Essays in Industrial Economics: Applications for Marketing and Management Decisions

Kohei Kawaguchi

k.kawaguchi@lse.ac.uk

Thesis submitted to

Department of Economics
London School of Economics and Political Science

for the award of the degree of

Doctor of Philosophy

in Economics

June 2015

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work except where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified).

 I declare that my thesis consists of 29,512 words.

Statement of Conjoint Work

I confirm that Chapter 2 was jointly co-authored with Professors Yasutora Watanabe and Kosuke Uetake and, all authors equally contributed to this work at every stage of the project. We jointly requested permission to conduct field experiments with the associated company, designed the experiment, analyzed the results, and wrote the paper.

Acknowledgment

I am greatly indebted to my supervisor Martin Pesendorfer for giving me all possible support. I would also like to thank my advisors Alessandro Gavazza and Pasquale Schiraldi their enthusiasm for my work. Particular thanks to my coauthors Yasutora Watanabe and Kosuke Uetake for their comments and advice. I enjoyed working with them on the second chapter of my thesis.

Many other researchers have provided feedback on several sections of the thesis: Mark Schankerman, John Sutton, Matthew Gentry, Pedro Ortega, Ricardo Alonso, Andrew Ellis, Ronny Razin, Taisuke Otsu, Hidehiko Ichimura, Toru Kitagawa, Mitsukuni Nishida, Ken Yamada, Yuichiro Waki as well as the seminar participants and speakers at the London School of Economics and other various conferences.

I would also thank Takayuki Ishidoya, Toshinari Sasagawa, Takatoki Izumi, Yukinori Kurahashi, Tomoki Mifune, and many others at JR East Water Business including the operations' staff for their support and insightful conversations throughout the project described in the first and second chapters. Daisuke Adachi, Masao Fukui, Takanori Nakagawa, and Yuta Yamauchi provided excellent research assistance.

Thanks to my PhD flatmates, Munir Squir, Sergio de Ferra, Thomas Carr, and other PhD students at the London School of Economis for making the PhD experience so enjoyable. I thank the PhD students with whom studied at the University of Tokyo but now study at other universities, Daisuke Hirata, Ken Onishi, Kentaro Tomoeda, Mari Tanaka, Takuya Ura, and Yusuke Kasuya. I thank Junichi Yamasaki for sharing the PhD experience at the University of Tokyo and the London School of Economics.

Finally, I thank Mikiko for sharing my exciting and tough moments as my partner. Thanks for always believing in me.

Abstract

IT innovation is allowing enterprises to find new ways to harness the power of information assets for decision making. This thesis presents three econometric method applications to marketing and management decisions.

The first chapter empirically studies retail network product assortment decisions under uncertain underlying demand parameters using structural estimation. I use detailed data from a beverage vending machine network in Tokyo and find that agents increase the expected total revenue of the network by 19.6% than the baseline, where 12.3% is attributable to learning from the sales data, and 7.3% is attributable to agents' informative initial belief. However, it is below the revenue when the demand parameters are known, which is 45.5% higher than the baseline. Furthermore, if the principal company could precisely process the sales data, the expected total revenue could be 39.6% higher even if the initial beliefs are no more informative than the rational expectation. The last observation indicates that there are some costs for the principal associated with the development and utilisation of sales data processing capabilities.

The second chapter studies the causal effects of product recommendation by conducting a field experiment using many vending machines in railway stations that programmatically offer recommendations for consumers after recognising their characteristics via a built-in camera. We study the effects of recommending popular products and unpopular products, and ask how the effects differ across times of day and consumer characteristics. We find that both popular and unpopular product recommendations increase vending machine sales and choice probability of recommended products. But unpopular product recommendations cause opposite effects in the morning. The negative effects are mainly from male customers in crowded vending machines. We attribute the decrease in morning vending machine sales to the congestion created by recommendations. We conjecture that the negative effect on choice probabilities in the morning is because of social pressure from the surrounding consumers.

In the third chapter, I derive a necessary condition for stochastic rationalisability by a set of utility functions with a unique maximiser, which I name the strong axiom of revealed stochastic preference (SARSP). I propose a test of rationality based on the SARSP that allows for any type of heterogeneity. The test can be implemented at low computational cost. Monte Carlo simulation shows that the test has an empirical size below the nominal level and relatively strong power.

Contents

List of Tables

List of Figures

Introduction

IT innovation is allowing enterprises to find new ways to harness the power of information assets for decision making. The empirical frameworks developed in industrial economics are particularly useful for this purpose. The frameworks allow managers to elicit information for counterfactual analysis from observational data possibly with the help of field experiments. The counterfactual analysis then allows managers to evaluate new policies for effective decision making.

This thesis describes three such attempts. Each attempt uses different empirical frameworks. Chapter 1 uses structural estimation to quantify how quickly agents in a retail network can learn and adapt to local demand and to detect the sources of friction preventing immediate adjustment. Chapter 2 uses a field experiment to study the causal effects of product recommendations in various situations. Chapter 3 suggests a rationalisability test that does not restrict preference heterogeneity across decision makers. Once data passes the test, the rationality assumption is used to obtain non-parametric bounds of the decision makers' response in a counterfactual situation.

Decentralised Learning in a Retail Network

The first chapter empirically studies retail network product assortment decisions in an uncertain underlying demand parameters environment. Specifically, I quantify how quickly beverage vending machine network agents can adapt to local demand. I also study the source of friction preventing immediate adjustment. I develop and estimate an empirical model of decentralised learning in a retail network in which agents assigned to local markets learn demand and decide part of product assortments. The modelling strategy follows the empirical structural analysis of learning literature,[1] but is extended to address the combinatorial nature of product assortment decisions by borrowing a boundedly rational solution concept from engineering literature. Using the estimated model, I find that agents increase expected total network revenue by 19.6% than the baseline, where 12.3% is attributable to learning from the sales data, and 7.3% to agents' informative initial belief. However, it is far below the expected total revenue when the demand parameters are known, which is 45.5% higher than the baseline. Moreover, if the principal company could process the sales data precisely, the expected total revenue could be 39.6% higher than the baseline even if the initial beliefs are no more informative than the rational expectation. The last observation indicates that there are some costs associated with the development and utilisation of sales data processing capabilities that do not rely on agents.

Although vast literature solves and characterises similar decision problems for firms that are actively learning demand,[2] actual firm behaviors are not extensively examined empirically. Among retailers' key managerial decisions, product assortment decisions receive less attention than price adjustments. However, product assortment decisions are is equally important because product selection and replacement is prevailing because of limited shelf space for product display and various motivations for product proliferation.[3] Ignoring this factor can cause significant bias in the calculation of price indexes (Nakamura and Steinsson, 2012). This chapter address the gap in the literature.

The retail network in this study partly delegates the task of learning and adjusting product assortments to the local demand to agents assigned to local markets. Organizational economics literature studies the implications of managers' bounded capacities on the advantages of decentralisation.[4] One strand of this literature considers managers that are bounded by the amount of information they possess (Groves and Radner, 1972; Arrow and Radner, 1979; Groves, 1983; Geanakoplos and Milgrom, 1988). In this case, decentralisation benefits the organisation by providing more information for solving managerial problems. The second strand of such literature distinguishes between the raw data and the processed data and assumes that managers

[1] For example, Jovanovich (1979); Miller (1984); Erdem and Keane (1996); Ackerberg (2003); Crawford and Shum (2005); Dickstein (2014).

[2] Early theoretical literature includes Prescott (1972); Grossman, Kihlstrom, and Mirman (1977); Trefler (1993); Mirman, Samuelson, and Urbano (1993); Harrington (1995); Keller and Rady (1999). The literature in management science and operations research include Caro and Gallien (2007); Rusmevichientong, Shen, and Shmoys (2010); Honhon, Gaur, and Seshadri (2010); Talebian, Boland, and Savelsbergh (2013).

[3] Bayus and Putsis (1999) empirically study the determinants of product proliferation, and provide an extensive survey of related literature.

[4] The survey is found in Garicano and Van Zandt (2013).

are bounded by what they can achieve in a certain amount of time (Malone and Smith, 1988; Mount and Reiter, 1990, 1996; Radner, 1993; Bolton and Dewatripont, 1994; Friedman and Oren, 1995; Reiter, 1996; Meagher and Van Zandt, 1998; Van Zandt, 1998; Beggs, 2001; Orbay, 2002). In this case, decentralisation allows the organisation to solve subproblems concurrently. My finding provides an evidence that decentralisation allowing the simultaneous solving of subproblems is important for the decision structure of the retail network.

This chapter contributes to several other strands of literature. First, this chapter contributes to the empirical structural analysis of learning literature (Jovanovich, 1979; Miller, 1984; Erdem and Keane, 1996; Ackerberg, 2003; Crawford and Shum, 2005; Dickstein, 2014; Covert, 2013) by developing empirical learning models to analyse complicated situations for which existing models do not work. Second, boundedly rational solution are increasingly used in the field of industrial organisation to solve complicated empirical models (Weintraub, Benkard, and Van Roy, 2008; Fershtman and Pakes, 2012). My modelling strategy is consistent with this trend. Third, introducing agents who are learning the environment under uncertainty but with limited information acquisition and processing capability is becoming common in macroeconomics (Mankiw and Reis, 2002; Sims, 2003; Woodford, 2009; Angeletos and La'O, 2009). This paper provides a motivating example for this assumption and an estimate for the appropriate degree of informational friction in a firm. Finally, establishing an empirical framework quantifying these issues based on managerial information has direct implications for organisational learning literature in management and marketing science (Easterby-Smith and Lyles, 2011).

Structural estimation is central to this analysis. First, I am interested in counterfactual analyses that change model parameters. A counterfactual analysis can only be conducted if a structural model is set up for which the parameter of interest has a unique economic interpretation. Second, I incorporate institutional restrictions imposed by the principal company on agents. These restrictions can be explicitly incorporated if the model is constructed at the raw level of the agent decision problem. The estimated model helps managers to identify the strengths and weaknesses of each agent and each aspect of the organisation and recommends appropriate measure for building an adaptive and responsive organisation.

Recommending (Un)popular Products: A Field Experiment using Vending Machines

Product recommendation is a major online and offline marketing tool used by retailers to influence customer choices. To measure the causal effects of product recommendation, we conduct a field experiment using many vending machines in railway stations that programmatically offer recommendations for consumers after recognising their characteristics via a built-in camera. We study the effects of recommending popular products and unpopular products, and ask how the effects differ across times of day and consumer characteristics. We find that both popular and unpopular product recommendations increase vending machine sales and choice probability of recommended products. But unpopular product recommendations cause opposite effects in the morning. The negative effects are mainly from male customers in crowded vending machines. We attribute the decrease in morning vending machine sales to the congestion created by recommendations. We conjecture that the negative effect on choice probabilities in the morning is because of social pressure from the surrounding consumers.

Our research contributes mainly to the small but growing literature on product recommendations and consumer choices.[5] An early work by Senecal and Nantel (2004) conducts a series of online experiments and finds that recommendations significantly affect demand. In behavioral work, Huang and Chen (2006) also find that other consumers' responses influence the choices of subjects. Using aggregate-level data, a recent paper by Kim, Albuquerque, and Bronnenberg (2010) estimate a sequential consumer search model for Amazon.com and simulate the effects of Amazon.com's recommendation system on consumer search behaviour. The authors find that the consumers benefit from recommendations because of the lower search cost. De, Hu, and Rahman (2010) use a server log file of an online company to uncover the relationship between consumers' recommendation system usage and online sales. The authors find that recommen-

[5]Increasing literature examines online advertisements. Sahni (2015) explores the effect of advertisement exposure frequencies on consumer purchase decisions using online randomized control field experiments. The author finds that purchase likelihood increases if advertisements are spread apart rather than bunched together. Lewis, Rao, and Reiley (2014) conduct a large field experiment, where they randomise display advertising. The authors' findings indicate that the effect of online display advertisements on offline sales is small. Goldfarb and Tucker (2011) use the data from an online field experiment and find that matching an advertisements to website content and increasing an advertisement's obtrusiveness increase a consumer's purchase intent. Blake, Nosko, and Tadelis (2015) implement a series of online field experiments for eBay and show heterogeneous effects of paid search advertisements on consumer behavior.

dations have positive impact on sales. Similarly, Tucker and Zhang (2011) investigate the effect of popularity information using a field experiment and find that ranking information affects consumer choice. Finally, Bodapati (2008) proposes that recommendations based on a higher sensitivity to recommendations are more effective than recommendations based on a higher probability of purchase. This is because consumers will purchase the products with a higher predicted purchase probability regardless of recommendations.

This study is also related to the consumer demand literature for the vending machine industry. Anupindi, Dada, and Gupta (1998) use data from multiple beverage vending machines in a US city to examine substitution patterns of beverage demand when a stock-out occurs. The authors find significance in the use of stock-out information to infer consumer demand. In a related paper, Conlon and Mortimer (2008) conduct a field experiment using snack vending machines and show that demand estimation is biased if product availability is not taken into account. We add to the literature by showing evidence of market intervention effectiveness in the vending machine industry because the existing research does not study the effects of the marketing mix.

The combination of field experiments and a large dataset about consumer behaviour analysis is central to this study. The field experiment creates exogenous variations in the status of product recommendations and allows us to understand otherwise unidentified causal effects of recommendations. The large dataset allows us to obtain accurate estimates on the causal effects and to further investigate the heterogeneity in the effects. Therefore, managers can tailor marketing instruments to the business environment of their organisations.

Testing Rationality Without Restricting Heterogeneity

The rationality assumption forms the core of economics. Testing the assumption is necessary to validate the empirical analysis. The classical axioms of revealed preference (Samuelson, 1938; Houthakker, 1950; Richter, 1966; Varian, 1982) provide the basic framework to test this assumption but do not address the inevitable heterogeneity in empirical studies. This chapter proposes a non-parametric test for rationality allowing for any type of preference heterogeneity across decision makers. I derive a necessary condition for stochastic rationalisability by a set of utility functions with a unique maximiser, which I name the strong axiom of revealed stochastic preference (SARSP). The test I propose is based on this condition. The test can be implemented at low computational cost. Monte Carlo simulation shows that the test has an empirical size below the nominal size and relatively strong power. I apply this method to the British Family Expenditure Survey (FES). The SARSP is not rejected at the 1% level of significance in this dataset.

Existing empirical studies conventionally introduce heterogeneity as an additive error in the choice function, and apply the axiom of revealed preference to the mean of the observed choice functions (Varian, 1985; Blundell, Browning, and Crawford, 2003, 2008). Lewbel (2001), however, shows that rationalisability at the individual level does not imply the rationalisability of the mean or vice versa, unless the heterogeneity in the population is strictly restricted.

McFadden (2005) introduces a notion of stochastic rationalisability. Given a set of utility functions $\mathcal{U}$, a sequence of observed choice probabilities is said to be stochastically $\mathcal{U}$-rationalisable if some probability law over $\mathcal{U}$ exists that can induce the observed choice probabilities as a result of utility maximization with random utilities following that law. The author also derives necessary and sufficient condition for the stochastic rationalisability. This notion allows us to test rationality without restricting heterogeneity. The concept is sufficiently general that it can be applied to any class of utility functions; however, checking all the conditions is often computationally demanding. Kitamura and Stoye (2013) circumvent this problem by restricting attention to a case with a finite number of linear budgets. I resolve this issue by focusing on a necessary condition for rationalisability.

Once data passes the test of rationalisability, I can use this restriction to obtain non-parametric bounds of decision makers' response in a counterfactual situation. For example, Manski (2014) uses the weak version of the axiom of stochastic revealed preference to obtain bounds on labour supply under hypothetical income tax policies. The advantage of this approach compared to structural estimation is that I can directly predict counterfactual responses from reduced-form parameters — choice probabilities — bypassing structural parameter estimation.

Chapter 1

Decentralised Learning in a Retail Network

1.1 Introduction

It is crucial for retail networks such as supermarkets and convenience stores to adjust prices and product assortments to adapt to changing demand. If uncertainty exists concerning underlying demand parameters, the adjustment process requires firms to actively learn demand. Such uncertainty obviously exists in a new market, but is is also significant in established markets where new products frequently enter or there is a non-trivial time-varying component in the demand parameters.

Although vast literature attempts to solve and characterize such decision problems for firms that are actively learning demand,[1] actual firm behaviors have not been extensively examined empirically. Among retailer's managerial decisions, much less attention is paid to product assortment decisions than price adjustments, but product assortment decisions are equally important because product selection and replacement prevails because of shelf space limitations for product display and various motivations for product proliferation.[2] Ignoring this managerial

[1] Early theoretical literature includes Prescott (1972); Grossman et al. (1977); Trefler (1993); Mirman et al. (1993); Harrington (1995); Keller and Rady (1999). The literature in management science and operations research include Caro and Gallien (2007); Rusmevichientong et al. (2010); Honhon et al. (2010); Talebian et al. (2013).

[2] Bayus and Putsis (1999) empirically study the determinants of product proliferation and provide an extensive survey of related literature.

behavior can cause a significant bias in the calculation of price indexes (Nakamura and Steinsson, 2012).

Understanding how good estimates are that firms use for underlying demand parameters and how quickly firms can learn and adapt to demand is a prerequisite in any investigation of the implications of changes in market institutions and the demand environment using high-frequency data. Detecting the sources of possible friction in this adjustment process is necessary for long-run prediction in which the technological environment can change the degree of friction. In the case of a retail network, for example, the network may struggle to collect, process, and analyze demand to derive optimal plans or implement the resulting assortment plans. All the tasks are carried out by a different set of workers and technologies with the organization and will be affected differently by technological progress.

This paper focuses on product assortment decisions of a beverage vending machine network in Tokyo. In the network, agents assigned to different local markets partially decide product assortments. I develop and estimate an empirical model of decentralised learning in a retail network. The modelling strategy follows the literature on the empirical structural analysis of learning,[3] but is extended to manage a problem caused by the combinatorial nature of product assortment decisions. Using the estimated model, I quantify how quickly agents in the network can adapt to local demand. I also study the source of friction preventing immediate adjustment.

The beverage vending machine business provides an ideal setting for the analysis of product assortment decisions because a vending machine has a physically well-defined set of slots in which products are assorted. Moreover, in this industry, prices are conventionally the same over time and across locations. Therefore product assortment is particularly important for raising profits. New products are introduced every season and learning demand on a seasonal basis is crucial. All of the network vending machines are installed in Japan Railway East (JRE) stations that have substantial daily passenger traffic and approximately ten times more sales per machine than other vending machine companies. This encouraged the network to introduce a point-of-sales system to monitor sales at the micro level in 2009 for the first time in this industry and significantly boosted the capability of the network to observe the demand data.[4] Additionally, the network provided me extensive administrative information that allows me to observe who makes which decisions. This administrative information facilitates the

[3]For example, Jovanovich (1979); Miller (1984); Erdem and Keane (1996); Ackerberg (2003); Crawford and Shum (2005); Dickstein (2014).

[4]By this time at most, weekly data collected by hand from operators represented the best sales information in this industry.

identification of parameters governing the decision of each agent instead of the joint decision of the principal and agents.

The empirical model treats each local market and agent separately. All the parameters are market- and agent-specific. Therefore, cross-sectional heterogeneity is completely recovered. Consumer's indirect utility from a product in a market is a product- and market- specific intercept plus product- and market-specific slope term multiplied by a log of temperature. These are the demand parameters of a market, and sales are assumed to be generated from a multinomial logistic model based on the demand parameters.

An agent assigned to a particular market does not know the exact value of the demand parameters of the market, but the agent does has an initial belief about them. The initial belief can be arbitrarily informative about the demand parameters. Before a season starts, the agent decides the precision level of his information processing technology, but it is costly to raise the precision. When a season starts, the agent decides the assortment of products for each vending machine in the market based on the agent's belief. Then, the sales are realised and his beliefs are updated according to the prespecified information processing technology. In the next period, the agent again decides the assortment of products.

I use a structural approach for two reasons. First, I am interested in counterfactual analyses that change model parameters. This can be done only if a structural model is set up in which the parameter of interest has a unique economic interpretation. Second, I would like to incorporate institutional restrictions imposed by the principal company on agents. These restrictions can be explicitly incorporated if the model is constructed at the raw level of the agent's decision problem.

Intuitively, I identify the key parameters as follows. The history of assortments reveals the belief process of an agent including his initial belief. From the history of sales under the assortment, I can compute the belief process of a Bayesian that processes the history of sales without cost. By comparing these belief processes, I can identify the precision of an agent's information processing. When I attempt to distinguish the contributions of private information and information processing by agents, I additionally assume that the initial belief of an agent reflects all the relevant private information of the agent.

This analysis is technically challenging. Any serious product assortment problem involves tens or hundreds of products, which renders the state space of the dynamic model — the space of beliefs over demand parameters of each product at each location — extremely high dimensionally. In the current application, the dimensionality exceeds one thousand. Additionally, a product assortment decision is a combinatorial optimisation problem. This first implies that the action space of the model is extremely large. Second, and more importantly, the multi-armed bandit problem becomes a so-called *restless* problem (Whittle, 1988) in which the belief state of an arm can change even when the other arm is taken because updating the belief about a product changes the belief of all the assortments including that product. This implies that efficient algorithms such as the dynamic allocation index (DAI)-based algorithm (Gittins, 1979) can no longer be used to find a fully rational solution that has been used in empirical analyses of the multi-armed bandit problem in economics.[5]

These factors make it impossible to compute a fully rational solution for the model. A fully rational solution is not only intractable but also unrealistic because it demands agents have infinite computational capability (Papadimitriou and Tsitsiklis, 1999). To resolve this problem, the computer science and engineering literature has sought workable and reliable boundedly rational decision rules (Vermorel and Mohri, 2005; Kuleshov and Precup, 2014). I follow this literature and employ a boundedly rational decision rules known as the *Bayesian control rule* (Ortega, 2011). A decision maker following the Bayesian control rule draws a sample from the decision maker's belief at the time and chooses the best alternative assuming that the sample represents the true state of the world. The decision maker mixes state-contingent optimal strategies with his belief about the state of the world. This decision rule is particularly intuitive in the context of product assortment decisions and has sound theoretical and empirical properties.

Once the decision rule and hence the choice probability have been specified, it is straightforward to write down the likelihood function. The agent's belief enters the model as a latent state variable and must be integrated out to evaluate the likelihood. This is achieved by applying a particle filter. Sampling from the posterior distribution uses a robust adaptive Metropolis-Hastings sampler (Vihola, 2011), and the point estimates of parameters are obtained by a maximum a posterior (MAP) estimator. The prior for the static parameters is non-informative, that is, the prior density is constant. So the inference is fully likelihood-based. Demand pa-

[5]One exception is the study of physician's drug choice by Dickstein (2014). The author classifies 19 drugs into six categories and assumes that a physician first chooses a category assuming homogeneity across drugs in the category and then chooses a drug in the category. He solves each step using Gittin's index. This approach, however, does not work in the current case. First, there are still tens of products after dividing products into categories, and there are no further obvious subcategories. Second, the author's approach allows belief updates to be dependent at the category level but does not allow complicated dependence in combinatorial multi-armed bandit problems.

rameters are estimated outside the dynamic model by non-linear least squares.

By simulating a series of sales and assortments from the estimated model, I first find that the expected total revenue under the actual parameter estimates is 19.6% higher than the revenue baseline, where no information processing is exerted, and the initial belief is no more informative than the rational expectation. Assuming that the principal's initial knowledge of the local demand parameters are no worse than the rational expectation, this implies that the upper bound of the agents' contributions to the network's expected total revenue is 19.6% of the revenue baseline. I also find that agents' informative initial beliefs and learning from the sales data explain 7.3% and 12.3% of the contribution, respectively.

However, agents are not perfect. By simulating sales series and assortments from a model with a hypothetical parameter setting, I find that the expected total revenue could be 45.5% higher than the revenue baseline if the demand parameters are known from the beginning. This implies that it is not appropriate to assume that retail companies know the demand parameters, for example, when we recover the marginal costs from their pricing decisions combined with demand models estimated using sales data. This concern is particularly serious when we analyse the short-run response of retailers to changes in the demand environment such as aggregate demand shocks caused by monetary policy or sudden change in taste caused by a product safety accident.

Finally, I compute expected total revenue under a hypothetical centralised assortment policy in which the principal company only uses the sales and temperature data without relying on agents' informative initial beliefs. In theory, the principal company could employ this policy with the current data collection system. If the revenue under this policy exceeds the actual revenue, the indication is that there is a cost associated with the development and utilisation of capabilities to process and analyse the sales data precisely. I find that the centralised policy could increase the expected total revenue by 39.5%. This exceeds the actual expected total revenue. Therefore, I cannot rationalise the delegation without considering the costs for developing and utilising information processing capabilities for the principal company. The costs exist and have to be at least as large as 39.5 - 19.6 = 19.9% of the baseline of gross profits[6] to rationalise the delegation. This means that by targeting the costs of processing information, we can significantly boost the productivity of the organisation although it will remove the task from agents and change the reward distribution.

[6]This is because the gross profits are approximately proportional to the revenue according to a company representative.

1.1.1 Related Literature

This paper is related to the organizational economics literature that studies the implications of managers' bounded capacities on the advantages of decentralization.[7] One strand of such literature considers managers that are bounded in the amount of information they may have (Groves and Radner, 1972; Arrow and Radner, 1979; Groves, 1983; Geanakoplos and Milgrom, 1988). In this case, decentralization benefits the organization by bringing more information for solving managerial problems. Second strand of such literature distinguishes raw data and information processing capability of mangers, and assumes that managers are bounded in what they can do in any amount of time (Malone and Smith, 1988; Mount and Reiter, 1990, 1996; Radner, 1993; Bolton and Dewatripont, 1994; Friedman and Oren, 1995; Reiter, 1996; Meagher and Van Zandt, 1998; Van Zandt, 1998; Beggs, 2001; Orbay, 2002). In this case, the benefit of decentralization mainly comes from the fact it allows the organization to solve subproblems concurrently.

This paper contributes to several other strands of literature. First, it contributes to the literature of empirical structural analysis of learning (Jovanovich, 1979; Miller, 1984; Erdem and Keane, 1996; Ackerberg, 2003; Crawford and Shum, 2005; Dickstein, 2014; Covert, 2013) by developing empirical models of learning to analyze a complicated situation in which existing models do not work. Second, it is increasingly popular in industrial organization to employ boudedly-rational solution concepts for solving complicated empirical models (Weintraub et al., 2008; Fershtman and Pakes, 2012). My modeling strategy is in line with this trend. Third, introducing agents who are learning the environment under uncertainty but with limited information acquisition and processing capability is becoming common in macroeconomics (Mankiw and Reis, 2002; Sims, 2003; Woodford, 2009; Angeletos and La'O, 2009). This paper provides a motivating example for this assumption and an estimate for the appropriate degree of informational friction in a firm. Finally, establishing an empirical framework quantifying these issues based on managerial information has a direct implication for the literature of organizational learning in management and marketing science (Easterby-Smith and Lyles, 2011). It helps managers to identify the strengths and weaknesses of each agent and each aspect of the organization, and guides to the appropriate measure for building an adaptive and responsive organization.

[7] The survey is found in Garicano and Van Zandt (2013).

1.2 Background and Data

1.2.1 Industry Background

Vending machines are important retail channels in Japan particularly for the beverage business. According to the Japan Vending Machine Association,[8] there were 2.6 million installed beverage vending machines at the end of 2013. In 2013, vending machine beverage sales were 2.3 trillion yen. This is approximately one third of the total annual sales of beverages in Japan and is equivalent to supermarket sales.

A unique feature of the industry is that retail prices are constant over time and locations except for unusual locations such as cinemas or mountain tops and only differ across product packages. The nominal price changed only when the consumption tax was raised from 3% to 5% in 1997 and from 5% to 8% in 2014. Between 1997 and 2014, the price of a 350 ml can was 120 yen, and the price of a 500 ml bottle was 150 yen across the industry. Because I use data from 2013, there was no nominal price change during the data period.[9]

This feature makes product assortment decisions — the choice of products to put in a limited number of slots in a vending machine — the only way to increase revenue in this industry. Therefore, beverage producers introduce many new products every season to attract consumers. Learning demand on a seasonal basis is important for retailers.

1.2.2 Institutional Background

JRE is the largest railway company in Japan. Originally part of the national railway, it was privatised in 1987. The company operates in the eastern part of mainland Japan including the Tokyo metropolitan area. In addition to the transportation business, the company operates a retail business inside stations to exploit the substantial number of passengers who regularly use the transportation service.[10] The beverage vending machine business is a branch of that retail business unit.

There are a few types of vending machines installed in JRE stations with between 30 and

[8] `http://www.jvma.or.jp/index.html`

[9] The reason for this price rigidity and homogeneity across locations and products is outside the scope of the paper. Possible explanations for time-invariance include the low inflation in the 2000s in Japan and menu costs. Homogeneity across locations may be attributed to retail price maintenance motivated by fire-sale fears among competitive retailers (Deneckere, Marvel, and Peck, 1996, 1997).

[10] The average number of daily passengers in a station in Tokyo ranges from tens to hundreds of thousands: `http://www.jreast.co.jp/passenger/`

Figure 1.1: A Vending Machine and An Agent

(a) A Vending Machine

(b) An Agent Refilling Products

42 slots. The most common type of machine has 36 slots. Figure 1.1 (a) shows a picture of a typical vending machine. A package of available products is displayed on the upper half of the front panel. Consumers choose the item they want to buy by pushing a button under the mock package. The consumer then inserts cash or touches the sensor with a JRE electric commuter card called SUICA in the middle of the panel. Then, the item drops into the box at the bottom of the machine. Because almost all passengers in Tokyo have this commuter card, a large percentage of transactions are made with the card. The machine and the system are almost identical to the vending machines of other companies outside stations.

JRE has a subsidiary firm managing the company's beverage business. This is the principal company in our analysis. The firm owns vending machines in a number of JRE stations. The principal company outsources maintenance to third-party operating companies.[11] In addition to maintenance, the principal company delegates part of the formal authority related to product assortment decisions for vending machines. Product selection for 70% of slots in the vending machines are directly determined by the principal company, but assortment for the remaining slots is the responsibility of the operating companies. The principal company expects and encourages the operating companies to use these slots for exploring and adapting to local demand. The product assortment decision by the principal company does not target this aspect. The principal company classifies vending machines into 31 types based on machine

[11] A typical vending machine and agent refilling products are shown in Figure 1.1.

type and market volume and sets the products they want to sell such as private-brand products or popular, well-known products such as Coke.

The operations of operating companies are divided into several local markets. A local market, on average, consists of three stations closely located, mostly along the same line with similar demographic characteristics. A local market has, on average, 30 vending machines. A local market is, in principle, operated by a single staff member or a small fixed number of staff during a season. This is the basic unit of decision in the business. Therefore, I call these units agents and focus on their decisions. Figure 1.1 (b) shows a picture of an agent refilling products at a vending machine at one of the stations.

The fees paid to the operating companies are proportional to the sales revenue from the allocated vending machines. Therefore, agents are concerned with top-line sales rather than profits. A representative of the principal company explains that this is because there is little variation in margins except for private brands, whose assortment is completely determined by the principal.

The business plan is determined by the principal company each season. April to September is the spring/summer season, and remaining months are fall/winter. By the beginning of each season, the principal company selects a list of products that consists of approximately 200 brands. Products to be inserted into the vending machines are chosen from this list. In the case of spring/summer 2013 on which I focus, there were 205 listed products, 38% of which were completely new, 9% of which had renewed packages, and the reminder were old items. The plan of the principal company for the spring/summer season is to find the best-selling products on the list at least by the end of June and before the true summer starts.[12]

Products are classified into categories such as green tea, black tea, other tea, coffee, mineral water, tea, sports drinks, and fruit juices. The number of slots allocated to each category at each vending machine is fixed during a season. This is verified as true both from the data and by a company representative. Therefore, the product assortment decision of focus is the choice of products for the fixed slots for each category in each vending machine.

Agents can collect sales data from vending machines when they visit the vending machines and check the record. This is the primary source of information for agents. The principal company has a point-of-sales system and monitors all the purchase records of the entire business area. Each week, the principal company sends a product-level normalised sales score, which is

[12]June is a rainy season in Japan when cold and warm air masses combine over Japan. As the rainy season passes, the true Japanese summer begins with temperatures of more than 35°C and humidity of almost 100%.

defined as the sample mean of the weekly unit sales of a product in a vending machine divided by the average weekly sales volume of the vending machine. This is the secondary source of information. I am not aware of the method for sharing information within an operating company. In the analysis, I allow information to be shared between the operating company and the agent in the same format as between the principal company and the agent.

Agents regularly visit the vending machines from a few times a day to a few times a week to refill stock depending on the market size of the vending machines. Thus, an out of stock vending machine is rare.[13] Because they regularly visit vending machines regardless of product changes, marginal costs of changing products on top of driving, visiting and refilling are not high.

1.2.3 Data

The dataset consists of the sales and product assortment history at the micro-level and administrative information. I focus on the product assortment decisions of 80 agents (assigned to 80 local markets) belonging to four operating companies about 22 products in the tea category across 2,483 vending machines in the Tokyo metropolitan area between March and September (the spring/summer season) 2013 with 2,162,812 purchase records.

The administrative information allows me to determine for each week and each vending machine slot which of the principal company and the agent chose that assortment for that week for that particular vending machine slot. This information is critically important to assess the ability of an agent. Otherwise, I could only identify the ability of the organization as a whole.

The auxiliary data provide the product characteristics such as price, package type, and the product release date. The data also contain the list of products for the season in question. I cite temperature data from the website of the Japan Meteorological Agency including the 10-year daily average temperatures. The principal company and the operating companies regularly check the same temperature data when finalising decisions.

The original sales and assortments data are recorded on a second basis. In the structural estimation, I aggregate them into weekly data because it is sufficient to capture the product assortment decisions variations over time, and it is convenient to compute and estimate the model. A week is defined to start on Monday. The assortment at the end of Thursday of a week is used as the assortment for the week. There are two reasons for this choice: i) the

[13]The principal company has a system to monitor the time of stock-out for each machine. According to this database, the time of stock-out is less than 2% per day even at the worst one percentile of vending machines.

Figure 1.2: Distribution of Average Daily Sales Volume Across Vending Machines

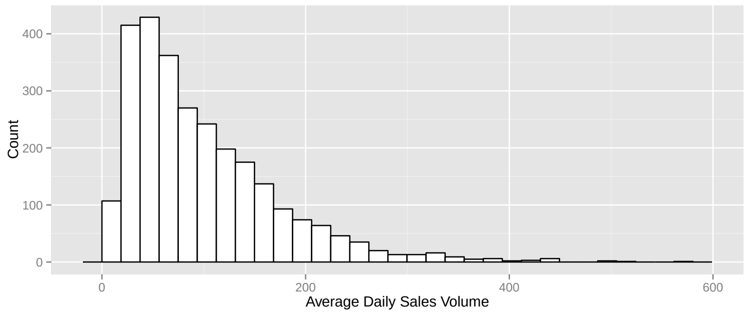

* The bandwidth of the histogram is set at 10. For each vending machine, I compute the daily sales volume and then take the average across days between March and September 2013. The sales of a vending machine include all the categories.

principal company sends the previous week's aggregate sales data on Wednesdays. Therefore, ii) the frequency of product changes is highest on Thursdays. To adjust the difference in days on weekly sales across products, the weekly sales of a vending machine product are defined as the daily product sales times seven. The size of potential consumers of a vending machine is defined as the average weekly sales of the vending machine times one hundred.[14] The choice probability of a product in a vending machine in a week is defined as the weekly sales of the product in the vending machine in the week as defined above per the size of consumers.

1.2.4 Descriptive Analysis

First, I illustrate the demand for the products in the tea category during the spring/summer season in 2013.

Figure 1.2 shows the distribution of the average daily sales volume across vending machines. The median is 78. The sales volume of the vending machines for this retail network is significantly larger than the vending machines outside their railway station network. This allows the network to learn local demand more accurately than other vending machine networks. Figure 1.3a shows the distribution of the total sales volume of each product. The best product sold 580 thousand bottles, whereas the median product sold only 36 thousand bottles. However, this classification does not adjust the days, the number, and the size of vending machine consumers

[14] The definition of the measure of the size of consumers of a vending machine is essentially the same as the measure used by the principal company, although they do not multiply it by one hundred.

Figure 1.3: Sales Volume and Average Choice Probabilities of Products

(a) Total Sales Volume of Products

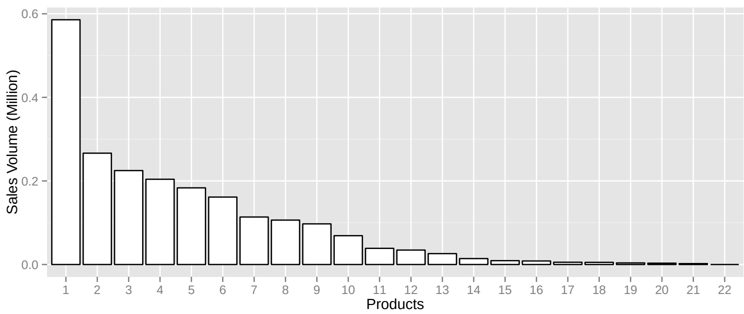

* I compute the total sales of each product between March and September 2013 in the Tokyo metropolitan area.

(b) Average Choice Probabilities of Products

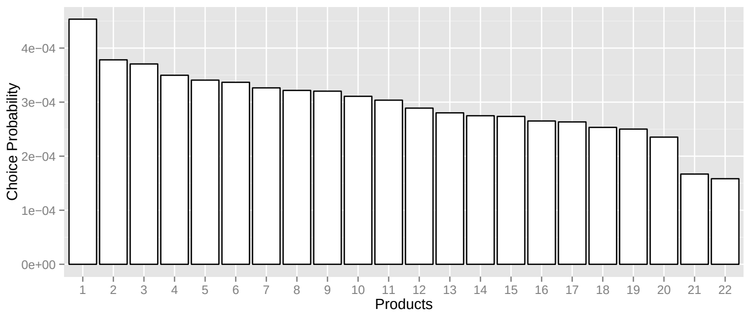

* First, I compute the daily sales volume of a product in a vending machine. Then, I divide the sales volume by the size of vending machine consumers as defined in the main text. Finally, I take the average across days between March and September 2013 and across vending machines in the Tokyo metropolitan area. The product identifications in this figure is different from those in Figure 1.3a.

for which each product was available. Figure 1.3b shows the average choice probabilities of products, which are adjusted for the days, the number, and the size of vending machine consumers for which each product was available. The difference decreases, but substantial heterogeneity in the popularity across products remains. The average choice probabilities of the most and the median popular products are 0.0045% and 0.0029%. Figure 1.4 shows that the popularity of a product can differ to a large extent across local markets. It is this heterogeneity in demand across products and local markets that motivates the retail network to learn local demand to optimise the product assortment of each location.

Figure 1.4: The Percentiles of Average Choice Probabilities of Products Across Local Markets

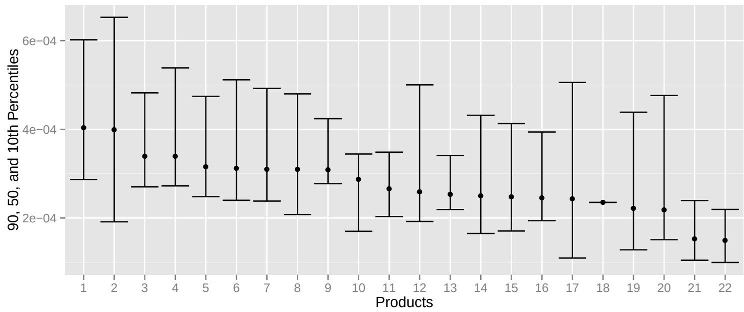

* I compute the average choice probability of a product as in Figure 1.3b but for each local market. Then, I compute the 90, 50, and 10th percentiles of the average choice probabilities of products across local markets. The product identifications in this figure is different from those in Figures 1.3a and 1.3b.

Second, I check the frequency of product changes in the tea category during the spring/summer seasons in 2013.

Figure 1.5a shows the distribution for the number of product changes in the tea category across vending machines. On average, products are changed 9.5 times per vending machine in the tea category during the season. Figure 1.5b shows the number of product changes in the tea category for each week during the season. Three peaks approximately correspond to the shift from winter assortments to spring, spring to summer, and spring to fall. On average, products are changed 856 times per week, and the peak and the bottom are 2,297 and 99 times per week, respectively.

Table 1.1a demonstrates that the timing of product changes are asynchronous across locations. The table implements logistic regressions of indicators of products' changed in a vending machine on a given day on various sets of dummy variables to study the source of the total variation. The table shows that the date dummies can explain only 6.8% of the total variation according to McFadden's R^2. Even the date $\times$ station dummies can explain only 31.0% of the total variation. Table 1.1b demonstrates that there is a spatial heterogeneity in product availability. The table implements logistic regressions of indicators of products available in a vending machine on a given day on various sets of dummy variables. The table shows that the date $\times$ good dummies can explain 31.2% of the total variation. Thus, aggregate seasonality matters as for product availability. Adding date $\times$ station dummies explains the additional 16.6% of the total variation in product availability. The reminder of the total variation in

Figure 1.5: The Number of Product Changes

(a) The Number of Product Changes across Vending Machines

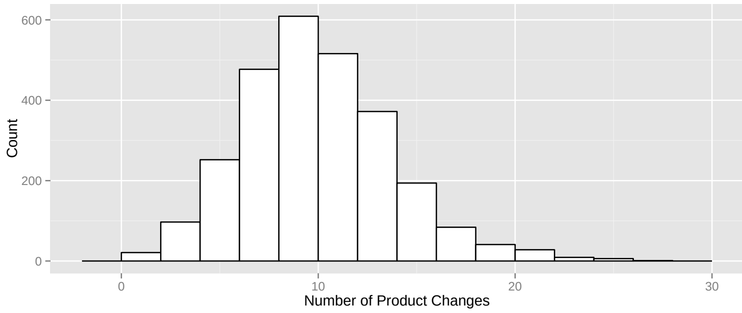

* I first construct the weekly history of product assortments for each vending machine as described in the main text. Then, I determine how many products are changed in a week in a vending machine. Finally, I counted the total number of product changes for each vending machine.

(b) The Number of Product Changes across Weeks

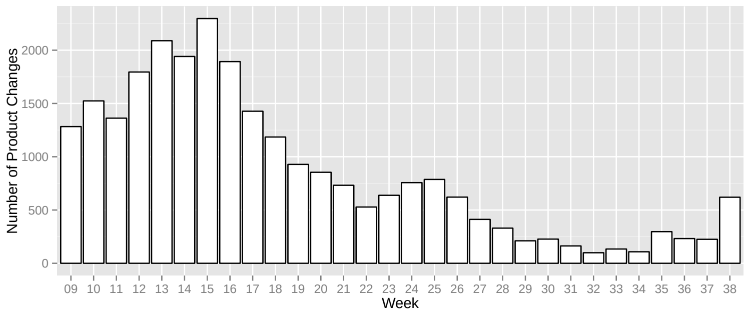

* I first construct the weekly history of product assortments for each vending machine as described in the main text. Then, I determine how many products are changed in a week in a vending machine. Finally, I counted the total number of product changes for each week.

product availability is from the variation across vending machines in a station.

Finally, I check that there is a positive correlation between the current product assortment decisions and past sales. To see this, I construct three variables: First, 'VM Share' of a product in a local market in a week is the number of vending machines in the market with the product over the number of vending machines in the market. Second, 'Exogenous VM Share' of a product in a local market in a week is the number of vending machines in the market in which the principal company requires to put the product over the number of vending machines in the market. Third, 'Past Choice Probability' of a product in a local market in a week is the average

Table 1.1: Determinants of Product Changes and Product Availabilities

(a) Timing of Product Changes

Dummies	McFadden's R^2
Date	0.0679
Date×Operator	0.0820
Date×Market	0.2016
Date×Station	0.3103

* I first construct a variable that takes one if tea products are changed in the vending machine on the day and takes zero otherwise. Then, I run logistic regressions on various sets of dummy variables. Finally, I compute McFaddden's R^2 as a measure of fit.

(b) Product Availability

Dummies	McFadden's R^2
Date×Good	0.3118
Date×Good×Operator	0.3589
Date×Good×Market	0.4364
Date×Good×Station	0.4774

* I first construct a variable that takes one if the tea product exists in the vending machine on the day and takes zero otherwise. Then, I run logistic regressions on various sets of dummy variables. Finally, I compute McFaddden's R^2 as a measure of fit.

choice probability of the product in the market in the last month. Then, I regress the log of 'VM Share' on the log of 'Exogenous VM Share' and 'Past Choice Probability' with and without fixed effects. Table 1.2 reports the regression results. First, the elasticities of 'VM Share' to 'Exogenous VM Share' are 0.86 and 0.85 in the left and right columns, indicating that agents are loyal to the order by the principal company. Second, the elasticity of 'VM Share' to 'Past Choice Probability' is 0.07 if only market $\times$ good-fixed effects are controlled, and 0.09 if both market $\times$ good- and week- fixed effects are controlled. They are statistically significant, but the magnitude is small. This suggests that agents are learning from past sales data to optimise product assortments, but there is large friction in the learning and adjustment process.

I explicitly incorporate the institutional restriction and the covariates such as temperatures

Table 1.2: Current Vending Machine Share and Past Choice Probabilities of Products

	log(VM Share)	log(VM Share)
log(Exogenous VM Share)	0.8663	0.8554
	(0.0040)	(0.0042)
log(Past Choice Probability)	0.0744	0.0930
	(0.0069)	(0.0072)
R-squared	0.8337	0.8364
Adj. R-squared	0.7635	0.7641
Market×Good-Fixed Effects	Yes	Yes
Week-Fixed Effects	No	Yes
Balanced	No	No
n	1081	1081
nT	12867	12867

* 'VM Share' of a product in a local market in a week is the number of vending machines in the market with the product over the number of vending machines in the market. 'Exogenous VM Share' of a product in a local market in a week is the number of vending machines in the market in which the principal company requires to put the product over the number of vending machines in the market. 'Past Choice Probability' of a product in a local market in a week is the average choice probability of the product in the market in the last month.

** n is the number of unique cross-sectional units, and nT is the total sample size.

*** I drop observations if either of 'VM Share', 'Exogenous VM Share', and 'Past Choice Probability' are zero or N.A.

**** Standard errors are in parentheses.

in the structural estimation, and quantify how effectively agents are learning from the sales data, and how much they contribute to the revenue of the retail network.

1.3 Model

1.3.1 A Model for Illustration

In this section, I illustrate the intuition of the model using a simple model. I also introduce the Bayesian control rule and provide rationales for employing this rule as a solution concept of the empirical model.

Demand

Consider an agent assigned in a local market. The agent has one vending machine with two slots. There are three distinct products in the list. The price of the products are one. The agent selects the assortments for the vending machines for a season. A season consists of two periods.

In a market, there are consumers with measure M. The indirect utilities of a consumer i are $v_{ij} = v_j + \nu_{ij}$ for $j = 0, 1, 2, 3$, where $j = 1, 2, 3$ are products $1, 2, 3$, and 0 represents the outside option of not buying anything. The preference shocks ν_{ij} are drawn from independent and identical type-I extreme-value distributions. The mean indirect utility of the outside option, v_0, is normalised at zero. When the consumer chooses an alternative that maximises his utility. The vector of mean indirect utilities $v = (v_1, v_2, v_3)'$ characterises the demand of the local market. I call this a local market demand parameter.

The Agent's Initial Belief and Information Processing Technology

The agent does not know the exact value of the demand parameter of his market. However, the agent has some initial belief about the demand parameters at the beginning of the season. Let $N(\mu, \Sigma)$ denote the agent's initial belief about the demand parameter v.

Before the season starts, the agent chooses the precision of his information processing technology by paying costs. Thus, the precision can be regarded as a reduced form parameter of the agent's costs of processing information. As a season starts, the agent decides the assortment for the vending machine. Then, the sales are realised and his information processing technology updates his belief. Particularly, letting $d = (d_1, d_2, d_3)'$ be the vector of indicators of product availability on the vending machine, the agent receives a signal as follows:

$$\underbrace{y_j}_{signal} = \ln \underbrace{\frac{\exp(v_j) d_j}{1 + \sum_{j'=1}^{3} \exp(v_{j'}) d_{j'}}}_{choice\ probability \equiv \frac{r_j(d,v)}{M}} + \underbrace{\eta_j}_{noise}, \tag{1.1}$$

where the noise η_j is drawn from the independent and identical normal distribution with a mean of zero and precision κ and updates the belief to $N(\mu', \Sigma')$ based on $y = (y_1, y_2, y_3)'$ according to the Bayes rule.[15] The precision of the agent's information processing technology

[15] Because the mean term of the measurement equation is non-linear in the demand parameters, the model has no obvious conjugate prior, thus the posterior is not analytically tractable even if a normal prior is used. I track the belief process using a first-order extended Kalman filter, and approximate the posterior belief with a normal belief.

κ determines the speed of learning. The noise terms capture all cognitive errors between sales data and the agent's action.[16]

The agent's initial belief $N(\mu, \Sigma)$ and the precision of information processing technology κ constitute the main parameters of interest in the model.

Agent's Decision Problem

Assume that the agent is risk-neutral.[17] A fully rational agent chooses a decision rule $d(\mu, \Sigma)$ and $d'(\mu', \Sigma')$ that maximises the expected revenue:

$$\mathbb{E}_v\left[\sum_{j=1}^{3} r_j(d(\mu,\Sigma),v)d_j(\mu,\Sigma)\middle|\mu,\Sigma\right] + \mathbb{E}_{(\mu',\Sigma')}\left[\mathbb{E}_v\left[\sum_{j=1}^{3} r_j(d'(\mu',\Sigma'),v)\middle|\mu',\Sigma'\right]\middle|\mu,\Sigma\right]. \quad (1.2)$$

This is a standard Markov decision process where the state variables are belief states (μ, Σ) and (μ', Σ'). In the decision problem, the agent faces a trade-off between exploitation and exploration: the agent wants to choose an assortment with a higher expected payoff (exploitation) but at the same time the agent wants to choose an assortment with a higher uncertainty (exploration).

The fully rational solution is, however, computationally infeasible. First, the dimensionality of the state space is extremely large. This simple model already has $3 + 3^2 = 12$ dimensions. The dimensionality exceeds one thousand in the application. Second, product assortment decisions are by nature combinatorial optimisation problems. This first means that the action space is also extremely large and, second, that the belief updating over payoffs of each assortment is highly dependent. The belief over the payoff of an assortment can change even if other assortments are chosen if they share a same product. If the belief updating is so complicated, then efficient algorithms based on Gittin's index, which have been used in the literature of empirical analysis of learning, no longer work.

[16] When I fit the demand model to the data, I include additional measurement errors. These errors in the sales data are distinguished from the agent's cognitive errors both conceptually and empirically. These measurement errors cannot be removed even if a perfect information processing technology ($\kappa = \infty$) is utilised.

[17] The fee is proportional to the sales, which suggests that the agent can be risk averse. However, I assume risk-neutrality. Because risk averseness and larger initial uncertainty have similar implications for the agent's assortment decisions, I normalise either the initial uncertainty at the level of rational expectation or the degree of risk-averseness at the risk-neutral level. My counterfactual analysis does not require distinguishing these two reasons for the agent's actions.

Bayesian Control Rule

Because of this computational problem, I borrow a decision rule called the Bayesian control rule (Ortega, 2011) from the engineering literature.

The algorithm works as follows: given a belief at period $N(\mu, \Sigma)$, i) the agent samples a state v^* from his belief $N(\mu, \Sigma)$, ii) chooses an action that maximises the objective assuming that sample v^* is the true state of the world, and iii) updates his belief based on the realised signals. The decision rule coincides with the ex-post optimal decision rule when there is no uncertainty. It converges to the ex-post optimal decision rule over trials.

The decision rule intuitively resolves the exploitation-exploration trade-off. Under the Bayesian control rule, assortments either with higher expected payoff or with higher uncertainty over the payoff are likely to be chosen. As the uncertainty is cleared, then the expected payoff dominates.

One drawback of the decision rule is that it is not forward-looking. Therefore, if there is an expectation that increases the value of exploration, then the decision rule fails to capture the effects of the expectation. For example, a fully rational agent will intensify exploration if he knows that more vending machines are allocated to his market in the next period. An agent following the Bayesian control rule fails to adjust the balance of exploitation-exploration trade-off in such a case. However, this problem should be negligible because the current setting does not include such occasions.

The decision rule works well under various settings (Vermorel and Mohri, 2005; Kuleshov and Precup, 2014). We later see that it also works well in the current application. Additionally, there is ample experimental evidence in psychology and experimental economics that this type of decision rule better describes human behaviors than the expected utility maximisation rule rule (Shanks, Tunney, and McCarthy, 2002).

The choice probabilities under the Bayesian control rule are:

$$\begin{aligned} p(d_0|\mu, \Sigma) &= \int 1\left\{ d_0 = \text{argmax}_d \sum_{j=1}^{3} r_j(d, v) d_j \right\} dN(v|\mu, \Sigma), \\ p(d_0'|\mu', \Sigma') &= \int 1\left\{ d_0' = \text{argmax}_{d'} \sum_{j=1}^{3} r_j(d', v') d_j' \right\} dN(v'|\mu', \Sigma'). \end{aligned} \tag{1.3}$$

Thus, state of the world optimal policies are mixed with belief at this point. This completes the description of the model.

1.3.2 Full Model

Now I modify the model to fit the reality. There are many agents in the retail network, and they are assigned to each local market. The demand parameters can be arbitrarily dependent across local markets. Every period, the principal company and operating companies send aggregate product-level sales data to each agent. However, I assume that agents' decisions are separated conditional on the demand parameters and the aggregate product-level sales data. Therefore, I suppress the index of agents in the followings. All variables and parameters are indexed by agents and markets.

The empirical model differs from the model for illustration in several aspects: i) There are K vending machines in a market, and J products are listed.[18] There are T weeks in a season. Agents discount the future by δ. The prices can differ across products (but are time-invariant). ii) Temperatures are entered as a covariate in the indirect utilities. iii) Every period, the principal and operating companies send aggregate product-level sales data to agents. iv) A hierarchical structure is imposed on the agent's initial belief to reduce the dimensionality of the parameter space. ii) Agents can change products in a vending machine only with the probability $\lambda \in [0, 1]$ because of the cost of implementing plans.

Demand

Consumers with measure M_k pass by vending machine k every week. The mean indirect utilities in a market are $v_{ijt}^L = \xi_j^L + \beta_j^L \ln(z_t)$ for $j = 0, 1, \cdots, J$, where z_t is a temperature of week t. I assume that the principal and agents have perfect foresight for temperatures.[19] The demand parameters of the outside option of not buying anything are normalized at zero: $\xi_0^L = 0$ and $\beta_0^L = 0$. The vector of parameters in the mean indirect utilities $\tau^L = (\tau_1^{L\prime}, \cdots, \tau_J^{L\prime})'$, where $\tau_j^L = (\xi_j^L, \beta_j^L)'$, characterise the demand of the local market, and are the demand parameters of the local market in the empirical model.

A consumer chooses an alternative that maximises utility. Then, the choice probabilities in vending machine $k = 1, \cdots, K$ in week $t = 1, \cdots, T$ are:

$$\frac{r_{jtk}(d, \tau^L)}{M_k} = \frac{\exp(\xi_j^L + \beta_j^L \ln(z_t))d_j}{1 + \sum_{j' \in J} \exp(\xi_{j'}^L + \beta_{j'}^L \ln(z_t))d_{j'}}, j = 0, 1, \cdots, J. \tag{1.4}$$

[18]The number of vending machines in a market and the set of products in the list change over time. This is reflected at the estimation stage. I abstract away from it in the description of the model.

[19]This is a safe assumption because the principal company and the agents regularly check 10-year average daily temperatures and short- and long-run weather forecasts released by the Japan Meteorological Agency.

Agent's Information Processing Technology

The agent has three sources of sales data. The first data come from the local market. The remaining data come from the principal and operating companies. At the beginning of the season, the agent chooses how precisely he processes each piece of information. Particularly, in week $t = 1, \cdots, T$, the agent receives signals such that:

$$
\begin{aligned}
&(\textit{Local}) && y^L_{jtk} = \ln \frac{r_{jtk}(d, \tau^L)}{M_k} + \eta^L_{jtk}, j = 1, \cdots, J, k = 1, \cdots, K, \\
&(\textit{Principal Company}) && y^P_{jt} = \frac{1}{N^P_{jt}} \sum_{k=1}^{N^P_{jt}} \ln \frac{r_{jtk}(d, \tau^L)}{M_k} + \eta^P_{jt}, j = 1, \cdots, J, \\
&(\textit{Operating Company}) && y^O_{jt} = \frac{1}{N^O_{jt}} \sum_{k=1}^{N^O_{jt}} \ln \frac{r_{jtk}(d, \tau^L)}{M_k} + \eta^O_{jt}, j = 1, \cdots, J,
\end{aligned}
\tag{1.5}
$$

where N^P_{jt} and N^O_{jt} are the numbers of vending machines that have product j in all the markets and in the markets belonging to the operating company, and $\eta^L_{jtk}, \eta^P_{jt}$, and η^O_{jt} are drawn from independent normal distributions with a mean of zero and precisions κ^L, κ^P, and κ^O, respectively.

These signals are related to the underlying demand parameters through measurement equations. The mean terms of the signals from the local market, y^L_{jtk}, are related to the local demand parameter τ^L through equation (1.4). The mean terms of the signals from the principal and operating companies are assumed approximately related to the aggregate demand parameters as:

$$
\begin{aligned}
&\frac{1}{N^P_{jt}} \sum_{k=1}^{N^P_{jt}} \ln \frac{r_{jtk}(d, \tau^L)}{M_k} \approx \xi^P_j + \beta^P_j \ln(z_t), j = 1, \cdots, J, t = 1, \cdots, T, \\
&\frac{1}{N^O_{jt}} \sum_{k=1}^{N^O_{jt}} \ln \frac{r_{jtk}(d, \tau^L)}{M_k} \approx \xi^O_j + \beta^O_j \ln(z_t), j = 1, \cdots, J, t = 1, \cdots, T.
\end{aligned}
\tag{1.6}
$$

The vectors $\tau^P = (\tau^{P\prime}_1, \cdots, \tau^{P\prime}_J)'$ with $\tau^P_j = (\xi^P_j, \beta^P_j)'$ and $\tau^O = (\tau^{O\prime}_1, \cdots, \tau^{O\prime}_J)'$ with $\tau^P_j = (\xi^O_j, \beta^O_j)'$ characterise the aggregate states of demand. I call them aggregate demand parameters and call $\tau = (\tau^{L\prime}, \tau^{P\prime}, \tau^{O\prime})'$ demand parameters. The agent updates the belief about τ based on the signals from the information processing technology defined as above according to the Bayes rule.

Agent's Initial Belief

At the beginning of the season, the agent has an initial belief $N(\mu_1, \Sigma_1)$ over τ. I impose a hierarchical structure on the initial belief to reduce the dimensionality of the parameter space.

The benchmark of the initial belief is that of the rational expectation. The belief under the rational expectation is the distribution of the vector of local demand parameters $\{\tau_j^L\}_{j \in J}$ across local markets. I assume that the principal company's belief is no worse than the belief under the rational expectation while agents may have better knowledge of their own local markets.

I assume that the mean terms of an agent's initial belief about their local demand parameters are distributed normally around the true local demand parameters. I allow for the dispersion to depend on the weeks after the release of relevant products at the beginning of the season because, at that point, the agent will already have some information.The mean terms of an agent's initial belief about a local demand parameter ξ_j^L, $\mu_1(\xi_j^L)$, for $j = 1, \cdots, J$, are distributed as:

$$\begin{aligned} &\mu_1(\xi_j^L) \sim N(\xi_j^L, \pi(\xi_j^L)^2), \\ &\ln \pi(\xi_j^L) \equiv \pi_{\xi,0} - \pi_{\xi,1} \ln(WeeksAfterRelease)_j, \end{aligned} \tag{1.7}$$

and the mean terms with respect to β_j^L for $j = 1, \cdots, J$ are given in the same way. The hyper parameters $\pi_{\xi,0}$ and $\pi_{\beta,0}$ determine the average biases of the agent's initial belief, and $\pi_{\xi,1}$ and $\pi_{\xi,1}$ are the sensitivity of the biases to the weeks after the release of the relevant product at the beginning of the season. If $\pi(\xi_j^L)$ and $\pi(\beta_j^L)$ are zero for $j = 1, \cdots, J$, then the agent has an unbiased belief about the local demand parameters.

I assume that the standard deviations of an agent's initial belief about local demand parameters are proportional to the standard deviations of the belief under rational expectation. I allow for the proportion to depend on the weeks after the release of relevant products at the beginning of the season. Particularly, the standard deviation terms of an agent's initial belief about a local demand parameter ξ_j^L, $\sigma_1(\xi_j^L)$, for $j = 1, \cdots, J$, are:

$$\begin{aligned} &\sigma_1(\xi_j^L) = \rho_j \sigma^{RE}(\xi_j^L), \\ &\ln(\rho_j) = \rho_0 - \rho_1 \ln(WeeksAfterRelease)_j, \end{aligned} \tag{1.8}$$

where $\sigma^{RE}(\xi_j^L)$ for $j = 1, \cdots, J$ are the standard deviations of the belief under rational expectation, and the standard deviations with respect to β_j^L for $j = 1, \cdots, J$ are given in the same way. The hyper parameter ρ_0 determines the average accuracy of the agent's initial belief, and ρ_1 determines the sensitivity of the biases to the weeks after the release of the relevant product at the beginning of the season. If the agent has an unbiased belief and ρ_j are zero for $j = 1, \cdots, J$, then the agent has complete knowledge about the local demand parameter.

The mean terms for the aggregate demand parameters $\mu_1(\xi_j^P), \mu_1(\beta_j^P), \mu_1(\xi_j^O)$, and $\mu_1(\beta_j^O)$ for $j = 1, \cdots, J$ are set at the actual average of $\{\xi_j^L\}_{j \in J}$ and $\{\beta_j^L\}_{j \in J}$ across entire local markets and across markets belonging to the operating company. The standard deviations $\sigma_1(\xi_j^P), \sigma_1(\sigma_j^P), \mu_1(\sigma_j^O)$, and $\sigma_1(\beta_j^O)$ are set at the same values as the standard deviations for the local parameters,$\sigma_1(\xi_j^L), \sigma_1(\sigma_j^L)$ for $j = 1, \cdots, J$.

The correlation in the initial belief between ξ_j^L and ξ_j^P or ξ_j^O are set at the actual correlation between $\{\xi_j^L\}_{j \in J}$ and the actual average of $\{\xi_j^L\}_{j \in J}$ across entire markets or across markets belonging to the operating company. The correlation for β_j^L terms are set in the same way. The correlation between terms with respect to different products or different types are set at zero.

This completes the description of an agent's initial belief and the information processing technology. I track the belief updating by the first-order extended Kalman filter, and approximate the belief using a normal belief.

Bayesian Control Rule

In week $t = 1, \cdots, T$, if the local demand parameters are $\{\tau_j^L\}_{j \in J} = \{(\xi_j^L, \beta_j^L)'\}_{j \in J}$, and the agent chose assortment d_{kt} for vending machine $k = 1, \cdots, K$, then the vending machine yields revenue:

$$R_{kt}(d_{kt}, \tau^L) \equiv \sum_{j \in J} p_j \frac{\exp(\xi_j^L + \beta_j^L \ln(z_t)) d_{ktj}}{\sum_{j' \in J} \exp(\xi_{j'}^L + \beta_{j'}^L \ln(z_t))}. \tag{1.9}$$

I assume that an agent can change vending machine products in a week only with probability $\lambda \in [0, 1]$ because of the cost of implementing plans. I assume that the arrivals of the chance of a move are independent across vending machines based on the evidence in the descriptive analysis. Therefore, in week $t = 1, \cdots, T$, if the local demand parameters are $\{\tau_j^L\}_{j \in J} = \{(\xi_j^L, \beta_j^L)'\}_{j \in J}$, and the agent chose assortment d_{kt} for vending machine $k = 1, \cdots, K$, the expected revenue from the vending machine is:

$$\begin{aligned} V_{kt}(d_{kt}, \tau) &\equiv \sum_{s=0}^{T-s} \lambda(1-\lambda)^s \sum_{u=0}^{s} \delta^u R_{k,t+u}(d_{kt}, \tau) + C_t(\tau) \\ &= \sum_{u=0}^{T-t} \delta^u R_{k,t+u}(d_{kt}, \tau) \sum_{s=u}^{T-t} \lambda(1-\lambda)^s + C_t(\tau) \\ &= \sum_{u=0}^{T-t} \delta^u (1-\lambda)^u [1 - (1-\lambda)^{T-t-u+1}] R_{k,t+u}(d_{kt}, \tau) + C_t(\tau), \end{aligned} \tag{1.10}$$

where $C_t(\tau)$ is a constant term that does not depend on the assortment d_{kt} that represents the sum of value functions at the point where the chance of a move arrives.

Given the deterministic choice-specific value functions $\{V_{kt}(d_{kt}, \tau)\}_{k \in K, t \in T}$ and the initial belief $N(\mu_1, \Sigma_1)$, the Bayesian control rule works as follows: for each $t = 1, \cdots, T$, for each vending machine $k = 1, \cdots, K$, if the chance of a move arrives, i) the agent samples a vector of local demand parameters τ^{L*} from $N(\mu_t, \Sigma_t)$, ii) chooses an assortment d_{kt} such that $d_{kt} = \text{argmax}_d V_{kt}(d, \tau^{L*})$, and after deciding assortments for all the vending machines, given realized signals y_t, iii) updates his belief to $N(\mu_{t+1}, \Sigma_{t+1})$ by the belief updating function $B(\mu_{t+1}, \Sigma_{t+1} | \mu_t, \Sigma_t; y_t, d_t)$. Then, the choice probabilities of an assortment d_{kt} for vending machine $k = 1, \cdots, K$ in week $t = 1, \cdots, T$ are:

$$p_{kt}(d_{kt} | \mu_t, \Sigma_t) = \int 1\{V_{kt}(d_{kt}, \tau^L) = \max_d V_{kt}(d, \tau^L)\} dN(\tau^L | \mu_t, \Sigma_t). \tag{1.11}$$

Identification

The demand parameters are identified outside the dynamic model from the history of sales and assortments. This gives the revenue functions. Intuitively, I identify the main parameter of interests as follows. First, the history of assortment decisions reveals the agent's belief process including his initial belief. Second, from the history of sales associated with the history of assortments, I compute the belief process of an agent without processing costs. Comparing these belief processes, I can identify the precision of agent's information processing technology. The identification strategy presumes that there are many vending machines in a local market. In reality, I borrow some identification powers of functional form restrictions to pin down the estimates.

1.4 Estimation

This section describes how to estimate the full empirical model.

All the parameters are agent- and market-specific, and estimation is done agent by agent and market by market. Hence, all the spatial heterogeneity is, in principle, recovered. However, the effective sample size to estimate dynamic parameters of an agent can be small because product change is infrequent. Therefore, the estimates of dynamic parameters of an agent can be inaccurate. However, errors in the estimates of the dynamic parameters at the agent level are washed out when I evaluate the expected total revenue of the network, my ultimate parameter of interest, because the expected total revenue is a sum of expected revenues across markets, and the expected revenue of a market is a statistic that is continuous in the dynamic parameters of an agent in the market.

I continue to suppress index of agent and market in this section.

1.4.1 Estimating Demand Parameters

I first estimate the demand parameters outside the dynamic model. I consider the following regressions:

$$
\begin{aligned}
&\ln \frac{r_{jtk}(d_{tk}, \tau^L)}{M_k} = \left\{ \xi_j^L + \beta_j^L \ln(z_t) - \ln \left(\sum_{j' \in J} \exp(\xi_{j'}^L + \beta_{j'}^L \ln(z_t)) d_{j'tk} \right) \right\} d_{jtk} + \epsilon_{jtk}^L, \\
&j = 0, 1, \cdots, J, t = 1, \cdots, T, k = 1, \cdots, K, \\
&\frac{1}{N_{jt}^P} \sum_{k=1}^{N_{jt}^P} \ln \frac{r_{jtk}(d_{tk}, \tau^L)}{M_k} = \xi_j^P + \beta_j^P \ln(z_t) + \epsilon_{jt}^P, j = 1, \cdots, J, t = 1, \cdots, T, \\
&\frac{1}{N_{jt}^O} \sum_{k=1}^{N_{jt}^O} \ln \frac{r_{jtk}(d_{tk}, \tau^L)}{M_k} = \xi_j^O + \beta_j^O \ln(z_t) + \epsilon_{jt}^O, j = 1, \cdots, J, t = 1, \cdots, T,
\end{aligned}
\tag{1.12}
$$

where $\epsilon_{jtk}^L, \epsilon_{jt}^P$, and ϵ_{jt}^O, are drawn from independent normal distributions with a mean of zero and standard deviations $\sigma_j^L, \sigma_j^P / \sqrt{N_{jt}^P}$, and $\sigma_j^O / \sqrt{N_{jt}^O}$. The standard deviations for the aggregate sales data decrease at the rate of $\sqrt{N_{jt}^P}$ and $\sqrt{N_{jt}^O}$ because the aggregate errors are the averages of the errors across N_{jt}^P and N_{jt}^O vending machines.

I run NLLS regressions to estimate the demand parameters τ and the standard deviations in the sales data $\sigma = (\sigma^{L\prime}, \sigma^{P\prime}, \sigma^{O\prime})'$ with $\sigma^L = (\sigma_1^L, \cdots, \sigma_J^L)'$, $\sigma^P = (\sigma_1^P, \cdots, \sigma_J^P)'$, and $\sigma^O = (\sigma_1^O, \cdots, \sigma_J^O)'$. These errors are fundamental in the sales data and are distinguished from the cognitive errors in the information processing technologies of agents.

The critical assumption for this estimation is that assortment decisions in week t are uncorrelated with the errors in the week. Because the assortment is correlated with past sales and so with past errors, the errors must be serially uncorrelated. Another critical assumption is the homogeneity of the demand parameters within a local market and a season. As long as the homogeneity assumption holds and a product is put in a vending machine for some periods in a local market, the local demand parameters regarding the product are, in principle, identified. In reality, few products are never tried in a local market. In such a case, I use the average of the estimates as the corresponding local demand parameter in the next step of the estimation and counterfactual analysis.

Figure 1.6: The Percentiles of Estimates of Products' Demand Parameters across Local Markets

(a) ξ: intercepts

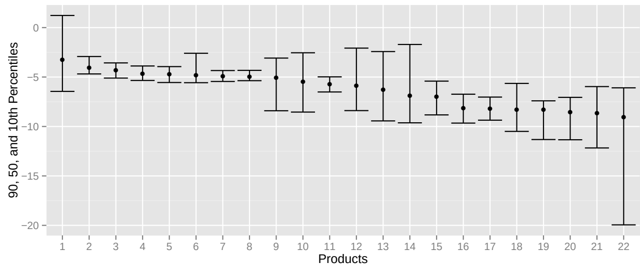

* For each product, I compute the 10, 50, and 90th percentiles of the estimates of the intercepts across local markets and sort by the value of the 50th percentile.

(b) β: coefficients on temperatures

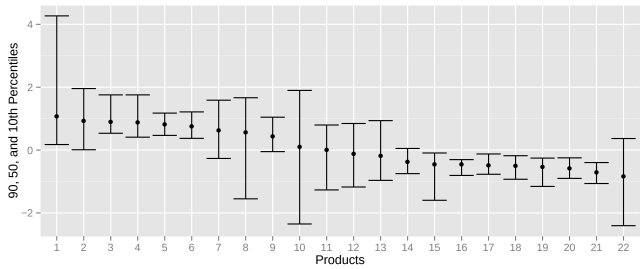

* For each product, I compute the 10, 50, and 90th percentiles of the estimates of the coefficients of temperatures across local markets and sort by the value of the 50th percentile.

Figure 1.6 displays the percentiles of estimates across local markets. Because these are tea products, most of the products react positively to temperatures. To test serial correlations, I apply the Breush-Godfrey test for the residuals of each cross-sectional unit (vending machine good), and to control the false discovery rate I adjust these p-values by the method of Benjamini and Hochberg (1995). Then, Figure 1.7 illustrates sorted adjusted p-values. The figure shows that only a few of the units reject the null hypothesis of no serial correlation. As a simple measure of fit of the demand models, Table 1.3 describes the summary statistics of the R-squared values of the NLLS regressions across local markets. On average, the value is 0.33.

Figure 1.7: Adjusted p-values of the Breush-Godfrey Tests

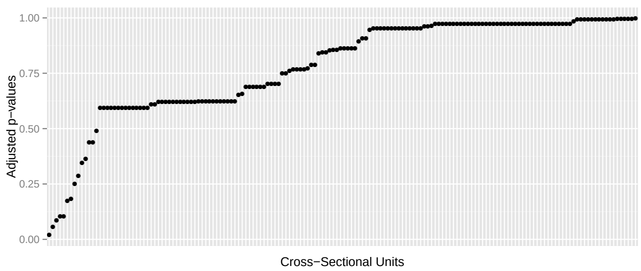

* I first apply the Breush-Godfrey test for the residuals of each cross-sectional unit (vending machine good) of the local demand models and then adjust the series of p-values by the method of Benjamini and Hochberg (1995). Then, I sort the p-values according to the size and plot them from the smallest to the largest.

Table 1.3: R^2 of Local Demand Models across Markets

	R^2
mean	0.3282
sd	0.0888
p25	0.2686
p50	0.3167
p75	0.3770
min	0.1280
max	0.5970

* For each local market, I first compute the R-squared values of the local demand models and then take their summary statistics across local markets.

1.4.2 Estimating Dynamic Parameters

A vector of dynamic parameters of interest, θ, consists of precisions of agent's information technologies $(\kappa^L, \kappa^P, \kappa^O)'$, hyper parameters in the agent's initial belief (ρ_0, ρ_1) and $(\pi_{\xi,0}, \pi_{\xi,1}, \pi_{\beta,0}, \pi_{\beta,1})$, the probability of a move λ, and the discount factor δ. A vector of observations consists of the history of sales r, assortments d, and temperatures z.

I let $B(\mu_{t+1}, \Sigma_{t+1}|\mu_t, \Sigma_t; r_t, d_t)$ be the transition probability of belief states conditional on assortments d_t defined in the previous section. Then, the likelihood function is:

$$
\begin{aligned}
&L(\theta|d,z) \\
&= \int \prod_{t=1}^{T} B(\mu_{t+1},\Sigma_{t+1}|\mu_t,\Sigma_t;r_t,d_t) \\
&\times \prod_{k=1}^{K} [(1-\lambda)1\{d_{kt}=d_{k,t-1}\} + \lambda p_{kt}(d_{kt}|\mu_t,\Sigma_t)] d\mu_{t+1} d\Sigma_{t+1}.
\end{aligned} \tag{1.13}
$$

Because belief states are hidden Markov state variables, I integrate them out by applying a particle filter where the proposal distribution of particles is given by $B(\mu_{t+1},\Sigma_{t+1}|\mu_t,\Sigma_t;r_t,d_t)$. Resampling is according to the normalised importance weight when the effective number of samples is below $N_{pf}/10$, where N_{pf} is the number of particles. The number of particles N_{pf} is set at 100, and the number of samples to evaluate each choice probability p_{kt}, N_{sim}, is set at 200.

The parameters are estimated by a maximum a posterior (MAP) estimator. The prior on the static parameters is non-informative, that is, the prior density is constant, and the inference is fully likelihood-based. I sample from the posterior distribution using a robust adaptive Metropolis-Hastings sampler (Vihola, 2011) to facilitate the convergence that targets mean acceptance probability 0.5 with the step-size tuning parameter γ set at 0.75. In total, 1000 samples are generated, and the first 500 samples are burned in.

To impose upper and lower bound restrictions, some parameters are reparameterised using monotonic mappings. First, because $\lambda,\delta \in (0,1)$, some parameters are reparameterised by a one to one mapping $\lambda = \exp(\lambda^*)/(1+\exp(\lambda^*))$ and $\delta = \exp(\delta^*)/(1+\exp(\delta^*))$, and the sampler is run on the space of λ^* and δ^*. $\kappa_0,\kappa_1,\kappa_2$ are reparameterised by a one to one mapping $\kappa = \exp(\kappa^*)$.

Table 1.4 summarises the estimation results. Because I obtain estimates for each local market, the table only shows the summary statistics of the estimates across local markets. I discuss the implications in the next section. To assess the fit of the model, I compute a pseudo R^2 as follows: let s_i be the number of slots allocated to a product in a local market in a week over the number of slots for which agents are responsible in the local market in the week. Then, let $\hat{s}_i$ be the expected share derived from the estimated model and $\overline{s}$ be the average of $\{s_i\}$. Then, the pseudo R^2 is defined as $1-\sum_i(s_i-\hat{s}_i)^2/\sum_i(s_i-\overline{s})^2$. The pseudo R^2 of the current model is 0.136. Some variations cannot be explained by the current model.

Table 1.4: Estimation Results: Summary Statistics Across Local markets

Across Markets.	mean	sd	p25	p50	p75	min	max
Signal Precision							
κ^L	40.2342	235.7781	3.1682	11.6261	22.9763	0.6521	2220.0666
κ^P	6.0493	12.8717	0.6902	1.6671	5.0235	0.0015	85.2308
κ^O	4.2772	10.3657	0.4695	1.3078	3.9344	0.0002	75.3124
Initial Belief							
ρ_0	-0.7215	1.2124	-0.9998	-0.5199	-0.0590	-8.3698	0.8102
ρ_1	0.0058	0.0221	0.0002	0.0009	0.0022	0.0000	0.1814
$\pi_{\xi,0}$	-0.3146	1.2897	-1.0829	-0.4189	0.5670	-4.0100	2.7819
$\pi_{\xi,1}$	0.0194	0.1053	0.0004	0.0012	0.0036	0.0000	0.9419
$\pi_{\beta,0}$	-0.8103	1.4291	-1.6066	-0.4265	0.1314	-4.9390	1.5277
$\pi_{\beta,1}$	0.0226	0.1751	0.0004	0.0010	0.0032	0.0000	1.6447
Others							
λ	0.0468	0.0291	0.0259	0.0415	0.0641	0.0004	0.1685
δ	0.9568	0.1201	0.9737	0.9910	0.9965	0.0493	1.0000

* I estimate the parameters market by market (agent by agent). Then, I take summary statistics of the estimates across local markets.

1.5 Analysis

1.5.1 Quantifying the Agent Contributions

I first study the agent contributions to the retail network. I define baseline revenue as the revenue when no information processing is exerted and no private information is exploited. No information processing is exerted means that the agents do not learn from past sales, that is, the signal precisions κ^L, κ^P, and κ^O are all zero. No private information is exploited means that the initial beliefs are those of the rational expectation.[20] The assumption behind this setting is that the principal's belief about local demand parameters is no worse than those of the rational expectation. Therefore, contributions of agents' effort compared to the revenue baseline should be interpreted as the upper bound of the agent contributions.

[20]The mean and the standard deviation of a parameter coincide with the mean and the standard deviation of the ex-post distribution of the parameter across local markets.

Figure 1.8: Simulated Distribution of the Total Revenue Relative to the Baseline

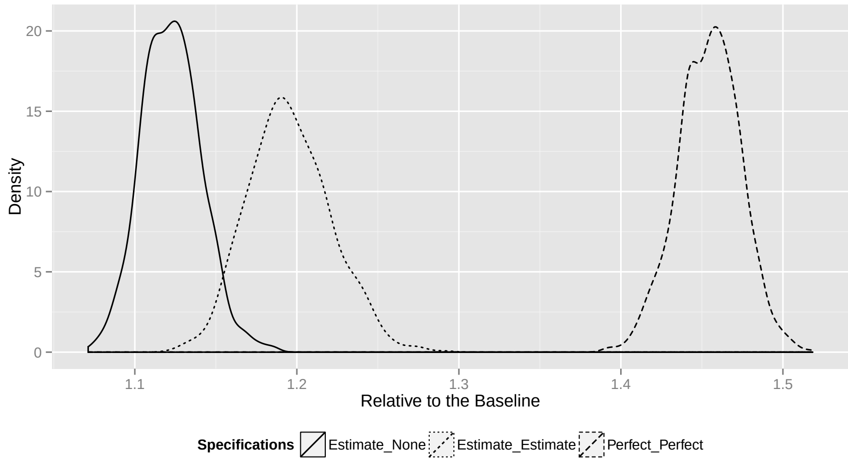

* I simulate 1000 paths for each local market under different parameter settings, compute the total revenue of the network, and divide the total revenues by the revenue baseline for each path.

Table 1.5: Expected Total Revenue Relative to the Baseline

Learning, Init. Belief	Estimate, None	Estimate, Estimate	Perfect, Perfect
Residual	1.2633	1.4190	1.9699
Entire	1.1226	1.1956	1.4545

* I simulate 1000 paths for each local market, compute the total revenue of the network for each path, and take the average of the total revenues to estimate the expected total revenue of the network under different degrees of informational friction. The expected revenues are divided by the revenue baseline. The revenue baseline is the revenue when no information processing is exerted and the initial beliefs are those of the rational expectation. The columns correspond to the different settings of the parameter: the middle column is the revenue under the actual estimates, the left column is the revenue when the contribution of the agent's initial belief is dropped, and the right column is the revenue when the agents knew the demand parameters from the beginning. The upper row is for the revenues from the slots for which the agents can decide assortments, and the lower row is the revenue from the entire slots.

To compute the expected total revenue of the retail network under different degrees of informational friction, I first simulate 1000 paths of sales and assortments under i) the baseline and ii) the actual parameter estimates. The distribution in the middle of Figure 1.8 is the distribution of total revenues under the actual parameter estimates relative to the revenue baseline. The middle column of Table 1.6 is the expected total revenue under the actual estimates relative to the revenue baseline. The upper row is the number for the slots for which agents are responsible, and the bottom row is the number for the entire slots. Table 1.6 shows that the agents' information processing and private information increases expected total revenue by 41.9% for the slots for which the agents are responsible and 19.6% for all slots. Thus, agents non-trivially contribute to the network by learning demand and adjusting the assortments.

However, this is far from perfect. I simulated 1000 paths of sales and assortments under the assumption that the local demand parameters are know from the beginning. The right distribution of Figure 1.8 is the distribution of the total revenues under the assumption relative to the revenue baseline. The distribution is far above the distribution under the actual parameter estimates. The right column of Table 1.6 is the expected total revenue under the assumption relative to the revenue baseline. The upper row shows that the network could increase the total revenue by 97.0% for the slots for which agents are responsible, and by 45.5% for all slots, if the local demand parameters are known from the beginning. This is the maximum attainable level of revenue in the current environment.

Next, to study the contributions of agents' information processing and initial belief, I simulate 1000 paths of sales and assortments under the assumption that the initial beliefs are at the baseline but the signal precisions are at the actual parameter estimates. The left distribution of Figure 1.8 is the distribution of the total revenues under the assumption relative to the revenue baseline. The left column of Table 1.6 gives the expected total revenue under the assumption relative to the revenue baseline. The upper row shows that the agents' information processing only increases the expected total revenue by 26.3% for the slots for which agents are responsible and by 12.3% for the entire slots. This implies that the agents' information processing and initial belief contributed to expected total revenue by 12.3% and 7.3% for all slots, respectively.[21]

1.5.2 Quantifying Informational Friction Favouring Delegation of Product Assortment Decisions

In this section, I study the implications of the estimation results for the principal's choice of the decentralised decision structure.

[21] We should carefully interpret this result. If I additionally assume that the agents' initial belief captures all the relevant private information of the agent, then the contribution of the initial belief is equal to the contribution of the agent's private information. If not, the contribution of the information processing quantified above captures the contribution of the learning from the sales data and some other potential private signals that is informative concerning the local demand and correlated with the sales data.

Figure 1.9: Simulated Distribution of the Total Revenue Relative to the Baseline

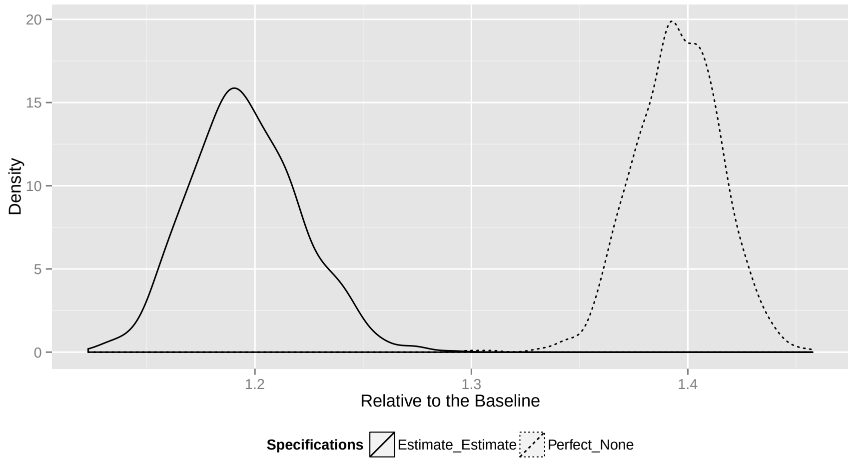

* I simulate 1000 paths for each local market under different settings of the parameter, compute the total revenue of the network and divide the total revenues by the revenue baseline for each path.

Table 1.6: Expected Total Revenue Relative to the Baseline

Learning, Init. Belief	Estimate, Estimate	Perfect, None
Residual	1.4190	1.8435
Entire	1.1956	1.3951

* I simulate 1000 paths for each local market, compute the total revenue of the network for each path, and take the average of the total revenues to estimate the expected total revenue of the network under different degrees of informational friction. The expected revenues are divided by the revenue baseline. The revenue baseline is the revenue when no information processing is exerted and the initial beliefs are those of the rational expectation. The columns correspond to the different parameter settings: the left column is the revenue under the actual estimates, and the right column is the revenue when the initial beliefs are those of the rational expectation, but the sales data are processed precisely. The upper row is the revenue from the slots for which the agents decide assortments, and the lower row is the revenue from all slots.

In theory, the principal company can process the sales data and decide all the product assortments without delegating the decisions to agents. However, as observed, the principal company delegates part of the product assortment decisions to agents based in local markets. There can be two reasons for this choice. First, there may be some information about the demand other than the sales data, and the information is dispersed across locations. Moreover, this information may be 'soft', that is, hard to communicate. Then, the principal may want to delegate decisions to exploit this private information at agents' hand. Second, it may be costly to develop and utilise information processing capability for the principal. Then, for a given allocation of information processing capabilities, it may be better for the principal to divide the big problem into small tasks, distribute the tasks across multiple agents, and solve them concurrently. This configuration may be better than increasing the principal's capability to process all the sales data and optimise product assortments.

I first ask whether I can rationalise the delegation without considering the second factor, the principal's costs for developing and utilising information processing capability. If not, I discern how large the costs must be to rationalise delegation. I can answer these questions by comparing expected total revenues under the actual policy and the hypothetical centralised assortment policy that only uses sales and covariates (temperatures) data but not the agents' initial beliefs. If the hypothetical centralised assortment revenue exceeds the actual revenue, I cannot rationalise the delegation. Because gross profits are roughly proportional to the revenue in their business, the costs for developing and processing information processing capability relative to the baseline of gross profits must be as large as the difference in the expected total revenues relative to the baseline of revenue between two policies.

Figure 1.9 displays the distributions of total revenues relative to the baseline of revenue under the actual and hypothetical assortment policies. Figure 1.9 shows that the distribution under the hypothetical assortment policy is above the distribution under the actual assortment policy. Table 1.8 shows the expected total revenues relative to the revenue baseline under the actual and hypothetical assortment policies. The upper row of the right column of the table shows that the hypothetical centralised assortment policy could increase the expected total revenue by 84.4% for the slots for which agents are responsible. From the bottom row, I can see that it could increase the expected total revenue by 39.5% for the entire slots. These increases are greater than the contributions of agents under the actual policy, 41.9% and 19.6%. Therefore, I cannot rationalise the delegation without considering costs that prevent the principal company from employing the hypothetical assortment policy. The costs must be as large as 39.5 - 19.6 = 19.9% of the baseline gross profit.

1.6 Concluding Remarks

This paper studied product assortment decisions of agents working for a beverage vending machine network in Tokyo. An uncertain environment concerning underlying demand parameters exists because of the release of new products. I found that agents have some informative initial belief about the demand parameters and are learning the demand parameters from the sales data. This informative initial belief and learning together increase the expected total revenue by 19.6% than the revenue baseline where no informative belief exists and no information processing of the sales data is exerted. Agent's informative beliefs and information processing contribute 7.3% and 12.3%, respectively. However, I also showed that this is far from the perfect: the expected total revenue could be 45.5% higher than the baseline if the demand parameters are known from the beginning. This raises concerns with the assumption that firms know the demand parameters when we recover marginal costs from the observed prices and the estimated demand functions. Finally, I found that the principal could increase expected total revenue by 39.5% if sales data are precisely analysed even if their initial belief is no more informative than the rational expectation. This indicates that there are costs for the principal company in developing and utilising information processing capabilities to centralise decisions and that the costs should be at least as large as the 19.9% of the baseline of gross profits.

This paper studies only one season of decisions of the beverage vending machine network. It is worth investigating how the speed of learning evolves over time and how growth is related to past experience. For example, an agent who worked in a large market or in a market that requires more effort for signal extraction may develop faster. Another limitation of the current paper is that it relies on a specific modelling assumptions of the agent decision process. I intend to conduct field experiments that change the assortment policy to an algorithmic policy to assess the validity of the last analysis in a future study. Although my analysis is specific to a beverage vending machine in Tokyo, the same argument can be applied to other retail formats such as supermarkets. Including pricing decisions in addition to product assortment decisions are conceptually straightforward and computationally not difficult, because pricing decisions typically can be solved easier than product assortment decisions requiring combinatorial optimization.

Chapter 2

Recommending (Un)popular Products: A Field Experiment using Vending Machines

2.1 Introduction

Product recommendations have become a major marketing strategy that companies use to induce consumer attention and increase purchase probability. Online retailers such as Amazon and Netflix use sophisticated machine-learning algorithms to recommend products. Product recommendations are also effective offline communication channels. In many purchasing occasions, such as choosing electronic appliances, financial products, and medical/health-related products, product recommendations affect the consumer's choice.

Although product recommendations attract substantial managerial and academic attention, measuring the effects of recommendations is not straightforward because online recommendation systems typically suggest already popular products. That is, recommendations based on popularity may simply reinforce the position of already popular products and the effect of recommendation is difficult to accurately estimate.

To overcome this endogeneity problem, we execute a field experiment designed to measure the causal impact of product recommendations. We use two treatments in our field experiment: recommending popular products (PP) and unpopular products (UP). We conduct the experiment using 459 vending machines selling beverages that are located in train stations across the Tokyo metropolitan area. The vending machines are typically placed on platforms and busy station corridors, and consumers are often in a hurry to make a purchase during rush hour. These circumstances may potentially influence the recommendation effectiveness. Thus, we design the experiment to measure the effects of treatment at different times of the day.

The vending machine we use has a particular feature that offers recommendations and is equipped with an a electronic touch-screen panel in front with a built-in camera on top. As a customer approaches, the machine is programmed to make recommendations according to age and sex based on the image captured by the camera. The machine recommends different products to customers with different characteristics at different times of the day. A business man in the morning, for example, may receive a recommendation to buy an energy drink or coffee, whereas a female teenager may receive a recommendation to purchase a bottle of mineral water in the afternoon.

Using this feature, we study the effect of PP and UP at different times of the day. Recommendations can strengthen consumer recognition or brand loyalty of already popular products, but recommendations can also be used to increase consumer attention to products with a small market share. Our design provides insights on the effects of time pressure or social impact, which are mostly examined in a lab in existing psychology literature. For instance, individuals tend to purchase according to routine in the morning when they are busy (under time pressure) and there are many other people around, whereas individuals are more likely to try something new after receiving a recommendation in the daytime when they have more time to consider the product. Thus, product recommendations of small-share products may become more effective during the day rather than the morning rush-hour.

Our field experiment uses 460 vending machines located across the Tokyo metropolitan area for a two-week period. We conduct a field experiment with two product recommendation treatments and a control of no recommendation. The two treatments are "Recommending Popular Products (PP)" and "Recommending Unpopular Products (UP)" together with the control group of "No Recommendation (NR)." We split a day into three periods (morning, afternoon, and evening), and we assign a different set of product recommendations for each period so that we can compare consumer demand and consumer choice behaviour in different situations. For example, NR is used in the morning of Day 1, and PP is assigned in the morning of Day 2. Then, comparing the consumer choice behaviour in the first case and in the second case, we infer the effects of popular product recommendations.

We begin our analysis by examining the impact of recommendation display on total beverage sales for each vending machine. The standard consumer model with limited attention assumes that recommendation increases attention and has no other effects on choice behaviour. Then, recommendations should increase the total sales. However, whether displaying recommendations can convince customers to buy or have any potential side-effects that reduce sales is an empirical question.

Our results show that the recommendations increase sales by 10.2% for UP and 3.7% for PP during the day. Both product recommendation treatments induce customers to buy during the day when stations are less crowded. Moreover, the effect is much stronger for UP than PP. The recommendation only appears after a customer stands in front of a machine; therefore a recommendation does not function as an advertisement to attract more customers to the vending machine. Hence, the result implies that the recommendation convinces customers to buy who otherwise would not purchase and the effect is stronger for UP.

In contrast to the daytime result, total sales decrease by approximately 3% for both PP and UP treatments during the morning rush hour. A potential explanation for the negative effects of the recommendations is a crowding-out effect because of congestion. Recommendation displays may increase the time that each consumer spends before making a decision, preventing potential consumers from purchasing a beverage before a train arrives. This interpretation is consistent with the stronger negative recommendation effects for machines with higher sales.

Next, to understand the effect of product level recommendations, we estimate a simple discrete-choice demand model where consumer preferences depend on product recommendations. We find that product recommendations, on average, increase the choice probability of recommended products among inside goods by 28.9% for popular products and 4.4% for unpopular products.

Interestingly, the effects are heterogeneous at different times of the day. We find that both PP and UP increase choice probability of recommended products by 16.2% and 18.1%, respectively, during the day. The result is similar in the evening, and the choice probability of recommended products increases by 33.3% and 17.3%.

In contrast to the results for daytime and night hours, we find that, in the morning, popular product recommendations increase the likelihood of purchasing those products by 35.1%, whereas unpopular product recommendations decrease by 11.9%. One interpretation of this result considering that choice probabilities decrease with recommendation in the morning is the following: product recommendation causes consumers to consider product choice longer before making a decision, but the recommended products may not necessarily be chosen during morning rush hour if the products are unpopular.

An interpretation based on consumer psychological motivation arises. The opposite effects of popular and unpopular product recommendations during morning rush-hour may indicate an effect of the presence of other passengers. The pattern is consistent with the finding of Nowlis (1995) that consumers under time pressure are more likely to choose top-of-the-line products. Additionally, a decrease in sales in the morning is consistent with the result of Dhar and Nowlis (1999) where time-pressured consumers choose to defer choice. Another possible explanation is social impact theory (SIT), which implies that a large presence of people in close proximity affects choice. Our result is consistent with Argo, Dahl, and Manchanda (2005)'s finding that consumers choose certain products to impress others and do so to a greater extent when more people are in close proximity.

The reminder of the paper is organised as follows. Section 2 explains the industry background and our experimental design. Section 3 presents the data obtained from the experiment and some preliminary analyses on the effects of recommendations on overall sales. We discuss our econometric model in Section 4, and the results and their managerial implications are explained in Section 5. Lastly, Section 6 concludes.

2.2 Industry Background and Experimental Design

2.2.1 The Vending Machine Industry

The vending machine is one of the main beverage sales channels in Japan; almost one-third of total beverage sales in Japan are generated by the vending machine industry.[1] There is one vending machine for every 50 people, and there were approximately 2.6 million vending machines in 2013.

We conduct a field experiment with a large beverage vending machine company (the company, hereafter), which is a subsidiary company of the largest railroad transportation company in Japan. The train company operates mainly in the Tokyo metropolitan area with more than 1,700 train stations and an average of 16 million daily passengers across all stations. The company owns 9,600 vending machines and places them in the train stations. Annual sales in 2013 were $260 million with an average annual growth rate of 30%. Most of the company vending

[1]Based on *Inryo Brand Book 2011* edited by Inryo Souken, 35% of total beverage sales are from vending machines, 36% from supermarkets, and 20% from convenience stores.

Figure 2.1: Product Recommendation

* The product recommendation (red bubble with the word “Recommended” that flashes) appears based on age and sex of the customer.

machines sell only beverages, but a small fraction of the vending machines sell some food such as snacks and fruits. This study solely focuses on beverage vending machines.

The 460 vending machines we use for the field experiment have several unique features. First, there is a large electronic touch-panel screen on the front of the machine and consumers purchase a product by touching their choice of product. Each vending machine offers approximately 40 different products, but some products, particularly popular products, occupy more than one column. Second, and most importantly, the vending machine is equipped with a camera in its front panel. When a consumer stands in front of a vending machine to purchase a beverage, the camera recognizes the consumer and captures consumer characteristics such as gender and age. Based on the characteristics of the consumer, the vending machine is programmed to make product recommendations based on the consumer's characteristics.[2] Consumers can easily identify which products are recommended through colourful flashing pop-ups displayed as the customer stands in front of the vending machine (see Figure 2.1).

The set of products recommended to consumers is pre-determined by the company and

[2]Because of privacy concerns, the company cannot collect the information collected by the camera, including consumer characteristics. The cameras are used only to detect consumer characteristics to make product recommendations.

depends on the consumer's age, gender, and the recommendation changes according to the time of day. Currently, the company does not have the ability to change the product recommendation list at the individual machine level. The set of available products carried by each vending machine, however, varies across stations and vending machines because each agent that manages machine has some discretion over product selection for each vending machine (see the previous chapter for details). Therefore, the variation in the number of recommendations and the set of recommended products for each vending machine is substantial.

The recommendation system that uses may function differently from the recommendation systems that online companies such as Amazon.com and Netflix use on their website. Amazon.com or Netflix have developed a complicated algorithm based on collaborative filtering and/or machine learning for product recommendation and use consumers' web-browsing history and past purchase information to predict their next-product-to-buy. In our context, the company maintains a large database that contains customers' past purchase information and demographic information similar to the databases of Amazon.com and Netflix, but the company currently does not customise the recommendation system for individual consumers using the customer information in the database. Thus, the company's recommendation system is purely for information provision. This simple system allows us to estimate the effect of recommendations that are otherwise difficult to tease out. The experimental design explained in the next section exploits the setting simplicity to infer the effect recommendations.

2.2.2 Experimental Design

To measure the effects of recommendation systems on total sales and consumer choices, we conduct a large field experiment using 460 vending machines. The experimental design used for measuring the effects of the recommendations is as follows:

We use two different treatments. Treatment PP is the normal system that recommends relatively popular products. In treatment UP, the set of recommended products are not popular. No product is recommended for the control group (denoted NP). In the standard vending machines without recommendations, products recommended in PP, on average, sell 14.5% more than products recommended in UP after adjusting the consumer size of vending machines during the experiment period. We execute these three treatments for three different times over two weeks. Morning is before 10 o'clock, daytime is between 10 and 18 o'clock, and night is after 18 o'clock. [3]

[3] Since the data is proprietary, it is not possible to disclose which products are included in the set of recom-

	15	16	17	18	19	22	23	24	25	26
	Mon	Tue	Wed	Thur	Fri	Mon	Tue	Wed	Thur	Fri
Morning	PP	UP	PP	UP	PP	NR	PP	UP	NR	PP
Daytime	UP	PP	UP	NR	UP	PP	UP	NR	PP	UP
Night	PP	UP	NR	PP	NR	UP	NR	PP	UP	NR

* We conduct the experiment with two different treatments, PP and UP, as well as a control, NR. In Treatment PP, the same set of products that the company is using is recommended. In treatment UP, unpopular products are recommended. No product is recommended for control NR.

Table 2.1: Experimental Design

The identification strategy is to compare the outcome from treatment PP (or treatment UP) and from the control (NP) conditional at the same time of day. By comparing the morning sales of Day 1 for which treatment PP is implemented with the mornings sales of Day 5 in which control NP is implemented, we infer the effect of the popular product recommendations on sales. Hence, the underlying assumption for identification is that consumer demand for beverages is invariant over a week conditional on the time of day and other observed characteristics, such as temperature. A limitation of our experimental design is that there is no consumer level randomisation in showing recommendation pop-ups, although such a randomisation could be the ideal way to identify the recommendation effects. We do not adapt such a strategy because of the technical limitations of the program that governs the recommendation system.

2.3 Data

In this subsection, we report the summary statistics of the vending machine data. In Table 2.2, we report two sets of summary statistics: station-level and vending machine-level characteristics.

The first six rows show the summary station-level characteristics. There are 187 train stations in our sample, and average daily sales (the number of cans and bottles sold at a station) are $4,400$, but there is a substantial variation across stations. The average number of products sold at a station is 39.2, and all train stations in the sample have at least 15 products. The field experiment uses an average of 2.5 vending machines per station, although some large stations have as many as 30 vending machines.[4]

mendation products.

[4]We use only 460 machines out of 9,600 vending machines operated by the company because the reminder of vending machines do not have the recommendation system. Hence, there is a much larger number of vending

	Obs	Mean	Std. Dev	Min	Max
Station Characteristics					
Daily # of Sales	187	$4,400.91$	$6,245.78$	105	$60,828$
# of Products	187	39.22	10.86	15	81
# of Machines	187	2.46	3.03	1	30
Average Temperature	187	25.97	1.56	20	29
Maximum Temperature	187	30.92	2.60	22	36
Minimum Temperature	187	22.55	1.52	18	26
Machine Characteristics					
Daily # of Sales	460	$1,789.07$	806.39	105	$5,442$
# of Products	460	30.64	2.58	15	39

* We report the number of observations, mean, standard errors, minimum and maximum of each characteristic.

Table 2.2: Summary Statistics of Stations and Vending Machines

From the fourth to sixth rows, we report the summary statistics for average temperature at the station level, daily maximum temperature, and daily minimum temperature. We obtain the detailed daily weather information from the Japan Meteorological Agency and match it to each station. Tokyo is typically hot and humid in July, but there is some temperature variation within a day and across stations. The average temperature is approximately 26 degree celsius, and average daily maximum temperature greater than 30 degrees celsius. This variation in temperature is likely to affect consumer demand for beverage

The last two rows report the summary vending machine characteristics. Average daily sales per machine are approximately $1,789$, which is approximately $200,000$ Japanese yen ($2,000 US dollars). The best selling vending machines sell products worth over $10,000,000$ Japanese yen ($1 million US dollars) per year. The average number of available products carried by a machine is approximately 30.6, which is more than the maximum number of slots per machine because some products (typically popular products such as mineral water or green tea) occupy more than one slot.

machines in each station.

If a consumer purchases beverages with a commuter card, then the customer ID will be recorded in the vending machine. In our sample, there are 384,762 uniquely identified consumers. The total number of sales from these customers is 798,300 during the experiment period. Among identified customers, only a fraction agreed to provide their demographic information such as age and gender. In total, demographic information for 37,144 unique consumers is available. We use only the data of consumers with available demographics for the analysis in Section 5, while use the entire sample for the analysis in Section 4 because demographic information is necessary to match the recommended products in treatment PP. This may cause a selection problem. However, the average age and the gender distribution in our sample are similar to the population averages. The average age of the consumers 40, and approximately 70% are male, while they are 37.4 and 66% in the population.[5]

Figure 2.2 shows the available products' price distribution across the entire vending machines. The price distribution ranges from 100 Japanese yen to 200 Japanese yen, but approximately 70% of products are priced at either 120 Japanese yen or 150 Japanese yen. Beverages contained in a small can (350 *ml*) typically cost 120 Japanese yen, whereas beverages in a large plastic bottle (500 *ml*) are sold for 150 Japanese yen. Some seasonal beverages with special flavor or taste are sold at a higher price than normal products, say 200 Japanese yen. Some products that cost less than 120 Japanese yen are sold only to the company's employees at special locations such as the company's office. The Figure 2.2 illustrates that limited price variation across products, indicating that the price of each product is mainly determined by rule of thumb. This vending machine industry characteristics causes a minor price endogeneity problem compared to other situations.

Figure 2.3 illustrates market share by category. The company classifies products into eleven categories. The largest category is soda, which accounts for 20% of products. The mineral water category is the second largest and accounts for approximately 12% of products. The company has three categories for different types of tea including green tea, black tea, and other tea (that is, jasmine tea or oolong tea), and total market share is over 25%. This sample distribution is similar to that of the entire beverage industry including other channels such as supermarkets and convenience stores.[6] Hence, consumer preference in our sample may not differ significantly

[5] The parent train company conducted a large-scale survey to understand passenger characteristics. See https://www.jreast.co.jp/development/tech/pdf_16/Tech-16-21-26.pdf

[6] The market share by category for the beverage industry is as follows: soda (18%), tea (28%), water (14%), coffee (15%). See the following link for details, http://www.ccwest.co.jp/pdf/ir/annualreview/ccw/an_2012_06.pdf

Figure 2.2: Price Distribution

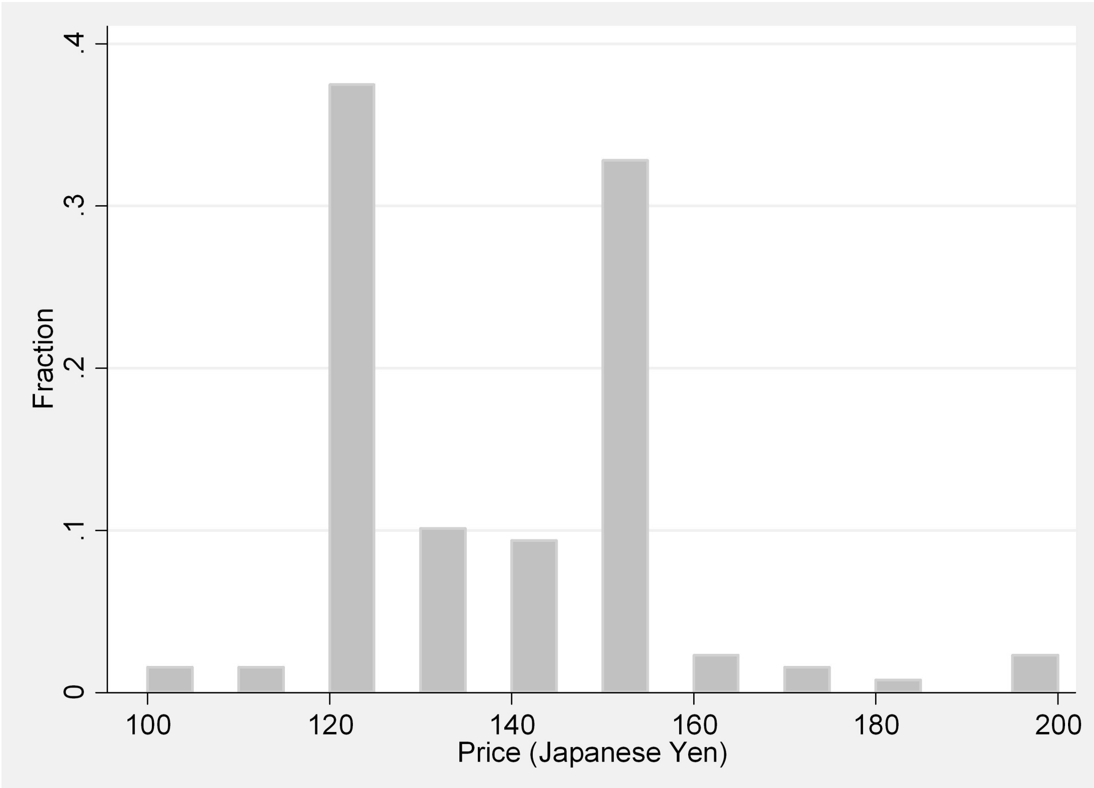

* The graph shows the probability distribution of the price of a beverage. The unit is Japanese Yen. Approximately $1 = 100 Yen.

from that of the entire beverage market.

2.4 Effects on Sales

We first examine the effects on total sales in order to design the optimal recommendation system. The standard consumer model with limited attention assumes that recommendation increases attention and has no other effects on choice behaviour. Then, recommendations should increase the total sales. However, whether displaying recommendations can convince customers to buy or have any potential side-effects that reduce sales is an empirical question.

We run regressions separately for three different timings, morning, daytime, and night. We also run regressions separately for treatment PP and treatment UP. We conduct the following panel linear regression model with vending machine fixed effects for each subsample to see how

Figure 2.3: Market Share by Category

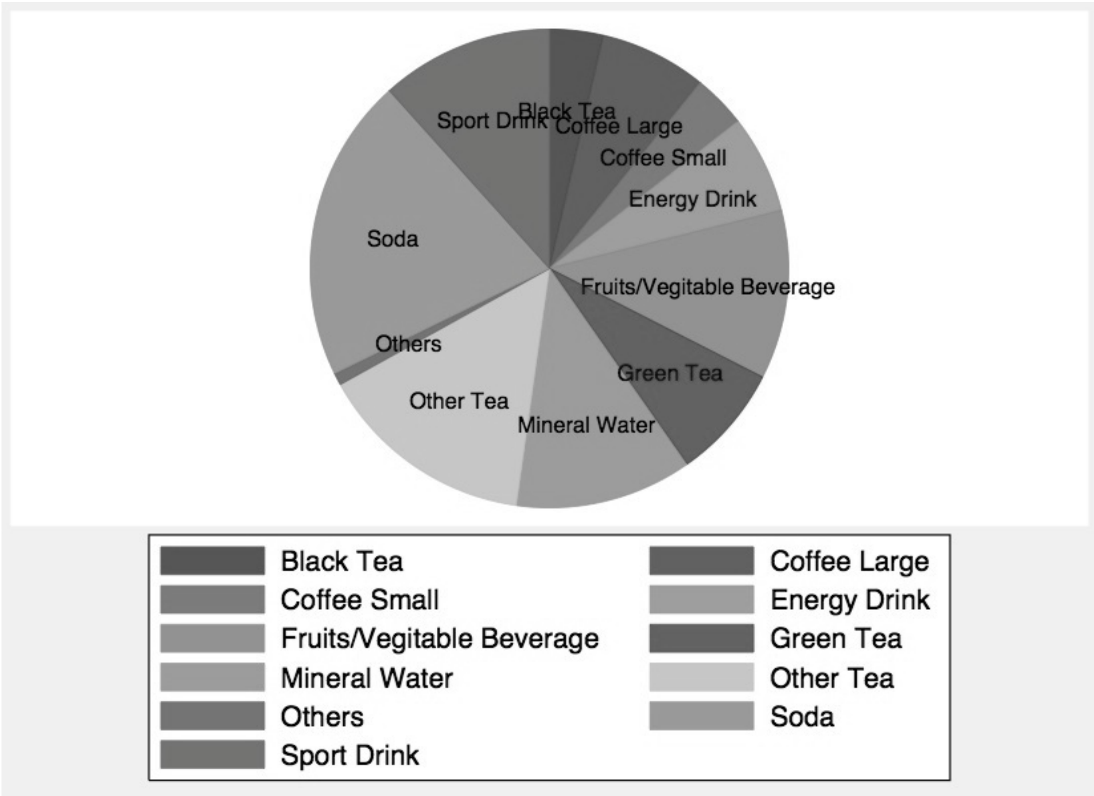

* The figure shows market shares in terms of the number of cans/bottles sold for each category. The company categorises more than 200 products into 11 categories.

sales of beverages are influenced by the recommendations:

$$\log(sales_{kt}) = \alpha_1 d_{kt} + \alpha_2 \times temp_{kt} + \alpha_3 \times temp_{kt}^2 + \mu_k + u_{kt}, \tag{2.1}$$

where $sales_{kt}$ is the total number of cans sold at vending machine k at period t, d_{kt} is a dummy variable indicating whether vending machine k is in one of the treatments, $temp_{kt}$ is the temperature for vending machine k at period t, μ_k is a vending machine fixed effect, and u_{kt} is an idiosyncratic random sales shock. Period t corresponds to the time of day (morning, daytime, and night) of a particular date, for example, the night of July 21.

The estimation results of Eq (2.1) are shown in Table 2.3. All regression results show that the model fits the data very well. The R^2 of all regression results are greater than 0.885.

Table 2.1 shows that, first, the effect of the recommendations on sales varies significantly

	Morning		Daytime		Night	
	PP	UP	PP	UP	PP	UP
Treatment	-0.0299^{***}	-0.0290^{***}	0.0370^{***}	0.1024^{***}	-0.1007^{***}	-0.0003
	(0.0075)	(0.0073)	(0.0086)	(0.0094)	(0.0118)	(0.0096)
Max Temp	0.0298^{***}	0.0103^{***}	0.0316^{***}	-0.0163^{***}	-0.0473^{***}	-0.0362^{***}
	(0.0030)	(0.0029)	(0.0031)	(0.0031)	(0.0039)	(0.0040)
(Max Temp)^2	0.0003^{*}	-0.0001	0.0005^{***}	0.0023	0.0031^{***}	0.0024^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0031)	(0.0001)	(0.0001)
FE	X	X	X	X	X	X
R^2	0.927	0.917	0.908	0.914	0.899	0.885
Obs	2,547	2,522	2,497	2,084	2,046	2,857

* The table show the results from the fixed effect linear regression estimates for Eq (2.1). We estimate the model by time of day (morning, day, and night). The numbers in parenthesis are standard errors of the estimates. *** indicates that the estimate is significant at 0.1%, ** 1%, and * 5%.

Table 2.3: Effects of the Recommendations on Sales

depending on when the recommendations are displayed.

During the day, our results show that the recommendations increase sales by 10.2% for UP and 3.7% for PP. Both product recommendation treatments induce customers to buy during the day, when stations are less crowded. Moreover, the effect is much stronger for unpopular products. The recommendation only appears after a customer stands in front of a machine. Therefore, a recommendation does not function as an advertisement to attract customers to the vending machine. Hence, the result implies that the recommendation convinces customers to purchase who would not purchase otherwise and the effect is stronger for UP.

In contrast to the daytime results, total sales decrease by approximately 3% under both PP and UP treatments during the morning rush hour. There are a number of potential explanations for this result. One explanation is a crowding-out effect because of congestion. Recommendation displays may increase the time that each consumer spends before making a decision and prevent potential consumers from purchasing a beverage before a train arrives. This interpretation is consistent with the stronger negative effect for machines with higher sales as we see below.

Another explanation is time pressure. Customer behaviour literature (for instance, Dhar and Nowlis (1999); Suri and Monroe (2003)) finds that consumers may defer choices or choose not to purchase under time pressure. The consumers we study are likely to be under time pressure, particularly during rush hour. The vending machines we use for the field experiment

	PP		UP	
	Small	Large	Small	Large
Morning	-0.025^{*}	-0.035^{***}	-0.024^{*}	-0.033^{***}
	(0.012)	(0.009)	(0.012)	(0.009)
Daytime	0.035^{*}	0.039^{***}	0.102^{***}	0.104^{***}
	(0.015)	(0.010)	(0.016)	(0.011)
Night	-0.105^{***}	-0.097^{***}	-0.005	-0.006
	(0.019)	(0.015)	(0.015)	(0.0121)

* Large indicates the vending machines with the top 50% of sales, while small indicates the machines with the bottom 50% of sales. The table shows the results from fixed effect linear regression estimates for Eq (2.1) including temperature and its square as controls. The numbers in parenthesis are standard errors of the estimates. *** indicates the estimate is significant at 0.1%, ** 1%, and * 5%.

Table 2.4: Effects of the Recommendations on Sales by Vending Machine Type

are mainly placed on train stations and, typically, on crowded stations where passengers transfer from one line to another line. The average number of daily passengers using the Tokyo Station is at least $400,000$,[7] and many other stations in the data have at least $50,000$ daily passengers. Hence, the platforms in these stations are extremely crowded during rush hour, and a consumer could be forced to wait behind a consumer who is making a purchasing decision. Additionally, trains are arriving and leaving constantly, which creates pressure for consumers. However, an extensive study of the effects of time pressure is beyond the scope of this paper.

Table 2.4 presents the results by the type of vending machine. We separate the vending machines into 'large' machines that have the top 50% of sales and to 'small' machines that have the bottom 50% of sales. The results show that the negative effect of a recommendation in the morning rush hour is stronger for 'large' machines. This result is consistent with the two possible explanations we discussed above because 'large' machines are typically in more crowded locations. The coefficients between 'small' and 'large' machines for daytime and night are not statistically different from one another.

Lastly, the effect of daily maximum temperature on sales is significantly positive. Since our experiment was conducted during the warmest times of the year in Tokyo, product assortments

[7] The number includes only the passengers who start their journey from the station. There are many other passengers who change trains at stations who are not included in this number.

are all cold drinks.

These regression results have some managerial implications. First, the company should consider potential negative effect that the recommendations may create. The recommendations that are intended to help consumers' decision making may have unintended consequences, and sales may decrease because of the congestion effect. The company should adjust its recommendation strategy to reduce such negative effects. In our context, the company should stop using the recommendation system during rush hour. Second, recommending products with different degree of popularity has different effect on sales. Our result imply that sales increase more when the company recommends relatively unpopular products rather than relatively popular products.

2.5 Effects on Consumer Choice Behaviour

2.5.1 Model

To see the effects of the recommendation on consumer choices, we estimate a discrete-choice demand model for beverages, for which recommendations can directly affect consumer choice probability. We start with a model with no consumer heterogeneity to examine the overall effects of the recommendations on choice probabilities by product popularity. Consumer i chooses a product from all of the available products in a vending machine. The set of available products is denoted by J_t. The set of available products is different across vending machines and across time. Product j's characteristics are denoted by x_{jt} and the price by p_{jt}. Hence, the indirect utility that consumer i obtains from purchasing product j is

$$u_{ijt} = -\alpha p_{jt} + x'_{jt}\beta + \delta^{PP} d_{jt}^{PP} + \delta^{UP} d_{jt}^{UP} + \xi_{jt} + d_c + \varepsilon_{ijt}, \tag{2.2}$$

where d_{jt}^{PP} is an indicator variable taking the value of 1 if product j is recommended in treatment PP at time t and d_{jt}^{UP} is similarly defined. The term ξ_{jt} is product j's fixed effect, d_c is the category fixed effect (which is ignored when product dummies are included), and ε_{ijt} is an idiosyncratic preference shock following a type-1 extreme value distribution. As in the typical discrete-choice demand model, ξ_{jt} and ε_{ijt} are unobserved to the econometrician. Lastly, the set of parameters is $\theta \equiv (\alpha, \beta, \delta^{PP}, \delta^{UP})$.

We estimate the model with heterogeneous effects for the time of day. In doing so, we interact d_{jt}^{PP} and d_{jt}^{UP} with the timing dummies (morning, daytime, and night) as follows;

$$u_{ijt} = -\alpha p_{jt} + x'_{jt}\beta + \sum_{s=AM,PM,N} \sum_{l=PP,UP} \delta_s^l \left(d_{jt}^l \times d_s\right) + \xi_{jt} + d_c + \varepsilon_{ijt}, \tag{2.3}$$

where d_s are the dummy variables for the different time of day, that is, $s \in \{AM, PM, N(ight)\}$. The set of parameters is now denoted by $\theta \equiv (\alpha, \beta, \delta_{AM}^{PP}, \delta_{AM}^{UP}, \delta_{PM}^{PP}, \delta_{PM}^{UP}, \delta_{N}^{PP}, \delta_{N}^{UP})$.

We estimate the parameters by the maximum likelihood estimation. Assuming ε_{ijt} follows a type-I extreme value distribution, we get

$$\Pr(d_{ijt} = 1; \theta) = \Pr(u_{ijt} \geq u_{ij't}, \forall j') = \frac{\exp(v_{jt})}{\sum_{k \in J_t} \exp(v_{kt})},$$

where $v_{jt} = -\alpha p_{jt} + x'_{jt}\beta + \delta_A d_{jt}^A + \delta_B d_{jt}^B + \xi_{jt}$ for the homogeneous model, and $v_{jt} = -\alpha p_{jt} + x'_{jt}\beta + \sum_{s=AM,PM,N} \sum_{l=PP,UP} \delta_s^l \left(d_{jt}^l \times d_s\right) + \xi_{jt} + d_c$ for the heterogeneous model. Then, the likelihood function to be maximized is written as

$$L(\theta) = \prod_i \prod_j \prod_t \Pr(d_{ijt} = 1; \theta).$$

2.5.2 Results

Base Results

We report our estimates of the homogeneous effect discrete choice model of Eq (2.2) in Table 2.5. The table gives our estimates of the price effect (α), treatment PP's effect (δ^{PP}), treatment UP's effect (δ^{UP}).

The homogeneous model of Eq (2.2) is estimated in five different specifications. In model (1), we include product category dummies and five variables discussed above. In model (2), we also include dummy variables for each product.

We begin our discussion with the estimates of δ^{PP} reported in the first row. The effect is strongly positive and statistically significant in both specifications, which implies that popular product recommendations induce consumer to select the recommended products.

The δ^{UP} estimate in the right column show that unpopular product recommendations also positively influence choice probability, but the effect less pronounced than it is for popular products.

To understand the estimates more quantitatively, Table 2.6 reports the odds ratios of the main treatment effects (δ^{PP} and δ^{UP}). Both models show similar results. For popular product recommendations, the likelihood of choosing the recommended product increases by more

	(1)	(2)
δ^{PP}	0.339***	0.254***
	(0.0198)	(0.0204)
δ^{UP}	-0.0380	0.0428*
	(0.0231)	(0.0254)
N	87,682	87,682
FE	Categories	Products

* The table reports the estimation results of the homogeneous specification in equation (2.2). The numbers in parenthesis are standard errors of the estimates. *** indicates the parameter estimate is significant at 0.1%, ** at 1%, and * at 5%.

Table 2.5: Estimation Results of Homogeneous Specification

	(1)	(2)
δ^{PP}	1.404***	1.289***
	(0.0278)	(0.0263)
δ^{UP}	0.963	1.044*
	(0.0223)	(0.0265)
FE	Categories	Products

* The table reports the odds ratio of the treatment effects δ^{PP} and δ^{UP}. The numbers in parenthesis are standard errors of the estimates. *** indicates the parameter estimate is significant at 0.1%, ** at 1%, and * at 5%.

Table 2.6: Odds Ratios of Homogeneous Specification

than 28.9%. For the unpopular product recommendations, the likelihood of the recommended products being chosen increases by approximately 4.4%.

	(1)	(2)
$\delta^{PP}_{Morning}$	0.552***	0.301***
	(0.0278)	(0.0288)
δ^{PP}_{Day}	0.211***	0.150***
	(0.0371)	(0.0383)
δ^{PP}_{Night}	0.0186	0.288***
	(0.0459)	(0.0470)
$\delta^{UP}_{Morning}$	-0.133***	-0.127***
	(0.0366)	(0.0383)
δ^{UP}_{Day}	0.0451	0.167***
	(0.0411)	(0.0433)
δ^{UP}_{Night}	0.00355	0.160***
	(0.0417)	(0.0451)
Observations	87,682	87,682
FE	Categories	Products

* The table presents the estimation results for the main treatment effects of the model (2.3). The numbers in parenthesis are standard errors of the estimates. *** indicates the estimate is significant at 0.1%, ** at 1%, and * at 5%.

Table 2.7: Estimation Results of Heterogeneous Specification

Results by Time of Day

Table 2.7 shows the estimates of Eq (2.3) that considers the heterogeneous effects by different time of day. We obtain estimates for the price effect (α), treatment PP's effect by time of day (δ_t^{PP}), treatment UP's effect by time of day (δ_t^{UP}). As in Table 2.5, we have two models: Model (1) includes product category dummies and variables discussed above. Model (2) includes dummy variables for each product.

The results show that popular product recommendations have positive effects across all times of day, and the effect is strongest during the morning rush hour and weakest during the daytime ($\delta^{PP}_{Morning} > \delta^{PP}_{Night} > \delta^{PP}_{Day} > 0$).

	(1)	(2)
	choice	choice
$\delta^{PP}_{Morning}$	1.738***	1.351***
	(0.0483)	(0.0389)
δ^{PP}_{Day}	1.235***	1.162***
	(0.0458)	(0.0446)
δ^{PP}_{Night}	1.019	1.333***
	(0.0468)	(0.0626)
$\delta^{UP}_{Morning}$	0.875***	0.881***
	(0.0320)	(0.0337)
δ^{UP}_{Day}	1.046	1.181***
	(0.0430)	(0.0512)
δ^{UP}_{Night}	1.004	1.173***
	(0.0419)	(0.0528)
FE	Categories	Products

* The table reports the odds ratios of the treatment effects for each of the heterogeneous effect model (2.3). The numbers in parenthesis are standard errors of the estimates. *** indicates the estimate is significant at 0.1%, ** at 1%, and * at 5%.

Table 2.8: Odds Ratio of the Heterogeneous Specifications

The pattern is quite different for unpopular product recommendations. Unpopular product recommendation lowers the choice probability in the morning of the recommended products ($\delta^{UP}_{Morning} < 0$). In the evening, the effect is positive but the magnitude is smaller than PP ($\delta^{PP}_{Night} > \delta^{UP}_{Night} > 0$). In contrast, the effect is positive and the magnitude is slightly larger than PP in the daytime. ($\delta^{UP}_{Day} > \delta^{PP}_{Day} > 0$).

Table 2.8 reports the odds ratio of the main treatment effects for both models in Table 2.8. In Model (2), the choice probability increases 35.1% in the morning for treatment PP, while it decreases by 11.9% for treatment UP. The magnitudes of the daytime increase are not very different between treatments PP and UP: an increase of 16.2% for PP and 18.1% for UP for

	PP		UP	
	Small	Large	Small	Large
Morning	0.330^{***}	0.285^{***}	-0.0975	-0.143^{***}
	(0.0471)	(0.0353)	(0.0612)	(0.069)
Daytime	0.195^{***}	0.126^{***}	0.128^{*}	0.186^{***}
	(0.0623)	(0.0472)	(0.0722)	(0.0521)
Night	0.300^{***}	0.280^{***}	0.110	0.190^{***}
	(0.0754)	(0.0588)	(0.0700)	(0.0553)

* Crowded indicates the vending machines with top 50% of the sales, while "Not Crowded" indicates the ones with bottom 50% sales. We estimate the model (2.3) with product fixed effects. The numbers in parenthesis are standard errors of the estimates. *** indicates the estimate is significant at 0.1%, ** at 1%, and * at 5%.

Table 2.9: Results by Vending Machine Type

model (2). The increase for the evening is higher under PP than under UP by approximately 33.3% and 17.3% for model (2).

Results by Vending Machine Type

The result that treatment UP significantly lowers the choice probabilities in the morning suggests a natural hypothesis that unpopular product recommendations under time pressure or in a crowded place could negatively affect choice. Although we cannot directly test this hypothesis, we can study how the results differ based on the types of vending machine. Consumers at vending machines with greater sales can be considered to experience more pressure, whereas consumers at vending machines with fewer sales are less likely to be under pressure. Thus, we estimate the same model as the previous model by dividing the sample into 'large' and 'small' where 'large' vending machines have sales greater than the machine with median sales, and 'small' machines have sales below the median.

Table 2.9 presents the results by vending machine type. The results show that the negative effects of treatment UP in the morning appear only statistically significantly on large vending machines. But the test of equality between the small vending machine effect and large vending machine effect is not rejected at 10% level.

	PP Small Female	PP Small Male	PP Large Female	PP Large Male	UP Small Female	UP Small Male	UP Large Female	UP Large Male
Morning	0.298^{***}	0.342^{***}	0.174^{***}	0.333^{***}	0.237^{**}	-0.278^{***}	0.172^{**}	-0.334^{***}
	(0.0891)	(0.0552)	(0.0644)	(0.0419)	(0.0970)	(0.0778)	(0.0713)	(0.0606)
Daytime	0.093	0.237^{***}	0.235^{***}	0.073	0.340^{***}	0.023	0.359^{***}	0.090
	(0.117)	(0.0799)	(0.0739)	(0.0576)	(0.0973)	(0.0657)	(0.0718)	(0.0531)
Night	0.474^{***}	0.248^{***}	0.305^{**}	0.273^{***}	0.384^{***}	0.015	0.535^{***}	0.052
	(0.151)	(0.0865)	(0.121)	(0.067)	(0.127)	(0.0817)	(0.0938)	(0.0660)

* We estimate the model (3) with both the product fixed effects. The numbers in parenthesis are standard errors of the estimates. *** indicates the estimate is significant at 0.1%, ** at 1% and * at 5%.

Table 2.10: Results by Customer Gender

Results by Customer Gender

An additional question we can ask is who drives the pattern we found above. This is also an interesting question in itself because we can study the difference in effects of recommendation by demographic characteristics. The result can have implications for customer group targeting of recommendations.

Table 2.10 presents the results by gender. The results show that the negative effects of treatment UP in the morning is attributed to male customers. The effect of UP in the morning is positive and statistically significant for both 'large' and 'small' vending machines for female customers. The result implies that managers should not recommend unpopular products to male customers especially when they are in time pressure, whereas they should aggressively interact with female customers to try unpopular products.

Discussion

Several key managerial implications emerge from the empirical findings. First, product recommendations, on average, can increase the likelihood of consumers choosing a recommended product regardless of its popularity. These findings confirm the two competing theories concerning the effect of product recommendations. Recommendations can reinforce the position of already popular products. However, recommendations also allow consumers to find products that they would otherwise ignore. Our findings indicate that both theories are the case as even unpopular product recommendations have a positive impact on consumer choice probabilities.

Second, the effect of recommendations on consumer choices is substantially heterogeneous according to the time of day. Moreover, the combination of which product to recommend and consumer characteristics plays a significant role. Recommendation of popular and unpopular products has a significantly different effect across time. Companies can exploit this characteristics by effectively modifying recommendations for targeted consumers and time.

In addition to these managerial implications, our finding that the effect differs depending on the time of day may indicate an effect of the time pressure or the pressure from the consumers around. This provides an interesting insight to the consumer behaviour literature like SIT (see, e.g., Griffit and Veitch (1971); Langer and Saegert (1977)). SIT theorises that just the presence of others affects a consumer's emotions and behaviour when making a purchase decision, even if another person is not interacting with the customer. A recent paper in the literature by Argo et al. (2005) finds that consumers choose more expensive products when others are around. A psychological explanation would be consumer desire to be regarded highly by others. Our finding that the effect of PP and UP are opposite during the morning rush hour (when there are many people around) while they are similar for other times of day is consistent with SIT.

2.6 Conclusion

This paper examines the effect of product recommendations on consumer behaviors by conducting a field experiment using a large number of a new type of vending machine in Tokyo. Our setting is unique because the new vending machines recommend different beverages for different consumers based on consumer characteristics, which are recognised by a video camera attached to the machine. We show that total sales decrease in the morning when consumers receive the recommendations, while total sales increase during the daytime. This may be caused by a congestion effect that the recommendations unintentionally create, that is, each single consumer spends more time making a purchase decision, which crowds out other consumers.

We demonstrate how consumer choice probabilities are influenced by recommendations. Our findings indicate that recommendations induce consumer choice for recommended products, but the effects are heterogeneous. Particularly, recommending unpopular products in the morning reduces choice probabilities among males, although unpopular product recommendations always increase among women.

There are some limitations to the study. First, this study does not allow us to infer potential mechanisms that lead to the results. The negative effect of unpopular product recommendations in the morning, for example, may be driven by some cognitive biases. Investigating similar problems might suggest useful product recommendation designs. Second, this paper does not address the long-term effects of product recommendations on consumer behaviors. Whether recommendations can stimulate consumer learning of new products is not well-known. We propose these subjects as valuable future research.

Chapter 3

Testing Rationality Without Restricting Heterogeneity

3.1 Introduction

The assumption of rationality forms the core of economics. Testing the assumption validates the empirical analysis. The classical axioms of revealed preference (Samuelson, 1938; Houthakker, 1950; Richter, 1966; Varian, 1982) provide the basic framework to test this assumption but do not address the heterogeneity inevitable in empirical studies. In this paper, I propose a non-parametric test for rationality allowing for any type of heterogeneity across decision makers. The test is easy to compute and functions in any choice situation.

Existing empirical studies introduce heterogeneity as an additive error in the choice function and apply the axiom of revealed preference to the mean of the observed choice functions (Varian, 1985; Blundell et al., 2003, 2008). However, Lewbel (2001) shows that rationalisability at the individual level does not imply rationalisability of the mean or vice versa unless the heterogeneity in the population is strictly restricted.

McFadden (2005) introduces the notion of stochastic rationalisability. Given a set of utility functions $\mathcal{U}$, a sequence of observed choice probabilities is considered to be stochastically $\mathcal{U}$-rationalisable if some probability law over $\mathcal{U}$ exists that can induce the observed choice probabilities as a result of utility maximisation with random utilities following that law. The author also derives the necessary and sufficient condition for stochastic rationalisability. This notion allows us to test the rationality without restricting heterogeneity.

The concept is sufficiently general that it is effective for any class of utility function but it is often computationally demanding to check all the conditions. Kitamura and Stoye (2013) circumvent this problem by restricting attention to the case with a finite number of linear budgets. I resolve this issue by focusing on a necessary condition for rationalisability.

I derive a necessary condition for stochastic rationalizability by a set of utility functions with a unique maximizer, which I name SARSP. I obtain the condition by sharpening the weak Axiom of revealed Stochastic preference (WARSP) proposed by Bandyopadhyay, Dasgupta, and Pattanaik (1999). The SARSP is a generalisation of the classic strong axiom of revealed preference (SARP). I propose a non-parametric test of rationalisability based on the SARSP and derive its asymptotic properties. The test can be implemented at low computational cost. The testing procedure is based on the bootstrap test for functional inequalities of Kitagawa (2010).

Monte Carlo simulation shows that the test has an empirical size below the nominal level and relatively strong power. Finally, I apply this method to the British FES. The SARSP is not rejected at the 1% level of significance for this dataset.

3.2 Setting

Let X be a universal metric space of the possible objects of choice, and $\mathcal{B}$ be the Borel σ-algebra of X. Let I be a set of choice situations and $\{B^i\}_{i\in I}, B^i \in \mathcal{B}$ be a set of choice sets for each choice situation $i \in I$.

A leading example is the standard consumer's problem. In this case, $X = \mathbb{R}^J_+$ is the consumption set, and a pair of price vector p^i and an income w^i, (p^i, w^i) is a choice situation. Given choice situation i, I can define the associated budget set B^i by

$$B^i = \{y \in \mathbb{R}^J_+ : \langle p^i, y\rangle \leq w^i\}. \tag{3.1}$$

Unless otherwise stated, general $(X, \mathcal{B})$ and $\{B^i\}_{i\in I}$ are considered below.

For each $B^i, i \in I$, let $\mathcal{B}^i$ be the Borel σ-algebra of the subsets of B^i. Let Π^i be the probability measure on the measurable set $(B^i, \mathcal{B}^i)$. For $C \in \mathcal{B}^i$, the interpretation of $\Pi^i(C)$ is the probability that the choice in situation $i \in I$ belongs to C. $\Pi = \{\Pi^i\}_{i\in I}$ is called the set of *choice probabilities.*

Let $(\mathcal{U}, \mathcal{A}(\mathcal{U}))$ be a hypothetical set of utility functions and the Borel σ-algebra of subsets of $\mathcal{U}$. For example, $\mathcal{U}$ would be a set of utility functions which represent locally non-satiated weak orders on the consumption set. For each $U \in \mathcal{U}$ and $i \in I$, let $d^i(U) \equiv \text{argmax}\{U(y) : y \in B^i\}$ denote a demand function associated with the utility function U. Let $\triangle(\mathcal{U})$ be the class of probability measures on the measurable space $(\mathcal{U}, \mathcal{A}(\mathcal{U}))$.

I define a special class of utility functions $\mathcal{U}_0$ and focus on this class in the following sections. Given a system of choice situations $\{B^i\}_{i\in I}$, I define $\mathcal{U}_0$ as a set of utility functions that achieves a unique maximum at every choice situation B^i in $\{B^i\}_{i\in I}$. The definition of $\mathcal{U}_0$ depends on the system of choice situations.

When discussing testability, it is assumed that $\{\Pi^i\}_{i\in I}$ is known by econometricians. Stochastic $\mathcal{U}$-rationalisability is defined in terms of choice probabilities as below:

Definition 1 (Stochastic $\mathcal{U}$-Rationalisability)**.** *A set of choice probabilities $\{\Pi^i\}_{i\in I}$ is stochastically rationalised by a class of utility functions $\mathcal{U}$ or is stochastically $\mathcal{U}$-rationalisable if a probability measure $\xi \in \triangle(\mathcal{U})$ exists such that*

$$\Pi^i(C) = \xi\{U \in \mathcal{U} : d^i(U) \in C\}, \forall i \in I, \forall C \in \mathcal{B}^i. \tag{3.2}$$

The rationalisability concept of standard revealed preference theory can be considered as a restrictive version of stochastic rationalisability if a deterministic choice function is identified with the following degenerate choice probability:

Definition 2 (Degenerate choice probabilities)**.** *A choice probability Π^i is degenerate if $x \in B^i$ exists such that $\Pi^i(A) = 0$ for any A such that $y \notin A \in \mathcal{B}^i$, and $\Pi^i(A) = 1$ for any A such that $y \in A \in \mathcal{B}^i$.*

3.3 Axioms of Revealed Stochastic Preference

This section discusses the axioms of revealed stochastic preference. The first subsection introduces the necessary and sufficient condition for general stochastic $\mathcal{U}$-rationalisability ($\mathcal{U}$-ARSP). The second subsection derives a necessary condition for the stochastic $\mathcal{U}_0$-rationalizability (SARSP).

3.3.1 $\mathcal{U}$-Axioms of Revealed Stochastic Preference ($\mathcal{U}$-ARSP)

McFadden (2005) provides the necessary and sufficient condition for stochastic $\mathcal{U}$-rationalisability of choice probabilities, which is called $\mathcal{U}$-axiom of revealed stochastic preference ($\mathcal{U}$-ARSP).

Given I, let $\mathcal{M}(I)$ and $\mathcal{M}^*(I)$ be a set of ordered subsets of I allowing repetition and not allowing repetition, respectively. For example, given $I = \{1, 2, 3, 4\}$, $\mathcal{M}(I)$ includes $\{1, 2\}, \{1, 2, 3, 3\}$, $\{1, 3, 2, 3\}$ and so on. $\mathcal{M}^*(I)$ does not include the second and third elements in the above example. $\mathcal{M}(I)$ is infinite but $\mathcal{M}^*(I)$ is finite if I is finite.

Definition 3 ($\mathcal{U}$-axiom of revealed stochastic preference ($\mathcal{U}$-ARSP))**.** *A set of choice probabilies $\{\Pi^i\}_{i\in I}$ satisfies the axiom of revealed stochastic preference with respect to the hypothetical utility functions $\mathcal{U}$ or $\mathcal{U}$-ARSP if for any $M \in \mathcal{M}(I)$ and for any $C^i \in \mathcal{B}^i, i \in M$,*

$$\sum_{i\in M} \Pi^i(C^i) \leq \sup_{U\in\mathcal{U}} \sum_{i\in M} \mathbf{1}(d^i(U) \in C^i). \tag{3.3}$$

We refer to a combination of $M \in \mathcal{M}(I)$ and $\{C^i\}_{i\in M}$ as a *trial*.

Theorem 1 (McFadden (2005), Theorem 5.2)**.** *Suppose Π^i is a finitely additive choice probability on $(B^i, \mathcal{B}^i)$ for each $i \in I$. Then $\mathcal{U}$-ARSP is necessary and sufficient for stochastic $\mathcal{U}$-rationalizability of a set of choice probabilities $\{\Pi^i\}_{i\in I}$.*

The necessity part of this theorem is straightforward and is from the following:

$$\sum_{i\in M} \Pi^i(C^i) = \sum_{i\in M} \xi(d^i(U) \in C^i) \leq \sup_{U\in\mathcal{U}} \sum_{i\in M} \mathbf{1}(d^i(U) \in C^i). \tag{3.4}$$

To show sufficiency, McFadden (2005) first applies the Hahn-Bannach theorem to show that $\mathcal{U}$-ARSP implies existence of some probability measure $\eta \in \triangle(\mathcal{U})$ such that for all $i \in I$ and $C^i \in \mathcal{B}^i$, $\Pi^i(C^i) \leq \eta(d^i(U) \in C^i)$ holds. Then, because $\Pi^i(C^{ic}) \leq \eta(d^i(U) \in C^{ic})$ also holds, $\Pi^i(C^i) = \eta(d^i(U) \in C^i)$ implying stochastic $\mathcal{U}$-rationalisability.

This $\mathcal{U}$-ARSP is sufficiently general that it can test any specification of random utility model *provided one can implement optimisation over the hypothetical spaces of $\mathcal{U}$ efficiently*: $\mathcal{U}$-ARSP can be written as follows:

$$\sup_{M\in\mathcal{M}(I)} \left\{ \sum_{i\in M} \Pi^i(C^i) - \sup_{U\in\mathcal{U}} \sum_{i\in M} \mathbf{1}(d^i(U) \in C^i) \right\} \leq 0, \tag{$\mathcal{U}$-ARSP}$$

This optimisation problem is not solvable in general. If the dimension of the hypothetical space of utility function $\mathcal{U}$ is finite, for example, if it is a set of Cobb-Douglas utility functions, the above optimization problem may be implemented. This will give a useful specification test and is worth pursuing in future research.

3.3.2 Strong Axiom of Revealed Stochastic Preference (SARSP)

This section focuses on $\mathcal{U}_0$-rationalisability and shows that a condition that is necessary for stochastic $\mathcal{U}_0$-rationalisability can be explicitly written down. This condition is weaker than $\mathcal{U}_0$-ARSP but offers a test that can be implemented at low computational cost. The condition is called SARSP because it is a stochastic generalisation of the classic SARP.

In the following, let $y^i R y^j$ and $y^i R^* y^j$ denote"y^i is weakly revealed preferred to y^j", and"y^i is indirectly revealed preferred to y^j", respectively.

Bandyopadhyay et al. (1999) suggest a condition called WARSP and show that it is necessary for stochastic $\mathcal{U}_0$-rationalisability. I obtain SARSP by reinforcing WARSP. The intuition behind WARSP and the relation between WARSP and $\mathcal{U}_0$-ARSP are carried over to SARSP. I begin with a discussion of WARSP.

Definition 4 (WARSP). *For all i and j in I and for all $A \subset B^i \cap B^j$, $\Pi^j(A) - \Pi^i(A) \leq \Pi^i(B^i \setminus B^j)$.*

Theorem 2 (Bandyopadhyay et al. (1999) , Proposition 3.5). *Suppose a set of choice probabilities $\{\Pi^i\}_{i \in I}$ is stochastically $\mathcal{U}_0$-rationalisable then, it satisfies WARSP, but the converse does not hold.*

First, WARSP implies the classical WARP. Suppose that choice probabilities are degenerate. Now identify degenerate choice probabilities as a choice function. Let y^1 and y^2 be degenerate points of B^1 and B^2 such that $y^1 R y^2$ and $y^1 \neq y^2$. Then, $y^2 \in B^1 \cap B^2$. Because $y^1 \neq y^2$, take $A \subset B^1 \cap B^2$ such that $y^1 \notin A$ but $y^2 \in A$. Then, for any such A, $\Pi^2(A) = 1$. By WARSP, $\Pi^1(B^1 \cap B^2 \setminus A) = 0$. Thus, $y^1 \notin B^1 \cap B^2$, that is $\neg y^2 R y^1$: WARP holds.

The intuition behind WARSP is as follows: suppose the choice probability of A increased by a shift in situations from i to j. Comparing two choice situations, alternatives in $B^j \setminus B^i$ are in the choice set, and alternatives in $B^i \setminus B^j$ are not in the choice set in situation j. New alternatives in the choice set do not raise the choice probability of the existing choices in A. Therefore, all the increase in the choice probabilities of A must be attributed to the change of behaviour of those who were choosing alternatives in $B^i \setminus B^j$ in situation i, and the increase in the choice probabilities in A must be no more than the choice probability of $B^i \setminus B^j$.

$\mathcal{U}_0$-ARSP actually implies WARSP. This argument clarifies another aspect of WARSP and helps to generalise it to SARSP which I will define later. Consider $M = \{i, j\}$ and $\{C^i, C^j\} = \{A, B\}$ with $A \subset B^i \cap B^j$ and $B = [B^i \cap B^j] \setminus A$. Then, $\mathcal{U}_0$-ARSP implies $\Pi^i(B) + \Pi^j(A) \leq \max_{U \in \mathcal{U}} [\mathbf{1}(d^i(U) \in B) + \mathbf{1}(d^j(U) \in A)]$. By rearranging, I get $\Pi^j(A) - \Pi^i(A) \leq$

$\max_{U\in\mathcal{U}}[\mathbf{1}(d^i(U)\in B)+\mathbf{1}(d^j(U)\in A)]-[\Pi^i(A)+\Pi^i(B)]=\max_{U\in\mathcal{U}}[\mathbf{1}(d^i(U)\in B)+\mathbf{1}(d^j(U)\in A)]-\Pi^j(B^i\cap B^j)$. The subsets A and B are disjoint and feasible under both i and j. Hence, the utility function such that $d^i(U)\in B$ does not satisfy $d^j(U)\in A$, and the utility function such that $d^i(U)\in A$ does not satisfy $d^j(U)\in B$. Therefore, $\max_{U\in\mathcal{U}}[\mathbf{1}(d^i(U)\in B)+\mathbf{1}(d^j(U)\in A)]=1$ and so $\Pi^j(A)-\Pi^i(A)\leq 1-\Pi^i(B^i\cap B^j)=\Pi^i(B^i\setminus B^j)$.

This argument illustrates the derivation of the axiom for stochastic $\mathcal{U}$-rationalisability. The inequality in $\mathcal{U}$-ARSP can bind only if the trial in question, $M\in\mathcal{M}(I)$ and $\{C^i\}_{i\in M}$, includes a sequence of choice sets all of which cannot be attained by any utility function in $\mathcal{U}$. In the case of stochastic $\mathcal{U}_0$-rationalisability, all the preferences are represented by some $U\in\mathcal{U}_0$ that attains a unique maximum in every $B\in\mathcal{B}$ and is transitive. WARSP (and the original WARP) exploits the property of the preferences having a unique maximum in every $B\in\mathcal{B}$. SARSP is obtained by exploiting transitivity consistent with Richter (1966) to obtain SARP from WARP.

The derivation of WARSP from $\mathcal{U}_0$-ARSP relies on the condition $\max_{U\in\mathcal{U}}[\mathbf{1}(d^i(U)\subset A)+\mathbf{1}(d^j(U)\subset B^i\cap B^j\setminus A)]=1$. This and $\mathcal{U}_0$-ARSP implies $\Pi^i(A)+\Pi^j(B^i\cap B^j\setminus A)\leq 1$, WARSP. This argument can be generalised to reinforce the restriction to be SARSP. Now SARSP is defined as:

Definition 5 (SARSP)**.** *For any $M\in\mathcal{M}^*(I)$ which is relabeled as $M=\{1,\cdots,n\}$, and for any $A\subset B^n\cap B^1$, the following inequality holds:*

$$\Pi^2(B^1\cap B^2)+\cdots+\Pi^n(B^{n-1}\cap B^n\setminus A)+\Pi^1(A)\leq n-1. \tag{3.5}$$

I check that SARSP generalises WARSP and SARP. The following propositions state this property of SARSP.

Proposition 1. *SARSP implies WARSP.*

Proposition 2. *Suppose a set of choice probabilities $\{\Pi^i\}_{i\in I}$ are degenerate and identify a degenerate choice probability as a deterministic choice function. Then, SARSP implies SARP.*

That SARSP implies WARSP can be checked by investigating the conditions of SARSP when $|M|=2$. That SARSP implies SARP when choice probabilities are degenerate can be checked using the same method employed to show that WARSP implies WARP. I state my main theorem in this section.

Theorem 3. *SARSP is necessary for stochastic $\mathcal{U}_0$-rationalisability.*

SARSP exploits transitivity in addition to the unique maximum property of the preferences. It is straightforward to obtain the necessity of SARSP from $\mathcal{U}_0$-ARSP. To see this, denote $U^*(C) \equiv \max_{c\in C} U(c)$ given a utility function U. Fix some $U \in \mathcal{U}_0$, and consider the corresponding demand function d_U. Suppose $d^i(U) \in B^{i-1} \cap B^i$ for $i = 1, \cdots, n$. This implies $U^*(B^1) \geq \cdots \geq U^*(B^n)$ holds. Moreover, suppose that $d^n(U) \in B^{n-1} \cap B^n \setminus A$ for $A \subset B^1 \cap B^n$. This implies that $U^*(B^n) > U^*(A)$ because U attains a unique maximum. Thus, $d^1(U) \in A$ cannot hold since this implies $U^*(B^n) > U^*(A) = U^*(B^1)$ and contradicts $U^*(B^1) \geq \cdots \geq U^*(B^n)$. Thus, the value of the right-hand side of $\mathcal{U}_0$-ARSP must be no greater than $n-1$. This gives the inequality condition in SARSP showing the necessity of SARSP for stochastic $\mathcal{U}_0$-rationalisability.

When additional restriction on the system of choice situations and observations exist, WARSP and SARSP can be reduced to simpler conditions. For example, consider the case where $B^i \equiv \{y \in \mathbb{R}^J_+ : \langle p^i, y\rangle \leq w^i\}$. Let $\overline{B}^i \equiv \{y \in \mathbb{R}^J_+ : \langle p^i, y\rangle = w^i\}$ and $\breve{B}^i \equiv \{y \in \mathbb{R}^J_+ : \langle p^i, y\rangle < w^i\}$. Suppose that $\Pi^i(\overline{B}^i) = 1$ for all $i \in I$. This holds if the utility functions are locally non-satiated. Or, if price vectors p and commodity vectors y are observed but income w was not, $w \equiv \langle p, y\rangle$. In this case, WARSP reduces to finitely many inequality conditions such that for all i and j in I and for all $A \subset B^i \cap B^j$, $\Pi^j(A\cap\overline{B}^j) - \Pi^i(A\cap\overline{B}^i) \leq \Pi^i(\overline{B}^i \setminus B^j)$. Then, there are only two candidates for A: $B^i \cap \overline{B}^j$ or $\breve{B}^i \cap \overline{B}^j$. Thus, the WARSP reduces to finite inequality conditions that for any $(i,j) \in I\times I$, $\max\{\Pi^j(B^i\cap\overline{B}^j) - \Pi^i(\overline{B}^i\cap\overline{B}^j), \Pi^j(\breve{B}^i\cap\overline{B}^j)\} \leq \Pi^i(\overline{B}^i\setminus B^j)$. The latter term in the max operator gives tighter restriction if and only if $\Pi^j(\overline{B}^i\cap\overline{B}^j) - \Pi^i(\overline{B}^i\cap\overline{B}^j) < 0$. If Π^i are continuous for all $i \in I$, the max operator disappears because both terms give the same value.

The same argument holds with SARSP. Consider again the case where $B^i \equiv \{y \in \mathbb{R}^J_+ : \langle p^i, y\rangle \leq w^i\}$ and $\Pi^i(\overline{B}^i) = 1$ for all $i \in I$. As in the case of WARSP, SARSP reduces to the condition that for any subset $A \subset B^n \cap \overline{B}^1$ the following inequality holds:

$$\begin{aligned} &\Pi^2(B^1 \cap \overline{B}^2) + \cdots + \Pi^n(B^{n-1} \cap \overline{B}^n \setminus A) + \Pi^1(A \cap \overline{B}^1) \leq n-1, \\ &\Leftrightarrow \Pi^2(B^1 \cap \overline{B}^2) + \cdots + \Pi^1(A \cap \overline{B}^1) - \Pi^n(A \cap \overline{B}^n) \leq n - 2 + \Pi^n(\overline{B}^n \setminus B^{n-1}). \end{aligned} \tag{3.6}$$

There are only two candidates for A: $B^n \cap \overline{B}^1$ or $\breve{B}^n \cap \overline{B}^1$, and SARSP reduces to finitely many inequality conditions that for $(i,j) \in I \times I$,

$$\Pi^2(B^1\cap\overline{B}^2)+\cdots+\max\{\Pi^1(B^n\cap\overline{B}^1)-\Pi^n(\overline{B}^1\cap\overline{B}^n), \Pi^1(\breve{B}^n\cap\overline{B}^1)\} \leq n-2+\Pi^n(\overline{B}^n\setminus B^{n-1}), \tag{3.7}$$

and the max operator disappears if Π^i is continuous for all $i \in I$. The method of Kitamura and Stoye (2013) relies on this additional restriction on the system of choice situations and observations.

Classical revealed preference theory shows that SARP is necessary and sufficient for rationalisability by utility functions with a unique maximum, and GARP is necessary and sufficient for locally non-satiated utility functions. Because I have assumed local non-satiation in the above simplified setting, SARSP implies GARP if I identify deterministic choice with degenerate choice probability. SARSP implies

$$\begin{aligned}&\Pi^2(B^1 \cap \overline{B}^2) + \cdots + \Pi^1(\breve{B}^n \cap \overline{B}^1) \le n - 2 + \Pi^n(\overline{B}^n \setminus B^{n-1}),\\ \Leftrightarrow\ &\Pi^2(B^1 \cap \overline{B}^2) + \cdots + \Pi^n(B^{n-1} \cap \overline{B}^n) + \Pi^1(\breve{B}^n \cap \overline{B}^1) \le n - 1\end{aligned} \tag{3.8}$$

Suppose that $\{\Pi^i\}_{i\in I}$, and let y^i be the degenerate point of $\overline{B}^i$. Suppose two degenerate points exist, for example, $y^1 \in \overline{B}^1$ and $y^n \in \overline{B}^n$ such that $y^1 R^* y^n$, that is, a sequence of degenerate points $y^i \in \overline{B}^i$ with $y^i \in B^{i-1}, i = 2, \cdots, n-1$ exists. Because $\Pi^2(B^1 \cap \overline{B}^2) = \cdots = \Pi^n(B^{n-1} \cap \overline{B}^n) = 1$ by the definition of a degenerate point. Then, by SARSP, we have $\Pi^1(\breve{B}^n \cap \breve{B}^1) = 0$. Therefore, $y^1 \notin \breve{B}^n$. Thus, GARP holds.

3.4 Hypothesis Testing

This section establishes a general testing procedure to test any type of stochastic rationalizability provided the corresponding $\mathcal{U}$-ARSP can be computed efficiently. Whether the following test is feasible in each situation depends on the specification of $\mathcal{U}$. The next section demonstrates that tests based on SARSP and WARSP are computationally feasible.

3.4.1 Estimation of Choice Probabilities

This section establishes an inference method similar to Kitagawa (2010). Let (Y, Q) be i.i.d. observation, where Y is a choice and Q is a choice situation. For example, Q is a pair of price vectors $P \in \{p^1, \cdots, p^I\}$ with $p^i \in \mathbb{R}^J_{++}$, and income $W \in \{w^1, \cdots w^I\}$ with $w^i \in \mathbb{R}_{++}$. Let $\Pi^i(C) = F_{Y|Q}(y \in C | Q = q^i)$. The data generating process is $\Pi = \{\Pi^i(\cdot)\}_{i\in I}$, and F is the probability distribution characterised by a value of (Y, Q). I divide the full sample into I subsamples based on the value of Q. Let N^i be the size of these subsamples. Let $\lambda^i = \mathbb{P}(Q = q^i) > \epsilon$ for some $\epsilon > 0$ and let $\lambda = (\lambda^1, \cdots, \lambda^I)$ and $\hat{\lambda} = (\hat{\lambda}^1, \cdots, \hat{\lambda}^I)$ where $\hat{\lambda}^i = N^i/N$. Note that $N \to \infty$ implies $N^i \to \infty$ and $\widehat{\lambda}^i \to \lambda^i$ for all $i \in I$.

The stochastic rationalisability conditions put restrictions on the choice probabilities Π. This is estimated by an empirical measure:

$$\widehat{\Pi}^i(C) = \frac{1}{N^i}\sum_{n=1}^{N} \mathbf{1}\{Y_n \in C, Q_n = q^i\}. \tag{3.9}$$

My application estimates the choice probability by smoothing over total expenditure. The following argument holds with this estimator also.

3.4.2 Test Statistics

The test statistics for WARSP, SARSP, and $\mathcal{U}$-ARSP based on the estimates of choice probabilities $\widehat{\Pi}(C; p, w)$ are defined as follows:

$$\widehat{T}_{\mathrm{W}} = \max_{(i,j)\in I\times I} \sup_{C\in B^i\cap B^j} \Big(\widehat{\Pi}^j(C) - \widehat{\Pi}^i(C) - \widehat{\Pi}^i(B^i \setminus B^j)\Big), \tag{WARSP}$$

$$\widehat{T}_S = \max_{M\in\mathcal{M}^*(I)} \sup_{A\subset B^n\cap B^1} \Big(\widehat{\Pi}^2(B^1\cap B^2)+\cdots+\widehat{\Pi}^n(B^{n-1}\cap B^n\setminus A)+\widehat{\Pi}^1(A)-(n-1)\Big), \tag{SARSP}$$

$$\widehat{T}_{\mathrm{A}} = \sup_{M\in\mathcal{M}(I)} \left\{ \sum_{i\in M} \widehat{\Pi}^i(C^i) - \sup_{U\in\mathcal{U}} \sum_{i\in M} \mathbf{1}(d^i(U)\in C^i) \right\} \tag{$\mathcal{U}$-ARSP}$$

where the set of events on which supremum is taken should be reduced so that uniform convergence of empirical measures hold. I discuss this below.

For all the statistics, the null hypothesis is $T \le 0$. Let $X_n \rightsquigarrow X$ represent weak convergence of X_n to X. Suppose that $\sqrt{N}(\widehat{T}-T)$ has an asymptotic distribution $\sqrt{N}(\widehat{T}-T) \rightsquigarrow J(\cdot;\Pi,\lambda)$, and let $c_{1-\alpha}(\Pi,\lambda)$ be the $(1-\alpha)$-th quantile of the limit distribution. Then, a procedure which rejects the null if $\widehat{T} - \hat{c}_{1-\alpha}/\sqrt{N} > 0$ yields a test with pointwise asymptotically correct sizes because for every Π satisfying the null $T \le 0$, I have

$$\begin{aligned}
\mathbb{P}(\widehat{T} - \hat{c}_{1-\alpha}/\sqrt{N} > 0) &\le \mathbb{P}(\widehat{T} - \hat{c}_{1-\alpha}/\sqrt{N} > T)\\
&= \mathbb{P}(\sqrt{N}(\widehat{T} - T) > \hat{c}_{1-\alpha})\\
&\underbrace{\rightarrow}_{N\to\infty} 1 - J(c_{1-\alpha}(\Pi,\lambda);\Pi,\lambda) = \alpha.
\end{aligned} \tag{3.10}$$

In the following, I derive the asymptotic distributions of the statistics and illustrate the estimation of the critical value of the distribution of test statistics.

3.4.3 Asymptotic Distribution

This section derives asymptotic distribution of $\widehat{T}_W$. A similar argument holds for $\widehat{T}_S$ and $\widehat{T}_S$. I follow his argument to establish asymptotic distribution of the test statistics.

For every ordered pair $(i,j) \in I$ and $C \in \mathcal{B}^i \cap \mathcal{B}^j$, define functions:

$$\begin{aligned} T_W^{i,j}(C) &\equiv \Pi^j(C) - \Pi^i(C) - \Pi^i(B^i \setminus B^j), \\ \widehat{T}_W^{i,j}(C) &\equiv \widehat{\Pi}^j(C) - \widehat{\Pi}^i(C) - \widehat{\Pi}^i(B^i \setminus B^j). \end{aligned} \tag{3.11}$$

Let $\mathbb{B}^1, \cdots, \mathbb{B}^I$ be some subclass of $\mathcal{B}^1, \cdots, \mathcal{B}^I$, and let $\mathbb{V}^{i,j} \equiv \mathbb{B}^i \cap \mathbb{B}^j$ and $\mathcal{V}^{i,j} \equiv \mathcal{B}^i \times \mathcal{B}^j$ which satisfy the following condition:

Assumption. 1 (Uniform convergence). *For each $i = 1, \cdots, I$, the set-indexed empirical processes $G_{\Pi^i,N^i}(\cdot) \equiv \sqrt{N^i}(\widehat{\Pi}^i(\cdot) - \Pi^i(\cdot))$ converge uniformly in law to tight mean zero Gaussian processes in $l^\infty(\mathbb{B}^i)$:*

$$G_{\Pi^i,N^i}(\cdot) \rightsquigarrow G_{\Pi^i}(\cdot), \tag{3.12}$$

where $Cov(G_{\Pi^i}(C), G_{\Pi^i}(C')) = \Pi^i(C \cap C') - \Pi^i(C)\Pi^i(C')$.

Assumption. 2 (Optimal event). *For every ordered pair $(i,j) \in I \times I$, a non-empty maximizer event class $\mathbb{V}^{i,j}_{\max}$ exists such that*

$$\mathbb{V}^{i,j}_{\max} = \{C \in \mathbb{V}^{i,j} : T_W^{i,j}(C) = \sup_{C' \in \mathcal{V}^{i,j}} \{T_W^{i,j}(C')\}\}. \tag{3.13}$$

Additionally, a non-empty subindex $(I \times I)_{\max} \subset I \times I$ exists such that for every ordered pair of $(i,j) \in (I \times I)_{\max}$, there exists $C \in \mathbb{V}^{i,j}$ such that

$$T_W^{i,j}(C) = \max_{(i,j) \in I \times I} \sup_{C' \in \mathcal{V}^{i,j}} \{T_W^{i,j}(C')\}. \tag{3.14}$$

Assumption. 3 (Existence of maximisers). *For every ordered pair $(i,j) \in I \times I$, events $\widehat{C}^{i,j} \in \mathbb{V}^{i,j}$ and $\widehat{C}^{i,j}_{\max} \in \mathbb{V}^{i,j}_{\max}$ exist such that*

$$\begin{aligned} \widehat{T}_W^{i,j}(\widehat{C}^{i,j}) &= \sup_{C \in \mathbb{V}^{i,j}} \{\widehat{T}_W^{i,j}(C)\}, \\ \widehat{T}_W^{i,j}(\widehat{C}^{i,j}_{\max}) &= \sup_{C \in \mathbb{V}^{i,j}_{\max}} \{\widehat{T}_W^{i,j}(C)\}. \end{aligned} \tag{3.15}$$

For any ordered pair $(i,j) \in I \times I$, define:

$$\widehat{T}_W^{i,j} \equiv \sup_{C \in \mathbb{V}^{i,j}} \{\widehat{T}_W^{i,j}(C)\}. \tag{3.16}$$

$$T_W^{i,j} \equiv \sup_{C \in \mathcal{V}^{i,j}} \{T_W^{i,j}(C)\}. \tag{3.17}$$

Assumption 1 ensures weak convergence of the empirical processes. Assumption 2 implies that $\sup_{C' \in \mathcal{V}^{i,j}} \{T_W^{i,j}(C')\}$ is attained on the class of subsets, $\mathbb{V}^{i,j}$. Assumption 3 implies that the same holds for sample analogues. Now define $\widehat{T}_W$ and T_W as follows:

$$\widehat{T}_W \equiv \max_{(i,j)\in I\times I} \sup_{C \in \mathbb{V}^{i,j}} \{\widehat{T}_W^{i,j}(C)\}. \tag{3.18}$$

$$T_W \equiv \max_{(i,j)\in I\times I} \sup_{C \in \mathcal{V}^{i,j}} \{T_W^{i,j}(C)\}. \tag{3.19}$$

The supremum for the population are taken over a Borel algebra $\mathcal{V}^{i,j}$, whereas for the sample are taken over some subset $\mathbb{V}^{i,j}$ of $\mathcal{V}^{i,j}$. This is because of Assumption 1

Given these assumptions, I have the following theorem:

Theorem 4. *Assume random sampling. Under assumptions 1, 2, and 3,*

$$\sqrt{N}(\widehat{T}_W - T_W) \rightsquigarrow \max_{(i,j)\in(I\times I)_{\max}} \sup_{C \in \mathbb{V}^{i,j}_{\max}} \{G_{T_W^{i,j}}(C)\}, \tag{3.20}$$

where $G_{T_W^{i,j}}(\cdot)$ is the mean zero tight Gaussian process in $l^\infty(\mathbb{V}^{i,j})$ where the covariance function for any $C, C' \in \mathbb{V}^{i,j}$ is,

$$\begin{aligned} Cov(G_{T_W^{i,j}}(C), G_{T_W^{i,j}}(C')) = {} & (\lambda^j)^{-1}[\Pi^j(C \cap C') - \Pi^j(C)\Pi^j(C')] \\ & + (\lambda^i)^{-1}[\Pi^i(C \cap C') - \Pi^i(C)\Pi^i(C')] \\ & + (\lambda^i)^{-1}\Pi^i(B^i \setminus B^j)[1 - \Pi^i(B^i \setminus B^j)], \end{aligned} \tag{3.21}$$

where $l^\infty(\mathbb{V}^{i,j})$ is the space of bounded functions on $\mathbb{V}^{i,j}$ with the sup norm metric.

3.4.4 Estimation of the Critical Value

The previous theorem shows that if the maximisers set is not a singleton, then the asymptotic distribution is not Gaussian. This non-pivotal feature of the asymptotic distribution invalidates the standard bootstrap (Andrews, 2000). A Monte Carlo simulation demonstrates this.

A two-step approach is employed to overcome this problem. In the first step, the maximiser set $(I \times I)_{\max}$ and $\mathbb{V}^{i,j}_{\max}$ are estimated. This is a set of points on which $\widehat{T}_W^{i,j}(C)$ takes a value close to the maximum of $\{\widehat{T}_W^{i,j}(C)\}$. A slackness parameter that determines the cutoff level is chosen to satisfy the appropriate rate. This step corresponds to a moment selection procedure in the inference based on moment inequalities. In the second step, a bootstrap analogue of $\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j}(C))\}$ is constructed and its maximum value over the maximiser set estimated in the first step is chosen. Iterating this procedure, I obtain the bootstrap estimate of the distribution of $\sqrt{N}(\widehat{T}_W - T_W)$.

In this paper, I follow the second approach using bootstrap. Let $\mathbf{Y}^i$ represents the original sample of Y under choice situation i. I compute the critical value from the following algorithm:

Algorithm 1 (Bootstrap)**.** *Input: data* $\{Y_n, Q_n\}_{n\in N}$. *Output:* $\{$*Accept, Reject*$\}$.

1. *Compute* $\widehat{T}^{i,j}_W(\cdot), \widehat{T}^{i,j}_W, \widehat{T}_W$.

2. *Let* η_N *be the slackness sequences that satisfy*

$$\frac{\eta_N}{\sqrt{N}} \to 0, \frac{\eta_N}{\sqrt{\log\log N}} \to \infty, \text{ as } N \to \infty. \tag{3.22}$$

3. *Estimate the maximiser set by*

$$\begin{aligned} \widehat{\mathbb{V}}^{i,j}_{\max} &= \{C \in \mathbb{V}^{i,j} : \sqrt{N}(\widehat{T}_W - \widehat{T}^{i,j}_W(C)) \leq \eta_N\}, \\ \widehat{(I \times I)}_{\max} &= \{(i,j) \in I \times I : \widehat{\mathbb{V}}^{i,j}_{\max} \neq \emptyset\}. \end{aligned} \tag{3.23}$$

4. *For each* $i = 1, \cdots, I$, *sample* N^i *observations from* $\mathbf{Y}^i$ *randomly with replacement to construct* $\widetilde{\Pi}^i(\cdot)$, *the empirical measure based on the bootstrapped sample. Using the constructed* $\{\widetilde{\Pi}^i(\cdot)\}_{i\in I}$, *obtain the bootstrap analogue of* $\widehat{T}^{i,j}_W(\cdot)$,

$$\widetilde{T}^{i,j}_W(C) = \widetilde{\Pi}^j(C) - \widetilde{\Pi}^i(C) - \widetilde{\Pi}^i(B^i \setminus B^j). \tag{3.24}$$

5. *Compute*

$$\max_{\widehat{(I\times I)}_{\max}} \sup_{C\in\widehat{\mathbb{V}}^{i,j}_{\max}} \{\sqrt{N}(\widetilde{T}^{i,j}_W(C) - \widehat{T}^{i,j}_W(C))\}. \tag{3.25}$$

6. *Iterate steps 3 and 4 many times and obtain* $\widetilde{c}_{1-\alpha}$ *as the sample* $(1-\alpha)$*-th quantile of the iterated statistics.*

7. *Reject the null hypothesis* $T_W \leq 0$ *if* $\widehat{T}_W - \widetilde{c}_{1-\alpha}/\sqrt{N} > 0$.

Kitagawa (2010) showed the validity of the above bootstrap procedure. Any choice of slackness parameter η_N satisfying the rate of convergence is asymptotically valid. However, in the small sample, the choice affects the rejection rate of the test. The larger the η_N, the larger the maximisers set. This implies greater bootstrap standard error and a higher rejection rate. When $\eta_N = 0$, the algorithm becomes a naive bootstrap. The adaptive choice of η_N is an open question.

Table 3.1: Generated Normalised Price Vectors for Monte Carlo Simulation

Choice Situation (i)	**Price (p_{i1})**	**Price (p_{i2})**	**Price (p_{i3})**
1	1.31	2.34	1.22
2	1.45	1.07	0.86
3	0.67	1.36	2.15
4	0.62	1.45	1.27
5	0.53	0.62	0.60
6	0.40	1.67	1.21
7	1.02	1.61	0.95
8	0.69	1.82	0.67

3.5 Monte Carlo Simulation

This section examines empirical sizes and powers. When budget sets are linear and all the choices are on the budget surfaces, the conditions to be tested are SARSP with finitely many inequality conditions. The regularity conditions for Theorem 4 trivially holds.

Consider a situation in which there are $J = 3$ goods and $I = 8$ choice situations. This subsection uses a normalised price vector throughout as in Table 3.1, which I generated according to i.i.d. distributions $p_{i1} \sim \Gamma(3, 0.2), p_{i2} \sim \Gamma(4, 0.3)$, and $p_{i3} \sim \Gamma(4, 0.2)$ where $\Gamma(k, \theta)$ is a gamma distribution with shape parameter k and scale parameter θ. There is no serious reason for the choice of these distributions. The number of replications is $K = 1,000$. In each simulation, I iterated bootstrap $L = 1,000$ times. The slackness parameter η_N for estimating the maximisers set is specified at $\eta_N = \log(\log(N))$. I implement the same exercises with $\eta_N = 0, 1\log(\log(N))$ and $10\log(\log(N))$.

We examine both testing procedures based on SARSP and WARSP. The testing procedure based on WARSP has less power than SARSP in theory but is more attractive computationally. As the number of choice situations, I, increases, the inequality conditions to be checked increase at a factorial rate. WARSP only requires inequality conditions checking for ordered pairs (i, j) among I choice situations, and therefore significantly reduces the computational burden when I is large. SARSP, however, requires that the conditions be checked for ordered k-tuples for all $k = 2, \cdots, I$. The conditions of SARSP show that the inequality conditions become harder to violate for higher k. Hence, for a problem with big I, it is logical not to check the conditions

Table 3.2: Null Specifications

Design	p_1	q_1	p_2	q_2	p_3	q_3	k	θ
Null 1	0.0100	0.0100	0.0100	0.0100	0.0100	0.0100	0.0100	0.0100
Null 2	0.0100	15.0000	0.0100	0.0100	15.0000	41.2500	15.0000	1.3747
Null 3	0.0100	0.0112	0.0100	0.9556	0.0118	15.0000	0.0100	14.9982

Table 3.3: Non-Rationalisable Models

Design	$(i = 1, \cdots, 4)$	$(i = 5, \cdots, 8)$	
		$\{\widetilde{\alpha}\}$	s
Not 1	Null 1	Null 3	Null 3
Not 2	Null 1	Null 3	Null 1
Not 3	Null 1	Null 1	Null 3

for all $k = 2, \cdots, I$ but rather to start by checking $k = 2$ and move forward only when the null is not rejected (although, rigorously speaking, sequential testing problem must then be solved). Monte Carlo simulation will show the decrease in power using the test based on WARSP instead of the test based on SARSP.

I check empirical sizes under three rationalisable models and powers under three non-rationalisable models with sample sizes in each choice situation set at $N^i = 1,000$.

Rationalisable models: I consider three CES-type random utility models as null data generating processes. Consumer n chooses a consumption bundle which maximises a CES utility function

$$u_n(c) = \left(\sum_{j=1}^{J} \alpha_{nj} c_j^{\frac{s_n - 1}{s_n}} \right)^{\frac{s_n}{s_n - 1}}. \tag{3.26}$$

The parameters $(\alpha_{n1}, \alpha_{n2}, \alpha_{n3}, s_n)$ represent heterogeneity in the utility function across consumers. Each null model is specified by a distribution of these parameters across consumers. Let $\alpha_{nj} = \widetilde{\alpha}_{nj}/(\widetilde{\alpha}_{n1} + \widetilde{\alpha}_{n2} + \widetilde{\alpha}_{n3})$. I parameterise the distribution of $(\widetilde{\alpha}_{n1}, \widetilde{\alpha}_{n2}, \widetilde{\alpha}_{n3}, s_n)$ as $\widetilde{\alpha}_{nj} \sim B(p_j, q_j)$ and $s_n \sim \Gamma(k, \theta)$ where $B(p, q)$ is a beta distribution with a parameter (p, q). Given a vector of normalised prices in Table 3.1, I compute three specifications of the distribution that are close to the boundary of the null space in the sense that the population T is close to zero.[1] The parameters of each null specification are listed in Table 3.3.

Non-rationalisable models: I consider three non-rationalisable models to examine empirical powers. Each model has the same distribution for situations $i = 1, \cdots, 4$ but is set to have different distributions for $i = 5, \cdots, 8$. In Not 1, the distribution becomes that of Null 3. In Not 2, only the distributions of $\widetilde{\alpha}$'s change to that of Null 3 while the distribution of s_n is fixed. In Not 3, however, only the distributions of $\widetilde{\alpha}$'s are fixed, but the distribution s_n changes to that of Null 3.[2]

The rejection rates under the nulls are listed in Table 3.4. The empirical sizes under the nulls are well controlled for the test with slackness parameters $\eta_N = \log(\log(N))$ and $10\log(\log(N))$. For example, the empirical sizes of the test with $\eta_N = \log(\log(N))$ for Null 3 are 10.3%, 5.2%, and 1.1% when the nominal sizes are 10%, 5%, and 1%, respectively. The empirical sizes are not statistically significantly larger than the nominal sizes. The empirical sizes of the tests with $\eta_N = 10\log(\log(N))$ are smaller than those with $\eta_N = \log(\log(N))$. This is obvious by construction. The empirical sizes of the tests with $\eta_N = 0$ are significantly upwardly biased, which demonstrates how inconsistent naive bootstraps are in this situation. Because the slackness parameter η_N is set close to zero, the result approximates that of the naive bootstrap.

The rejection rates under the alternatives are summarised in Table 3.5. Not 1 and 3 are rejected 100%, whereas Not 2 are hardly rejected. This shows that the power of the test is not uniform, and a deviation to some direction is hard to detect. This is because SARSP is a necessary condition for $\mathcal{U}_0$-rationalisability and can hold under some non-rationalisable models.

Overall, the result shows that the test based on WARSP has similar performance compared to the test based on SARSP. This shows that the test based on WARSP in large I situations, in which the test based on SARSP becomes computationally unattractive. In this exercise with $I = 8$, the computation of the test based on SARSP took only seconds for each replication. However, as I increases, it will become harder to implement.

[1] I implemented the same exercise for other null data generating processes but the rejection rates were entirely zero for all such DGPs because they are in the interior of the null space in the sense that the true T is bounded below zero. For this reason, I compute parameters as in this section to make the DGPs approximately the least favorable.

[2] Similar transitions from Null 1 to Null 2 were 100% rejected for every η_N.

Table 3.4: Simulated Rejection Rate for Rationalisable Models (CES utility, 1,000 resampling, 1,000 MC replications)

Slackness ($\eta_N = \log(\log(N))$)	**Size**	**Null 1**	**(s.e.)**	**Null 2**	**(s.e.)**	**Null 3**	**(s.e.)**
SARSP	0.100	0.022	(0.0046)	0.000	(0.0000)	0.103	(0.0096)
	0.050	0.014	(0.0037)	0.000	(0.0000)	0.052	(0.0070)
	0.010	0.003	(0.0017)	0.000	(0.0000)	0.011	(0.0033)
WARSP	0.100	0.064	(0.0077)	0.000	(0.0000)	0.106	(0.0097)
	0.050	0.028	(0.0052)	0.000	(0.0000)	0.056	(0.0073)
	0.010	0.008	(0.0028)	0.000	(0.0000)	0.013	(0.0036)

Slackness ($\eta_N = 10\log(\log(N))$)	**Size (α)**	**Null 1**	**(s.e.)**	**Null 2**	**(s.e.)**	**Null 3**	**(s.e.)**
SARSP	0.100	0.022	(0.0046)	0.000	(0.0000)	0.011	(0.0033)
	0.050	0.009	(0.0030)	0.000	(0.0000)	0.004	(0.0020)
	0.010	0.003	(0.0017)	0.000	(0.0000)	0.000	(0.0000)
WARSP	0.100	0.056	(0.0073)	0.000	(0.0000)	0.076	(0.0084)
	0.050	0.024	(0.0048)	0.000	(0.0000)	0.034	(0.0057)
	0.010	0.003	(0.0017)	0.000	(0.0000)	0.008	(0.0028)

Slackness ($\eta_N = 0$)	**Size (α)**	**Null 1**	**(s.e.)**	**Null 2**	**(s.e.)**	**Null 3**	**(s.e.)**
SARSP	0.100	0.419	(0.0156)	0.000	(0.0000)	0.603	(0.0155)
	0.050	0.190	(0.0124)	0.000	(0.0000)	0.378	(0.0153)
	0.010	0.024	(0.0048)	0.000	(0.0000)	0.095	(0.0093)
WARSP	0.100	0.416	(0.0156)	0.000	(0.0000)	0.602	(0.0155)
	0.050	0.187	(0.0123)	0.000	(0.0000)	0.378	(0.0153)
	0.010	0.024	(0.0048)	0.000	(0.0000)	0.095	(0.0093)

* (s.e.) are MC standard errors.

Table 3.5: Simulated Rejection Rate for Non-Rationalisable Models (CES utility, 1,000 resampling, 1,000 MC replications)

Slackness $\eta_N = \log(\log(N))$	**Size** (α)	**Not 1**	**(s.e.)**	**Not 2**	**(s.e.)**	**Not 3**	**(s.e.)**
SARSP	0.100	1.000	(0.0000)	0.021	(0.0045)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.010	(0.0031)	1.000	(0.0000)
	0.010	1.000	(0.0000)	0.003	(0.0017)	1.000	(0.0000)
WARSP	0.100	1.000	(0.0000)	0.054	(0.0072)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.027	(0.0051)	1.000	(0.0000)
	0.010	1.000	(0.0000)	0.006	(0.0024)	1.000	(0.0000)

Slackness $\eta_N = 10\log(\log(N))$	**Size** (α)	**Not 1**	**(s.e.)**	**Not 2**	**(s.e.)**	**Not 3**	**(s.e.)**
SARSP	0.100	1.000	(0.0000)	0.020	(0.0153)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.008	(0.0137)	0.998	(0.0014)
	0.010	1.000	(0.0000)	0.002	(0.0099)	0.935	(0.0078)
WARSP	0.100	1.000	(0.0000)	0.054	(0.0153)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.028	(0.0137)	1.000	(0.0000)
	0.010	1.000	(0.0000)	0.014	(0.0099)	0.990	(0.0010)

Slackness $\eta_N = 0$	**Size** (α)	**Not 1**	**(s.e.)**	**Not 2**	**(s.e.)**	**Not 3**	**(s.e.)**
SARSP	0.100	1.000	(0.0000)	0.356	(0.0151)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.151	(0.0113)	1.000	(0.0000)
	0.010	1.000	(0.0000)	0.015	(0.0038)	1.000	(0.0000)
WARSP	0.100	1.000	(0.0000)	0.352	(0.0151)	1.000	(0.0000)
	0.050	1.000	(0.0000)	0.148	(0.0112)	1.000	(0.0000)
	0.010	1.000	(0.0000)	0.015	(0.0038)	1.000	(0.0000)

* (s.e.) are MC standard errors.

Table 3.6: The Retail Price Index for Foods, Non-durables, and Services

Year	**Foods**	**Nondurables**	**Services**
1988	105.8	105.3	105.8
1989	111.9	110.4	112.8
1990	121.0	117.9	122.0
1991	128.8	128.1	133.7
1992	132.9	134.9	142.8
1993	136.6	139.8	150.6
1994	139.2	142.8	155.6
1995	144.8	147.7	158.7
1996	150.0	152.2	162.0
1997	152.3	155.9	168.4
1998	155.8	160.0	175.2
1999	158.8	162.5	183.0

January 1987=100.

3.6 Application

3.6.1 British Family Expenditure Survey

In this section, I apply my method to household consumption data from the British FES for the period 1988 to 1999 and test whether it accepts or rejects SARSP. FES is a repeated cross-sectional survey consisting of approximately $7,000$ households each year. To compare the results with those of Blundell, Browning, and Crawford (2003, 2008), I restrict my sample to couples with children who own a car. This provides approximately $1,500$ observations on average each year. I use data for the period 1988 to 1999. Expenditures for each good are summarised into three categories: food, services and other non-durables. The definition is the same as that of Blundell, Browning, and Crawford (2003, 2008). The relative prices of food, services, and non-durables are quoted and calculated with the annual retail price index and associated weights (Table 3.6). I assume the same RPI for every household.

One problem with my FES application is that the choice situation in FES is continuous because total expenditure is continuous, whereas my method assumes that it is discrete. We discretise expenditures as a compromise to solve this problem. I focus on the subsample with total expenditure near its annual median and recompute the annual expenditures for each category by multiplying the original expenditure shares to the median annual total expenditure. The hypothesis testing does not correct for measurement errors associated with this procedure. I draw subsamples with total expenditure between 40 and 60 percentiles. To check robustness, I also test with subsamples between 45 and 55 percentiles.

We find that SARSP is not rejected with subsamples between the 40 and 60 percentiles in the 1% tests. However, SARSP is rejected at 10% with subsamples between the 45 and 55 percentiles.

3.7 Proofs

Proof of Proposition 1. Consider the case $\{1, 2\}$. Take $A \subset B^1 \cap B^2$. These sets satisfy the hypothesis part of SARSP. Thus, I have $\Pi^2(B^1 \cap B^2 \setminus A) + \Pi^1(A) \leq 1$. This represents the WARSP.

$\blacksquare$

Proof of Proposition 2. Let y^i be degenerate points of B^i, $i \in I$. Suppose $y^1 R^* y^n$ and $y^1 \neq y^n$. Then, $y^i \in B^{i-1} \cap B^i, i = 2, \cdots, n$. (i) $y^n \in B^{n-1} \cap B^n \setminus B^1$. Then, for any $A \subset B^n \cap B^1$, then $y^n \in B^{n-1} \cap B^n \setminus A$. Because y^i's are degenerate points, $\Pi^2(B^1 \cap B^2) + \cdots + \Pi^n(B^{n-1} \cap B^n \setminus A) = n-1$. These sets satisfy the hypothesis part of SARSP, therefore, it must be $\Pi^1(A) = 0$. This holds for any $A \in B^n \cap B^1$, $y^1 \notin B^n \cap B^1$. Thus, $\neg y^n R y^1$. (ii) $y^n \in B^{n-1} \cap B^n \cap B^1$. Then, $y^1 R y^n$. Since SARSP implies WARSP by Proposition 1 and WARSP implies WARP, we get $\neg y^n R y^1$. Thus, SARP holds.

$\blacksquare$

Proof of Theorem 3. Let $M = \{1, \cdots, n\} \in \mathcal{M}^*(I)$ and let $A \subset B^n \cap B^1$. Assume that U satisfies $d^{i_k} \subset B^{k-1} \cap B^k, k = 2, \cdots, n$. This implies $U \in \bigcap_{k=2,\cdots,n} \{U \in \mathcal{U} : U^*(B^{k-1}) \geq U^*(B^k)\} \subset \{U \in \mathcal{U} : U^*(B^1) \geq U^*(B^n)\}$. Moreover, because $d^n(U) \in B^{n-1} \cap B^n \setminus A$, and $A \in B^n$, $U^*(B^n) > U^*(A)$. Now, if $d^1(U) \in A$, then $U^*(A) \geq U^*(B^1)$, which implies $U^*(B^n) > U^*(A) \geq U^*(B^1)$. This contradicts the transitivity of U. Thus, $\max_{U \in \mathcal{U}} [\mathbf{1}(d^2(U) \subset B^1 \cap B^2) + \cdots + \mathbf{1}(d^n(U) \subset B^{n-1} \cap B^n \setminus A)) + \mathbf{1}(d^1(U) \subset A)] \leq n-1$. This also holds straightforwardly in the case where there are some $k = 2, \cdots, n-1$ such that $d^k(U) \notin B^{k-1} \cap B^k$ or $d^n(U) \in B^{n-1} \cap B^n \setminus A$. Thus, SARSP holds by this and by $\mathcal{U}_0$-ARSP.

$\blacksquare$

Before presenting the proof of Theorem 4, I equip $\mathcal{B}$, the Borel σ-filed of $\mathbb{R}^J_+$, with a seminorm $d_\rho(C, C') = \rho(C \triangle C')$, where ρ denotes a finite non-negative measure on $\mathcal{B}$ such that ρ is absolutely continuous with respect to some measure μ on $\mathcal{B}$ and $\rho(B) \geq \max_{i \in I}\{\Pi^i(B)\}$ holds for any $B \in \mathcal{B}$.

Note that such ρ always exists by the definition of Π^i. In the following, $\mathbb{P}^*$ and $\mathbb{P}_*$ refer to outer and inner measures.

Lemma 1. *Suppose Assumptions 1 through 3 hold. Then, for every ordered pair* $(i,j) \in I \times I$*, a sequence* $\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}$ *exists which represents the maximisers as defined in Assumption 3 such that* $d_\rho(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}) \to 0$ *as* $N \to \infty$ *a.s.*

Proof of Lemma 1. First, show that $|T^{i,j}_W(\widehat{C}^{i,j}) - T^{i,j}_W| \to 0$ a.s. hold for any sequence of maximisers $\widehat{C}^{i,j}$. By Assumption 2, $\mathbb{V}_{\max}$ is non-empty. Fix some $C^* \in \mathbb{V}^{i,j}_{\max}$. Then, because I have

$$T^{i,j}_{\max}(C) = \widehat{T}^{i,j}_{\max}(C) - [\widehat{\Pi}^j(C) - \Pi^j(C)] + [\widehat{\Pi}^i(C) - \Pi^i(C)] + [\widehat{\Pi}^i(B^i \setminus B^j) - \Pi^i(B^i \setminus B^j)], \tag{3.27}$$

we get

$$\begin{aligned}
0 &\leq T^{i,j}_W - T^{i,j}_W(\widehat{C}^{i,j}) \\
&= T^{i,j}_W(C^*) - T^{i,j}_W(\widehat{C}^{i,j}) \\
&= \widehat{T}^{i,j}_W(C^*) - \widehat{T}^{i,j}_W(\widehat{C}^{i,j}) - [\widehat{\Pi}^j(C^*) - \Pi^j(C^*)] + [\widehat{\Pi}^i(C^*) - \Pi^i(C^*)] + [\widehat{\Pi}^i(B^i \setminus B^j) - \Pi^i(B^i \setminus B^j)] \\
&\leq -[\widehat{\Pi}^j(C^*) - \Pi^j(C^*)] + [\widehat{\Pi}^i(C^*) - \Pi^i(C^*)] + [\widehat{\Pi}^i(B^i \setminus B^j) - \Pi^i(B^i \setminus B^j)] \to 0, \text{ a.s.}
\end{aligned} \tag{3.28}$$

as $N \to \infty$ by Assumption 1 and the Gilivenko-Canteli theorem. Thus $|T^{i,j}_W(\widehat{C}^{i,j}) - T^{i,j}_W| \to 0$ a.s. The function $T^{i,j}_W(\cdot)$ is continuous on $\mathbb{V}^{i,j}_{\max}$ with respect to the semimetric d_ρ, because for $C, C' \in \mathbb{V}^{i,j}$,

$$\begin{aligned}
|T^{i,j}_W(C) - T^{i,j}_W(C')| &\leq |\Pi^j(C) - \Pi^j(C')| + |\Pi^i(C) - \Pi^i(C')| \\
&\leq \Pi^j(C \triangle C') + \Pi^i(C \triangle C')| \\
&\leq 2\rho(C \triangle C')| \\
&= 2d_\rho(C, C').
\end{aligned} \tag{3.29}$$

Now suppose to the contrary, for any sequences $\widehat{C}^{i,j} \in \mathbb{V}^{i,j}$ and $\widehat{C}^{i,j}_{\max} \in \mathbb{V}^{i,j}_{\max}$, $d_\rho(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}) \not\to 0$. Then, for any sequences $(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max})$, ϵ and ξ exist such that $\mathbb{P}(d_\rho(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}) > \epsilon \text{ i.o. }) > \xi$. Now, fix a sequence of $\widehat{C}^{i,j}$ and let $\{\widehat{C}^{i,j}_{\max}(\lambda)\}_{\lambda \in \Lambda}$ be the set of possible sequences of $\widehat{C}^{i,j}_{\max}$. Then, for every $\lambda \in \Lambda$, $\epsilon(\lambda) > 0$ and $\xi(\lambda) > 0$ exist such that $\mathbb{P}(d_\rho(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}(\lambda)) > \epsilon(\lambda) \text{ i.o. }) > \xi(\lambda)$. Then $\inf_{C \in \mathbb{V}^{i,j}_{\max}} d_\rho(\widehat{C}^{i,j}, C) \to 0$ a.s. cannot hold. However, then by the continuity of $T^{i,j}_W(\cdot)$ with respect to d_ρ, $\eta > 0$ and $\xi > 0$ exist such that $\mathbb{P}(|T^{i,j}_W - T^{i,j}_W(\widehat{C}^{i,j})| > \eta) > \xi$. This contradicts the previous argument. Hence, a sequence $\widehat{C}^{i,j}$ and $\widehat{C}^{i,j}_{\max}$ exist such that $d_\rho(\widehat{C}^{i,j}, \widehat{C}^{i,j}_{\max}) \to 0$ a.s.

■

Proof of Theorem 4. The proof is a modification of Proposition 3.1 and 3.2 by Kitagawa (2010). I first show that for each ordered pair $(i,j) \in I \times I$, $G_{T^{i,j}_W, N}(\cdot) \equiv \sqrt{N}(\widehat{T}^{i,j}_W(\cdot) - T^{i,j}_W(\cdot))$ converges uniformly to some tight Gaussian process $G_{T_W^{i,j}}(\cdot)$ in $l^\infty(\mathbb{V}^{i,j})$. Then, by definition,

$$\begin{aligned}
G_{T_W^{i,j},N}(C) &= (\widehat{\lambda}^j)^{-1/2}\sqrt{N^j}\widehat{\Pi}^j(C) - (\widehat{\lambda}^i)^{-1/2}\sqrt{N^i}\widehat{\Pi}^i(C) - (\widehat{\lambda}^i)^{-1/2}\sqrt{N^i}\widehat{\Pi}^i(B^i \setminus B^j) \\
&\quad - (\widehat{\lambda}^j)^{-1/2}\sqrt{N^j}\Pi^j(C) - (\lambda^i)^{-1/2}\sqrt{N^i}\Pi^i(C) - (\widehat{\lambda}^i)^{-1/2}\sqrt{N^i}\Pi^i(B^i \setminus B^j) \\
&= (\widehat{\lambda}^j)^{-1/2}G_{\Pi^j,N^j}(C) - (\widehat{\lambda}^i)^{-1/2}G_{\Pi^i,N^i}(C) - (\widehat{\lambda}^i)^{-1/2}G_{\Pi^i,N^i}(B^i \setminus B^j)
\end{aligned} \tag{3.30}$$

By Assumption 1, the asymptotic distribution of $G_{T_W^{i,j},N}(C)$ for a fixed C is a Gaussian:

$$G_{T_W^{i,j},N}(C) \rightsquigarrow (\lambda^j)^{-1/2}G_{\Pi^j}(C) - (\lambda^i)^{-1/2}G_{\Pi^i}(C) - (\lambda^i)^{-1/2}G_{\Pi^i}(B^i \setminus B^j). \tag{3.31}$$

By Theorem 1.5.4 and 1.5.7 of van der Vaart and Wellner (1996), it suffices to show that a semimetric d on $\mathbb{V}^{i,j}$ exists such that $(\mathbb{V}^{i,j}, d)$ is totally bounded and $G_{T_W^{i,j},N}(\cdot)$ is asymptotically uniformly d-equicontinuous in probability. For such a semimetric, I will use (the restriction of) d_ρ (on $\mathbb{V}^{i,j}$) as defined above.

First, show that $G_{T_W^{i,j},N}(\cdot)$ is asymptotically uniformly d_ρ-equicontinuous in probability. Under Assumption 1, $G_{\Pi^i,N^i}(\cdot)$ and $G_{\Pi^j,N^j}(\cdot)$ are asymptotically uniformly d_ρ-equicontinuous in probability, that is, for arbitrary $\epsilon > 0$ and $\eta > 0$, $\alpha^i > 0$ and $\alpha^j > 0$ exists such that for $k = i, j$,

$$\lim_{N\to\infty} \inf \mathbb{P}^*\left\{ \sup_{\Pi^k(C\triangle C')<\alpha^k} |G_{\Pi^k,N^k}(C) - G_{\Pi^k,N^k}(C')| \le \sqrt{\widehat{\lambda}^k}\frac{\epsilon}{2} \right\} > 1 - \eta \tag{3.32}$$

holds. Let $\alpha \equiv \min\{\alpha^i, \alpha^j\}$. By the definition of d_ρ and ρ, whenever $d_\rho(C, C') \le \alpha$ for $C, C' \in \mathbb{V}^{i,j}$, I have $\Pi^k(C\triangle C') \le \rho(C\triangle C') \le d_\rho(C, C') \le \alpha$ for $k = i, j$. From this, I get

$$\begin{aligned}
&\sup_{d_\rho(C,C')\le\alpha} |G_{T_W^{i,j},N}(C) - G_{T_W^{i,j},N}(C')| \\
&\le \sum_{k=i,j} (\widehat{\lambda}^k)^{-1/2} \sup_{\Pi^k(C\triangle C')} |G_{\Pi^k,N^k}(C) - G_{\Pi^k,N^k}(C')|.
\end{aligned} \tag{3.33}$$

From the above inequality, I obtain

$$\begin{aligned}
&\lim_{N\to\infty} \inf \mathbb{P}^*\left\{ \sup_{d_\rho(C,C')\le\alpha} |G_{T_W^{i,j},N}(C) - G_{T_W^{i,j},N}(C')| \le \epsilon \right\} \\
&\ge \lim_{N\to\infty} \inf \mathbb{P}^*\left\{ \sup_{\Pi^k(C\triangle C')\le\alpha^k} |G_{\Pi^k,N^k}(C) - G_{\Pi^k,N^k}(C')| \le (\widehat{\lambda}^k)^{1/2}\frac{\epsilon}{2}, k = i, j \right\} \\
&\ge \lim_{N\to\infty} \inf \prod_{k=i,j} \mathbb{P}^*\left\{ \sup_{\Pi^k(C\triangle C')\le\alpha^k} |G_{\Pi^k,N^k}(C) - G_{\Pi^k,N^k}(C')| \le (\widehat{\lambda}^k)^{1/2}\frac{\epsilon}{2} \right\} \\
&\ge \prod_{k=i,j} \lim_{N\to\infty} \inf \mathbb{P}^*\left\{ \sup_{\Pi^k(C\triangle C')\le\alpha^k} |G_{\Pi^k,N^k}(C) - G_{\Pi^k,N^k}(C')| \le (\widehat{\lambda}^k)^{1/2}\frac{\epsilon}{2} \right\} \\
&\ge (1 - \eta)^2.
\end{aligned} \tag{3.34}$$

Because η is arbitrary, $G_{T_W^{i,j},N}(\cdot)$ is asymptotically uniformly d_ρ-equicontinuous in probability.

$\mathbb{V}^{i,j}$ equipped with d_ρ is totally bounded, therefore, I can conclude that $G_{T^{i,j},N}(\cdot)$ converges uniformly to a tight Gaussian Process $G_{T_W^{i,j}}(\cdot)$. The covariance function is calculated from the marginal distribution as

$$\begin{aligned}\mathrm{Cov}(G_{T_W^{i,j}}(C), G_{T_W^{i,j}}(C')) = {} & \Pi^j(C \cap C') - \Pi^j(C)\Pi^j(C') \\ & + \Pi^i(C \cap C') - \Pi^i(C)\Pi^i(C') \\ & + \Pi^i(B^i \setminus B^j)[1 - \Pi^I(B^i \setminus B^j)].\end{aligned} \tag{3.35}$$

Next, I show that the asymptotic distribution of $\sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j})$ and $\sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j}(C))\}$ are asymptotically equivalent. Because $T_W^{i,j}(C) = T_W^{i,j}$ if $C \in \mathbb{V}_{\max}^{i,j}$, by Assumption 2, I have

$$\begin{aligned}\sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j}(C))\} &= \sup_{C \in \mathbb{V}_{\max}^{i,j}} \sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j})\} \\ &\leq \sup_{C \in \mathbb{V}^{i,j}} \sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j})\} \\ &= \sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j}).\end{aligned} \tag{3.36}$$

By Assumption 3, maximisers $\widehat{C}^{i,j}$ and $\widehat{C}_{\max}^{i,j}$ of $\widehat{T}_W^{i,j}$ on $\mathbb{V}^{i,j}$ and $\mathbb{V}_{\max}^{i,j}$ exist, I have

$$\begin{aligned}&0 \leq \sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j}) - \sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j})\} \\ &= \sqrt{N}(\widehat{T}_W^{i,j}(\widehat{C}^{i,j}) - \widehat{T}_W^{i,j}(\widehat{C}_{\max}^{i,j})) \\ &= (\widehat{\lambda}^j)^{-1/2}[G_{\Pi^j,N^j}(\widehat{C}^{i,j}) - G_{\Pi^j,N^j}(\widehat{C}_{\max}^{i,j})] - (\widehat{\lambda}^i)^{-1/2}[G_{\Pi^i,N^i}(\widehat{C}^{i,j}) - G_{\Pi^i,N^i}(\widehat{C}_{\max}^{i,j})] \\ &+ [T_W^{i,j}(\widehat{C}^{i,j}) - T_W^{i,j}(\widehat{C}_{\max}^{i,j})] \\ &\leq (\widehat{\lambda}^j)^{-1/2}[G_{\Pi^j,N^j}(\widehat{C}^{i,j}) - G_{\Pi^j,N^j}(\widehat{C}_{\max}^{i,j})] - (\widehat{\lambda}^i)^{-1/2}[G_{\Pi^i,N^i}(\widehat{C}^{i,j}) - G_{\Pi^i,N^i}(\widehat{C}_{\max}^{i,j})].\end{aligned} \tag{3.37}$$

The last inequality follows from $T_W^{i,j}(\widehat{C}_{\max}^{i,j}) = \sup_{C \in \mathcal{V}_{\max}^{i,j}} T_W^{i,j}(C) \geq T_W^{i,j}(\widehat{C}^{i,j})$ because $\widehat{C}_{\max}^{i,j} \in \mathbb{V}_{\max}^{i,j}$.

Now by Lemma 1, $\widehat{C}^{i,j}$ and $\widehat{C}_{\max}^{i,j}$ could be chosen so that $\Pi^k(\widehat{C}^{i,j} \triangle \widehat{C}_{\max}^{i,j}) \to 0$ a.s. for $k = i, j$. Then, the asymptotic stochastic equicontinuity of $G_{\Pi^k,N^k}(\cdot)$ for $k = i, j$ implies that $G_{\Pi^k,N^k}(\widehat{C}^{i,j}) - G_{\Pi^k,N^k}(\widehat{C}_{\max}^{i,j}) \to 0$ in outer probability for $k = 1, 2$, and so is $\sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j}) - \sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j})\} = o_P^*(1)$. Thus, the asymptotic distribution of $\sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j})$ is identical to $\sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\sqrt{N}(\widehat{T}_W^{i,j}(C) - T_W^{i,j})\} = \sup_{C \in \mathbb{V}_{\max}^{i,j}}\{G_{T_W^{i,j},N}(C)\}$. Because the supremum functional $\sup_{C \in \mathbb{V}_{\max}^{i,j}}\{\cdot\}$ is continuous on $l^\infty(\mathbb{V}^{i,j})$, the continuous mapping theorem and the uniform convergence of $G_{T_W^{i,j},N}(\cdot)$ yields

$$\sqrt{N}(\widehat{T}_W^{i,j} - T_W^{i,j}) \rightsquigarrow \sup_{C \in \mathbb{V}_{\max}^{i,j}}\{G_{T_W^{i,j}}(C)\}. \tag{3.38}$$

Finally, I show $\sqrt{N}(\widehat{T}_W - T_W)$ weakly converges to $\max_{(i,j)\in(I\times I)_{\max}} \sup_{C\in\mathbb{V}^{i,j}_{\max}} \{G_{T^{i,j}_W}(C)\}$. Note that

$$
\begin{aligned}
&\sqrt{N}(\widehat{T}_W - T_W) \\
&= \max_{(i,j)\in I\times I} \sup_{C\in\mathbb{V}^{i,j}} \{\sqrt{N}(\widehat{T}^{i,j}_W(C) - T^{i,j}_W(C)) + \sqrt{N}(T^{i,j}_W(C) - T_W)\} \\
&= \max_{(i,j)\in I\times I} \left\{ \sup_{C\in\mathbb{V}^{i,j}} [\sqrt{N}(\widehat{T}^{i,j}_W(C) - T^{i,j}_W(C)) + \sqrt{N}(T^{i,j}_W(C) - T^{i,j}_W)] + \sqrt{N}(T^{i,j}_W - T_W) \right\} \\
&= \max_{(i,j)\in I\times I} \left\{ \sup_{C\in\mathbb{V}^{i,j}_{\max}} \sqrt{N}(\widehat{T}^{i,j}_W(C) - T^{i,j}_W(C)) + o_p(1) + \sqrt{N}(T^{i,j}_W - T_W) \right\},
\end{aligned} \tag{3.39}
$$

and $\sup_{C\in\mathbb{V}^{i,j}_{\max}} \sqrt{N}(\widehat{T}^{i,j}_W(C) - T^{i,j}_W(C))$ is a.s. bounded but $\sqrt{N}(T^{i,j}_W - T_W) \to -\infty$ if $(i,j) \notin (I\times I)_{\max}$ and $T^{i,j}_W - t_W = 0$ if $(i,j) \in (I\times I)_{\max}$. Therefore,

$$
\sqrt{N}(\widehat{T}_W - T_W) = \max_{(i,j)\in(I\times I)_{\max}} \sup_{C\in\mathbb{V}^{i,j}_{\max}} \{G_{T^{i,j}_W}(C)\} + o_p(1). \tag{3.40}
$$

■

Bibliography

Ackerberg, D. A. (2003): "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination," *International Economic Review*, 44, 1007–1040.

Andrews, D. (2000): "Inconsistency of the Bootstrap when a Parameter is on the Boundary of the Parameter Space," *Econometrica*, 68, 399–405.

Angeletos, G. and J. La'O (2009): "Noisy Business Cycles," *manuscript*.

Anupindi, R., M. Dada, and S. Gupta (1998): "Estimation of Consumer Demand with Stock-out based Substitution: An Application to Vending Machine Products," *Marketing Science*, 17, 406–423.

Argo, J. J., D. W. Dahl, and R. V. Manchanda (2005): "The Influence of a Mere Social Presence in a Retail Context," *Journal of Consumer Research*, 32, 207–212.

Arrow, K. J. and R. Radner (1979): "Allocation of Resources in Large Teams," *Econometrica*, 47, 361–385.

Bandyopadhyay, T., I. Dasgupta, and P. Pattanaik (1999): "Stochastic Revealed Preference and the Theory of Demand," *Journal of Economic Theory*, 84, 95–110.

Bayus, B. L. and W. P. Putsis, Jr. (1999): "Product Proliferation: An Empirical Analysis of Product Line Determinants and Market Outcomes," *Marketing Science*, 18, 137–153.

Beggs, A. W. (2001): "Queues and Hierarchies," *The Review of Economic Studies*, 68, 297–322.

Benjamini, Y. and Y. Hochberg (1995): "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of The Royal Statistical Society Series B-Methodological*, 57, 289–300.

Blake, T., C. Nosko, and S. Tadelis (2015): "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment," *Econometrica*, 83, 155–174.

Blundell, R., M. Browning, and I. Crawford (2003): "Nonparametric Engel Curves and Revealed Preference," *Econometrica*, 71, 205–240.

——— (2008): "Best Nonparametric Bounds on Demand Responses," *Econometrica*, 76, 1227–1262.

Bodapati, A. V. (2008): "Recommendation Systems with Purchase Data," *Journal of Marketing Research*, 45, 77–93.

Bolton, P. and M. Dewatripont (1994): "The Firm as a Communication Network," *The Quarterly Journal of Economics*, 109, 809–839.

Caro, F. and J. Gallien (2007): "Dynamic Assortment with Demand Learning for Seasonal Consumer Goods," *Management Science*, 53, 276–292.

Conlon, C. T. and J. H. Mortimer (2008): "Demand estimation under incomplete product availability," .

Covert, T. (2013): "Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale," *manuscript*.

Crawford, G. S. and M. Shum (2005): "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73, 1137–1173.

De, P., Y. Hu, and M. S. Rahman (2010): "Technology usage and online sales: An empirical study," *Management Science*, 56, 1930–1945.

Deneckere, R., H. P. Marvel, and J. Peck (1996): "Demand Uncertainty, Inventories, and Resale Price Maintenance," *The Quarterly Journal of Economics*, 111, 885–913.

——— (1997): "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition," *American Economic Review*, 87, 619–641.

Dhar, R. and S. M. Nowlis (1999): "The Effect of Time Pressure on Consumer Choice Deferral," *Journal of Consumer Research*, 25, 369–384.

Dickstein, M. J. (2014): "Efficient provision of experience goods: Evidence from antidepressant choice," *manuscript*.

Easterby-Smith, M. and M. A. Lyles (2011): *Handbook of Organizational Learning and Knowledge Management*, John Wiley & Sons.

Erdem, T. and M. P. Keane (1996): "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1–20.

FERSHTMAN, C. AND A. PAKES (2012): "Dynamic Games with Asymmetric Information: A Framework for Empirical Work," *The Quarterly Journal of Economics*, 127, 1611–1661.

FRIEDMAN, E. J. AND S. S. OREN (1995): "The Complexity of Resource Allocation and Price Mechanisms under Bounded Rationality," *Economic Theory*, 6, 225–250.

GARICANO, L. AND T. VAN ZANDT (2013): "Hierarchies and the Division of Labor," in *The Handbook of Organizational Economics*, ed. by R. Gibbons and J. Roberts, Princeton University Press, 604–654.

GEANAKOPLOS, J. AND P. MILGROM (1988): "A Theory of Hierarchies Based on Limited Managerial Attention," *Journal of the Japanese and International Economies*, 5, 205–225.

GITTINS, J. C. (1979): "Bandit Processes and Dynamic Allocation Indices," *Journal of The Royal Statistical Society Series B-Methodological*, 41, 148–177.

GOLDFARB, A. AND C. TUCKER (2011): "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30, 389–404.

GRIFFIT, W. AND R. VEITCH (1971): "Hot and Crowded: Influence of Population Density and Temperature on Interpersonal Affective Behavior." *Journal of Personality and Social Psychology*, 17, 92.

GROSSMAN, S. J., R. E. KIHLSTROM, AND L. J. MIRMAN (1977): "A Bayesian Approach to the Production of Information and Learning by Doing," *The Review of Economic Studies*, 44, 533–547.

GROVES, T. (1983): "The Usefulness of Demand Forecasts for Team Resource Allocation in a Dynamic Environment," *The Review of Economic Studies*, 50, 555–571.

GROVES, T. AND R. RADNER (1972): "Allocation of Resources in a Team," *Journal of Economic Theory*, 4, 415–441.

HARRINGTON, JR, J. E. (1995): "Experimentation and Learning in a Differentiated-Products Duopoly," *Journal of Economic Theory*, 66, 275–288.

HONHON, D., V. GAUR, AND S. SESHADRI (2010): "Assortment Planning and Inventory Decisions under Stockout-based Substitution," *Operations Research*, 58, 1364–1379.

HOUTHAKKER, H. (1950): "Revealed Preference and the Utility Function," *Economica*, 17, 159–174.

HUANG, J.-H. AND Y.-F. CHEN (2006): "Herding in Online Product Choice," *Psychology and Marketing*, 23, 413–428.

JOVANOVICH, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972–990.

KELLER, G. AND S. RADY (1999): "Optimal Experimentation in a Changing Environment," *The Review of Economic Studies*, 66, 475–507.

KIM, J. B., P. ALBUQUERQUE, AND B. J. BRONNENBERG (2010): "Online Demand Under Limited Consumer Search," *Marketing Science*, 29, 1001–1023.

KITAGAWA, T. (2010): "Testing for Instrument Independence in the Selection Model," *manuscript*.

KITAMURA, Y. AND J. STOYE (2013): "Nonparametric Analysis of Random Utility Models: Testing," *manuscript*.

KULESHOV, V. AND D. PRECUP (2014): "Algorithms for Multi-Armed Bandit Problems," *manuscript*.

LANGER, E. J. AND S. SAEGERT (1977): "Crowding and Cognitive Control," *Journal of Personality and Social Psychology*, 35, 175–182.

LEWBEL, A. (2001): "Demand Systems with and without Errors," *The American Economic Review*, 91, 611–618.

LEWIS, R., J. M. RAO, AND D. H. REILEY (2014): "Measuring The Effects of Advertising: The Digital Frontier," .

MALONE, T. W. AND S. A. SMITH (1988): "Modeling the Performance of Organizational Structures," *Operations Research*, 36, 421–436.

MANKIW, N. G. AND R. REIS (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

MANSKI, C. F. (2014): "Identification of Income-Leisure Preferences and Evaluation of Income Tax Policy," *Quantitative Economics*, 5, 145–174.

MCFADDEN, D. (2005): "Revealed Stochastic Preference: A Synthesis," *Econometric Theory*, 26, 245–264.

MEAGHER, K. AND T. VAN ZANDT (1998): "Managerial Costs for One-Shot Decentralized Information Processing," *Review of Economic Design*, 3, 329–345.

MILLER, R. A. (1984): "Job Matching and Occupational Choice," *Journal of Political Economy*, 92, 1086–1120.

MIRMAN, L. J., L. SAMUELSON, AND A. URBANO (1993): "Monopoly Experimentation," *International Economic Review*, 34, 549–563.

MOUNT, K. R. AND S. REITER (1990): "A Model of Computing with Human Agents," manuscript.

——— (1996): "A Lower Bound on Computational Complexity Given by Revelation Mechanisms," *Economic Theory*, 7, 237–266.

Nakamura, E. and J. Steinsson (2012): "Lost in Transit: Product Replacement Bias and Pricing to Market," *American Economic Review*, 102, 3277–3316.

Nowlis, S. M. (1995): "The Effect of Time Pressure on the Choice between Brands that Differ in Quality, Price, and Product Features," *Marketing Letters*, 6, 287–295.

Orbay, H. (2002): "Information Processing Hierarchies," *Journal of Economic Theory*, 105, 370–407.

Ortega, P. A. (2011): "A Unified Framework for Resource-Bounded Autonomous Agents Interacting with Unknown Environments," *PhD Thesis*.

Papadimitriou, C. H. and J. N. Tsitsiklis (1999): "The Complexity of Optimal Queuing Network Control," *Mathematics of Operations Research*, 24, 293–305.

Prescott, E. C. (1972): "The Multi-Period Control Problem under Uncertainty," *Econometrica*, 40, 1043–1058.

Radner, R. (1993): "The Organization of Decentralized Information Processing," *Econometrica*, 61, 1109–1146.

Reiter, S. (1996): "Coordination and the Structure of Firms," manuscript.

Richter, M. (1966): "Revealed Preference Theory," *Econometrica*, 34, 635–645.

Rusmevichientong, P., Z.-J. M. Shen, and D. B. Shmoys (2010): "Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint," *Operations Research*, 58, 1666–1680.

Sahni, N. (2015): "Effect of Temporal Spacing between Advertising Exposures: Evidence from Online Field Experiments," .

Samuelson, P. A. (1938): "A Note on the Pure Theory of Consumer's Behaviour," *Economica*, 5, 61–71.

Senecal, S. and J. Nantel (2004): "The Influence of Online Product Recommendations on Consumers' Online Choices," *Journal of Retailing*, 80, 159–169.

Shanks, D. R., R. J. Tunney, and J. D. McCarthy (2002): "A Reexamination of Probability Matching and Rational Choice," *Journal of Behavioral Decision Making*, 15, 233–250.

Sims, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

Suri, R. and K. B. Monroe (2003): "The Effects of Time Constraints on Consumers' Judgments of Prices and Products," *Journal of Consumer Research*, 30, 92–104.

Talebian, M., N. Boland, and M. Savelsbergh (2013): "Pricing to Accelerate Demand Learning in Dynamic Assortment Planning," *manuscript*.

Trefler, D. (1993): "The Ignorant Monopolist: Optimal Learning with Endogenous Information," *International Economic Review*, 34, 565–581.

Tucker, C. and J. Zhang (2011): "How Does Popularity Information Affect Choices? A Field Experiment," *Management Science*, 57, 828–842.

van der Vaart, A. and J. A. Wellner (1996): *Weak Convergence and Empirical Processes: With Applications to Statistics*, Springer Series in Statistics, Springer.

Van Zandt, T. (1998): "The Scheduling and Organization of Periodic Associative Computation: Efficient Networks," *Review of Economic Design*, 3, 93–127.

Varian, H. (1982): "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50, 945–973.

——— (1985): "Non-parametric Analysis of Optimizing Behavior with Measurement Error," *Journal of Econometrics*, 30, 445–458.

Vermorel, J. and M. Mohri (2005): "Multi-Armed Bandit Algorithms and Empirical Evaluation," in *Machine Learning: ECML 2005*, Springer, 437–448.

Vihola, M. (2011): "Robust adaptive Metropolis Algorithm with Coerced Acceptance Rate," *Statistics and computing*, 22, 997–1008.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008): "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, 76, 1375–1411.

Whittle, P. (1988): "Restless Bandits: Activity Allocation in a Changing World," *Journal of Applied Probability*, 25, 287–298.

Woodford, M. (2009): "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics*, 56, 100–124.