ESSAYS ON PRICE AND QUALITY TRADEOFFS

A Dissertation
Presented to
The Academic Faculty

By

Zoey Hu

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Scheller College of Business

Georgia Institute of Technology

December 2022

ESSAYS ON PRICE AND QUALITY TRADEOFFS

Thesis committee:

Dr. Adithya Pattabhiramaiah
Scheller College of Business
Georgia Institute of Technology

Dr. Douglas Bowman
Goizueta Business School
Emory University

Dr. Xiaojing Dong
Leavey School of Business
Santa Clara University

Dr. Samuel Bond
Scheller College of Business
Georgia Institute of Technology

Dr. Mike Palazzolo
Graduate School of Management
University of California, Davis

Date approved: August 22, 2022

Tradeoffs have been with us ever since the late unpleasantness in the Garden of Eden.

Thomas Sowell

To my parents.

ACKNOWLEDGMENTS

“Try to ascend the mountain’s crest / It’d dwarf all peaks under our feet”[1] – my path towards a Ph.D. is a memorable journey like this, with the crest still far away but new realms tapped, hard questions answered, old myths demystified, and fresh doubts again sowed for future endeavors. All along the way, I have received numerous help, guidance, support and love from my committee members, the faculty at Scheller, my family and friends. Without them, I would have never come this far.

I would like to express my deepest gratitude to my advisor, Adithya Pattabhiramaiah. Adi is always a good “scholar” example to learn from: logical in thinking, elegant in writing, diligent in working. From now and then, I find it beneficial to put on an “Adi Hat”. I thank Adi for his flexibility in allowing me to explore wide interests in quantitative modeling, but always providing me a safety net when I stumble, involving me in various projects that suit me, and introducing me to his reliable co-authors. There is no way to complete this list, but I am forever grateful for his humongous input to guide me in every stage of the journey.

I am deeply indebted to my other advisor, Xiaojing Dong. I sincerely thank Xiaojing for offering her hand when I lost, especially during my first year. Her well-timed encouragement and unshakable confidence in me have guided me through dark periods and will continue to light up my future career. I thank her for being extremely generous in providing all kinds of data opportunities, introducing me to meet faculty members at other universities, and finding me stages to practice talks on different occasions. Xiaojing is not only a mentor but also a dear friend, with whom I could talk about anything. I feel lucky to have her in my life.

[1] “Gazing on Mountain Tai” by Du Fu

My heartfelt thanks to Samuel Bond for offering a different perspective on my research projects, and giving me valuable and detailed feedback after each of my presentations. Besides the thesis, I have profited from Sam's professional advice and consideration on every aspect of my life at Georgia Tech, from selecting courses, developing research projects, preparing for teaching, and making career choices. I sincerely thank Sam for taking care of me and other Ph.D. students with the utmost patience, no matter how busy he is as a coordinator.

I would like to thank my other committee members, Mike Palazzolo and Douglas Bowman, for their advice and help throughout my approach towards a doctoral degree and beyond. As an advisor/co-author for one of my thesis essays, Mike has patiently helped me with my puzzles and questions, and guided me to learn more in the process. I am grateful to him for offering learning materials and suggesting workshops to fortify my expertise. I thank Doug for his insightful modeling seminar, which gave me extensive exposure to not only modeling approaches but also the research community of marketing. Doug has also provided me with invaluable guidance since my first year. I sincerely thank him for reminding me of the big picture from time to time.

I appreciate all the support from the Scheller faculty. I thank Ajay Kohli for the training on theoretical construction. Ajay is a philosopher who inspires me to dive deep into the very foundation of a concept, a research question, and the common sense. I thank Koushyar Rajavi and Michael Lowe for involving me in various cooperation opportunities, being both patient advisors and reliable collaborators. I thank Lizhen Xu for many inspiring conversations on modeling details. I would like to thank Michael Lowe, Michael Buchanan, Donna Kantak, and Deborah Turner - for their unreserved help and encouragement in my preparation for teaching. I also want to thank our supporting staff of the Ph.D. program and the IT office for providing me with all the necessary facilities and conveniences.

For the last year and a half, I worked remotely from Ann Arbor. I am incredibly grateful to Puneet Manchanda from Ross Business School at the University of Michigan, who has been an irreplaceable mentor to me. I have benefited tremendously from conversations with Puneet, which helped me gradually form a taste for research. I thank him for all the food for thoughts, tips for surviving the snow, and for being my first friend and the bridge to many other friends in a new environment.

Special thanks are due to the marketing family - Yuly Hong, Kimberly Hyun, Merve Uzunogullari, Cheng He, Deborah Abrams, Dionne Nickerson, Iman Paul. Because of you, my Ph.D. life has not been lonely at all. I thank Yuly for being the sweetest "twin sister" I have ever wanted, going through every milestone together with me, holding my hand before dawn, exploring Atlanta side-by-side. I thank Kimberly for being a wonderful friend to share anecdotes and play tennis, introducing me to many of her friends. I thank Cheng for arranging hotpot dinners at his house, connecting us as a community. I thank Merve for being a quasi-quant companion for me, and for sharing many happy moments about her dog Rico. Thank you for making my life at Georgia Tech full of joy. Thanks should also go to my friends outside of the marketing group. I thank Da Young Kim for attending many courses with me, sharing hotels at conferences, and exploring new cities and museums together. I thank Congshan Li for being a good hiking pal. I thank Liuyi Meng, Xinlei Duan, Ruiqi Jiang, and Na Wang for being great roommates and sharing many delightful moments with me.

Words cannot describe my gratitude to my family. I could not have undertaken this journey without the unconditional support from my parents and my brother. They always stand behind my decisions and encourage me to pursue my dreams. I would like to extend my sincere thanks to Carole Wooten, who has been taking good care of me like a parent during my whole program. I thank her for inviting me to Thanksgiving gatherings, Christmas ballets, and regular brunches and dinners, making me feel

at home. I also appreciate her opening the door for me to understand the American culture and lifestyle, which started with showing me how to distinguish different types of cheese on menus. I cannot imagine my life without her in Atlanta. I especially thank Haopeng Xiao, who went to the same high school, college, and graduate school with me, also close to a family member. I thank him for driving me to groceries routinely, celebrating Chinese festivals with me, and sharing his Hulu account with me for many years. The last and also the most important family member I would like to thank is my husband Cheng Li. Dozens of cross-continental flight tickets have witnessed his support for me throughout this journey. On countless nights, he accompanied me virtually via phone calls when I was working on a homework assignment, reading a paper, tuning a model, or refining a presentation. Exposure to his research on planetary science has also extended my horizon on research designs in general, inspired me to explore many cross-disciplinary possibilities. I thank him for introducing me to the Python package “emcee” for Markov chain Monte Carlo (MCMC) in astrophysics, which I have adopted in modeling consumer search behaviors in my thesis. After so many years of being in long distance, we have finally united to start a new chapter in life.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Price and product quality are among the most important factors that influence consumers' choices. This dissertation comprises two essays that examine the price and quality trade-off among consumers in two distinct settings.

The first essay develops a Bayesian dynamic decision model aimed at capturing consumers' price-quality tradeoffs that can be mapped to their preferences. I combine real-time information shown to consumers and their search actions to develop an innovative framework for integrating the learning processes of both consumers and online platforms. Consumers often use search filters for navigating a large pool of alternatives online, routinely updating their beliefs about the availability of price and quality bundles along the way, and using such information to inform their subsequent search strategies. Online platforms are uniquely positioned to infer consumer preferences on the scale of price and quality owing to their ability to observe and control the information set shown to consumers at each stage of the consumer search journey.

I derive a closed-form solution using a utility form that customizes to the price-quality tradeoffs faced by consumers to aid a transparent interpretation of consumer search dynamics. Using both simulations and real-world data from a large travel website, I show that my model can afford online platforms with powerful prediction benefits that accrue from its enhanced capability to discover price-quality tradeoffs from search traces in real-time. The proposed approach contributes to the existing empirical search literature by incorporating the list of hotel options into a consumer search model that dynamically updates consumers' understanding of the available options. The magnitude of performance benefits suggests that my model can be a potentially promising utility to help firms with contextual targeting, with overcoming cold-start problems, and with recommendation generation tasks in a world with increasingly

strict consumer privacy regulations.

The second essay investigates the relationship between price and quality by focusing on the impact on households' nutrition outcomes following an important policy legislation passed by Congress in 2010: the Healthy, Hunger-Free Kids Act (HHFKA). The HHFKA instituted key reforms to the National School Lunch Program and School Breakfast Program, updating the nutritional guidelines for foods served in all public schools in the United States. Utilizing the Nielsen Homescan panel data, I document the potential spillovers of this policy on household food purchases. By comparing purchasing activities of a matched set of households with and without kids (who I argue are respectively likely and unlikely to have benefited from the treatment), I find out that the former group sizably reduces its calories purchased from grocery stores while keeping the nutritional composition of calories purchased materially unchanged even after the HHFKA went into effect, potentially leveraging school food programs that now provided a higher quality of food.

Moreover, the overall calorie reduction is driven primarily by changes in the behaviors of a specific sub-segment of households that, prior to policy implementation, purchased less food, and food of lower nutritional quality than did the median household. I also find patterns supporting the view that this sub-segment of households began substituting their at-home food consumption with school meals presumably owing to two key types of resource constraints — time and nutritional awareness. Additionally, I find that a different sub-segment of households that purchased more food but of lower nutritional quality than the median household now exhibited deterioration of its dietary health from grocery purchases, suggesting an unintended licensing effect of this legislation on some households' food purchases. I discuss the key implications for marketing and public policy from my main findings from both essays.

Both essays reflect my empirical analysis training in the doctoral program, using distinct approaches, from Bayesian statistics and its application to a structural approach, to Econometrics and its application to causal analysis. The overarching theme is to gain insights into consumers' decision-making process by leveraging rich consumer data and building reasonable models, in order to derive guidance for either business practice or policy making.

CHAPTER 1

UNVEILING CONSUMER PREFERENCE FROM REAL-TIME SEARCH TRACES USING A BAYESIAN DYNAMIC APPROACH

1.1 Abstract

Making inferences about consumer preferences has been instrumental in targeting and personalized recommendations. Traditional methods have relied on access to historical consumer-level data. Such data become less available given the rise of regulatory privacy protections. In this study, we develop a Bayesian dynamic approach to unveil consumer preferences by leveraging their real-time search traces. This approach opens a window for firms to improve consumers' search experience and recommendations at an early search stage with minimal information burden. To achieve that, we build a Bayesian model to capture real-time information and evaluate individual level price-quality tradeoffs in a dynamic manner. Using a utility form that tailors to price and quality tradeoffs, we offer a closed-form solution, which enables a transparent understanding of how important factors in the process together shape consumer search decisions. We test the model performance in a controlled simulation and a real-world dataset from a large travel website using the Markov chain Monte Carlo method. Both simulated and real-world applications of this new approach demonstrate powerful prediction advantages, indicating that firms can effectively discover

consumer preferences leveraging the real-time search data.

1.2 Introduction

With the development of gadgets and tracking technologies, it is much easier to acquire customer information nowadays with all types of digital footprints. Online platforms keep track of order history much more conveniently than a brick-and-mortar grocery store; smart watches listen to how many steps we have walked and how intense they are; mobile navigation knows our itinerary better than our family members. In the past decades, what raised up with eCommerce was data science, mining digital fortunes from rich personal information. Purchase panel data help discover individual preferences through choice models (Rossi et al., 1996; Danaher et al., 2015); sampling order in search sequences sheds light on roles of both product features and factors contributing to search costs (Weitzman, 1979; Chen and Yao, 2017; Dong et al., 2018); rating or purchase histories, connecting both consumer attributes and product item characteristics, inspire a stream of recommender system algorithms (Goldberg et al., 1992; Resnick et al., 1994; Gomez-Uribe and Hunt, 2016; Zou et al., 2019).

However, it is also often the case when legions of data scientists make efforts to distill gold from bits and bytes, key information stands behind a heavy fog, due to multiple reasons. First of all, recent privacy regulations have created hurdles for firms collecting personal data. Both the General Data Protection Regulation (GDPR, effective in 2018, EU) and the California Consumer Privacy Act (CCPA, effective in 2020, US) enhance individuals' control over their personal data. Although both regulations apply regionally, the impact extends globally. In 2021, Apple released a new feature to its operating systems – the App Tracking Transparency, for additional protection of consumer privacy, causing nearly $10 billion lost in ad revenue for big tech firms in only two quarters (McGee, 2021). In 2022, Google announced similar plans to curtain app tracking on Android phones and phase out third-party cookie tracking in a year or two. Meanwhile, Mozilla and Meta together developed the Interoperable Private Attribution (IPA), an aggregate system that prevents links to individual user behaviors. In the meantime, consumers are becoming more privacy-aware than ever, leveraging tools like private windows in the browser, AdChoice programs and "Do Not Track" requests, etc. to escape the ubiquitous data tracking net. Another reason for the absence of consumer information is an initiative by firms to reduce customer efforts. According to the Internet Retailer Top 500 database by 2018, 66% of the top 100 retail sites provide the "guest checkout" option, and consumers are 1.2 times as likely to use this option than logging in. For firms, providing a better and trustable customer experience while obtaining more customer data is always a tradeoff. For consumers, convenience and privacy protection are both expected from the firm they interact with.

With privacy protection becoming a recent and future theme, traditional models and algorithms which rely heavily on purchase or rating histories face increasing challenges. For example, the Dutch Banking Association received a warning from the Dutch Data Protection Authority in 2019 that historical transaction information was not allowed for the purpose of direct marketing. This incident created an influential ripple effect among a wide range of marketing-related practices which involve mining personal transaction data to understand customer-level differential preferences. The new situation calls for new approaches to restore consumer preferences.

Given the challenges of storing and mining historical data, the natural question is can we instead discover consumer preferences by leveraging the rich, in-session dynamic data while the consumer is navigating the company web pages? In particular, could we leverage the recordings of a customer's choices of search filters to reveal her preference?

Information like this is under-exploited in the current practice, but is quite revealing, especially considering that consumers keep updating the choice of search filters after reviewing the search results.

Imagine that a consumer is making a purchase decision, it is quite common that she does not really have a good idea about her willingness to pay without seeing what is available on the market. She may gather information about available products and possible price ranges that are suitable to her preference through a sequence of searches. For example, when renting an apartment in a new city, a young couple may spend quite some time learning about the available options at a certain distance from the workspace, a certain number of bedrooms, and of course, the possible price. Similarly, when looking for a hotel room to book online for a family vacation, a consumer may first specify a range of price, and gauge the offerings at that price, and adjust the price range multiple times, to find an acceptable hotel at the right price level. This process generally happens under two conditions. First, the potential number of alternatives to choose from is quite large, which allows the consumer to adjust one of her criteria, say price range, multiple times in order to gain information and formulate the expectations of the possible offerings as a bundle of features, in addition to the price. Second, the cost of search is negligible, compared to the price of the product to purchase. This allows the consumer to pay little or no attention to the search cost, while focusing on learning about the feature bundles of the product and finding the alternative of the right bundle.

To model such a dynamic decision process, we need to consider both sides. On one hand, we need to have data about how consumers adjust their search criteria in a dynamic manner within a search session. On the other hand, we need to examine the records of alternatives displayed to the consumer that satisfy her current search criteria. Leveraging both sides of the data, we can study how a consumer formulates her understanding of the market level availability, and in the meanwhile how a platform hosting such a search process updates the knowledge about consumer preferences based on the sequential search criteria specified by consumers.

In this study, we develop a TRAdeoff SEarch model (TRASE) that is tailored to the dynamics from both the consumer and the platform side. The model relies only on on-site, real-time search results and refinements, which large platforms usually collect by terabytes every day yet often find it too noisy and dynamic to make use of. Our empirical context is hotel search on one of the largest online travel agency (OTA) in China. The data recorded detailed search information on both sides – search actions carried out by each consumer, mainly the search aids deployed; hotel information displayed as a result – but no information about consumers' demographics, past purchase histories, auxiliary data or browser fingerprints, conforming to modern privacy protection standards.

More specifically, we focus on understanding a particular preference structure that is of key interest in marketing: the price-quality tradeoffs. When making complex decisions, consumers usually need to deal with conflicting objectives. Low price and high quality are such a pair of conflicting objectives that are hard to maximize simultaneously. The situation often gets even more complicated when uncertainties accompany the problem: the tradeoff issue remains, and the ambiguity comes into play (Keeney et al., 1993). Our model builds on the tradeoff decision literature, assuming that (1) consumers are clear with how much they are willing to trade off the price with quality, and (2) they are uncertain about the market level price-quality offerings and update this knowledge based on the revealed product options. Figure 1.7 shows a sample of consumer hotel search sequences.

We notice that many consumers start the hotel search with very specific price ranges, but once engaged in the search process, exposed to available options, they finally make reservations that are often outside of the initial price ranges. This short-term price shift – before and after search – points to the updated evaluation of price and quality offerings on the market, given the observed hotel options when they explore different price ranges in between. This means that the search process is also an uncertainty-resolving, information-revising process.

Apart from random walking, a natural cause of adjustments in search filters is a revised projection of what price could buy a product with what level of quality on the market. Intuitively, the platform should be able to infer consumer preferences with their search query tweaks as well as information displayed to them based on the sequential decisions made by consumers. To build a quantitative model, we use a utility function to directly measure consumer tradeoffs between price and quality on the individual level. The utility form, inspired by the literature of "value map" (Besanko et al., 2009), allows us to capture the essence of the price-quality tradeoffs using only one parameter (see Equation 1.3). Combining utility optimization and Bayesian inference, we are able to derive a closed-form solution for an individual consumer's price-quality tradeoff preference. This utility form could theoretically be applied to any case where a tradeoff relationship is desired.

We test the model using simulated data first, before applying it to a real dataset from a major online travel website in China. Both the simulation and empirical results show that the prediction of search actions improves as more search queries come in. In particular, compared to the baseline model, which does not consider the dynamic adjustment of the search queries, the prediction precision is doubled using data with 4 search queries; and the improvement can reach as high as 5-7 times better after 7 queries. This performance shows that the proposed model, simple as it is, captures the basic dynamics of how consumers' price-quality tradeoffs, as well as the market level offerings, together influence their decisions on search actions. This introduces an opportunity for platforms to offer desired alternatives at an earlier stage with real-time input and output, saving consumer search cost while personalizing the search experience.

The rest of this chapter is organized as follows. Section 1.3 reviews the relevant streams of literature. Section 1.4 introduces the model framework. Section 1.5 describes the simulation study aiming at testing the validity of the model and its capability of recovering true preference parameters. Section 1.6 presents the empirical study using data from a major online website. Section 1.7 concludes the paper and highlights managerial implications.

1.3 Related Literature

Our paper contributes to several streams of the current literature. First, our study is related to the price-quality tradeoffs literature. Choosing the option with the best combination of price and quality has been a field of keen interest since half a century ago. Previous studies mainly focus on how consumers perceive the relation between price and quality and how they make choices when faced with a set of options by trading off the desired levels of price vs. quality. Consensus has been reached stating that in the value function, price and quality should be represented in a subtractive model rather than a ratio model (Hagerty, 1978; Levin and Johnson, 1984). This setting is adopted in our model framework. There is also plenty of work on how the trade-off structure improves understanding of consumer choices, especially compared to other models like proportional matching (Tversky, 1972; Simonson and Tversky, 1992; Carmon and Simonson, 1998). However, most of the early research leverages data from experiments. One major reason for the gap in empirical studies is that it is very hard to observe the dynamic tradeoff process from mere scanner data or field surveys.

The lack of data, however, has been addressed by the availability of consumer search data on platforms, in which we could observe both the product attributes and the continuous price adjustments reflected in the usage of price filters. In such a scenario, we are able to fill this gap by exploring how the price-quality tradeoff preference, as well as consumers' projection of the market level price-quality relation, together shape their search actions.

Second, we also lay bricks on the consumer search literature. Consumer search has been of great interest to researchers since the 1960s. When Stigler, as an initiator, called for attention to "search", he defined it as consumers canvassing various sellers to ascertain the most favorable price (Stigler, 1961). This definition focusing on price has been adopted by many subsequent researchers (Diamond, 1971; Rothschild, 1974; Rosenfield and Shapiro, 1981), in which the objective of search is the price possible and the information consumers collect is also price. This is a reasonable simplification in the brick-and-mortar setting, since price was among the few variables that are quantifiable and, in the meantime, objective[1]. In the digital environment, however, both the availability of product attributes and the search mechanism have changed. It is self-evident that consumers also make efforts to evaluate the quality of a product in addition to the price.

Recent development of search models favors the reward-cost structure by Weitzman (1979). Great details have been woven into the tapestry of search literature, as this structure easily accommodates various search environment settings. Existing literature has expanded the basic structure to allow for feature importance of the

[1]distance to the store may qualify as another such variable

search platform facilities, interactions between search cost and product preferences, influence of search filters on the search process, highly complex demand patterns at the market level, etc. (Kim et al., 2010; Chen and Yao, 2017; De los Santos and Koulayev, 2017; Dong et al., 2018). The Weitzman framework does, however, hold strong assumptions, one of which presumes that consumers first evaluate and rank all options, and then sample one by one the option with the highest reservation utility among all unexplored alternatives. In an online shopping scenario, there could still be reasonable fixes. For example, the Weitzman search rule may be applied within each page of results, which usually ranges from 20 to 50 options. However, as automatic loading of more options becomes a norm in the mobile shopping environment, as well as more and more PC shopping websites, the empirical feasibility of this search rule decreases. In the current environment, consumers are faced with, as Greminger (2022) puts, a "search and discovery problem" rather than a "sampling problem".

In recent years, researchers keep looking for alternative frameworks that better suit the online shopping scenario, in which consumers are faced with a really large set of selections of products to choose from. Among these attempts, Padilla et al. (2019) proposes a stochastic framework, in which queries, clicks, and purchases are modeled separately, allowing rich heterogeneity considerations, and are combined later. This approach helps circumvent the need to optimally decide the search order, and thus does not require an overview of the product alternative pool. Ursu et al. (2020) adopts a sophisticated sampling approach to allow consumers to collect extra information until the uncertainty shrinks within a threshold. The stopping rule thus depends on the uncertainty level rather than the optimality among all. Greminger (2022), on the other hand, develops a framework where consumers make both inspecting decisions and discovery decisions, hence permitting optimal search without evaluating all available options first.

Different from the above schemes, we assume that when searching for a product among numerous alternatives, the biggest challenge for consumers is not to identify the optimal one from many, but to navigate to a safe subset that they believe is satisfactory. For example, when purchasing a home, if there are thousands of potential properties, the natural behavior is to deploy a set of screening rules and then choose among the refined choice set. Further, empirically we found that when consumers search for a product with rich characteristics, such as a hotel, they turn to highly quantitative features like price to refine results and navigate through the alternative pool. This is also consistent with the trade-off decision making literature that quantitative information increases evaluability (Hsee, 1996; Nowlis and Simonson, 1997; Shaddy et al., 2021). Based on this, we develop a framework where consumers utilize search filtering tools, mainly the more evaluable "price range", to sail through endless product alternatives. The feedback that helps them adjust search directions is thus the market level availability of price and quality bundles. Such adjustment decisions are rich in data and do not depend on external input, which facilitate applications on the fly, even with strict privacy regulations.

Another stream of literature that utilizes consumer real-time search process involves recommendation algorithms. The system recommends a list of products to a consumer based on either product similarities or user similarities (Resnick and Varian, 1997; Resnick et al., 1994; Adomavicius et al., 2008; Gomez-Uribe and Hunt, 2016). Recommendation algorithms not only encourage product exploration by exposing options consumers may not actively know of and seek for, but also help reduce search cost to a certain degree (Kim et al., 2010). Traditional recommender systems, including content-based, user-based and collaborative filtering-based algorithms, usually require a long history for effective recommendations, thus are not for real-time practice (Zhang et al., 2021). Recently, the rise of Recurrent Neural Network (RNN) has started a revolution by enabling short session-based recommendation (Hidasi et al., 2015; Li et al., 2017). However, the shortcoming of this approach resides in the nature of Neural Networks that it is not very interpretive. TRASE contributes to this field by offering a preference-based, utility optimization approach that achieves both intelligibility and efficiency.

1.4 The Tradeoff Search Model (TRASE)

In this section, we develop a parsimonious model with a flexible, yet special formulation that presents the model with three unique features. First, it directly evaluates the relative weight each decision maker puts on price versus quality, which reflects the individual level tradeoff between achieving lower price and desiring higher quality. Second, it reflects the nonlinear relationship between weighing price versus quality in a utility maximization framework, as documented by the literature (see Ding et al. (2010) as an example). Finally, it allows dynamic evaluations as the feedback from previous searches are provided and additional information on the possible range of price-quality combination is gathered.

1.4.1 Utility Function, U_{ij}

We start with a parametric form of consumer i's utility function (U_{ij}) characterized by two features of product j, price p_j, quality q_j, and a parameter α:

$$U_{ij} = U_{ij}(p_j, q_j; \alpha). \tag{1.1}$$

Since our work focuses on how consumers make tradeoffs, to put price and quality on the same weighing scale, we normalize both variables based on a reference product on the market with price $\bar{p}$ and quality $\bar{q}$ (we set them to be the median of both variables in empirical testing; see Table 1.1 for a simple example), such that $\tilde{p_j} = p_j/\bar{p}$ and $\tilde{q_j} = q_j/\bar{q}$. This setting facilitates us to capture the idea that consumers are making tradeoffs – they care about price and quality being relatively higher or lower, free of the unit scale of both factors. In this case, a hotel room tagged 100 euros won't be different in utility value from being 109 dollars to a consumer. Note that including $\bar{p}$ and $\bar{q}$ in the utility function is a weaker assumption than supposing a consumer knows the distribution of a variable, since the former assumes less prior information of a consumer.

Specifically, following the economics and marketing literature, consumers have a quasi-linear preference over price and quality (Mas-Colell et al., 1995; Besanko et al., 2009; Ding et al., 2010), and the utility function is specified as,

$$\begin{aligned} \tilde{U}_{ij} &= A_i\tilde{q_j} - B_i\tilde{p_j}^{\alpha} + \eta_i \\ &= C_i(\sin\theta_i\tilde{q_j} - \cos\theta_i\tilde{p_j}^{\alpha} + \epsilon_i), \end{aligned} \tag{1.2}$$

where $C_i = \sqrt{A_i^2 + B_i^2}$, $\sin\theta_i = A_i/\sqrt{A_i^2 + B_i^2}$ and $\cos\theta_i = B_i/\sqrt{A_i^2 + B_i^2}$ with $\theta_i \in [0, \frac{\pi}{2}]$, $\epsilon_i = \eta_i/C_i$ is a random shock. Note that we use a positive sign for $\tilde{q}$ and a negative sign for $\tilde{p}$ because "quality" is a *good* and "price" is a *bad*. The parameter α measures the curvature of the indifference curve, affecting the marginal rate of substitution of quality for price. α is constrained to be greater than 1 to satisfy the convexity requirement of the utility function. In principle, α could vary across consumers but we find the choice of α largely inconsequential within a range from the real data robustness check (see Appendix B). Therefore, we let α stay common for all consumers, while assuming price and quality weigh differently (θ_i) for everyone.

In our context, since a consumer optimizes her utility given only her own price-quality tradeoff preference and the information she gathers at each step, we omit the constant scale C_i in Equation 1.2. Thus, we update the utility function to be

$$U_{ij} = \tilde{q_j} \sin\theta_i - \tilde{p_j}^{\alpha} \cos\theta_i + \epsilon_i \tag{1.3}$$

This utility form allows us to characterize a consumer's tradeoff between price and quality with one parameter θ_i, which varies across individuals in an intuitive way:

- when $\theta_i = 0, \quad U_{ij} = -\tilde{p_j}^{\alpha} + \epsilon_i$. This is an extreme case in which consumer i only cares about price; therefore, she seeks the option with the lowest price possible.
- when $\theta_i = \frac{\pi}{2}, \quad U_{ij} = \tilde{q_j} + \epsilon_i$. This is the other extreme case when consumer i only cares about quality; therefore, she seeks the option with the highest quality possible.
- in most cases, a consumer balances the conflicting objectives of price and quality; therefore, θ_i is between 0 and $\frac{\pi}{2}$. Figure 1.1 shows the utility curves when $\alpha = 2$ and $\theta = \frac{\pi}{4}$.

The transformation from the common additive utility form to the trigonometric form is an innovative and necessary setting. Later on, we will use a Markov Chain Monte Carlo (MCMC) method to estimate the distribution of θ_i and the MCMC method works best with bounded parameters. Though the common additive form of the utility is equivalent to the trigonometric form mathematically, the parameters in the additive form are unbounded. In our setting, the tradeoff preference, θ_i, is naturally bounded to the range of $(0, \pi/2)$. Bounded parameters are favoured in many economic settings (e.g. Lewbel 1998; Matsuda 2006; Deaton 2008; Lee and Rivera 2021).

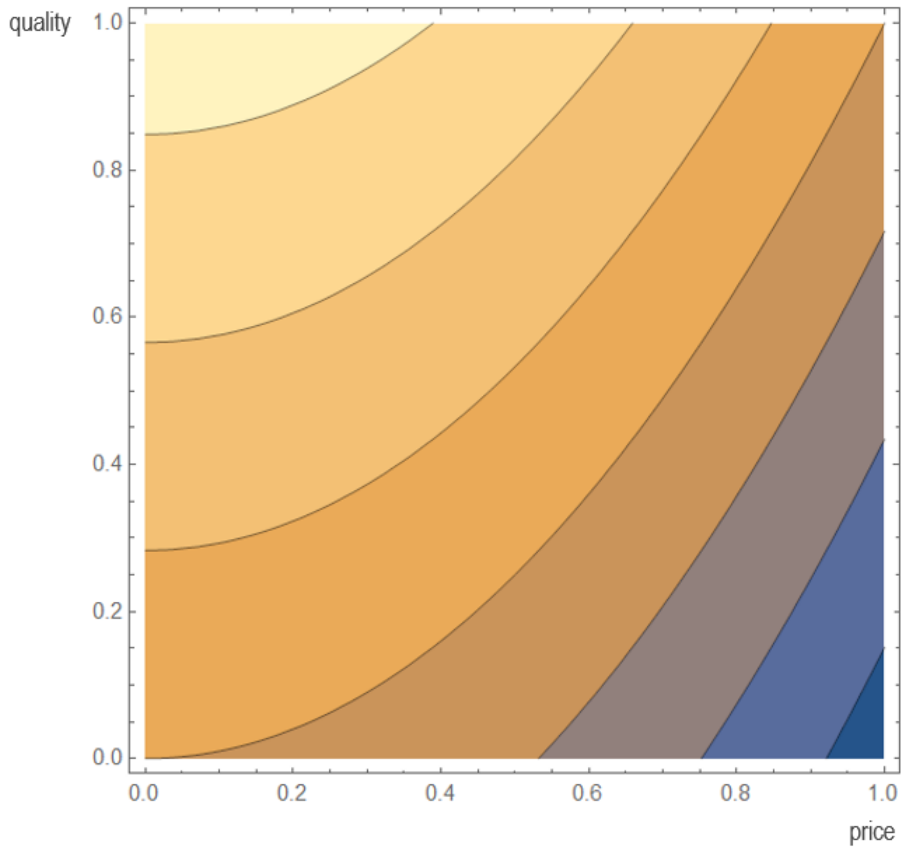

Figure 1.1: Utility Curves When $\alpha = 2$ and $\theta = \frac{\pi}{4}$

1.4.2 Tradeoff Parameter, θ_i

Using a numerical example, we describe how θ_i affects a consumer's preferences. Assume that consumers are considering the four hotels showing in Table 1.1 with a standard product characterized as $(\bar{p} = 312.5, \bar{q} = 3.4)$.

	Hotel 1	Hotel 2	Hotel 3	Hotel 4
Price, p	\$300	\$200	\$400	\$350
Quality, q	3	3	4.5	3.2
Normalized Price, $\tilde{p} = p/\bar{p}$	0.96	0.64	1.28	1.12
Normalized Quality, $\tilde{q} = q/\bar{q}$	0.88	0.88	1.31	0.93

Table 1.1: Price and Quantity Trade-off Alternatives.

In this artificial example, we created 13 consumers on the market with varying values of θ_i, representing their different trade-off preferences. Setting $\alpha = 2$ the results are listed in Table 1.2, with each row calculating the values for each consumer. The first column lists their values of θ_i, the next two columns display the calculated values of $\sin\theta_i$ and $\cos\theta_i$. Plugging in these values and the price and quality values for each hotel (in Table 1.1), we can calculate the utility for each one of the four hotels for each consumer. These utilities are listed in the last four columns. The utility values are highlighted with different shades of red color, the darker one indicates the preferred option by each customer.

The values of θ_i increase as the consumer ID increases, which shifts from price-focused to quality-focused. The consumer in the first row cares about price the most but quality the least, and the one in the last row cares about the quality the most but quality the most.

Apparently, among the four hotels we fabricated, Hotel 2 has the lowest price, and therefore attracted those in the top rows who care more about price than quality until it comes to consumer 10. As we go down the table, the importance of quality increases. Even though Hotel 2 is attractive in price, starting at consumer 11 and the rest of the two consumers care about quality enough, and Hotel 2 is not their preferred option anymore. Instead, Hotel 3, by offering much higher quality at a much higher price, become the favorite for the three consumers at the bottom of the table.

Consumer i	θ_i	$\sin\theta_i$	$\cos\theta_i$	$U_{i,Hotel1}$	$U_{i,Hotel2}$	$U_{i,Hotel3}$	$U_{i,Hotel4}$
1	0.00	0.00	1.00	-0.92	-0.41	-1.64	-1.25
2	0.01	0.01	1.00	-0.91	-0.40	-1.62	-1.24
3	0.21	0.21	0.98	-0.72	-0.21	-1.32	-1.03
4	0.30	0.30	0.95	-0.62	-0.13	-1.17	-0.92
5	0.45	0.43	0.90	-0.45	0.01	-0.91	-0.73
6	0.62	0.58	0.81	-0.24	0.18	-0.57	-0.48
7	0.85	0.75	0.66	0.04	0.38	-0.10	-0.13
8	0.91	0.79	0.61	0.13	0.44	0.03	-0.03
9	1.01	0.85	0.53	0.25	0.52	0.24	0.13
10	1.19	0.93	0.37	0.48	0.66	0.62	0.41
11	1.24	0.95	0.33	0.53	0.69	0.71	0.47
12	1.40	0.98	0.17	0.70	0.79	1.01	0.70
13	1.57	1.00	0.00	0.88	0.88	1.31	0.93

Table 1.2: Utility Changes Along with θ_i

1.4.3 Decision-Making Process

One of the biggest challenges that consumers face in online shopping is choice proliferation. If a consumer were capable of examining all available alternatives, like in the example in Section subsection 1.4.2, she would simply pick the alternative with the highest expected utility. In reality, consumers are often swamped with hundreds of options in the market, which makes it impossible to scrutinize every one of them. This is the fundamental difference between our TRASE model and more conventional models like Weitzman (1979). The latter assumes that each consumer is *able* to evaluate all alternatives in the market albeit with uncertainty on the reward, while we assume that each consumer is *unable* to examine all product options. Instead, consumers gather information at each step and approach their optimal choices with the help of a compass – filtering criteria.

While most online shopping platforms offer multiple filtering criteria, such as price, rating scores, brands, etc., we abstract product features into price-quality pairs and focus on the most frequently used and generalizable filter: price range. Theoretically, our model could be adapted to center around any other vital filters in a given context. For example, "distance" might be the crucial filter for Yelp. The tradeoff could thus be distance versus rating, rather than price versus quality. We will continue our discussion with price-quality tradeoff as our primary objective and price filter as the navigating tool.

After specifying a price range, consumers would be presented with a list of search results, among which they carefully examine the quality of product alternatives. More importantly, based on price and quality of the examined options, they form a projection about the market level price-quality relationship, which further guides them to adjust price range for the next search query. For instance, in searching for a hotel room in San Francisco, a consumer may filter results to be within \$200-\$300. She expected to find a satisfactory hotel within this price range. However, after browsing a list of hotels presented to her, she noticed that they were of lower quality than she could accept. She revised her knowledge about hotels in San Francisco and realized that a decent hotel costs more than she first anticipated, and move the price range up to \$300-\$400 for the next query.

We assume that a consumer makes decisions following this process, which involves - First, setting an initial price range and examine alternatives presented to her, – Then, making projections of market-level price and quality relationship based on the search results she saw, – Finally, revising the price range for the next search query by adjusting the price range selections, until they find an optimal alternative.

The essential part is that consumers learn about the availability of the possible bundles of price and quality on the market during the search process. We feed our model with the flow of feedback acquired with each search query and adjustments made by the consumer that reflects her updates about the available price-quality bundles.

1.4.4 Target Price at Query n, $\hat{p}_i^n$

Suppose that a consumer has examined M_n products up to query[2] n, and each j product is denoted as a quality-price pair $\{(q_j, p_j) \mid 1 \leq j \leq M_n\}$. The additional products revealed in query n are represented as $\{(q_j, p_j) \mid M_{n-1} < j \leq M_n\}$ with M_0 defined to be 0. Note that M_n is an strictly increasing function of query n.

We assume that a consumer has very limited cognitive bandwidth for the non-linearity between q and p (limited rationality). Therefore, they grasp only a linear relation

[2] We will see more clearly the meaning of a search query later. Right now, a search query is defined by its literal meaning.

between price and quality of the market offerings given $\{(q_j, p_j) \mid 1 \le j \le M_n\}$:

$$q_j = q_0^n + \gamma^n p_j, \tag{1.4}$$

where both q_0^n and γ^n are random variables estimated from the list of M_n products the consumer has seen so far. Given the linear relationship specified in (Equation 1.4), an intuitive search strategy is to extrapolate this relation to all unseen alternatives and estimate the expected utility:

$$\begin{aligned} U_{ij} &= \frac{q_j}{\bar{q}} \sin\theta_i - \left(\frac{p_j}{\bar{p}}\right)^{\alpha} \cos\theta_i + \epsilon_i \\ &= \frac{q_0^n + \gamma^n p_j}{\bar{q}} \sin\theta_i - \left(\frac{p_j}{\bar{p}}\right)^{\alpha} \cos\theta_i + \epsilon_i \end{aligned} \tag{1.5}$$

A simple optimization yields that the choice with the largest expected utility should occur at price:

$$\begin{aligned} \hat{p}_i^n &= \operatorname*{argmax}_{p_j} EU_{ij} \\ &= \operatorname*{argmax}_{p_j} \left(\frac{q_0^n + \gamma^n p_j}{\bar{q}} \sin\theta_i - \left(\frac{p_j}{\bar{p}}\right)^{\alpha} \cos\theta_i + \epsilon_i \right) \\ &= \left(\frac{\gamma^n \bar{p}^2}{\alpha \bar{q}} \tan\theta_i \right)^{\frac{1}{\alpha-1}}, \end{aligned} \tag{1.6}$$

where $\alpha = 2$ is a special case yielding $\hat{p}_i^n$ proportional to γ^n.

$$\hat{p}_i^n = \frac{\gamma^n \bar{p}^2}{2\bar{q}} \tan\theta_i, \tag{1.7}$$

The notation $\hat{p}_i^n$ in Equation 1.6 is denoted as the target price at query n. It is the price at which a consumer maximizes the expected utility given all the products she has seen up to search query n. The distribution of $\hat{p}_i^n$ depends on two parameters: 1) γ^n, the price-quality relationship among the available options, formulated up to query n; and 2) θ_i consumer i's relative preference towards quality vs. price. As the number of queries increases, consumer i gathered more information regarding the possible price-quality bundles of available products, the uncertainty in estimating γ^n, quality-price relation of the market offerings, decreases and so does the uncertainty of $\hat{p}_i^n$ according to Equation 1.6. A consumer with a larger θ_i tends to have a higher $\hat{p}_i^n$ because such a consumer values quality more than price (see the example in Section subsection 1.4.2).

We postulate that $\hat{p}_i^n$ is a latent variable that underlies a consumer's price range selection after examining the hotel list generated by the website based on the last price range selected by the consumer. Why would a consumer elect to search at a higher price for the next query? To answer that question, the model believes it is because that $\hat{p}_i^n$ is higher than the specified price range at the current query. Similarly, a consumer would elect for a lower price range in the next query if $\hat{p}_i^n$ is lower than the price range in the current query. This is the basic mechanism under which the consumers choose the price range, update their knowledge about the feasible price-quality bundles based on the list from the last query, and adjust the search range. Since $\hat{p}_i^n$ encodes the information of θ_i and the uncertainty of $\hat{p}_i^n$ shrinks with the number of queries, more queries will lead to the reduction of the uncertainty of θ_i, the preference reveals gradually.

1.4.5 Search Queries and Actions, $\mathcal{U}_i^n, \mathcal{D}_i^n$

In our context, a consumer conducts a search query by specifying a price range $\left[p_{min}, p_{max}\right]$ to refine search results. She then examines alternatives and learns about the price-quality relation of the market offerings before she initiates another search query.

Here, we formally define a "search query" and a "search action". We define a search query as a price query range $\left[p_{min}, p_{max}\right]$ selection by a consumer during a search process. A consumer sets a price range as the filter and in return gets a list of alternatives that satisfy such price range as feedback to her search query. She can then adjusts her search query to reflect her desired price range based on her updated knowledge about the price-quality bundles available on the market. In this process, we observe the sequence of the price ranges a consumer chooses over time, and the list of alternatives presented to her.

In addition, we define two actions by a consumer. We define a "search up" action for consumer i from query n to query $n+1$ as

$$\mathcal{U}_i^n = \left[p_{min}^n, p_{max}^n\right] \rightarrow \left[p_{min}^{n+1}, p_{max}^{n+1}\right] \quad \Big| \quad p_{max}^{n+1} > p_{max}^n \,,\, p_{min}^{n+1} \geq p_{min}^n; \tag{1.8}$$

and "search down" action as

$$\mathcal{D}_i^n = \left[p_{min}^n, p_{max}^n\right] \rightarrow \left[p_{min}^{n+1}, p_{max}^{n+1}\right] \quad \Big| \quad p_{min}^n < p_{min}^o \,,\, p_{max}^n \leq p_{max}^o. \tag{1.9}$$

As discussed in Section subsection 1.4.4, when a consumer searches up, she expects the target price obtained after seeing all the hotel options listed by search query n, denoted as $\hat{p}_i^n$, to be higher than the price range she specified in this query. Therefore, the probability of taking the action of searching up is the same as the probability of $\hat{p}_i^n$ to being greater than p_{max}^n, i.e.

$$\text{Prob}(\mathcal{U}_i^n) = \text{Prob}(\hat{p}_i^n > p_{max}^n) = 1 - F_{\hat{p}_i^n}(p_{max}^n), \tag{1.10}$$

where F_{p^*} is the cumulative distribution function (CDF) of $\hat{p}_i^n$. Similarly, when the consumer searches down, she expects $\hat{p}_i^n$ to locate at a price lower than p_{min}^n. Thus,

$$\text{Prob}(\mathcal{D}_i^n) = \text{Prob}(\hat{p}_i^n < p_{min}^n) = F_{\hat{p}_i^n}(p_{min}^n). \tag{1.11}$$

Note that Equation 1.10 and Equation 1.11 do not necessarily add up to one, to allow other possible search actions, such as broadening or narrowing the price ranges. These other actions are not directly modeled, as they involve mostly adjustment on the number of alternatives rather than revealing $\hat{p}_i^n$, and hence they do not contribute to the likelihood function.

Using the notations of $\mathcal{U}_i^n$ and $\mathcal{D}_i^n$, a consumer's search journey becomes a sequence of search actions, for example:

$$\mathcal{U}_i^1 \mathcal{D}_i^2 \mathcal{U}_i^3 \mathcal{U}_i^4 \mathcal{D}_i^5 \dots \tag{1.12}$$

The probability of observing the search sequence in the above example is a result of the consumer decision dynamics, influenced by her tradeoff preference θ_i, her sequential selections of price ranges $\left[p_{min}^n, p_{max}^n\right]$, and the sets of products displayed to her by the platform represented as quality-price pairs (q_j, p_j), denoted as:

$$\{(q_j, p_j), \left[p^n_{min}, p^n_{max}\right]\theta_i, 1 \leq j \leq M_n, 1 \leq n \leq T\} \longrightarrow \text{Prob}(\mathcal{U}^1_i \mathcal{D}^2_i \mathcal{U}^3_i \mathcal{U}^4_i \mathcal{D}^5_i \ldots), \tag{1.13}$$

where T refers to the total number of queries, M_n indicates the total number of product options examined up to the query number n. On the left-hand side of Equation 1.13, both (q_j, p_j) and $\left[p^n_{min}, p^n_{max}\right]$ are known to the platform and recorded in the data; θ_i is the unknown parameter for each individual that we are trying to discover.

The details of the inference process are discussed in the next section.

1.4.6 Preference Inference

To infer the model parameters and the consumer level preference, in this section, we first derive consumers' price-quality preferences using a single search query. We then incorporate the time dimension and expand it to the dynamic search process across multiple search queries.

1.4.6.1 A Single Query

Assume that a consumer is looking for a hotel room for a family vacation at an online travel agency (such as hotels.com). If She knew exactly the price range and the quality a certain price can buy, she has the option of selecting the search criteria directly, and then finding the best option for her trip as a bundle of price and quality. In most cases, when going to a new place, or even to somewhere she visited before, changes could happen due to the time differences. She would need to learn about the available hotel rooms in terms of their prices and the corresponding quality level. She starts by specifying a price range $\left[p^n_{min}, p^n_{max}\right]$ in query n. The platform responds by showing a list of hotel rooms with price-quality combinations $\{(p_j, q_j)M_n < j \leq M_{n+1}\}$. Given her preference identified by the parameter θ_i and the hotel lists she has received from all previous queries $\{(q_j, p_j)1 \leq j \leq M(n+1)\}$, she can formulate a simple relationship between price and quality of the hotel rooms that are likely to be available to her. This relationship formulation by the consumer will guide her choice of the next price range specification. For example, if the hotels from the previous feedback indicate a strong relationship between price and quality, the consumer will likely choose to increase (decrease) her price range specification in the next search query when she cares more (less) about quality than price. With a shorter list of hotels observed by the consumer, the inference is less certain. As the number of hotels increases with additional search queries, the inference becomes more and more certain. This suggests that using the change in consumer price range specifications, combined with the list of hotels from the feedback of the previous search queries, we can identify her preference for the price-quality tradeoffs. Formally, we cast this decision process into a utility-maximizing framework. To connect with the data where we observe the consumer's search up or search down choices, we assume those actions directly reflect the consumer's formulation of the optimal price with the updated information about the price-quality relationship on the market

To facilitate the mathematical formulation of the decision at each search query, we define each price range by two metrics, the lower and upper bounds. Search up ($\mathcal{U}^n_i$) is defined as when the upper bound increases and the lower bound increases or stays the same. Define the next price range as $\left[p^{n+1}_{min}, p^{n+1}_{max}\right]$, mathematically, search up is defined as $p^{n+1}_{max} > p^n_{max}$ and $p^{n+1}_{min} \geq p^n_{min}$. The probability of making this choice is:

$$\text{Prob}(\mathcal{U}^n_i \theta_i, (q_j, p_j), \left[p^n_{min}, p^n_{max}\right]) = 1 - F_{\hat{p}^n_i}(p^n_{max}), \tag{1.14}$$

where $\hat{p}_i^n$ is the target price at query n and is determined by quality-price pairs (q_j, p_j) according to Equation 1.6. Since $\hat{p}_i^n$ is related to a Gaussian random variable γ^n, by change of variables, we are able to derive a closed-form expression of $F_{\hat{p}_i^n}(p_{max}^n)$ using Gaussian distribution ϕ_{γ^n} (See details in Appendix A):

$$F_{\hat{p}_i^n}(p) = G_{\gamma^n}(p, \theta_i; \alpha, \bar{p}, \bar{q}) = \cot\theta_i \int_{-\infty}^{\frac{\alpha\bar{q}}{\bar{p}^2}p^{\alpha-1}} \phi_{\gamma^n}(x' \cot\theta_i) dx'. \tag{1.15}$$

Bayes' theorem states that the posterior probability of θ_i given the search action is proportional to the prior probability $\text{Prob}(\theta_i)$ multiplied by the probability of making the decision (likelihood), i.e.

$$\text{Prob}(\theta_i \mathcal{U}_i^n, (q_j, p_j), \left[p_{min}^n, p_{max}^n\right]) = \frac{\text{Prob}(\mathcal{U}_i^n \theta_i, (q_j, p_j), \left[p_{min}^n, p_{max}^n\right])\text{Prob}(\theta_i)}{\text{Prob}(\mathcal{U}_i^n)}, \tag{1.16}$$

Therefore, the closed-form solution for the posterior probability distribution of θ_i given $\mathcal{U}_i^n$ is

$$\text{Prob}(\theta_i \mathcal{U}_i^n) = \frac{\text{Prob}(\theta_i)\big(1 - G_{\gamma^n}(p_{max}^n, \theta_i)\big)}{1 - \int_0^{\frac{\pi}{2}} G_{\gamma^n}(p_{max}^n, \theta_i)\text{Prob}(\theta_i) d\theta_i}. \tag{1.17}$$

We have omitted hyperparameters in the function arguments to simplify the notation. Similarly, the posterior distribution for θ_i given $\mathcal{D}_i^n$ as the decision is:

$$\text{Prob}(\theta_i \mathcal{D}_i^n) = \frac{\text{Prob}(\theta_i) G_{\gamma^n}(p_{min}^n, \theta_i)}{\int_0^{\frac{\pi}{2}} G_{\gamma^n}(p_{min}^n, \theta_i)\text{Prob}(\theta_i) d\theta_i}. \tag{1.18}$$

1.4.6.2 Multiple Queries

The inference above for a single query demonstrates how the platform could update the knowledge of a consumer's price-quality preference θ_i given a search action $\mathcal{A} \in$

$\{\mathcal{U}, \mathcal{D}\}$. In this subsection, we discuss a sequence of search actions.

A search action sequence executed by consumer i is defined as an ordered collection of search actions, denoted as $(\mathcal{A}_i^n)_{n=1}^T = \mathcal{A}_i^1\mathcal{A}_i^2\mathcal{A}_i^3\ldots\mathcal{A}_i^T$. The probability of observing the search action sequence is then:

$$\text{Prob}\big((\mathcal{A}_i^n)_{n=0}^{T-1}\big) = \text{Prob}(\mathcal{A}_i^1\mathcal{A}_i^1\mathcal{A}_i^2\ldots\mathcal{A}_i^T). \tag{1.19}$$

Assuming that each search action is executed independently, the probability of observing a sequence is the product of the probabilities of individual ones:

$$\text{Prob}(\mathcal{A}_i^0\mathcal{A}_i^1\mathcal{A}_i^2\ldots\mathcal{A}_i^T) = \prod_{n=1}^{T}\text{Prob}(\mathcal{A}_i^n). \tag{1.20}$$

Therefore, for multiple queries, the posterior probability function is:

$$\text{Prob}\big(\theta_i(\mathcal{A}_i^n)_{n=1}^T\big) = \frac{\text{Prob}\big((\mathcal{A}_i^n)_{n=1}^T\theta_i\big)\text{Prob}(\theta_i)}{\text{Prob}\big((\mathcal{A}_i^n)_{n=1}^T\big)}, \tag{1.21}$$

with the likelihood function being

$$\text{Prob}\big((\mathcal{A}_i^n)_{n=1}^T\theta_i\big) = \prod_{n=1}^{T}\text{Prob}(\mathcal{A}_i^n\theta_i). \tag{1.22}$$

Terms in equation (Equation 1.22) are evaluated by either Equation 1.10 or Equation 1.11 depending on whether the search action $\mathcal{A}$ is $\mathcal{U}$ or $\mathcal{D}$.

1.4.6.3 Markov Chain Monte Carlo method

Finally, we estimate the posterior probability of θ_i given a search action sequence $(\mathcal{A}_i^n)_{n=1}^T$ (Equation 1.21) using an ensemble Markov Chain Monte Carlo (MCMC) sampler (Goodman and Weare, 2010). For each step, we draw θ_i randomly from a prescribed prior probability density $\mathrm{Prob}(\theta_i)$ and evaluate the likelihood of generating a search sequence $(\mathcal{A}_i^n)_{n=1}^T$ according to Equation 1.22. Whether a new state is accepted to the chain or discarded is according to the law of detailed balance. The final statistics are gathered by the distribution of all states, θ_i, in the chain. We usually initialize 10 parallel walkers and execute the MCMC simulation for 1000 steps. An example of the simulation is shown in the next section.

1.5 Simulation

We test the TRASE Model in simulation, in which we imitate consumer search processes with their true price-quality tradeoff preference θ_i known to us. The data generation is based on the given tradeoff parameter and our definition of search actions. Then we apply the model to infer θ_i and compare the distribution to the true value. In this way, we get to understand how well the model performs and the uncertainty associated with it. We have tried to make the simulation as close to reality as possible.

1.5.1 Search Behavior Simulation

First, a hotel directory with 240 alternatives has been generated. The price ranges from \$5 to \$1200, whereas quality stays between 0.13 to 4.84. The two variables are linearly correlated and normally distributed. We then create 500 users with their price-quality tradeoff parameters θ_i drawn from a scaled $\mathcal{B}eta$ distribution $\mathcal{B}eta(a, b; \frac{\pi}{2})$. The parameter $\pi/2$ scales the range of θ_i such that $\theta_i \in [0, \frac{\pi}{2}]$.

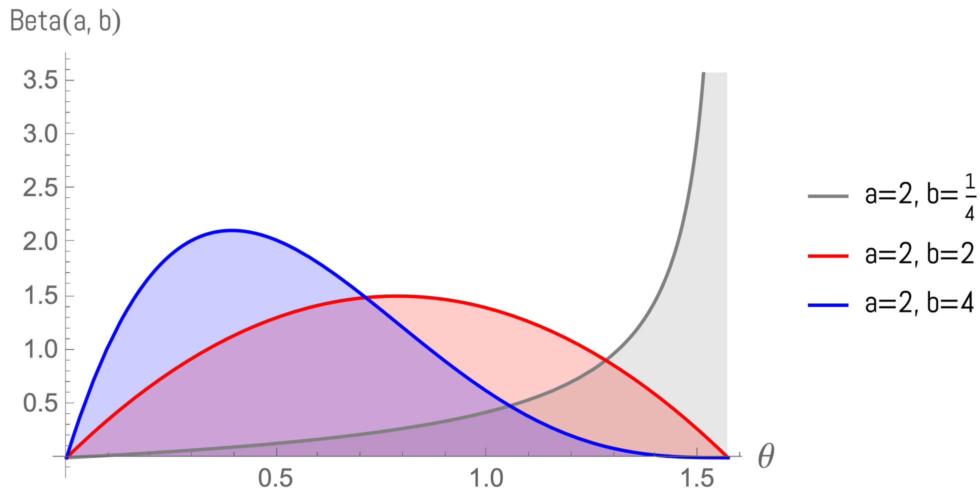

Figure 1.2: θ Follows a Scaled Beta Distribution

Figure 1.2 illustrates three probability density functions (PDF) from the $\mathcal{B}eta$ distribution, $(a = 2, b = 1/4), (a = 2, b = 2)$ and $(a = 2, b = 4)$. It shows that $\mathcal{B}eta$ distribution is a versatile distribution that can represent symmetric, asymmetric and extreme distributions of θ_i by adjusting the hyperparameters a and b.

We then simulate search queries for these users. In order to mimic real search as close as possible, we also allow search actions such as broadening and narrowing price ranges, which adds a decent amount of noise in simulated behaviors. Since our model doesn't have a stopping rule, we ask users in our simulation to search for $T = 10$

queries, whether they have found their optimal choice or not.

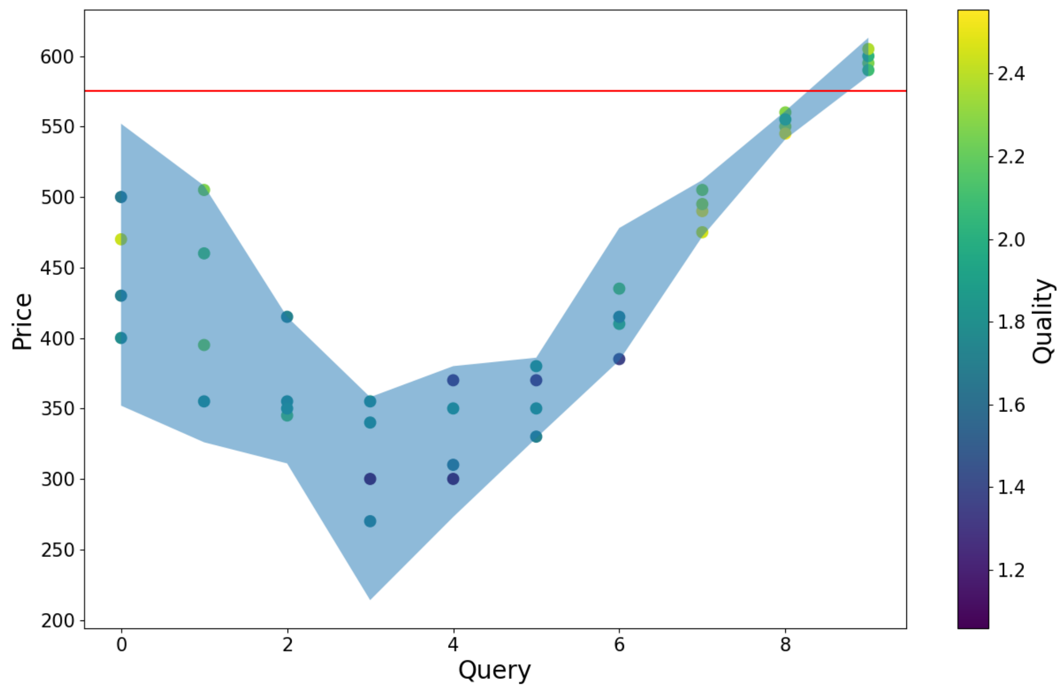

Figure 1.3: Simulated Search Process of a User

Figure 1.3 shows the search process of one user, as well as alternatives revealed to her. She starts the search by confining hotel prices between \$350-\$550, and in response receives four hotels for her to view – an appropriate number of alternatives visible at the same time on a mobile phone screen. The yellow-ish dot denotes a hotel with a price \$468 and quality above 2.2. This indicates that high-quality hotels do not necessarily come with very high prices, which we could easily find in the real-world dataset as well. This triggers some bargain-hunting queries afterward: she continues to search down for three more queries, until she has been convinced by the later revealed hotels that price and quality are, after all, more positively correlated than she first expected. This user, according to our profile, is much more concerned about quality than price, with the tradeoff parameter being 1.19 out of the range 0 to $\frac{\pi}{2}$.

The red horizontal line denotes the optimal hotel from our fabricated hotel list for this user. The search process of this user shows that although one could be biased by what the platform displays in search impressions, with more data, she is able to converge to the optimal choice, with continuous adjustment of the price search tool. In fact, most users in our simulation are smart enough to approach the optimal hotel within 10 search queries. Some, however, marches away due to the unrepresentative samples encountered in their early exploration. Overall, all these cases constitute a reasonable set of search processes.

1.5.2 Tradeoff Parameter Recovery

With search actions generated, we then feed the TRASE model with the simulated observations to recover each consumer's true preference, using MCMC simulation.

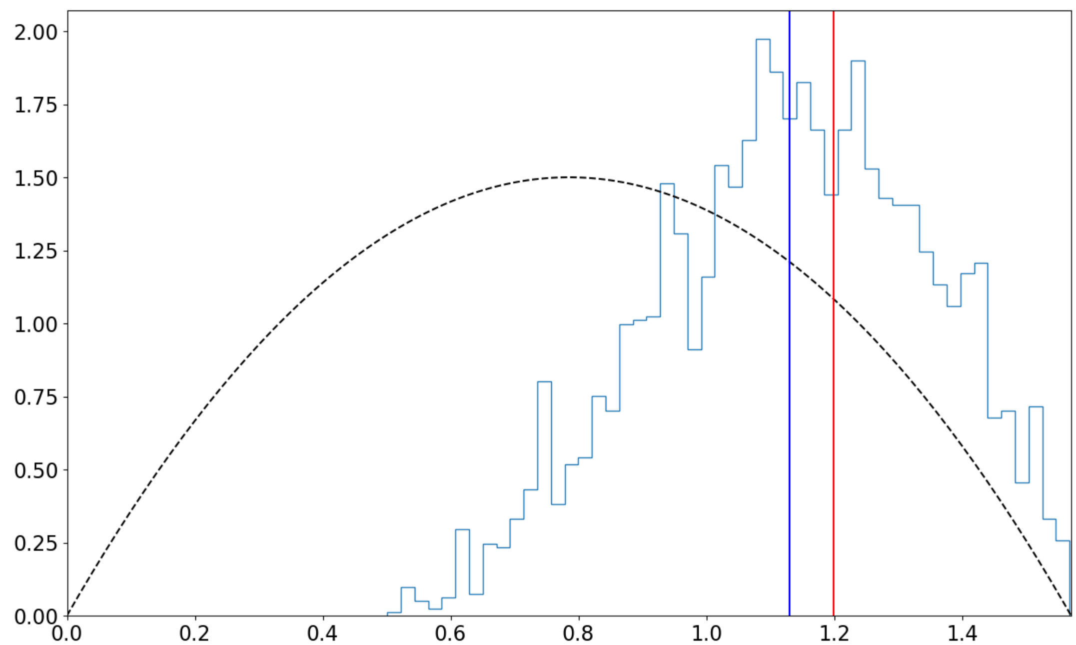

Figure 1.4: Inferred Tradeoff Parameter θ For the Same User in Figure 1.3

Figure 1.4 demonstrates the inferred posterior distribution of the tradeoff parameter of this user, given search actions of all ten queries. The red vertical line denotes the true tradeoff preference, while the blue vertical line denotes the average posterior value, which is 1.165. The dashed arc describes the shape of the prior distribution $\mathcal{B}eta(2, 2; \frac{\pi}{2})$, while the blue histogram gives the posterior distribution. With ten search actions to guide the extrapolation, the distribution in the MCMC chains gradually shifts right, drawing near to the true parameter.

Session by session inference, which does not show in this figure, displays that after inputting the first "search down" query, the model guesses the tradeoff to be 0.79; with another two "search down" queries, the estimation keeps dropping to 0.60, moving away from the actual preference. This corresponds to the first few queries steering away from the optimal option (red line). Given the many "search up" sessions afterward, the estimation direction turns around. The model reaches to 1.17 in the end, which is very close to the true tradeoff parameter 1.19. Overall, our model proves to be powerful and accurate in uncovering consumer preferences.

1.5.3 Prediction Performance

With θ_i recovered from the search sequence, a platform can now predict the search actions of the consumer with higher accuracy and thus make real-time recommendations better tailored to the consumer's needs. In this section, we use the prediction accuracy of search actions, random and elusive as they appear, as a signal of the improvement our model could bring to a better comprehension of consumer preferences.

Moreover, since we do not know the true price-quality trade-off of each consumer in the real world, it is impossible to test the model performance using real-world data by comparing the inferred distribution against the truth. However, we do observe their search actions, a series of decisions to search up or search down. With confidence in the ability of preference restoration, we take it one step further to prediction, which allows verification of the model performance in the real world as well.

We first test the prediction accuracy using simulated data, on the basis of the inferred tradeoff parameter θ_i. Later on, we apply it to predict real consumer search actions. The prediction performance of simulation data could be used as a benchmark for real-world data. Figure 1.5 shows the overall prediction performance of simulated search actions. The two curves below the red horizontal line denote the prediction precision of the baseline model and the TRASE model, query by query. The curve above the red horizontal line denotes the ratio of accuracy between our model and the baseline model.

Since our model is the first to use micro-decisions such as search up and search down actions to dynamically reveal consumer preferences, it is hard to find a classic model to compare with. Potentially, we could compare with models using the Weitzman framework. For example Chen and Yao (2017) and De los Santos and Koulayev (2017) both consider how the employment of search refinement tools could affect consumer search. However, their models mostly focus on implications of the number of searches, click-through rates, consumer surplus, the market structure, etc. While our model focuses on how to quickly and accurately infer consumer preferences through dynamic feedback and product impressions, providing platforms a swift gadget to adjust product recommendations at an early search state. We thus compare the performance of our model to one which relies on random guess of consumers' tradeoff preference, but also feeds on information about the last search action. The difference between TRASE and the baseline model is that the former is built on the step-by-step learning and updating process.

In Figure 1.5, we could see that the first query does not give much information to start the engine. However, the accuracy improves steadily with consumers contributing more search queries. Prediction precision of TRASE is doubled after three to four queries, which is a huge boost given how little information we have utilized. After four queries, the precision has risen up to above 80%. With even more decisions in tuning the price filters, the performance continues to improve, compared to the baseline model. One thing that catches our attention is that the performance of the baseline model deteriorates, rather than improves. A closer look at individual cases gives useful insight into this. Unlike usual prediction objectives, such as clicks and conversions, price search decisions (up or down) are more dynamic measurements in nature. The decision at one point heavily depends on previous paths and knowledge updates, thus creating very noisy traces. Ignoring the middle processes and relying only on the static information, which the baseline model does, would mask the true preference and intention of a consumer. This might be one of the reasons why search steps are often out of sight in many economic and machine learning models. Our model hence comes in handy when platforms hope to make use of this readily available data.

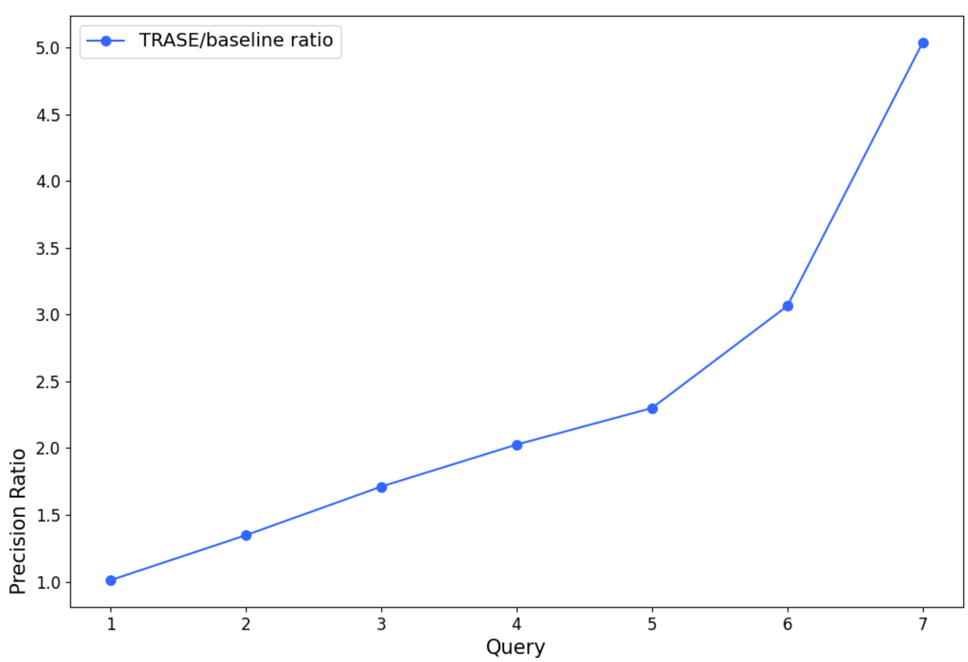

Figure 1.5: Overall Prediction Performance in Simulation

In sum, the prediction performance shows that the TRASE model is powerful in inferring consumer price-quality tradeoff preferences, and the precision increases with more search information coming in.

1.6 Empirical Application

1.6.1 Data

We get the dataset from a major online travel company in China. The company keeps records of consumer search behaviors on the mobile application. This platform provides 10+ filters, and Price Range is the most frequently used one (see Figure 1.6). Consumers in our dataset search for a hotel in Shanghai. There are in total over 8000 different hotels showing up in the impressions, which indicates a highly homogeneous and competitive market. This highlights the need for consumers to use a filter tool to narrow down search results.

Consumers learn about the price-quality structure of the market during their search process. Based on the updated belief at one point, they change their price filter range accordingly. In fact, out of the 60,413 users in the dataset who make use of the Price Range tool, 67% of them utilize and adjust Price Range 3+ times. To test the TRASE model, we randomly select 500 such users as our testing sample. Table 1.3 shows basic summary statistics of the sample.

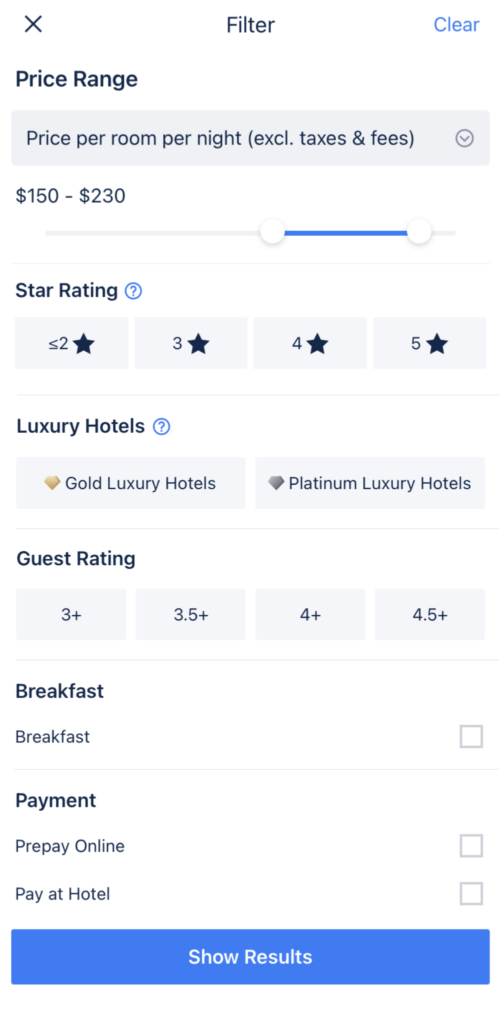

Figure 1.6: Search Filters the Firm Provides

Table 1.3: Summary Statistics of Consumer Search Behaviors in the Sample

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
number of hotels	500	129.61	119.51	13	57	152.5	970
number of clicks	500	7.01	7.73	0	3	8.5	50
number of orders	500	0.11	0.35	0	0	0	2
number of price ranges	500	21.89	29.23	3	7	45	105
number of sessions	500	3.47	2.87	1	2	4	23

In the model, consumers know their own price-quality tradeoff preference, and learn about the market price-quality structure through search queries. While testing the model using real-world data, one task is to calculate the quality of each hotel. Luckily, the platform evaluates each hotel with a "diamond level", a comprehensive score on hotel quality ranging from 0 to 5, which takes into account hotel facilities, services, brand reputation, star ratings from the National Tourism Bureau, user rating scores and rating numbers, etc. We use this "diamond level" as a proxy of hotel quality.

We further verify the reliability of this proxy by regressing 72 hotel attributes of all hotels in our dataset, including labels from the platform for each hotel[3], facilities, quotes from consumer reviews, etc., on diamond levels. Among all the significantly contributing independent variables, the top 5 variables, which show strong *positive* correlation with diamond levels, are "mansion", "designer hotel", "gym", "family-friendly", "vacation"; the top 5 variables showing strong *negative* correlation with diamond levels are "sale everyday", "farm house", "village house", "youth hostel", "hotel inn". These results exhibit, from an explanational view, that "diamond level" is a reasonable proxy for quality.

1.6.2 Model Testing

For real data, we do not know the true tradeoff parameters of consumers. Thus, we do not have a direct way to test the preference-recovering ability of the model. However, we do see how consumers search up or down after seeing a set of hotels. We instead use their search actions as our prediction targets to reflect the model performance. We could potentially compare the prediction performance of the model in the real world (Figure 1.8) to that in the simulation (Figure 1.5) as a reference with regard

[3] such as "family-friendly", "ocean view", "free upgrade", etc.

to preference unveiling capabilities.

Unlike in simulation, consumers in the real world have more flexibility in deciding how many alternatives to check before changing the price range filter. Also, they could freely decide how many queries and sessions to carry out. This adds more variation to the search action observations overall. In the simulation, it is hard to set more complicated rules in addition to what we have imposed, since it requires more subjective judgments or complementary models in the data generating mechanism. The hotel search dataset, fortunately, could help make up for this shortcoming of simulation data, and reinforce the credibility of our model performance.

Figure 1.7 displays the search process of a real user in real data. This consumer conducts in total of six queries. He first sets the price range to \$200-\$300. After only two hotel views, perhaps as a quick correction, he quickly changes it to "\$250 and below". Within this price range, he browses ten hotels, with a decent amount of variation in both price and quality. Afterwards, he tunes the range back to \$200-\$300, with the updated information, and goes up once more, and then back to the original range again. Behind the curtain, as he searches down, up and down again, the model estimate of his tradeoff parameter also changes to 0.72, 0.76, and 0.70. As we do not know his true preference, we couldn't tell if the estimation moves toward the answer. However, this example helps in giving a rough idea of how consumers search in the real world, and how in general the model estimation changes along with search actions.

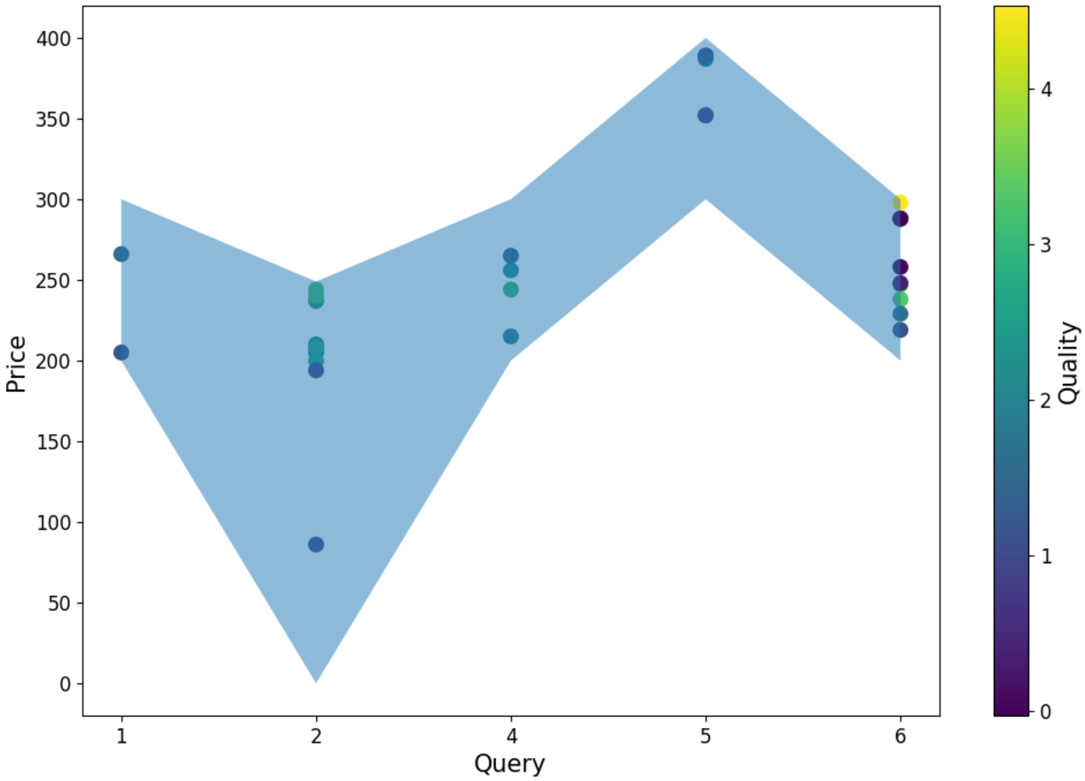

Figure 1.7: Search Process of a Real User

Figure 1.8 depicts the overall prediction performance. Similar to Figure 1.5, the two lines below give the prediction precision of real search actions – the gray line for the baseline model and the blue line for TRASE. The baseline model still performs unsatisfactorily in predicting dynamic search sequences, while the TRASE model is built for this setting and hence performs better. The overall precision ratio follows a similar pattern to that in the simulation, with gradual improvement for the first few queries and a significant boost with more actions available. The improvement lift, however, kicks in later in real data. After four queries, the precision ratio only reaches to 1.69, compared to 2.03 in simulation. But after five queries, the ratio improves to the same level as in simulation, which is 3.1 times as good as the baseline model. After seven queries, the ratio climbs to an even higher level than that in simulation.

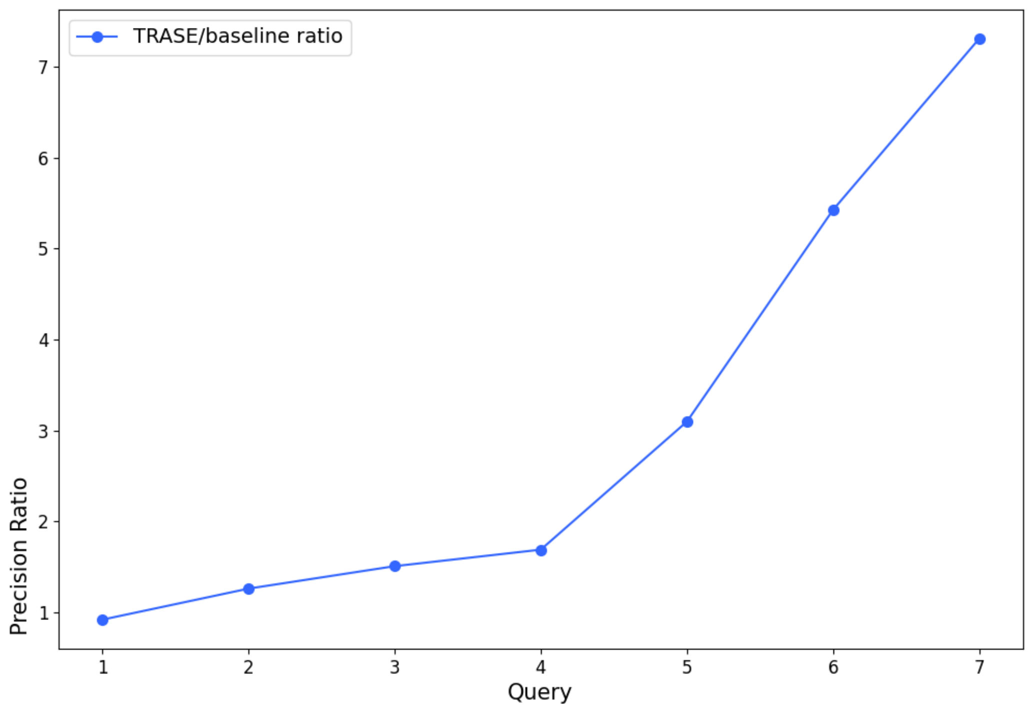

Figure 1.8: Overall Prediction Performance in Real Data

Legit questions could arise from the readers about the empirical meaning of the above model performance. If we need to accumulate three to four queries to arrive at a satisfactory performance, does the real search context support such luxury? This is a good chance for us to distinguish “search session” and “search query”. The former is usually defined as a sequence of consumer actions delimited by the inactivity of 30 minutes. The duration of a search session could last from one minute to several hours, depending on the shopping scenario, but often between 30-60 minutes for searches on travel websites. A search query, however, is defined by a set of query inputs (such as check-in/out dates, price, brand) and the action of clicking on the "search" button. Multiple search queries could take place within a single search session. In our sample, a consumer could carry out nearly 7 queries per session. In addition, according to the dataset provided to us, more than half of consumers utilize "price ranges" to screen search results. And among them, 67% adjust the price range for more than three times, which means four queries already. In this sense, our model could apply to around forty percent of such customers. Thus, a burn-in of 3-4 search queries could be feasible in reality.

Overall, the rising curve demonstrates that TRASE is powerful in prediction search actions in the real world as well. This is an indirect indicator that our model does a good job of capturing the tradeoff preferences.

1.7 Concluding Remarks

The past four years have witnessed the flourish of privacy awareness and privacy regulations. From GDPR to CCPA, from Apple, Google, Meta to Mozilla, governments and big tech firms are leading a revolution in personal data collection. Marketing models should also adapt to the challenge and evolve. Inspired by the new development, we build a model that is not addicted to consumer historical profiles or complementary data, but only sits on on-site real-time search actions to reveal consumer preferences.

Traditionally, product choices were the major source of information for marketing modelers to discover consumer preference on price, quality, and their relevant importance, in both quantitative models and behavioral experiments. Recently, as consumer search data prior to their purchases become readily available, leveraging such data to understand the consumer decision process becomes pertinent. In this paper, we develop a new modeling approach to take advantage of the additional information embedded in the search process, focusing on the dynamic nature of consumer search. Specifically, we focus on a preference structure that is of key interest in the marketing field – the tradeoff between price and quality.

We base our model on the assumption that consumers are clear about their own tradeoff preferences and need to learn about the market price-quality structure step by step. Utilizing their frequent decisions to search up or down, combined with the product alternatives revealed to them, we are able to update our knowledge about consumers' inner tradeoff inclination towards price and quality. An innovative utility form allows us to represent this tradeoff with a single parameter. Following utility optimization and Bayes' Law, we arrive at a closed-form solution for an individual consumer's price-quality tradeoff preference. The simplicity of an analytical solution enables platforms to embed our model into real-time applications on the fly. We test the model performance in both simulation and a real-world dataset from a large travel website. Both results show that our model is powerful in recovering consumer preferences. In the simulation, the model proves efficient in restoring the true tradeoff parameters; in both simulation and real data, compared to the baseline model, which does not consider the dynamic adjustment of search actions, the prediction precision is doubled given four search queries, 5-7 times better after seven queries. Moreover, evidence from the empirical data also verifies that it is feasible and affordable to have three to four burn-out search queries to reinforce model performance. This performance indicates that the TRASE model, simple as it is, captures the basic dynamics of how consumers' price-quality tradeoffs influence their decisions on search actions.

Since our model helps unveil consumer preferences using real-time search data, it kills three birds with one stone: (1) it enables managers to ride the tide of privacy awakening as it does not depend on consumer profiles and histories; (2) it helps address the cold start problem for new customers or new contexts for old customers alike; (3) it cashes in on the ubiquitous, massive and noisy search sequence data. The application could be fruitful. The ideal application scenario is real-time recommendation: with two to three search queries as the burn-in stage to accumulate information, platforms could effectively personalize search experience with optimal products pushed to higher ranks in subsequent queries. Moreover, it could be applied to general marketing purposes. For example, when a customer comes to a platform and stops searching after a few sessions due to some external reason, the early learned preference could be applied to direct marketing within an effective time window. Additionally, as preferences inferred by the model cling to the present shopping task, it could act as a reinforcement source for existing models, injecting task-and-context-specific information. For example, it is feasible to combine the preference inferred from TRASE with model-based collaborative filtering algorithms, which is versatile in taking in rich side information (Shi et al., 2014; Zhang et al., 2021).

In addition, as our model builds on search queries and product impressions only, it is especially advantageous in shopping scenarios involving infrequent purchases, context heterogeneity, or group decision-making. The common hurdle for traditional choice models is that they try to elicit consumer preferences from historical data, assuming a stable and representative disposition in front of the same shopping task. However, in the above scenarios, the task or even the person is hardly consistent. For example, a business person could search for flight tickets for a family vacation on one occasion, and for a business trip on another occasion[4]; a housewife may search for a silverware set for herself on one day, and as a wedding gift for her friend on another day; behind a string of search behaviors, there might be a single shopper or a group of

[4]A recent study has focused on this "context heterogeneity" (Padilla et al. (2019))

people as well[5]. What's more, the recorded data might not be representative. In the travel business, an average American adult takes 2-3 flights per year (Soulier (2018)). Given the competitive market, each firm might get an even smaller share of data from a customer.

These scenarios, however, highlight the importance and benefits of a real-time search model. The preferences, of whomever the individual or group, aiming for whatever tasks, stay coherent in a short search period. Moreover, consumer search tends to be more extensive in infrequent or context-specific purchases, which could feed our model with even more search queries to improve performance. According to Statista, grocery and fashion shoppers worldwide search for an average of 32 pages at a buying session alone on platforms (Chevalier (2021)). While searching for a digital camera on Amazon, BestBuy, etc., a typical household generates 13.9 sessions (Bronnenberg et al. (2016)). On a U.S. travel website, an average user searches for a hotel room 33.5 days in advance of their stay (De los Santos and Koulayev (2017)). People normally spend far more days on Zillow, Carvana, Expedia, etc. than on eBay and Amazon, until they arrive at a satisfactory choice. The extensive search gives away trains of thoughts behind consumer decisions.

We build this model that tailors to the trade-off between price and quality, including the utility form, as well as the derivation in utility optimization. However, it could expand beyond the price-quality structure and be applied in a more general setting, as we do not put strict restrictions on the characteristics of trade-off variables while developing both the utility format and the search action strategy, except that 1) the two variables are conflicting with regard to consumer choice decisions, with one variable a normal "good" in economic settings and the other a "bad" (Varian, 2014), and that 2) they are continuous variables in our derivation of optimization. For the

[5]Related work see Cao and Zhu (2019)

first assumption, it is the key to generalization to capture the overarching pair of conflicting variables in a specific context. For example, when consumers use Yelp for restaurant reviews, the centering trade-off would not be price and quality, as we do not need to pay on Yelp and the website emphasizes more on reviews, which is a critical part of quality. Instead, the pair of conflicting goals could be the rating score of a restaurant and how convenient to get to the place (e.g. distance). Often, preliminary analysis of consumer search actions in the empirical context could offer valuable clues for the most important trade-off structure. For the second assumption, there are also practicable solutions to resolve the continuity requirement for discrete variables. We could either smooth a discrete variable like the star level of ratings, or we could exhaust the possible outcomes and directly replace it with discrete probability distributions in optimization, the downside of the latter is that we may or may not have a closed-form solution. It would be interesting to see how effective this model is in other contexts, which could be a new essay topic in the future.

We offer a new approach to feeling consumers' pulse in a privacy-sensitive, data-restricted context using search data. Serving this purpose, the challenge is substantial but the research potential is also vast. We want to point out that we only focus on the price-quality tradeoff preference recovery through search traces. It still awaits investigation with regard to what other general information we could distill from the readily available search sequences or other similar data types. Steer toward another direction, it would also be interesting to see how the new tradeoff utility structure could be applied to contexts where tradeoffs are salient and central.

CHAPTER 2

DO SCHOOL NUTRITION MANDATES AFFECT HOUSEHOLD FOOD PURCHASES?

2.1 Abstract

In 2012, a provision of the Healthy, Hunger-Free Kids Act (HHFKA) placed strict requirements on what types of foods could be offered at schools, increasing the healthiness of food available to students. Did changes to the healthiness of school meals, even in the absence of more generous meal subsidies, lead parents to substitute towards school meals, and away from purchasing meals for their kids from supermarkets? This provides a unique opportunity to observe how households make healthy eating choices, when a food option improves the healthiness. We document a small but significant change in food quantity at home, but not in food quality. Moreover, much of the reduction in food purchased from the supermarket appears to come from food categories traditionally associated with breakfast or lunch–the two meals for which food can easily be purchased from schools. To investigate the heterogeneous differences, we group households into quadrants according to them being "low" and "high" on food quantity and quality prior to the HHFKA taking into effect. Surprisingly, the reduction in quantity is mainly driven, not by households more nutritiously conscious or those who eat at home more, but by households who purchase less food quantity with low quality, with calorie purchases decreasing by 19%. We find out

that households with time pressure, financial constraint and relative food literacy contribute to them being in this quadrant, making further use of the improved school meals. In addition, households who previously purchase low-quality food significantly further decreased the healthiness of food at home after the treatment, indicating an unintended licensing effect of the HHFKA on household food purchases.

2.2 Introduction

Research in marketing and a broad array of other fields has, in recent years, made efforts to help answer an increasingly important question: "How can we help consumers eat healthier?" Studies from different approaches try to gradually piece together a solution chart, including access to healthy food, implementing sugar taxes, marketing communications like advertising bans and nutrition labels, as well as various forms of behavioral nudges. For example, higher access to healthy food on varying scales – from restaurant menus, refrigerators or lunch lines in a cafeteria, to grocery stores with fresh produce available – is proven helpful in promoting healthy food intake (Downs et al., 2009; Thorndike et al., 2012; Hanks et al., 2012; Althoff et al., 2022). Related, researchers also find that refraining the physical access to food could also encourage choices towards healthier items. Since contextual factors like the sight, smell and touch of food could increase the desire for eating more (Rogers and Hill, 1989; Cornell et al., 1989; Peck and Childers, 2006), purchasing food online or pre-order it before getting hungry and exposed to sensory stimuli could successfully reduce the selection of unhealthy food and increase healthy food choice (Miller et al., 2016; Huyghe et al., 2017). Behavioral nudges are generally considered effective in reinforcing health eating, such as larger portion size and plate size for fruits and vegetables(DiSantis et al., 2013; Miller et al., 2015), healthy-eating calls or recommendations (Mayer et al., 1986; Reicks et al., 2012; Miller et al., 2016; Anzman-Frasca et al., 2018), nutritional rating systems (Levin, 1996; Nikolova and Inman, 2015; Olstad et al., 2015; Ensaff et al., 2015). In addition, government-initiated legislation and regulations also aim to create a more health-friendly food consumption environment such as enforcing nutrition labels on food categories, advertisement bans on fast food targeting children, implementing soda tax and fat tax, etc., with the effects noticed varying among different demographic groups (Moorman, 1996; Dhar and Baylis, 2011; Bollinger and Sexton, 2018; Dubois et al., 2020).

Despite all the research findings and efforts like the aforementioned, however, public dietary health is still an alarming issue today in our society, the most significant reflection of which is the climbing obesity rates among both children and adults. From 1971 to 2018, the prevalence of obesity among children and adolescents rose from 6.2% to 25.4%; and that of adults increased from 14.3% in 1960 to 52.4% in 2018[1]. Researchers also have devoted to examine why some households fail to eat healthily. An often-cited explanation, of course, is an inability to afford nutritious food (Kearney and McElhone, 1999; Nicklas et al., 2013; de Mestral et al., 2017; Palazzolo and Pattabhiramaiah, 2021). Other explanations include a lack of access to healthy foods (Jetter and Cassady, 2006; Althoff et al., 2022), a lack of "food literacy"(Wijayaratne et al., 2018), and the simple fact that unhealthy food tastes quite good(Drewnowski, 1997; Epstein et al., 2006).

One challenge researchers face in learning about consumer eating choices is that we rarely observe large, exogenous shocks to the options available to consumers. The Healthy, Hunger-Free Kids Act (HHFKA) provides such a unique opportunity to improve our understanding of how households make healthy eating choices. HHFKA has been signed into law in 2010, as a centerpiece of Michelle Obama's "Let's Move!"

[1]National Center for Health Statistics, National Health and Nutrition Examination Surveys, 1971–1974, 1976–1980, 1988–1994, and 1999–2018, U.S.

initiative to combat childhood obesity. It has allowed USDA to reform the National School Lunch Program (NSLP) and School Breakfast Program (SBP) by improving, for the first time, the nutrition standards for food sold at school on a nationwide scale. In 2012, the nutritional standards provision of the HHFKA went into effect, making the existing food options at school markedly healthier. Under the HHFKA, minimum standards for food elements served in the school cafeteria are mandated, setting the milk to be fat-free or 1% fat, requiring all grains to be whole-grain rich. In addition, caps and floors of micro-nutrients are enforced to guarantee food quality. For example, the meal programs have to serve healthy nutrients like fruits, vegetables, and whole grains satisfying a minimum level; unhealthy components such as sodium, sugar, and fat are limited by a maximum. In order to alleviate overeating, portion sizes are also reduced in meals. With meal quality at school more guaranteed than ever, a natural question is, "Does it spillover to household food purchases?" This paper examines *whether* households also changed their food purchase behaviors when they were exposed to more nutritious food for their children at school after the HHFKA, if yes — *which* households altered their food shopping results, and *how* their shopping baskets have changed.

Using a Difference-in-Difference design, we measure changes to both the quantity and quality of shopping baskets for households with school-age children in response to the new nutritional standards for school meals set by the HHFKA. We get information on the food quantity they purchase leveraging the Nielsen Homescan panel data, and the quality of food utilizing a complement dataset on nutritional information from the Label Insight. We compare the shopping baskets of these households to that of households without kids, since the former are more likely to benefit from the HHFKA, while the latter is unlikely to be impacted by this legislation. We ensure that the comparison group is as similar as possible to the treated group on both demographics and pre-treatment purchase patterns.

Our result shows a small but statistically significant change in food quantity purchased among households with school-age kids, with an average 4.1% reduction in calorie purchases per adult equivalent, but no significant change in food quality. Additional analyses help build confidence that these changes are due to children in these households having their food needs better met at school, with breakfast or lunch-related categories declining in grocery shopping. They could potentially either purchase school meals more often or finish more portions of the meals (Cohen et al., 2014), both leading to their calorie needs reduced at home. Further evidence shows that the quantity changes are driven primarily by the reduction in purchases of "kids-friendly" UPCs.

However, not all households should be expected to respond to the healthier meals at school in the same way. Prior policy research has often investigated heterogeneous differences across demographic dimensions. For instance, Kinderknecht et al. (2020) has found, through a national survey, that the Health Eating Index (HEI) for lunch meals of students from low-to-middle-income households demonstrated a significantly larger improvement after the HHFKA. Similarly, Kenney et al. (2020) presented a substantial decline in the risk of obesity only for children in poverty after the legislation. As these families, due to financial limitations and potentially a lack of food literacy (Wijayaratne et al., 2018), are more likely to eat less healthy foods in general, they are also possibly the ones who benefit more greatly from the healthier school meals.

We examine the effect brought by the HHFKA on household shopping baskets from another angle of view. We aim to find out whether the quantity and quality of a household's pre-treatment shopping basket, which represents the behavioral food purchase patterns rather than pure demographic properties, are predictive of how they respond to the implementation of nutritional mandates on school meals. Using a simple median split, we construct a two-by-two matrix: households that are either "low" or "high" on food purchase quantity and quality prior to the nutritional standards provisions of the HHFKA taking into effect.

Interestingly, we find that the reduction in food quantity purchases at retailers is driven by a single quadrant — households that used to purchase both a low quantity of food and low quality of food. The reduction in calories purchased by this quadrant is quite large, ranging from 13% to 19%, depending on the model specification used.

With regard to the food quality purchased by more involved households, the effect could go either way. According to the Spillover-Crossover Model (Bakker and Demerouti, 2013; Westman, 2013; Dolan and Galizzi, 2015), we tend to behave consistently with our prior actions and beliefs. If kids receive nutritional education and eat healthy at school, they will probably feel more consistent with keeping the same dietary habit at home, bringing potential conversations and changes to parents. With this mechanism, the quality of household food purchases should be enhanced after the HHFKA took place. The licensing effect could, however, endorse the contrary story, stating that if our prior choice helps boost a positive self-concept, we could subsequently license the choice of a more self-indulgent option (Khan and Dhar, 2006; Sachdeva et al., 2009). The moral licensing predicts that eating healthier at school may lead to the perception of "health mission completed" and thus indulge in more tasty food, which is often unhealthier.

We find a statistically significant decrease in the nutritional quality of the shopping basket for both the above-mentioned quadrant and another quadrant that also buys food of low quality, but in large quantity. This potentially supports a licensing story in the spillover effect of the HHFKA to household food purchases, especially for the high-quantity, low-quality quadrant, since they do not appear to be changing their total quantity of calories purchased. These two quadrants, who, pointed by our evidence, that might have already been participating in the school meal programs to some extent prior to the treatment, may have marked the healthy eating goal "completed" from the improvement in nutrition standards at school, and relaxed the extent to which they eat healthily at home. While these effects are significant, their impact is thankfully dwarfed by the improved healthiness of school meals.

A household's pre-treatment purchase habit, including both the quantity and the quality of food purchased at the retailers, may reflect several underlying drivers. Financial constraint could be one potential driver — households that purchase low quantities and low quality of food may be unable to afford a full, healthy diet (Palazzolo and Pattabhiramaiah, 2021). Time constraints may be another reason behind this — busy households may not have time to construct healthy meals at home and prefer to outsource their food consumption away from home, either by eating out more in restaurants, or via higher pre-treatment participation in school meal programs. According to Farris et al. (2016), the top two motivational factors for NSLP participants across all schools are "convenience" and "saving time". Moreover, the low quality of food choices at groceries could potentially be explained by a lack of food literacy(Wijayaratne et al., 2018).

To examine which of these drivers are likely to play a role in governing which households reduce their purchases of food for home consumption, we investigate which demographic characteristics are predictive of the membership in each of the quadrants, especially the ones with low quantity and low quality prior to the HHFKA, since they reduced their quantity purchases the most and also decreased the food quality. We find that low-income households are not especially likely to belong to the "low quantity, low quality" quadrant. We do, however, find that single-parent households, households with kids spanning multiple age groups, and households with two working parents are much more likely to be in this quadrant, providing some support for the notion that these households are time constrained. Moreover, we find some evidence that the education of the female head of a household is also predictive: households with a female head that lack a college degree are also especially likely to belong to this quadrant. This suggests that a lack of food literacy may also drive membership to this group, and contribute to higher utilization of the healthier school meal programs. Taken together, our analyses suggest that the nutritional standards implemented as part of the HHFKA made school meals more attractive to time-constrained, food illiterate households — households that perhaps found the prospect appealing to outsource their children's meals to a trusted provider of healthy meals.

Our research contributes to both (*i*) the stream of marketing literature aiming to understand what leads households to make healthy or unhealthy eating choices, and how to encourage healthier diets, and (*ii*) the policy literature on the HHFKA. With respect to the marketing literature, we find encouraging evidence that at least some households, when presented with an opportunity to provide healthier meals for their children, took this opportunity. Interestingly, the households that did so were not necessarily the households one might have expected — it was not the healthiest households, but the least healthy households that gravitated towards school meals. With respect to the policy literature, we provide evidence that one potential unintended consequence of the HHFKA is the reduction of healthy eating at home, perhaps due to a licensing effect. Fortunately, it was trivially small compared to the improved healthiness of school meals. Additionally, while other research has debunked detractors' claims that kids would dislike healthier meals and not consume them in full, leading to more waste (Schwartz et al., 2015; Cohen et al., 2014), our research goes further, suggesting that a non-trivial proportion of children had more of their caloric needs met at school, such that less food was needed at home.

The rest of the paper is organized as follows. Section 2.3 introduces the HHFKA and the new standards mandated by it. We then present the data we use for the research questions in Section 2.4. In Section 2.5, we provide model details and analysis of the impact of the nutritional changes in school meals under the HHFKA on household food purchase behaviors. Section 2.6 concludes our findings and elicits implications.

2.3 Healthy, Hunger-Free Kids Act of 2010

The Healthy, Hunger-Free Kids Act (HHFKA) provides a unique opportunity to improve our understanding of how households make healthy eating choices. The HHFKA has been signed into law in 2010, as a centerpiece of Michelle Obama's "Let's Move!" initiative to combat childhood obesity. It has allowed USDA to reform the National School Lunch Program (NSLP) and School Breakfast Program (SBP) by improving, for the first time, the nutrition standards for food sold at school on a nationwide scale. The NSLP and the SBP were introduced back in 1946 and have been lauded as the nation's two most successful programs. They granted access to federally assisted meals to millions of households meeting financial eligibility criteria. The main focus of these programs historically was to improve households' access to food.

In 2012, the nutritional standards provision of the HHFKA went into effect, making the existing food options at school markedly healthier. Under the HHFKA, minimum standards for food elements served in the school cafeteria are mandated, setting the milk to be fat-free or 1% fat, requiring all grains to be whole-grain rich. In addition, caps and floors of micro-nutrients are enforced to guarantee food quality. For example, the meal programs have to serve healthy nutrients like fruits, vegetables, and whole grains satisfying a minimum level; unhealthy components such as sodium, sugar, and fat are limited by a maximum. In order to alleviate overeating, portion sizes are also reduced in meals. With meal quality at school more guaranteed than ever, a natural question is, "Does it spillover to household food purchases?" This paper examines *whether* households also changed their food purchase behaviors when they were exposed to more nutritious food for their children at school after the HHFKA, if yes — *which* households altered their food shopping results, and *how* their shopping baskets have changed.

The United States Congress authorized special funding for child nutrition programs in 2010 via the HHFKA. This provision authorized the US Department of Agriculture (USDA) to initiate significant reforms to nutrition policies that apply to food served in public schools. As a major initiative of its kind in almost 30 years, the HHFKA tightened nutritional standards governing the SBP and NSLP, strictly regulating the nutritional composition of foods served as part of these meals. These changes officially took into effect at the start of the 2012 school year (August 2012). Table 2.1 summarizes the key changes as they applied to each component of these school meals.

The HHFKA also tightened the nutritional standards for "competitive foods" (food and beverages offered for purchase in school cafeterias, vending machines, etc.). However, these revisions did not take effect until the start of the 2014 school year (July 2014). Under the guidelines of HHFKA, the USDA significantly tightened the nutrition guidelines with a view to combating the childhood obesity problem. For the first time, the new standards mandated a ceiling on total calories per meal, the total sodium for both meals and a la carte items

Component / nutrient	**Requirement**
Milk	Must be 1% fat or fat-free; flavored must be fat-free
Fruits	Increased servings: must be served daily, no more than half of servings can be juice
Vegetables	More variety, with weekly minimum requirements for dark green, red/orange vegetables; beans/peas (legumes); starchy and other vegetables to meet weekly amount required
Grains	At least one-half of all grains must be whole-grain rich; in school year 2014-15, all grains offered must be whole-grain rich.
Meat (or alternate)	Minimum servings required
Calories	Maximums as well as minimums established for weekly average meal served
Fat	Limit on total fat removed Saturated fat <10% of calories Trans-fat 0g per serving
Sodium	New limits phasing in with final targets to be met by 2022

Table 2.1: Summary of New Nutritional Requirements for School Meals

The HHFKA included two major sections. Section 104 included changes to who is eligible to participate in subsidized school meal programs, as part of the Community Eligibility Provision (CEP). Section 104 eased the qualification criteria for households' eligibility in such meal programs, automatically rendering entire school districts eligible so long as at least 40% of enrolled students met certain income-based qualification criteria. Changes to the CEP were rolled out in a staggered fashion.

In contrast, a separate section of the HHFKA - Section 208 focused only on revising the nutritional standards, with a focus on improving the quality of food provided in public schools. Our study focuses on investigating the impact of revisions to nutritional quality on households' food choices in grocery stores, while acknowledging and controlling for the staggered implementation of changes to CEP guidelines that only impacted access to subsidized meals (but not nutrition guidelines). We discuss this in more detail in the next section.

2.4 Data

We combine two datasets together to answer our research questions. One is the Nielsen Homescan Panel data, provided by the Kilts Center at the University of Chicago. The other is the Label Insight's open data, which offers nutritional information for food sold in the U.S. on the UPC level. We utilize the two data sets to investigate the policy impact of the HHFKA outside of schools — in household food purchases to be specific.

The Nielsen panel data records grocery shopping activities on the household-daily level with samples throughout the nation. We are able to locate what food and how much they have purchased during their cooperating periods. Since the HHFKA went into effect in 2012, we utilize records between 2011 to 2013 to study the effects. In total, we have observed grocery shopping activities from 42,125 households in 41 states. We exclude seven states (Illinois, Kentucky, Michigan, New York, Ohio, West Virginia, the District of Columbia) where CEP was also effective during our two-year time window. As we mainly want to examine the influence of the quality change in food, the expanded access to free/subsidized meals served in schools is a potential confounding factor to eliminate. In addition, the Homescan data does not cover Hawaii and Alaska.

We then match the nutritional information from the Label Insight's data, so that we know the proportion of macro-nutrients in the food they purchased. We compare the differences between the treated and control groups on mainly two sets of variables – food quantity purchased (including calories and servings) and the overall health index of foods (including two widely adopted health indices in academia).

2.4.1 Treatment assignment

We set households without school-age kids as the control group, since they have little chance to receive influences from the HHFKA of 2010. Ideally, households who actually participate in school meal programs should be the treated group as they directly engage with the new nutrition standards for meals. However, since we don't have the exact participation information on the household level, we instead use households who have school-age children as the treated group, for they as a group are more likely to be affected by the HHFKA, compared to the control group. To be more specific, we use households with kids between 6-17 years old as the treated group.

Based on the above treatment assignment method, one potential issue is that some households may switch group status during our examination period. For example, a household could have a kid aged five in 2011 which makes them part of the control group; in 2012 and 2013, however, this household fell into the treated group since the kid went into school age in 2012. The contrary situation could happen as well – a household might have a kid aged 17 in 2012 so they are likely to be influenced by the HHFKA. At some time in 2013, as the kid grew over 18 years old, this household then switched to the control group. Considering that households like these receive mixed treatment assignments, we have excluded them from our study. A closer look at the excluded households shows that they are mostly households with both a male head and a female head in these families, middle-aged. More than half of them have a 13-17 adolescent (who probably ages 17, but we only have age groups instead of accurate age in the data) in house and 43% only have one kid in this age range.

We further improve the estimation of causal effects by applying the Coarsened Exact Matching method (CEM), which is a common way to reduce the imbalance in demographic variables between the treated and the control group. We match both groups using covariates including household income levels, the number of parents in the household, age of the male head and female head of a household, the location (state) of the household, as well as the time they participate in the panel.

2.4.2 Dependent Variables

We measure the influences of the HHFKA on household food purchases using both quantity changes and quality changes. For quantity changes, we use "calories per adult equivalent" and "servings per adult equivalent" of packaged foods; for quality changes, we examine both Scores from the Nutrient Profiling Model (NPS) and Healthy Eating Index (HEI).

Calories per adult equivalent Calories are a critical measure of whether the HH-FKA at school affects household food purchases, encouraging meal substitution from home to school. We constructed a log variable of calories called the "log calories per adult equivalent" (abbreviated as $ln(cal)$). We calculate this variable following two steps as below.

First, we calculate the number of "adults equivalents" of each household, following prior work (Allcott et al., 2019b; Palazzolo and Pattabhiramaiah, 2021). The number of children is converted to a proportioned number of an adult based on the ratio of their calorie needs to that of a nationwide average adult. For example, a child under age five is factored as a half adult, with regard to calorie needs.

Second, with "adult equivalents" per household calculated, we could easily calculate the "calories purchased per adult equivalent" for each household, using the total calories in the food they have bought on the daily level, divided by the "adult equivalents". We then take the natural log value of this variable to suppress extreme value influences. In this way we get our first dependent variable "log calories per adult equivalent". Our data shows that the average household size is 2.3, including adults and children. Based on the food they purchased, the daily average calorie for a household is 755.5.

Servings per adult equivalent An intuitive measure of food quantity purchased is the number of servings of packaged UPC goods. While calories are a crucial assessment in diets, changes in categories like vegetables can be hardly detected since they are very low in calorie density. We thus use "servings per adult equivalent" as a complementary estimation of food quantity. We calculate this variable using similar procedures as in calculating "calories per adult equivalent". The average daily servings a household buys is 12.9.

NPS We calculate a score based on the Nutrient Profiling Model (NPM), developed by the Food Standards Agency in the UK to calculate the nutritional quality of foods and drinks. This model has been widely accepted as a standard to calculate nutritional quality (Poon et al., 2018; Dubois et al., 2021). We follow the approach by Palazzolo and Pattabhiramaiah (2021) and build this health index by calculating a score between -40 to 15 for each UPC, according to the volume of three healthy components (protein, fiber, and fruits/vegetables/nuts) and four unhealthy components (saturated fat, sugar, sodium, and calories) per 100 grams. This score, since it is originated from the Nutrient Profiling Model, is abbreviated as *NPS* in this paper.

From our data set, the NPS for an average household is -5.47, which is important to denote here. A positive change means that household food purchases get more healthy than before, and negative means more unhealthy.

HEI The HEI is a tool to measure diet quality, developed by USDA in 1995, which gets updated every few years. We use the latest version — the HEI-2015 score — for this paper. It is mainly designed to assess how well Americans follow key dietary recommendations. For every 1000 calories, it recommends a minimum level of healthy macronutrients (1.2 cups of fruits, 1.3 cups of vegetables, 1.3 cups of dairy, 1.5 ounces of whole grains, 3.3 ounces of protein) and a maximum for unhealthy macronutrients (1.1 grams of sodium, 6.5% of energy from added sugar, 8% of energy from saturated fat).

We calculate the HEI based on the same method adopted by Palazzolo and Pattabhiramaiah (2021). In general, if a household increases the purchase of unhealthy components in food or drops that in healthy components, the HEI will decrease; if a household increases the purchase of healthy components or reduce that of unhealthy macronutrients, the HEI will increase accordingly.

2.5 Model

We are interested in finding out whether the HHFKA of 2010 at school also has influenced household food purchases or not, and along which direction if it has. With the households with school-age children as the treated group and those without kids at school as the control, we measure the difference before and after the treatment. We adopt a Difference-in-Difference based identification strategy to measure the causal effect of the HHFKA on both the quantity of household food purchases and the overall healthiness of food.

Further, we take the analysis down to heterogeneity effects. Different from previous research which focused on demographic variances, we categorize households by their pre-treatment purchase patterns, as we would like to see if their behavioral characteristics play an important role in whether they leverage such an improved meal option. To be more specific, we categorize the households by a two-by-two matrix: whether they purchase "low" or "high" on quantity and quality prior to the HHFKA.

2.5.1 Model-free Evidence

Before we apply the Difference-in-Difference framework, we first plot the monthly difference between the treated and the control for the four dependent variables (Figure 2.1). From both calories and servings per adult equivalent, the treated group overall purchase less quantity of food than the control. There is a clear gap between the control and the treated group and the difference is getting larger over time. For the healthiness of foods, NPS and HEI slightly differ from each other, with the latter showing closer levels of food quality. However, the dropping trends agree with each other that the two groups are getting more disparate after the HHFKA.

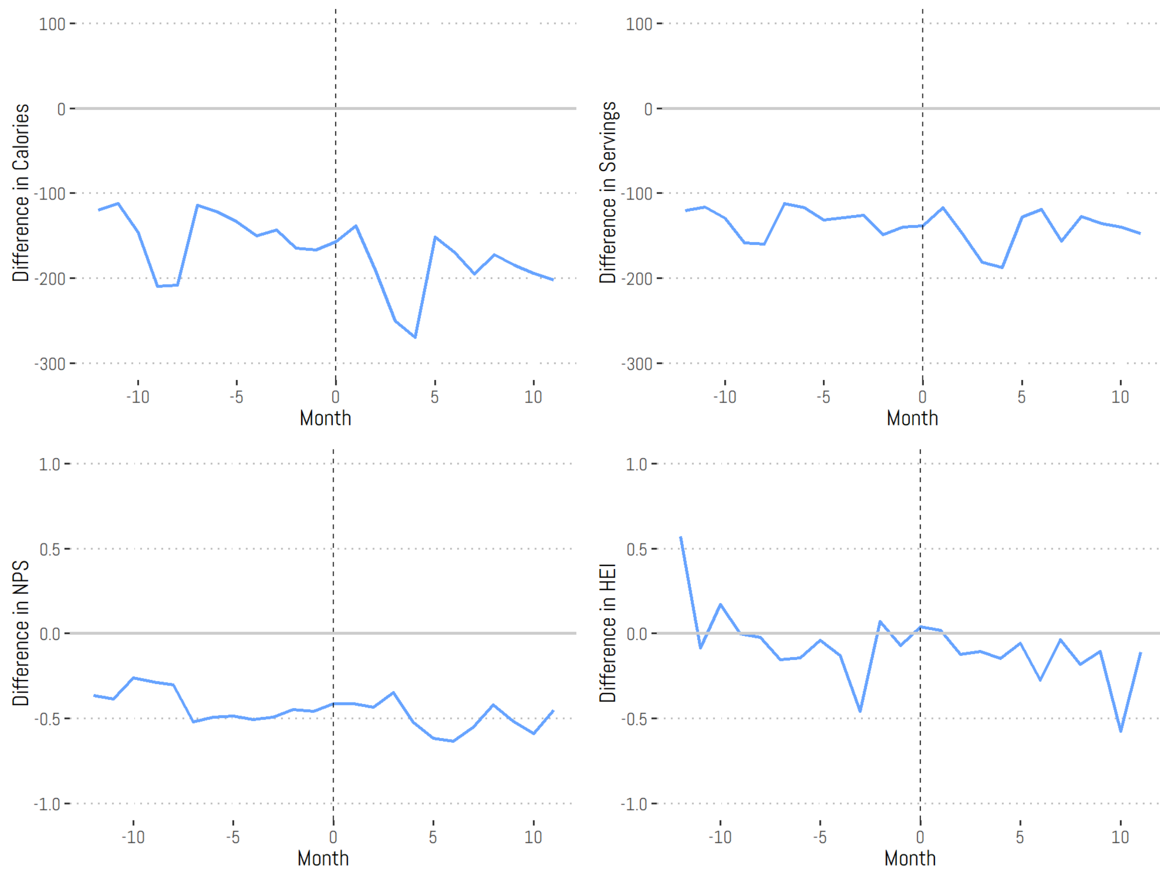

Figure 2.1: Trends of Quantity and Quality Differences for the Treated and Control Groups

In summary, the model-free evidence tells us that, compared to the control group, the treated group with school-age kids reduced their food quantity purchase; in the meanwhile, the overall healthiness of foods deteriorated as a result as well.

2.5.2 Pre-period Trends

A closer look at Figure 2.1 could easily identify seasonality differences for the treated and control groups, which might be caused by school holidays. Thus, it is hard to have the pre-period trends parallel displayed without any pre-processing. To address this seasonal pattern difference, we control the seasonality by taking a first-difference of dependent variables for the same calendar month.

Figure 2.2 shows that after controlling seasonality, the pre-trend differences are not distinguishable across groups.

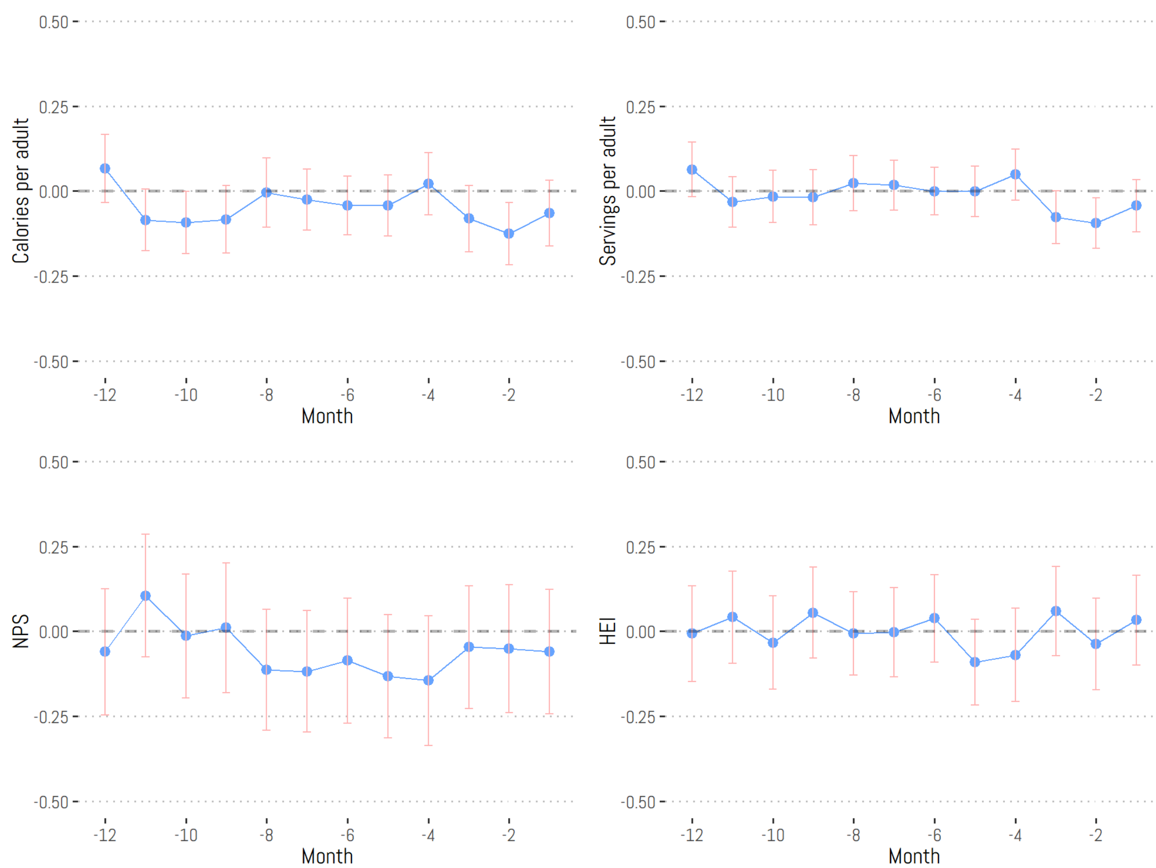

Figure 2.2: Pre-trend difference of DVs after first differencing

2.5.3 The Main Effect

Our Difference-in-Difference model specification for household calorie purchases is as follows.

$$ln(Cal_{ht}) = \alpha_{hl} + \beta(SchoolKids_h \times AfterAct_t) + \lambda_{lt} + \epsilon_{ht}, \tag{2.1}$$

where Cal_{ht} denotes calories per adult equivalent and $SchoolKids_h$ denotes whether this household has schooling kids (treated) or not (control). $AfterAct_t$ is a time indicator of whether the HHFKA requirements are mandated or not. The coefficient β captures the causal effect of the new nutrition standards of the HHFKA on household calorie purchases. As we estimate the model effect with the dependent variable in log value, β represents the semi-elasticity, which means the percentage change in calories when households are treated.

We hold the household-locality fixed effects to control for heterogeneity on the household level and locality level (α_{hl}) which may introduce bias in households' location choices. We also hold the locality-time fixed effect (λ_{lt}) to control for different evolution along time at one location. The main effects on all dependent variables are measured using the same estimation approach as in Equation 2.1.

To further improve the estimation of causal effects, we apply the Coarsened Exact Matching method to make the treated and the control group more comparable. We match the two groups on key variables that depict a silhouette of a household, including household income levels, whether it is a single-father, single-mother family, or one with both parents present, the age of the household male/female heads, the state they reside in, and the time they participate in the panel. In this way, although the treated and the control are different in whether they have school-age kids, the composition of households on both sides is similar to each other.

As Table 2.2 demonstrated, we estimate that there is a 4.1% decrease in calories per adult equivalent, and a 3.6% of reduction in the number of servings per adult

Table 2.2: Main Effects of HHFKA on Household Food Purchases

	(1) Calories	(2) Servings	(3) NPS	(4) HEI
β	-0.041^{**}	-0.036^{***}	-0.050	0.002
	(0.016)	(0.014)	(0.034)	(0.022)
Constant	6.086^{***}	1.866^{***}	-5.463^{***}	-1.758^{***}
	(0.001)	(0.001)	(0.003)	(0.002)

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

equivalent. A lower decrease in servings than calories probably indicates that the reduction mainly takes place with foods of higher calorie density. As the decrease in calories and servings are pretty close to each other, we will focus on calories in the analysis, while keeping servings as a consistency measure. Translating to actual food quantity measures, the numbers aforementioned mean that parents purchase 31 fewer calories each day for each family member. For a family with both parents and two kids, the total daily calorie reduction is close to 124.

If an average kid needs 350 calories for breakfast, the reduction equals to 18% of the calorie needs of breakfast for the two kids. Of course, in our data we estimate the effects for all parents with school-age kids, which may include both households participating in school meal programs and those who do not. Considering that the average SBP participation rate prior to HHFKA was 12% and that for NSLP was 32%[2], we could reasonably say that the calorie reduction reflects about one breakfast for a kid on average. Our results show indirect evidence that parents do substitute meals away from home for kids, given that food quality at school improves, even if there's no change in subsidies. The overall healthiness, measured by both the NPS and the HEI, of food they purchase does not significantly vary from previous.

[2]Data source: USDA Food and Nutrition Service — Child Nutrition Tables

Of course, the above deduction only gives a rough idea of what the calorie deduction size might mean. To find out in detail what households have decreased purchasing, we further examine in which food departments and food categories has the reduction taken place in the next section.

2.5.4 Changes by Food Types

We label food UPCs according to which departments they usually belong to in a supermarket: dairy, deli, dry grocery, fresh produce, frozen food, and packaged meat. As displayed in Table 2.3, the reduction in calories mainly takes place in dairy (milk, eggs, yogurt, dough products, cheese, desserts dairy, dressings, and spreads, etc.), dry grocery (bread, baked goods, flour, packaged drinks, breakfast food, pasta, soup, candy, cookies, snacks, etc.) and frozen foods (prepared foods, pizza, snacks, unprepared meat, frozen vegetables, etc.). The number of servings also gets declined, which is basically consistent with calories, except for a marginally significant reduction in Deli (dressings, salads, prepared foods - deli, etc.), which may be less dense in calories and thus does not show a significant reduction in "calories per adult equivalent". However, food with a closer relation to dinners, such as fresh produce and package meat, has no significant changes in both quantity and quality. From the list above that shows significant differences, we could see that parents mostly decrease the purchase of breakfast-and-lunch-related foods and snacks. It is reasonable to speculate that they were taking advantage of the school meals with higher standards under the HHFKA.

Table 2.3: Departmental Heterogeneous Effects of HHFKA

	Calories	Servings	NPS	HEI
Dairy	-0.33^{***}	-0.23^{***}	-0.13^{*}	-0.02
	(0.107)	(0.071)	(0.067)	(0.028)
Deli	-0.12	-0.09^{*}	-0.08	0.02
	(0.075)	(0.046)	(0.076)	(0.054)
Dry Grocery	-0.74^{***}	-0.58^{***}	-0.01	0.01
	(0.266)	(0.185)	(0.042)	(0.030)
Fresh Produce	-0.06	-0.02	0.02	-0.41^{**}
	(0.062)	(0.031)	(0.016)	(0.168)
Frozen	-0.44^{***}	-0.27^{***}	0.05	0.03
	(0.145)	(0.087)	(0.058)	(0.048)
Packaged Meat	-0.01	-0.01	-0.10	-0.01
	(0.080)	(0.05)	(0.088)	(0.019)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In Table 2.3, we also calculate an NPS index for the healthiness of food purchased in each department. Results show that the main deterioration happens in the dairy department. A simple decline in the purchase of dairy proportionately may not fully explain the decline in healthiness, as the NPS index is calculated per 100 grams. However, if households mainly reduced purchases of milk and eggs, which are closely related to breakfast, then it is possible that the portion of healthy components (like protein) goes down and the unhealthy components (such as added sugar) goes up, ending up with a slide-down of the NPS in dairy overall.

For dry groceries and frozen foods, although calorie purchases diminish, the healthiness of these categories does not change significantly. HEI shows a different angle that the healthiness of fresh produce part decreases. This is probably because the portion size of fresh produce falls below the recommended consumption rate (1.2 cups of fruit, 1.3 cups of vegetables per 1,000 calories). However, since the calorie density is very low compared to other food departments, the change in fresh produce is not detected in calories. What's more, the measure of servings only applies to UPC, which excludes quite a large portion of fresh produce from the calculation.

While changes in food purchases at different food departments help us get a picture of what households are purchasing less, and which categories are getting unhealthier, it is still a few steps far from betraying the real-life show centering on the HHFKA. To have a closer peak, we further categorize foods into "sub-departments", with clearer utilizing purposes. Table 2.4 offers a more detailed description of food purchase changes in households after the HHFKA. Reductions in both calories and servings demonstrate that parents buy less soda, breakfast, and prepared food from supermarkets, which are indicators of substitution for both breakfasts and lunches. Additionally, they purchase fewer cooking ingredients, which is a sign of less cooking at home.

Estimates of the NPS show that they pay less attention to healthiness when they purchase unprepared food and snacks, although the total quantity in both calories and servings does not change. This tells a potential story of self-licensing. However, the HEI presents a different story, that their purchase of snacks is getting healthier. The different evaluation of NPS and HEI on snacks is worth investigation.

Table 2.4: Sub-Departmental Heterogeneous Effects of HHFKA

	Calories	Servings	NPS	HEI
Beverage: Juice	−0.06	−0.04	0.08	0.02
	(0.056)	(0.037)	(0.07)	(0.055)
Beverage: Soda	-0.11^{**}	-0.10^{**}	0.03	0.01
	(0.053)	(0.035)	(0.161)	(0.034)
Beverage: Other	0.00	−0.02	0.14	0.05
	(0.05)	(0.038)	(0.151)	(0.067)
Breakfast	-0.08^{**}	-0.06^{**}	−0.04	0.00
	(0.037)	(0.026)	(0.054)	(0.03)
Cooking Ingredients	-0.13^{**}	-0.08^{**}	0.11	0.19
	(0.056)	(0.04)	(0.087)	(0.165)
Prepared Food	-0.09^{**}	-0.06^{**}	−0.05	0.00
	(0.04)	(0.026)	(0.04)	(0.013)
Unprepared Food	−0.07	−0.04	-0.14^{*}	−0.19
	(0.048)	(0.030)	(0.077)	(0.018)
Fruits & Vegetables	−0.04	−0.03	0.01	0.03
	(0.042)	(0.028)	(0.044)	(0.084)
Snacks	0.01	−0.01	-0.09^{*}	0.03^{**}
	(0.031)	(0.022)	(0.046)	(0.016)
Toppings	−0.05	−0.05	0.06	0.01
	(0.05)	(0.035)	(0.063)	(0.054)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In addition to the above analysis, we have also examined whether the changes are with respect to children eating content. We labeled UPCs as “kids-friendly” or not. “Kids-friendly” UPCs include items like Frosted Flakes, Berry Colossal Crunch, Marshmallow Mateys, Fruit Loops, Cocoa Pops, etc. Our results (Table 2.5) exhibit that calories of “kids-friendly” food products have significantly decreased by 5.0%, and servings declined by 7.4%; while those labeled as “non-Kids-friendly” do not show significant changes in food quantity purchased. Moreover, the NPS of purchased “kids-friend” UPCs at home also declined to a larger degree than those in the departmental and sub-departmental categories, though the effect on HEI is not significant. Interestingly, the healthiness of “non-Kids-friendly” food also deteriorated.

Table 2.5: Heterogeneous Effects on (non-)Kids-friendly Food Purchases

	(1) Calories	(2) Servings	(3) NPS	(4) HEI
kids-friendly UPCs	-0.050^{**}	-0.074^{***}	-0.168^{***}	-0.006
	(0.025)	(0.016)	(0.052)	(0.028)
non-kids-friendly UPCs	-0.005	-0.007	-0.199^{**}	-0.063
	(0.016)	(0.014)	(0.085)	(0.063)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Taken together, this set of analyses shows that the reduction of calories did focus on “Kids-friendly” food, indicating that households are outsourcing kids’ food rather than that for adults, putting a strong tie between household food reduction to school meal participation. The unintended consequence of self-licensing on food quality also convincingly relates to "Kids-friendly" food. What's more, parents also lean towards more unhealthy choices on "non-Kids-friendly" UPCs due to the "completed" health eating goal brought by the HHFKA.

2.5.5 Heterogeneous Effects

Another interesting finding from our results is that the HHFKA shakes the habit of food purchases in some households further away than others. We group households by their food purchasing habits prior to the HHFKA, including the amount of food they usually purchase and the overall healthiness (NPM) of the food, which indicate whether they are prone to buy more healthy food or less healthy food. On both dimensions, we split the population by the median, and thus get a two-by-two categorization. This grouping method enables us to see how their food purchase patterns have been revised by the quality-improved meal options at school under the HHFKA. Table 2.6 displays the summary statistics of the four groups, in which "$\downarrow$" denotes "low" levels of the variable ensues and "$\uparrow$" denotes "high" levels for the sake of simplicity.

Table 2.6: Summary Statistics of Household Groups

	All	↓Cal ↓NPM	↑Cal ↓NPM	↓Cal ↑NPM	↑Cal ↓NPM
Calories	755.51	409.37	948.85	425.95	937.46
	(345.45)	(165.97)	(275.43)	(151.82)	(268.31)
Servings	12.88	6.94	15.45	7.75	16.27
	(9.52)	(5.75)	(8.36)	(6.27)	(10.57)
NPS	−5.47	−7.65	−7.27	−3.92	−4.03
	(2.33)	(2.04)	(1.55)	(1.63)	(1.37)
HEI	−1.74	−2.40	−2.43	−1.12	−1.28
	(1.35)	(1.27)	(1.08)	(1.37)	(1.19)

Table 2.7 presents the heterogeneous effects of the four groups based on their food purchase habits prior to the HHFKA taking into effect. Model (1) demonstrates that the overall effect of the reduction in calorie purchases is driven by one single quadrant – households that often buy less food quantity with low overall healthiness level prior to the HHFKA. Post the HHFKA, they are encouraged to purchase much less from grocery stores and supermarkets, substituting in a much greater degree of food calories away from the home purchase than the average.

An unintended consequence we find of the HHFKA is that households with less healthy food habits, according to their purchase records, lean towards an even less healthy eating habit afterwards. Among them, those who purchase less and unhealthy calories take the largest negative shock, with the NPS lowered by -0.31 and HEI declined by -0.12. This result could be a sign of self-licensing effect (Khan and Dhar, 2006; Sachdeva et al., 2009), that consumers may feel licensed to refrain from good behaviors when they have accrued a surplus of moral currency. Here, the good behavior is to keep a healthy diet and the surplus of moral currency is brought in by the more nutritious meals at school. Knowing that kids eat more healthily at school, parents may feel they could slack back and lower the healthiness standards for food consumed at home. This is an interesting behavioral result of the HHFKA, of which the psychological mechanism worth further investigation.

As a robustness check, we tried another version of estimate dividing households on the healthiness dimension using HEI instead of NPM, which is displayed in Model (2) and (4) in Table 2.7. The estimation results are consistent in the changing direction of effects on calories purchased and the health index, although the effect sizes shrink slightly. There are possibly two reasons for the differences in effect sizes. First, it could be explained by the definition of the two indices. In NPS, calorie is also factored as something "bad" that would bring the score down. Thus, the change in NPS also reflects the reduction in calories, while HEI does not show such a direct influence. Second, despite similar ranges of the two indices in our data (with HEI a bit tighter), HEI shows a smaller variance (2.58 compared to 3.82 of NPS). Households that consume the full recommended value of all healthy macronutrients will receive a value of 1 on each, and a value of -1 on each when they reach the recommended limit of the unhealthy ones. While NPS is calculated for each UPC first and then weighted-averaged together.

Table 2.7: Heterogeneous Effects of HHFKA on Household Food Purchases

	(1)	(2)	(3)	(4)
	Calories	Calories	NPS	HEI
Low cal, Low NPM	-0.187^{***}		-0.309^{***}	
	(0.052)		(0.090)	
High cal, Low NPM	-0.032		-0.182^{***}	
	(0.022)		(0.058)	
Low cal, High NPM	-0.018		-0.022	
	(0.044)		(0.072)	
High cal, High NPM	-0.021		-0.035	
	(0.021)		(0.043)	
Low cal, Low HEI		-0.118^{**}		-0.121^{*}
		(0.054)		(0.062)
High cal, Low HEI		-0.023		-0.079^{**}
		(0.022)		(0.036)
Low cal, High HEI		-0.067		0.042
		(0.043)		(0.048)
High cal, High HEI		-0.028		0.002
		(0.020)		(0.028)
Constant	6.088^{***}	6.088^{***}	-5.455^{***}	-1.754^{***}
	(0.002)	(0.002)	(0.003)	(0.002)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2.5.6 Who are the households driving the changes?

Finding out that households in the "low quantity, low quality" quadrant are the driving force of the effect, we are curious about who are these households, and why they respond to the HHFKA to such a large extent. Previous research has indicated some potential reasons for such households making use of the improved school meals. Relating to work in Allcott et al. (2019a); Palazzolo and Pattabhiramaiah (2021), one may easily notice that households with "food insecurity" seem to overlap with this quadrant in our case. They may be limited by their financial condition and could only afford a small quantity of food with low nutritional contents. Time constraint might be a second guess, since recent survey shows that the top two factors that motivating households to participate NSLP are "convenience" and "saving time" (Farris et al., 2016). In addition, work has shown that "lack of time to plan, shop, prepare and cook healthy foods" is also among the top barriers of healthy eating(de Mestral et al., 2017; Munt et al., 2017). Another possible reason is a lack of food literacy that prevents them from preparing healthy food at home (Munt et al., 2017; Wijayaratne et al., 2018).

To investigate which of these reasons play a role in governing which quadrant of households reduces food purchases at home, we inspect which demographic characteristics of the households are predictive of membership in each quadrant through a multinomial logistic regression. As Table 2.8 shows, households with large household sizes, single-parent households, and households with both parents fully employed are much more likely to be in the "low quantity, low quality" quadrant. This evidence supports the idea that people with time pressure are more likely to outsource food consumption away from home. With the HHFKA providing healthier food, they may potentially substitute towards school meals for their children.

Financial constraint is not predictive of households being in this quadrant, but is contributing to being in the "high quantity, low quality" group. In addition, high household income households, even purchasing less food, are more likely to purchase food of high quality. Moreover, we find that households with female heads highly educated (graduated from college or higher) are not likely to be in the "low quantity, low quality" quadrant, but in the "low quantity, high quality" group. This shows that a lack of food literacy may be another driving factor of the membership of these quadrants.

Combining the evidence in Table 2.8, we may conclude that the new nutritional standards mandated by the HHFKA made school meals more attractive to households with time pressure and food illiteracy. School meal programs provide them the convenience of food consumption outsourcing, since they have limited time to prepare food at home. In addition, although they do not have the knowledge to prepare a healthy diet and make healthy food choices, the national mandated nutritional standards make school meals a trusted vendor of healthy food.

Table 2.8: Multinomial Logit Regression of Quadrant Membership

	(1) ↓Cal ↓NPM	(2) ↑Cal ↓NPM	(3) ↓Cal ↑NPM
$<$ 3 household members	-1.036^{***}	-0.018	-0.485^{**}
	(0.219)	(0.191)	(0.218)
$>$ 4 household members	0.599^{***}	0.071	0.367^{***}
	(0.081)	(0.079)	(0.082)
household income \$40k-\$100k	-0.127	-0.253^{***}	0.019
	(0.094)	(0.095)	(0.103)
household income $>$ 100k	-0.079	-0.186^{*}	0.296^{***}
	(0.099)	(0.099)	(0.105)
no male household head	0.287^{***}	0.085	0.179^{***}
	(0.052)	(0.054)	(0.056)
no female household head	0.269^{***}	0.030	0.131
	(0.086)	(0.092)	(0.095)
female head fully employed	0.122	-0.090	0.009
	(0.134)	(0.133)	(0.142)
male head fully employed	-0.002	-0.107	-0.084
	(0.120)	(0.119)	(0.127)
both heads fully employed	0.352^{**}	0.253	0.322^{**}
	(0.153)	(0.155)	(0.161)
male head highly educated	-0.116	-0.353^{***}	0.103
	(0.073)	(0.076)	(0.075)
female head highly educated	-0.123^{*}	-0.040	0.121^{*}
	(0.070)	(0.072)	(0.073)
race: non-white	0.609^{***}	0.101	0.571^{***}
	(0.079)	(0.086)	(0.081)
Constant	-0.611^{***}	0.162	-0.847^{***}
	(0.132)	(0.126)	(0.140)

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

2.5.7 Robustness Checks

2.5.7.1 Public School Availability, Meal Program Participation

We supplement the results with a battery of systematic robustness checks and falsifications. One obvious confounding factor is that the HHFKA only affects meal programs in public schools. As a falsification, we exploit variation at the geographic level to conduct falsification tests which should help us alleviate the influence of unobservable confounding factors.

Specifically, supplementing with a nationwide data set from *GreatSchools*, we acquire data dimensions closely related to properties of schools as well as school meal participation rates. The rationale behind the falsification test is that: 1) if the enrollment in private schools is low, then the possibility of households with school-age kids participating in the meal programs is higher, and the higher the potential carry-over effects to household food purchases; 2) if the participation rate of school meal programs is higher, then similarly, the higher probability that households in this area would be affected by the HHFKA.

We match the GreatSchool data set with Nielsen's Homescan data as well as Label Insight's nutritional information data on the zip code level. For each zip code area, we calculate the enrollment rate of public schools. And for each public school, we calculate the meal program participation rate. Among the 23,303 zip code areas, since most of the areas do not have any private schools, the average enrollment rate of private schools is very low. The median private enrollment rate is 0% and the 75% quartile is 8.67%. The average meal program participation rate in public schools overall is 51.8%.

Our results find that calorie purchases have reduced to a larger degree in high public-school enrollment (12%) than in high private-school enrollment rate districts (6%). Effects on health index also get attenuated in high-private areas. In addition, calorie purchases have reduced to a greater degree (14%) in high meal participation areas than in the low equivalent areas (9%). Besides calories, effects on health index get also attenuated in these districts. Both falsification tests support our findings that calorie reduction and deteriorated health eating habits do highly correlates with public school meal program participation.

2.5.7.2 Restaurant Spending and Visits

Another factor that may confound our findings is that households may, for some reason, reduce their calorie purchases from grocery stores because they increase their food purchase away from home. For example, they might eat more at restaurants instead. To exclude this possibility, we utilized the Consumer Expenditure Surveys from 2011 to 2017, which contains consumer expenditure on food away from home (restaurants, steak houses, etc.) also on the zip code area level, the percent of households who visit these places where they might fulfill their need on food other than schools and at home, as well as the percent of households who visit, decided by the children, fast food restaurants more than once in the last 30 days. We match the survey data on the above-mentioned three dimensions depicting restaurant food consumption with our existing data sets. Figure 2.3 gives the trend of expenditure on food away from home for both areas with high school-meal-program-participation and low participation, which differ very little from each other.

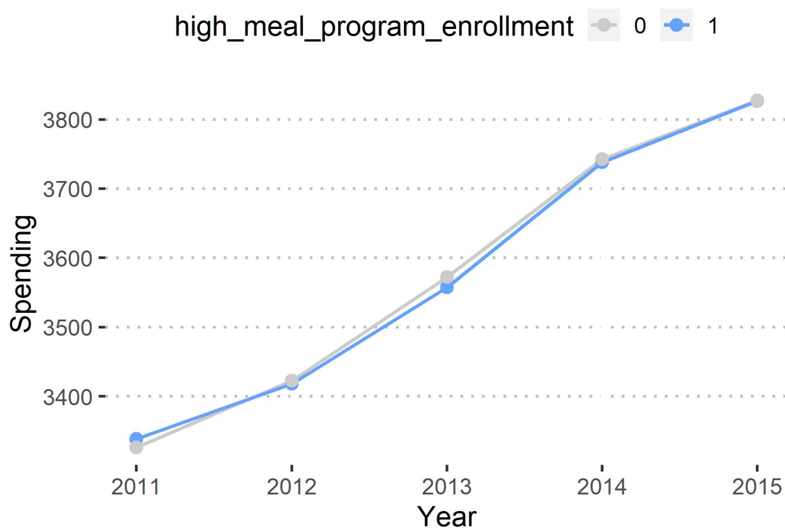

Figure 2.3: Trends of Household Hood Spending Away From Home

We then apply the Difference-in-Difference analysis to the effect of HHFKA on household spending on restaurants, yearly visit percent, and monthly visit percent influenced by children on the zip code level, with the control group being zip code areas with 1) a lower proportion of households in the "low-quantity, low quality" quadrant, 2) lower private school enrollment, and 3) lower meal program participation rate. accordingly. The idea behind these designs is that:

1. Since the effect is driven solely by households in the "low-quantity, low quality" quadrant, if an area has more such households and households do visit and spend more on restaurants systematically after the HHFKA, then this area shall display larger effects.

2. Since the HHFKA does not cover meals provided in private schools, if an area has a higher enrollment rate in private schools, the hypothesized effect should be smaller in such areas.

3. Higher meal program participation rate means potentially more households exposed to the direct influence of the HHFKA, and such areas should expect a higher effect on spending and visits away from home, if that should be a confounding factor.

The results in Table 2.9 display that household spending and visits to restaurants did not change significantly in areas with more "low-quantity, low-quality" households, and areas with more private school enrollment rates. It did show that household spending in restaurants did reduce significantly. However, instead of being a factor that cancels the effect of calorie substitution towards school, it offers evidence from another perspective that households may reduce food consumption at home and in restaurants, with the only possible growing channel to be at school for children's calorie needs.

Table 2.9: Spending and Visits on Restaurants

	(1) Spending	(2) Yearly Visit %	(3) 30d Visit %: Kids' Decisions
More LL Households	−6.608	0.000	0.000
	(5.100)	(0.000)	(0.000)
Higher Private Enrollment	−0.762	0.001	0.000
	(5.477)	(0.000)	(0.000)
Higher Lunch Participation	−13.446**	0.000	−0.000
	((6.485)	(0.000)	(0.000)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.10: Effects of HHFKA on Quantity and Quality of Food Purchase

	ATT	S.E.	z-score	C.lower	CI.upper
Calories	−0.0934***	0.0015	−62.27	−0.0963	−0.0904
Servings	−0.1056***	0.0011	−98.62	−0.1077	−0.1035
NPS	−0.1105***	0.003	−36.61	−0.1164	−0.1046
HEI	−0.0465***	0.0019	−23.92	−0.0503	−0.0427

2.5.7.3 Trajectory Balancing

Trajectory Balancing (Hazlett and Xu, 2018) offers several re-weighting schemes of the control units to achieve a better match for the treated units. We used the "mean balancing" method to re-weight the control group units in our data set to align pre-treatment outcomes and covariates.

Table 2.10 displays stronger and more significant effect of HHFKA on household food purchases, possibly due to the benefits of the Trajectory Balancing method. The mean balancing method we used re-weights the control households based on not only covariates but also pre-treatment outcomes, to ensure that the treated and the (re-weighted) control are approximately equal (see Figure 2.4-Figure 2.7). Results of Trajectory Balancing show a reduction of both calories purchased and the number of servings by around ten percent.

In addition, the overall healthiness, measured by NPS or HEI, deteriorates after the new nutritional standards of school meals went into effect. Figure 2.4 - Figure 2.7 also reveal that, unlike calories and servings, the effect on health indices does not ring immediately, but instead is delayed by a month or two. During Month 5 of the academic calendar, the healthiness of the treated and control gets overlapped with each other, since it is December in the natural calendar. This indirectly points out that the changes in healthiness is related to school meals.

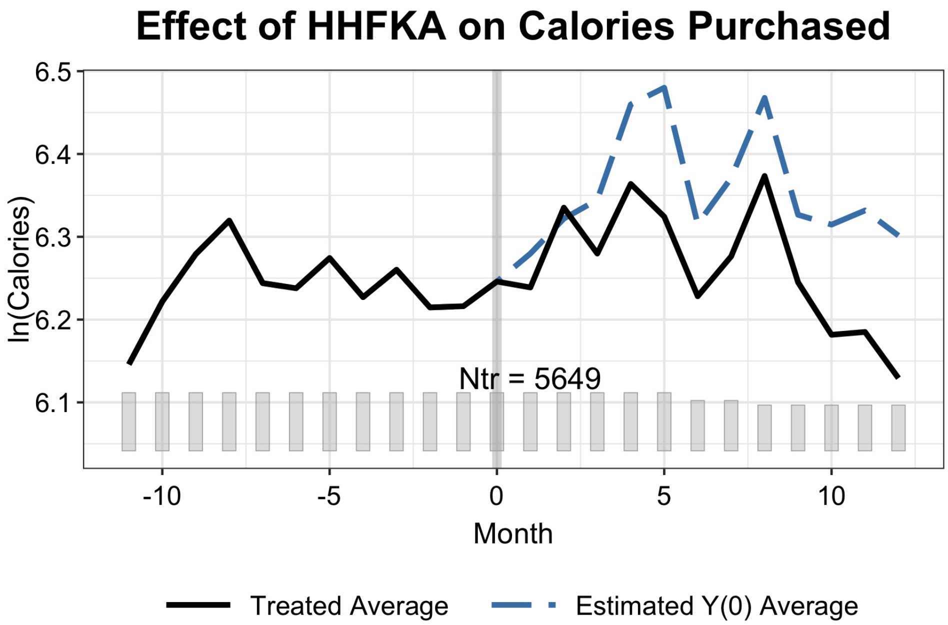

Figure 2.4: Effects of HHFKA on Calories Purchased

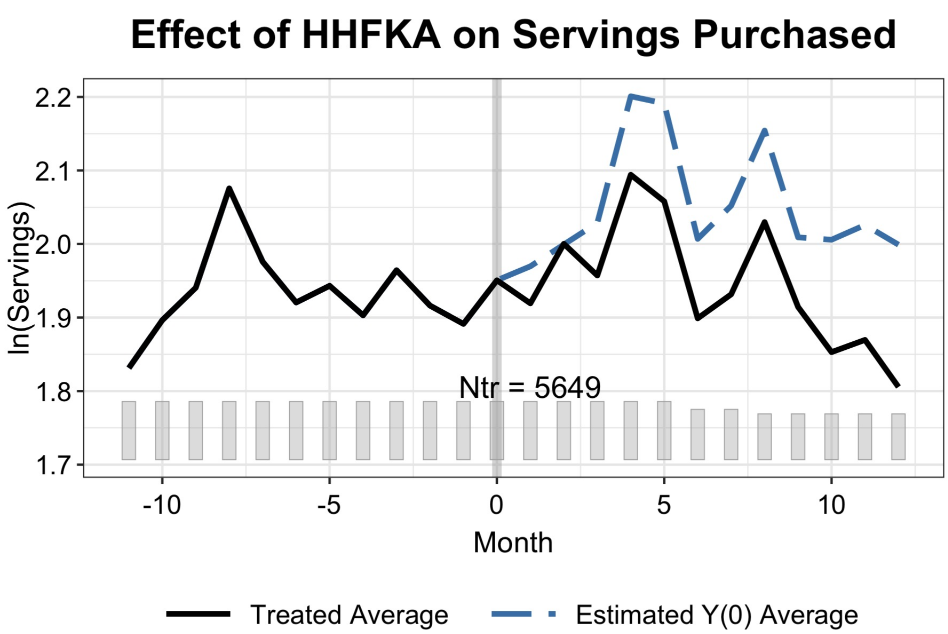

Figure 2.5: Effects of HHFKA on Food Servings Purchased

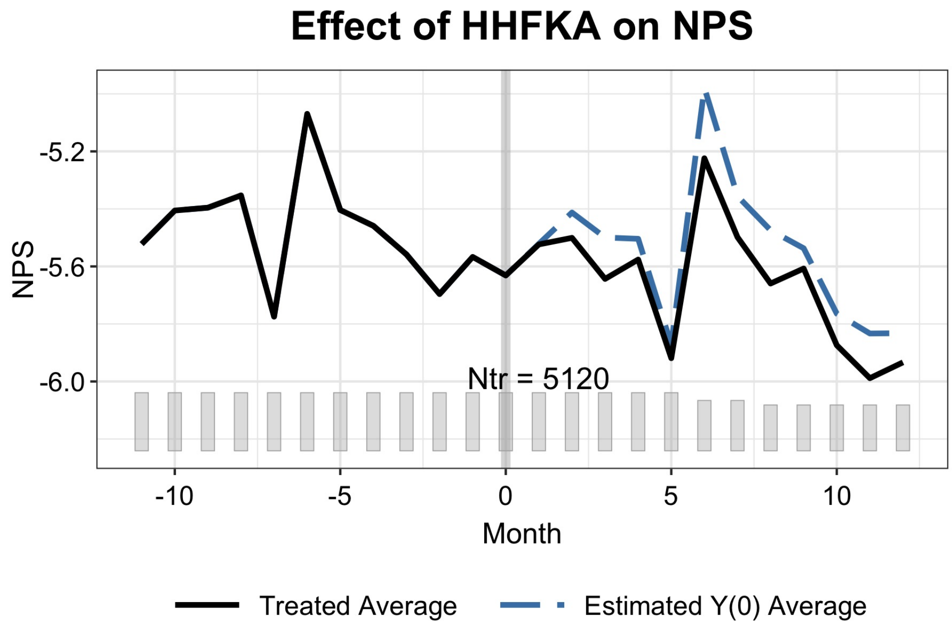

Figure 2.6: Effects of HHFKA on NPS of Food Purchased

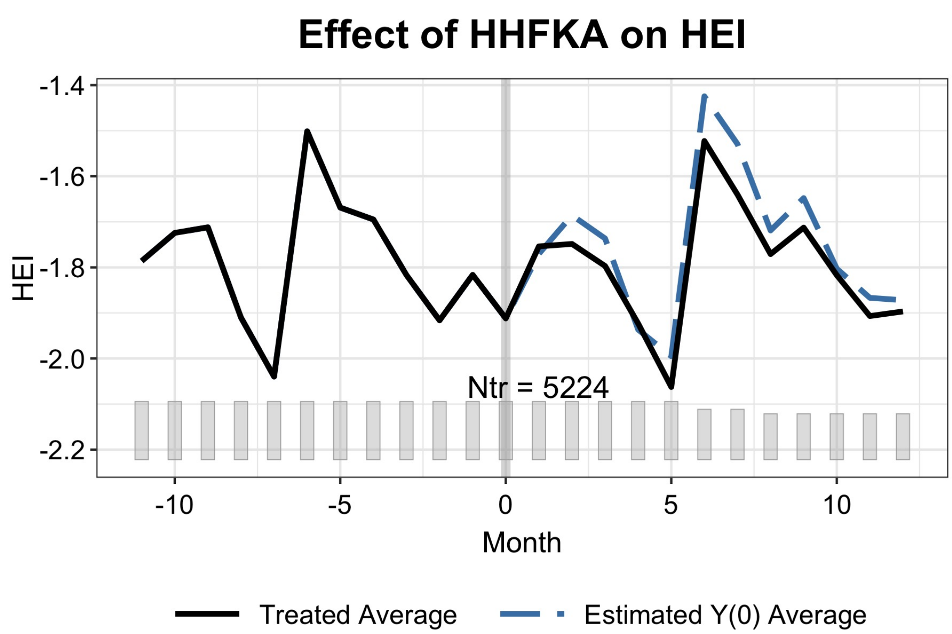

Figure 2.7: Effects of HHFKA on HEI of Food Purchased

2.6 Concluding Remarks

As an important part of the Healthy, Hunger-free Kids Act of 2010, new nutritional standards have brought unprecedented changes to school meal programs. For the first time, food sold at schools, including meals, entrees, à la carte items, etc. was mandated a ceiling on total calories and total sodium per meal/item. A systematic series of restrictions (minimums and maximums) were set for all food sold in school cafeterias. In this article, we try to answer three research questions using a Difference-in-Difference design. (1) Do parents respond by substituting towards the high-quality school meals under the HHFKA? (2) If yes, how does the HHFKA affect both the quantity and quality of household food purchases? (3) Which households change their shopping baskets after the treatment?

We compare the food purchases of households with school-age kids, who are more influenced by the HHFKA, to that of households with no school-age children, while ensuring that the two groups are as similar to each other as possible via Coarsened Exact Matching. Our results show that parents do take advantage of the improved school meals, and reduce calorie purchases as well as the number of servings at home. In specific, they purchase less food related to breakfast and lunch meals, including soda, breakfast food, prepared food and cooking ingredients from the departments of dairy, dry grocery and frozen foods. We further label food UPCs as "kids-friendly" or not, and found out that the changes mainly took place with regard to children food contents. Such changes indicate that the HHFKA is stimulating positive participating responses among parents. However, the quality of food purchases at home remains the same level.

Among all households, we find that it is households who eat less and unhealthily that drive the response to the HHFKA. They drop calorie purchases by 18.7% and servings by 11.8%. We examine which household characteristics are predictive of the membership of this quadrant, revealing the mechanism that drives their membership and behavioral changes. We find that households with time pressure — those with large household sizes, single-parent households, and households with two working parents are more likely to be in this quadrant. Moreover, the education level of the female head indicates that a lack of food literacy in households is also predictive of them being in this group. As they have limited time and knowledge to prepare healthy diets at home, the quality-improved school meals became an attractive and trusty provider of healthy foods, which they could reliably outsource to.

We also discover some unintended consequences of the HHFKA. For people who have unhealthy food consumption habits, the overall healthiness deteriorates rather than enhances after the HHFKA. They may reply on schools to provide more healthy meals while do not apply the healthy eating concepts to household foods. It is more of a substitution than a positive spillover. Evidence from the NPS shows that the health level of snacks also declines. It indicates that new standards on school meals may trigger the self-licensing effect. Kids or parents think that the "healthy eating" task has been accomplished by the more nutritious meals at school, so they could slack back and have less healthy food, although they do reduce the calorie purchase. Hopefully, the size of deterioration is not large and may be dwarfed by the healthier meals at school.

The self-licensing effect of HHFKA on household food purchases is slightly alarming. From signing the bill into law to allocating budgets to help schools improve food nutrition, USDA has invested hundreds of millions of dollars into this national project. However, if household food consumption gets deteriorated as a result, then the efforts to combat childhood obesity are still a distant hope rather than a near future. It is thus very important to take household behaviors into the big picture. For example, policy makers could promote nutritional education to society as a whole, or cooperate with sin tax policies to culture a better societal environment for children and adolescent eating habits. The consistent effect of calorie reduction at home does denote a great achievement of HHFKA. The quality itself, when the quantity of food is controlled, could inspire parents to make more use of school meal programs.

Our research also provides a natural experiment for marketing practices. While previous research has examined healthy eating diets primarily with respect to informational nudges, increase access, taxes, etc., we study whether changes in the nutritional composition of an option affect purchases of other alternatives. We are the first to show that changing the nutritional content of food options can, by itself, be attractive to consumers, and encourage substitution.

There are certain limitations of our research. The most prominent one is that we do not have access to school meal participation of households in our panel. Our estimates are based on the assumption that households with school-age kids as a whole, despite that some of them go to private schools and some do not participate in school meal programs, have a higher probability to be affected by the new nutritional standards brought by the HHFKA. The estimation accuracy would be further improved with individual school enrollment and meal participation data. Another interesting behavior worth further investigation is the potential self-licensing effect associated with the HHFKA. We speculate it as one potential mechanism behind our findings, but it would be valuable to investigate the effect of the HHFKA in a lab environment.

Appendices

APPENDIX A

DERIVATION OF THE POSTERIOR DISTRIBUTION (SINGLE QUERY)

Based on Equation (Equation 1.16), we further derive the expression of $F_{\hat{p}_i^n}(p^n_{max})$ by letting $\bar{s} = \frac{\alpha \bar{q}}{\bar{p}^2}$.

$$\begin{aligned}
F_{\hat{p}_i^n}(p^n_{max}) &= \mathrm{Prob}_{\hat{p}_i^n}(p < p^n_{max}) \\
&= \mathrm{Prob}_{\hat{p}_i^n}[(\frac{\gamma^n \tan\theta_i}{\bar{s}})^{\frac{1}{\alpha-1}} < p^n_{max}] \\
&= \mathrm{Prob}_{\hat{p}_i^n}[\frac{\gamma^n \tan\theta_i}{\bar{s}} < (p^n_{max})^{\alpha-1}] \\
&= \mathrm{Prob}_{\gamma^n}(\gamma^n < \frac{\bar{s}(p^n_{max})^{\alpha-1}}{\tan\theta_i}).
\end{aligned} \tag{A.1}$$

Since $\gamma^n \sim \mathcal{N}(\mu, \sigma^2)$, we then have

$$\begin{aligned}
\mathrm{Prob}_{\gamma^n}(\gamma^n < \frac{\bar{s}(p^n_{max})^{\alpha-1}}{\tan\theta}) &= \int_{-\infty}^{\frac{\bar{s}(p^n_{max})^{\alpha-1}}{\tan\theta_i}} \phi_{\gamma^n}(x)dx \\
&= \int_{-\infty}^{\bar{s}(p^n_{max})^{\alpha-1}} \phi_{\gamma^n}(x'\cot\theta_i) d\frac{x'}{\tan\theta_i} \\
&= \cot\theta_i \int_{-\infty}^{\bar{s}(p^n_{max})^{\alpha-1}} \phi_{\gamma^n}(x'\cot\theta_i)dx'.
\end{aligned} \tag{A.2}$$

The last but one step of Equation (Equation A.2) comes from letting $x' = x\tan\theta_i$.

Let $G_X(p, \theta) = \cot\theta \int_{-\infty}^{\bar{s}p^{\alpha-1}} \phi_X(x'\cot\theta)dx'$, then

$$\mathrm{Prob}(\mathcal{U}_i^n \theta_i) = 1 - F_{\hat{p}_i^n}(p^n_{max}) = 1 - G_{\gamma^n}(p^n_{max}, \theta_i) \tag{A.3}$$

$$\text{Prob}(\mathcal{U}_i^n) = \int \text{Prob}(\mathcal{U}_i^n \theta_i)\text{Prob}(\theta_i)d\theta_i = 1 - \int_{-\infty}^{\frac{\pi}{2}} G_{\gamma^n}(p_{max}^n, \theta)\text{Prob}(\theta_i)d\theta_i. \quad \text{(A.4)}$$

As a result, we arrive at the posterior distribution of θ_i:

$$\begin{aligned} \text{Prob}(\theta_i \mathcal{U}_i^n) &= \frac{\text{Prob}(\mathcal{U}_i^n \theta_i)\text{Prob}(\theta_i)}{\text{Prob}(\mathcal{U}_i^n)} \\ &= \frac{\text{Prob}(\theta_i)\big(1 - G_{\gamma^n}(p_{max}^n, \theta_i)\big)}{1 - \int_{-\infty}^{\frac{\pi}{2}} G_{\gamma^n}(p_{max}^n, \theta)\text{Prob}(\theta_i)d\theta_i}. \end{aligned} \quad \text{(A.5)}$$

If the consumer decides to search down, in which case $p_{max}^{n+1} \leq p_{max}^n$ and $p_{min}^{n+1} < p_{min}^n$, similarly we have,

$$\text{Prob}(\mathcal{D}_i^n \theta_i) = F_{\hat{p}_i^n}(p_{min}^n) = G_{\gamma^n}(p_{min}^n, \theta_i) \quad \text{(A.6)}$$

$$\text{Prob}(\mathcal{D}_i^n) = \int_{-\infty}^{\frac{\pi}{2}} G_{\gamma^n}(p_{min}^n, \theta_i)\text{Prob}(\theta_i)d\theta_i, \quad \text{(A.7)}$$

which leads to

$$\begin{aligned} \text{Prob}(\theta_i \mathcal{D}_i^n) &= \frac{\text{Prob}(\mathcal{D}_i^n \theta)\text{Prob}(\theta_i)}{\text{Prob}(\mathcal{D}_i^n)} \\ &= \frac{\text{Prob}(\theta_i)G_{\gamma^n}(p_{min}^n, \theta_i)}{\int_{-\infty}^{\frac{\pi}{2}} G_{\gamma^n}(p_{min}^n, \theta_i)\text{Prob}(\theta_i)d\theta_i} \end{aligned} \quad \text{(A.8)}$$

APPENDIX B

ROBUSTNESS CHECKS

Theoretically, α could be any value as long as it satisfies the convexity requirement, which is $\alpha > 1$. The analytical solution also facilitates easy functionalization in programming, since the value could be directly passed to the model as a parameter. However, whether a certain α value accords with a practical case shall be tested and determined by the specific shopping environment and industry.

Specifically, in addition to ensuring convexity of the utility function, α adds one degree of freedom in shaping the indifference curve. α enters the marginal rate of substitution of quality for price, with larger α decreasing the substitution ratio and smaller α otherwise. In both simulation and real data model testing, we adopt $\alpha = 2$, in consideration of both simplicity and performance.

Figure Figure B.1 shows the prediction performance with different α values.

- $\alpha = 1.5$ shows great prediction improvement given only three queries and performs better than $\alpha = 2$ until the seventh query. Thus, $\alpha = 1.5$ or around could also be a good choice if our goal is to interfere at an earlier stage, or if consumers do not apply price filters as frequently as expected in the search process.
- Setting α to be 2.1 leads to prediction results very close to 2, which is within expectation. This also demonstrates that α could be a fractional number from a continuous number line, which extends more flexibility to the model.

- $\alpha = 2.5$ turns out to be less powerful and less stable in performance.

- When $\alpha = 3$, the prediction performance of the model still improves in general, with more consumer search queries accumulated. The prediction accuracy ratio approaches towards 2 given only two queries, which worthies highlighting. However, the model becomes less stable and robust in the sense that it does not necessarily increase prediction power with more information. For example, prediction performance after four queries is even worse than with three queries.

- When $\alpha = 4$, the performance gets even less powerful and more unstable than when $\alpha = 3$. For most of the queries, the prediction ratio stays under 2.0. Very occasionally, it could rise to around 3, for instance, after two queries. However, the overall prediction performance is unsatisfactory. This is mainly because as the value of α gets too large, the emphasize on price gets exaggerated. Information about the quality and the market level price-quality relation could hardly shake away the gigantic power of price in the utility function. This could also be verified in the case when $\alpha = 5$ – for most of sessions, the prediction accuracy will be either zero or one. Consequently, the prediction ratio gets beyond extreme, which could hardly be presented through a plot.

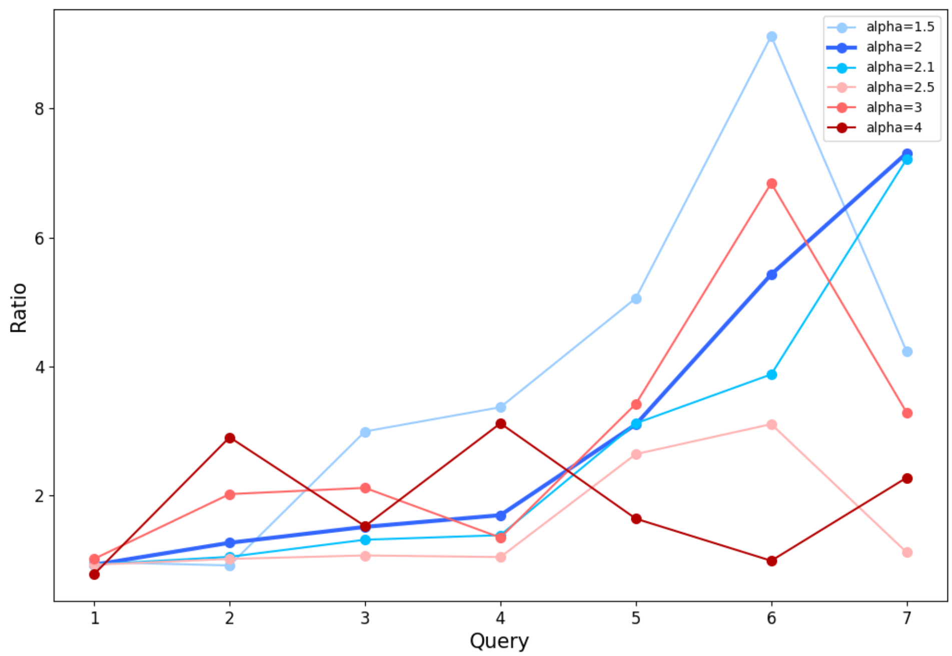

Figure B.1: Overall Prediction Performance in Real Data When α Varies

In summary, the robustness check of α helps provide a safe range in which the performance is satisfactory. In a shopping scenario similar to ours, $1 < \alpha < 2.5$ should be examined while testing the waters. In our practical case, $\alpha = 2$ stands out in both prediction precision and performance stability overall.

REFERENCES

Adomavicius, G., Huang, Z., and Tuzhilin, A. (2008). Personalization and Recommender Systems. In *State-of-the-Art Decision-Making Tools in the Information-Intensive Age*, INFORMS TutORials in Operations Research, pages 55–107.

Allcott, H., Diamond, R., Dubé, J.-P., Handbury, J., Rahkovsky, I., and Schnell, M. (2019a). Food deserts and the causes of nutritional inequality. *The Quarterly Journal of Economics*, 134(4):1793–1844. Publisher: Oxford University Press.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019b). Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, 134(3):1557–1626. Publisher: Oxford University Press.

Althoff, T., Nilforoshan, H., Hua, J., and Leskovec, J. (2022). Large-scale diet tracking data reveal disparate associations between food environment and diet. *Nature communications*, 13(1):1–12. Publisher: Nature Publishing Group.

Anzman-Frasca, S., Braun, A. C., Ehrenberg, S., Epstein, L. H., Gampp, A., Leone, L. A., Singh, A., and Tauriello, S. (2018). Effects of a randomized intervention promoting healthy children's meals on children's ordering and dietary intake in a quick-service restaurant. *Physiology & behavior*, 192:109–117. Publisher: Elsevier.

Bakker, A. B. and Demerouti, E. (2013). The spillover-crossover model. *New frontiers in work and family research*, pages 54–70.

Besanko, D., Dranove, D., Shanley, M., and Schaefer, S. (2009). *Economics of Strategy*. John Wiley & Sons.

Bollinger, B. and Sexton, S. (2018). Local excise taxes, sticky prices, and spillovers: evidence from Berkeley's soda tax. *Sticky Prices, and Spillovers: Evidence from Berkeley's Soda Tax (January 12, 2018).*

Bronnenberg, B. J., Kim, J. B., and Mela, C. F. (2016). Zooming in on choice: How do consumers search for cameras online? *Marketing science*, 35(5):693–712.

Cao, X. and Zhu, Y. (2019). Group Search Strategy. *NYU Stern School of Business*. Available at SSRN: https://ssrn.com/abstract=3437993.

Carmon, Z. and Simonson, I. (1998). Price–quality trade-offs in choice versus matching: New insights into the prominence effect. *Journal of Consumer Psychology*, 7(4):323–343. Publisher: Wiley Online Library.

Chen, Y. and Yao, S. (2017). Sequential search with refinement: Model and application with click-stream data. *Management Science*, 63(12):4345–4365.

Chevalier, S. (2021). Number of pages seen at buying session worldwide in 2019, by vertical. https://www.statista.com/statistics/1106617/number-of-pages-seen-buying-session/.

Cohen, J. F., Richardson, S., Parker, E., Catalano, P. J., and Rimm, E. B. (2014). Impact of the new US Department of Agriculture school meal standards on food selection, consumption, and waste. *American journal of preventive medicine*, 46(4):388–394. Publisher: Elsevier.

Cornell, C. E., Rodin, J., and Weingarten, H. (1989). Stimulus-induced eating when satiated. *Physiology & behavior*, 45(4):695–704. Publisher: Elsevier.

Danaher, P. J., Smith, M. S., Ranasinghe, K., and Danaher, T. S. (2015). Where, When, and how Long: Factors that Influence the Redemption of Mobile Phone

Coupons. *Journal of Marketing Research*, 52(5):710–725. Publisher: SAGE Publications Inc.

De los Santos, B. and Koulayev, S. (2017). Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science*, 36(4):542–564.

de Mestral, C., Khalatbari-Soltani, S., Stringhini, S., and Marques-Vidal, P. (2017). Fifteen-year trends in the prevalence of barriers to healthy eating in a high-income country. *The American journal of clinical nutrition*, 105(3):660–668. Publisher: Oxford University Press.

Deaton, A. (2008). *Essays in the Theory and Measurement of Consumer Behaviour: In Honour of Sir Richard Stone*. Cambridge University Press. Publication Title: Cambridge Books.

Dhar, T. and Baylis, K. (2011). Fast-food consumption and the ban on advertising targeting children: the Quebec experience. *Journal of Marketing research*, 48(5):799–813. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, 3(2):156–168. Publisher: Academic Press.

Ding, M., Ross Jr, W. T., and Rao, V. R. (2010). Price as an indicator of quality: implications for utility and demand functions. *Journal of Retailing*, 86(1):69–84. Publisher: Elsevier.

DiSantis, K. I., Birch, L. L., Davey, A., Serrano, E. L., Zhang, J., Bruton, Y., and Fisher, J. O. (2013). Plate size and children's appetite: effects of larger dishware on self-served portions and intake. *Pediatrics*, 131(5):e1451–e1458. Publisher: American Academy of Pediatrics Elk Grove Village, IL, USA.

Dolan, P. and Galizzi, M. M. (2015). Like ripples on a pond: behavioral spillovers and their implications for research and policy. *Journal of Economic Psychology*, 47:1–16.

Dong, X., Morozov, I., Seiler, S., and Hou, L. (2018). Estimation of Preference Heterogeneity in Markets with Costly Search. Stanford University Graduate School of Business Research Paper No. 18-34, Available at SSRN: https://ssrn.com/abstract=3202534 or http://dx.doi.org/10.2139/ssrn.3202534.

Downs, J. S., Loewenstein, G., and Wisdom, J. (2009). Strategies for promoting healthier food choices. *American Economic Review*, 99(2):159–64.

Drewnowski, A. (1997). Taste preferences and food intake. *Annual review of nutrition*, 17:237. Publisher: Annual Reviews, Inc.

Dubois, P., Albuquerque, P., Allais, O., Bonnet, C., Bertail, P., Combris, P., Lahlou, S., Rigal, N., Ruffieux, B., and Chandon, P. (2021). Effects of front-of-pack labels on the nutritional quality of supermarket food purchases: evidence from a large-scale randomized controlled trial. *Journal of the Academy of Marketing Science*, 49(1):119–138. Publisher: Springer.

Dubois, P., Griffith, R., and O'Connell, M. (2020). How well targeted are soda taxes? *American Economic Review*, 110(11):3661–3704.

Ensaff, H., Homer, M., Sahota, P., Braybrook, D., Coan, S., and McLeod, H. (2015). Food choice architecture: An intervention in a secondary school and its impact on students' plant-based food choices. *Nutrients*, 7(6):4426–4437. Publisher: MDPI.

Epstein, L. H., Handley, E. A., Dearing, K. K., Cho, D. D., Roemmich, J. N., Paluch, R. A., Raja, S., Pak, Y., and Spring, B. (2006). Purchases of food in youth: influence of price and income. *Psychological Science*, 17(1):82–89. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Farris, A. R., Misyak, S., Duffey, K. J., Atzaba-Poria, N., Hosig, K., Davis, G. C., and Serrano, E. (2016). Elementary parent perceptions of packing lunches and the National School Lunch Program. *J Child Nutr Manag*, 40(1):1–10.

Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70.

Gomez-Uribe, C. A. and Hunt, N. (2016). The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Transactions on Management Information Systems*, 6(4):13:1–13:19.

Goodman, J. and Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80. Publisher: Mathematical Sciences Publishers.

Greminger, R. P. (2022). Optimal search and discovery. *Management Science*, 68(5):3904–3924. Publisher: INFORMS.

Hagerty, M. R. (1978). Model testing techniques and price-quality tradeoffs. *Journal of Consumer Research*, 5(3):194–205. Publisher: The University of Chicago Press.

Hanks, A. S., Just, D. R., Smith, L. E., and Wansink, B. (2012). Healthy convenience: nudging students toward healthier choices in the lunchroom. *Journal of public health*, 34(3):370–376. Publisher: Oxford University Press.

Hazlett, C. and Xu, Y. (2018). Trajectory balancing: A general reweighting approach to causal inference with time-series cross-sectional data. *Available at SSRN 3214231*.

Hidasi, B., Karatzoglou, A., Baltrunas, L., and Tikk, D. (2015). Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.

Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational behavior and human decision processes*, 67(3):247–257. Publisher: Elsevier.

Huyghe, E., Verstraeten, J., Geuens, M., and Van Kerckhove, A. (2017). Clicks as a healthy alternative to bricks: how online grocery shopping reduces vice purchases. *Journal of Marketing Research*, 54(1):61–74. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Jetter, K. M. and Cassady, D. L. (2006). The availability and cost of healthier food alternatives. *American journal of preventive medicine*, 30(1):38–44. Publisher: Elsevier.

Kearney, J. M. and McElhone, S. (1999). Perceived barriers in trying to eat healthier–results of a pan-EU consumer attitudinal survey. *British Journal of Nutrition*, 81(S1):S133–S137. Publisher: Cambridge University Press.

Keeney, R. L., Raiffa, H., and Meyer, R. F. (1993). *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.

Kenney, E. L., Barrett, J. L., Bleich, S. N., Ward, Z. J., Cradock, A. L., and Gortmaker, S. L. (2020). Impact Of The Healthy, Hunger-Free Kids Act On Obesity Trends: Study examines impact of the Healthy, Hunger-Free Kids Act of 2010 on childhood obesity trends. *Health Affairs*, 39(7):1122–1129.

Khan, U. and Dhar, R. (2006). Licensing effect in consumer choice. *Journal of marketing research*, 43(2):259–266. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Kim, J. B., Albuquerque, P., and Bronnenberg, B. J. (2010). Online demand under limited consumer search. *Marketing science*, 29(6):1001–1023.

Kinderknecht, K., Harris, C., and Jones-Smith, J. (2020). Association of the healthy, hunger-free kids act with dietary quality among children in the US national school lunch program. *Jama*, 324(4):359–368. Publisher: American Medical Association.

Lee, S. and Rivera, A. (2021). Extrapolation Bias and Robust Dynamic Liquidity Management. *Management Science*.

Levin, I. P. and Johnson, R. D. (1984). Estimating price-quality tradeoffs using comparative judgments. *Journal of consumer research*, 11(1):593–600. Publisher: The University of Chicago Press.

Levin, S. (1996). Pilot study of a cafeteria program relying primarily on symbols to promote healthy choices. *Journal of Nutrition Education*, 28(5):282–285. Publisher: Elsevier.

Lewbel, A. (1998). An Exactly aggregable trigonometric engel curve demand system. *Econometric Reviews*, 7(1):97–102. Publisher: Taylor & Francis.

Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., and Ma, J. (2017). Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1419–1428.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic theory*, volume 1. Oxford university press New York.

Matsuda, T. (2006). A trigonometric flexible consumer demand system. *Canadian Journal of Economics/Revue canadienne d'économique*, 39(1):145–162.

Mayer, J. A., Heins, J. M., Vogel, J. M., Morrison, D. C., Lankester, L. D., and Jacobs, A. L. (1986). Promoting low-fat entree choices in a public cafeteria. *Journal of Applied Behavior Analysis*, 19(4):397–402. Publisher: Wiley Online Library.

McGee, P. (2021). Snap, Facebook, Twitter and YouTube lose nearly $10bn after iPhone privacy changes. *Financial Times*. https://www.ft.com/content/4c19e387-ee1a-41d8-8dd2-bc6c302ee58e.

Miller, G. F., Gupta, S., Kropp, J. D., Grogan, K. A., and Mathews, A. (2016). The effects of pre-ordering and behavioral nudges on National School Lunch Program participants' food item selection. *Journal of economic psychology*, 55:4–16. Publisher: Elsevier.

Miller, N., Reicks, M., Redden, J. P., Mann, T., Mykerezi, E., and Vickers, Z. (2015). Increasing portion sizes of fruits and vegetables in an elementary school lunch program can increase fruit and vegetable consumption. *Appetite*, 91:426–430. Publisher: Elsevier.

Moorman, C. (1996). A quasi experiment to assess the consumer and informational determinants of nutrition information processing activities: The case of the nutrition labeling and education act. *Journal of Public Policy & Marketing*, 15(1):28–44. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Munt, A. E., Partridge, S. R., and Allman-Farinelli, M. (2017). The barriers and enablers of healthy eating among young adults: A missing piece of the obesity puzzle: A scoping review. *Obesity reviews*, 18(1):1–17. Publisher: Wiley Online Library.

Nicklas, T. A., Jahns, L., Bogle, M. L., Chester, D. N., Giovanni, M., Klurfeld, D. M., Laugero, K., Liu, Y., Lopez, S., and Tucker, K. L. (2013). Barriers and facilitators for consumer adherence to the dietary guidelines for Americans: the HEALTH study. *Journal of the Academy of Nutrition and Dietetics*, 113(10):1317–1331. Publisher: Elsevier.

Nikolova, H. D. and Inman, J. J. (2015). Healthy choice: the effect of simplified point-of-sale nutritional information on consumer food choice behavior. *Journal of Marketing Research*, 52(6):817–835. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Nowlis, S. M. and Simonson, I. (1997). Attribute–task compatibility as a determinant of consumer preference reversals. *Journal of marketing research*, 34(2):205–218. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Olstad, D. L., Vermeer, J., McCargar, L. J., Prowse, R. J., and Raine, K. D. (2015). Using traffic light labels to improve food selection in recreation and sport facility eating environments. *Appetite*, 91:329–335. Publisher: Elsevier.

Padilla, N., Ascarza, E., and Netzer, O. (2019). The customer journey as a source of information. Working paper, Columbia Business School.

Palazzolo, M. and Pattabhiramaiah, A. (2021). The Minimum Wage and Consumer Nutrition. *Journal of Marketing Research*, page 00222437211023475.

Peck, J. and Childers, T. L. (2006). If I touch it I have to have it: Individual and environmental influences on impulse purchasing. *Journal of business research*, 59(6):765–769. Publisher: Elsevier.

Poon, T., Labonté, M.-, Mulligan, C., Ahmed, M., Dickinson, K. M., and L'Abbé, M. R. (2018). Comparison of nutrient profiling models for assessing the nutritional quality of foods: A validation study. *British Journal of Nutrition*, 120(5):567–582. Publisher: Cambridge University Press.

Reicks, M., Redden, J. P., Mann, T., Mykerezi, E., and Vickers, Z. (2012). Photographs in lunch tray compartments and vegetable consumption among children in elementary school cafeterias. *Jama*, 307(8):784–785. Publisher: American Medical Association.

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, J. (1994). Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186.

Resnick, P. and Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40(3):56–58. Publisher: ACM New York, NY, USA.

Rogers, P. J. and Hill, A. J. (1989). Breakdown of dietary restraint following mere exposure to food stimuli: interrelationships between restraint, hunger, salivation, and food intake. *Addictive behaviors*, 14(4):387–397. Publisher: Elsevier.

Rosenfield, D. B. and Shapiro, R. D. (1981). Optimal adaptive price search. *Journal of Economic Theory*, 25(1):1–20.

Rossi, P. E., McCulloch, R. E., and Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4):321–340.

Rothschild, M. (1974). Searching for the lowest price when the distribution of prices Is unknown: a summary. In *Annals of Economic and Social Measurement, Volume 3, number 1*, pages 293–294. NBER.

Sachdeva, S., Iliev, R., and Medin, D. L. (2009). Sinning saints and saintly sinners: The paradox of moral self-regulation. *Psychological science*, 20(4):523–528.

Schwartz, M. B., Henderson, K. E., Read, M., Danna, N., and Ickovics, J. R. (2015). New school meal regulations increase fruit consumption and do not increase total plate waste. *Childhood Obesity*, 11(3):242–247. Publisher: Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA.

Shaddy, F., Fishbach, A., and Simonson, I. (2021). Trade-offs in choice. *Annual Review of Psychology*, 72:181–206. Publisher: Annual Reviews.

Shi, Y., Larson, M., and Hanjalic, A. (2014). Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)*, 47(1):1–45. Publisher: ACM New York, NY, USA.

Simonson, I. and Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of marketing research*, 29(3):281–295. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Soulier, S. (2018). Air Travel More Accessible in 2017, According to Latest Air Travelers in America Report. https://www.airlines.org/news/air-travel-more-accessible-in-2017-according-to-latest-air-travelers-in-america-report/.

Stigler, G. J. (1961). The economics of information. *Journal of political economy*, 69(3):213–225. Publisher: The University of Chicago Press.

Thorndike, A. N., Sonnenberg, L., Riis, J., Barraclough, S., and Levy, D. E. (2012). A 2-phase labeling and choice architecture intervention to improve healthy food and beverage choices. *American journal of public health*, 102(3):527–533. Publisher: American Public Health Association.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological review*, 79(4):281. Publisher: American Psychological Association.

Ursu, R. M., Wang, Q., and Chintagunta, P. K. (2020). Search duration. *Marketing Science*, 39(5):849–871. Publisher: INFORMS.

Varian, H. R. (2014). *Intermediate microeconomics: a modern approach: ninth international student edition*. WW Norton & Company.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654. Publisher: JSTOR.

Westman, M. (2013). Crossover of positive states and experiences. Issue: 4 Pages: 263–265 Publication Title: Stress and Health Volume: 29.

Wijayaratne, S. P., Reid, M., Westberg, K., Worsley, A., and Mavondo, F. (2018). Food literacy, healthy eating barriers and household diet. *European Journal of Marketing*, 52(12):2449–2477. Publisher: Emerald Publishing Limited.

Zhang, Q., Lu, J., and Jin, Y. (2021). Artificial intelligence in recommender systems. *Complex & Intelligent Systems*, 7(1):439–457. Publisher: Springer.

Zou, L., Xia, L., Ding, Z., Song, J., Liu, W., and Yin, D. (2019). Reinforcement Learning to Optimize Long-term User Engagement in Recommender Systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 2810–2818, New York, NY, USA. Association for Computing Machinery.