dictator (proposer) by rejecting the allocation decision, and then both participants get nothing. Baquero et al. (2013) find that responders punish same unfair allocations much more frequently in a treatment when a pair is instructed to "share losses" compared to a treatment where they are instructed to "share gains"; proposers anticipate it and increase responders' shares. Dictators' generosity is also increased in a loss domain; the extent is much smaller, but statistically

[7] Bystander effect mens that a person is less likely to help someone in need when there are other potential helpers

significant. However, the researchers adopted within-subject experimental design, and direct comparison of loss and gain treatments might explain at least part of the observed difference in dictators' behavior. Another study of the domain effect on dictator giving by Bardsley (2008) includes three treatments: dictators were suggested to "give to", to "take to", or to "either give to or take from" another player. They find that when the option of "taking" is added to the option of "giving", the frequency of positive giving is decreased compared to the treatment where only "giving" is possible. Their results may imply that dictator giving can be influenced by a so-called "context effect" – a violation of the "independence of irrelative alternatives" principle of the individual choice theory (see Özer and Zheng (2012) for more information on this effect). Our experiment can provide us further evidence of the domain effect on the dictator choice and add to this stream of literature.

Now we circle back to PWYW pricing literature and describe controlled laboratory experiments designed specifically for this framework – so far there were very few of them. The study by Mak et al. (2015) adopts a public goods prospective: in their model a seller provides a product for PWYW price in each following period only if the total revenue collected in current period exceeds some pre-determined value. They argue, that with a threat of service abruption and sufficient coordination among buyers, PWYW pricing can sustain even in absence of fairness considerations towards the seller, and demonstrate it in a controlled laboratory experiment.

The closest work to our study is the recent experiment by Schmidt et al. (2015). Their experiment aims to elicit which of the fairness-related mechanisms drive the buyers' choice of a price in a PWYW market: outcome-based fairness (buyers are averse towards inequitable payoffs), reciprocity (buyers award the seller's choice of pro-social pricing and hence pay more than zero), or strategic concerns (buyers are selfish, but pay positive prices to keep the seller in the market). They conclude that strategic concerns drive the buyer-determined prices up, but buyers pay positive prices even in the absence of selfish incentives. They don't find sufficient evidence for reciprocal considerations. Prices tend to go up as the sellers' costs and buyers' valuation go up, which is interpreted as a support for outcome-based fairness considerations.

Our paper aims to clarify the impact of the seller's price further by suggesting and testing the analytical framework of the buyer's behavior. We estimate our model using the data from the Schmidt et al. experiment and design another experiment specific to our model.

## 2.3 Model Development

In this section we formally describe the sequence of events in a buyer-seller interaction on a PWYW market. To analyze the effect of the seller's cost and losses associated with it, on the buyer-determined price, we attempt to strip our model from other elements that might affect the behavior, such as positive reciprocity induced by switching to PWYW or sunk costs fallacy.

### 2.3.1 Problem Specification

We build our model as a one-time interaction between a single seller who produces a product at marginal cost $c$ and a single buyer who has a valuation $v$ for the product. At the beginning of the game the seller produces one unit of a product at marginal cost $c$ and suggests it to a buyer under PWYW pricing modality. Both $v$ and $c$ are common knowledge.

The buyer decides whether to purchase the product or to forgo the purchase. If the buyer decides to buy the product, he or she chooses a price $p \geq 0$, pays it and gets the product. The resulting payoffs from the interaction are $v - p$ for the buyer and $p - c$ for the seller. If the buyer decides not to purchase the product, the seller disposes the product and gets a salvage value $c$, which recovers the initial investment. The resulting payoffs from the interaction are $(0, 0)$. This sequence of events is shown in Figure 2.1.

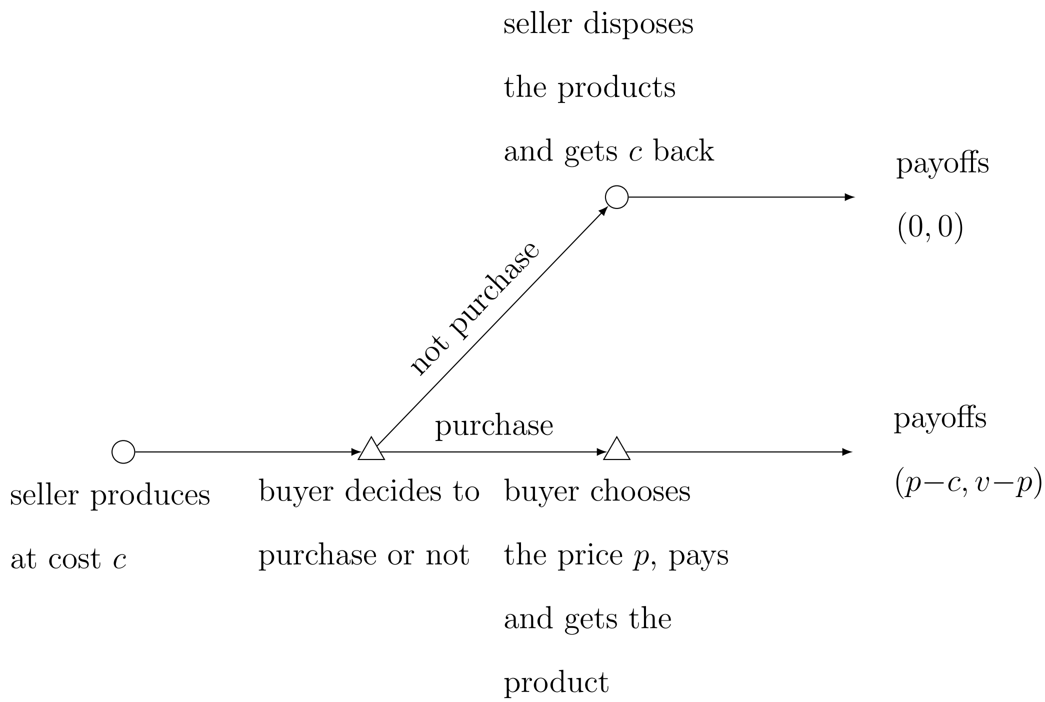


FIGURE 2.1: Sequence of events in a PWYW interaction

Several features of our conceptual framework need to be emphasized. First, the seller has no option to refuse a buyer's offer even if $p < c$, which means that the payoff of both the buyer and the seller in this interaction are fully determined by the buyer. Second, our framework leaves out the seller's choice decision to enter the market and the choice of pricing modality – since the reciprocity considerations are out of the focus of our study, we omit these steps.[8] Third, we allow the buyer to forego the purchase opportunity. Though as long as $v > c$ buyers have neither monetary nor fairness incentives to do it, if they don't have an option to forego the purchase, they may perceive the seller's costs as sunk and adjust their behavior accordingly. To leave out the sunk costs considerations completely, we allow the

[8] To further motivate leaving the seller's choice of pricing modality out of the framework, we note that many real-world PWYW businesses do not have this choice. For instance, many New York museums get support from the city, and they are obliged to ensure access for all the New Yorkers (Kadet, A. (2013, Nov 09). City News – Metro Money: To be cheap, or not to be. Wall Street Journal). Though buyers may positively reciprocate a seller's switching to PWYW, common intuition suggests that if all businesses stick to PWYW, reciprocal considerations must vanish over time. It is unlikely that restaurant customers reciprocate PWYW format of the waiter's service when choosing the amount of tip.

seller to dispose the product and receive exactly the same value as the costs incurred during the production.

Finally, for a purpose of expositional simplicity we use one-buyer-one seller setting instead of one-buyer-multiple-sellers. One-to-one setting rules out the bystander effect from our model – as we mentioned in the literature review, this effect have shown to distort the behavior in similar economic games. We believe that the bystander effect is very likely to take place in the real-world PWYW interactions, but for the purpose of our study we aim to strip the hypothesized seller's cost effect from other possible impacts.

Our proposed conceptualization can be linked to the Dictator game that was described in the previous sections. There are two essential differences between our PWYW game setting and the standard Dictator game setting. The first difference is the context: a typical Dictator game is context-free; PWYW game is context-rich.[9] The second crucial difference is the payoff range: in a typical Dictator game the payoffs are restricted to be nonnegative; in PWYW game the seller's payoff from an interaction with a buyer can be as low as $-c$.

### 2.3.2 The Fairness Model

In our setting two players – a buyer and a seller – split a value of $v - c$, and the buyer can allocate some value – positive or negative – to the seller. Buyer's utility from purchasing the product is determined by two factors: a private surplus (the difference between buyer's private valuation of the product $v$ and the price he or she pays to the seller $p$) and disutility from unfairness (the difference between the price $p$ and the fair price). We build our model on the (Fehr and Schmidt, 1999) parametrized inequity aversion model by adding the possibility of buyer's empathy on the seller's loss aversion:

[9] Nevertheless, there are context-rich Dictator games in the literature. (Hoffman et al., 1994) in their experimental instructions refer to the roles as a "buyer" and a "seller". As in our PWYW setting, they frame the transaction as a purchase: a seller (dictator) sets the price and a buyer (recipient) is forced to purchase at that price. In our setting the roles are reversed.

$$U(v, p, c) = (v - p) - \alpha((p - c) - (v - p))^{+} - \beta((v - p) - (p - c))^{+} - \gamma(c - p)^{+}, \quad (2.1)$$

where $(\cdot)^{+}$ denotes $\max\{\cdot, 0\}$; $\alpha$ and $\beta$ are buyer-specific parameters that determine his or her attitude towards the disadvantageous and advantageous inequality respectively, parameter $\gamma \geq 0$ determines buyer's sensitivity to the seller's losses. That is, when a price is decreased by one dollar, the buyer's psychological penalty decreases by $2\alpha$ if his or her surplus is lower than seller's one, and increases by $2\beta$ if his or her surplus is higher than the seller's, but the seller's payoff is still positive. $\alpha$ will be further referred to as an "envy factor"; $\beta$ will be referred to as a "guilt factor". The model assumes $\alpha > \beta$ and $\beta \in [0, 1)$ (see (Fehr and Schmidt, 1999) for justification of these assumptions). If a seller incurs losses, a buyer incurs a penalty of $2\beta + \gamma$ for every dollar decrease of the price. $\gamma = 0$ means that the buyer evaluates seller's gains and losses equally, and our model reduces to a classic Fehr and Schmidt model.

Let $r$ be the "fair price" – a price that results in an equitable split of the surplus created between the seller and the buyer: $r = 0.5(v + c)$. When $p = r$, the buyer does not incur any unfairness-related disutility. Let us further assume that the price is limited to be $p \geq 0$. Then the buyer-determined price is:

$$p = \begin{cases} 0, & \text{if } \beta < 0.5 \text{ and } 2\beta + \gamma < 1; \\ c, & \text{if } \beta < 0.5 \text{ and } 2\beta + \gamma > 1; \\ r, & \text{if } \beta > 0.5. \end{cases} \quad (2.2)$$

if $\beta = 0.5$, the buyer is indifferent between any price in the interval $[c, r]$; if $\beta < 0.5$ and $2\beta + \gamma = 1$, the buyer is indifferent between any price in the interval $[0, r]$. Note, that if $\gamma = 0$, like in the original Fehr and Schmidt model, $\beta < 0.5$ implies $2\beta + \gamma < 1$, and the option $p = c$ can never be chosen.

Now we show how the described model can explain the decrease in quantity demanded under PWYW pricing modality, which has been reported by some observers.

PROPOSITION 1. *If the buyer's valuation $v$ is less than the seller's cost $c$, and $2\beta + \gamma > 1$, a buyer will not purchase the product under PWYW pricing.*

*Proof.* According to 2.2, a buyer with $2\beta + \gamma > 1$ achieves the maximum utility from purchasing the product either at $p = c$ or at $p = r$. Since $v < c < r$, the buyer's resulting payoff is negative at the point of maximum. Because of this, the buyer's first best option will be to forego the purchase decision and stay with a payoff of zero. □

In other words, Proposition 1 means that if a sufficiently fair-minded buyer believes that the seller's costs may be higher than his or her private valuation, he or she will avoid purchasing the product.

### 2.4 Study 1. Data from Schmidt et al. (2015)

The experimental framework of the study by (Schmidt et al., 2015) is relevant to our setting, and their data may be appropriate for a preliminary test of our hypothesis. Their experiment was conducted in controlled laboratory environment with z-tree software (see (Fischbacher, 2007)); the z-tree output files were provided in the online appendix for the original paper. Subjects were undergraduate students from the University of Munich and the Technical University of Munich. The experiment consisted of 3 treatments; we use only the first (Base) treatment in our analysis.

In the base treatment subjects were divided into groups of four; one subject was assigned a seller role, three other subjects were assigned the buyer roles. Then they played five periods together as follows: in the beginning of each period a seller decided whether to enter the market or not. Then, if she decided to enter, she decided whether to invest two units of resources into product quality or not. The quality investment costed the seller 2

Experimental Units (EU), and it doubled the buyer's valuation of the product. Then the seller's production cost and the buyers' base valuation were realized. The seller's production cost was drawn at random from the set $\{0, 1, 2, 3, 4\}$ EU, the buyer's base valuation was drawn from $\{2, 6, 10\}$ and doubled if the seller's decision was "to invest". Every buyer knew his own valuation, the seller's production cost and whether the seller decided to invest or not. Every buyer made a decision to purchase the product or not, then, if the decision was "to purchase", a buyer determined her price for the product. The pecuniary payoff of buyer $b$ from this interaction was:

$$M^b = \begin{cases} 0, \text{ if the seller } s \text{ did not enter the market or the buyer } b \text{ decided not to buy,} \\ (1 + I^s)v^b - p^b \text{ otherwise,} \end{cases}$$

where $I^s$ is the indicator variable which takes the value 1 if the seller invested into quality and 0 otherwise, $v^b$ is the buyer's base valuation of the product, $p^b$ is the price that she decided to pay.

The pecuniary payoff of the seller $s$ is:

$$M^b = \begin{cases} 0, \text{ if the seller } s \text{ did not enter the market,} \\ \sum_{b=1}^{3} B^b(p^b - c) - 2I^s \text{ otherwise,} \end{cases}$$

where $B^b$ is the indicator variable which takes the value 1 if the buyer $b$ decided to purchase the product and 0 otherwise.

This sequence repeated 5 times, and then the participants were rematched. Every experimental session consisted of 4 such blocks, roles were fixed during all the session. The base treatment consisted of 4 sessions with 72 buyers, total of 1192 observations.

Figure 2.2 depicts the prices paid by the buyers, grouped by levels of seller's variable costs. We reject the hypothesis of the presence of nonstationarity in the data (Dicker-Fuller

-10.314, p-value = 0.01), so we do not need to control for the buyers' learning trends in our data analysis.

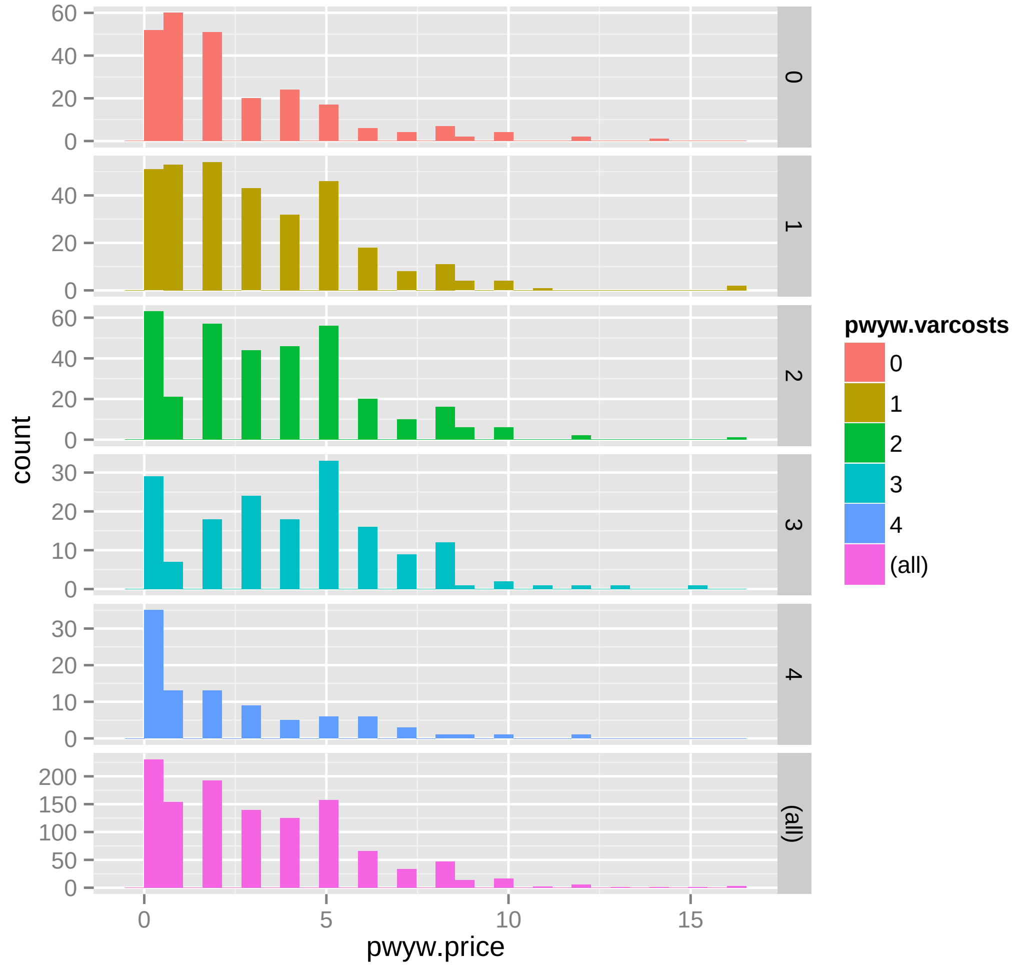


FIGURE 2.2: Buyer-determined price frequencies in the study by (Schmidt et al., 2015), grouped by the production cost value.

The option of the seller's investment into quality was originally introduced in order to elicit the influence of buyers' reciprocity considerations: if such considerations were present, a seller who voluntarily chose to invest into quality would get on average higher prices than a seller whose decision to invest was imposed endogenously (exogenous investment was imposed in one of the remaining two treatments). Fehr et al. come to the conclusion that the reciprocity considerations *per se* did not influence the buyer-determined prices, but the additional investment expenditures might have had an effect as a part of the seller's costs. To accommodate the investment expenditures in our model, we split the investment cost equally among the three buyers. In other words, we assume that if the seller has invested, the perceived cost of each buyer will be increased by $\frac{2}{3}$ EU.

Result 1: Seller's production cost serves as a tipping point for the buyer-determined prices distribution. The distribution parameters may suggest that buyers experience more guilt for reducing price by 1 EU when the price is below the seller's costs than when the price is above.

To test our hypothesis of empathized loss aversion, we fit the regression model with individual fixed effects.

$$pwyw.price = f(\beta guilt + \gamma spline),$$

where *guilt* and *spline* correspond to the third and the fourth items in (2.1) respectively. Note that we don't add the "envy" factor $\alpha$ since similar studies have shown that subjects virtually never treat themselves unfairly.

Since the buyers were restricted to discrete nonnegative prices, we use the Poisson regression model for count data. An important benefit of the Poisson functional form is its close correspondence to the Quantal Equilibrium paradigm which is well established in the behavioral economics literature (see (McKelvey and Palfrey, 1995)).

TABLE 2.1: Confidence intervals: Robust estimators

| | 5% | 95 % |
|---|---|---|
| (Intercept) | 1.4887496 | 1.8004650 |
| guilt | -0.1394454 | -0.1131403 |
| spline | -0.7536681 | -0.5847453 |

The results of this regression are in Table 2.1. The coefficient of interest, which corresponds to $\gamma$ in our model, is negative and significant, which suggests that buyers may indeed experience empathic concern regarding the seller's loss aversion.

An eyeball examination of Figure 2.2 suggests that there might be a spike at the price of 5 EU. This goes in line with the evidence from (Lynn et al., 2013) who shows that buyers prefer round prices when they choose their prices themselves. To account for this effect, we add a dummy variable for the price of 5 EU. We also want to control for the possible anchoring effect: buyers may chose the price equal for the seller's production cost just because they were primed with this number, so we add a corresponding dummy variable as well (note: for the anchoring effect we use the original variable costs, not adjusted for the investment cost). We also add a dummy variable for the last period in a block, which was shown to have a significant effect on the price in the original study by Schmidt et al.

Table 2.2: Confidence intervals: Robust estimators

| | 5% | 95 % |
|---|---|---|
| (Intercept) | 1.4617766 | 1.7751074 |
| guilt | -0.1432192 | -0.1172132 |
| spline | -0.6758242 | -0.5346733 |
| last.period | 0.1143907 | 0.2294298 |
| I(pwyw.price == 5)TRUE | 0.2473939 | 0.3583301 |
| I(pwyw.price == pwyw.varcosts)TRUE | -0.6973038 | -0.4736757 |

As we see in Table 2.2, though the additional variables have a significant effect, our coefficient of interest remains significant. Unexpectedly, the indicator variable for the seller's cost is negative and significant, which is the opposite to the proposed anchoring effect, and may need further investigation.

## 2.5 Study 2. Pay-What-You-Want Experiment

The data from Schmidt et al. (2015) provides us with some helpful preliminary insights. However, since this experiment was designed with another research question in mind, the data interpretation is problematic. In addition to that, their data does not allow us to test how the buyers will behave when the buyer's valuation is salient. Hence, we design and conduct our own experiment.

### 2.5.1 Experimental Design and Procedure

In our experiment, subjects are assigned to computer terminals and separated by blinds. Before the experiment they read the paper instruction, then listen to the presentation accompanied by slides which repeat and restate the instructions. After that the participants answer a short quiz to ensure their understanding.

Roles are randomly assigned and kept constant during the treatment. The subjects play two practice rounds and 30 playing rounds. After each round, the participants are randomly rematched. In each round a buyer is required, first to decide whether to buy the product or not to buy, second, if the decision is "to buy", to submit the price for this product. The buyer's decision appears on the screen of the seller he or she is matched with, but a buyer does not get any feedback.

The buyers' valuations vary from \$3 to \$40, and the seller's costs varied from \$0 to \$10. Each buyer faces the same set of cost-value pairs, but in different sequences. From the point of view of each subject, his/her sequence of values and costs should look as IID from discrete uniform distribution, because we want to avoid any second-guessing of the patterns or experimenter's goals on the buyers' side. Among 30 rounds, 25 had positive surplus (i.e. buyer's valuation was higher than the seller's cost), 4 had negative surplus and 1 had zero surplus; in 5 rounds the seller's cost was zero, in 25 rounds it was positive.

Both the buyer and the seller know the cost. With regard to the buyer's value, we have two conditions: *Base treatment*, where both the buyer and the seller observe it, and *Asymmetric information*, where the buyer knows her valuation, but the seller knows only that it is in the range from \$3 to \$40. Every subject is exposed to only one condition.

In the end one round is randomly selected, and the prices submitted in that round are paid. In addition to that, each subject gets a show up fee of \$10. If a price received by the seller in the selected round is lower than the production cost, the difference is subtracted from the show up fee. For our parameters, the subjects' earnings range from \$2 to \$50.

Here we describe the rationale behind our procedures. First we address the multi-round structure of our experiment: every dictator made exactly the same decision nine times. While it is reasonable to expect that the decision must be the same in every round, in (Bolton et al., 1998) many subjects in similar setting made very different decisions in every round, distributing money in a "capricious", possibly random manner. Nevertheless, when pooled, the distribution of the allocation decisions in a multi round treatment turned out to be statistically indistinguishable from the one-round treatment. Because of this, though our conceptual framework assumes only one interaction, we want to have several rounds in order to decrease the random noise in our data.

Second, we justify the random round selection for the final cash payment. The purpose of it is eliminating the bystander effect: if every round gets paid, one seller receives money from several buyers, and it may lead to the behavior distortion as observed by in (Panchanathan et al., 2013). Some concerns about the random payment structure in decision making games have been previously raised in the literature: (Sefton, 1992) and (Stahl and Haruvy, 2006) find that random payoff structure on average increases the subject generosity. However, in (Sefton, 1992) it were subjects to be paid who were chosen at random: a dictator got the money he or she allocated to self only with 25% chance. In (Stahl and Haruvy, 2006) participants were divided in groups, everyone made a decision as if he or she was in charge for the final allocation; in the end one participant in each group was *ex post* chosen to be a “realized” decision maker. They argue that the random payment structure diminishes individual responsibility for the final allocation: since there was only a small chance that a participant will get the monetary payoff allocated to self, the expected cost of demonstrating generosity was decreased. In our setting one decision of a buyer is paid for sure, so the buyers must be less inclined to pretend being more generous.

Third, in our setting the buyers can only pay round prices: the payments must be discrete. It is commonly acknowledged that individuals tend to prefer round numbers. For instance, (Lynn et al., 2013) shows that restaurant guests typically either choose a round number as a tip amount or select an amount that will round their bill to a whole number. The effect of round prices is out of the focus of our study, so we would like to avoid it. To do it, we allow the subjects to make payments only in increments of $1.

Fourth, we give a lot of attention to the instruction procedure. We want to ensure the buyers’ understanding of the seller’s payoff structure instead of referring only to the part of the instruction relevant to their role. Therefore, in addition to the printed instructions, our subjects go through an interactive “demo stage”, where they get to see the experiment from both the seller’s and the buyer’s prospective and then take the instruction quiz. Only after they successfully pass the quiz, they get to know their roles in the experiment. A screenshot from the demo stage is provided in Appendix A.

### 2.5.2 Results

We have conducted four experimental sessions in August-September 2015, with 12-18 subjects (6-9 buyers) in each session. The experiments were conducted in Center and Laboratory for Behavioral Operations and Economics (CLBOE) at The University of Texas at Dallas (UTD).

UTD hosts a large population of international students, and academic literature shows that cultural background plays a significant role in fairness-related decisions (e.g., Roth

Table 2.3: Subject demographics in PWYW experiment

| | Full | Asymmetric | Total |
|---|---|---|---|
| USA | 4 | 5 | 9 |
| International | 10 | 9 | 19 |
| Total | 14 | 14 | 28 |
| Female | 3 | 4 | 7 |
| Male | 11 | 10 | 21 |
| Total | 14 | 14 | 28 |

et al. (1991) and Özer et al. (2014)). To control for the cultural background, in the post-questionnaire we ask the subjects about their country of birth. We also ask those who reported being born outside of the USA how long they have been in the USA. Among those who reported being born outside of the USA, no one had spent in the USA for more than 10 years, and most of then had spent 1-2 years, so we believe that the country of birth reflects the cultural background well enough.

Among 28 subjects who participated in our study in the role of buyers, 9 were from the USA, 14 from India, and 1 each from Pakistan, South Korea, Sweden and China. One subject did not specify the country and answered "Asia". For the purposes of our analysis, we divide the subjects into two groups, and call those who reported being born in the USA as *USA students*, and those who reported being born in other countries as *international students*. Table 2.3 shows the summary of the demographic data.

Figure 2.3 shows the distribution of the surplus share allocated to the seller for the positive surplus values and positive costs. We observe that there is a considerable difference between the results for USA and international students. Information asymmetry, however, does not appear to have a negative effect that we expected to see.

We regress the price paid by subjects on treatment and demographic variables using censored regression with random effects and report the results in Table 2.4. It confirms that

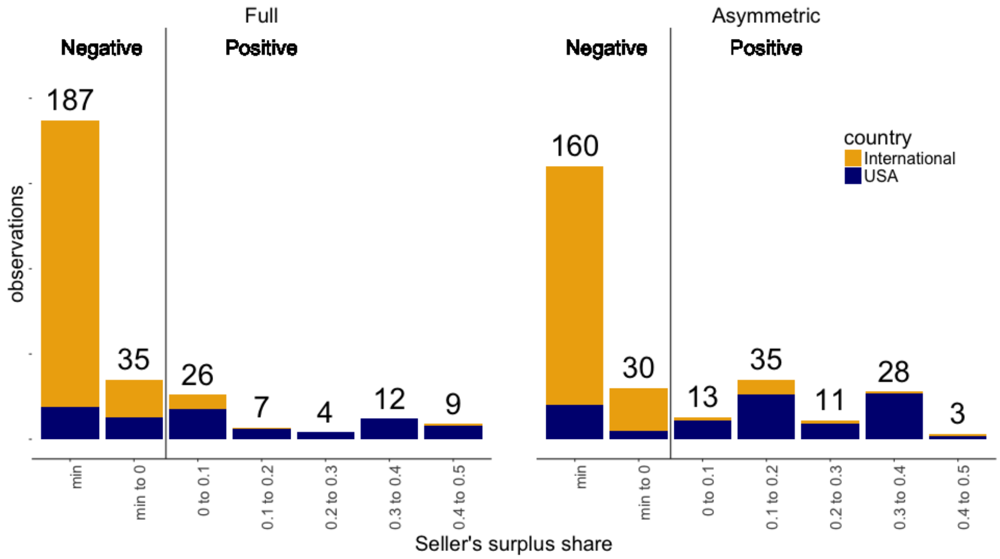


FIGURE 2.3: Pooled results for rounds with positive seller's cost and positive surplus

TABLE 2.4: Censored panel regression on PWYW prices for positive surplus values

| | Price |
|---|---|
| (Intercept) | $-4.252^{***}$ (0.332) |
| BuyerValue | $0.137^{***}$ (0.007) |
| SellerCost | $0.434^{***}$ (0.030) |
| InfoAsymmetric | $1.280^{***}$ (0.281) |
| USA | $6.322^{***}$ (0.279) |
| GenderMale | $-1.135^{***}$ (0.333) |
| Observations | 697 |
| Left Censored | 312 |
| *Note:* | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 |

USA students indeed pay significantly higher prices than international ones. Moreover, it appears that information asymmetry as the opposite effect from what we expected: when the value was unobserved, buyers on average paid higher prices.

While this appears counterintuitive, it is consistent with the results of Gneezy et al. (2012), who found that people in a PWYW restaurant paid higher prices when the payment was anonymous.

## 2.6 Summary and Further Work

In this paper we investigate the impact of the seller's costs and the buyer's valuation of the product on the buyer's choice of the price in a PWYW market. We have provided a conceptual analysis of the PWYW pricing modality by means of a model which is built on a classic inequality aversion model by adding a loss aversion parameter. We demonstrate how inequality aversion can explain a decrease in quantity demanded that has been observed by previous studies of PWYW.

Our model incorporates the possibility of an "empathized" loss aversion effect: a buyer may feel more guilt for a payment decrease by one unit if this price is below the production cost than when price is above it. Thus, even if he or she does not pay the "fair" price, he or she may still attempt to cover the seller's cost. We fit this model using a data from a past experiment.

To study the effect of buyer's valuation, we design and conduct another experiment. Curiously, we find that the buyer's valuation privacy does not appear to negatively affect the prices. More than that, buyers on average pay higher prices when their true valuation is unobserved.

# CHAPTER 3

# A BEHAVIORAL STUDY OF RETAIL ASSORTMENT PLANNING

## 3.1 Introduction

Assortment planning, that is, selecting which products to offer to the consumer, is one of the key managerial responsibilities in retail operations. In recent years, many stores keep increasing the variety of each product category in order to secure a larger market share. However, high variety comes at a cost: it increases inventory, leads to more overhead and exacerbates product cannibalization. The multitude of factors which need to be considered, along with the combinatorial nature of the problem, explains the extremely high computational complexity of the assortment optimization question (Kök et al., 2009). Given this, it may seem paradoxical that this task is often performed by managers without any support from computerized optimization algorithms. Though these managers often have plenty of sales data for each SKU, they typically lack analytical tools which would allow them to forecast sales of products which are not currently in the assortment and the effects of resulting substitution. In words of Fisher and Vaidyanathan (2012), “The tools do little more than facilitate a manual planning process that relies on the judgment of managers for key inputs.”

Because assortment planning decisions heavily rely on managers’ judgments, we believe that it is important to understand how efficient human decision makers are in solving this problems, how exactly they arrive to their decisions and ultimately, to develop procedure recommendations that will help to improve these decisions. Despite this, to our best knowledge, at the current date there have been no systematic studies of assortment planners’ behavior in academic literature.

In this chapter, we set out to obtain initial insights into this problem with a controlled laboratory experiment where human subjects assume the roles of assortment planners. Our subjects are asked to assume the role of a retailer on a competitive market and decide which products to include into their assortment. Each consumer on the market has a certain utility associated with each product and the outside option, and randomly chooses between the available options according to the Multinomial Logit (MNL) model. Products differ only by their expected utility for the customer, which are common knowledge. Each additional product included into the assortment increases the retailer's expected sales, which must be weighed against the operational costs increase.

As an analytical benchmark we use the model introduced in the seminal paper by van Ryzin and Mahajan (1999), which has a very simple and intuitive structure of the optimal assortment: it must be a subset of products with the highest expected utility values. We are interested whether the human subjects arrive to this conclusion, and if so, whether they include the optimal number of products in their assortment.

In addition to that, we attempt to investigate the effect of decision support tools on the efficacy of assortment planners. To do so, we present subjects with intermediate metrics derived from the initial problem inputs, which are supposed to simplify the task of finding the optimal assortment.

In the next section we present the assortment problem setup and its analytical properties. The research questions and hypotheses, along with the supporting literature background, are detailed in Section 3.3. In Section 3.4 we describe our experimental design and laboratory protocol. In Section 3.5 we present the results of our experiment, along with the additional conjectures arising from our observations. In Section 3.6 we summarize our findings and propose some future extensions of our work.

## 3.2 The Assortment Planning Problem

Our experiment uses a computerized simulation of a market environment, where subjects act as managers in charge of assortment selection at a retail store. The timing of the events is as follows. First, the retailer chooses the products to offer on the market from a set of possible products. Second, the potential consumers observe the assortment and either buy from the retailer or choose an outside option.

Customer demand for the product is random and generated assuming that each consumer chooses a product according to the Multinomial Logit (MNL) model. This demand mode was introduced into the assortment planning literature in the seminal paper by (van Ryzin and Mahajan, 1999) who showed that, under certain assumptions, the optimal assortment has a very simple and intuitive structure. The subjects in our experiments receive detailed explanation on the demand generating process and on how their assortment decisions affect their profit. Below, we introduce the detailed problem formulation and discuss the properties of its optimal solution.

### 3.2.1 Market Demand Model

Consider a product category with potential products $N = \{1, 2, \ldots, n\}$. A retailer must choose a subset $S \subseteq N$ as her *assortment* to sell to a market of $\lambda$ potential consumers (where $\lambda$ is a fixed exogenous parameter). When the assortment is determined, each potential consumer observes $S$ and either buys one unit of some product $i \in S$ or selects an outside option, denoted by 0, which can be interpreted as buying from a firm's competitor or not buying anything at all. The choices of potential consumers are probabilistic; they are independent and identically distributed according to the MNL model. In this model, the probability of choosing any option $j \in S$ is proportional to its *popularity index*, denoted by $v_j$ as follows. Let $q_j$ be the probability of choosing a product $j \in S \cup \{0\}$, which is given by:

$$q_j = \frac{v_j}{\sum_{i \in S} v_i + v_0}, \tag{3.1}$$

and $q_j = 0$ for $j \notin S \cup \{0\}$.

The retailer knows the popularity indices of all the products in the set $N$ she is choosing from, as well as the popularity index of the outside option $v_0$, which is independent of her decisions. Without loss of generality we order the products in $N$ from the highest popularity index to the lowest one, that is, $v_1 \geq v_2 \geq \cdots \geq v_n$.

Let $Y_j$ be the random demand for a product $j \in S$ and $Y_0$ be the random number of potential consumers who pick the outside option. Since the total number of consumers is $\lambda$, the number of potential consumers who pick an option $j \in S \cup \{0\}$ is a random variable with marginal distribution $Y_j \sim B(\lambda, q_j)$. Therefore, the expected number of consumers who pick an option $j \in S \cup \{0\}$ is $E[Y_j] = \lambda q_j$.

Adding any product to an assortment $S$ increases expected total demand for the assortment, equal to $\frac{\lambda \sum_{j \in S} v_j}{\sum_{i \in S} v_i + v_0}$, but decreases the expected individual demand for each existing product in the assortment.

### 3.2.2 Retailer's Profit Function

As in van Ryzin and Mahajan (1999) we assume that all products have identical profit margins denoted by $r$. We modify their model by assuming that the retailer does not carry inventory of the products (i.e., make-to-order setting). We consider two versions of the retailer's profit function.

We use a simple structure for the cost of offering assortment $S$: for each product the retailer includes in her assortment, she incurs a fixed operational cost $K > 0$. In this case, the retailer's expected profit from an assortment $S \subseteq N$ can be written as:

$$\mathbb{E}[\pi(S)] = r\lambda \left(1 - \frac{v_0}{\sum_{i \in S} v_i + v_0}\right) - |S| \times K, \tag{3.2}$$

It turns out that an optimal solution to this problem can always be found among assortments of the form $\{1, 2, \ldots, k\}$ for some $k = 0, 1, ..., n$, which are commonly referred to as *popular sets* (see van Ryzin and Mahajan (1999), Cachon et al. (2005) and Kök and Xu (2011)). To see this, note that the expected profit from an assortment of a fixed cardinality $k$

is increasing in the sum of the popularity indices for the products included in the assortment, and is therefore maximized by picking the $k$ products with the highest popularity indices.

Given this result, the space of solution candidates is reduced from $2^n$ to $n+1$ subsets of $N$. Moreover, comparing all the $n+1$ solutions is typically unnecessary: one can start with an empty assortment and keep adding the products in order of decreasing popularity index as long as doing so is increasing the expected profit. The proof that this greedy procedure will yield an optimal solution is given in Appendix B.1.

### 3.3 Research Question and Hypotheses

The primary goal of our study is to assess the overall performance of human decision makers at assortment optimization and to explore possible systematic deviations in their solutions. The main structural property of the model we use in our experiment is the optimality of a popular assortment, and we want to find out if the subjects would guess to search only among such sets. We expected that a substantial number of them would do so: since all products have identical profit margins, ranking them by popularity indices appears an intuitive thing to do. Some reviews of industry practice indicate that managers tend to approach the assortment optimization task in a similar way: they identify and remove products that have the lowest demand (e.g., Kök et al. (2009), Bernales et al. (2017), van Hoek and Pegels (2006)).

Given that a decision maker selects a popular set, the size of it can be optimal, too large or too small. We wanted to check if the deviation from the optimal solution would have a systematic pattern, that is, whether the subjects would consistently offer too large or too small assortments. We conjectured two possibilities. On one hand, it is possible that decision makers tend to offer larger-than-optimal assortments: many reviews of retail industry argue that the excessive variety is a common problem (see, for example, (Boatwright and Nunes, 2001)). On the other hand, decision makers may have a tendency to lean towards "average"

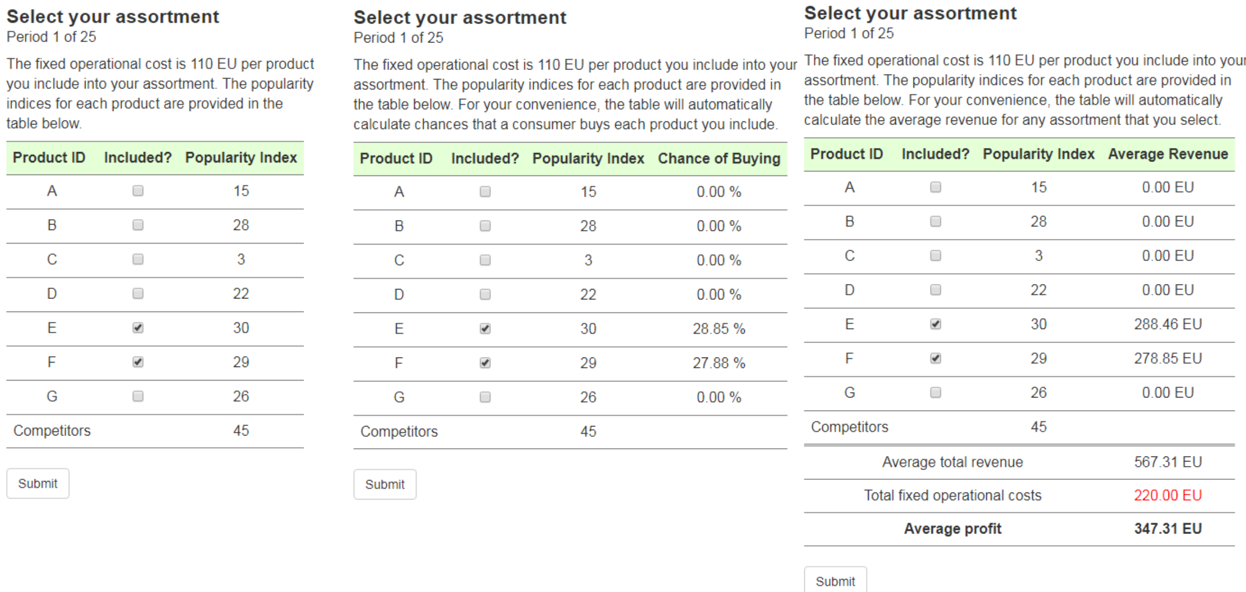

**Select your assortment**
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below.

| Product ID | Included? | Popularity Index |
|---|---|---|
| A | ☐ | 15 |
| B | ☐ | 28 |
| C | ☐ | 3 |
| D | ☐ | 22 |
| E | ☑ | 30 |
| F | ☑ | 29 |
| G | ☐ | 26 |
| Competitors | | 45 |

Submit

**Select your assortment**
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below. For your convenience, the table will automatically calculate chances that a consumer buys each product you include.

| Product ID | Included? | Popularity Index | Chance of Buying |
|---|---|---|---|
| A | ☐ | 15 | 0.00 % |
| B | ☐ | 28 | 0.00 % |
| C | ☐ | 3 | 0.00 % |
| D | ☐ | 22 | 0.00 % |
| E | ☑ | 30 | 28.85 % |
| F | ☑ | 29 | 27.88 % |
| G | ☐ | 26 | 0.00 % |
| Competitors | | 45 | |

Submit

**Select your assortment**
Period 1 of 25

The fixed operational cost is 110 EU per product you include into your assortment. The popularity indices for each product are provided in the table below. For your convenience, the table will automatically calculate the average revenue for any assortment that you select.

| Product ID | Included? | Popularity Index | Average Revenue |
|---|---|---|---|
| A | ☐ | 15 | 0.00 EU |
| B | ☐ | 28 | 0.00 EU |
| C | ☐ | 3 | 0.00 EU |
| D | ☐ | 22 | 0.00 EU |
| E | ☑ | 30 | 288.46 EU |
| F | ☑ | 29 | 278.85 EU |
| G | ☐ | 26 | 0.00 EU |
| Competitors | | 45 | |
| | Average total revenue | | 567.31 EU |
| | Total fixed operational costs | | 220.00 EU |
| | **Average profit** | | **347.31 EU** |

Submit

(A) No decision support (NS) (B) Probabilities of buying (PB) (C) Expected Profits (EP)

FIGURE 3.1: Decision stage screens for different levels of decision support information
*Note: parameters shown are for low optimal variety (LOV) condition*

solutions: that is, given a category of $N$ products, where the optimal assortment size is $n$, they would tend to choose popular assortments of the size between $n$ and $\frac{N+1}{2}$, which is the median value of possible assortment sizes. Preference for "average" solutions is a common pattern documented in behavioral operations and economics research, which manifests itself as the "compromise effect" (Simonson, 1989) and "extremeness aversion" (Simonson and Tversky, 1992) in consumer behavior and the "pull-to-center" effect in Newsvendor problem (Schweitzer and Cachon, 2000).

## 3.4 Experimental Design and Procedure

### 3.4.1 Design of Experiment

In our experiment, decision makers select from seven products by checking and unchecking a box next to a product label in a table format and hitting a "Submit" button whenever they have completed the selection (see screenshots in Figure 3.1). When the chosen assortment

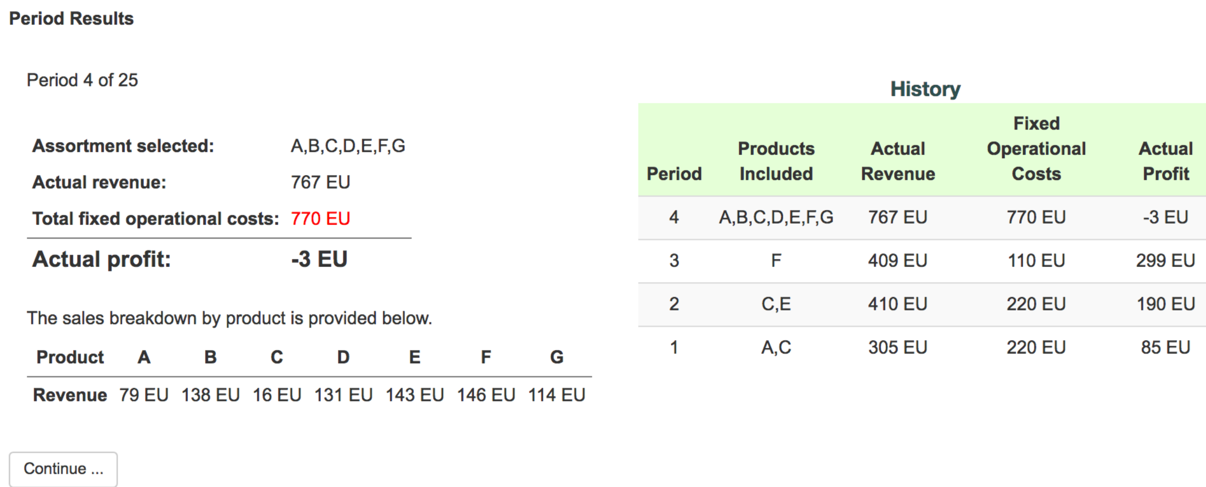


FIGURE 3.2: Results stage screen

is submitted, demand for each selected product is generated according to the MNL model described in 3.2.1.[1] After that the resulting sales for each product, the total cost and the profit appear on the screen, as shown in Figure 3.2.

This procedure is repeated for 25 independent periods, each having identical parameters. The subjects can also see the full history of their past decisions and profit realizations in a form of history table that was displayed on the right side of the screen on both the assortment selection and the results stages. An example of the history table is shown in the screenshot of the results stage (Figure 3.2).

We label the products $A, B, C, D, E, F, G$ and list them in alphabetic order with corresponding popularity indices, which are show in Table 3.1. Note that the product ordering does not correspond to increasing or decreasing popularity. The profit margin for each product is $r = 1$ EU (Experimental Unit), and the market size is $\lambda = 1,000$ potential customers. These parameters are identical across all treatments.

[1]We use the PHP `rand()` function to generate the demand after the subject submits an assortment, but we specify a random seed as a function of round number. That is, if two subjects who are assigned to the same optimal variety condition choose the same assortments in the same round, they get identical results.

TABLE 3.1: Popularity indices parameters

| Product | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| Popularity index | 15 | 28 | 3 | 22 | 30 | 29 | 26 |

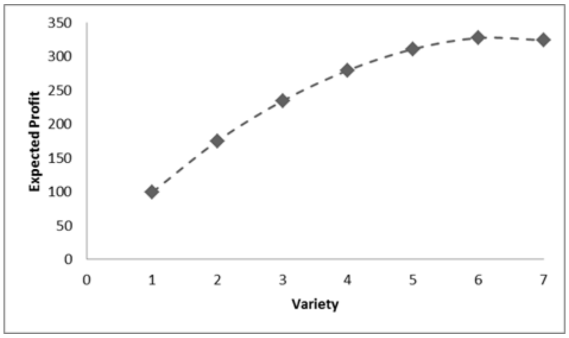


(A) High optimal variety (HOV)

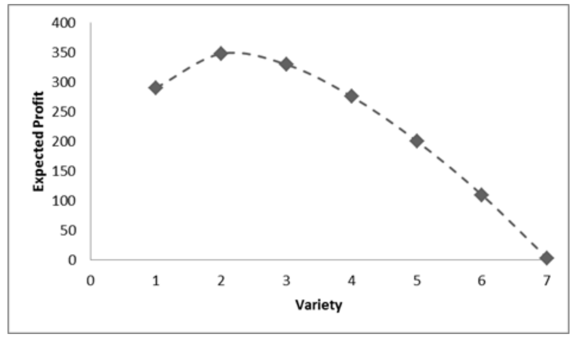


(B) Low optimal variety (LOV)

FIGURE 3.3: Expected profits for popular assortments

In this experiment we consider two factors. The first factor is *the optimal variety*, that is, the size of the optimal assortment with two levels: the *high optimal variety (HOV)* condition and the *low optimal variety (LOV)* condition. Under the *HOV* condition the optimal assortment consists of six products (i.e., products A, B,D, E,F and G) and under the *LOV* condition the optimal assortment consists of two products (i.e., products E and F). We purposely avoid extreme optimal solutions (i.e., assortments of size one or seven) so that decision makers can deviate from the optimal assortment size in either direction.

We achieve the different optimal assortment sizes by varying the attractiveness of the outside option $v_0$ and the cost parameter, that is, the fixed operational cost $K$. The exact values of the parameters are given in Table 3.2. The values of expected profit for the popular assortments of different sizes are shown in Figure 3.3.

TABLE 3.2: Problem parameters for each optimal variety condition

| | $v_0$ | $K$ |
|---|---|---|
| *HOV* | 250 | 8 EU |
| *LOV* | 45 | 110 EU |

Our second condition manipulation is meant to analyze the effect of analytical decision support tools on the performance of assortment planners. We call this factor *decision support information* and consider three levels: *no decision support information (NS)*, *probabilities of buying (PB)* and *expected profits (EP).*

The baseline condition is *no decision support information*, where decision makers can see only the problem parameters (i.e., the popularity indices values, market size, profit margin, attractiveness of outside option and cost parameter) and the demand realizations at the end of each round. Under the *purchase probabilities* condition, decision makers have all the baseline condition information, plus some intermediate calculations: as they click to select products in their assortment, the screen automatically displays the "chance of buying" (i.e., probability of choosing $q_i$) for each product in the currently selected assortment. Under the *expected profits* condition, decision makers see the baseline condition information along with the expected revenue for each product and the expected profit from the whole assortment.[2] The decision screens for all three decision support information conditions are shown in Figure 3.1.

The *expected profits (EP)* condition has what we consider the most advanced[3] decision support information: decision makers can directly observe expected profit values for all possible assortments and compare them in order to make a selection. However, doing this for all $2^7$ would likely be quite time-consuming though focusing on the seven which are "popular sets" may be doable. Note that in all three decision support information conditions the decision makers have all the necessary information to obtain these exact expected profit values themselves. To ensure they know how to do this, we make them take a pre-experimental

[2]In the experiment we use the terms "average revenue" and "average profits" which we believe resonate better with subjects.

[3]Note that the chances of buying each product which are included in the *PB* condition are not shown to subjects in the *EP* condition, so that the information provided under *PB* is not a subset of that under *EP*. Yet we argue that *EP* is the most advanced decision support system as it provides expected profit calculations, which is the metric that should be ultimately needed in order to make decisions.

quiz, where they have to calculate chance of buying, expected revenue and expected profit for a two-product exercise.

To summarize, we use a 2 (optimal variety) $\times$ 3 (decision support information) full factorial design for a total of six treatments. We use a between-subject design: each subject participates only in one experiment and is exposed only to one treatment.

### 3.4.2 Experiment Implementation

We implemented our experiment in a web application on a SoPHIE platform (Hendriks, 2012) and recruited the subjects through the Amazon Mechanical Turk (mTurk). mTurk is an online labor market that specializes on what they call human intelligence tasks (HITs), such as categorizing images or transcribing podcasts. This marketplace has gained a wide popularity as a tool for recruiting subjects for social science surveys and experiments (Schumann and Dweck (2014), Scopelliti et al. (2015) and Paolacci et al. (2015) are some recent examples). It has multiple benefits over experiments in a physical laboratory, such as large and diverse subjects pool that increases external validity of the results, and the absence of travel time that eliminates the need for a show up fee. The absence of travel time also enhances the desirable tradeoff between the effort and the earnings: a subject can choose to spend more time on the task in expectation of higher payoff, or pay minimum possible attention and get to the next available job as soon as possible. The main notable downside associated with an online experiment is the lack of control compared to what we would have in a brick-and-mortar laboratory dedicated to behavioral research. Nevertheless, it was shown that behavior of mTurk workers in classic behavioral economics and operations games is very similar to the behavior observed in a traditional behavioral laboratory (see, for example, Paolacci et al. (2010), Horton et al. (2011) and Goodman et al. (2013)).

The invitation to participate in our experiment was listed among other HITs on Amazon mTurk. After clicking on the listing, the subjects saw a page with a brief experiment description and a consent form. Before starting the experiment, the subjects were asked to download the instructions and answer three open-ended questions meant to ensure their understanding of the problem setup. The questions required the subjects to solve a two-product version of the problem: calculate purchase probabilities, costs and profits for some of the possible assortments. Subjects were not allowed to proceed until they entered the correct numbers.

The subjects' decision making procedure is described in the beginning of 3.4.1. In addition to recording the submitted assortments, we attempted to get an insight into the subjects' decision process by collecting their clicks history: the timing of checking and unchecking the boxes. After the experiment we asked the subjects to fill a short demographic questionnaire.

In the end of the experiment the total profit earned during all 25 periods was converted from experimental units to US dollars; in Experiment 1 the rates were $1,000$ EU to \$1. The resulting amount was paid as a "worker's bonus" through Amazon mTurk system.

## 3.5 Results

In total, 88 human subjects participated in Experiment 1; the breakdown per treatment is given in Table 3.3. We collected self-reported demographic data for 69 subjects. Among them 21 were female and 48 were male; 61 resided in the USA, 4 in India, 3 in Canada and 1 in Poland; 11 were between 18 and 24 years old, 37 between 25 and 34, 16 between 35 and 44, 5 between 45 and 54. We did not find significant differences in performance between demographic groups. Average earnings per subject, including \$1 fixed participation fee, were about \$9. On average, it took a subject about 30 minutes to complete the experiment.

### 3.5.1 Overall Performance Analysis

We first compare the subjects' performance in terms of profits across treatments. Since the subjects' choice is combinatorial, there are multiple possible ways to map their decisions into

Table 3.3: Number of subjects per treatment

| | Low Optimal Variety (LOV) | High Optimal Variety (HOV) | Total |
|---|---|---|---|
| No Decision Support (NS) | 12 | 12 | 24 |
| Probabilities of Buying (PB) | 21 | 15 | 36 |
| Expected profits (EP) | 12 | 16 | 28 |
| Total | 45 | 43 | 88 |

a linear performance metric. In our analysis we use two dependent variables: (1) a binary variable indicating whether the assortment which was selected in a given period is optimal or not and (2) the expected total profit corresponding to subjects' assortment, normalized by the expected profit from the optimal assortment in the corresponding treatment. For example, in the $LOV$ condition, the optimal assortment is $\{E, F\}$ which yields an expected profit of 347 EU per period. Suppose that in period 2 a subject chooses assortment $\{D, E, F\}$ which yields an expected profit of 329 EU. In this case, the metric (2) is equal to $\frac{329}{347} = 94.81\%$, even though the actual payoff received by the subject is likely different from 329 as it depends on the realized value of demand. We argue that this metric is a better indicator of the subjects' performance as it is stripped from the "luck" dimension associated with the randomly generated demand values in the experiment.

Figure 3.4 shows the average expected profits per period and the average numbers of optimal assortments per subject. We observe that for every decision support condition both performance metrics are higher in the $LOV$ treatments, where the optimal assortment consists of the two products with the highest popularity indices, compared to the $HOV$ treatments, where the optimal assortment consists of the six products with the highest popularity indices. Subjects perform best in the $LOV \times EP$ treatment, where the average expected profits in all periods is equal to 98.11%.

We use regression analysis to evaluate effects of the condition manipulations jointly across the treatments. Our primary tool for data analysis are two panel regressions with random

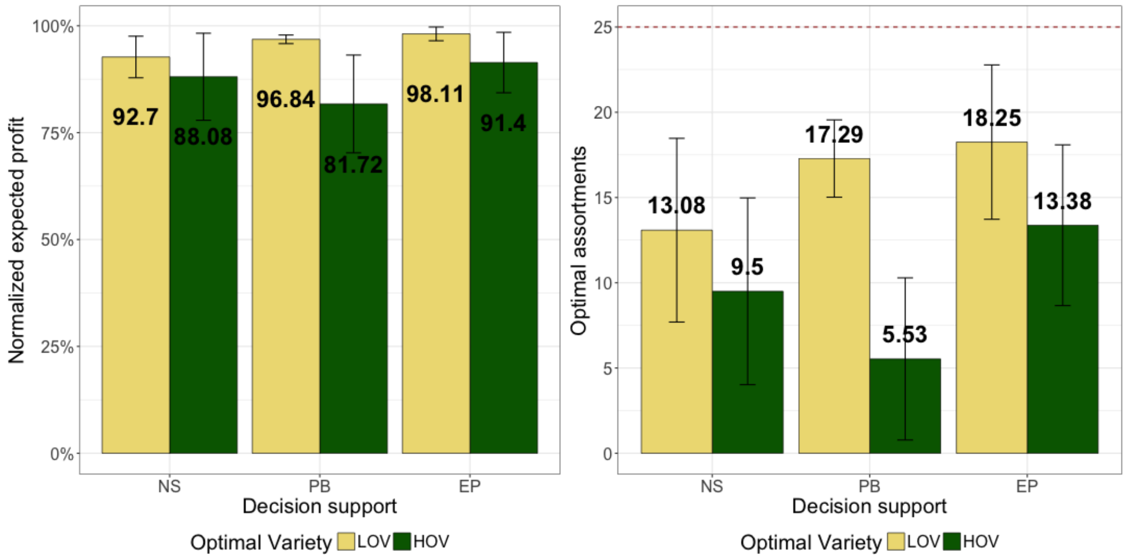


(A) Average normalized expected profits (B) Average number of optimal assortments per subject

Figure 3.4: Performance metrics summary

effects[4] The dependent variables of these two regressions are normalized expected profits and the binary variable indicating the optimality of the selected assortment. The independent variables are the period index and the factors with their interactions. The outputs of these regressions are shown in Table 3.4.[5]

We observe that the implications of performance metrics are consistent with each other: all coefficients have the same signs. The coefficients for $HOV$, $PB \times HOV$ and $EP \times HOV$ are negative for both dependent variables. All three coefficients are significant in the binary regression, and $PB \times HOV$ is also significant in the censored regression. We conjecture that this is due to the fact that reaching the optimal assortment in the $HOV$ treatments required

[4]Fixed effects model for panel data encounters incidental parameter problem, which results in a downward bias in the maximum likelihood estimate of standard errors and, consequently, impossibility of statistical inference (see Greene (2008)).

[5]Outputs of additional regressions are shown in Appendix B. Table B.3 indicates that we can disregard the factor interactions with the period index: though the coefficients $Period \times EP$ and $Period \times EP \times$HOV are significant, the plot in Figure 3.5 suggests that apparent lack of improvement in $EP \times$LOV treatment, implied by those coefficients, can be explained by censoring: since the performance is close to ideal in the very first period, there is no room for improvement over time.

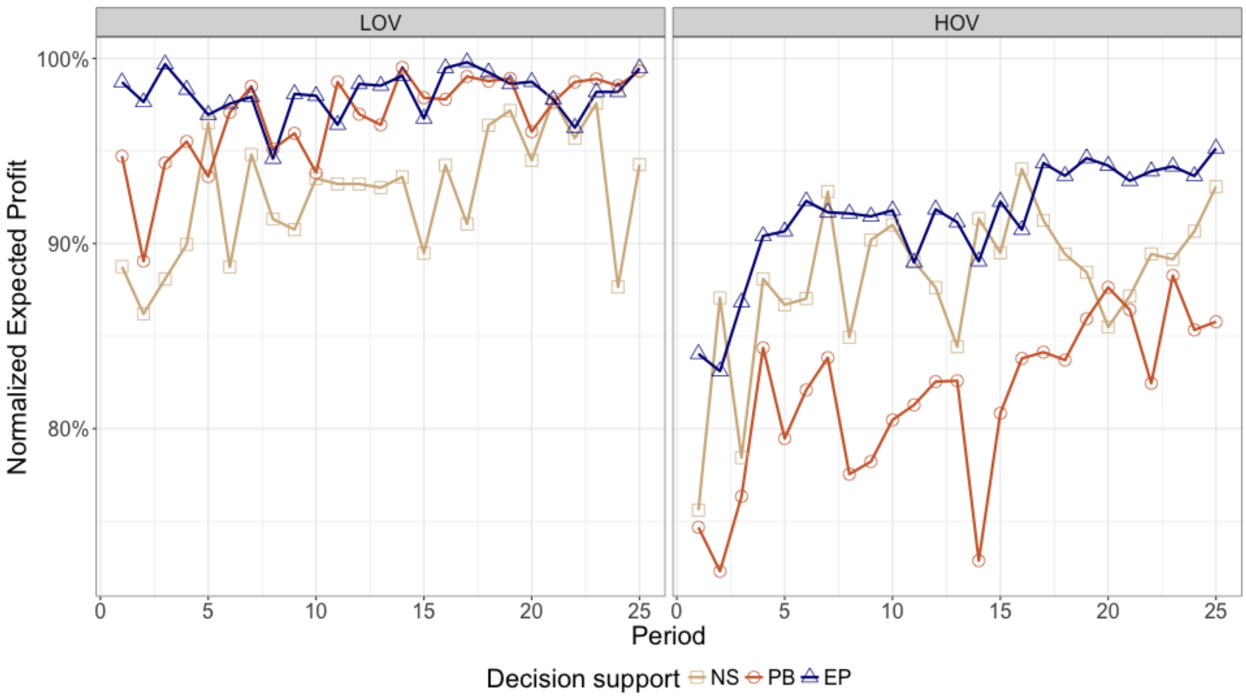


Figure 3.5: Normalized expected profits

more effort from the subjects: more products to keep in mind, more arithmetic operations in *NS* and *PB* conditions and more clicks in the *EP* condition.

### 3.5.2 "Pull-to-Center" Effect

Our data show that the subjects make intelligent decisions: their choices clearly cannot be explained by randomness as evidenced by the frequency of choosing the optimal assortments. Within every decision support condition, the subjects in choose the assortment $\{E, F\}$ more often in *LOV* condition, where it is the optimal assortment, than in *HOV* condition, where the optimal assortment is $\{A, B, D, E, F, G\}$, and vice versa (two-sided rank-sum test p-value <0.01 for each decision support condition).

We further analyze their behavior by looking for patterns of deviation from the optimal assortment. To do this, we break down their assortment choices into the following four categories: (i) the *optimal* assortment (which is a popular set) (ii) *too small*: a popular

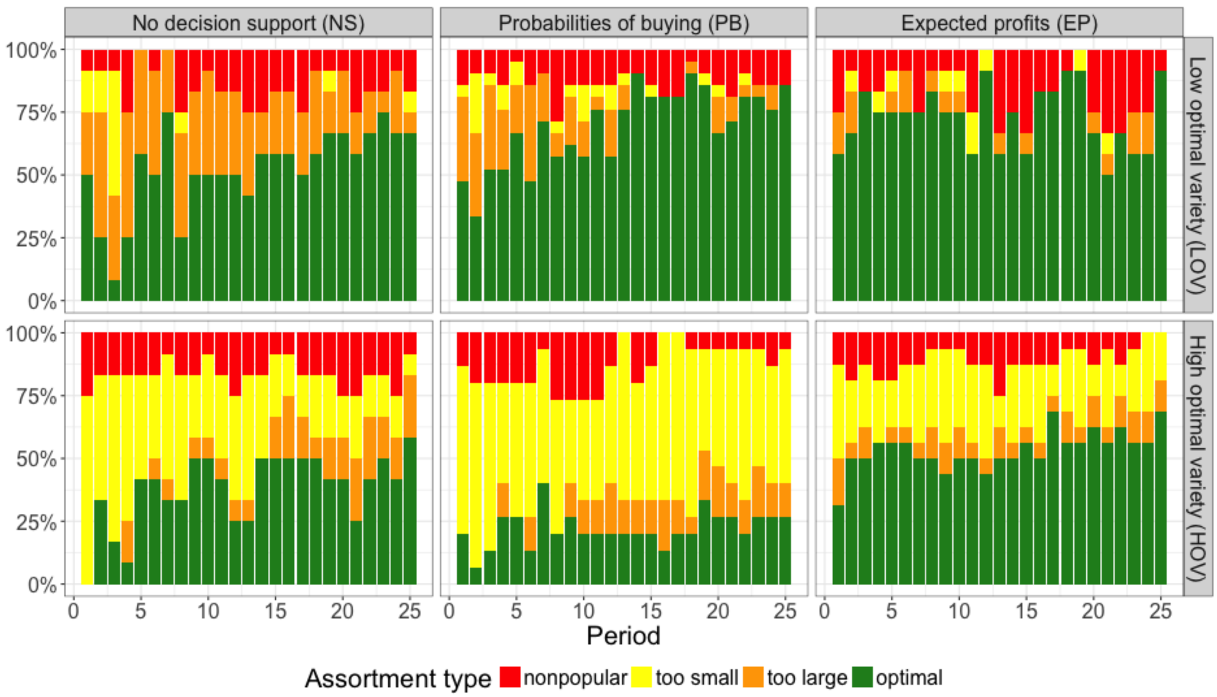


FIGURE 3.6: Assortment types over time

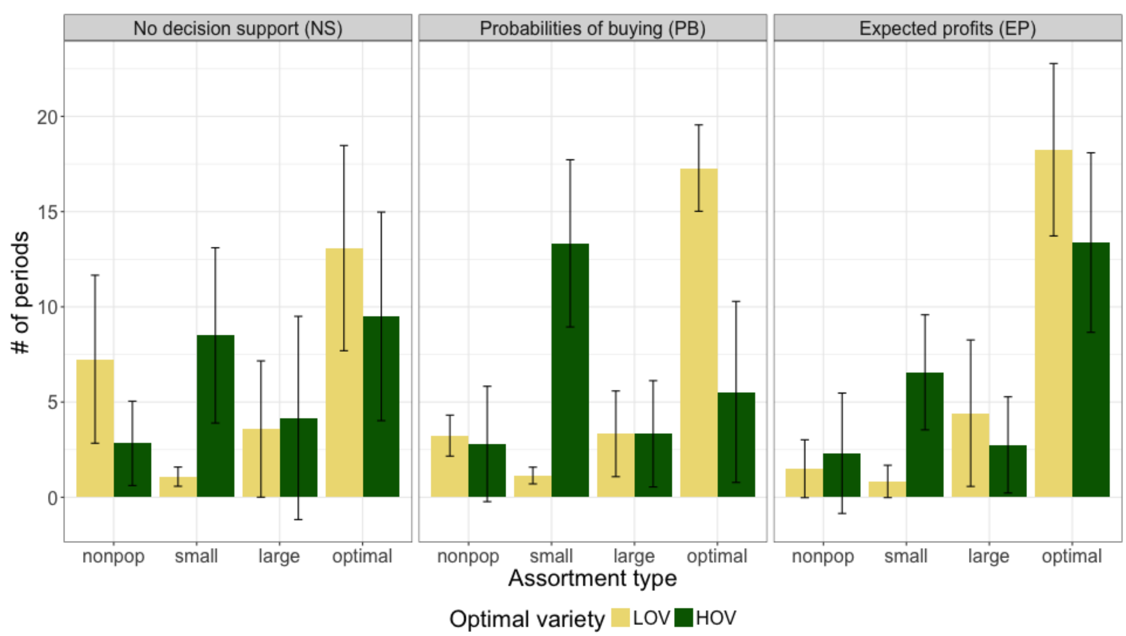


FIGURE 3.7: Average number of solutions of each type with 95% confidence intervals

Table 3.4: Panel regressions with factor interactions

| | *Dependent variable:* | |
|---|---|---|
| | Normalized expected profit | Is optimal? |
| | *censored* | *binary logistic* |
| (Intercept) | 83.776$^{***}$ (1.010) | −1.1051$^{***}$ (0.2643) |
| Period | 0.569$^{***}$ (0.027) | 0.0706$^{***}$ (0.0086) |
| PB | 12.744$^{***}$ (1.284) | 1.1186$^{***}$ (0.3250) |
| EP | 7.969$^{***}$ (1.680) | 1.8213$^{***}$ (0.3442) |
| HOV | −0.533 (1.454) | −0.8730$^{***}$ (0.3178) |
| PB×HOV | −17.291$^{***}$ (1.906) | −2.0752$^{***}$ (0.5019) |
| EP×HOV | −1.420 (2.196) | −0.8553$^{**}$ (0.4262) |
| logSigmaMu | 2.672$^{***}$ (0.022) | |
| logSigmaNu | 2.684$^{***}$ (0.008) | |
| Sigma | | 3.1381$^{***}$ (0.2150) |
| Observations | 2,200 | 2,200 |
| Right Censored | 1150 | |
| Log Likelihood | −4,903.067 | −984.6782 |
| *Note:* | | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 |

The normalized expected profit regression is estimated with the R package 'censreg' by Henningsen (2017) that implements Butler and Moffitt (1982) method. For the binary panel regression, we use the R package 'pglm' by Croissant (2017).

assortment, which is smaller than optimal (iii) *too large*: a popular assortment, which is larger than optimal (iv) a *nonpopular* assortment. The average number of chosen assortments from each category per subject is shown in Figure 3.7; a box-and-whisker plot in Figure 3.8 complements it with medians and the overall dispersion picture of the data.

From Figure 3.7 we see that in all treatments most subjects choose popular assortments, and that the average number of non-popular assortments is about the same across the *LOV* and *HOV* treatments. We also observe that for *NS* and *PB* decision support conditions the subjects in the *LOV* treatments offer excessive variety significantly more often than the subjects in the *HOV* treatment (two-sided rank-sum test p-values 0.04826 *(NS)*, 0.0151 *(PB)*, and 0.8743 *(EP)* ), and in all the decision support conditions the subjects in the *HOV*

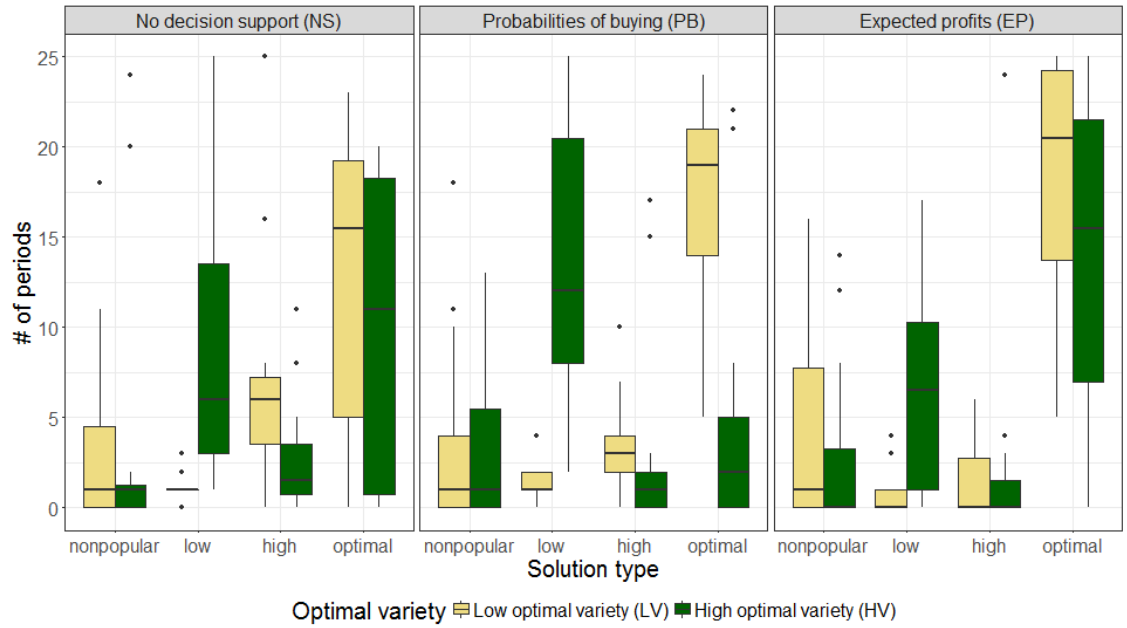


FIGURE 3.8: A box plot of number of solutions of each type per subject
*Note: whiskers are at 1.5 interquartile range*

treatments offer insufficient variety significantly more often than the subjects in the *LOV* condition (p-values 0.0002 *(NS)*, <0.0001 *(PB)* and 0.0021 *(EP)*).

Table 3.5 shows the output of regression analyses on the number of particular deviations from the optimal assortments. In the regression on the number of *too small* assortments per subject, which is shown in the second column, the coefficient for *HOV* is positive and significant, and the coefficients for interaction variables are also positive, which implies that relationship is significant for all decision support conditions. In the third column (the first regression on the number of *too large* assortments) the coefficient for *HOV* is negative and significant, but the coefficients for the interaction variables are positive. To confirm that the difference between the number of *too large* popular assortments in *HOV* and *LOV* conditions is only significant in the *NS* condition, we re-run the regression treating *EP* as a base level factor and report the resulting coefficients in the fourth column of Table 3.5.

So far we considered the data on all four categories of chosen assortments, and the results might be confounded by presence of *nonpopular* assortments, which are the third way to deviate from the *optimal* assortment. In the last column of Table 3.5 we only considered popular suboptimal assortments: the dependent variable is the proportion of *too large* assortments among the subject's *too large* and *too small* assortments, which takes values from 0 to 1. Note that this regression includes only 81 subjects out of 88, because the remaining seven did not choose any popular suboptimal assortments. In this regression, the base factor levels are *EP* and *LOV*. This analysis confirms that the effect of *HOV* condition is negative and significant in all three decision support conditions: *HOV* is negative and significant, and the coefficients for the interaction variables are also negative.

TABLE 3.5: Censored regressions on the number of *too small* and *too large* popular assortments per subject

| | *Dependent variable:* | | | |
|---|---|---|---|---|
| | too small | too large | | too large (conditional) |
| | (1) | (2) | (3) | (4) |
| (Intercept) | 0.371 (1.644) | 7.145*** (1.801) | −2.454 (2.188) | 0.491*** (0.175) |
| NS | | | 9.599*** (2.832) | 0.285 (0.213) |
| PB | −0.268 (2.071) | −4.290* (2.260) | 5.309** (2.571) | 0.254 (0.199) |
| EP | −2.657 (2.497) | −9.599*** (2.832) | | |
| HOV | 8.260*** (2.294) | −5.397** (2.588) | 1.874 (2.761) | −0.457** (0.212) |
| NS×HOV | | | −7.271* (3.788) | −0.136 (0.273) |
| PB×HOV | 5.104* (2.982) | 3.142 (3.403) | −4.129 (3.551) | −0.282 (0.258) |
| EP×HOV | 0.242 (3.273) | 7.271* (3.788) | | |
| logSigma | 1.709*** (0.087) | 1.820*** (0.099) | 1.820*** (0.099) | −0.892*** (0.112) |
| Observations | 88 | 88 | 88 | 81 |
| Left-censored | 16 | 30 | 30 | 23 |
| Right-censored | 2 | 1 | 1 | 9 |
| Log Likelihood | −232.946 | −206.903 | −206.903 | −54.696 |

*Note:* *p<0.1; **p<0.05; ***p<0.01

*In (3) and (4) EP is treated as a base factor level for sake of coefficients interpretation*

Estimated with the R package 'censreg' by Henningsen (2017) that implements Butler and Moffitt (1982) method.

### 3.5.3 Effect of Decision Support

Next we study the impact of the decision support on the subjects' performance. As shown in Figure 3.4, on average, subjects choose the optimal assortments more often and achieve higher payoffs in the *EP* condition compared to the *NS* condition. This observation is supported by regression analysis in Table 3.4. This is as expected, because in the *EP* condition, subjects are provided with expected profit calculations for each assortment which they selected so they only have to compare the expected profits values across possible assortments in order to make a choice. Though there are $2^7 = 128$ possible assortments, only seven of them are popular ones, so provided they realize the optimality of popular assortments, they will be comparing no more than seven numbers.

However, the effect of *PB* decision support condition is ambiguous: under $LOV$, subjects on average perform better under *PB* than *NS*, but under $HOV$ the opposite holds. Regression analysis supports this observation: Table 3.4 indicates that the effect of *PB* treatment is significant in both optimal variety conditions, but the direction of this effect is different. This is unexpected, because the subjects in *PB* condition have *more information* than the subjects in the *NS* condition; in particular, they are given purchase probabilities (termed "chances of buying") for each product in each assortment they consider. The apparent adverse effect in the *PB* condition implies that the additional information somehow confuses and/or misguide the subjects. We currently consider three possible underlying mechanisms.

1. Additional information reduces the subjects' motivation to think through the problem.

2. The display of purchase probabilities makes the cannibalization effect (i.e., the fact that adding a product "steals" demand from previously included products) more salient, which causes the subjects to assign excessive weight to it.

3. The display of purchase probabilities makes subjects realize that the problem is nontrivial, and they are less likely to go with an intuitive solution, which would be a correct one.

**Reason #1** means that in *PB* condition the subjects put less effort into solving the problem than in *NS* condition, and in *HOV* condition this effect outweighs the positive effect of additional information. In *LOV* condition, on the other hand, this adverse effect is outweighed, because the additional information is easier to interpret and therefore more valuable: to get the expected revenue for the optimal assortment one has to sum two displayed probabilities of buying and multiply the resulting number by the market size, which was $1,000$. If this conjecture is correct, the subjects probably would spend more time on their decision in the *NS* condition than in *PB*, at least in the first round.

**Reason #2** means that that by the direct observation of purchase probabilities, subjects realize how each additional product decreases the probabilities of buying for each product that is already in assortment. This observation somehow concerns them and as a result, they add less products than they would had they not been made aware of this phenomenon. In the *LOV* condition this effect would counterbalance the "pull-to-center" effect, but in *HOV* it may exacerbate the bias.

**Reason #3** means that the subjects notice that the products with smaller popularity indices have smaller cannibalization effect, and thus may choose nonpopular assortments. If this is the case, *PB* may improve the subjects performance in the case when the optimal assortment is indeed nontrivial.

### 3.5.4 Analysis of the Clickstream Data

In this section we analyze the subjects' clicking behavior in order to get some insights into their decision making process. In addition to the choices submitted by the subjects, we have

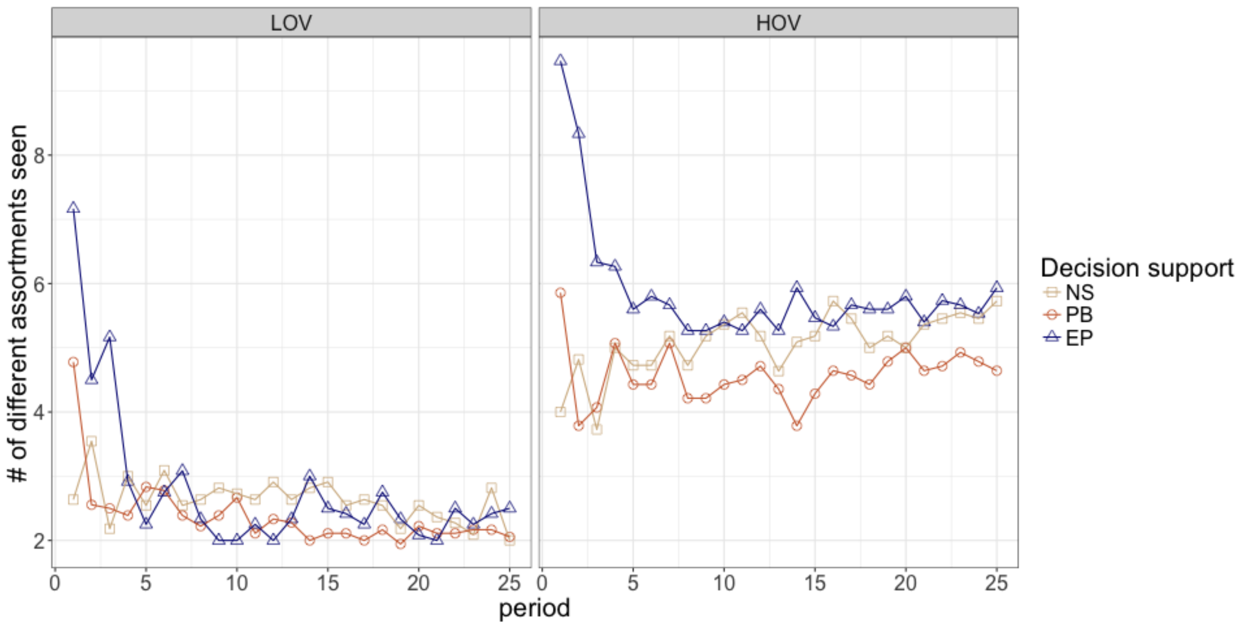


Figure 3.9: Average number of different assortments seen per period

collected their clickstream data: the boxes they were checking and unchecking before hitting the "Submit" button and the time stamps of those clicks.[6]

We say that an assortment is *seen* by a subject if at some point during a period the corresponding products were checked simultaneously and that an assortment is *chosen* by a subject if it was submitted in a given period. Figure 3.9 shows averages for the number of different assortments seen by the subjects in each period. Table 3.6 gives mean and median statistics on the different assortments seen and chosen by each subject over the whole experiment. Note that the subjects in the *HOV* condition on average see more different assortments than the subjects in the *LOV* condition, but this does not necessarily imply that they are more inclined to explore. Recall that in the *PB* and *EP* conditions decision support

[6]We use JavaScript to collect the clickstream data, which may sometimes give inaccurate results due to technical issues with the user's browser that we cannot control. Therefore, we verified the consistency of the collected clicks data with the final decision each subject submitted to the server and discarded the observations that were not consistent with the submitted selections. We discarded a total of 22 out of 2200 observations for 7 out of 88 subjects. When reporting period averages, we include all subjects; when reporting cumulative values (e.g., the number of different assortments seen in the whole experiment), we include only 81 subjects who have complete entries for all 25 periods.

Table 3.6: Number of different assortments seen and chosen by subjects over the entire experiment, mean and median

| | Seen | Chosen | Number of subjects |
|---|---|---|---|
| HOV | 28.10 (26) | 6.05 (4.5) | 40 |
| NS×HOV | 28.91 (29) | 7.09 (5) | 11 |
| PB×HOV | 22.64 (22)** | 6.43 (5) | 14 |
| | 0.0372 | 1 | |
| EP×HOV | 32.6 (31) | 5 (4) | 15 |
| | 0.7551 | 0.2942 | |
| LOV | 13.66 (12) | 5.34 (5) | 41 |
| NS×LOV | 13.91 (12) | 6.73 (5) | 11 |
| PB×LOV | 11.72 (10) | 5.39 (4.5) | 18 |
| | 0.7348 | 0.2901 | |
| EP×LOV | 16.33 (16) | 4 (3)* | 12 |
| | 0.1465 | 0.06299 | |
| Grand Total | 20.79 (19) | 5.70 (5) | 81 |

*Note:* *p<0.1; **p<0.05; ***p<0.01

*Median values are given in brackets, p-values at the bottom (Mann-Whitney U test for equality in ranks with NS condition)*

information updates every time a subject is checking or unchecking a box corresponding to some product. Therefore, a subject who selects an assortment of six products (which was the size of expected profit maximizing assortment in the *HOV* condition), will see at least six different assortments in that period.

To make judgements about the subjects' exploration behavior we compare the number of assortments seen in *EP* and *PB* conditions to the ones seen in the *NS*, where subjects did not get any information from clicking. In Table 3.6 we see that subjects in the *EP* condition see more different assortments but end up submitting fewer of them than in the *NS* condition, which suggests that they indeed take advantage of the forms of decision support that we

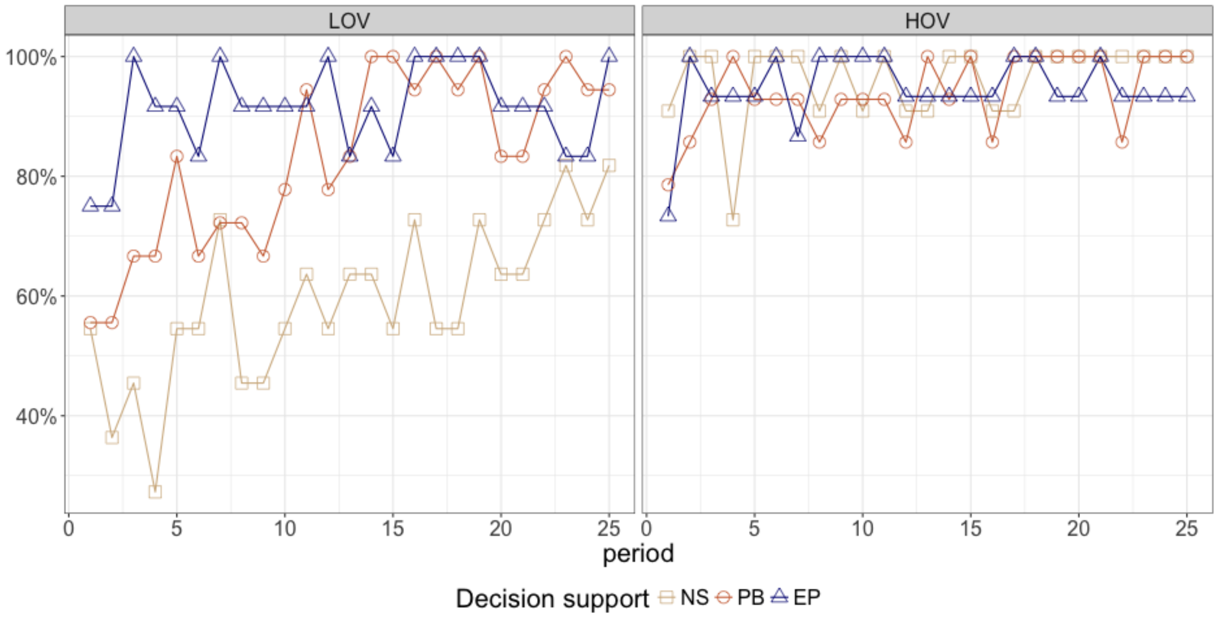


FIGURE 3.10: Percentage of subjects who picked the best assortment among the ones seen

provide. Mann-Whitney $U$ tests, however, do not indicate significance, except for one case (chosen assortments within $LOV$ condition). Curiously, it appears that subjects in $PB$ condition see significantly fewer assortments than subjects in the $NS$ condition.Combined with the subjects' poor performance in the $PB \times HOV$ treatment, this observation further supports our conjecture that in the $HOV$ condition $PB$ information hinders the subjects rather than helps them.

One of the potential concerns in our experimental design could be that the subjects did not trust the decision support information: they could suspect some kind of deception[7] or they could think that choosing the assortment with the highest expected profit is not the optimal strategy due to risk aversion. Another possible explanation for not choosing the assortment with the highest expected profit among the seen ones could be the complexity:

[7]Our experiment does not involve any deception, and the experiment consent form clearly says so (see Appendix B.3). Use of deception is strongly discouraged in the scholarly field of behavioral operation management (Katok, 2011); however, Amazon mTurk workers are likely to have previous exposure to studies in other fields, for example, psychology, where ethical guidelines allow deception under certain conditions (American Psychological Association, 2017).

Table 3.7: Assortments seen and choices in the first period

| | % chose max seen | % saw optimal | % chose optimal |
|---|---|---|---|
| NS | 75.00% | 41.67% | 25.00% |
| HOV× NS | 91.67% | 0.00% | 0.00% |
| LOV×NS | 58.33% | 83.33% | 50.00% |
| PB | 65.71% | 57.14% | 37.14% |
| HOV×PB | 78.57% | 28.57% | 21.4% |
| LOV×PB | 57.14 % | 76.17% | 47.62% |
| EP | 75.00% | 50.00% | 42.85% |
| HOV ×EP | 75.00% | 31.25% | 31.25% |
| LOV×EP | 75.00% | 75.00% | 58.33% |
| Grand Total | 71.26% | 50.57% | 35.22% |

subjects might have problems remembering the expected profits for all the assortments they have seen. Finally, it is possible that even after successfully passing the instructions quiz some subjects had problems understanding the experiment, which could be either grasping the concept of "Average profit" (recall that this is how we refer to expected profit in the instructions) or getting used to the experiment interface.

To investigate this, we take a look at the subjects' choice conditional on them seeing the optimal solution while checking and unchecking the boxes in a given period. Table 3.7 provides the summary statistics for the first period, where we see that not all the subjects in the *EP* condition chose the assortment with the maximum expected profit value among the those they have seen: only 75% of them did so. However, as we see in Figure 3.10, in later periods almost all subjects in the *EP* condition chose the assortment with the maximum expected profit that they see in that period. Therefore, it seems that even if the subjects indeed have reservations trusting the decision support information, their wariness diminish over time.

We now look at the timing of subjects' clicks. We refer to the difference between the time when a subject checks a box corresponding to some product for the first time during

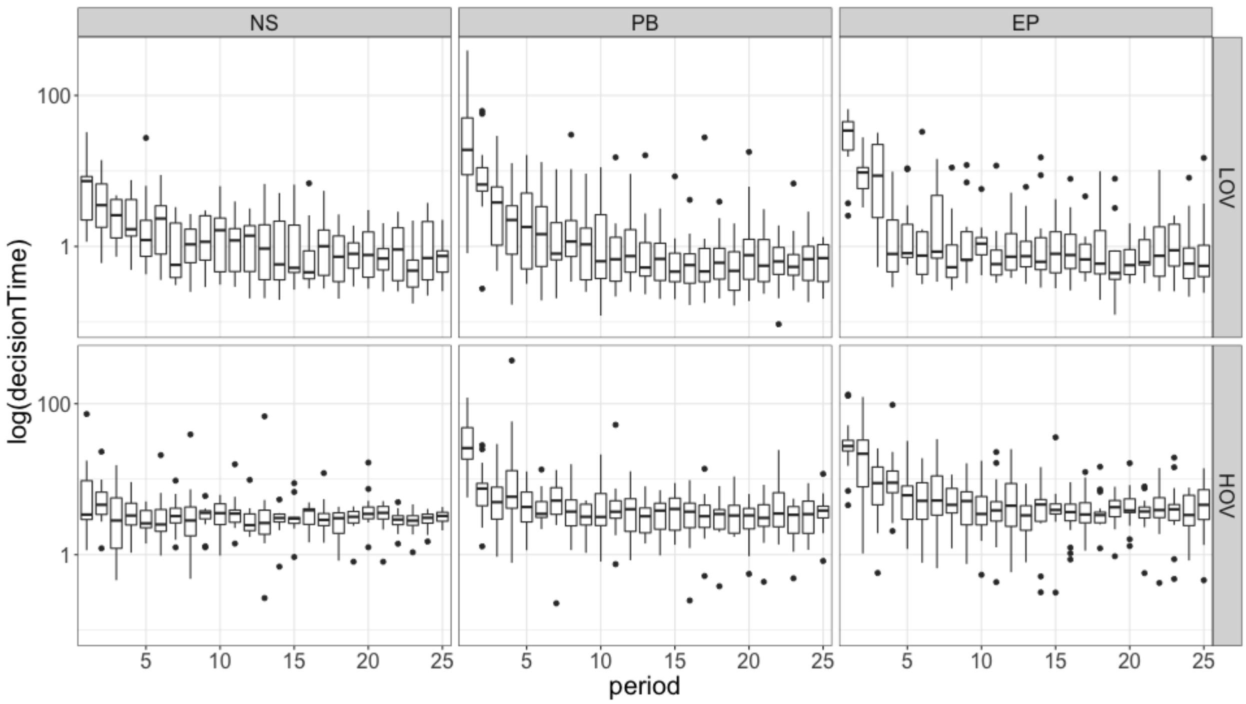


FIGURE 3.11: Decision time (in seconds – logarithm scale)

the period and the time she checks or unchecks a box right before hitting the "Submit" button as her *decision time*, which we measure in seconds.[8] Figure 3.11 shows the box plots of decision time logarithm in each time period. Note that the data is highly variable: even after applying a logarithm transform we observe a considerable amount of outliers.[9] Both Figure 3.11 and Figure 3.9 suggest that the subjects in the *EP* and *PB* conditions do most of their exploration in the first few periods of the experiment. In Table 3.8 we report average and median decision times for the first, the fifth and the last period of the experiment for each treatment. We compare the decision times in the *EP* and *PB* conditions to the decision times in *NS* condition by means of the rank test. We see that in the first period the decision

[8]We only have the timestamps of checkbox clicks, so we cannot say how long a subject spends contemplating the problem before she started clicking or how long a subject hesitates before clicking the "Submit" button. Nevertheless, checkbox timestamps allow us to judge how long the subjects are engaged with the decision support that we provided.

[9]For unscaled view of this plot, see Appendix B.5, Figure B.3.

Table 3.8: Decision time, mean and median

| | Period 1 | Period 5 | Period 25 |
|---|---|---|---|
| HOV | 29.44 (21.39) | 5.34 (3.89) | 4.18 (3.63) |
| NS×HOV | 11.28 (3.41) | 2.93 (2.60) | 3.18 (3.26) |
| PB×HOV | 33.99 (25.77)$^{**}$ | 4.61 (4.23) | 3.63 (3.74) |
| | 0.0157 | 0.2363 | 0.5418 |
| EP×HOV | 32.51 (27.44)$^{***}$ | 2.59 (6.12)$^{**}$ | 2.03 (4.58)$^{*}$ |
| | 0.0003 | 0.02569 | 0.0661 |
| LOV | 35.44(15.43) | 3.38(0.94) | 1.08(0.63) |
| NS×LOV | 7.46 (5.37) | 3.81 (1.25) | 0.76 (0.74) |
| PB×LOV | 53.10 (19.04)$^{***}$ | 3.58 (1.48) | 0.72 (0.70) |
| | 0.0018 | 0.6943 | 0.8683 |
| EP×LOV | 38.77 (34.29)$^{***}$ | 7.82 (0.78) | 5.43 (0.55) |
| | 0.0009 | 0.4776 | 0.9327 |
| Grand Total | 32.50 (18.99) | 4.33 (2.55) | 2.59 (1.30) |

*Note:* $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

*Median values are given in brackets, p-values at the bottom (Mann-Whitney U test for equality in ranks with NS condition)*

times in the *EP* and *PB* conditions is significantly higher than in *NS* condition, but in the fifth and the last periods the difference gets smaller, and might even disappear in some treatments.

In contrast with our observations for the number of different assortments seen by subjects, decision time values suggest that subjects in *PB* indeed spend considerable time on the problem, at least during the first period. This evidence goes against Reason #1 ("less effort") we suggested in the previous section as an explanation for subjects' poor behavior in $PB \times HOV$ treatment.

Finally, we take a look at the relation between the clicking behavior and the expected profit outcomes. We want to see if the subjects who interacted with the decision support system more (in the *PB* and *EP* conditions) or those who spent more time at the problem performed better. As a proxy for the subjects' exploration behavior we use the number of different assortments seen in the first period, when all the subjects are seeing the problem for the first time, and there is no knowledge carried over from the previous periods. Figure 3.12 shows the normalized total expected profit per subject plotted against the number of different assortments they have seen in the first period, and Figure 3.13 shows the same profits plotted against the logarithm of decision time.

From the plots it appears that higher number of assortments seen in the first period is associated with higher expected profit not only in the *PB* and *EP* conditions, but also in *NS* condition, where subjects did not receive any information from clicking. However, as we discussed, subjects may be seeing more different assortments simply because they are selecting larger assortments. In particular, in $HOV \times NS$ treatment a small number of assortments seen in the first period may indicate a solution which does not maximize the expected profit, which was carried into the later periods. Therefore, to distinguish between assortment size effects and exploration behavior, we think of the *NS* condition as a "baseline".

With this in mind, we run a censored regression with total normalized expected profit as a dependent variable, where we interact each of the variables # *ast seen* (number of different assortments seen in the first period) and $\log(time)$ (logarithm of decision time) with the treatment variables. Table 3.9 provides the regression output.

The obtained coefficients suggest that higher number of assortments seen is associated with lower overall performance in $LOV \times NS$ treatment. Similar to the positive effect in $HOV \times NS$, this effect is probably due to the fact that in *LOV* condition, where the EP-maximizing assortment can be achieved in two clicks, more clicks indicate a solution that yields smaller expected profit. This negative effect seems to disappear in $LOV \times PB$ and

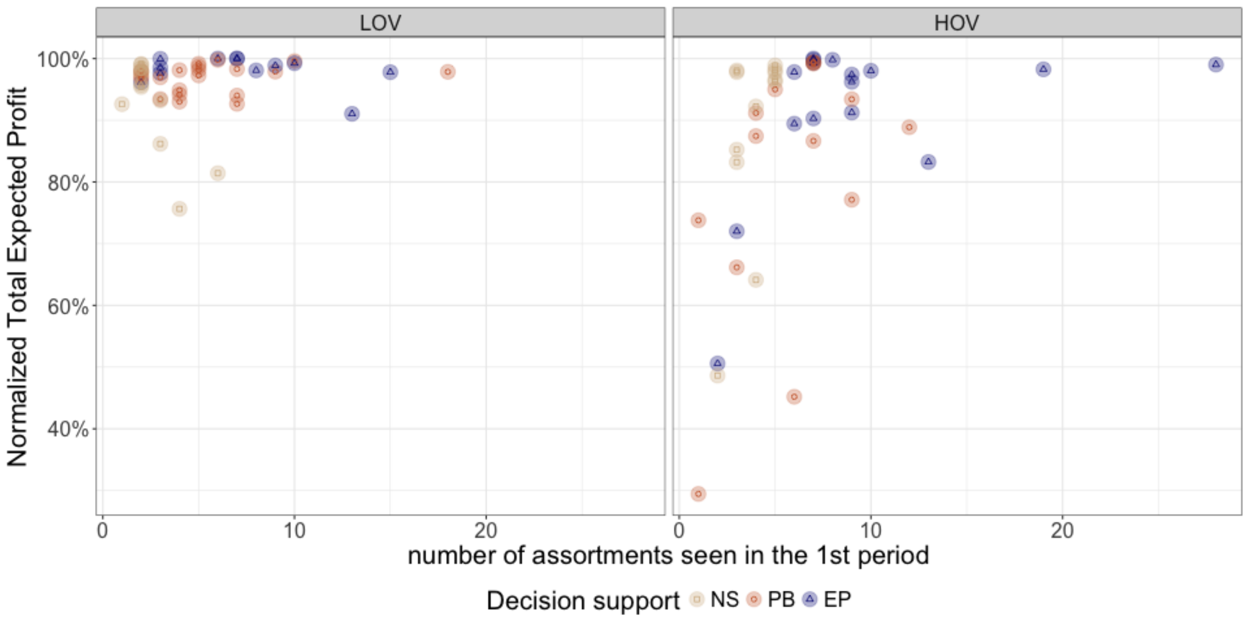


FIGURE 3.12: Exploration in the first period and normalized expected profit : assortments seen

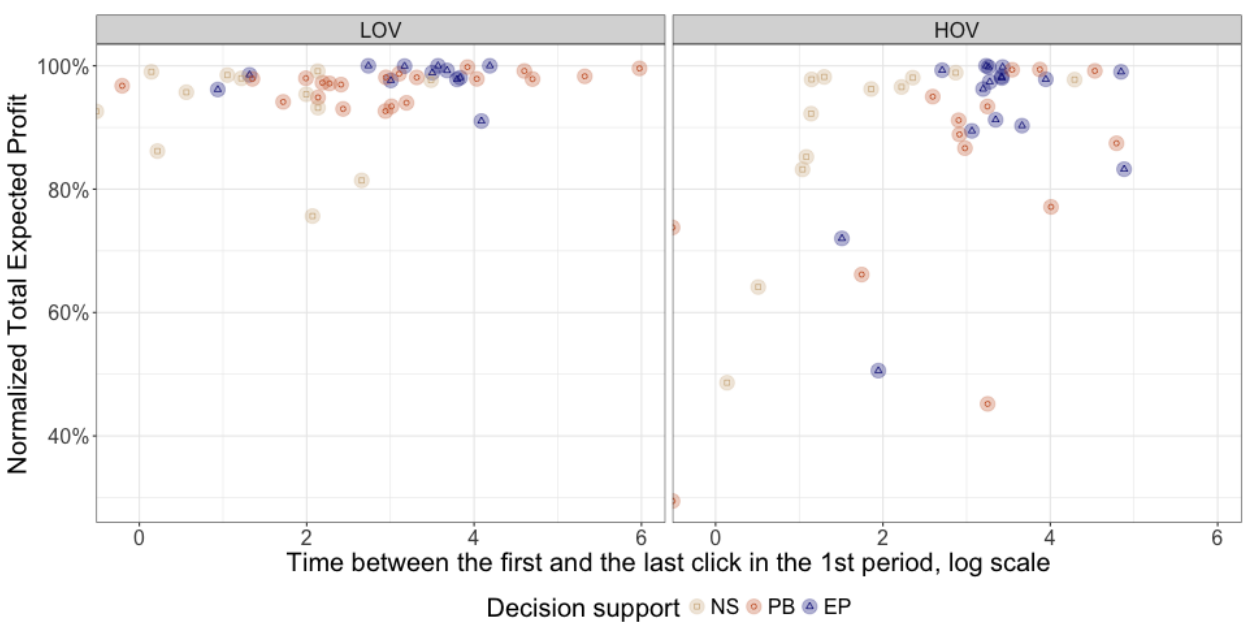


FIGURE 3.13: Exploration in the first period and normalized expected profit in the experiment: decision time

TABLE 3.9: First period exploration and performance in the whole experiment: censored regression

| | *Total Normalized Expected Profit:* |
|---|---|
| (Intercept) | $1.060^{***}$ (0.064) |
| HOV | $-0.468^{***}$ (0.126) |
| PB | $-0.110$ (0.081) |
| EP | $-0.096$ (0.113) |
| log(time) | 0.011 (0.012) |
| # ast seen | $-0.055^{**}$ (0.024) |
| PB×HOV | 0.214 (0.149) |
| EP×HOV | 0.163 (0.189) |
| log(time)×HOV | $0.056^{*}$ (0.034) |
| PB×log(time) | $-0.003$ (0.023) |
| EP×log(time) | 0.019 (0.042) |
| PB× # ast seen | $0.054^{**}$ (0.025) |
| EP× # ast seen | $0.047^{*}$ (0.026) |
| # ast seen×HOV | $0.100^{**}$ (0.041) |
| PB×log(time)×HOV | $-0.036$ (0.040) |
| EP×log(time)×HOV | $-0.013$ (0.066) |
| PB×# ast seen×HOV | $-0.089^{**}$ (0.043) |
| EP×# ast seen×HOV | $-0.091^{**}$ (0.042) |
| logSigma | $-2.370^{***}$ (0.078) |
| Observations | 87 |
| Right-censored | 4 |
| Log Likelihood | 76.013 |
| *Note:* | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 |

$LOV \times EP$ conditions, which may indicate that exploration is associated with better performance. Time spent making the decision, on the other hand, is associated with better performance in both optimal variety conditions within $NS$ condition (both log($time$) and log($time$) $\times$ $HOV$ variables have positive coefficients).

Since we do not control the subject exploration behavior, we cannot draw any conclusions about the causality between the dependent and independent variables. There are multiple possible confounding factors: aside from the ones we already discussed, subject's interest in mathematics can affect both her exploration patterns and the profit outcome. However, based on these observations, we believe that investigating the possible causal link between the amount of exploration and the overall subject performance is a promising research direction.

### 3.6 Summary and Conclusions

The data from Experiment 1 shows that subjects who are facing the combinatorial problem of assortment planning make intelligent decisions. Most of them choose popular assortments, as the theory would suggest. Moreover, in all treatments with low optimal variety ($LOV$ condition), the subjects choose the optimal assortment more than 50% of the time, including the treatment when they have no assistance in a form of decision support information ($LOV \times NS$ treatment).

Subjects perform significantly worse in the $HOV$ condition, when the optimal assortment consists of six products. Since in practice assortment planners have to deal with assortments of dozens, or even hundreds of products, the profits losses due to the flaws in the human assortment planners decision making might be much higher than our experiment seems to suggest. However, from the theoretical point of view our results are very encouraging: since human subjects decisions are overall consistent with implications of an analytical assortment planning model, even on a relatively small scale, the prospectives of further investigation of their decision making patterns are open, and our results can serve as a foundation for future research directions.

We find the support for the "pull-to-center" effect that we hypothesized: in the $LOV$ condition subjects offer excessive variety more often than the subjects in $HOV$ condition, and vice versa. In particular, in $HOV$ treatment very few subjects ever choose to include all seven products, though this assortment has the second-highest expected profit, of all possible assortments and an expected profit value which is less than 5% smaller than that of the optimal assortment.

There are multiple possible underlying drivers of this effect. First, pulling to center may be driven by random errors: given that an assortment is popular, and that the optimal assortment includes two or six out of seven products, there are five ways to err on one side of the optimal one, and only one way to err on the other side. Though the subjects in our experiment clearly did not make choices in arbitrary fashion, it is possible that some of them used a mental heuristic similar to "consideration set" heuristic, which is frequently used as a consumer choice model in the marketing literature to explain why consumers of packaged goods tend to rotate between some small portfolio of brands in a seemingly arbitrary fashion (see, for example Roberts and Lattin (1997) and Ratchford (2009)). According to this theory, a person who is choosing from multiple alternatives, does so in two stages: first, she forms a small portfolio of candidates, and then makes a choice among these candidates. In our case this choice model would imply that a decision maker forms a "consideration set" based on some hunch (e.g., "popular assortments larger than four") and them picks randomly among them.

Another possible mechanism is a so-called "compromise effect": people in general prefer "middle-ground" solutions. In the $HOV$ case, the tendency to choose lower-than-optimal assortments can be exacerbated by risk aversion: since the fixed operational costs are certain, and the revenue is variable, larger assortments are associated with higher risk.

A somewhat surprising result of Experiment 1 is the direction of the decision support effect: in the $HOV$ condition providing subjects with probabilities of buying ($PB$) results in a worse performance, compared to the treatment where no decision support was provided ($NS$). The possible explanation for this effect and their implications are discussed in 3.5.3).

# CHAPTER 4

# A MODEL FOR PERSONALIZED ASSORTMENT PLANNING IN PRESENCE OF CONSUMER LEARNING

## 4.1 Introduction

In this chapter, we model the behavior of a consumer who learns about her idiosyncratic preferences by repeatedly purchasing from a certain product category during a finite selling season. Each time the consumer picks a product that maximizes her current expected utility. After each purchase, the consumer gets a noisy signal about the product's utility and updates her beliefs in a Bayesian fashion. If the expected utility of a product becomes lower than the utility of the outside option, the consumer stops buying.

The firm chooses an assortment before the beginning of the selling selling season and cannot change it later. The firm does not incur any additional operational costs for offering a larger assortment. However, different products may have different profit margins, so the firm faces the problem of product cannibalization: it may be profitable to exclude products with a low profit margin from the assortment so that it does not cannibalize the demand from the one with a higher profit margin. On the other hand, offering more products in the assortment decreases the probability that the consumer deems all the products unsuitable and switches to the outside option.

To capture this tradeoff, we compare the firm's profit from offering one versus two products in the assortment. First, we study one-product problem, where the consumer repeatedly purchases the same product as long as her expected utility is above the threshold determined by the outside option. We show that under certain assumptions on the utility signal distribution the consumer's beliefs evolution is characterized by a one-dimensional random walk and can be linked to a certain variation of a stochastic process known as a gambler's ruin problem (Ross, 1995, p. 185). We then obtain the expression for the firm's expected profit for this scenario.

Second, we consider the problem where the firm offers an assortment of two products in each period. We characterize the consumer purchasing behavior as a sequence of switchings between the products offered by the firm, that can eventually be terminated by the consumer's switching to the outside option. We show that the sequence of the consumer's switching times follows a certain alternating pattern.

## 4.2 Consumer Learning Model

### 4.2.1 Problem Formulation

We consider the setting where the consumer repeatedly buys a product, updating her beliefs about it after each purchase. Each time the consumer buys the product, she receives a random utility from her experience with the product. Let $\{U_1, U_2, \dots\}$ denote the sequence of utility values from these experiences. The $t$'th experience can be either satisfactory, which generates the utility of 1 (i.e., $U_t = 1$) or unsatisfactory, which generates the utility of zero (i.e., $U_t = 0$), for $t = 1, 2, \dots$. There are two possible states of the product: either *good* ($G$) or *bad* ($B$). Conditional on the state of the product $S \in \{G, B\}$, $\{U_t\}$ is a Bernoulli process (i.e., the variables $U_1, U_2, \dots$ are i.i.d.). A good product gives a satisfactory experience with the probability $q_G$, a bad product gives a satisfactory experience with the probability $q_B$, where $q_G > q_B$. The consumer knows the values $q_G$ and $q_B$, but she is uncertain about the state of the product $S$, and her prior belief that the product is *good* is $p$, that is, $P(S = G) = p$ and $P(S = B) = 1 - p$.

In each period the consumer either purchases a product from the firm or picks the outside option that gives a certain utility $\underline{u}$. We assume that the consumer is myopic; in each period she picks the option with the highest expected utility.

Each time after the consumer buys the product and experiences it, she updates her beliefs in a Bayesian fashion. The consumer's probability that the first experience will be satisfactory (i.e., $U_1 = 1$) is

$$P(U_1 = 1;\, p) = pq_G + (1 - p)q_B \tag{4.1}$$

Consequently, $\mathbb{E}\left[U_1;\, p\right] = 1 \times P(U_1 = 1;\, p) + 0 \times P(U_1 = 0;\, p) = pq_G + (1-p)q_B$. Observe that if $p = 1$ (i.e., the consumer is certain that the product is *good*), then $\mathbb{E}\left[U_1\right] = q_G$, and if $p = 0$ (i.e., the consumer is certain that the product is *bad*), then $\mathbb{E}\left[U_1\right] = q_B$. If both *good* and *bad* product types *a priori* dominate the outside option (i.e., $\mathbb{E}\left[U_1;\, 0\right] = q_B > \underline{u}$), the model becomes trivial: the consumer always buys the firm's product in every period, regardless of the product true state and the consumer's prior beliefs on it. Similarly, if both product types are *a priori* dominated by the outside option (i.e., $\mathbb{E}\left[U_1;\, 1\right] = q_G < \underline{u}$), the consumer never buys any product from the firm. To avoid these trivial scenarios, we make an assumption $q_B < \underline{u} < q_G$: that is, the consumer prefers a *good* product over the outside option, but she prefers an outside option over a *bad* product.

We now consider how the consumer beliefs change after the first purchase. When the consumer buys the product and observes the utility $U_1 = 1$, her posterior probability that the product is *good* is:

$$P(S = G \mid U_1 = 1;\, p) = \frac{pq_G}{pq_G + (1 - p)q_B} \tag{4.2}$$

Similarly, when the consumer purchases the product and observes the experience $U_1 = 0$, her posterior probability that the product is *good* is

$$P(S = G \mid U_1 = 0;\, p) = \frac{p(1 - q_G)}{p(1 - q_G) + (1 - p)(1 - q_B)} \tag{4.3}$$

To track the evolution of the consumer's beliefs over time, it is more convenient to express them as an odds ratio. We take the initial odds ratio, $R_0$, to be

$$R_0 = \frac{p}{1 - p}. \tag{4.4}$$

Given the odds ratio $R_0$, the consumer's initial prior that the product is *good* can be written as

$$p = \frac{R_0}{1 + R_0}. \tag{4.5}$$

If the consumer observes a satisfactory experience after the first purchase, her odds ratio will become $R_1 = \frac{q_G}{q_B} R_0$, and if she observes an unsatisfactory experience, it will become $R_1 = \frac{1-q_G}{1-q_B} R_0$.

In the remaining of this section we will show that the consumer's beliefs evolution can be expressed in terms of a random walk with i.i.d. increments. In the analysis that follows we use the random walk properties to characterize the firm's profit.

### 4.2.2 Consumer Beliefs Evolution

By induction, the consumer's odds ratio, at time $t$, $R_t$, is described by the expression:

$$R_t = \left(\frac{q_G}{q_B}\right)^{\sum_{i=1}^{t} U_i} \left(\frac{1-q_G}{1-q_B}\right)^{t-\sum_{i=1}^{t} U_i} R_0, \tag{4.6}$$

where $\sum_{i=1}^{t} U_i$ is the random variable corresponding to the number of satisfactory experiences after buying the product $t$ times, and $t - \sum_{i=1}^{t} U_i$ corresponds the number of unsatisfactory ones.

By applying an order-preserving logarithm transform, to both sides of (4.6), we obtain:

$$\log R_t = \log R_0 + \sum_{i=1}^{t} \left( U_i \log \frac{q_G}{q_B} + (1 - U_i) \log \frac{1-q_G}{1-q_B} \right)$$

Let $q_s \in \{q_G, q_B\}$ denote the true probability of satisfactory experience. Since there is a one-to-one correspondence between the number of successes out of $t$ experiences and the value of $R_t$, the state space of the process $\{\log R_t, t \geq 0\}$ is countable. Therefore, $\{\log R_t, t \geq 0\}$ is a Markov chain, where the transition probabilities are as follows:

$$\log R_t \mid \log R_{t-1} = \begin{cases} \log R_{t-1} + \log \frac{q_G}{q_B} & \text{w.p. } q_s, \\ \log R_{t-1} + \log \frac{1-q_G}{1-q_B} & \text{w.p. } 1-q_s, \end{cases}$$

In general the Markov chain $\{\log R_t, t \geq 0\}$ is not irreducible. To see this, note that for $\log R_t = \log R_{t+s}$ for some $t \geq 0$, $s > 0$, we need $k \log \frac{q_G}{q_B} + (s-k) \log \frac{1-q_G}{1-q_B} = 0$ for some $k = 1, 2, \ldots, s-1$. That is,

$$\frac{k}{s-k} = \frac{\log \frac{1-q_B}{1-q_G}}{\log \frac{q_G}{q_B}}$$

Since $\frac{k}{s-k}$ is a rational number, so must be the expression on the right side of the equation. However, this does not necessarily hold. For example, given values $q_G = \frac{1}{2}$ and $q_B = \frac{1}{4}$, the right side of the equation becomes

$$\frac{\log\left(\frac{3}{4} \times \frac{2}{1}\right)}{\log\left(\frac{1}{2} \times \frac{4}{1}\right)} = \frac{\log 3 - \log 2}{\log 2} = \log 3 - 1,$$

which is not a rational number.

To simplify our analysis, we impose restrictions on $q_G$ and $q_B$ that would make the Markov chain $\{\log R_t, t \geq 0\}$ irreducible. Similarly to Banks and Sundaram (1992), we make an assumption $q_B = 1 - q_G$, which also implies $q_G > 0.5$ and $q_B < 0.5$. Figure 4.1 shows examples of possible values of $q_G$, $q_B$ and $\underline{u}$ that satisfy our assumptions. Note that *good* and *bad* product states are not necessarily "symmetrical" around the outside option: for example, in the middle example in Figure 4.1 a *good* product is only slightly better than the outside option, while a *bad* one is much worse than it.

With this assumption, the Markov chain $\{\log R_t, t \geq 0\}$ has a period of two, and therefore becomes a random walk with the increments

$$\log R_t - \log R_{t-1} = \begin{cases} \log \frac{q_G}{1-q_G} & \text{w.p. } q_s, \\ -\log \frac{q_G}{1-q_G} & \text{w.p. } 1-q_s, \end{cases}$$

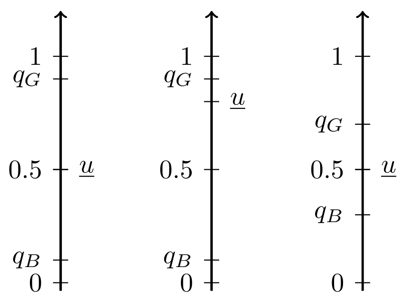


FIGURE 4.1: Examples of problem setup parameters that satisfy the assumption $q_B = 1 - q_G$

Finally, we apply a shift-and-scale transformation to $\{\log R_t, t \geq 0\}$ to obtain $\{Y_t, t \geq 0\}$ – a non-symmetric one dimensional random walk with the increments of $\pm 1$. Formally, we define $Y_t = \psi(R_t, C)$, where

$$\psi(x, c) = \left(\log \frac{q_G}{1 - q_G}\right)^{-1} \log x + c. \tag{4.7}$$

In this transform $c$ can be any arbitrary constant that we can manipulate to shift the starting point $Y_0$. Without this constant (i.e. with $c = 0$) we would have $Y_0 = \psi(R_0, 0) = \left(\log \frac{q_G}{1-q_G}\right)^{-1} \log R_0$. Consequently, to obtain a random walk that starts at zero and takes integer values, we would choose $c = -\left(\log \frac{q_G}{1-q_G}\right)^{-1} \log R_0$. Formally, for some constant $C$, the random walk $\{Y_t, t \geq 0\}$ is defined as:

$$Y_t = Y_0 + W_1 + W_2 + \cdots + W_t, \quad Y_0 = \psi(R_0, C) \tag{4.8}$$

Note that $W_t = 1$ when the consumer observes a satisfactory experience, and $W_t = -1$ when the consumer observes an unsatisfactory experience, so $W_t = 2U_t - 1$. Consequently, $Y_t - Y_0$ is equal to the difference between satisfactory and unsatisfactory experiences observed by the consumer after buying the product $t$ times. $\{W_t, t \geq 1\}$ is a sequence of i.i.d. random variables distributed as $P_s(W_t = 1) = q_s$, $P_s(W_t = -1) = 1 - q_s$; $s \in \{G, B\}$ being the state of the product. We can therefore express (4.6) in terms of $\{W_t, t \geq 1\}$ as

$$R_t = \left(\frac{q_G}{1-q_G}\right)^{\sum_{i=1}^{t} W_i} R_0. \tag{4.9}$$

or, equivalently, we can express it in terms of the initial prior $p$ and the observed experiences $\{U_t, t \geq 1\}$:

$$R_t = \left(\frac{q_G}{1-q_G}\right)^{\sum_{i=1}^{t}(2U_i - 1)} \frac{p}{1-p} \tag{4.10}$$

Since $\psi(x)$ is an order-preserving mapping, $R_i > R_j$ if and only if $Y_i > Y_j$. We use this property in the following sections to characterize the consumer's switching patterns between the firm's products and the outside option.

## 4.3 One-Product Case

In this section, we study how the consumer learning dynamics affects profits of a firm that offers a consumer the same unique product in each period. Recall that the outside option with certain utility $\underline{u}$ is available in each period, therefore at some point the consumer may prefer it over the firm's product, in which case she never returns to buy the product again. We call the last period when the consumer buys the product the *abandonment time*, saying that the *abandonment time is infinite* in a case when the consumer never stops buying the product. In the next section, we derive the distribution of this abandonment time.

### 4.3.1 Consumer's Abandonment Time

We first obtain the distribution of number of purchases made by a consumer conditionally on her prior beliefs and the product's true state in an infinite horizon problem.

We assumed that the consumer always picks an option with the highest expected utility. In one-product case this implies that she keeps on buying the product as long as the expected value of her experience is greater than the certain utility $\underline{u}$ she can derive from an outside option. After the consumer buys the product $t$ times, her expected utility from purchasing it one more time is:

$$
\begin{aligned}
E[U_{t+1} \mid U_1, \ldots, U_t; p] =& E[U_{t+1} \mid S = G] \times P(S = G \mid U_1, \ldots, U_t; p) \\
&+ E[U_{t+1} \mid S = B] \times P(S = B \mid U_1, \ldots, U_t; p) \\
=& \left(1 \times q_G + 0 \times (1 - q_G)\right) \frac{R_t}{1 + R_t} + \left(1 \times (1 - q_G) + 0 \times q_G\right) \frac{1}{1 + R_t} \\
=& q_G \frac{R_t}{1 + R_t} + (1 - q_G) \frac{1}{1 + R_t},
\end{aligned}
$$

where $R_t$ is given in (4.10). Therefore, given the odds ratio $R_t$, the consumer's expected value from sampling the product can be expressed as

$$
E[U_{t+1} \mid R_t] = \frac{q_G R_t + 1 - q_G}{1 + R_t}. \tag{4.11}
$$

Let $\underline{p}$ be the probability of the product being *good* that would make the consumer indifferent between buying the product and picking the outside option, and $\underline{R} = \frac{\underline{p}}{1-\underline{p}}$ be the corresponding odds ratio. Then we have

$$
\begin{aligned}
& \underline{u} = \mathbb{E}\left[U_{t+1} \mid R_t = \underline{R}\right] \\
& \Rightarrow \underline{u} = \frac{q_G \underline{R} + 1 - q_G}{1 + \underline{R}} \\
& \Rightarrow \underline{u} + \underline{u} \cdot \underline{R} = q_G \underline{R} + 1 - q_G \\
& \Rightarrow \underline{R} = \frac{q_G - 1 + \underline{u}}{q_G - \underline{u}}
\end{aligned} \tag{4.12}
$$

The consumer will abandon the product when her odds ratio $R_t$ falls below $\underline{R}$ for the first time. Since the mapping $\psi(x, c)$ defined in (4.7) is order preserving, $R_t \geq \underline{R}$ if and only if $Y_t \geq \psi(\underline{R}, C)$. Therefore, the consumer's abandonment time is $\min\{t : Y_t < \psi(\underline{R}, C)\}$.

The random walk $\{Y_t, t \geq 0\}$ can take values only from the set $Y_0 \pm t$, $t \in \mathbb{N} \cup 0$, so there is only one random walk position when the abandonment can occur, which is a maximum admissible $Y_t$ satisfying the stopping condition. Formally,

$$Y_\tau = Y_0 + k^*; \quad k^* = \max\{k : Y_0 + k < \psi(\underline{R}, C), k \in \mathbb{Z}\} \tag{4.13}$$

Since $R_0 > \underline{R}$ we must have $Y_0 > \psi(\underline{R}, C)$, and the value of $k^*$ must be negative. That is, the consumer needs to see more unsatisfactory experiences than satisfactory ones for her odds ratio to fall below $\underline{R}$.

Note that by manipulating the constant $C$ in (4.7) we can set $Y_0$ in such way that the abandonment occurs when the random walk $\{Y_t, t \geq 0\}$ hits the zero. Let $Q = -k^*$, $Y_\tau = Q > 0$ such that $Y_\tau = 0$. Then the constant $C$ in (4.7) must satisfy

$$\begin{aligned} Q &= \left(\log \frac{q_G}{1 - q_G}\right)^{-1} \log R_0 + C \\ \implies C &= \left(\log \frac{q_G}{1 - q_G}\right)^{-1} \log R_0 - Q \end{aligned} \tag{4.14}$$

Since $Y_0 - Y_\tau = \tau - 2\sum_{t=1}^{\tau} U_t$, $Q$ can be interpreted as the number of unsatisfactory minus satisfactory experiences that characterizes the state of beliefs at which the consumer abandons the product. From (4.13), we obtain:

$$-Q = \max\left\{k : Y_0 + k < \psi\left(\underline{R}, C\right), k \in \mathbb{Z}\right\}. \tag{4.15}$$

From here we can find $Q$ as a minimum of a corresponding set. By rearranging the terms and applying the logarithm transform we obtain an alternative expression for $Q$:

$$\begin{aligned} Q &= \min\left\{k \in \mathbb{Z} : \left(\log \frac{q_G}{1 - q_G}\right)^{-1} \log R_0 + C - k < \left(\log \frac{q_G}{1 - q_G}\right)^{-1} \log \underline{R} + C\right\} \\ &= \min\left\{k \in \mathbb{Z} : k > \frac{\log \underline{R} - \log R_0}{\log(1 - q_G) - \log q_G}\right\}. \end{aligned}$$

From here we can express $Q$ as a function of $\underline{R}$ and $R_0$:

$$Q = \left\lceil \frac{\log R_0 - \log \underline{R}}{\log q_G - \log(1 - q_G)} \right\rceil, \tag{4.16}$$

where $\lceil x \rceil$ is the smallest integer larger or equal to $x$.

And by substituting $c$ in (4.7), we get

$$\psi(x) = \frac{\log x - \log R_0}{\log q_G - \log(1 - q_G)} - Q \tag{4.17}$$

**Example 1.** *Consider an example with the parameters $q_G = 0.55$, $q_B = 1 - 0.55 = 0.45$ and $\underline{u} = 0.5$. The consumer starts with a prior parameter $p = 0.6$. Let us calculate $Q$.*

*First, we use (4.4) and (4.12) to calculate the initial odds ratio $R_0$ and the threshold odds ratio $\underline{R}$ respectively:*

$R_0 = \frac{0.6}{1-0.6} = 1.5$;

$\underline{R} = \frac{0.55-1+0.5}{0.55-0.5} = 1$.

*We can calculate $Q$ by counting the minimum number of unsatisfactory experiences that it will take to achieve the value of odds ratio below $\underline{R} = 1$. From (4.9), assuming that all the consumer's experiences are unsatisfactory (i.e., $\sum_{i=1}^{t} W_i = -t$), we get the following sequence:*

- $R_1 = \left(\frac{0.55}{1-0.55}\right)^{-1} \times 1.5 = \frac{9}{11} \times \frac{3}{2} = \frac{27}{22} > 1$;
- $R_2 = \frac{9}{11} \times \frac{27}{22} = \frac{243}{242} > 1$;
- $R_3 = \frac{9}{11} \times \frac{243}{242} = \frac{2187}{2662} < 1$;

*Therefore, $Q = 3$. We get the same result by using (4.16):*

$Q = \left\lceil \frac{\log 1.5 - \log 1}{\log 0.55 - \log 0.45} \right\rceil \approx \left\lceil \frac{0.405}{-0.598+0.799} \right\rceil \approx \lceil 2.020 \rceil = 3.$

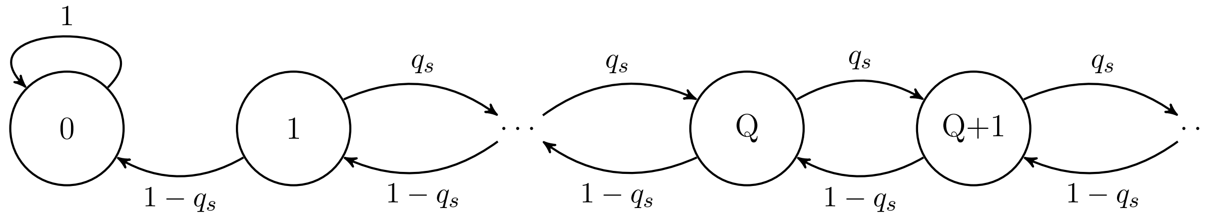


FIGURE 4.2: Evolution of consumer's beliefs as a Markov Chain $\{Y_t\}$, conditional on the product state $S = s$

The mapping $Y_t = \psi(R_t)$ given in (4.17) completes our construction of random walk $\{Y_t : t \geq 0\}$. The consumer's abandonment time is $K$ such that

$$K = \min\{t : Y_t = 0\} \tag{4.18}$$

That is, the consumer's beliefs evolution is a Markov chain with a single absorbing state (i.e., state 0), as is illustrated in Figure 4.2.

To obtain the distribution of the consumer's quitting time we will use the *Hitting time theorem* (see, for example, (Van Der Hofstad and Keane, 2008)). According to this theorem, for a random walk starting at $Q \geq 1$ with i.i.d. steps $\{W_i\}_{i=1}^{\infty}$ satisfying $W_i \geq -1$ almost surely, the distribution of the walk's first hitting time of the origin is given by

$$\mathbb{P}_{s,Q}(K = t) = \frac{Q}{t}\mathbb{P}_{s,Q}(Y_t = 0),$$

where $\mathbb{P}_{s,Q}$ is the law of the random walk characterized by parameter $s$, which is starting at $Q$.

First, we obtain the probability on the right hand side.

$$\begin{aligned}
\mathbb{P}_{s,Q}(Y_t = 0) &= P\left(Q + \sum_{i=1}^{t} W_i = 0 \mid s\right) \\
&= P\left(Q + \sum_{i=1}^{t}(2U_i - 1) = 0 \mid s\right) \\
&= P\left(\sum_{i=1}^{t} U_i = \frac{t-Q}{2} \mid s\right) \\
&= \begin{cases} \binom{t}{(t-Q)/2} q_s^{(t-Q)/2}(1-q_s)^{(t+Q)/2}, & \text{if } \frac{t-Q}{2} \in \mathbb{N}_0, \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}$$

The condition $\frac{t-Q}{2} \in \mathbb{N}_0$ means that if $Q$ is odd, the consumer can only stop after odd number of samples, and if $Q$ is even, the consumer can only stop after even number of samples. That is, we must have $t = 2m + Q$ for some $m \in \mathbb{N}_0$. Note that $m$ corresponds to the number of satisfactory experiences out of $t$ samples. With this notation, the expression above can be reformulated as follows

$$\mathbb{P}_{s,Q}(Y_{2m+Q} = 0) = \binom{2m+Q}{m} q_s^m (1-q_s)^{m+Q}.$$

And the distribution of the consumer's abandonment time can be expressed as

$$\mathbb{P}_{s,Q}(K = 2m+Q) = \frac{Q}{2m+Q}\binom{2m+Q}{m} q_s^m (1-q_s)^{m+Q}. \tag{4.19}$$

**Example 2.** *(Continued from Example 1) Let $q_G = 0.55$, $\underline{u} = 0.5$ and $p = 0.6$. In Example 1 we have shown that for these parameters $Q = 3$.*

*Assume that the product state is good, that is, $s = G$. Then the probabilities that the consumer will abandon the product after purchasing it $1, 2, 3, 4, 5, 6$ and $7$ times will be, respectively*

- $P(K = 1 \mid s = G, Q = 3) = 0$;

- $P(K = 2 \mid s = G, Q = 3) = 0$;
- $P(K = 3 \mid s = G, Q = 3) = P(K = Q + 2 \times 0 \mid s = G, Q = 3) = \frac{3}{3}\binom{3}{0}0.55^0(1-0.55)^3 \approx 0.091$;
- $P(K = 4 \mid s = G, Q = 3) = 0$;
- $P(K = 5 \mid s = G, Q = 3) = P(K = Q + 2 \times 1 \mid s = G, Q = 3) = \frac{3}{5}\binom{5}{1}0.55^1(1-0.55)^4 \approx 0.068$;
- $P(K = 6 \mid s = G, Q = 3) = 0$;
- $P(K = 7 \mid s = G, Q = 3) = P(K = Q + 2 \times 2 \mid s = G, Q = 3) = \frac{3}{5}\binom{7}{2}0.55^2(1-0.55)^5 \approx 0.050$.

*If the product state is bad, that is, $s = B$, those probabilities will be*

- $P(K = 1 \mid s = B, Q = 3) = 0$;
- $P(K = 2 \mid s = B, Q = 3) = 0$;
- $P(K = 3 \mid s = B, Q = 3) = P(K = Q + 2 \times 0 \mid s = B, Q = 3) = \frac{3}{3}\binom{3}{0}0.45^0(1-0.45)^3 \approx 0.166$;
- $P(K = 4 \mid s = B, Q = 3) = 0$;
- $P(K = 5 \mid s = B, Q = 3) = P(K = Q + 2 \times 1 \mid s = B, Q = 3) = \frac{3}{5}\binom{5}{1}0.45^1(1-0.45)^4 \approx 0.124$;
- $P(K = 6 \mid s = B, Q = 3) = 0$;
- $P(K = 7 \mid s = B, Q = 3) = P(K = Q + 2 \times 2 \mid s = B, Q = 3) = \frac{3}{7}\binom{7}{2}0.45^2(1-0.45)^5 \approx 0.092$.

*These distributions for values up to $n = 23$ are plotted in Figure 4.3.*

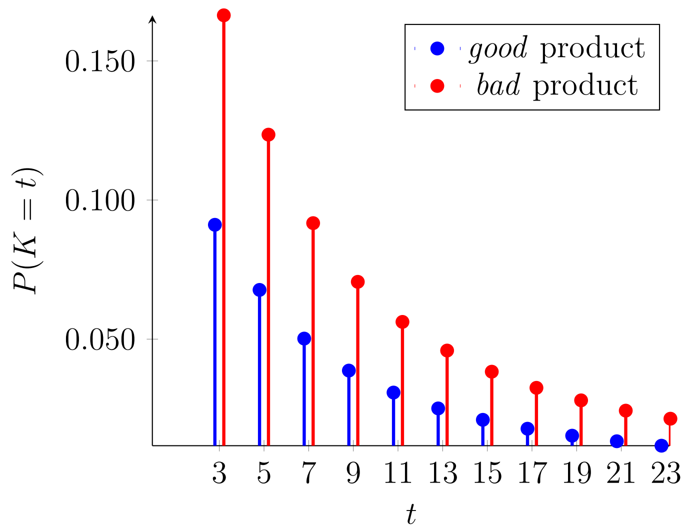


FIGURE 4.3: Abandonment time distribution illustration for Example 2

### 4.3.2 Firm's Profit

When obtaining the distribution for the customer's abandonment time we assumed that the firm will keep offering the product indefinitely. Now we relax this assumption and set the duration of the firm's presence on the market to $T$ periods, that is, $K \leq T$. A finite horizon $T$ can be interpreted as the time of a selling season, time until the planned retirement of a product line, or time until the next assortment policy revision.

Let $r > 0$ be the profit margin for the product and $\rho < 1$ the discount factor. The firm's total discounted profit $V$ at time 1 is a function of the random variable $K$, which is the consumer's abandonment time:

$$V = \begin{cases} r\sum_{t=0}^{k-1}\rho^t, & \text{if } K = k < T; \\ r\sum_{t=0}^{T-1}\rho^t, & \text{otherwise .} \end{cases}$$

Noting that $\sum_{t=0}^{k-1}\rho^t = \frac{1-\rho^k}{1-\rho}$, we obtain the expression for the conditional expectation of the firm's profit:

$$\mathbb{E}\left[V \mid s, Q\right] = r \sum_{k=1}^{T-1} \frac{1-\rho^k}{1-\rho} \mathbb{P}_{s,Q}(K=k) + r\frac{1-\rho^T}{1-\rho}\left(1 - \sum_{k=1}^{T-1} \mathbb{P}_{s,Q}(K=k)\right)$$

$$= r \sum_{k=1}^{T-1} \frac{\rho^T-\rho^k}{1-\rho} \mathbb{P}_{s,Q}(K=k) + r\frac{1-\rho^T}{1-\rho} \quad (4.20)$$

Recall that $Q$ is defined in 4.16 as a function of $R_0, \underline{R}$, which in turn are functions of the problem input parameters $p, \underline{u}$.

For a given $Q$, $K$ is restricted to take values from the set $\{Q, Q+2, \ldots, Q+2M, \text{ T }\}$, where $M = \frac{T-Q-2}{2}$ if $T-Q$ is even, and $M = \frac{T-Q-1}{2}$ if $T-Q$ is odd. For a general case, we can write $M = \lfloor \frac{T-Q-1}{2} \rfloor$, where $\lfloor x \rfloor$ stands for the largest integer less than or equal to $x$. Using this and (4.19), we can rewrite the previous equation as:

$$\mathbb{E}\left[V \mid s, Q\right] = r \sum_{m=0}^{\lfloor \frac{T-Q-1}{2} \rfloor} \frac{\rho^T - \rho^{2m+Q}}{1-\rho} \frac{Q}{2m+Q} \binom{2m+Q}{Q} q_s^m (1-q_s)^{m+Q} + r\frac{1-\rho^T}{1-\rho}$$

By collecting the terms, we get:

$$\mathbb{E}\left[V \mid s, Q\right] = \frac{r}{1-\rho}\left[1 - \rho^T - \sum_{m=0}^{\lfloor \frac{T-Q-1}{2} \rfloor} (\rho^{2m+Q} - \rho^T) \frac{Q}{2m+Q} \binom{2m+Q}{m} q_s^m (1-q_s)^{m+Q}\right]. \quad (4.21)$$

**Example 3.** *(Continued from Examples 1 and 2) Here we will calculate the profits for the problem parameters we introduced in the Examples 1 and 2:* $q_G = 0.55$, $\underline{u} = 0.5$, $p = 0.6$ *and* $s = G$. *For these parameters,* $Q = 3$ *(see Example 1).*

*Assume* $T = 6$, $\rho = 0.9$ *and* $r = 1$. *Using (4.21), we obtain:*

$$
\begin{aligned}
\mathbb{E}\left[V \mid s=G, Q=3\right] &= \frac{1}{1-0.9}\left[1-0.9^6-\sum_{m=0}^{\lfloor\frac{6-3-1}{2}\rfloor}(0.9^{2m+3}-0.9^6)\frac{3}{2m+3}\binom{2m+3}{m}0.55^m(1-0.55)^{m+3}\right] \\
&= 10\times\left[1-0.9^6-(0.9^3-0.9^6)\times 0.45^3-(0.9^5-0.9^6)\times\frac{3}{5}\times\binom{5}{3}\times 0.55^1\times 0.45^4\right] \\
&\approx 10\times\left[1-0.531-(0.729-0.531)\times 0.091-(0.590-0.531)\times 0.068\right] \\
&\approx 4.466
\end{aligned}
$$

### 4.4 Two-Product Case

We now consider the case where the firm offers two products in each period – product $i$ and product $j$ with the respective profit margins $r_i$ and $r_j$, states (*good* or *bad*) $s_i$ and $s_j$ and the consumer's initial beliefs on these states $p_i$ and $p_j$. We denote the probability of satisfactory experience from products $i, j$ as $q_i, q_j$ respectively, where $q_i = q_G$ if $s_i = G$ and $q_i = 1 - q_G$ if $s_i = B$. Similarly, $q_j = q_G$ if $s_j = G$ and $q_j = 1 - q_G$ if $s_j = B$.

As in the one-product scenario, the consumer has outside option with the certain utility $\underline{u}$ and in each period the consumer chooses the option with the highest current expected utility among the three available options: firm's products $i$ and $j$ and the outside option.

Like in the one-product scenario, we assume that the firm will be offering the products for the time horizon $T$, and in our analysis we will use the notation $T \to \infty$ to indicate that a certain statement holds only in absence of time horizon restrictions. Recall that $\rho < 1$ is a discount factor. Then the firm's total discounted profit is

$$
V = \sum_{t=1}^{T} \rho^{t-1}\left(r_i I_{it} + r_j I_{jt}\right), \tag{4.22}
$$

where $I_{it}$ and $I_{jt}$ are indicator random variables that take the value 1 if the corresponding product was purchased in period $t$.

We assume that the consumer's initial beliefs on the products states are independent, so after experiencing a product the consumer will only update her beliefs on the state of that product. This implies that, like in a one-product case, if in a period $t$ the consumer prefers the outside option to the firm's product $i$ or $j$, then she never buys this product in any of the subsequent periods $t+1, t+2, \ldots$, and her beliefs on this product stay unchanged. Since both products are offered in each period, once the consumer chooses the outside option, she never buys from the firm again. We will refer to the time frame of the consumer buying from the firm as the *consumer lifetime*. For example, if the consumer chooses the outside option for the first time in period 6, then the consumer lifetime is 5 periods.

Just like in the one-product case, the firm's profit depends on whether or not the consumer switches to the outside option before the end of the time horizon $T$, and if she does, in what period it occurs. Having a second product in the assortment gives the firm some "insurance": before switching to the outside option, the consumer needs to buy and reject two products instead of one.

We assume that the firm's assortment policy is fixed for the whole time horizon $T$, and in this section we are looking into the case when the firm is offering both products in every period. In order to characterize the firm's profit, we want to know how many times the consumer buyw product $i$, how many times she buys product $j$ and, in presence of discounting factor $\rho < 1$, we also want to know how these purchases are spread over the problem time horizon.

We represent the consumer's beliefs as odds ratios $R_{i0} = \frac{p_i}{1-p_i}$ and $R_{j0} = \frac{p_j}{1-p_j}$, and we define $R_{iv}$ and $R_{jw}$ as the consumer's odds ratios after buying product $i$ for $v$ times and product $j$ for $w$ times respectively. The consumer can buy at most one product at a time, therefore if $R_{iv}$ and $R_{jw}$ are the odds ratios in $t$ periods from the beginning of time horizon, $v + w \leq t$, where the equality holds if the consumer does not switch to the outside option before the period $t$.

The consumer abandons product $i$ after buying it $v$ times if $R_{iv} < \underline{R}$ (and subsequently, she abandons product $j$ when $R_{jw} < \underline{R}$), where $\underline{R}$ is given in (4.12). Consequently, the consumer lifetime is $t = v + w$ such that $R_{iv} < \underline{R}$ and $R_{jw} < \underline{R}$ for some $v, w \in \mathbb{N}$. Note that if either $R_{i0} < \underline{R}$ or $R_{j0} < \underline{R}$, the consumer will never buy one of the products, and the case is equivalent to one-product case. To avoid this, as well as the trivial case where the consumer never buys any product, we assume that $R_{i0} > \underline{R}$ and $R_{j0} > \underline{R}$.

We apply the transform $\psi(\cdot)$ to obtain two stochastic processes $\{Y_{iv}, v \geq 0\}$ and $\{Y_{jw}, w \geq 0\}$ such that $Y_{iv} = \psi_i(R_{iv})$ and $Y_{jn} = \psi_j(R_{jw})$, where $\psi_i(\cdot)$ and $\psi_j(\cdot)$ are obtained from (4.17) by substituting $R_0$ with $R_{i0}$ and $R_{j0}$ respectively. We denote the starting points of these random walks as $Y_{i0} = Q_i$ and $Y_{j0} = Q_j$, where $Q_i, Q_j$ are obtained by substituting $R_{i0}$ and $R_{j0}$ for $R_0$ in (4.16). Like in the one-product scenario, $Q_i$ and $Q_j$ can be interpreted as difference between the number of unsatisfactory and satisfactory experiences for product $i$ and $j$ respectively that results in the abandonment of this product.

The obtained stochastic processes have independent increments that we denote $W_{ik}, W_{jk} \in \{-1, +1\}$, where $P(W_{ik} = 1) = q_i$ and $P(W_{jk} = 1) = q_j$ (subsequently, $P(W_{ik} = -1) = 1 - q_i$ and $P(W_{jk} = -1) = 1 - q_j$).

$$Y_{iv} = Y_{i0} + \sum_{k=1}^{v} W_{ik}; \qquad Y_{jv} = Y_{j0} + \sum_{v=1}^{k} W_{jk}, \quad v,\, w = 0, 1, \ldots \tag{4.23}$$

Because we assumed that the consumer is myopic, it follows that she buys the product with the highest probability of being *good* first. We use the index $H$ (i.e., *High*) for this product and the index $L$ (i.e., *Low*) for the other product: $H = \arg\max_{\{i,j\}}[p_i, p_j]$, $L = \arg\min_{\{i,j\}}[p_i, p_j]$. We assume there is no ties, so $\{H, L\} = \{i, j\}$. The consumer then keeps buying product $H$ as long as her odds ratio for this product stays above the odds ratio for product $L$.

After $t$ periods of this successive buying of product $H$ the consumer's beliefs are $(R_{Ht}, R_{L0})$, that is, the beliefs for product $L$ stay unchanged, and $R_{Ht} = \psi_H^{-1}(Y_{Ht})$, where $Y_{Ht} \in \{Y_{it}, Y_{jt}\}$

is the element of one of the two random walks defined in (4.23). Consider the sequence of two-dimensional vectors $(R_{H0}, R_{L0}), (R_{H1}, R_{L0}), \ldots, (R_{Hv}, R_{Lw}), \ldots)$, where $(R_{Hv}, R_{Lw})$ are the consumer's odds ratios at time $t$ such that $t = v + w$.

Let $\tilde{\tau}$ be the number of periods that it takes for the random walks to get into the same position:

$$\tilde{\tau} = \min\{t : Y_{Ht} = Y_{L0}\}. \tag{4.24}$$

Observing the analogy with one-product problem, we derive the distribution for $\tilde{\tau}$ by substituting $Y_{H0} - Y_{L0} = |Q_i - Q_j|$ for $Q$ in (4.19)

$$P(\tilde{\tau} = 2m + |Q_i - Q_j|) = \frac{|Q_i - Q_j|}{2m + |Q_i - Q_j|} \binom{2m + |Q_i - Q_j|}{m} q_H^m (1 - q_H)^{m + |Q_i - Q_j|}, \tag{4.25}$$

where $q_H = q_G$ if product $H$ is *good*, and $q_H = 1 - q_G$ if product $H$ is *bad*.

Recall that Q defined in (4.16) is a discrete-valued function of consumer's odds ratio, and therefore a function of $p$. Thus, depending on the exact values of parameters $p_i$ and $p_j$, with $Y_{H\tilde{\tau}} = Y_{L0}$ we may have $R_{H\tilde{\tau}} > R_{L0}$, in which case the consumer continues buying product $H$ in period $\tilde{\tau} + 1$, or $R_{H\tilde{\tau}} < R_{L0}$, in which case $\tilde{\tau}$ is a time of switching to product $L$. Since product labels $i$ and $j$ can be used interchangeably, we say that the product which is bought at the period $\tilde{\tau} + 1$ is product $i$.

We define a *switching time* as the last period of consecutive uninterrupted sampling of the same product. For example, if the consumer buys product $i$ in periods 1 and 4, product $j$ in periods 2 and 3 and resorts to the outside option starting from period 5 then the switching times are $1, 3$ and 4. We show that starting from period $\tilde{\tau} + 1$ the consumer switching times and her corresponding buying behavior follow a certain pattern.

Let $J_n$ be the time $n$th switching observed in the time interval $[\tilde{\tau}, T]$, where $T$ can potentially be infinite, and let $\tau_n = J_n - J_{n-1}$ for $n \geq 1$, where $J_0 = \tilde{\tau}$. We refer to $\tau_1, \tau_2, \ldots$ as *inter-switching times*, by analogy with interarrival times in queuing models.[1]

$$J_n = J_0 + \sum_{k=1}^{n} \tau_k, \quad J_0 = \tilde{\tau},\ n = 1, 2, \ldots, 2\min\{Q_i, Q_j\} \tag{4.26}$$

We also define $N(t)$ as the number of switchings observed in the time interval $(\tilde{\tau}, t)$, or, equivalently, $[\tilde{\tau}+1, t-1]$, where $N(t) = 0$ if $t < \tilde{\tau}$. $N(t)$ is defined in terms of the process $\{J_n, n \geq 0\}$ as

$$N(t) = \max\{n : J_n < t\} \tag{4.27}$$

**Example 4.** *Here we look at the sample path of the consumer beliefs for two products shown in Figure 4.4. For consistency with the examples for one-product case, we use the parameters* $q_G = 0.55$ *and* $\underline{u} = 0.5$. *Let* $T = 15$, $p_H = 0.75$ *and* $p_L = 0.6$. *Recall that we use the label* $i$ *for the product that consumer buys at time* $\tilde{\tau}+1$, *where* $\tilde{\tau}$ *is defined in (4.24). As we show below, for these parameters* $i = H$, *so from here we omit the* $\{H, L\}$ *notation and write* $p_i = 0.75$ *and* $p_j = 0.6$.

*For this example we use the following sample sequences of experiences from the products:* $\mathbf{U}_i = (0,0,0,1,0,0,0,0)$ *and* $\mathbf{U}_j = (1,0,0,0,1,1,1)$. *The top of Figure 4.4 shows a sample sequence* $(R_{i0}, R_{j0}), (R_{i1}, R_{j0}), \ldots, (R_{iv}, R_{jw}), \ldots)$ *plotted against the time horizon* $t$, *and the bottom shows the corresponding random walk values* $(Y_{i0}, Y_{j0}), (Y_{i1}, Y_{j0}), \ldots, (Y_{iv}, Y_{jw}), \ldots)$ *plotted against the same time horizon. Below are the step-by-step calculations for the first few values of these sequences.*

*As shown in Example 1, the values* $q_G = 0.55$ *and* $\underline{u} = 0.5$ *result in threshold odds ratio* $\underline{R} = 1$. *Since* $p_j = 0.6$, *which is equal to the value of* $p$ *we used in Examples 1-3, the values*

[1] Note that $\tau_1$ is not necessarily an inter-switching time in a literary sense, because, as we discussed, the switching does not necessarily occur at time $J_0 = \tilde{\tau}$.

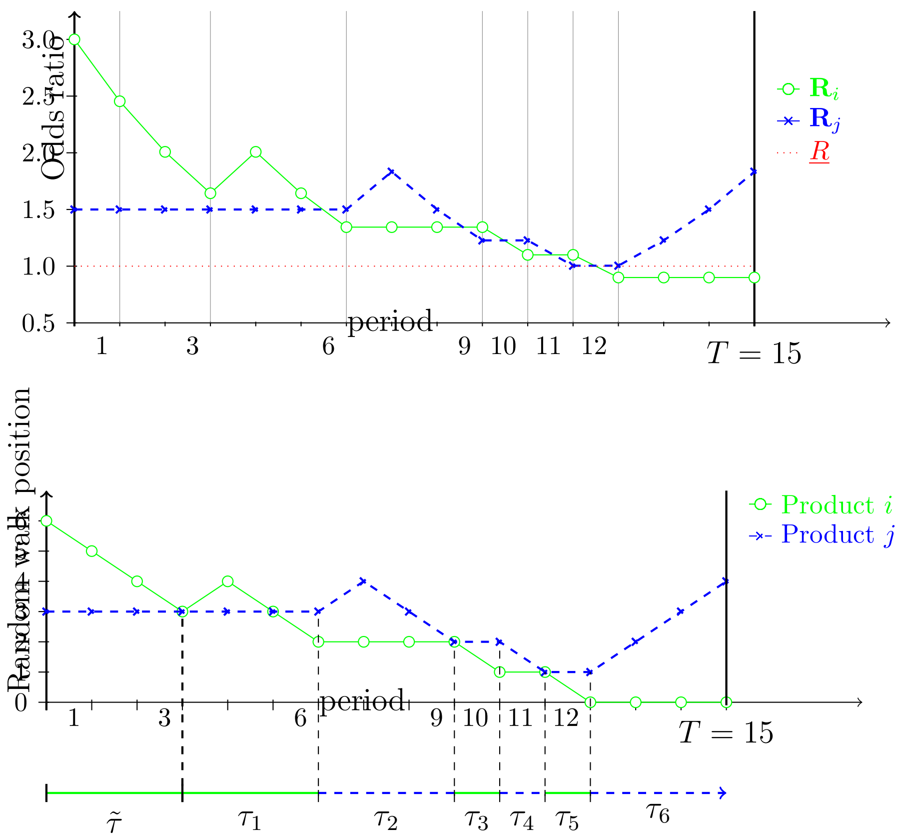


FIGURE 4.4: Sample path of consumer beliefs evolution for two-product case

*of* $R_{j0}$ *and* $Q_j$ *are equal to the values of* $R_0$ *and* $Q$ *we obtained in those examples:* $R_{j0} = 1.5$ *and* $Q_j = 3$.

*For product* $i$*, we get* $R_{i0} = \frac{0.75}{1-0.75} = 3$ *and* $Q_i = \left\lceil \frac{\log 3 - \log 1}{\log \frac{0.55}{0.45}} \right\rceil = \left\lceil \frac{0.477}{0.087} \right\rceil = \lceil 5.483 \rceil = 6$. *Therefore,* $Q = \min\{Q_i, Q_j\} = 3$

*By setting* $Y_{i0} = Q_i = 6$ *and using the values of* $\mathbf{U}_i$*, we find that the elements of the random walk* $\{Y_{iv}, v \geq 0\}$ *are* $(6, 5, 4, 3, 4, 3, 2, 1, 0)$*: that is, the walk hits the origin after 8 steps. Similarly, by setting* $Y_{j0} = Q_j = 3$ *and using the values of* $\mathbf{U}_j$*, we find the first 7 elements of the random walk* $\{Y_{jw}, w \geq 0\}$ *are* $(3, 4, 3, 2, 1, 2, 3, 4)$*. Note that it does not hit the origin yet, so the consumer can potentially buy product* $j$ *more than 7 times.*

*We use* $R_{iv} = \left(\frac{q_G}{1-q_G}\right)^{Y_{iv}-Y_{i0}} R_{i0}$ *to obtain*

- $R_{i0} = 3$

- $R_{i1} = \left(\frac{0.55}{0.45}\right)^{-1} \times 3 \approx 2.454,$
- $R_{i2} = \left(\frac{0.55}{0.45}\right)^{-2} \times 3 \approx 2.008,$
- $R_{i3} = \left(\frac{0.55}{0.45}\right)^{-3} \times 3 \approx 1.643,$
- $R_{i4} = \left(\frac{0.55}{0.45}\right)^{-2} \times 3 \approx 2.008,$
- $R_{i5} = \left(\frac{0.55}{0.45}\right)^{-3} \times 3 \approx 1.643,$
- $R_{i6} = \left(\frac{0.55}{0.45}\right)^{-4} \times 3 \approx 1.344,$
- $R_{i7} = \left(\frac{0.55}{0.45}\right)^{-5} \times 3 \approx, 1.1$
- $R_{i8} = \left(\frac{0.55}{0.45}\right)^{-6} \times 3 \approx 0.9$

*In the same way, for product $j$ we get $R_{jw} = 1.500, 1.833,\ 1.500, 1.227, 1.004, 1.227,$ $1.5000,\ 1.833\ldots.$*

*From (4.24) we obtain $J_0 = \tilde{\tau} = \min\{t : Y_{it} = Y_{j0})\} = 3$, that is, after 3 periods the random walks $\{Y_{vi}, v \geq 0\}$ and $\{Y_{jw}, w \geq 0\}$ get into the same position for the first time. Since $R_{i3} = 1.643 > R_{j0}$, at time $\tilde{\tau}+1 = 4$ the consumer buys product $i$, which confirms that the labeling of the products is as intended. Note that there is no actual switching at time $\tilde{\tau}$.*

*The consumer switches from product $i$ to product $j$ for the first time at period $J_1 = 6$, because $R_{i6} = 1.344 < R_{j0}$, that is, the odds ratio for product $i$ falls below the odds ratio for product $j$ for the first time. Therefore, $\tau_1 = 6 - 3 = 3$. Observe that at this point $Y_{i6} = 2 = Y_{j0} - 1$. The random walk $\{Y_{iv}, v \geq 0\}$ stays at its initial position, and the consumer beliefs remain at the state $R_{i6} = 1.344$ until the random walk $\{Y_{jw}, w \geq 0\}$ gets into the position $R_{j3} = 1.227 < R_{i6}$. Observe that at this occurs at time $J_2 = 9$ and $Y_{i6} = Y_{3j} = 2$. As we state in Theorem 1, this pattern of the positions $(Y_{iw}, Y_{jw})$ at the switching times $t = v + w$ holds everywhere after the time $\tilde{\tau}$.*

*In this example, we observe $N(T) = 5$ switchings, and the last one occurs at time $J_{N(T)} = 12$; by this time product $i$ is bought 8 times and $Y_{i8} = 0$, so product $i$ gets abandoned. Therefore, in absence of time horizon limit we could observe at most one more switching: from product $j$ to the outside option, for a total of 6 switchings, which happens to be equal to $2Q_j$. In Theorem 1 we argue that the maximum number of switchings that can be observed after time $\tilde{\tau}$ is indeed $2\min\{Q_i, Q_j\}$.*

THEOREM 1. *Let $T \to \infty$. Then for any $t \geq 0$ we have $N(t) \leq 2\min\{Q_i, Q_j\}$, where $N(t)$ is defined in (4.27). Furthermore, $\tau_1, \tau_2, \ldots$ are independent random variables that can take only odd values, and for $m = 0, 1, \ldots$ we have*

$$P(\tau_n = 2m+1) = q_n \times \frac{1}{2m+1}\binom{2m+1}{m}\left(q_G(1-q_G)\right)^m, \quad n = 1, 2, \ldots, 2\min\{Q_i, Q_j \tag{4.28}$$

*where $q_n = q_i$ if $n$ is odd, and $q_n = q_j$ if $n$ is even.*

In Theorem 1 we essentially argue that in absence of time horizon limits (i.e., $T \to \infty$), inter-switching times $\tau_1, \tau_3, \tau_5, \ldots$ are distributed as a nonnegative random walk with unit increments with the probability of an upwards step $q_i$ (and, consequently, probability of a downwards step $1 - q_i$), and inter-switching times $\tau_2, \tau_4 \ldots$ are distributed as a nonnegative random walk with unit increments with the probability of an upwards step $q_j$ (and, consequently, probability of a downwards step $1 - q_j$). Here, by nonnegative random walk we understand a random walk that interrupts when in gets into a negative position for the first time (Bertoin and Doney, 1994).

The formal proof of Theorem 1 is provided in Appendix C.1, and here we provide some intuition for it. Note that $J_0 = \tilde{\tau}$ is the stopping time for the random sequence $(R_{iv}, R_{jw})$. Thus, for simplicity of notation, we set $\tilde{\tau} = 0$, which implies $Y_{i0} = Y_{j0}$ and $R_{i0} > R_{j0}$. In this case, $J_1 = \tau_1$, that is, $\tau_1$ is the first switching time. In the first period the consumer buys product $i$, and she keeps buying it as long as the odds ratio for $i$ stays above odds ratio for $j$. That is, until the first switching, product $j$ plays the same role as an outside option in one-product case. We show that $\tau_1 = \min\{t : R_{it} < R_j 0\}$ is such that $Y_{i\tau_1} = Y_{j0} - 1$, that is, the random walk gets one step down from the initial position for the first time. This is identical to the one-product case abandonment time, so we obtain (4.28) by substituting $\tau_1$ for $K$, $q_i$ for $q_s$ and 1 for $Q$ in (4.19).

Further, we argue that $J_1$ is also a stopping time for $(R_{iv}, R_{jw})$, that is, the time of the first switching has no effect on the duration between the subsequent switchings. In period $\tau_1 + 1$ the consumer buys product $j$ and keeps buying it as long as the odds ratio for $j$ stays above odds ratio for $i$. We show that $\tau_2 = \min\{t : R_{jt} < R_i \tau_1\}$ is such that $Y_{j\tau_2} = Y_{i\tau_1} = Y_{j0} - 1$. This alternating pattern repeats again and again: consumer switches from product $i$ to product $j$ when $Y_{i\tau_1+\tau_3+\cdots+\tau_{2n+1}} = Y_{j\tau_2+\tau_4+\cdots+\tau_{2n}} - 1$ and from product $j$ to product $i$ when when $Y_{i\tau_1+\tau_3+\cdots+\tau_{2n+1}} = Y_{j\tau_2+\tau_4+\cdots+\tau_{2n}}$. Since with every switching one of the random walks gets one step closer to the outside option, $N(t) = 2\min\{Q_i, Q_j\}$ implies that at time $t$ both products are already abandoned, so $2\min\{Q_i, Q_j\}$ is the maximum number of switchings we can observe after $\tilde{\tau}$.

Observation 1. *In this observation we confirm the following Gambler's Ruin property: for $T \to \infty$ probability that the consecutive buying of the product ever gets interrupted is* 1 *for a* bad *one, and* $\frac{1-q_G}{q_G}$ *for a* good *one (Ross (1995, p. 344)).*

$$
\begin{aligned}
P_s(\tau_n < \infty) =&(1-q_n)\sum_{m=0}^{\infty}\frac{1}{2m+1}\binom{2m+1}{m}\left(q_G(1-q_G)\right)^m && (4.29)\\
=&(1-q_n)\frac{2}{1+\sqrt{1-4\left(q_G(1-q_G)\right)}} && (4.30)\\
=&(1-q_n)\frac{2}{1+\sqrt{1-4q_G+4q_G^2}}\\
=&(1-q_n)\frac{2}{1+2q_G-1}\\
=&\frac{1-q_n}{q_G} && (4.31)\\
=&\begin{cases} 1 & \textit{if } q_n = 1-q_G,\\ \frac{1-q_G}{q_G} & \textit{if } q_n = q_G. \end{cases} && (4.32)
\end{aligned}
$$

*To make the step (4.29)-(4.30) we use a binomial series generating function (Graham et al., 1994, §5.4).*

*In other words, in absence of time horizon limit any inter-switching time when the consumer is buying a* good *product has a nonzero probability of being is singular.*

It is important to stress that for a *good* product $\tau_n$ is a "random variable in extended sense" (Miller (1961)), because it takes values from the extended real line $\mathbb{R}\cup\infty$. In particular, for a good product $\int_{\mathbb{R}} f_{\tau_n}(x)dx < 1$, where $f_{\tau_n}(x)$ is a generalized probability density function of $\tau_n$. We have to keep this in mind when operating with this variable and the functions of it, including the firm's expected profit.

## 4.5 Summary and Conclusions

In this chapter we have developed a mathematical model for the problem of a firm that repeatedly sells a product to a consumer who learns about her preferences over time. We have shown that in a case when the firm is offering only one product, the problem setup can be linked to a specific type of random walk known as a gambler's ruin problem. Using this property, we have derived the expressions for the expected product lifetime and the firm's expected profit in the case when the firm offers one product.

We have also provided some initial analysis of consumer's beliefs evolution and switching patterns for the case when the firm offers two products. The next step in our analysis will be to obtain the expression for expected profit when the firm offers the two products in every period of the finite planning horizon. Finally we will consider the case where the firm can decide which of two products to offer in every period.

The mathematical framework developed in this chapter, in particular, Theorem 1 establishes a ground for evaluating firm's assortment policies for the case when the consumer's preferences change over time due to learning process. The properties of the consumer's beliefs evolution opens the possibilities for generalizing our model for assortments that include multiple products. The other potentially interesting future research directions is relaxing the assumption on the firm's knowledge of the consumer's beliefs and the product state and incorporating the possibility for the firm to revise the assortment after the beginning of the time horizon.

# APPENDIX A

# APPENDIX FOR CHAPTER 2

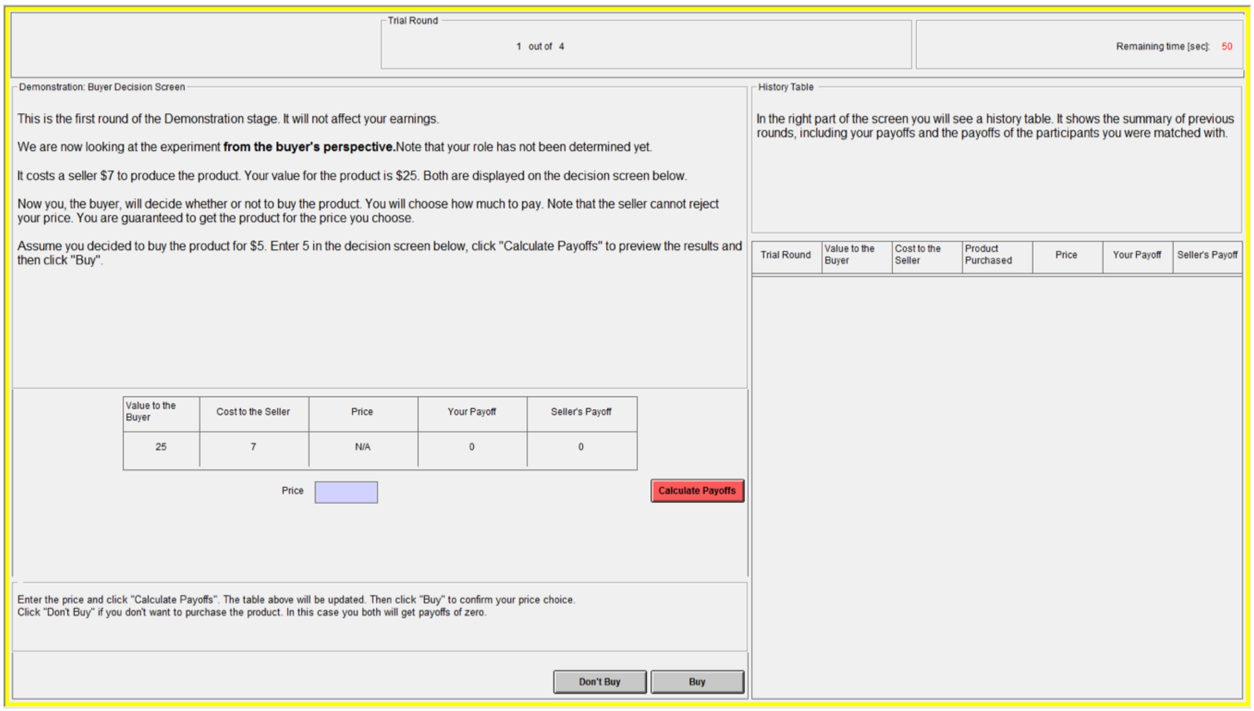


FIGURE A.1: Subject interface for PWYW experiment, instructional stage

# APPENDIX B

# APPENDIX FOR CHAPTER 3

## B.1 Proof of Greedy Algorithm Optimality for Assortment Planning Model

To prove the optimality of the greedy algorithm, we will show that when the set of possible assortments is restricted to popular sets only, the expected profit is discrete concave in the assortment variety. Let $A_k$ be a popular assortment with variety $k$, that is $A_k = \{1, 2, \ldots, k\}$, where $k \leq n$. We want to show that $\mathbb{E}[\pi(A_{k-1})] + \mathbb{E}[\pi(A_{k+1})] \leq 2\mathbb{E}[\pi(A_k)]$ for any admissible $k$.

$$
\begin{aligned}
\mathbb{E}[\pi(A_{k-1})] + \mathbb{E}[\pi(A_{k+1})] = &\ \lambda \left(1 - \frac{v_0}{\sum_{i=1}^{k-1} v_i + v_0}\right) - (k-1)K \\
&+ \lambda \left(1 - \frac{v_0}{\sum_{i=1}^{k+1} v_i + v_0}\right) - (k+1)K \\
&= \lambda \left[2 - v_0 \left(\frac{1}{\sum_{i=1}^{k} v_i + v_0 - v_k} + \frac{1}{\sum_{i=1}^{k} v_i + v_0 + v_{k+1}}\right)\right] - 2kK \\
&\leq 2\lambda \left(1 - \frac{v_0}{\sum_{i=1}^{k} v_i + v_0}\right) - 2kK \\
&= 2\mathbb{E}[\pi(A_k)],
\end{aligned}
$$

where the inequality follows from the convexity of the function $x^{-1}$.

Since $\mathbb{E}[\pi(A_k)]$ is discrete convex, an optimal popular assortment can be found via a greedy algorithm. Since the set of popular assortments must contain an optimal assortment, the obtained solution will be globally optimal.

## B.2 Subject Instructions

# Assortment Planning Game

In this experiment, you will assume a role of an assortment planner: that is, you will be picking the products that your company is offering on the market. You will be making this decision for 25 consecutive periods. The periods are *independent*, which means that your decisions and outcomes in one period will not affect outcomes in any other period. At the end of the game the profits in EU (Experimental Units) you have accumulated over all the periods will be converted to a US dollars amount and paid to your murk account as a HIT bonus.

### How your sales are calculated

In each period there are 1,000 potential consumers. Each consumer buys exactly one unit: either from your assortment or from your competitors'.

Your marketing manager has conducted a survey and calculated a *popularity index* for each product that you can potentially include into your assortment. In addition to that, she has computed an aggregated popularity index for the products of your competitors. The chance that a consumer will buy a particular product is proportional to this product's popularity index.

Consider the example below: you can include two products into your assortment: product A and product B.

| Product | Popularity Index |
|---|---|
| A | 4 |
| B | 2 |
| Competitors | 6 |

If you include only product A in your assortment, the chance that a consumer buys it is:

- $4/(4+6) = 40\%$ for product A
- $6/(4+6) = 60\%$ for your competitors' products.

If you include both products A and B in your assortment, the chance that a consumer buys it is:

- $4/(4+2+6) = 33\%$ for product A
- $2/(4+2+6) = 17\%$ for product B
- $6/(4+2+6) = 50\%$ for your competitors' products.

In each period, each one of the 1,000 potential customers randomly chooses between the available products – yours and your competitors' assortments. The resulting sales for each product in your assortment have a normal[1] (bell-shaped) distribution with a mean of:

$$Average\ sales\ =\ Chance\ of\ buying\ \times\ 1{,}000$$

*Your actual sales can be more or less than the average sales and will vary even if you select the same assortment in every period.*

For example, if the chance that a customer buys product A is 40%, the average sales will be $40\% \times 1{,}000\ =\ 400$. The distribution of *actual sales* for this case is shown on the picture below.

**Figure 1: Example of actual sales distribution for a single product**

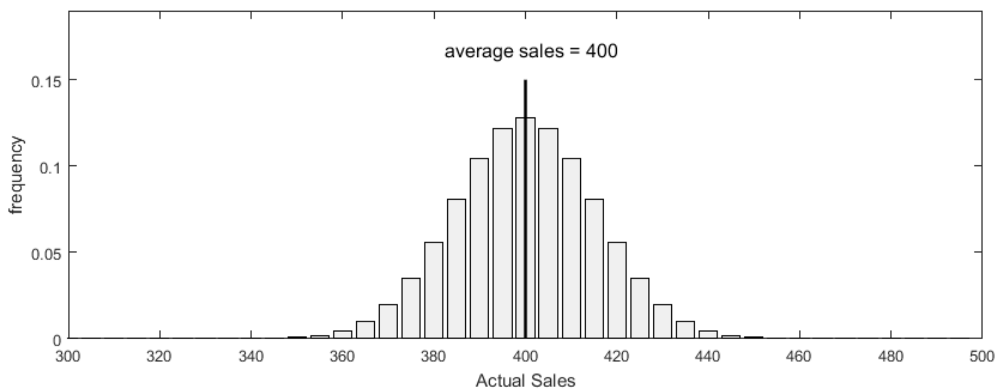


## How you make money

You have an abundant supply of all products available immediately upon request, so you can satisfy any demand and you do not incur any inventory costs. However, there is a *fixed operational cost* for every product you include into your assortment. For example, if your assortment includes three products, and the operational cost is 50 EU, your total fixed operational costs are $50\ EU \times 3\ =\ 150\ EU$, regardless of your actual sales.

For each unit of product you sell, you earn **1 EU**. Your objective is to maximize your profit in each period, which is given by:

$$Profit\ = (1\ EU\ \times\ Actual\ Sales) - (Fixed\ Operational\ Cost\ \times\ Number\ of\ Products\ in\ the\ Assortment)$$

## How your bonus payment is calculated

The game will end after 25 periods. You will then see your total profit (in EU), which is simply the sum of the profits you have accumulated across all 25 periods. Your total profit will be converted to a US dollars amount at the rate of **1000 EU = $1**.

[1] Actually, the distribution is binomial with a mean of $N \times p$ and a standard deviation of $\sqrt{N \times p \times (1-p)}$, where $N$ is the number of potential consumers and $p$ is the probability of buying. However, since $N$ is very large, it very much resembles a normal distribution.

## B.3 Consent Form

The following text was displayed as an HIT description on Amazon mTurk.

Welcome to "Assortment Planning Study," a web-based experiment that examines human decision process under uncertainty. Participation in the study typically takes about 30 minutes. Participants will receive an HIT bonus of up to about $8, based on performance. We expect that most participants will earn a bonus around $6.

Before taking part in this study, please read the consent form below and click on the "Accept HIT" button at the bottom of the page if you understand the statements and freely consent to participate in the study.

**Consent Form**

This study involves a web-based experiment designed to improve our understanding of how humans make decisions in uncertain environments. The study is being conducted by Yulia Vorotyntseva, Dr. Elena Katok and Dr. Dorothee Honhon of The University of Texas at Dallas, and it has been approved by The University of Texas at Dallas Institutional Review Board. No deception is involved, and the study involves no more than minimal risk to participants (i.e., the level of risk encountered in daily life).

Participation in the study typically takes about 30 minutes and is strictly confidential. Participants will act as retailers on a made-up market where they will need to select an assortment of products to sell. Participants will see products characteristics associated with the potential demand and costs associated with including each product in the assortment. After selecting the products, participants will see randomly generated demand and corresponding profit.

All responses are treated as confidential, and in no case will responses from individual participants be identified. Rather, all data will be pooled and published in aggregate form only. Participants should be aware; however, that the experiment is not being run from a "secure" https server of the kind typically used to handle credit card transactions, so there is a small possibility that responses could be viewed by unauthorized third parties (e.g., computer hackers). Your Mechanical Turk Worker ID will be used to distribute payment to you but will not be stored with your [survey responses/data]. Please be aware that your MTurk Worker ID can potentially be linked to information about you on your Amazon public profile page, depending on the settings you have for your Amazon profile. We will not be accessing any personally identifying information about you that you may have put on your Amazon public profile page.

Many individuals find participation in this study enjoyable, and no adverse reactions have been reported thus far. Participants will receive an HIT bonus of up to about $8, based on performance.

Participation is voluntary, refusal to take part in the study involves no penalty or loss of benefits to which participants are otherwise entitled, and participants may withdraw from the study at any time without penalty or loss of benefits to which they are otherwise entitled.

If participants have further questions about this study, they may contact the Principal Investigator, Yulia Vorotyntseva at yulia.vorotyntseva@utdallas.edu, or the UT Dallas Faculty Sponsor, Elena Katok at ekatok@utdallas.edu; Participants who want more information about their rights as a participant or who want to report a research related concern may contact The University of Texas at Dallas Institutional Review Board at (972) 883-4579.

**If you are 18 years of age or older, understand the statements above, and freely consent to participate in the study, click on the "Accept HIT" button to begin the experiment.**

## B.4 Additional Regression Models Output

Table B.1: Subject totals: censored regressions with full interactions

| | *Dependent variable:* | |
|---|---|---|
| | Normalized expected profit | # of optimal assortments |
| (Intercept) | 92.703*** (3.503) | 12.765*** (2.513) |
| PB | 4.135 (4.391) | 4.521 (3.143) |
| EP | 7.879 (5.084) | 6.800* (3.605) |
| HOV | −4.625 (4.954) | −4.415 (3.591) |
| PB×HOV | −10.497 (6.432) | −9.250** (4.686) |
| EP×HOV | −4.100 (6.888) | −1.997 (4.934) |
| logSigma | 2.496*** (0.077) | 2.157*** (0.086) |
| Total observations | 88 | 88 |
| Left-censored | 0 | 11 |
| Right-censored | 4 | 4 |
| Log Likelihood | −331.834 | −276.362 |
| *Note:* | *p<0.1; **p<0.05; ***p<0.01 | |

Table B.2: Cross sectional regression: last period

| | *Dependent variable:* | |
|---|---|---|
| | Normalized expected profit | Is optimal? |
| | *censored* | *binary logistic* |
| (Intercept) | 112.564*** (10.164) | 0.693 (0.612) |
| PB | 20.336 (13.685) | 1.099 (0.874) |
| EP | 27.664 (17.746) | 1.705 (1.210) |
| HOV | −4.718 (13.452) | −0.357 (0.847) |
| PB×HOV | −37.279** (18.383) | −2.447** (1.203) |
| EP×HOV | −20.771 (21.699) | −1.253 (1.449) |
| logSigma | 3.327*** (0.142) | |
| Observations | 88 | 88 |
| Log Likelihood | −161.171 | −46.479 |
| *Note:* | *p<0.1; **p<0.05; ***p<0.01 | |

Table B.3: Panel regressions with full interactions

| | *Dependent variable:* | |
|---|---|---|
| | Normalized expected profit | Is optimal? |
| (Intercept) | 82.674*** (1.624) | −1.626*** (0.420) |
| Period | 0.662*** (0.081) | 0.109*** (0.023) |
| PB | 10.963*** (2.180) | 1.265** (0.507) |
| EP | 15.626*** (2.935) | 3.296*** (0.562) |
| HOV | 0.380 (2.384) | −0.573 (0.554) |
| Period*PB | 0.152 (0.114) | −0.007 (0.028) |
| Period*EP | −0.610*** (0.186) | −0.115*** (0.033) |
| Period*HOV | −0.078 (0.135) | −0.021 (0.033) |
| PB*HOV | −14.682*** (2.915) | −1.689** (0.0802) |
| EP*HOV | −8.689** (3.583) | −1.829** (0.721) |
| Period*PB*HOV | −0.216 (0.166) | −0.036 (0.046) |
| Period*EP*HOV | 0.581*** (0.218) | 0.076* (0.044) |
| logSigmaMu | 2.672*** (0.023) | |
| logSigmaNu | 2.682*** (0.008) | |
| sigma | | 3.145*** |
| Observations | 2,200 | 2.200 |
| Right Censored | 1150 | |
| Log Likelihood | −4,895.628 | −974.370 |

*Note:* *p<0.1; **p<0.05; ***p<0.01

The normalized expected profit regression is estimated with the R package 'censreg' by Henningsen (2017) that implements Butler and Moffitt (1982) method. For the binary panel regression, we use the R package 'pglm' by Croissant (2017).

## B.5 Additional Graphs

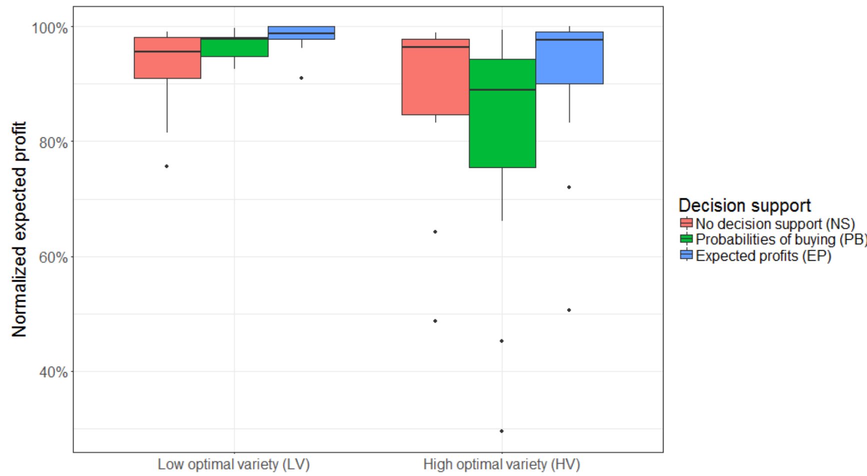


Figure B.1: Normalized expected profits per subject with quartiles
*Note: Whiskers at 1.5 interquartile range*

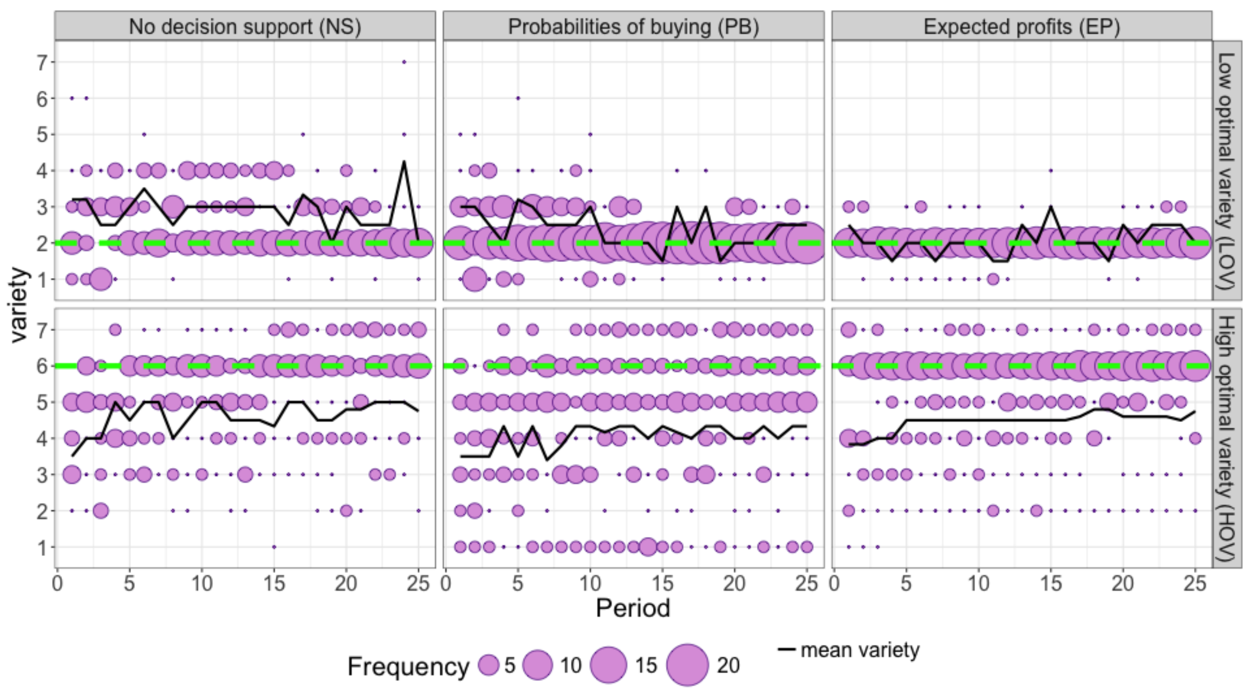


Figure B.2: Assortment sizes over time. Both popular and nonpopular assortments are included

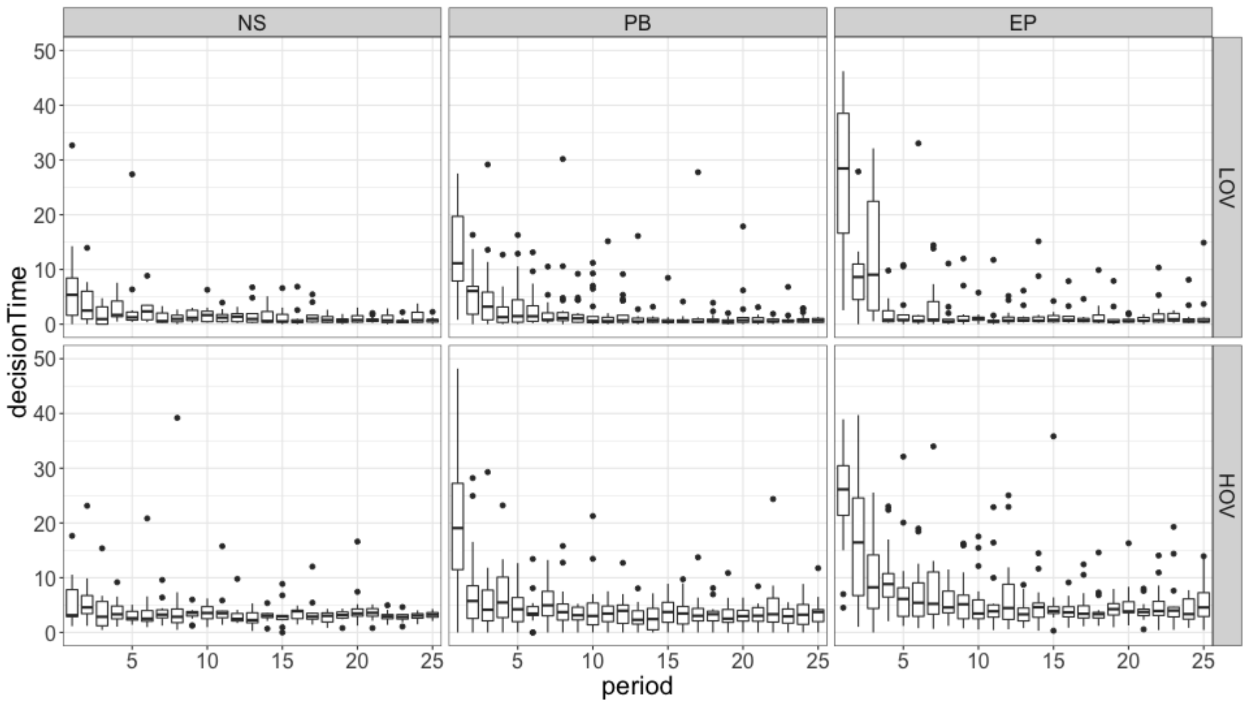


FIGURE B.3: Time between the first and the last box clicks, seconds
*Note: some outliers are omitted due to scaling*

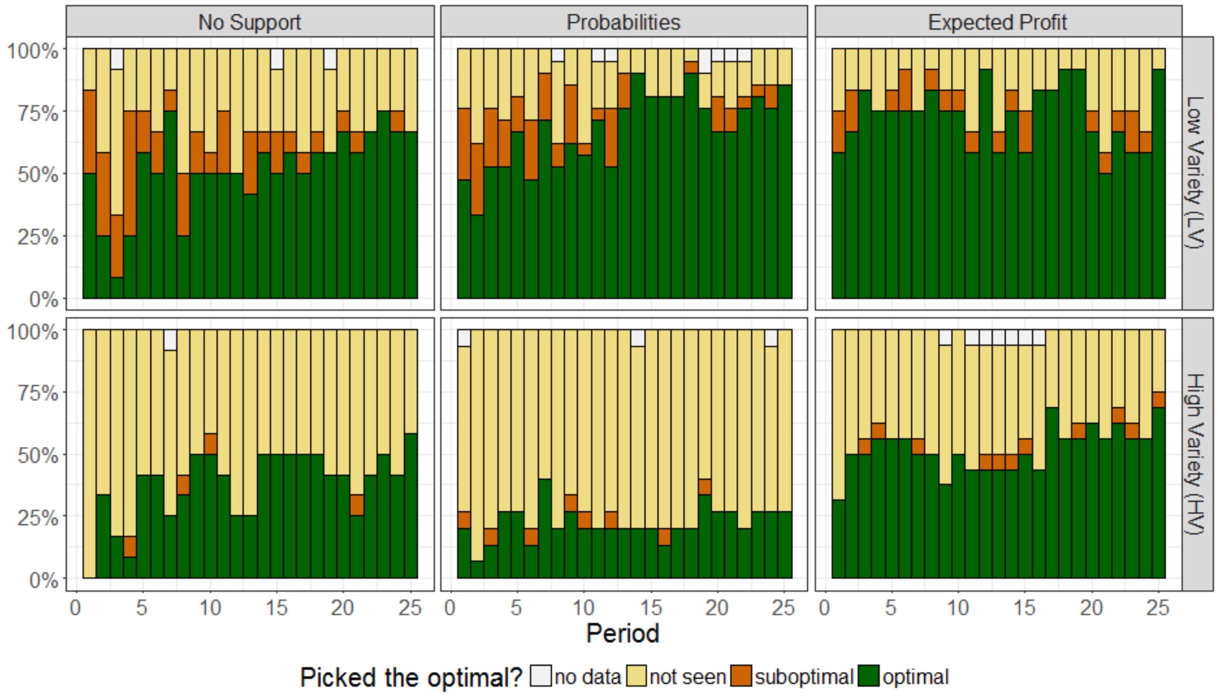


FIGURE B.4: Solution type conditional on the optimal assortment being selected at some point during the period

# APPENDIX C

# APPENDIX FOR CHAPTER 4

## C.1 Proof or Theorem 1

Since this proof is concerned only with the part of the consumer beliefs evolution starting from the time $\tilde{\tau}$, for sake of notation simplicity we will assume here that $\tilde{\tau} = 0$, which implies $Q_i = Q_j = Q$ and $p_i > p_j$.

Since product $j$ is never during the first $\tau_1$ periods, consumer's beliefs on product $j$ at time $\tau_1$ are characterized by $R_{j0} = Q$. Observe that $\Delta_j - \Delta_i < 1$, so we get $R_{j0} = \left(\frac{q_G}{1-q_G}\right)^{-(\Delta_j - \Delta_i)} R_{i0} > \psi_i^{-1}(Q-1)$. According to our assumptions product $j$ initially dominates the outside option ($R_{j0} > \underline{R}$) so at period $\tau + 1$ the consumer buys the product $j$.

We will first show that given the starting conditions $R_{i0} > R_{j0}$ and $Y_{i0} = Y_{j0} = Q$, for any $Y_{in} > 0$ and $Y_{jk} > 0$ we have

$$Y_{in} \geq Y_{jk} \iff R_{in} > R_{jk}, \tag{C.1}$$

Since $Y_{in}$ is increasing in $R_{in}$ and $Y_{in}$ is increasing in $R_{in}$, we only need to show that $Y_{in} = Y_{jk} \to R_{in} > R_{jk}$ and $Y_{in} = Y_{jk} - 1$ implies $R_{in} < R_{jk}$. Let $Y_{in} = Y_{jk} = z$ . Then $R_{in} = \psi_i^{-1}(Q-z)$ and $R_{jk} = \psi_j^{-1}(Q-z)$, and

$$\frac{R_{in}}{R_{jk}} = \left(\frac{q_G}{1-q_G}\right)^{z-Q} R_{i0} \times \left(\frac{q_G}{1-q_G}\right)^{Q-z} R_{j0}^{-1} = \frac{R_{i0}}{R_{j0}} > 1$$

Now assume that $Y_{jk} = z$ and $Y_{in} = z - 1$. Then

$$\frac{R_{in}}{R_{jk}} = \left(\frac{q_G}{1-q_G}\right)^{-1} \frac{R_{i0}}{R_{j0}} = \left(\frac{q_G}{1-q_G}\right)^{\Delta_j - \Delta_i - 1} < 1,$$

where the inequality follows from $\frac{q_G}{1-q_G} > 1$ and $\Delta_j - \Delta_i - 1 < 0$.

From (C.1) it follows that when the consumer switches to the product $i$ (or $j$) in period $t$, her beliefs on this product being $Y_{in}$ ($Y_{jn}$) she will switch to the other product or, if the other product is already abandoned, to the outside option, at time $t + \tau_k$, where $\tau_k = \min\{l \geq 0 : Y_{i(n+l)} = Y_{in} - 1\}$ (or, if the product $j$ is being buy d, $\tau_k = \min\{l \geq 0 : Y_{j(n+l)} = Y_{jn} - 1\}$)

We can find that the number of unsatisfactory experiences gets larger than the number of satisfactory ones after $\tau_k = 2m + 1$ periods of successive sampling by substituting $Q = 1$ into (4.19) to obtain (4.28).

Note that since each interval of successive sampling ends with a random walk being one step closer to the origin, the consumer switches to the outside option after the interval $\tau_{2Q}$. Therefore, if $\sum_{k=1}^{2Q} \tau_k = t < T$, then at the period $t$ the consumer abandons the firm and switches to the outside option.

## REFERENCES


Alptekinoğlu, A. and J. H. Semple (2016). The exponomial choice model : A new alternative for assortment and price optimization. *Operations Research 64*(1), 79–93.

American Psychological Association (2002, amended June 1, 2010 and January 1, 2017)). Ethical principles of psychologists and code of conduct. Retrieved from `http://www.apa.org/ethics/code/index.aspx`.

Banks, J. S. and R. K. Sundaram (1992). A class of bandit problems yielding myopic optimal strategies. *Journal of Applied Probability 29*(3), 625–632.

Baquero, G., W. Smit, and L. Wathieu (2013). The generosity effect: fairness in sharing gains and losses (working paper).

Bardsley, N. (2008). Dictator game giving: altruism or artefact? *Experimental Economics 11*(2), 122–133.

Batson, C. D., N. Ahmad, and D. A. Lishner (2009). Empathy and altruism. In S. J. Lopez and C. R. Snyder (Eds.), *Oxford handbook of positive psychology* (2 ed.)., pp. 417–426. Oxford University Press.

Ben-Ner, A., L. Putterman, F. Kong, and D. Magan (2004). Reciprocity in a two-part dictator game. *Journal of Economic Behavior & Organization 53*(3), 333–352.

Bernales, P. J., Y. Guan, H. P. Natarajan, P. S. Gimenez, and M. X. A. Tajes (2017). Less is more: Harnessing product substitution information to rationalize skus at intcomex. *Interfaces 47*(3), 230–243.

Bertoin, J. and R. A. Doney (1994). On conditioning a random walk to stay nonnegative. *The Annals of Probability 22*(4), 2152–2167.

Boatwright, P. and J. C. Nunes (2001). Reducing assortment: an attribute-based approach. *Journal of Marketing 65*(3), 50–63.

Bolton, G., E. Katok, and R. Zwick (1998). Dictator game giving: Rules of fairness versus acts of kindness. *International Journal of Game Theory 27*(2), 269–299.

Butler, J. and R. Moffitt (1982). A computationally efficient quadrature procedure for the one-factor multinomial probit model. *Econometrica 50*(3), 761–764.

Cachon, G. P., C. Terwiesch, and Y. Xu (2005). Retail assortment planning in the presence of consumer search. *Manufacturing & Service Operations Management 7*(4), 330–346.

Camerer, C. F. and E. Fehr (2004). Measuring social norms and preferences using experimental games: A guide for social scientists. In *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Chapter 3, pp. 55–95. Oxford University Press.

Croissant, Y. (2017). pglm: Panel generalized linear models. R package version 0.2-1, `http://CRAN.R-project.org/package=pglm`.

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

Dana, J., D. M. Cain, and R. M. Dawes (2006). What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes 100*(2), 193–201.

Elberse, A. and J. Bergsman (2008). *Radiohead: Music at your own price (A).* Harvard Business School.

Engel, C. (2011). Dictator games: a meta study. *Experimental Economics 14*(4), 583–610.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, 817–868.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Fisher, M. and R. Vaidyanathan (2012). Which products should you stock? *Harvard Business Review 90*(11), 108–+.

Forsythe, R., J. Horowitz, N. Savin, and M. Sefton (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior 6*(3), 347–369.

Gneezy, A., U. Gneezy, L. D. Nelson, and A. Brown (2010). Shared social responsibility: A field experiment in pay-what-you-want pricing and charitable giving. *Science 329*(5989), 325–327.

Gneezy, A., U. Gneezy, G. Riener, and L. D. Nelson (2012). Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences of the United States of America 109*(19), 7236–40.

Goodman, J. K., C. E. Cryder, and A. Cheema (2013). Data collection in a flat world: The strengths and weaknesses of mechanical turk samples. *Journal of Behavioral Decision Making 26*(3), 213–224.

Graham, R. L., D. E. Knuth, and O. Patashnik (1994). *Concrete Mathematics: A Foundation for Computer Science* (2 ed.). Addison-Wesley Professional.

Greene, W. H. (2008). *Econometric analysis* (6th ed.. ed.). Upper Saddle River, N.J.: Pearson/Prentice Hall.

Hendriks, A. (2012). SoPHIE – software platform for human interaction experiments. University of Osnabrück, `http://www.sophie.uni-osnabrueck.de`.

Henningsen, A. (2017). censReg: Censored regression (tobit) models. R package version 0.5-26, `http://CRAN.R-project.org/package=censReg`.

Hoffman, E., K. McCabe, K. Shachat, and V. Smith (1994). Preferences, property rights and anonymity in bargaining games. *Games and Economic Behavior 7*(3), 346–380.

Honhon, D., E. Katok, and Y. Vorotyntseva (2018). Can managers plan assortments? an experimental study. Working paper.

Honhon, D., C. Ulu, and Y. Vorotyntseva (2018). Tell me what I want: assortment planning for learning consumers. Working paper.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics 14*(3), 399–425.

Jagabathula, S. and G. Vulcano (2017). A partial-order-based model to estimate individual preferences using panel data. *Management Science null*(null), null.

Johnson, J. W. and A. P. Cui (2013). To influence or not to influence: External reference price strategies in pay-what-you-want pricing. *Journal of Business Research 66*(2), 275–281.

Katok, E. (2011). Using laboratory experiments to build better operations management models. *Foundations and Trends® in Technology, Information and Operations Management 5*(1), 1–86.

Katok, E. and V. Pavlov (2013). Fairness in supply chain contracts: A laboratory study. *Journal of Operations Management 31*(3), 129–137.

Kim, J.-y., M. Natter, and M. Spann (2009). Pay What You Want: A new participative pricing mechanism. *Journal of Marketing 73*(1), 44–58.

Kök, A. G., M. L. Fisher, and R. Vaidyanathan (2009). *Assortment Planning: Review of Literature and Industry Practice*, pp. 99–153. Boston, MA: Springer US.

Kök, A. G. and Y. Xu (2011). Optimal and competitive assortments with endogenous pricing under hierarchical consumer choice models. *Management Science 57*(9), 1546–1563.

Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *The American Economic Review 90*(4), 1072–1091.

Lynn, M. (1990). Choose your own price – an exploratory study requiring an expanded view of prices functions. *Advances in consumer research 17*, 710–714.

Lynn, M., S. M. Flynn, and C. Helion (2013). Do consumers prefer round prices? evidence from pay-what-you-want decisions and self-pumped gasoline purchases. *Journal of Economic Psychology 36*, 96–102.

Mak, V., R. Zwick, A. R. Rao, and J. A. Pattaratanakun (2015). "Pay what you want" as threshold public good provision. *Organizational Behavior and Human Decision Processes 127*, 30–43.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and economic behavior 10*(1), 6–38.

Miller, H. (1961). A generalization of Wald's identity with applications to random walks. *The Annals of Mathematical Statistics 32*(2), 549–560.

Oxoby, R. J. and J. Spraggon (2008). Mine and yours: Property rights in dictator games. *Journal of Economic Behavior & Organization 65*(3-4), 703–713.

Özer, Ö. and Y. Vorotyntseva (2015). Inequity and loss aversion in Pay-What-You-Want. Working paper.

Özer, O. and Y. Zheng (2012). Behavioral issues in pricing management. In O. Özer and R. Phillips (Eds.), *Oxford Handbook of Pricing Management*, pp. 415—464. Oxford: Oxford University Press.

Özer, Ö., Y. Zheng, and Y. Ren (2014). Trust, trustworthiness, and information sharing in supply chains bridging china and the united states. *Management Science 60*(10), 2435–2460.

Panchanathan, K., W. E. Frankenhuis, and J. B. Silk (2013). The bystander effect in an N-person dictator game. *Organizational Behavior and Human Decision Processes 120*(2), 285–297.

Paolacci, G., J. Chandler, and P. G. Ipeirotis (2010). Running experiments on amazon mechanical turk. *Judgment and Decision Making 5*(5).

Paolacci, G., L. M. Straeter, and I. E. Hooge (2015). Give me your self: Gifts are liked more when they match the giver's characteristics. *Journal of Consumer Psychology 25*(3), 487–494.

Raju, J. and Z. Zhang (2010). Pay As You Wish pricing. In J. Raju and J. Z. Zhang (Eds.), *Smart Pricing: How Google, Priceline, and Leading Businesses Use Pricing Innovation for Profitability* (1st edition ed.)., pp. 19–40. Pearson Prentice Hall.

Ratchford, B. T. (2009). Consumer search behavior and its effect on markets. *Foundations and Trends® in Marketing 3*(1), 1–74.

Riener, G. and C. Traxler (2012). Norms, moods, and free lunch: Longitudinal evidence on payments from a Pay-What-You-Want restaurant. *The Journal of Socio-Economics 41*(4), 476–483.

Roberts, J. H. and J. M. Lattin (1997). Consideration: Review of research and prospects for future insights. *Journal of Marketing Research*, 406–410.

Ross, S. M. (1995). *Stochastic Processes* (2 ed.). Wiley.

Roth, A. E. (1995a). Bargaining experiments. In J. H. Kagel and A. E. Roth (Eds.), *The handbook of experimental economics*, pp. 253–291. Princeton University Press.

Roth, A. E. (1995b). Introduction to experimental economics. In J. H. Kagel and A. E. Roth (Eds.), *The handbook of experimental economics*, pp. 3–110. Princeton university press.

Roth, A. E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir (1991). Bargaining and market behavior in jerusalem, ljubljana, pittsburgh, and tokyo: An experimental study. *The American Economic Review*, 1068–1095.

Ruffle, B. J. (1998). More is better, but fair is fair: tipping in dictator and ultimatum games. *Games and Economic Behavior 23*(2), 247–265.

Schmidt, K. M., M. Spann, and R. Zeithammer (2015). Pay what you want as a marketing strategy in monopolistic and competitive markets. *Management Science 61*(6), 1217–1236.

Schumann, K. and C. S. Dweck (2014). Who accepts responsibility for their transgressions? *Personality and social psychology bulletin 40*(12), 1598–1610.

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

Scopelliti, I., C. K. Morewedge, E. McCormick, H. L. Min, S. Lebrecht, and K. S. Kassam (2015). Bias blind spot: Structure, measurement, and consequences. *Management Science 61*(10), 2468–2486.

Sefton, M. (1992). Incentives in simple bargaining games. *Journal of Economic Psychology 13*, 263–276.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research 16*(2), 158–174.

Simonson, I. and A. Tversky (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing 29*(3), 281–295.

Stahl, D. O. and E. Haruvy (2006). Other-regarding preferences: Egalitarian warm glow, empathy, and group size. *Journal of Economic Behavior & Organization 61*(1), 20–41.

Van Der Hofstad, R. and M. Keane (2008). An elementary proof of the hitting time theorem. *The American Mathematical Monthly 115*(8), 753–756.

van Hoek, R. and K. Pegels (2006). Growing by cutting SKUs at Clorox. *Harvard Business Review 84*(4), 22.

van Ryzin, G. and S. Mahajan (1999). On the relationship between inventory costs and variety benefits in retail assortments. *Management Science 45*(11), 1496–1509.

## BIOGRAPHICAL SKETCH

Yulia Vorotyntseva was born in Voronezh, Russia. She received her Bachelor's degree in Management with Marketing concentration from Voronezh State University. She obtained a Master's degree in Applied Mathematics and Computer science from National Research University Higher School of Economics in Moscow, Russia. She joined the Ph.D. program in Operations Management at The University of Texas at Dallas (UTD) in August 2012. During her pursuit of Ph.D., she earned an M.S. degree in Supply Chain Management from UTD.

CURRICULUM VITAE

# Yulia Vorotyntseva

The University of Texas at Dallas
Jindal School of Management

Office: JSOM 3.228
Email: yulia.vorotyntseva@utdallas.edu

## Research Interests

Assortment planning, Behavioral Operations Management, Bayesian Learning Models, Discrete Choice Modeling, Stochastic Dynamic Programming

## Education

Ph.D., Management Science, 2018
Operations Management Concentration, GPA 3.88/4.00
The University of Texas at Dallas, Richardson, TX
Dissertation title: *Essays on individual choice modeling and analysis*
Advisors: Elena Katok and Dorothée Honhon

M.S., Supply Chain Management (with Distinction), GPA 3.88/4.00 2016
The University of Texas at Dallas, Richardson, TX

M.S., Applied Mathematics and Computer Science, 2012
Mathematical Modeling Concentration (with Distinction), GPA 9.21/10.00
Higher School of Economics (National Research University), Moscow, Russia
MS thesis title: *Comparative Analysis of Mathematical Models and Developing the Structure of a Decision Support System for Planning Advertising Campaigns*
Advisor: Alexander Belenky

B.S., Management (with Distinction), GPA 4.97/5.00 2010
Voronezh State University, Voronezh, Russia

## Working Papers

Can Managers Plan Assortments? An Experimental Study *(with Katok and Honhon)*

Tell Me What I Want: a Study of Personalized Assortment Planning *(with Honhon and Ulu)*

Effect of Private Valuation in Pay What You Want *(with Özer)*

## Conference Presentations

Can Managers Plan Assortments? An Experimental Study

*POMS 29th Annual Conference*, May 2018

*48th Annual Meeting of the Decision Sciences Institute*, November 2017

*INFORMS Annual Meeting Houston*, October 2017

*Young Scholars Workshop at 12th Annual Behavioral Operations Conference*, July 2017

Tell Me What I Want: a Study of Personalized Assortment Planning

*48th Annual Meeting of the Decision Sciences Institute*, November 2017
*INFORMS Annual Meeting Houston*, October 2017
*INFORMS Annual Meeting Nashville*, November 2016
*POMS 27th Annual Conference*, May 2016

Effect of Private Valuation in Pay What You Want

*INFORMS Annual Meeting Philadelphia*, October 2015

Inequity and Loss Aversion in Distribution Games wit Negative Payoffs:
The Case of Pay-What-You-Want

*Young Scholars Workshop at 10th Annual Behavioral Operations Conference*, July 2015
*POMS 26th Annual Conference*, May 2015

## Teaching Experience

**Instructor** UT Dallas, Spring 2016
Operations Management (*Undergraduate Core*). Enrollment: 45
Instructor Evaluation: 4.58 / 5.00

**Teaching Assistant** UT Dallas, 2012-2016

*Undergraduate courses:*
Operations Management
Managerial Methods in Decision Making under Uncertainty
Quantitative Business Analysis

*Doctoral courses:*
Stochastic Dynamic Programming

*MS/MBA courses:*
Operations Management
Quantitative Introduction to Risk and Uncertainty in Business
Retail Operations
Project Management
Product Lifecycle Management
Supply Chain Strategy

## Industry Experience

**Lead community manager** 2009
Astrum Nival LLC, Voronezh, Russia
*Industry:* Video games publishing
*Duties:* Supervised a newly established community team. Participated in hiring activities, trained new employees, coordinated activities with upper management.

**PR-manager** 2007-2008
Skyfallen Entertainment LLC, Voronezh, Russia
*Industry:* Video games development
*Duties:* Prepared marketing materials, moderated game forums, published press-releases, wrote game dialogues, maintained corporate website, coordinated marketing activities with the publisher.

**Sales representative** 2006-2007
Oasis-s LLC, Voronezh, Russia
*Industry:* IBM sales and service center
*Duties:* created and maintained a database of existing and prospective customers, reached out prospective buyers through cold calling, dealt with customers' problems and complaints.

## Selected Coursework

**UT Dallas**
Behavioral Operations Management
Consumer Search Models
Deterministic Models in Operations Research
Econometrics
Game Theory
Optimal Control
Pricing Management (PhD Seminar)
Probability and Stochastic Processes
Statistical Inference
Stochastic Models in Operations Research
Supply Chain Strategy

**NRU HSE**
Convex Analysis
Data Analysis
Data Mining Methods and Systems
Decision Analysis and Support
Discrete Mathematics
Models of Utility Theory
Non-classical Logics
Ordered Sets in Data Analysis
Stochastic Modeling

## Other Activities

Post Proposal PhD Student Consortium at the annual DSI meeting, 2017

MIT Professional Education's program
*Discrete Choice Analysis: Predicting Demand and Market Shares*, 2017

INFORMS Business Analytics Professional Colloquium, 2017

Doctoral Student Colloquium at the INFORMS Annual Meeting, 2015

PhD Summer Academy at the MIT Zaragoza Logistics Center, 2013

## Skills

*Languages:* English (fluent), Russian (native)

*Computing Packages:* Git, Mercurial, Matlab, R (extensive experience); SAS, STATA (as a part of coursework)

*Specialized Software:* zTree, SoPHIE with related PHP and AngularJS

## Honors and Awards

Betty and Gifford Johnson Travel Award, 2017

Half-tuition Scholarship from MIT Discrete Choice Analysis Program, 2017

Women in OR/MS Professional Colloquium Award, 2017

UT Dallas Dean's Excellence Scholarship, 2017

Women in OR/MS Monsanto Student Travel Award Women in OR/MS, 2016

UT Dallas Graduate Studies Scholarship, 2012- 2016

## Professional Associations

The Institute for Operations Research and Management Sciences (INFORMS)

Behavioral Operations Management (BOM) Section of INFORMS

The Forum for Women in OR/MS (WORMS)

Production and Operations Management Society (POMS)

Decision Sciences Institute (DSI)

## References

Elena Katok
Ashbel Smith Professor of Operations Management
Naveen Jindal School of Management
The University of Texas at Dallas
Tel: (972) 883-4047
Email: `ekatok@utdallas.edu`

Dorothée Honhon
Associate Professor of Operations Management
Naveen Jindal School of Management
The University of Texas at Dallas
Tel: (972) 883-5947
Email: `dorothee.honhon@utdallas.edu`

Canan Ulu
Associate Professor
Operations and Information Management Area
McDonough School of Business
Georgetown University
Tel: (202) 687 6438
Email: `canan.ulu@georgetown.edu`