Essays on Quantitative Marketing and Empirical Industrial Organization

by

Rudolf-Harri Oberg

Department of Economics
Duke University

Date: ______________________________

Approved:

____
Carl Mela, Supervisor

____
Yi (Daniel) Xu, Chair

____
Peter Arcidiacono

____
Bryan Bollinger

____
Andres Musalem

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University
2020

ABSTRACT

Essays on Quantitative Marketing and Empirical Industrial Organization

by

Rudolf-Harri Oberg

Department of Economics
Duke University

Date: ______________________

Approved:

____
Carl Mela, Supervisor

____
Yi (Daniel) Xu, Chair

____
Peter Arcidiacono

____
Bryan Bollinger

____
Andres Musalem

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2020

Abstract

This dissertation presents two essays at the intersection of quantitative marketing and empirical industrial organization. The first essay, based on a joint research project with Andres Musalem, studies how consumers respond to subsidy designs with different spending restrictions in the context of SNAP (commonly known as the Food Stamps Program), with implications to the optimal design of consumer subsidies. In the second essay, I empirically examine how supply-side heterogeneity on a crowd-based platform moderates policy outcomes of platform initiatives.

To Kristi and Albert.

Contents

List of Tables

List of Figures

Acknowledgements

I am very grateful to my committee members Carl Mela, Daniel Xu, Andres Musalem, Bryan Bollinger, and Peter Arcidiacono, for the help and advice they have shared with me over the years. I want to thank Carl in particular for supervising my doctoral studies, being patient and reassuring when it took me time to develop research ideas, providing moral and financial support on this journey, helping overcome various obstacles, and having shared his advice and experience when I prepared for the job market. I also want to thank Andres Musalem, whose feedback and collaboration have been valuable for developing the work in the first essay. Finally, I want to thank the Department of Economics, the Fuqua School of Business, and the Duke Graduate School for providing me with funding during my time at Duke.

Chapter 1

Introduction

This dissertation presents two essays at the intersection of quantitative marketing and empirical industrial organization. The first essay studies how consumers respond to subsidy designs with different spending restrictions, with implications to the optimal design of consumer subsidies. In the context of SNAP (commonly known as the Food Stamps Program), we analyze the effect of food stamp benefits from the perspective of recipients (who care about overall consumer welfare) and the policymaker (who prefers that funds are used to buy food). Studying the effect of different subsidies in the context of SNAP is interesting since there are active proposals to change food stamp benefits in opposite ways: (i) to lift restrictions and provide cash with no strings attached, (ii) to add restrictions and exclude certain items (e.g., soda). To simulate consumer behavior under different subsidies, we develop a structural model of consumer demand which integrates consumer decisions for brands, categories, and stores. Our main finding is that expanding food stamp benefits to include household goods would be preferred by both benefit recipients and the policymaker. The mechanism driving this result is that flexible benefits give access to a wider selection of items, which provides greater incentives to visit stores. In addition, we quantify trade-offs between different benefit designs and study the effect of banning benefit use on sweetened soda. Finally, our model of brand, category and store choice makes a technical contribution that could be interesting beyond the application considered in this paper.

In the second essay, I empirically examine how supply-side heterogeneity on a crowd-based platform moderates ad-policy outcomes. I study this question in the setting of a large live-streaming platform, where some content producers work full-time, and others produce content as a hobby. The main distinction between professional and amateur content providers is the ability and motivation to earn revenue. I examine how content creator response affects the impact of a policy change in site-wide monetization.

Preliminary evidence suggests that amateur content creators respond to the policy by creating more content and changing their streaming schedules to include fewer but longer streams. Amateur behavior does not seem to be associated with changes in their performance or direct effects of the ad policy. Instead, amateurs may change behavior because the ad policy benefits professionals both directly and indirectly, raising the prospects of becoming a professional streamer. Since many platforms rely heavily on professional market participants to provide services, understanding "professionalization" of the supply side could yield valuable insight into platform decision making.

Chapter 2

Should Consumer Subsidies Be More Flexible? A Structural Analysis and Case Study of the Food Stamps Program

The contents of this chapter are based on a joint research project with Andres Musalem. I was in charge of writing the manuscript, estimating the model, performing counterfactual experiments and analyzing results; Andres was in charge of preparing and cleaning the data and advising other parts of the project.

2.1 Introduction

The Supplemental Nutrition Assistance Program (SNAP) – formerly and more familiarly known as the Food Stamps Program – provides monthly benefits for food spending. SNAP is the second-largest means-tested[1] program in the United States after Medicaid, providing benefits to 40 million Americans in the fiscal year of 2018 (USDA (2018)). The program's focus on subsidizing food expenditure, as opposed to supporting the purchase of household items or giving cash, reflects its goal to offer nutrition assistance to low-income individuals and families. At the same time, there is a growing debate over whether SNAP and other "in-kind"[2] style social benefit programs should be more flexible, with some proposing

[1]Means-tested benefits are available to individuals who can demonstrate that their means (income, savings, and other capital) are below a certain threshold.

[2]In-kind transfer programs give benefits in the form of goods and services (e.g., SNAP, Medicare). This is opposed to in-cash transfer programs which give benefits in the form of a cash transfer (e.g., U.S. Social Security and unemployment benefits).

replacing in-kind benefits with simple cash benefits.[3] On the one hand, economic theory implies that giving cash over restrictive subsidies is always more effective from a welfare perspective since recipients are free to pick the bundle that maximizes their utility. On the other hand, the policymaker may prefer that recipients only use benefits on certain kinds of goods and as a result sets in place place different rules and restrictions on benefit spend in an attempt to achieve program goals. This difference in the perspective of recipients and the policymaker is important for evaluating SNAP and other in-kind transfer programs, as proposals to change program implementation may be derived from a singular perspective. As a result, there are proposals to change the current SNAP implementation in opposing ways, with some proposals calling for lifting restrictions on food stamp benefits, and other proposals calling for further restrictions on food stamp benefits by excluding certain items (e.g., soda).

In this paper, we[4] analyze the effects of food stamp benefits relative to alternative subsidy designs from the perspective of the policymaker and benefit recipients. We ask how key outcomes like consumer welfare or marginal propensity to consume food items (MPCF)[5] differ under (i) the usual food stamp restrictions (food only), (ii) grocery restrictions (food and household items only), and (iii) tax rebate (cash transfer). In order to simulate consumer behavior under different benefit systems, we develop a structural model of consumer demand which integrates consumer decisions for brands, categories, and stores. Using a structural model to account for direct and indirect effects of consumer behavior

[3] For articles in the popular press see Kenny (2015), Reinhardt (2011), Cannon (2016).

[4] This chapter is based on joint research project with Andres Musalem.

[5] MPCF is a key indicator which reflects the main goal of the SNAP program and is also the main object of interest in many research papers. Therefore, we identify preferences of the policymaker to be in line with shifting benefit dollars toward food expenditure, i.e., maximizing MPCF out of provided benefits.

is crucial for understanding the true effect of different subsidies. For example, restricting consumer subsidies to food only, as opposed to including household items, may raise food expenditures conditional on store visit but may reduce the overall prospect of visiting a store, leaving it an open question which policy shifts more dollars towards food expenditure.

In addition to studying trade-offs between different subsidies, we also consider how excluding items from benefit use in one category affects purchases in other categories. The argument that some food categories – especially the category of sweetened beverages which includes sweetened soda – should be made ineligible for purchase with SNAP benefits has been gaining momentum.[6] The idea to restrict spending on soda comes from the fact that more money is spent on soft drinks than any other item (USDA (2016)). Those who propose banning soda from SNAP argue that the ban could change purchase patterns and reduce soda consumption, whereas those who oppose the ban predict no change in consumption patterns since SNAP recipients could just buy the same amount of soda with their own cash. In addition, shopping frequency may change when subsidy use is restricted and it is an empirical question to what extent consumers substitute to other categories within a store and to what extent they substitute to the outside good (and decrease shopping frequency). We use our demand model to simulate a policy in which SNAP-eligible households face a restricted version of SNAP where spending on soda is banned. We establish that banning soda from SNAP affects the budget line in the same way as mandating a tax on soda which prompts us to compute the extent of soda taxation needed to reduce soda consumption to the level recorded under the ban.

We use household panel data from the IRI marketing data set (Bronnenberg et al. (2008)) to estimate our structural model. The data contain information on weekly store

[6]See for example Lane (2017), O'Connor (2017), Bittman (2012) for articles in the popular press.

trips and purchase decisions for a panel of over four thousand consumers, making it an ideal setting to study how consumer behavior changes in response to policy measures on all three levels (store, category, and brand choice). Our data are augmented with store locations and consumers' pre-tax income and location.

We find evidence that expanding food stamp benefits to include grocery items (food and household goods) would be preferred by both benefit recipients *and* the policymaker. Our estimates imply that SNAP eligible consumers spend more money on food items using grocery benefits compared to regular food stamp benefits. Grocery benefits give access to a wider selection of items which provides greater incentives to visit stores. As a result, shopping frequency increases by 6.3% and serves as the main mechanism driving increased food spending (5.2%) relative to the regular food stamp subsidy. In addition, our simulations show that consumers would trade $100 in regular benefits for $95 in grocery benefits, which implies that an equal amount of grocery benefits would improve consumer welfare. Altogether, this finding provides a positive answer to the main question posed in the title of this paper and provides a concrete example of how benefits could be more flexible. This finding is interesting because it shows that a new design – which is not observed in practice – can improve both outcomes of interest without any compromise. Our simulations also confirm existing results about the trade-offs of regular food stamp policy and cash benefit, yielding further credibility to our main finding. Second, we find that consumers respond to a 10% soda tax and a soda ban in a similar way, reducing soda consumption by 23% and 20%, respectively. We discover that the ban policy has a differential impact on consumers' substitution patterns and find evidence that differences in response are mainly generated by differences in preferences, as opposed to differences in travel costs.

Our results contribute to the literature on the effects of SNAP benefits on food spending. While most papers in the literature use reduced-form type of analysis to estimate key parameters like the MPCF, ours is the only paper (to the best of our knowledge) that takes a structural approach. Our structural model yields estimates that are in-line with existing estimates in the literature and has the benefit of providing new (counterfactual) estimates for policies that are often part of the policy debate but are not observed in practice. Hoynes and Schanzenbach (2009), Castner and Mabli (2010), Beatty and Tuttle (2015) and Hastings and Shapiro (2018) all provide estimates for MPCF out of cash and SNAP and our estimates for these two quantities agree with what has been reported in these papers. Our results are also related to the literature on the effect of soda taxes (see e.g., Wang (2015), Bollinger and Sexton (2018) for overview). Most papers in this literature use event studies and/or structural approaches to predict the reduction in the consumption of sugar-sweetened beverages induced by such a tax. We approach the topic of soda taxation from the perspective of restricting food stamp benefits and calculate how the ban would be similar to a soda tax that is universally mandated or enacted only upon low income consumers. One caveat of our analysis is that our modeling framework does not account for the possibility that firms could reoptimize prices in response to new taxes or restricted benefits. However, recent research has demonstrated limited retailer price response to demand shocks and soda tax (Arcidiacono et al. (2018), Bollinger and Sexton (2018)). At the very least, our results provide a useful ballpark estimate of the effect of policy change.

Our modeling framework contributes to the literature on store choice and consumer demand. The main novelty of our model is integrating consumer demand for brands, categories and stores in a direct utility framework (see table 2.1). Few papers have modeled all three decisions and to our knowledge ours is the first to be based on a direct utility model of consumer behavior (Chandukala et al. (2008), Chintagunta and Nair (2011)). The primary benefit of considering all three levels of decision making in our application is that the policy of restricting or expanding food stamp benefits is often targeted at specific items and the effectiveness of the policy depends on changes in consumer demand. For example, if sugar-sweetened sodas are banned from food stamp benefits but consumers substitute to other goods that are deemed unhealthy (e.g., other sodas) then such a restriction to benefits is less effective in serving its purpose. Including store choice in the model is important for measuring how consumers substitute between the inside good (items sold at stores) and the outside good (expenditure outside of stores) in response to changes in their budget constraint. In terms of joint consideration of how much demand is allocated across goods, categories, and stores, our work is closest to Bell et al. (1999). We build on this work by using a direct utility model where a single primitive (consumer consumption utility) drives all three decisions. Our direct utility model builds on the work of Lee and Allenby (2009) by nesting their model of brand-category choice into a model of store choice. Direct utility models have recently been applied to explain store and category choice in a variety of applications, e.g., Thomassen et al. (2017), Allcott et al. (2019), and Shriver and Bollinger (2017). We build on this research by adding more dimensions (brand choice) to the model, thereby enabling a richer set of substitution patterns, as noted above. Although not considered in the current application, the addition of brands is also useful for exploring the concept of loss leader pricing which rests on the premise that a given brand or category can drive demand for all goods within the store (Walters and MacKenzie (1988)). To conclude, our model makes a technical contribution which is relevant beyond the application considered in this paper.

The rest of the paper is organized as follows. In section 2.2 we give an overview of the

Table 2.1: Store Choice Literature

	Brand	Category	Store	Direct Utility
Bell and Lattin (1998)	✓	✓	✓	
Lee and Allenby (2009)	✓	✓		✓
Kim (2012)		✓	✓	✓
Thomassen et al. (2017)		✓	✓	✓
Shriver and Bollinger (2017)		✓	✓	✓
Briesch et al. (2009)		✓	✓	
Schiraldi et al. (2012)		✓	✓	
Pakes (2010)			✓	
Figurelli (2013)			✓	
Zhu et al. (2011)			✓	
Singh et al. (2006)			✓	

data. Section 2.3 develops the model and section 2.4 discusses estimation. In section 2.5, we perform counterfactual experiments and analyze results. We summarize our findings in section 3.4.

2.2 Data

We use household panel data from the IRI marketing data set (Bronnenberg et al. (2008)). The data consist of weekly store trips of individuals who are residents of either Eau Claire, Wisconsin, or Pittsfield, Massachusetts, and who participate in IRI's BehaviorScan program. Our final data set follows the store, category and brand choice for 4002 consumers for 52 weeks. Weekly data include info on consumer spending, purchased quantities, store choice, distance to stores, product prices and characteristics. We focus on 10 most purchased categories with 3–6 most purchased brands in each category. We will now give details on how we process the data.

2.2.1 Data Set Construction

There are a number of common tasks that need to be completed to make the data usable for choice modeling. We outline our choices below. In most cases, we closely follow the methods used by Gordon et al. (2013).

The IRI data set tracks 30 product categories but we focus on the same 19 product categories that are used by Gordon et al. (2013). We rank the categories in terms of sales volume and pick the top ten categories for analysis. The remaining 9 categories are grouped into a composite category where the weight of each category is given by its sales volume. The top ten categories in terms of sales volume are carbonated beverages, frozen pizza, yogurt, frozen dinner, potato chips, toilet tissue, coffee, laundry detergent, paper towel, tortilla chips. The composite category consists of hot-dog, spaghetti sauce, margarine and butter, mayonnaise, peanut butter, shampoo, ketchup, deodorant, and mustard categories.

To make the consumer's choice problem tractable, we aggregate UPC's in each category into a set of brands. First, UPC's that serve a specific market niche, have very low sales or are otherwise irrelevant to the analysis are removed. Among the UPC's that remain, only most popular package types and sizes are used in the analysis. For example, only UPC's of liquid detergents and their most popular package sizes are included in the category of laundry detergent since they make up more than 95 percent of the category sales. The same principle of UPC selection is applied in each category which on average causes a 10 percent reduction in the number of UPC's over different categories. Second, the selected UPC's are aggregated into brands following the decision rules discussed in Gordon et al. (2013). Since private labels cannot be identified with a specific retail chain from the data, all private labels are grouped under the same brand regardless of the chain they belong to. The resulting brands are sorted by market share and those brands that yield a cumulative market share of at least 80 percent or have an individual market share of at least 4 percent are included in the analysis. The remaining brands are grouped into a single composite brand with an average market share of 18.3 percent. Next, prices are aggregated to the brand level as an weighted average of UPC prices. The weight of an UPC is equal to its share of sales volume in the brand it belongs to where the sales volume is calculated at the store level in the given year.

In selecting the sample of households, we include all households who have visited at least two stores and for whom location and income data exists. This rule selects 4210 panelists from an initial sample of 11322. We focus on each consumer's two stores with the highest expenditure shares. First, shopping outside the top two stores is minimal (see table 2.3). Second, the computational cost of accommodating additional stores are severe which is why the literature on multi-store multi-category choice has typically focused on the case of two stores.[7] We approximate each consumer's weekly budget by dividing their pre-tax income by 52. If weekly spending is greater than the approximated budget, we increase the weekly spending by 5 percent and take this to be the new budget for that week. To calculate the distance to visited stores, we use latitude-longitude information provided for each consumer's residence. Since information on consumer's exact location is sensitive, the data has been disguised so that the true location lies within 225 meters from the provided data. We calculate distance to visited stores by executing a Python script which queries Google for driving distance. To make the sample more compact, we

[7] With K stores the state-space has 2^K elements, with the same number of optimization exercises for each consumer.

get rid of outliers and extreme values. For each consumer's two most visited stores, one store is closer and the other store is further away. We drop 2.5 percent of consumer who live very close to their closer store and 2.5 percent of consumer who live very far from their other store. This rule ensures that each panelists distance to their two most visited stores is at least 0.4 miles and not more than 14 miles and brings the sample to 4002 panelists.

Finally, we use federal guidelines to simulate whether consumers in our sample are eligible for SNAP benefits. Using criteria from the 2018 fiscal year, we say that a consumer is SNAP eligible when his yearly reported income falls below the eligibility threshold.[8] We find that around 24 percent of our sample is SNAP eligible using these guidelines. Since direct SNAP participation data is notoriously hard to get, most research relies on indirect measures of participation.[9] For example, Hastings and Shapiro (2018) use method of payment to infer SNAP participation. Therefore, we believe that simulating SNAP participation in our sample using federal guidelines is sufficient for analyzing how target consumers respond to different subsidies.

2.2.2 Descriptive Statistics

Table 2.2 presents descriptive demographic and income data for the estimation sample. The high share of outside good is worthy of note as it illustrates the magnitude of economic choices relative to the reported income that we are able observe in the data. In particular, when consumers do not visit any stores in a given week, then by construction they spend

[8]The threshold is 130 percent of the federal poverty line. In federal fiscal year of 2018, the federal poverty line was $1702 for a household of three, making the eligibility threshold $2213 per month or about $27k per year. (USDA 2018)

[9]Meyer and Goerge (2010) and Meyer and Mittag (2015) show that common surveys underreport participation by 30 percent or more. In addition, SNAP eligibility is a complex function of many types of income that cannot be measured precisely in any existing data set.

all their income on the outside good. Even when consumers visit a store in a given week, observed spending is relatively small compared to unobserved spending. One concern is that we may be missing important expenditure information by only focusing on two stores. Table 2.3 shows that this concern is minimal and limiting our attention to two stores captures 95 percent of expenditure among all shopping outcomes. Consumer spending is in line with visit patterns and on average around two thirds of all store trips are done at store 1 (the top store in terms of total consumer expenditure). In sum, this shows that most consumers in our sample use one main store for most of their purchases and supplement their purchases by occasional visits to another store; outside of these two stores, consumer spending is minimal. Table 2.4 gives an overview of consumer spending across all categories that form the inside good in our analysis. All categories in the sample are actively purchased as they are top categories by construction. The most popular category is the carbonated beverages category. This observation is consistent with spending patterns reported by the USDA (USDA (2016)). Another observation consistent with the latter report is that there are only small differences in spending habits across income, and therefore we only report spending habits across all consumers. The extent of multi-brand purchasing is low, meaning that consumers tend to pick only one brand when deciding to purchase in a category – this will guide our modeling choices which we outline in the next section.

2.3 Model

In this section, we develop our demand model for brand, category, and store choice. We first outline the challenges in modeling brand-category choice and show our specification

Table 2.2: Summary of Consumer Data.

	Mean	SD	$P_{0.1}$	$P_{0.25}$	$P_{0.5}$	$P_{0.75}$	$P_{0.9}$
Income ($/week)	996	528	433	577	962	1346	1682
Distance to the closer store (mi)	2.7	1.7	1	1.6	2.3	3.4	4.7
Distance to the store further away (mi)	5.2	2.4	2.5	3.6	4.7	6.4	8.6
Consumer's weekly share of outside good (cond. on visiting stores in the set $\{1,2\}$) (%)	98.05	2.3	96.3	97.9	98.7	99.1	99.4

Table 2.3: Use of Stores.

	Mean	St Dev
Prob. of store visit in any given week	0.58	0.19
Prob. of visiting store 1 (cond. on visiting stores in the set $\{1,2\}$)	0.68	0.22
Prob. of visiting store 2 (cond. on visiting stores in the set $\{1,2\}$)	0.21	0.17
Prob. of visiting store 1, 2 (cond. on visiting stores in the set $\{1,2\}$)	0.11	0.11
Expenditure share in store 1 by weekly spending (across *all* stores)	0.69	0.17
Expenditure share in store 2 by weekly spending (across *all* stores)	0.26	0.15

Table 2.4: Summary of Category Data.

	Percent of consumers ever buying in a category (%)	Percent of consumers buying on any given week (%)	Share of category expenditure (%)	Percent of multi-brand purchases (conditional on buying) (%)
Carbonated beverages	97	28	30	30
Composite category	99	31	15	0
Frozen pizza	76	9	12	3
Frozen dinner	67	7	8	12
Yogurt	82	16	6	8
Toilet tissue	78	8	6	1
Laundry detergent	70	6	6	1
Potato chips	89	14	5	6
Coffee	63	6	5	2
Paper towel	72	7	4	2
Tortilla Chips	75	8	3	5

for a brand-category model which is similar to the model specified by Lee and Allenby (2009). We then show how to nest our brand-category model into a store choice framework. Second, we develop a statistical specification of our proposed model and derive the likelihood function.

2.3.1 Direct Utility Model

There are a number of challenges in modeling brand-category demand. First, multiple discreteness, or the purchase of multiple units of multiple alternatives, is a common characteristic of market basket data. Consumers make purchases in multiple categories which means that the utility specification must be able to accommodate corner and interior solutions. Selection of multiple alternatives is also present in store choice data as consumers may visit several stores. Second, the decision to buy from a category may be related to the purchase decision in other categories which means that the model should be able to capture potential cross-category dependencies. For categories that are complements, a joint purchase gives higher utility than the combined utility from separate purchases and for categories that are substitutes the opposite holds true. Third, since we aim to model brand-level purchase quantities rather than category-level purchase incidence, we must adopt a specification that keeps the number of parameters at a reasonable level.

Model of Bundle Choice

We define utility functions that are relevant for bundle choice and set up the consumer problem.

We model consumer i's utility for bundle $\mathbf{q} = (\mathbf{q_{in}}, q_{\text{out}})$ as

$$u(\mathbf{q}) = \Psi^T \mathbf{q_{in}} - \mathbf{H}^T B \mathbf{H} + \frac{q_{\text{out}}^{\beta_0}}{\beta_0} \tag{2.1}$$

where $\mathbf{q_{in}}$ is the vector of inside (store) goods, q_{out} is the quantity of the outside good, and $\mathbf{H} = (u_c)_{c=1}^{C}$ is the vector of category utilities. Utility of category $c \in \{1, \ldots, C\}$ is denoted by u_c and is computed as the sum of category-specific brand-level marginal utilities:

$$u_c = \sum_{k \in \mathscr{I}(c)} \psi_k q_{\text{in},k}, \tag{2.2}$$

where ψ_k is the marginal utility of brand k, $q_{\text{in},k}$ is the quantity of brand k, and $\mathscr{I}(c)$ denotes the index set of brands that belong to category c. The consumer problem for bundle choice is

$$\max_{\mathbf{q}} \quad u(\mathbf{q}) \quad \text{s.t. } \mathbf{p}^T \mathbf{q_{in}} + q_{\text{out}} = E, \tag{2.3}$$

where $\mathbf{p}$ is the price vector of brands and E is the consumer's budget constraint. Notice that the price of the outside good is normalized to one.

The model specification for bundle utility in equation 2.1 is similar to the model in Lee and Allenby (2009) to the extent that both models feature a "linear part" which is an additive function of all the brand preferences, and a "quadratic part" which is an interaction term between category-specific sub-utilities. We specify a separate model for the outside good, whereas Lee and Allenby (2009) treat the outside good as one of the categories. Finally, our statistical specification for the model uses the same strategy but different shock distributions, see section 2.3.2 for details.

Key parameters in the utility function defined in equation 2.1 are Ψ and B. The parameter vector Ψ controls the linear part of the utility function and the matrix B controls the quadratic part of the utility function. These parameters can be interpreted in greater detail by looking at the marginal utility with respect to one of the brands:

$$MU_{cb} := \frac{\partial u(\mathbf{q})}{\partial q_{cb}} = \Psi_{cb}\left(1 - 2\cdot\left(\beta_{C_c,C_c}u_{C_c} + \sum_{k\neq c}^{C}\beta_{C_c,C_k}u_{C_k}\right)\right). \tag{2.4}$$

In equation 2.4, we see that the marginal utility of brand b in category c depends on the sub-utilities of all other categories. When categories c and k are complements, then $\beta_{C_c,C_k} < 0$, and purchases in category k increase the marginal utility, and when the categories are substitutes, then $\beta_{C_c,C_k} > 0$, and purchases in category k decrease the marginal utility. The off-diagonal terms of B, therefore, capture cross-category relationships, while the diagonal terms capture each category's degree of satiation. The parameter Ψ_{cb} can be thought of as the *baseline marginal utility* for brand b in category c which is then either shifted up or down depending on other purchases the consumer makes.

Finally, note that the model implies choice of one brand when the consumer decides to purchase from a category. Equation 2.4 implies that differences between any two alternatives only depend on the baseline marginal utility parameter, the Ψ_{cb} term, and so only one brand will be chosen. This choice behavior is consistent with choice behavior observed in the dataset, in section 2.2.2 we noted that multi-brand purchasing is fairly low across all categories.

A Simple Example

To illustrate the utility function related to bundle choice with a simple example, suppose that the number of stores, categories, and brands is two. The utility function is then

$$
\begin{aligned}
u(\mathbf{q_{in}}, q_0) &= \begin{pmatrix} \Psi_{s_1c_1b_1} \\ \Psi_{s_1c_1b_2} \\ \Psi_{s_1c_2b_1} \\ \Psi_{s_1c_2b_2} \\ \Psi_{s_2c_1b_1} \\ \Psi_{s_2c_1b_2} \\ \Psi_{s_2c_2b_1} \\ \Psi_{s_2c_2b_2} \end{pmatrix}^T \begin{pmatrix} q_{s_1c_1b_1} \\ q_{s_1c_1b_2} \\ q_{s_1c_2b_1} \\ q_{s_1c_2b_2} \\ q_{s_2c_1b_1} \\ q_{s_2c_1b_2} \\ q_{s_2c_2b_1} \\ q_{s_2c_2b_2} \end{pmatrix} - \begin{pmatrix} u_{C_1} \\ u_{C_2} \end{pmatrix}^T \begin{pmatrix} \beta_{C_1,C_1} & \beta_{C_1,C_2} \\ \beta_{C_1,C_2} & \beta_{C_2,C_2} \end{pmatrix} \begin{pmatrix} u_{C_1} \\ u_{C_2} \end{pmatrix} + \frac{q_0^{\beta_0}}{\beta_0} \qquad (2.5) \\
&= \Psi^T \mathbf{q_{in}} - \beta_{C_1,C_1} u_{C_1}^2 - 2\beta_{C_1,C_2} u_{C_1} u_{C_2} - \beta_{C_2,C_2} u_{C_2}^2 + \frac{q_0^{\beta_0}}{\beta_0}.
\end{aligned}
$$

The marginal utility with respect to $q_{s_1c_1b_1}$ is

$$\frac{\partial u}{\partial q_{s_1c_1b_1}} = \Psi_{s_1c_1b_1} \left(1 - 2\left(\beta_{C_1,C_1} u_{C_1} + \beta_{C_1,C_2} u_{C_2}\right)\right), \qquad (2.6)$$

which shows that the marginal utility is non-constant and allows for corner and interior solutions. Figure 2.1 plots the indifference curves and shows that the curves are convex to the origin, allowing for an interior solution, and the curves also intersect the two axes, allowing for a corner solution.

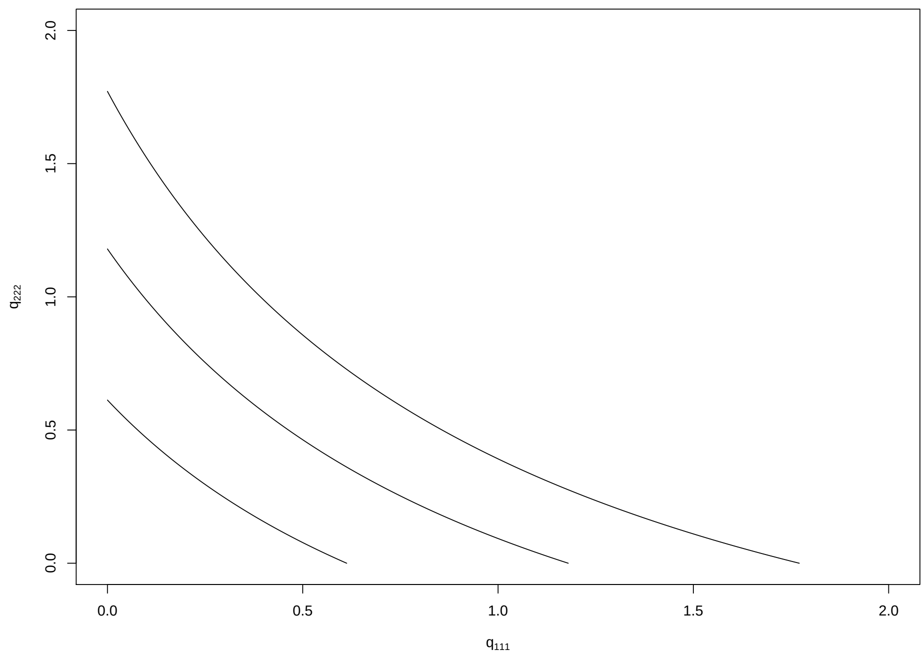

Figure 2.1: Indifference Curves for the Case of Two One-Brand Categories (Two Products) and One Store.

Model of Store Choice

We model consumers' store visit utility as a function of store-specific bundle utility and consumer-store distance and fixed effects. Let $\mathbf{q}_\mathbf{s}^*$ denote the solution to the consumer problem defined in equation 2.3 when bundle choice is restricted to store s. Utility for store s is given by

$$U(s) = u(\mathbf{q}_\mathbf{s}^*) - \tilde{RE}_s - \rho \text{dist}_s + \eta_s, \tag{2.7}$$

where $u(\mathbf{q}_\mathbf{s}^*)$ is the bundle utility, $\tilde{RE}_s$ is the store fixed effect, $\rho \cdot \text{dist}_s$ is travel cost to store s, and η_s is the store error term that is known to the consumer but unknown to the econometrician. The store choice problem is given by

$$\max_s U(s). \tag{2.8}$$

In our empirical specification, we use data on each consumer's two most visited stores. We model the options of no visit and joint visit separately which means the model described in equation 2.7 is a four-alternative model where $s \in \{s_0, s_1, s_2, s_{12}\}$ denotes either the option of no store visit, first store visit, second store visit, or visiting both stores.

2.3.2 Statistical Specification

We use the following decomposition of joint probability of bundle and store choice as the basis for forming the likelihood function later on:

$$P(\mathbf{q}, s) = P(\mathbf{q}|s)P(s). \tag{2.9}$$

The decomposition in equation 2.9 outlines our strategy to compute the likelihood function in two parts. First, to calculate $P(\mathbf{q}|s)$, we introduce stochastic shocks into the marginal utility specification and use Kuhn-Tucker first order conditions for constrained utility maximization to form the likelihood for the observed bundle. Second, to calculate $P(s)$, we define an error distribution of stochastic store shocks.

Computing likelihood of bundle choice, $P(\mathbf{q}|s)$

The likelihood of observed bundle choice is formed based on Kuhn-Tucker or KT (1951) first order conditions for constrained utility maximization. The KT approach assumes that the utility function is random (from the analyst's perspective), implying that the optimal consumption vector is random too. Under this approach, probabilities for corner and interior solutions can be derived by solving the constrained utility maximization problem using KT conditions.

Following a common practice in the literature, we introduce the stochastic element into the marginal utility specification:

$$MU_{scb,\varepsilon} = MU_{scb} \exp\left(\varepsilon_{scb}\right)$$

$$\varepsilon_{scb} \sim_{iid} \text{EVT1}(0, \sigma_\varepsilon) \quad \forall s \in \{1, \ldots, S\}, \forall c \in \{1, \ldots, C\}, \forall b \in \{1, \ldots, B\}. \tag{2.10}$$

The error term captures unobserved (from the analyst's perspective) information that impacts the marginal utility for each good. Our statistical specification for product shocks relies on the extreme value distribution with unknown scale parameter, whereas Lee and Allenby (2009) use a fixed multivariate normal distribution with an identity matrix. The error term is exponentiated to guarantee positivity of the marginal utility. The underlying utility function changes in the following way:

$$u(\mathbf{q}) = (\Psi * \exp(\varepsilon))^T \mathbf{q_{in}} - \mathbf{H}^T B \mathbf{H} + \frac{q_{\text{out}}^{\beta_0} \exp(\varepsilon_{\text{out}})}{\beta_0}, \tag{2.11}$$

where the operation $*$ denotes element by element multiplication between two vectors. Equation 2.11 shows that the stochastic element affects the vector of baseline marginal utility, Ψ; error terms do not affect the vector of category utility, $\mathbf{H}$, and thus there are no squared error terms. Taking derivative of equation 2.11 gives back the expression of marginal utility in equation 2.10. In addition, to guarantee that $\Psi_{scb} > 0$, we parameterize the non-stochastic part of baseline marginal utility as

$$\Psi_{scb} = \exp\left(\psi^T x_{scb}\right) \tag{2.12}$$

where x_{scb} is a set of attributes that characterize the alternative and ψ is the corresponding preference paramater which describes consumer preferences for these characteristics. In the empirical specification, brand characteristics are brand-level dummy variables and a measure of feature and display at the week level.

Optimal Quantity Allocation

Observed purchase quantities are associated with KT conditions that are derived from the following Lagrangian:

$$L = (\Psi * \exp(\varepsilon))^T \mathbf{q_{in}} - \mathbf{H}^T B \mathbf{H} + \frac{q_{\text{out}}^{\beta_0} \exp(\varepsilon_{\text{out}})}{\beta_0} + \lambda \left(E - \mathbf{p}^T \mathbf{q_{in}} - q_{\text{out}}\right), \tag{2.13}$$

where λ is the Lagrangian multiplier associated with the budget constraint. The KT first-order conditions for optimal quantity allocation are given by

$$\frac{\partial L}{\partial q_{scb}} = \Psi_{scb}\left(\exp(\varepsilon_{scb}) - 2\cdot\left(\beta_{C_c,C_c}u_{C_c} + \sum_{k\neq c}^{C}\beta_{C_c,C_k}u_{C_k}\right)\right) - \lambda p_{scb} = 0, \quad \text{if } q^*_{scb} > 0 \tag{2.14}$$

$$\frac{\partial L}{\partial q_{scb}} = \Psi_{scb}\left(\exp(\varepsilon_{scb}) - 2\cdot\left(\beta_{C_c,C_c}u_{C_c} + \sum_{k\neq c}^{C}\beta_{C_c,C_k}u_{C_k}\right)\right) - \lambda p_{scb} < 0, \quad \text{if } q^*_{scb} = 0 \tag{2.15}$$

$$\frac{\partial L}{\partial q_{\text{out}}} = q_{\text{out}}^{\beta_0-1}\exp(\varepsilon_{\text{out}}) - \lambda = 0, \quad \text{since } q_{\text{out}} > 0 \text{ by assumption.} \tag{2.16}$$

The optimal quantity allocation $\mathbf{q}^* = \left(\mathbf{q}^*_{\mathbf{in}}, q^*_{\text{out}}\right)$ satisfies KT first-order conditions in equations 2.14–2.16 and the budget constraint $\mathbf{p}^T\mathbf{q}_{\mathbf{in}} + q_{\text{out}} = E$. We assume that the outside good is always consumed, which makes it possible to substitute $\lambda = q_{\text{out}}^{\beta_0-1}\exp(\varepsilon_{\text{out}})$ into the other first-order conditions and rewrite the KT first-order conditions as follows:

$$\varepsilon_{scb} = \ln\left(2\cdot\left(\beta_{C_c,C_c}u_{C_c} + \sum_{k\neq c}^{C}\beta_{C_c,C_k}u_{C_k}\right) + \exp\left(\varepsilon_{\text{out}}\right)\frac{p_{scb}}{\Psi_{scb}}q_{\text{out}}^{\beta_0-1}\right) =: V_{scb}, \quad \text{if } q^*_{scb} > 0 \tag{2.17}$$

$$\varepsilon_{scb} < \ln\left(2\cdot\left(\beta_{C_c,C_c}u_{C_c} + \sum_{k\neq c}^{C}\beta_{C_c,C_k}u_{C_k}\right) + \exp\left(\varepsilon_{\text{out}}\right)\frac{p_{scb}}{\Psi_{scb}}q_{\text{out}}^{\beta_0-1}\right) =: V_{scb}, \quad \text{if } q^*_{scb} = 0. \tag{2.18}$$

Likelihood Formation

The likelihood for a vector of observed demand $\mathbf{q}^* = \left(\mathbf{q}^*_{\mathbf{in}}, q^*_{\text{out}}\right)$ can be formed by using the distributional assumptions of ε-shocks and the change-of-variables technique. In equation 2.10, we specified the shocks to follow extreme value type-1 distribution with location 0 and scale σ_ε. Without loss of generality, assume that the consumer purchased the first K of M goods. The likelihood for the observed bundle is then given by

$$P\left(q_1,\ldots,q_K,0,\ldots,0|\varepsilon_{\text{out}}\right) = \prod_{j=1}^{K} \frac{1}{\sigma_\varepsilon} g\left(\frac{V_j}{\sigma_\varepsilon}\right) |J| \cdot \prod_{j=K+1}^{M} G\left(\frac{V_j}{\sigma_\varepsilon}\right), \tag{2.19}$$

where $g(\cdot)$ and $G(\cdot)$ are the pdf and cdf of the standard extreme value type-1 distribution. We can see that the likelihood is composed of a density component for all positive goods and a probability mass component for all items that are not purchased. The Jacobian is generated from the change-of-variables technique and is a $K \times K$ matrix with elements

$$J_{scb,s'c'b'} = \frac{\partial \varepsilon_{scb}}{\partial q_{s'c'b'}} = \frac{2\beta_{C_c,C_{c'}} \Psi_{s'c'b'} - \exp(\varepsilon_{\text{out}})\, q_{\text{out}}^{\beta_0-2} \frac{p_{scb}}{\Psi_{scb}} (\beta_0 - 1) p_{s'c'b'}}{\exp(V_{scb})}. \tag{2.20}$$

In order to form the likelihood function, we also have constrain the support of the parameters. We want the utility function to be quasi-concave and all marginal utilities to be positive. The likelihood is given by

$$\begin{aligned} \mathscr{L}\left(\Theta|q_1,\ldots,q_K,0,\ldots,0\right) \quad = \quad & \int P\left(q_1,\ldots,q_K,0,\ldots,0|\varepsilon_{\text{out}}\right) d\varepsilon_{\text{out}} \qquad (2.21) \\ & \times 1\left(\frac{\partial u}{\partial q_{scb}} > 0 \quad \text{for all}\, s,c,b\right) \\ & \times 1\left(u(\mathbf{q_{in}}, q_{\text{out}})\, \text{quasiconcave on the budget constraint}\right) \end{aligned}$$

Quasi-concavity of the utility function ensures sufficiency of the KT first-order conditions, positivity of marginal utility is needed to ensure a proper utility function.

Computing likelihood of store choice, $P(s)$

To compute likelihood of store choice, we need to specify a distribution for store-specific error terms, η_s. We use multi-variate normal distribution for error shocks which means that our statistical specification for store choice is the probit model:

$$U(s) = u(\mathbf{q}_\mathbf{s}^*) - R\tilde{E}_s - \rho \text{dist}_s + \eta_s \tag{2.22}$$

$$\eta \sim MVN\left(\mathbf{0}, \tilde{\Sigma}\right) \tag{2.23}$$

where $\tilde{\Sigma}$ is the full covariance matrix of store-specific error terms. We use two stores in our empirical specification and model the options of visiting both stores and visiting no stores as separate alternatives. This means the store indicator $s \in \{s_0, s_1, s_2, s_{12}\}$ can take on four values. We will discuss normalization for level and scale under identification in the next section.

2.4 Estimation

2.4.1 Identification

We will now discuss identification of model parameters. First, in the model for bundle choice, it is necessary normalize preference parameters for scale. Either one of the psi parameters (describing baseline marginal utility) or one of the beta parameters (describing satiation) needs to be normalized at one for each consumer and each category. We follow the convention used by Lee and Allenby (2009) and set the diagonal terms of the satiation matrix B equal to one which corresponds to normalizing the within-category satiation term for each consumer. Second, in the model for store choice, it is necessary to normalize the model for both level and scale. We use the typical normalization scheme for probit discrete choice models (e.g., Train (2009)). We normalize the four-alternative model described in equation 2.7 for utility levels by taking utility differences with respect to the no visit alternative, and for scale by setting the top-left element of the covariance matrix of error differences to one. Finally, we cannot separately identify individual-specific store fixed-effects and the effect of travel costs as these quantities are constant over time for any consumer-store pair. We define a generalized individual-specific fixed effect as $RE_s = \tilde{RE}_s + \rho \text{dist}_s$ which is identified for $s \in \{s_1, s_2, s_{12}\}$ and let the no visit alternative serve as the baseline. To study the effect of travel costs, we regress estimated fixed effects against store distance.

2.4.2 Heterogeneity

We allow for heterogeneity in shopping behavior by assuming that individual parameters come from population-level distributions with hyperparameters. In the model for bundle choice, each household is characterized by their vector of baseline marginal utilities ψ^h, matrix of within-category and cross-category satiation parameters B^h, and satiation of outside good denoted parameter β_0. In the model for store choice, each household is characterized by a set of generalized store fixed effects, RE_s^h. To sum up, we model household specific parameters with the following population-level distributions:

$$\psi_i^h \sim \text{log-normal}\left(\bar{\psi}_i, \bar{\sigma}_{\psi_i}\right) \quad i = 1, \ldots, K \tag{2.24}$$

$$\beta_{C_i,C_j}^h \sim N\left(\bar{\beta}_{ij}, \bar{\sigma}_{\beta_{ij}}\right) \quad i, j = 1, \ldots, C \tag{2.25}$$

$$\frac{\beta_0^h}{1-\beta_0^h} \sim \text{log-normal}\left(\bar{\beta}_0, \bar{\sigma}_{\beta_0}\right) \tag{2.26}$$

$$RE_s^h \sim N\left(\bar{RE}_s, \bar{\sigma}_{RE_s}\right) \quad s \in \{s_1, s_2, s_{12}\} \tag{2.27}$$

This set-up makes the model hierarchical with population-level hyperparameters at the top level, individual parameters at the next level, and choice data at the lowest level. Markov Chain Monte Carlo estimation is used to estimate all model parameters. We use various methods, including adaptive proposal rules, and give the details of the algorithm in the appendix.

2.4.3 Estimates and Model Fit

We will now discuss parameter estimates and model fit. Figures A.1 and A.2 show estimates of hyperparameters of baseline marginal utility coefficients (the ψ - coefficients in equations 2.4 and 2.24). We focus on the hyperparameters of log-normally distributed individual coefficients since it is easier to spot a trend in the ranking of brands using hyperparameters, rather than individual coefficients which can be close to zero for many brands. Each category has 3–6 brands where the first brands represent top brands in the category and the last brand represents the composite brand. Figure A.1 shows that the mean parameters of brand baseline marginal utility generally follow the pattern predicted by brand market shares, with top brands being associated with higher mean parameters. The ranking of the composite brand depends on the extent of market concentration in the category and may be among the lowest or highest ranking brands. Figure A.2 illustrates spread among the individual coefficients of baseline marginal utility. The general pattern in figure A.2 shows that higher ranking brands are associated with less spread which is an intuitive pattern since more consumers have to value a brand relatively highly for it to be a top brand. We use the concept of Bayesian credible interval – defined as the 2.5th and 97.5th percentile of the posterior distribution – to illustrate precision of estimates. We see that most coefficients are precisely estimated, the credible interval associated with the composite category's brand appears larger because it is the only brand in that category, in addition, the underlying coefficients of baseline marginal utility are extremely close to zero, resulting in a wide range of parameters to fit the data well. Next, figure A.3 shows estimates of cross-category interactions (the β parameters in matrix B in equations 2.4 and 2.25). We focus on individual coefficients since they are directly interpretable and consumers display heterogeneity in whether two categories are treated as subsitutes or complements. Figure A.3 shows boxplots for each cross-category interaction (the elements below the diagonal of the matrix B). Figure A.3 reveals that most category interactions demonstrate a slight degree of substitutability. Interactions of household categories with food categories display a greater degree of substitutability but also show that for some consumers household and food categories can be complements.

2.5 Results

2.5.1 Background on SNAP

We now use the estimated model to conduct counterfactual experiments and study the effect of various subsidy designs. Our population of interest consists of consumers who are eligible for SNAP benefits and we are interested how this population behaves in response to a policy change. We use federal guidelines to simulate whether consumers in our sample are eligible for SNAP benefits. Using criteria from the 2018 fiscal year, we say that a consumer is SNAP eligible when his yearly reported income falls below the eligibility threshold.[10] In section 2.2 we discussed why observing SNAP participation is notoriously difficult. Therefore, we use federal guidelines as a proxy for consumers' SNAP eligibility and assume no SNAP participation as the baseline status.

Consumers receive SNAP benefits via an EBT card (Electronic Benefit Transfer) which is similar to a debit card. The amount of benefits depends on household size and household's own contribution. We use the estimated average monthly benefit for a household of three in the fiscal year of 2018 as the benefit amount we use in counterfactual experiments – that amount is \$384 per month (CBPP (2018)). How do we implement the effect of receiving a food stamp subsidy? Food stamps act as an extra income which can only be spent on food. Therefore, when formulating the consumer's problem, food stamps affect the budget constraint. As a result, the consumer problem for bundle choice in equation 2.3

[10]The threshold is 130 percent of the federal poverty line. In federal fiscal year of 2018, the federal poverty line was \$1702 for a household of three, making the eligibility threshold \$2213 per month or about \$27k per year. (USDA 2018)

is reformulated as follows:

$$\max_{\mathbf{q}} \quad u(\mathbf{q}) \quad \text{s.t. } \mathbf{p}^T \mathbf{q}_{\mathbf{in}} + q_{\text{out}} = E + \Delta_{\text{SNAP}}, \tag{2.28}$$

$$\mathbf{p}^T \mathbf{q}_{\text{food}} \leq E + \Delta_{\text{SNAP}}, \tag{2.29}$$

$$\mathbf{p}^T \mathbf{q}_{\text{non-food}} + q_{\text{out}} \leq E. \tag{2.30}$$

Equations 2.28—2.30 show how the budget constraint changes when the consumer is receiving SNAP benefits, denoted by Δ_{SNAP}. In addition to the general budget constraint (equation 2.28), two other budget constraints need to be accounted for. First, equation 2.29 specifies that food spending must remain below the combined amount of the subsidy and the original budget. Second, equation 2.30 specifies that non-food expenditure and outside good expenditure are still constrained by the original budget. Similarly, restricting subsidy use to food and household items changes the consumer problem in the following way:

$$\max_{\mathbf{q}} \quad u(\mathbf{q}) \quad \text{s.t. } \mathbf{p}^T \mathbf{q}_{\mathbf{in}} + q_{\text{out}} = E + \Delta_{\text{SNAP}}, \tag{2.31}$$

$$\mathbf{p}^T \mathbf{q}_{\text{food}} + \mathbf{p}^T \mathbf{q}_{\text{non-food}} \leq E + \Delta_{\text{SNAP}}, \tag{2.32}$$

$$q_{\text{out}} \leq E. \tag{2.33}$$

Equation 2.32 specifies that the subsidy now applies to food and non-food (household) items and that in this case spending on the outside good is still limited by the original budget (equation 2.33). For cash transfers, no additional constraints are used and the consumer problem remains standard with one budget constraint as specified in equation 2.3.

2.5.2 Marginal Propensity to Consume Food

In this section, we estimate marginal propensity to consume food (MPCF) under different subsidy designs. MPCF is a key indicator which reflects the main goal of the SNAP program and is also the object of main interest in many research papers, allowing to compare our estimates to other estimates previously reported in the literature. Our target population of interest are food stamp eligible consumers who make up around 24% of our sample. Our definition of marginal propensity to consume certain items out of a subsidy benefit relies on comparing expenditure levels on target items before (the no subsidy regime) and after (the subsidy regime) the receipt of benefits. We compare MPCF under (i) normal food stamp benefits (food only), (ii) expanded food stamp benefits (food and household items), (iii) cash subsidy (e.g., tax rebate). To simulate behavior under different policy regimes, we use $S = 20$ draws from the posterior distribution of model estimates. For each draw, we set up a separate optimization exercise corresponding to the subsidy design and record the outcome. We measure total food expenditure across the observation window (52 weeks) with and without a specific subsidy design and compute incremental effect of the subsidy on food expenditure.

Table 2.5: Incremental Impact of Different Subsidies.

Mean Marginal Propensity of Spending per $1	SNAP Benefits	Grocery	Tax Rebate
Food (MPCF)	0.232	0.244	0.126
Household Items	0.077	0.119	0.066
Outside Good	-0.021	-0.033	0.613

Table 2.5 demonstrates how different subsidies affect consumer spending across various categories. The average marginal propensity to consume food (MPCF) out of tax rebate benefits across SNAP eligible consumers is 0.126. This estimate is in line with existing

evidence that MPCF out of cash income for low-income individuals is around 0.1 (Castner and Mabli (2010), Hoynes and Schanzenbach (2009)). We estimate that MPCF out of food stamps is 0.232 or 84% larger than MPCF out of cash which is also consistent with reported results in the literature.[11] Our estimate that MPCF out of grocery benefits is 0.244 or 94% higher than MPCF out of cash is a novel and interesting result. The result is novel because it refers to a counterfactual scenario where SNAP spending restrictions are relaxed to allow spending on household items. In addition, the result is interesting because it is not obvious how extending spending to one set of categories affects spending in another set of categories. In our model, there are two mechanisms that tie together spending between different categories. First, the utility function in equation 2.1 includes cross-category interaction effects which are captured in the quadratic part of the utility function. When spending in one category increases, then spending in another category will also increase (decrease) when the two categories are complements (substitutes). Second, utility from store goods is a key part in the decision to visit a particular store or to not visit any stores (equation 2.7). When one category is subsidized, the overall attractiveness of visiting *any* store is increased. As a result, even in the absence of any cross-category effects, spending in one category may increase when another category is subsidized since consumers are likely to visit more stores and have more opportunities to make purchases from all existing categories. In the context of our counterfactuals, we find evidence that store visit decisions indirectly affect marginal propensity of spending on various categories. Both SNAP and extended SNAP benefits incentivize consumers to visit stores and increase yearly store visits by 4.8% and 5.1%, respectively. Extended SNAP benefits are more effective at

[11]Most papers in the literature estimate an MPCF out of food stamps between 0.16–0.32 (e.g., Bruich (2014),Hoynes and Schanzenbach (2009)). Recently, Hastings and Shapiro (2018) reported higher estimate of 0.5–0.6.

driving recipients to visit stores because they can spend benefits on a wider range of items which in turn increases spending on all categories, including food. Tax rebate, on the other hand, allows recipients to use benefits without visiting any store and encourages to minimize store visits to avoid travel costs. As a consequence, yearly store visits under tax rebate benefits decrease by 4.1% which means that more purchases are concentrated on fewer trips.

2.5.3 Consumer Welfare

In this section, we measure cost-effectiveness of different subsidy designs from the perspective of consumer welfare. We treat consumer welfare under the usual SNAP subsidy as the baseline and ask each consumer

"How many dollars of subsidy X are needed to achieve the utility under the SNAP subsidy?".

We define this concept of equivalent variation using the following notation. Each period, consumers solve the problem of store *and* bundle choice. Let $(\mathbf{q}^*_{s^*}, s^*)$ denote the generic solution to this problem where s^* solves the store choice problem

$$s^* = \arg\max_s U(s) = \arg\max_s \; u(\mathbf{q}^*_\mathbf{s}) - R\tilde{E}_s - \rho \text{dist}_s + \eta_s \tag{2.34}$$

and $\mathbf{q}^*_{s^*}$ solves the bundle choice problem in store s^* subject to restrictions denoted by R:

$$\mathbf{q}^*_{s^*} = \arg\max_{\mathbf{q}} \quad u(\mathbf{q}) \quad \text{s.t. } (R_1)\; \mathbf{p}^T \mathbf{q_{in}} + q_{\text{out}} = E. \tag{2.35}$$

Define

$$\mathbb{U}\left(\{\mathbf{p}_t\}_{t=1}^T, E, R\right) := \sum_{t=1}^T u\left(\mathbf{q}^*_{s^*,\text{t}}\right) \tag{2.36}$$

as the "total consumption utility" over the observation window when solving the problem of bundle and store choice each period facing prices $\{\mathbf{p}_t\}_{t=1}^{T}$, budget E and spending restrictions R. Note that the definition in equation 2.36 on the preceding page only involves the "bundle utility" and ignores any utility components that affect store choice. That is because we treat utility from direct consumption of goods as the natural baseline measure of consumer welfare in the context of our setting.

To compare the cost-effectiveness of different subsidy designs, we are interested in estimating the size of hypothetical benefits $\hat{\Delta}_{\text{SNAP extended}}$ and $\hat{\Delta}_{\text{tax break}}$ which will make a consumer indifferent between receiving the full SNAP subsidy and receiving benefits under a design where spending is restricted to food and household items (extended SNAP) or when there are no spending restrictions (tax rebate), respectively. Using the notation above, we are interested in estimating the size of benefits $\hat{\Delta}_{\text{SNAP extended}}$ and $\hat{\Delta}_{\text{tax break}}$ such that the following holds:

$$\mathbb{U}\left(\{\mathbf{p}_t\}_{t=1}^{T}, E+\hat{\Delta}_{\text{SNAP extended}}, R_{\text{SNAP extended}}\right) = \mathbb{U}\left(\{\mathbf{p}_t\}_{t=1}^{T}, E+\Delta_{\text{SNAP}}, R_{\text{SNAP}}\right) \tag{2.37}$$

$$\mathbb{U}\left(\{\mathbf{p}_t\}_{t=1}^{T}, E+\hat{\Delta}_{\text{tax rebate}}, R_{\text{tax rebate}}\right) = \mathbb{U}\left(\{\mathbf{p}_t\}_{t=1}^{T}, E+\Delta_{\text{SNAP}}, R_{\text{SNAP}}\right). \tag{2.38}$$

The usual SNAP restrictions, denoted by R_{SNAP}, limit *subsidy spending* to food categories only. The extended SNAP subsidy restrictions, denoted by $R_{\text{SNAP extended}}$, restrict subsidy spending to food and household categories. There are no restrictions on subsidy spending when receiving a tax rebate ($R_{\text{tax rebate}}$).

We estimate the size of hypothetical benefits for all consumers who qualify for the SNAP subsidy in our sample. We find significant heterogeneity in benefit sizes which is driven by individual-level differences in consumer preferences and shopping frequency which determine the value of more flexible spending under alternative subsidy designs. We find that at mean-levels the "exchange rate" between different subsidies is the following:

$$\$100 \text{ SNAP benefits} = \$93 \text{ extended SNAP benefits} = \$16 \text{ tax break.} \tag{2.39}$$

The exchange rate in equation 2.39 reflects the extent to which different spending restrictions affect consumer welfare. The substantial difference between tax benefits and the usual SNAP benefits reveals that consumers place a high priority on spending on the outside good. This estimate also describes the incentive to trade SNAP benefits for cash (a form of benefit fraud). Our estimate that on average \$100 in SNAP benefits is equivalent to a \$16 tax rebate is in line with existing evidence that SNAP benefits are traded for cash at the rate of \$0.50 in cash for \$1 in SNAP benefits on online bulletin boards and auction sites.[12] The relatively similar exchange rate between the normal SNAP subsidy and extended SNAP subsidy benefits reveals that while consumers prefer extended benefits which allow subsidy spending on household items, the relative impact of restricting spending to food only is low.

There are two channels in our modeling framework that drive the results above: (i) travel costs and (ii) preference parameters for store goods and outside good. Since there is heterogeneity across consumers, it would be interesting to know whether there are any systematic differences across consumers in the mechanism that drives the "exchange rate" results. First, figure A.4 shows that there are systematic differences among income groups as consumers in the bottom half of the income distribution are willing to accept less cash to be indifferent between receiving the tax rebate subsidy and the food stamp subsidy. As

[12] See Smith (2019) for references to articles in the popular press.

pointed out above, there are two possible channels for this discrepancy. Figure A.5 shows that travel costs are indeed an important factor in driving exchange rate results, however, there are no significant differences in travel costs across the two income groups (figure A.6). As a result, the source of discrepancy in the exchange rate between the two income groups is driven by systematic differences in preferences.

We now discuss implications of these findings for tax policy, benefit design, and consumer welfare. First, government spending on in-kind programs is significant and there are active debates about the desirability of flexible benefits (Lieber and Lockwood (2019)). We quantify the relative trade-offs between normal benefits and flexible benefits in the context of SNAP. Our results show that most of the benefits from lifting restrictions on spending are accrued on the margin of the outside good and not on the inside good. For considering policy change, this result can be viewed from two angles. From the perspective of cost-effectiveness, SNAP benefits could be replaced by much smaller cash transfers in the form of tax rebates, while not affecting consumer welfare. From the perspective of consumer welfare, offering recipients a choice between SNAP benefits or a smaller cash transfer would improve consumer welfare and improve cost-effectiveness of the program. In sum, we quantify differences between various subsidy designs from the perspective of consumer welfare and point out how these results can be used to inform policy.

2.5.4 Restricting Soda Purchases with Soda Tax or Ban

The fact that SNAP is the largest nutrition assistance program and at the same time its benefit recipients spend more money on soft drinks than any other item[13] has spurred

[13] See USDA (2016) for an overview of spending habits of SNAP recipients.

debates on how benefit spending should be restricted[14]. While several states and medical groups have urged changes to the SNAP, the Department of Agriculture (USDA) has denied every request so far. Aside from potential administrative costs of implementing additional restrictions, it is uncertain how restricting benefits would change buying habits. Since most households supplement benefits with their own income, restricting soda purchases may have no effect since households can simply use their own cash instead of the benefits to buy soda, as also pointed out in a recent expert testimony (Schanzenbach (2017)). However, there are circumstances when this prediction does not hold. Recent empirical evidence suggests that consumers reduce spending in target category in response to reduced allowances and only partially substitute reduced benefits with personal funds (Andreyeva et al. (2013)[15]). First, there may be psychological reasons why consumers treat benefits and cash separately (e.g., Hastings and Shapiro (2018) find evidence of mental accounting among SNAP recipients). Second, if consumers follow traditional demand theory then restricting benefit use is equivalent to a price change: when depicting the consumer problem graphically, the budget line pivots in response to the policy change in both cases. Consistent with the latter explanation, consumers in our model respond to subsidy restriction and price increase in the same way. Since the policy maker may have particular preferences to mandate either solution, it would be interesting to know under what circumstances enacting a soda tax or banning soda from SNAP yield the same outcomes in terms of soda consumption.

[14] See for example Lane (2017), O'Connor (2017), Bittman (2012) for articles in the popular press.

[15] Andreyeva et al. (2013) look at recipient response to reduced allowances of sugary juices in the context of the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC). They find that after implementation of the revisions, WIC juice purchases were reduced on par with allowance changes and this reduction was only partially compensated in the form of an increase in juice purchases using personal and other non-WIC funds.

We simulate consumer behavior under two scenarios: enacting a soda tax and banning soda from SNAP. We find that these two policies yield similar outcomes in terms of reducing soda consumption. Figure A.7 shows mean response under three different soda tax policies and a ban policy where soda is excluded from SNAP benefits. Consumer response is calculated with respect to the baseline policy (the full SNAP benefit). Figure A.7 shows that consumers reduce soda consumption by 20 percent and 23 percent under the soda ban and 10% soda tax, respectively. How does rest of the consumer basket change? We find that consumers substitute between different categories within a store, increasing their non-soda food spending substantially (figure A.8) with little changes to the aggregate share of the outside good or shopping frequency. Looking at the effect of soda ban across income groups (figure A.9), we notice that the ban policy has a differential impact on two income groups. In particular, the lowest income consumers only partially substitute to other food items and instead increase their outside good spending relative to the baseline policy. In contrast, consumers whose income is above the median of benefit-eligible consumers do not change their share of outside good purchases and and instead substitute to alternative store goods that are still allowed under the SNAP subsidy with soda excluded. Similar to results in section 2.5.3, consumer response depends on two channels, including (i) travel costs and (ii) preferences for store goods and outside good. Since there are no significant differences in travel costs across the two income groups (figure A.6), we conclude that most of the discrepancy in response is driven by differences in preferences. This contrast among different groups is interesting as it shows how a simple ban policy can have a differential impact through different substitution patterns generated by differences in underlying preferences.

2.6 Conclusion

Many government funded welfare programs provide its recipients in-kind transfers which are subject to different spending restrictions. Since the budgetary and administrative costs of implementing such programs are large, there is a growing debate over whether in-kind benefits should be replaced with simple cash benefits. In the context of SNAP (formerly known as the Food Stamps Program), we analyze the effect of SNAP benefits from the perspective of the policymaker (who prefers that funds are used to buy food) and recipients (who care about overall consumer welfare). The existing literature has documented that marginal propensity to consume food (MPCF) out of food stamps is higher than out of cash, thereby rationalizing the use of food stamps from the perspective of the policymaker. We develop a structural model of consumer demand for brands, categories, and stores and study how SNAP benefits affect spending on different categories. Our structural model yields estimates that are in line with existing results in the literature and provides novel estimates for alternative subsidy designs that are not observed in practice but are often part of the policy debate. Our main finding is that expanding SNAP benefits to grocery items (food plus household goods) would yield outcomes that are preferred by both benefit recipients *and* the policymaker. This finding is interesting because it shows that the new design improves *both* measures of interest (food spending measured in MPCF and consumer welfare) without any compromise, thereby providing a positive answer to the main question posed in this paper. While expanding benefits to household goods would improve consumer welfare, we find that most benefits are accrued on the margin of the outside good and demonstrate that the exchange rate of cash subsidy to SNAP subsidy is around one fifth. Second, we study the other lever which the public has often prompted for the policymakers to pull and examine the effects of banning benefit use on certain goods. We establish that restricting benefit use has similar effects to a tax and in the context of soda category find that a ban would have similar effects on reducing soda consumption as enacting a ten percent tax on soda. We find evidence that excluding soda from SNAP benefits has a differential impact on the substitution patterns of low and high income consumers, which is mainly driven by differences in preferences.

Results and analysis in this paper are relevant for policymakers who need to consider both sides – recipients and program goals set by the policymaker – when considering policy change. Our empirical results highlight the relative trade-offs of existing and alternative policies often considered in the public debate and offer concrete mechanisms to explain the differential impact of subsidy policies.

Chapter 3

The Role of Professional and Amateur Suppliers in Crowd-Based Platforms

3.1 Introduction

In many crowd-based platforms, it is increasingly common to see the supply side consist of a mix of professional and non-professional (amateur) market participants. Professional and amateur sellers co-exist in many common marketplaces, for example, Youtube, AppStore, Etsy, Shutterstock. Since third-party suppliers are vital to a crowd-based platform's functioning and ability to scale rapidly, it is important to understand what role professionals and amateurs play in platform outcomes and policymaking.

Despite behavioral and operational differences between professional and amateur suppliers in many crowd-based markets (e.g., Li et al. (2016)), there is limited research that looks at how supply-side heterogeneity affects policy outcomes of platform initiatives. Past research has focused on the importance of having a large number of suppliers to achieve network effects, abstracting away from implications driven by the heterogeneity of suppliers. Nevertheless, in practice, many platforms design policies that affect professional and amateur suppliers in different ways. For example, changes in minimum quality standards (e.g., Apple's App Store[1]) or pricing tools (e.g., Airbnb[2]) target amateur suppliers, whereas

[1]Moore (2017) gives a discussion of how changes in app design standards may disproportionately affect amateur App Builders and small businesses in the App Store marketplace.

[2]See Airbnb (2017) for description of Smart Pricing in the context of vacation rentals.

changes in monetization options (e.g., Youtube[3], Twitter[4]) mainly benefit professional suppliers. When a policy affects one type of supplier more than the other, the policy may indirectly prompt some suppliers to "switch sides," causing non-incremental changes in response when amateurs become professional or vice versa. As an example, increasing monetization options on a media content platform may directly benefit professional content creators, but may also convince some amateur content creators to start producing content full-time. How would the changing composition of full-time versus hobbyist content creators alter the effect of policy? In the case of monetization through ads, the direct effect on user experience may be negative when consumers find ads annoying. However, the indirect effect on user experience might be positive when more content creators start producing content full-time and raise the quantity and quality of content available. Since changes in the supply-side composition may change the user experience in the opposite direction, it is an empirical question as to whether an ad monetization policy benefits the platform.

The goal of this paper is to empirically examine how supply-side heterogeneity on a crowd-based platform moderates policy outcomes of platform initiatives. We study this question in the setting of a large live-streaming platform, where some content producers work full-time, and others produce content as a hobby. The main distinction between professional and amateur content providers is the ability and motivation to earn revenue. We examine how content creator response affects the impact of a policy change in site-wide monetization.

[3]The announcement made by YouTube (2012) describes extending their Partner program to include more content creators.

[4]Twitter (2015) announces opportunity to place targeted ads.

Preliminary evidence suggests that amateur content creators respond to the policy by creating more content and changing their streaming schedules to include fewer but longer streams. Amateur behavior does not seem to be associated with changes in their performance or direct effects of the ad policy. Instead, amateurs may change behavior because the ad policy benefits professionals both directly and indirectly, raising the prospects of becoming a professional streamer. Since many platforms rely heavily on professional market participants to provide services, understanding "professionalization" of the supply side could yield valuable insight into platform decision making, which we plan as a topic for further research.

Our work is related to the broad literature on multi-sided platforms (Armstrong (2006), Caillaud and Jullien (2003), Parker and Alstyne (2005), and Rochet and Tirole (2003)) and more specifically to the strand of this literature that studies supply-side strategies. The main insight in the theoretical literature is that the strength of cross-group externalities may not only depend on the number of interactions, but also on the quality of interactions, which in our context is captured by differences between professional and amateur suppliers. In this perspective, the platform can control the composition of the supply-side through fees, barriers to access, and other governance rules and thus affect incentives for innovation and investment (Rysman (2009), Hagiu (2009), Wang (2018), Casadesus-Masanell and Zhu (2010), Belleflamme and Peitz (2010), Hagiu and Spulber (2013), Hagiu (2014)). While the theoretical literature is well developed, there is limited empirical research in studying supply-side response and heterogeneity. Li et al. (2016) document differences in performance and behavior between professional and non-professional agents on the Airbnb marketplace. Boudreau (2018) looks at how decreasing minimum development costs may lead to sudden increases in the number of developers, giving rise to the "bottom-falling-out" phenomenon in the marketplace for mobile applications; Boudreau (2012) studies investment incentives of developers in the same market. In the context of user-generated content (UGC) platforms, our paper relates to research that considers "crowd" reactions to policy changes. Sun and Zhu (2013) find that blog writers shift their content toward popular content in response to monetization options with ads, and Wu and Zhu (2018) study suppliers' production speed and product novelty in response to regulatory changes on a novel-writing platform. Ahn et al. (2016) study the interplay of sponsored and user content in the framework of a dynamic structural model of UGC creation. Our paper is different from previous work in that we emphasize how differences in market participants' response moderate policy outcomes. Finally, the specific policy change in our setting is related to the literature on freemium models (Sato (2019), Zennyo (2019), Casadesus-Masanell and Zhu (2010), Bapna et al. (2018)).

The rest of the paper proceeds as follows. In the next section, we describe the institutional setting and the policy change as well as the data we use for our analysis. In section 3.3, we discuss preliminary reduced-form evidence of the effect of the policy and present associated details in the appendix. Section 3.4 concludes.

3.2 Empirical Setting and Data

3.2.1 Background

The empirical setting for this study is Twitch.tv, the market-leading live streaming platform. Twitch, which is owned by Amazon, is approximately the 13th most-viewed website in the US, with over 15 million daily visitors who spend an average of 95 minutes watching content from 3 million monthly content creators[5]. The core of the website relies on user-generated gaming content (UGC), which is provided by individual content creators who use the platform to stream their activities and engage with their audiences in real-time (see figure 3.1). There are various ways content creators can earn a living or supplementary income via live streaming. Twitch creates various opportunities to monetize content and shares the revenues with content creators. Streamers and the platform share revenue in three main ways. First, subscriptions are channel-specific monthly payments, which give viewers channel-specific benefits, including custom emoticons and custom badges, which denote their status in the associated stream chat window. The platform acts as an intermediary, taking a variable cut of the subscription cost. A subscription costs the viewer $4.99, which is usually shared equally between the streamer and the platform, with popular content creators negotiating a higher share. Second, the platform has a donation system centered around an in-platform currency called "bits." Viewers can purchase "bits" and use them to donate to a streamer, with exclusive emoticons available depending on the size of the donation. The "bits" currency intends to make donating fun, entertaining, and straightforward. The streamer receives approximately $1 for every $1.40 given in the in-platform currency. Third, streamers earn a share of ad-revenue from ads shown on their stream. Ads are shown right after viewers click on a channel (onboarding ads) and during the stream (midroll ads). Twitch uses unique video technology[6] to insert ads directly into the broadcast, reducing the effect of third-party services that bypass ads and increasing the advertising revenue pool available. Finally, for options that do not include the platform,

[5]Twitch provides basic user statistics on its website, e.g., on https://www.twitch.tv/p/press-center/.

[6]Twitch calls the technology SureStream, and has been using it since late 2016; see https://blog.twitch.tv/en/2016/11/02/introducing-sure-stream-for-a-better-video-ad-experience-on-twitch-3ca5ce3287c/

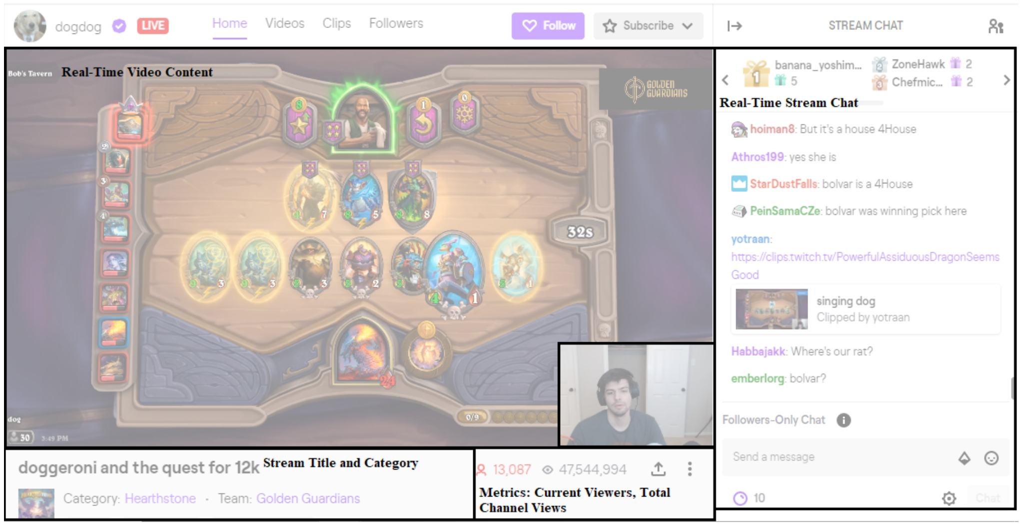

Figure 3.1: Website Layout: A Live Stream on Twitch

streamers can earn revenue from company sponsorships and direct donations on external websites.

It is important to note that not all monetization options are available to all streamers. Depending on their relationship with the platform, streamers belong to one of three categories. Table 3.1 gives an overview of how the availability of monetization options depends on status. Most prominent streamers are given the Partner status, which has a wide number of benefits, including the ability to earn revenue from subscriptions, bits, and ads[7] and the ability to customize these monetization opportunities to channel content. There are around 27000 Partners, corresponding to 1 – 2 percent of all active streamers, and the status is granted on an application basis. As a stepping stone to Partner status, the Affiliate Program allows qualified streamers to start earning revenue from subscriptions

[7] In addition to access to all three monetization opportunities, Partners also enjoy other benefits like priority support, premium technical features, and access to exclusive events. For a complete description of benefits see https://www.twitch.tv/p/partners.

Table 3.1: Monetization Options by Status on the Platform

	Partner	Affiliate	No Status
Subscriptions	✓	✓	
"Bit" donations	✓	✓	
Ad-Revenue	✓		

and bits[8]; however, with fewer customization options available. In contrast to Partner status, entry requirements for the Affiliate Program are relatively low[9], and the onboarding process is automated. There are more than 220000 Affiliates, corresponding to 8 – 10 percent of active streamers. Finally, for steamers with no status, Twitch does not offer any monetization opportunities. However, content creators can link to external websites on their channel page and receive donations outside of the platform.

3.2.2 Policy Change

In studying the role of supplier response, we focus on a significant policy change on the platform. In particular, we study the effect of disabling ad-free viewing for Twitch Prime members, and how supplier response moderated the effect of this policy.

Amazon launched Twitch Prime in 2016 as an exclusive benefit for its Amazon Prime members. Twitch Prime membership is only obtainable with Amazon Prime membership, and not as a standalone service. When launched, the main benefits of Twitch Prime membership were (i) ad-free viewing of the entire website, (ii) one free monthly subscription, (iii) monthly in-game benefits for various games. On August 20, 2018, Twitch announced

[8]While not implemented during our observation period, Affiliates now also earn a share of ad-revenue since November 2019; see https://help.twitch.tv/s/article/ads-experience-updates?language=en_US

[9]Twitch Affiliates need have at least 50 followers, 500 total minutes broadcast, and an average of 3 or more concurrent viewers in the past 30 days (https://affiliate.twitch.tv/).

plans to monetize Twitch Prime members and discontinue ad-free viewing as a Twitch Prime benefit while keeping other benefits. The purpose of the policy was to "strengthen and expand that advertising opportunity for creators."[10] The policy rolled out for new Prime members on September 14, 2018. Existing members with monthly Prime subscriptions continued to receive full benefits until October 15, and members with annual Prime subscriptions continued receiving ad-free viewing until their next renewal date. As a result, the policy change affected the Prime member population[11] in batches which means that the impact on viewers may happen several months after the policy announcement. For annual subscribers who renewed or signed up for Amazon Prime before the September 14, 2018 implementation date, the policy took effect on September 13, 2019, at the latest. After the policy change, the only way to get ad-free viewing is being a member of Twitch Turbo, a monthly subscription service primarily targeted for non-Prime members and countries where no Prime membership is available. While Twitch Turbo may serve as an alternative to some Prime members as a way to get ad-free viewing, we argue that the substitution is likely minimal due to differences in benefits and cost. Twitch Turbo ($8.99 month) is a standalone service only for the Twitch platform, offering ad-free viewing and priority customer support. Amazon Prime ($12.99 month) is a bundled benefits package for a wide variety of Amazon services, including benefits for shipping, music and video streaming, photo storage, shopping online and in brick and mortar stores. Therefore, we expect the decision to be a Prime member not to be tied to the changes in Twitch Prime benefits and use this as an identifying

[10]The full company statement can be found at https://blog.twitch.tv/en/2018/08/20/changes-to-twitch-prime-a986f0d8c9a9/.

[11]In 2018, in a letter to shareowners, Amazon CEO Jeff Bezos revealed that Amazon Prime has more than 100 million subscribers globally. There is some evidence that the share of monthly subscribers has been growing, making up approximately half of all subscribers (see, e.g., https://fortune.com/2020/01/16/amazon-prime-subscriptions/).

assumption in the empirical analysis.

3.2.3 Data Description

The data used in this study originate from several analytics websites[12] that track and aggregate publicly available information on the Twitch platform. We use four data sources, which we discuss below: (i) panel data of 46000 streamers, (ii) clips data, (iii) subscription and "bit" donation data, (iv) use of emoticons data.

First, the panel data set follows streamers' activity and performance statistics for 18 months from May 2018 to October 2019. These data contain information on streamer performance statistics (followers, views, the average number of concurrent viewers), content creation activity (minutes streamed, number of streams, content streamed), and demographics (channel creation date, language, mature content indicator, Partner status). We focus on three significant games[13] that are streamed on the platform and select streamers who streamed at least one of the games during our observation window. Further, we filter streamers who streamed at least one of the games for 30 hours over the observation period and had an average of five or more concurrent viewers. Setting activity requirements helps limit the size and noise of the data. As a result, our analysis focuses on the subpopulation of streamers who meet these minimal activity requirements. Second, the clips data consist of all clips with more than five views, and data are collected for every streamer with more than an average of five concurrent viewers. Clips are viewer-made short video clips of memorable moments in stream, allowing viewers to share unique content and

[12]We received data directly from twitchstats.net, twitchanalysis.top, and sullygnome.com.

[13]We choose the most popular game in three different categories. Fortnite is a battle royale game, League of Legends is a battle arena game, and Hearthstone is a collectible card game.

enable content creators to grow their channels through social sharing. Third, subscription and "bit" donation data are based on mining event notifications from the chat window. When viewers subscribe or donate "bits," a notification is shared with the chat, alerting the content creator and the audience of the event. Chat window contents are recorded for the top 8000 live channels and offline channels with more than a hundred average concurrent viewers past month. Finally, data on emoticon use are collected similarly to subscription and donation data. Used emoticons are recorded in the stream chat window and aggregated to the daily level. Emoticons play a significant role in the chat experience, allowing viewers to interact with the community and express their feelings.

Since the selection criteria for the panel data and clips data are the same, we can match these two datasets, and call this our primary dataset. We use subscription, "bit" donation, and emoticon data as complementary to our analysis.

3.3 Reduced-Form Evidence

How do we expect the change in ad policy to affect the platform ecosystem? In what follows, we discuss how the affected population of Amazon Prime users on the demand-side and content creators on the supply-side of the platform may react to the change in ad policy. We refer to content creators with Affiliate status as "Amateurs," and use "Professionals" to denote content creators with Partner status. Table 3.2 gives an overview of the expected impact with a focus on differences between Professionals and Amateurs.

On the supply side, the direct effect of the policy is increased ad inventory, which directly benefits content creators with Partner status[14]. Since the prospect of becoming

[14] See table 3.1 for an overview of monetization benefits depending on status.

Table 3.2: Expected Impact of the Ad Policy

	Amateurs	Professionals
Content Quality	Substantial increase as an effort to achieve Partner status	Minimal increase since most Partners are already creating high-quality content
Broadcasting Time	Substantial increase as an effort to achieve Partner status	Minimal increase since most Partners are already engaged in full-time streaming
Channel Visits	Decline in incoming traffic since on-boarding ads discourage browsing of smaller channels	Increase in incoming traffic since on-boarding ads discourage browsing of smaller channels
Concurrent Viewership and Engagement	Indeterminate effect since (i) mid-roll ads tend to decrease user engagement, (ii) predicted higher content quality may increase engagement	Indeterminate effect since (i) mid-roll ads tend to decrease user engagement, (ii) predicted higher content quality may increase engagement
Subscriptions and Bit Donations	Indeterminate effect since (i) incoming user traffic tends to decrease, (ii) effect on viewership and engagement is indeterminate	Indeterminate effect since (i) incoming user traffic tends to increase, (ii) effect on viewership and engagement is indeterminate
Advertising Revenue	Increase in ad revenue since the target population did not see any ads before	Increase in ad revenue since the target population did not see any ads before

a professional streamer increases, we expect more content creators to take steps toward making streaming their career. Existing Partners may increase their broadcasting time and content quality to some degree. However, we expect the adjustment to be minimal since many Partners are already engaged in full-time streaming. Amateurs who are considering full-time streaming, on the other hand, have increased incentives to achieve Partner status in order to transition to full-time streaming. Since achieving Partner-status requires fulfilling criteria pertaining to broadcasting activity and audience size, we expect those Amateurs who want to transition to professional streaming to increase their broadcasting time and content quality substantially.

On the demand side, the direct effect of the ad policy on viewer experience is likely to be negative. Viewers may find ads annoying when advertising content interferes with the main content that viewers are interested in (Goldstein et al. (2014), Wilbur (2008)). Ad content on the platform appears in the form of on-boarding and mid-roll ads, affecting both browsing and viewing experience. On-boarding ads make the experience of visiting several content creators' channels less fluent, increasing browsing costs. If site visitors weigh expected utility of visiting a channel against the costs of watching an on-boarding ad, we expect the number of browsed channels per visitor to decrease. As a result, we expect increased traffic toward Partnered channels, which offer a higher expected utility on average and are more visible in the directory of channels when channels are sorted by the number of concurrent viewers. Mid-roll ads may affect user engagement and viewing experience of channel visitors who stay and watch channel content for a longer period. Since ads intervene with the main content of viewers' interest, we expect viewer engagement (concurrent viewership, follows, subscriptions) to decrease. However, since viewer reaction may also depend on content quality, which is predicted to increase, the overall effect of the ad policy on consumer browsing and content engagement is ambiguous.

Finally, what is the expected benefit of the ad policy from the platform's perspective? The platform shares revenue with content creators from two sources: (i) subscriptions and "bit" donations, (ii) advertising revenue. The amount of subscriptions and "bit" donations depends on the interplay of how many users visit a channel and what proportion of the visiting users find the content engaging enough to donate. Since the presence of ad content and higher content quality are predicted to shift user engagement in opposing directions, the effect on subscription and "bit" donation revenue from Amateurs and Professionals is ambiguous. Ad revenue, however, will certainly increase since the affected user population did not see any ads before the policy change.

To test the predictions in table 3.2, we will now outline a series of empirical tests and present summary statistics. Table 3.3 provides an overview of data summary statistics before and after the policy change. First, looking at the share of different content creators, we see that Amateur content creators are becoming more prevalent on the platform. This fact is consistent with the prediction that more Amateur streamers may be interested in streaming as a full-time career. Second, and related to the first point, looking at the amount of content produced on the platform, we see that the share of Amateur content is increasing. However, Partners still produce the majority of watched content on the platform. In particular, Affiliate content is becoming more popular, with the share of watched content more than doubling in the period after the policy. Third, we can see that changes in content popularity are, to some extent, driven by changes in broadcasting effort: Affiliates stream more hours after the policy, and more of their content is watched, which is in line with predictions in table 3.2. In sum, changes in streaming effort and composition

Table 3.3: Data Summary Statistics (Before and After Policy)

	Content Creators by Status (%)		Minutes Watched (%)		Hours Streamed (hrs/week)	
	Before	After	Before	After	Before	After
No Status	4.2	5.7	0.9	0.9	8.11	7.84
Affiliate	81.2	82.7	5.5	11.8	15.1	16.3
Partner	14.6	11.6	93.6	87.3	24.4	23.9

Notes: This table reports summary statistics for the data sample. We use four months of data before the policy (May – August 2018) and four months of data after the policy (November 2018 – February 2019). Policy implementation happened in September 2018 (new subscribers) and October 2018 (existing subscribers). Minutes watched is calculated as the product of average viewers and time streamed.

of active streamers suggest some reasons why Amateur content is becoming more popular. Since table 3.3 provides simple averages and does not condition on underlying trends, we look at reduced-form evidence next (see appendix B.1 for an overview of details).

Figure B.3 shows estimated changes in content output at months around the ad policy change, measured in minutes streamed. We see that Amateurs increase content output shortly after the policy, with the effect falling off after about six months later. Compared to changes in Amateur content, Partner content remains relatively stable shortly after the policy period. These broad patterns on Amateur and Partner behavior shortly after the policy change are consistent with predicted changes in table 3.2. The overall decline in streaming activity after six months since the policy introduction could be related to dynamic effects due to the policy or undocumented trends. Looking more closely at what is driving changes in content output, figures B.1 and B.2 show changes along the intensive and extensive margin, respectively. Figure B.1 shows that while all content creators stream more per occasion, the trend is most significant for Affiliate streamers. The number of streaming occasions (figure B.2) remains stable after the policy. However, after about six months, all streamers broadcast content on fewer occasions, which coincides with the trend of decreasing total content output. In conclusion, on the supplier side, we see a persistent increase in the intensive margin (minutes per stream). In the extensive margin (number of streams), we observe little to no response after the policy and then a downward trend. Consistent with our expectations in table 3.2, Affiliate streamers respond the most, with a non-incremental shift in broadcasting activity. Since the number of minutes streamed per occasion remains consistently high after the policy, we can interpret the changes in total content output as streaming schedules becoming more "professional," with fewer but longer streams. However, streamers may also adjust their effort in response to viewer behavior, which we analyze next.

Figure B.4 shows estimated changes in channel visits at months around ad policy change. The "visit" variable reflects browsing habits because "visits" are recorded after clicking on a channel in the game directory or loading the channel page directly. The ad policy affects visits to Partnered channels, which get significantly higher traffic, with no effect on Amateur channels. This observation is consistent with predicted browsing behavior, where the presence of on-boarding ads discourages users from clicking on channels with fewer viewers in anticipation of low expected utility. Looking at the average number of concurrent viewers (figure B.5), we see large gains for Partners and close to zero-size effects for Amateurs. Interestingly, the ability to convert channel visitors to engaged viewers is not constant during the observation window. Figure B.6 shows estimated changes in average viewership of a channel while controlling for the amount of visitor traffic. More channel visits turn into concurrent viewers in the period right after the policy, with the trend reversing in periods of higher visitor traffic. Interpreting the trend of the visit-to-viewer ratio over time shows that the gained traffic resulted in more viewers but at a lower rate. These patterns illuminate the equilibrium outcome of two opposing forces: (i) ad-annoyance effects, and (ii) content creator incentives to offer higher content quality. While the initial prediction for user engagement on Partnered and Amateur channels was indeterminate (table 3.2), figures B.5 and B.6 show how these effects played out in practice. For Amateurs, ad-annoyance effects and increased content quality seem to cancel each other out, resulting in close to zero-size effects. Partners, however, benefit considerably, indicating that content quality effects outweigh the ad-annoyance effects. The effect of "higher Partner content quality" can be explained in two ways: (i) Partners offering higher quality content than before, (ii) new users being exposed to Partner content as a result of substitution in browsing toward Partnered channels.

In conclusion, when looking at how content creators react to the change in ad policy, then Amateurs seem to be the most responsive in terms of adjusting their content output, consistent with the predicted behavior of "wanting to become Professional." Since the ad policy offers direct monetization benefits to Partners, the prospect of becoming a Partner is more valuable for those who do not yet have Partner status. As a result, Amateurs put in extra effort in an attempt to achieve Partner status, which requires fulfilling requirements that pertain to audience size and streaming activity. This result is interesting because it captures a regime shift for a group of agents that goes beyond the incremental response to the policy. Since Partnered content creators represent the platform brand and create the most engaging content, the question to what extent may platform policies induce "professionalization" among Amateur streamers and its attendant implications to platform outcomes is important for decision making.

The large standard errors of Partner effects – which result in wide confidence intervals

– indicate that the effect varies substantially across different Partners, pointing to the need to account for heterogeneity among the Partner population with more detailed modeling, which we leave for future research.

3.4 Conclusion

We have examined how heterogeneity in supply-side composition in the form of professional and amateurs may alter the effect of an ad monetization policy on a large live-streaming platform. Preliminary evidence suggests that amateurs respond to the policy with a regime switch by creating more content and changing their streaming schedules to include fewer but longer streams. Amateur behavior does not seem to be associated with changes in their performance or direct effects of the ad policy. Instead, amateurs may change behavior because the ad policy benefits professionals both directly and indirectly, raising the prospects of becoming a professional streamer. Since many platforms rely heavily on professional market participants to provide services, the question to what extent may platform policies induce "professionalization" among amateur streamers and its attendant implications to platform outcomes is important for decision making. The empirical results presented in this paper provide evidence of professionalization among amateur streamers, and we plan to study this topic in future research.

Our study has several limitations. We have not currently incorporated all available data into the analysis; most importantly, we are yet to include data that could illustrate shifts in content quality (clips data, emoticons used data). Another direction for future research is to set up an explicit empirical model to evaluate the impact of the ad policy.

Chapter 4

Conclusion

This dissertation presents two essays at the intersection of quantitative marketing and empirical industrial organization. We now conclude the work presented in the two chapters.

The first chapter is motivated by the debate over the question of whether the in-kind benefits of government welfare programs should be replaced with simple cash benefits. In the context of SNAP (formerly known as the Food Stamps Program), we analyze the effect of SNAP benefits from the perspective of the policymaker (who prefers that funds are used to buy food) and recipients (who care about overall consumer welfare). We develop a structural model of consumer demand for brands, categories, and stores and study how SNAP benefits affect spending on different categories. Our structural model yields estimates that are in line with existing results in the literature and provides novel estimates for alternative subsidy designs that are not observed in practice but are often part of the policy debate. Our main finding is that expanding SNAP benefits to grocery items (food plus household goods) would yield outcomes that are preferred by both benefit recipients and the policymaker. This finding is interesting because it shows that the new design improves both measures of interest (food spending measured in MPCF and consumer welfare) without any compromise, thereby providing a positive answer to the main question posed in the chapter. We also study the effects of banning benefit use on certain goods. We establish that restricting benefit use has similar effects to a tax and, in the context of the soda category, find that a ban would have similar effects on reducing soda consumption as enacting a ten percent tax on soda. Interestingly, we find evidence that excluding soda from SNAP benefits has a differential impact on the substitution patterns of low and high-income consumers, which is mainly driven by differences in preferences. Results and analysis in the first chapter are relevant for policymakers who need to consider both sides – recipients and program goals set by the policymaker – when considering policy change. Our empirical results highlight the relative trade-offs of existing and alternative policies often considered in the public debate and offer concrete mechanisms to explain the differential impact of subsidy policies.

The second chapter is motivated by the presence of heterogeneous suppliers (professionals and amateurs) present on many common crowd-based platforms. I examine how heterogeneity in supply-side composition in the form of professional and amateurs may alter the effect of an ad monetization policy on a large live-streaming platform. Preliminary evidence suggests that amateurs respond to the policy with a regime switch by creating more content and changing their streaming schedules to include fewer but longer streams. I also find evidence that the ad-policy raises prospects of becoming a professional streamer. Since many platforms rely heavily on professional market participants to provide services, the question to what extent may platform policies induce "professionalization" among amateur streamers is important for platform strategy.

Appendix A

Appendices for chapter 2

A.1 Overview of Markov Chain Monte Carlo Estimation

This section gives an overview of the estimation algorithm. For individual coefficients, we use random walk Metropolis algorithm (for β_0^h) and Adaptive Metropolis (AM) algorithm of Haario et al. (2001) for vector components (ψ^h and B^h). To update the hyperparameters, we use posterior distributions based on the updated individual coefficients. In the first case, a new individual parameter value β_0^{h*} is proposed by perturbing the current logit ratio

$$\text{logit}\left(\beta_0^{h*}\right) = \text{logit}\left(\beta_0^{h}\right) + \sigma_{MET,\beta_0}\bar{\sigma}_{\beta_0} N(0,1),$$

and then backing out the implied value by $\beta_0^{h*} = \frac{\exp\left(\text{logit}\left(\beta_0^{h*}\right)\right)}{1+\exp\left(\text{logit}\left(\beta_0^{h*}\right)\right)}$. The value for the current iteration is set to β_0^{h*} if

$$U < \frac{P(\mathbf{q}|\beta_0^{h*})P\left(\beta_0^{h*}|\bar{\beta}_0, \bar{\sigma}_{\beta_0}\right)}{P(\mathbf{q}|\beta_0^{h})P\left(\beta_0^{h}|\bar{\beta}_0, \bar{\sigma}_{\beta_0}\right)} \quad \text{where } U \sim \text{Uniform}(0,1),$$

and left unchanged otherwise. The metropolis standard deviation, σ_{MET,β_0}, is individual specific and is modified after every 500 iterations to make the acceptance rate of the chain close to 44 percent, the optimal rate of acceptance in a one-dimensional context (Roberts and Rosenthal (2001)). The metropolis standard deviation is increased if the acceptance rate is greater than the target acceptance rate, and decreased if the acceptance is smaller than the target acceptance rate. In the second case, we use vector jumping and accept or reject the entire vector. For updating ψ^h, the proposal is given by

$$Q\left(\psi^{h*}\right) = (1-\beta)N\left(\psi^h, (2.38)^2\Sigma_{\text{emp}}/d\right) + \beta N\left(\psi^h, (0.1)^2 I_d/d\right),$$

where Σ_{emp} is the current empirical estimate of the covariance structure of target distribution based on the run so far, and where $\beta = 0.05$ is a small positive constant. The idea is that the proposal $N\left(\cdot, (2.38)^2\Sigma/d\right)$ is optimal in a particular large-dimensional context (Roberts and Rosenthal (2001)), and the empirical version is an effort to approximate this. A small amount of normal noise is added to avoid the chain from getting stuck. The algorithm can be adapted during the warm-up phase by specifying H, the length of memory, and c, the coefficient in front of Σ_{emp}. We take $H = 200$ and update the empirical covariance structure after every H iterations and fix the covariance structure after the warm-up period is over. For the coefficient, we start with $c = 2.38$ as shown above and modify it after every 50 iterations to make the acceptance rate close 23 percent, the optimal acceptance rate in large-dimensional contexts (Roberts and Rosenthal (2001)).

To initialize the sampler, we draw initial values for the hyperparameters and then generate individual values from their corresponding distributions. To generate initial utility shocks, we use equations:

$$\varepsilon^0_{scb,t} = \ln\left(2\cdot B^0_{sc,\cdot}\cdot u^0_t + \frac{p_{scb,t}}{\psi^0_{scb}}\cdot q_{0,t}^{\beta^0_{out}-1}\exp(\varepsilon^0_{0,t})\right) \quad q^*_{scb} > 0$$

$$\varepsilon^0_{scb,t} \sim EVT1(0,\sigma^0_\varepsilon) \quad \text{s.t} \quad \varepsilon^0_{scb,t} < \ln\left(2\cdot B^0_{sc,\cdot}\cdot u^0_t + \frac{p_{scb,t}}{\psi^0_{scb}}\cdot q_{0,t}^{\beta^0_{out}-1}\exp(\varepsilon^0_{0,t})\right) \quad q^*_{scb} = 0.$$

A.2 Figures

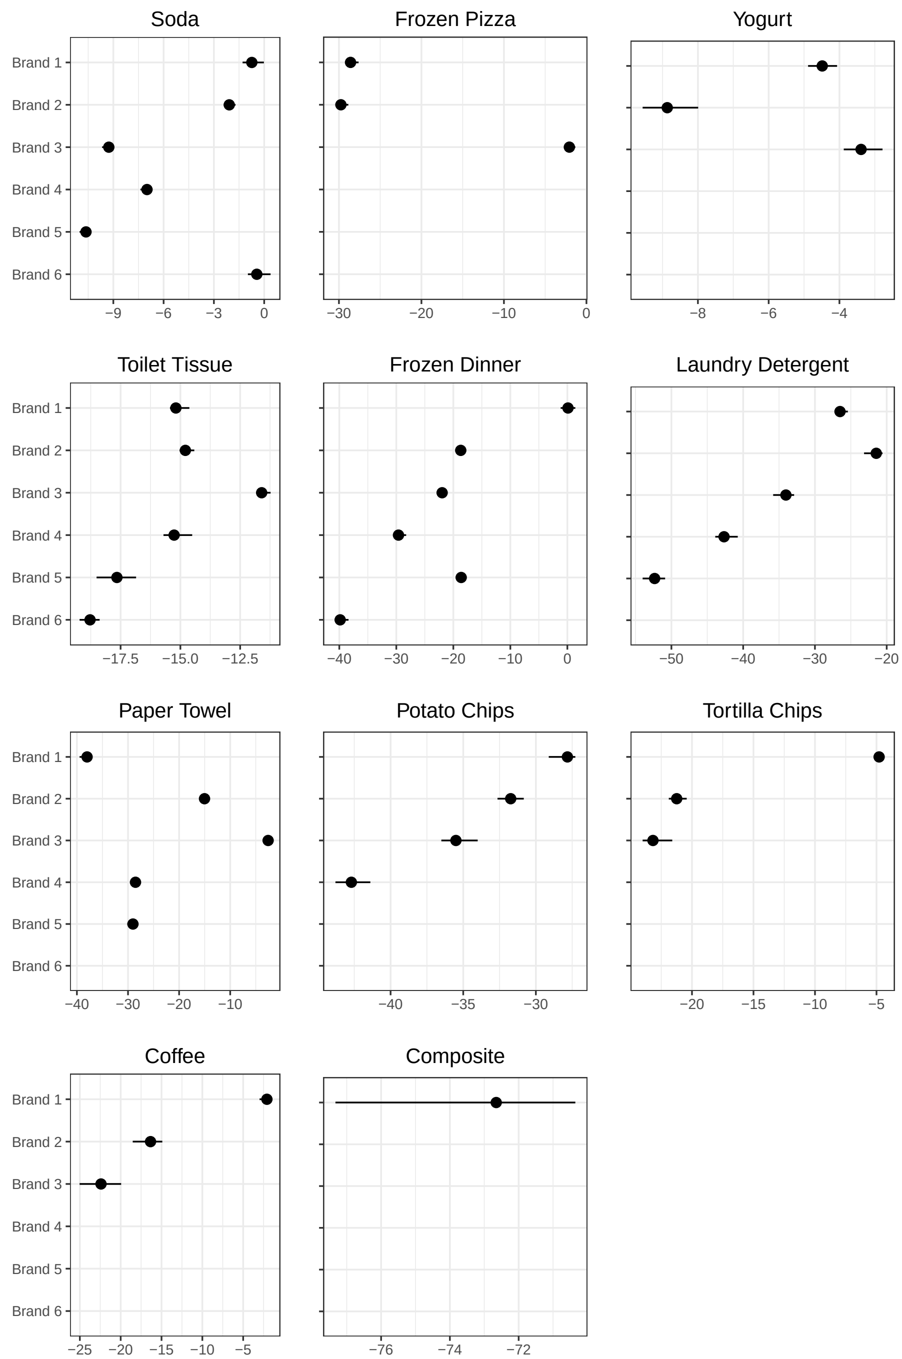

Figure A.1: Estimates of the Mean Parameter of Log-Normally Distributed Baseline Marginal Utility ($\bar{\psi}$).

Notes: This figure presents posterior mean estimates and 95% Bayesian credible intervals for population means of log-normally distributed brand preferences (the $\bar{\psi}_i$ - parameters in $\psi_i^h \sim$ log-normal$\left(\bar{\psi}_i, \bar{\sigma}_{\psi_i}\right) \quad i = 1, \ldots, K$). There are 3–6 brands in each category, with the first brands representing top brands in the category and the last brand representing the composite brand. Calculations are based on $S = 100$ random samples from the thinned MCMC chain where every 300th iteration is saved.

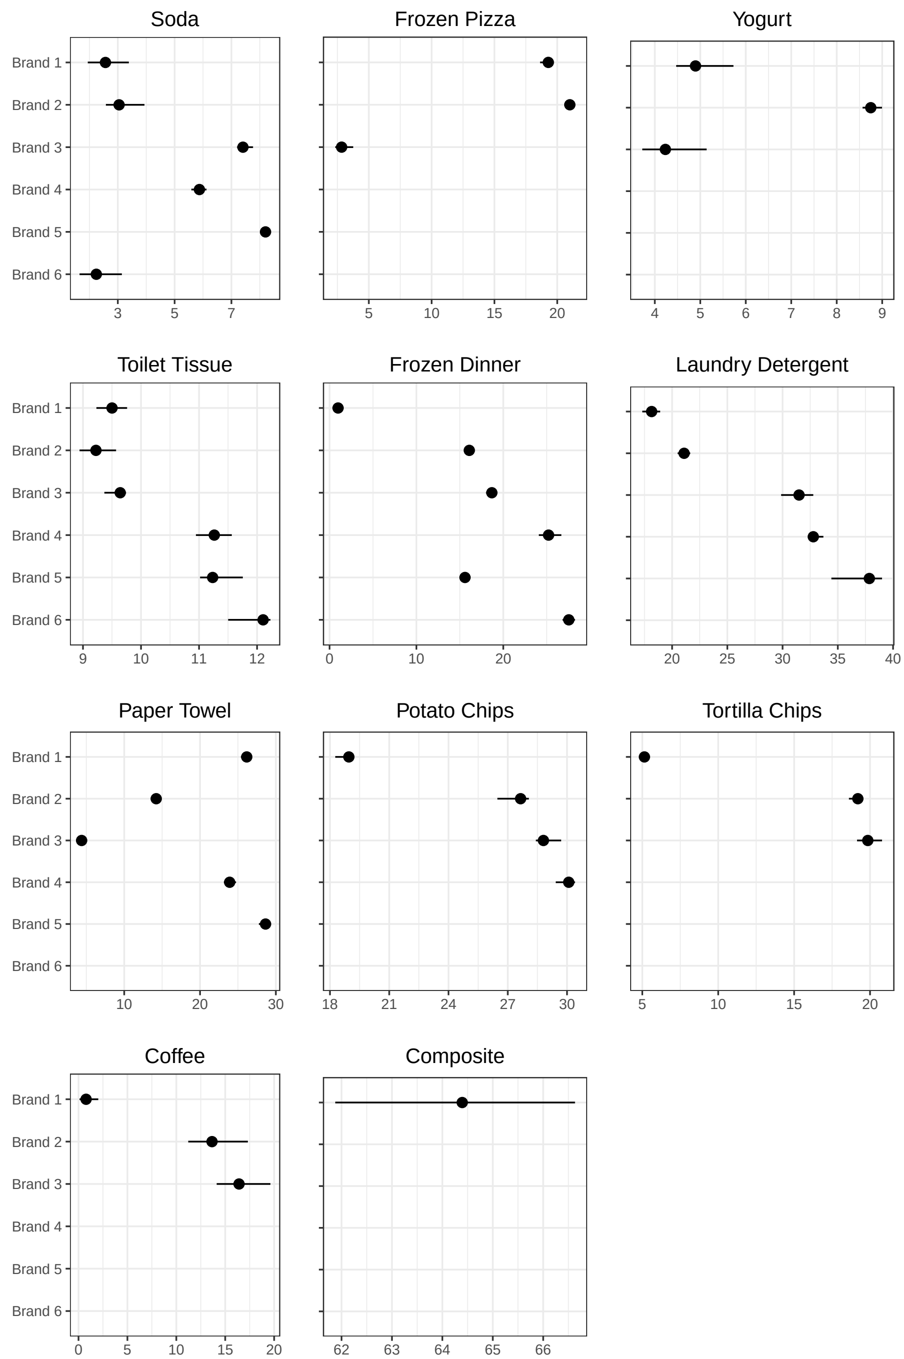

Figure A.2: Estimates of the Standard Deviation Parameter of Log-Normally Distributed Baseline Marginal Utility ($\bar{\sigma}_{\psi}$).

Notes: This figure presents posterior mean estimates and 95% Bayesian credible intervals for population standard deviations of log-normally distributed brand preferences (the $\bar{\sigma}_{\psi}$ - parameters in $\psi_i^h \sim \text{log-normal}\left(\bar{\psi}_i, \bar{\sigma}_{\psi_i}\right) \quad i = 1, \ldots, K$). There are 3–6 brands in each category, with the first brands representing top brands in the category and the last brand representing the composite brand. Calculations are based on $S = 100$ random samples from the thinned MCMC chain where every 300th iteration is saved.

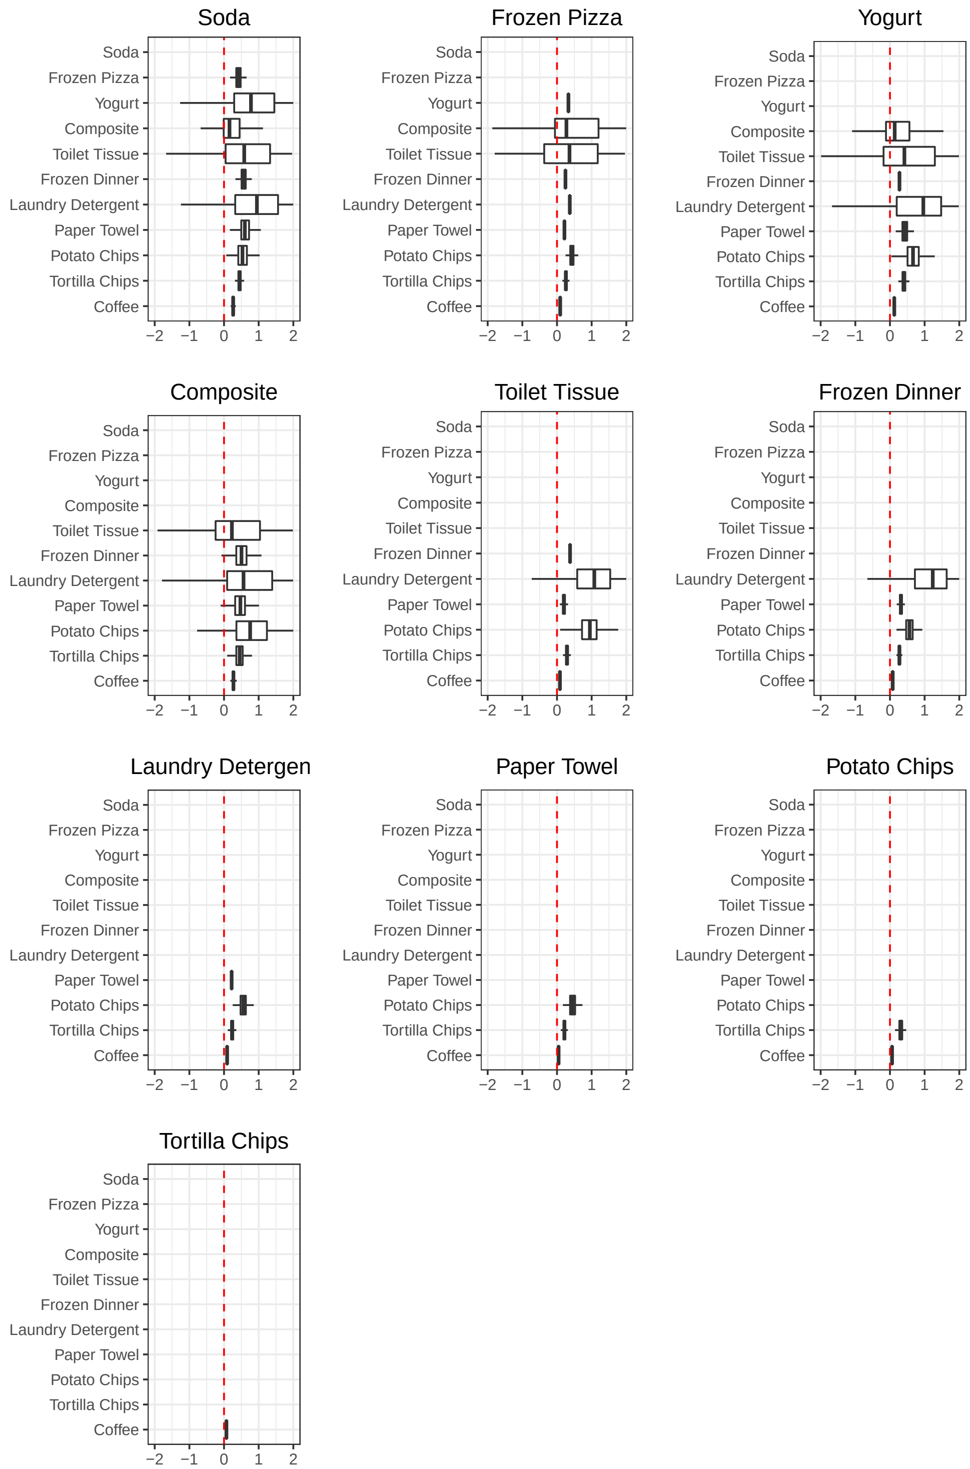

Figure A.3: Category Interaction Effects.

Notes: This figure presents boxplots of individual coefficients of category interaction effects. Positive coefficients indicate degree of substitutability and negative coefficients indicate degree of complementarity. Estimate of each individual's coefficient is based on the posterior mean of $S = 100$ random samples from the thinned MCMC chain where every 300th iteration is saved.

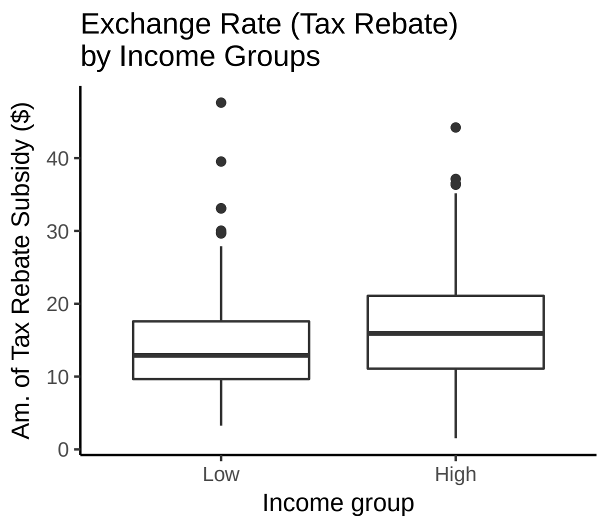

Figure A.4: Exchange Rate by Income Groups.

Notes: This figure presents boxplots of subsidy "exchange rate" across two income groups. Two income groups are generated by dividing eligible consumers (monthly income below \$2203) into two groups based on the median income (monthly income \$1346). The exchange rate on the y-axis is the amount of dollars in cash subsidy (no restrictions) that would make consumers indifferent between choosing the cash subsidy in the respective amount or a \$100 food stamp subsidy.

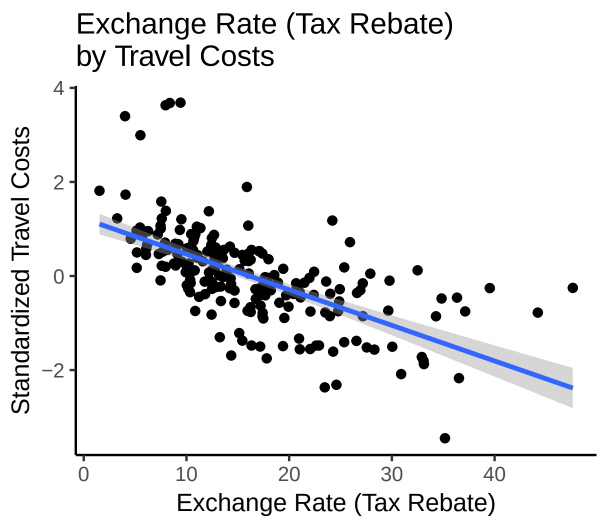

Figure A.5: Exchange Rate by Travel Costs.

Notes: This figure presents a scatterplot of subsidy "exchange rate" across standardized measure of travel costs. The exchange rate on the x-axis is the amount of dollars in cash subsidy (no restrictions) that would make consumers indifferent between choosing the cash subsidy in the respective amount or a $100 food stamp subsidy. Travel costs comprise of individual-store fixed effect and a distance measure (see equation 2.7 for details) and are standardized (Z-score).

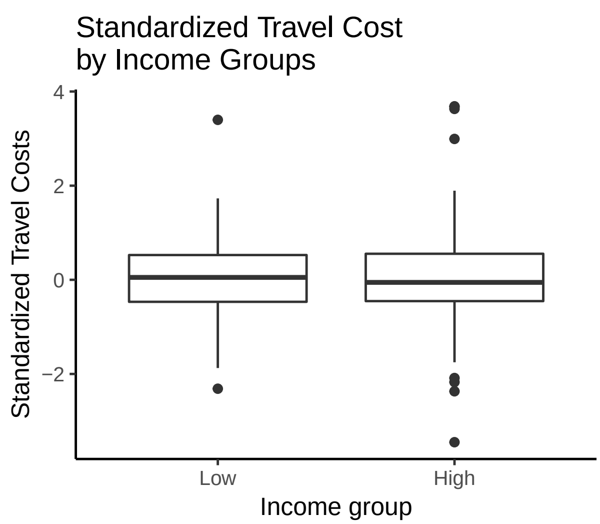

Figure A.6: Travel Costs Across Two Income Groups.

Notes: This figure presents boxplots of standardized travel costs across two income groups. Two income groups are generated by dividing eligible consumers (monthly income below $2203) into two groups based on the median income (monthly income $1346). Travel costs comprise of individual-store fixed effect and a distance measure (see equation 2.7 for details) and are standardized (Z-score).

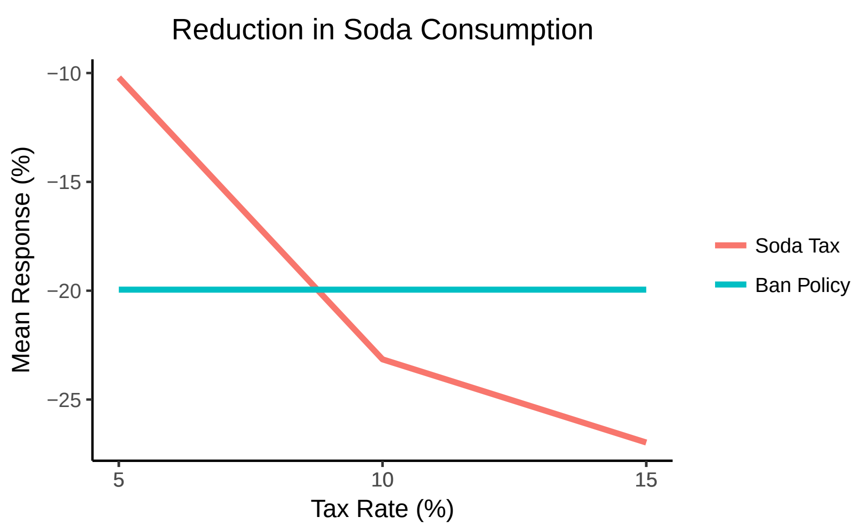

Figure A.7: Reduction in Soda Consumption Across Different Policies Relative to SNAP Policy.

Notes: This figure presents a line plot of reduction in soda consumption (relative to the baseline SNAP policy) across different policies. The tax policy is computed at three points, corresponding to enacting a 5% , 10% , and 15% tax on soda while still keeping consumers subject to the usual SNAP policy. The ban policy excludes soda from SNAP benefits, while keeping benefits otherwise the same.

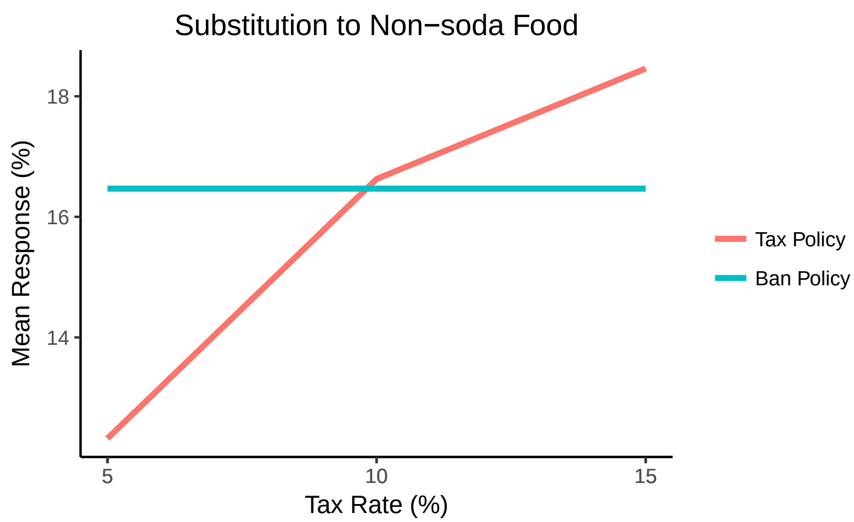

Figure A.8: Substitution to Non-soda Food.

Notes: This figure presents a line plot of increase in non-soda food consumption (relative to the baseline SNAP policy) across different policies. The tax policy is computed at three points, corresponding to enacting a 5% , 10% , and 15% tax on soda while still keeping consumers subject to the usual SNAP policy. The ban policy excludes soda from SNAP benefits, while keeping benefits otherwise the same.

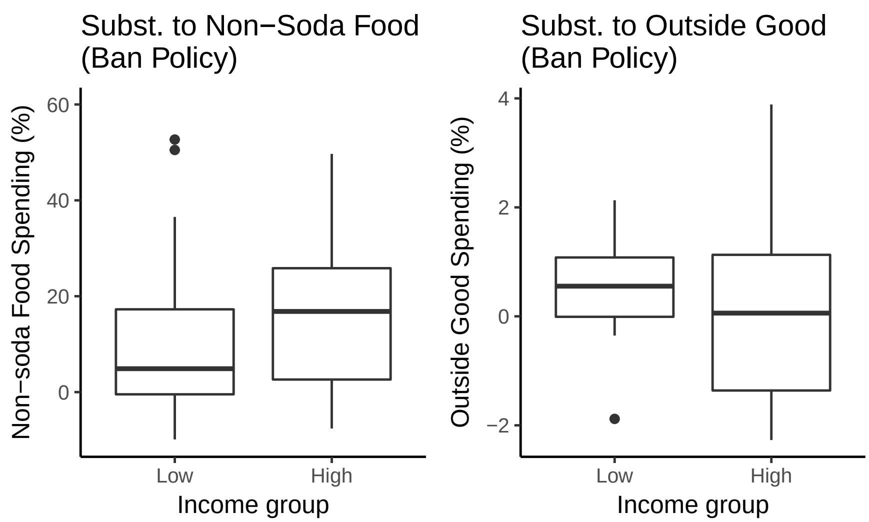

Figure A.9: Comparison of Substitution Between Non-soda and Outside Good Across Two Income Groups.

Notes: This figure presents a boxplot comparison of two variables across two income groups. The figure compares substitution patterns between non-soda and outside good across two income groups. Two income groups are generated by dividing eligible consumers (monthly income below $2203) into two groups based on the median income (monthly income $1346).

Appendix B

Appendices to Chapter 3

B.1 Details on Reduced-Form Regressions

In this subsection, we give details of econometric specifications that we describe in the section on reduced-form evidence. Our econometric specifications have the following event-study form:

$$Y_{igt} = \sum_{k=-4}^{13} \delta_k^g \cdot \mathbb{1}_k + X_{it} + \kappa_i, \tag{B.1}$$

where Y_{igt} is the outcome variable for content creator i with status group $g \in$ {Partner, Affiliate, No Status} at time t. The $\mathbb{1}_k$'s are indicator variables for the number of k months before or after ad-free viewing is discontinued for new Amazon Prime members. The content creator specific fixed-effect is denoted by κ_i and X_{it} contains various control variables, such as main game played and channel age in month t. We describe different types of control variables below and table B.1 gives an overview of all specifications.

1. Baseline Controls

   (a) $MainGame_{it}$ – a categorical variable that denotes the monthly main game. Game g is the main game at month t if streamer i spends more than 50 percent of all streaming time on game g. Categories are (i) Fortnite, (ii) League of Legends, (iii) Hearthstone, (iv) Other.

   (b) $ChannelAge_{it}$ – a categorical variable that denotes content creator maturity on the website. Categories are (i) up to 3 months, (ii) 3+ – 6 months, (iii) 6+ – 12 months, (iv) more than 12 months.

2. Community Activity

   (a) $InitialFollowers_{it}$ – a continuous variable that denotes the number of followers at the beginning of month t.

   (b) $FollowerGain_{it}$ – a continuous variable that denotes the number of new followers in month t.

3. Streamer Activity

   (a) $MinStreamed_{it}$ – a continuous variable that denotes the number of minutes streamed in month t.

4. Directory Effect

   (a) $MainGame_{it} \times lag\left(AvgViewers_{it}\right)$ – an interaction effect that captures predicted placement in this month's main game directory based on past month's viewership.

Table B.1: Overview of Reduced-Form Regressions

Control Variables		Dependent Variable: *Streamer Output* — Number of Streams	*Streamer Output* — Total Minutes Streamed	*Streamer Output* — Total Minutes Streamed	*Channel Performance* — Channel Visits	*Channel Performance* — Average Number of Viewers	*Channel Performance* — Average Number of Viewers	*Channel Performance* — Follower Gain
Regression Specific Controls				Number of Streams			Channel Visits	
Baseline Controls	$MainGame_{it}$ $ChannelAge_{it}$	✓	✓	✓	✓	✓	✓	✓
Community Activity	$InitialFollowers_{it}$ $FollowerGain_{it}$				✓	✓	✓	✓
Streamer Activity	$MinStreamed_{it}$				✓	✓	✓	✓
Directory Effect	$MainGame_{it} \times$ lag $(AvgViewers_{it})$				✓	✓	✓	✓
Streamer Fixed Effects		✓	✓	✓	✓	✓	✓	✓

B.2 Figures

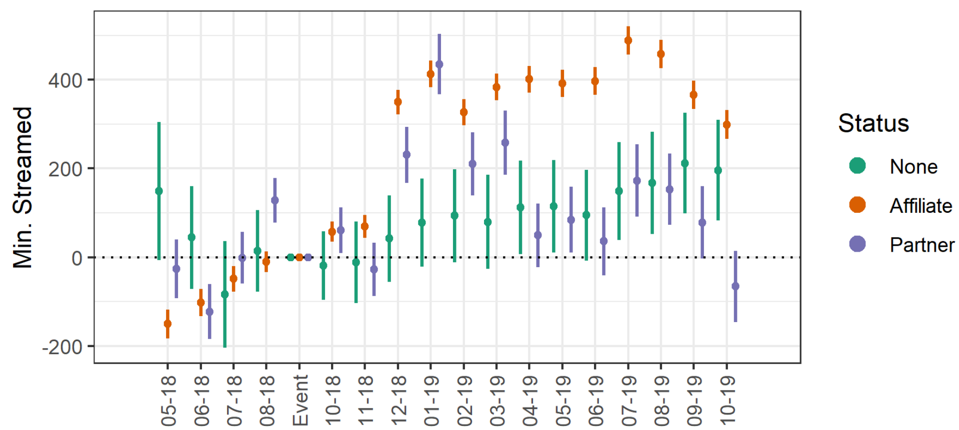

Figure B.1: Estimated Changes in Minutes Streamed Per Stream at Months around Ad Policy Change

Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

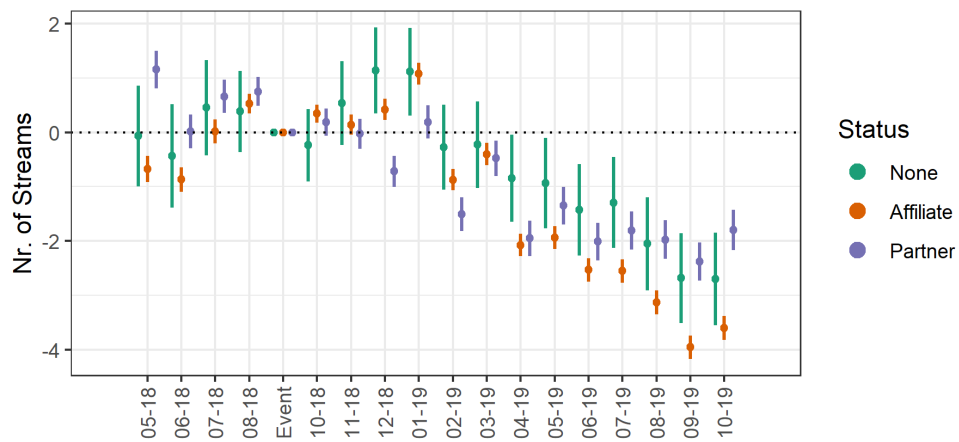

Figure B.2: Estimated Changes in the Number of Streams at Months around Ad Policy Change

Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

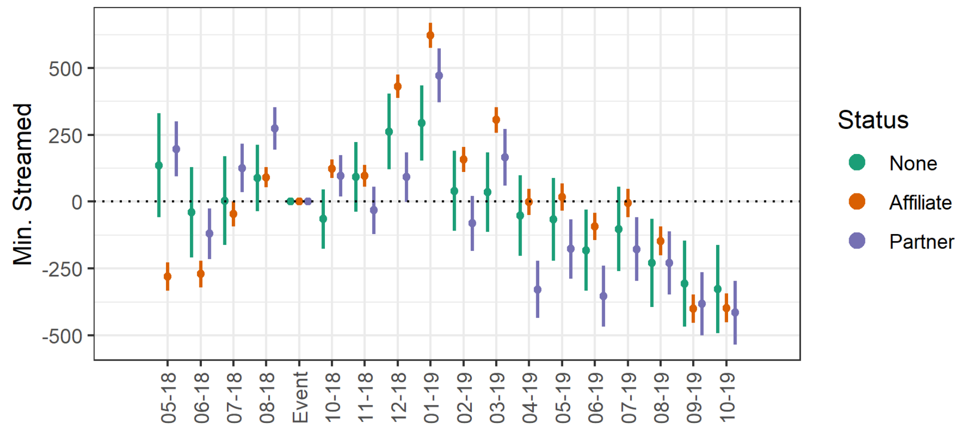

Figure B.3: Estimated Changes in Total Minutes Streamed at Months around Ad Policy Change

Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

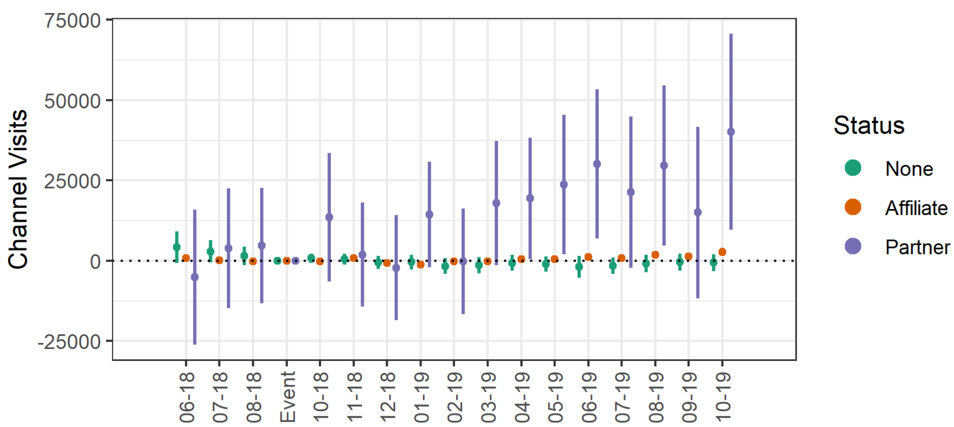

Figure B.4: Estimated Changes in Channel Visits at Months around Ad Policy Change
Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

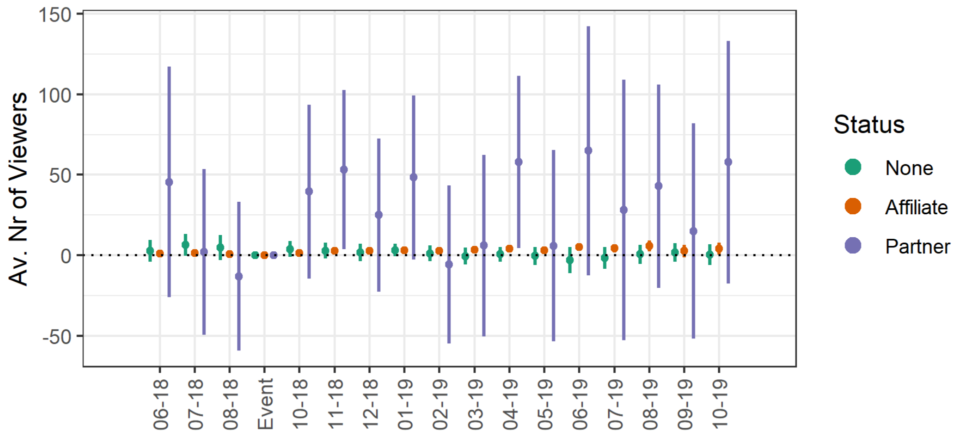

Figure B.5: Estimated Changes in the Average Number of Viewers at Months around Ad Policy Change

Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

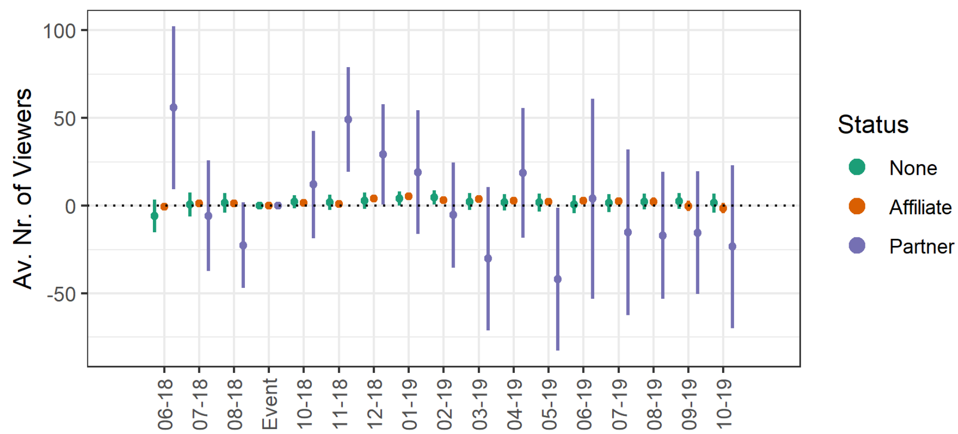

Figure B.6: Estimated Changes in the Average Number of Viewers Per Channel Visits at Months around Ad Policy Change

Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

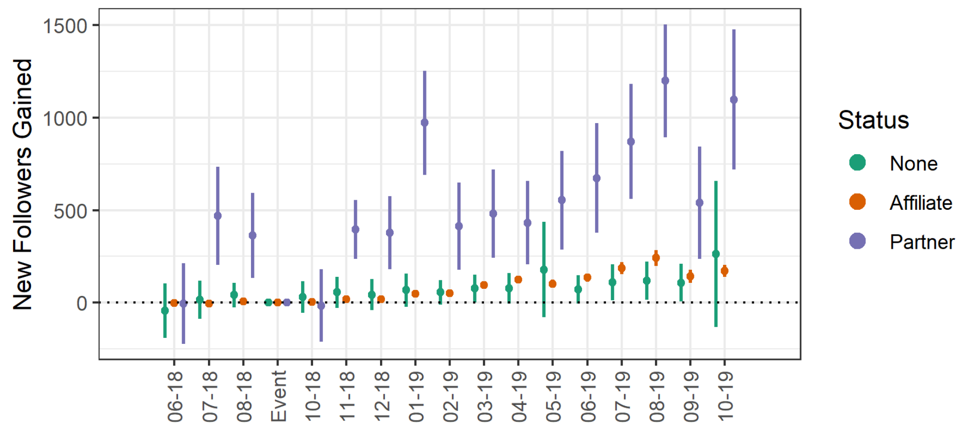

Figure B.7: Estimated Changes in Follower Gain at Months around Ad Policy change
Notes: Figure plots estimates of coefficients δ from equation B.1 and 95 percent confidence intervals. Changes in ad policy are implemented in September 2018 ("event") for new subscribers, and in October 2018 for existing monthly subscribers; annual subscribers are subject to the policy after their renewal date. Standard errors are clustered at the content creator level.

Bibliography

Ahn, D.-Y., J. A. Duan, and C. F. Mela (2016). Managing user-generated content: A dynamic rational expectations equilibrium approach. *Marketing Science*.

Airbnb (2017, November). What's smart about smart pricing? Airbnb Blog.

Allcott, H., R. Diamond, J.-P. Dubé, J. Handbury, and I. Rahkovsky (2019). Food deserts and the causes of nutritional inequality. *Quarterly Journal of Economics*.

Andreyeva, T., J. Luedicke, A. S. Tripp, and K. E. Henderson (2013). Effects of reduced juice allowances in food packages for the women, infants, and children program. *Pediatrics*.

Arcidiacono, P., P. B. Ellickson, C. F. Mela, and J. D. Singleton (2018). The competitive effects of entry: Evidence from supercenter expansion.

Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*.

Bapna, R., J. Ramaprasad, and A. Umyarov (2018). Monetizing freemium communities: Does paying for premium increase social engagement? *MIS Quarterly*.

Beatty, T. K. M. and C. J. Tuttle (2015). Expenditure response to increases in in-kind transfers: Evidence from the supplemental nutrition assistance program. *American Journal of Agricultural Economics*.

Bell, D. R., J. Chiang, and V. Padmanabhan (1999). The decomposition of promotional response: An empirical generalization. *Marketing Science*.

Bell, D. R. and J. M. Lattin (1998). Shopping behavior and consumer preference for store price format: Why "large basket" shoppers prefer edlp. *Marketing Science 17*(1), 66–88.

Belleflamme, P. and M. Peitz (2010). Platform competition and seller investment incentives. *European Economic Review*.

Bittman, M. (2012). Stop subsidizing obesity. *The New York Times*.

Bollinger, B. and S. E. Sexton (2018). Local excise taxes, sticky prices, and spillovers: Evidence from berkeley's soda tax.

Boudreau, K. J. (2012). Let a thousand flowers bloom? an early look at large numbers of software app developers and patterns of innovation. *Organization Science*.

Boudreau, K. J. (2018, April). Amateurs crowds & professional entrepreneurs as platform complementors. Working Paper 24512, NBER.

Briesch, R. A., P. K. Chintagunta, and E. J. Fox (2009). How does assortment affect grocery store choice? *Journal of Marketing Research 46*(2), 176 – 189.

Bronnenberg, B., M. W. Kruger, and C. F. Mela (2008). The iri marketing data set. *Marketing Science 27*(4), 745–748.

Bruich, G. A. (2014). The effect of snap benefits on expenditures: New evidence from scanner data and the november 2013 benefit cuts.

Caillaud, B. and B. Jullien (2003). Chicken & egg: Competition among intermediation service providers. *The RAND Journal of Economics*.

Cannon, M. F. (2016). Distribute medicare's budget in cash so recipients can choose care. *The New York Times*.

Casadesus-Masanell, R. and F. Zhu (2010). Strategies to fight ad-sponsored rivals. *Management Science*.

Castner, L. and J. Mabli (2010). Low-income household spending patterns and measures of poverty. Technical report, Mathematica Policy Research.

CBPP (2018). A quick guide to snap eligibility and benefits.

Chandukala, S. R., J. Kim, T. Otter, P. E. Rossi, and G. M. Allenby (2008). Choice models in marketing: Economic assumptions, challenges and trends. *Foundations and Trends in Marketing 2*(2), 97–184.

Chintagunta, P. K. and H. S. Nair (2011). Structural workshop paper - discrete-choice models of consumer demand in marketing. *Marketing Science 30*(6), 977–996.

Figurelli, L. (2013). Store choice in spatially differentiated markets. Technical report, Boston College.

Goldstein, D. G., S. Suri, R. P. Mcafee, M. Ekstrand-Abueg, and F. Diaz (2014). The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research*.

Gordon, B. R., A. Goldfarb, and Y. Li (2013). Does price elasticity vary with economic growth? a cross-category analysis. *Journal of Marketing Research*, 4–23.

Haario, H., E. Saksman, and J. Tamminen (2001). An adaptive metropolis algorithm. *Bernoulli 7*(2), 223–242.

Hagiu, A. (2009). Quantity vs. quality and exclusion by two-sided platforms. Working Paper 09-094, Harvard Business School Strategy Unit.

Hagiu, A. (2014). Strategic decisions for multisided platforms. *MIT Sloan Management Review*.

Hagiu, A. and D. Spulber (2013). First-party content and coordination in two-sided markets. *Management Science*.

Hastings, J. and J. M. Shapiro (2018). How are snap benefits spent? evidence from a retail panel. *American Economic Review*.

Hoynes, H. W. and D. W. Schanzenbach (2009). Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Applied Economics*.

Kenny, C. (2015). Give poor people cash. *The Atlantic*.

Kim, H. (2012). Estimation of store choice model with endogenous shopping bundles. *Available at SSRN 2361305*.

Kuhn, H. W. and A. W. Tucker (1951). *Nonlinear Programming*. Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability. University of California Press, Berkeley, CA.

Lane, C. (2017). Getting soda out of snap. *The Washington Post*.

Lee, S. and G. M. Allenby (2009). A direct utility model for market basket data. *Fisher College of Business Working Paper* (1443390).

Li, J., A. Moreno, and D. J. Zhang (2016, August). Pros vs joes: Agent pricing behavior in the sharing economy. Working Paper 1298, Michigan Ross School of Business.

Lieber, E. M. and L. M. Lockwood (2019). Targeting with in-kind transfers: Evidence from medicaid home care. *American Economic Review*.

Moore, M. (2017, August). App store review guidelines causing big problems for smbs, diy app builders. SD Times.

O'Connor, A. (2017). In the shopping cart of a food stamp household: Lots of soda. *The New York Times*.

Pakes, A. (2010). Alternative models for moment inequalities. *Econometrica 78*(6), 1783–1822.

Parker, G. G. and M. W. V. Alstyne (2005). Two-sided network effects: A theory of information product design. *Management Science*.

Reinhardt, U. E. (2011). Provide cash, or benefits in kind? *The New York Times*.

Roberts, G. O. and J. S. Rosenthal (2001). Optimal scaling for various metropolis-hastings algorithm. *Statistical Science 16*, 351–367.

Rochet, J.-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*.

Rysman, M. (2009). The economics of two-sided markets. *Journal of Economic Perspectives*.

Sato, S. (2019). Freemium as optimal menu pricing. *International Journal of Industrial Organization*.

Schanzenbach, D. W. (2017). Pros and cons of restricting snap purchases.

Schiraldi, P., S. Seiler, and H. Smith (2012). Supermarket retail competition with multi-stop shopping: who competes with a big box store. Technical report, Working Paper.

Shriver, S. and B. Bollinger (2017). Structural analysis of multi-channel demand.

Singh, V. P., K. T. Hansen, and R. C. Blattberg (2006). Market entry and consumer behavior: An investigation of a wal-mart supercenter. *Marketing Science 25*(5), 457–476.

Smith, T. (2019). Mental accounting and grocery shopping.

Sun, M. and F. Zhu (2013). Ad revenue and content commercialization: Evidence from blogs. *Management Science*.

Thomassen, O., H. Smith, S. Seiler, and P. Schiraldi (2017). Multi-category competition and market power: A model of supermarket pricing. *American Economic Review*.

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press.

Twitter (2015, March). Introducing partner audiences. Twitter Blog.

USDA (2016). Foods typically purchased by supplemental nutrition assistance program (snap) households.

USDA (2018). Snap participation and costs 1969-2018. https://www.fns.usda.gov/pd/supplemental-nutrition-assistance-program-snap.

Walters, R. G. and S. B. MacKenzie (1988). A structural equations analysis of the impact of price promotions on store performance. *Journal of Marketing Research 25*(1), pp. 51–63.

Wang, E. Y. (2015). The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity. *The RAND Journal of Economics*.

Wang, J. (2018). Quality screening and information disclosure in two-sided markets. *Economics Letters*.

Wilbur, K. C. (2008). A two-sided, empirical model of television advertising and viewing markets. *Marketing Science*.

Wu, Y. and F. Zhu (2018, August). Competition, contracts, and creativity: Evidence from novel writing in a platform market. Working Paper.

YouTube (2012, April). Being a youtube creator just got even more rewarding. Youtube Official Creator Blog.

Zennyo, Y. (2019). Freemium competition among ad-sponsored platforms. *Information Economics and Policy*.

Zhu, T., V. Singh, and A. Dukes (2011, 06). Local competition, entry, and agglomeration. *Quantitative Marketing and Economics 9*(2), 129–154.

Biography

Rudolf-Harri Oberg completed his Ph.D. in economics at Duke University in May 2020. Before coming to Duke, Rudolf earned his B.A. in economics and a side-degree in mathematics from the University of Tartu (Estonia) in 2013. His research considers questions in consumer behavior and public policy, advertising, and online marketplaces. He uses methods from structural modeling, causal inference, and Bayesian econometrics.