Essays on Food Security Measurement, Food Security Programs and Substance Use

ABSTRACT

ESSAYS ON FOOD SECURITY MEASUREMENT, FOOD SECURITY PROGRAMS AND SUBSTANCE USE

By

JOSÉ ODAIR DE ALMEIDA XILAU

August, 2025

Committee Chair: Dr. Rusty Tchernis

Major Department: Economics

This dissertation consists of three chapters examining food security and the economics of teenage substance use.

My first chapter investigates the benefit spending behavior of households who participate in the Supplemental Nutrition Assistance Program (SNAP), the largest nutrition assistance program in the United States. I utilize a restricted administrative panel dataset to investigate several aspects of the Supplemental Nutrition Assistance Program benefit cycle and uncover potential unobserved heterogeneity. I find evidence supporting the existence of two types of SNAP households, which I denote as slow and fast spenders. Fast spenders comprise the majority of SNAP households, 77%, and spend approximately 70% of their benefit in the first week. In addition, fast spenders shop more frequently throughout the month, stay longer in the program, 3 years compared to 1 year and 4 months for slow spenders, and receive benefits that are almost twice as large, on average.

In my second chapter, I propose a novel Bayesian Graded Response Model (BGRM) for food security measurement. The BGRM produces continuous food security estimates and measures of estimation uncertainty at the household level. I utilize data from the 2017-18 National Health and Nutrition Examination Survey (NHANES) to estimate the BGRM for responses to the 10 adult core Food Security Module (FSM) questions. I find non-trivial uncertainty in household-level estimates, overlap in BGRM estimates across USDA-defined food security categories and significant within-category variation.

In my third chapter, I examine the effect of macroeconomic conditions on teenage marijuana, cigarette and alcohol use. Using data from the Youth Risk Behavior Surveillance System, I find evidence of counter-cyclical teenage alcohol and cigarette use among females and 12 to 14 year olds. These results are consistent with evidence in the literature.

ESSAYS ON FOOD SECURITY MEASUREMENT, FOOD SECURITY PROGRAMS AND SUBSTANCE USE

BY

JOSÉ ODAIR DE ALMEIDA XILAU

A Dissertation Submitted in Partial Fulfillment
of the Requirements for the Degree
of
Doctor of Philosophy
in the
Andrew Young School of Policy Studies
of
Georgia State University

GEORGIA STATE UNIVERSITY
2025

ACCEPTANCE

This dissertation was prepared under the direction of the candidate's Dissertation Committee. It has been approved and accepted by all members of that committee, and it has been accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Economics in the Andrew Young School of Policy Studies of Georgia State University.

Dissertation Chair:	Dr. Rusty Tchernis
Committee:	Dr. Will Davis
	Dr. Jonathan Smith
	Dr. Keith Teltser

Electronic Version Approved:

Thomas J. Vicino, Dean
Andrew Young School of Policy Studies
Georgia State University
AUGUST, 2025

Acknowledgments

I am extremely grateful to several people who supported and encouraged me throughout the completion of my dissertation.

First and foremost, I would like to thank my dissertation committee chair, trusted advisor and friend, Rusty Tchernis. You have been an invaluable source of guidance both professionally and personally. I deeply appreciate your support, kindness and excellent mentorship throughout my PhD years. I would also like to deeply thank George William Davis, for his excellent support of my dissertation, and his guidance throughout the academic publication process. Additionally, I would also like to deeply thank the other members of my dissertation committee, Jonathan Smith, and Keith Teltser for their support.

Completing my dissertation would not have been possible without the support and love of my family and my partner. Therefore, I would like to thank my parents, André and Isabel, and my siblings Naira, Amandio, Paulo and Roberto for their unending support and faith in me. Your support and encouragement has made this long journey, an enriching and fulfilling experience. I would also like to thank my partner, Erica Mtenga. You have been a great source of inspiration throughout this journey. I am extremely grateful to have you by my side and happy to complete this journey together.

Last, but not least, I would like to thank my esteemed colleagues and friends in the program. Your friendship, and support have made the PhD experience a pleasure and I am excited for what is to come in the next stage of our careers.

Table of Contents

List of Tables

List of Figures

Introduction

This dissertation consists of three essays related to food security and the economics of substance use. This research aims to shed additional light on the behavior of households who participate in food assistance programs, the measurement of food security and the relationship between unemployment and substance use.

My first chapter examines the Supplemental Nutrition Assistance Program (SNAP) Benefit Cycle. The Supplemental Nutrition Assistance Program (SNAP) is the largest nutrition assistance program in the United States. SNAP benefits are issued on a monthly basis, and recipients tend to redeem a larger share of benefits shortly after issuance, a pattern of behavior known as the SNAP benefit cycle. Previous research documenting the SNAP benefit cycle and its various impacts have often relied on cross-sectional survey data, with relatively small sample sizes, short periods of analysis, and self-reported SNAP participation and spending. In this study, I leverage a unique administrative panel dataset to investigate several aspects of the SNAP benefit cycle and uncover potential unobserved heterogeneity. The data include the universe of SNAP benefit transaction activity from 2011 to 2015 for 1.8 million unique households in Georgia, and detailed information regarding benefit receipt and redemption, including timing, dollar amount, and location of transactions. Using a panel data finite mixture model, I find evidence supporting the existence of two types of SNAP households, which I denote as slow and fast spenders. Fast spenders comprise the overwhelming majority of SNAP households, 77%, and spend approximately 70% of their benefit in the first week. In addition, fast spenders shop more frequently throughout the month, stay longer in the program, 3 years compared to 1 year and 4 months for slow spenders, and receive benefits that are almost twice as large, on average. The two groups show similar patterns of rural or urban residence, and the types of stores where they spend SNAP benefits. Finally, I also find that shopping frequency, tenure in SNAP, household SNAP benefit, and household store preferences are not strong predictors of household spending type. My results add to a more comprehensive understanding of the SNAP benefit cycle and inform welfare improving policies that promote smoother spending of SNAP benefits throughout the month.

In my second chapter, co-authored with Dr. Tchernis, Dr. Davis, and Dr. Gregory, we examine the measurement of food security. We propose a novel Bayesian Graded Response Model (BGRM) for food security measurement. Our BGRM has several attractive features. It produces continuous food security estimates and measures of estimation uncertainty at the household level. Unlike the USDA's official measurement model, the standard BGRM can be used with binary and polytomous items. Furthermore, we modify our BGRM to include any combination of binary, ordered polytomous, and continuous variables. With data from the 2017-18 National Health and Nutrition Examination Survey (NHANES), we estimate our BGRM for responses to the 10 adult core Food Security Module (FSM) questions. We find non-trivial uncertainty in household-level estimates, emphasizing the inherent uncertainty of latent trait estimation. We observe overlap in BGRM estimates across USDA-defined food security categories and significant within-category variation. We estimate our model using Current Population Survey (CPS) data as a robustness check, finding qualitatively similar results to those from the NHANES. We also use our continuous food security estimates to calculate three Foster, Greer, and Thorbecke (FGT) indices which capture the prevalence, depth, and severity of food insecurity in our NHANES sample. Finally, as a demonstration of the BGRM's flexibility, we include an additional continuous variable, household-level Monthly Food Spending (MFS), into the model, capturing economic access/affordability of food and food security information in a novel latent construct. The adaptability of our BGRM positions it as a versatile tool for measuring food security in certain applications, but may also prove useful for the measurement of other latent traits requiring a diverse range of variable types like nutrition security.

In my third chapter, I examine the relationship between teenage substance use and macroeconomic conditions. In particular, I examine how macroeconomic conditions affect teenage marijuana, cigarette and alcohol use. The findings are based on state and year fixed effects models, using data from the Youth Risk Behavior Surveillance System, spanning 1991-2019. In contrast to the recent literature, I find no statistically significant evidence of a relationship between economic conditions and teenage marijuana use. However, I find evidence of counter-cyclical teenage alcohol and cigarette use among females and 12 to 14 year olds. The results for teenage alcohol use are consistent with evidence in the literature.

Chapter I: Examining Heterogeneity in the Supplemental Nutrition Assistance Program Benefit Cycle: A Finite Mixture Approach

1.1 Introduction

The Supplemental Nutrition Assistance Program (SNAP) is the largest food assistance program in the United States disbursing nearly $114 billion in benefits in fiscal year (FY 2022) to over 41.2 million individuals (Leftin et al. (2024)). Previous studies examining the SNAP program, reveal that SNAP recipients spend a significant amount of their monthly benefit early in the benefit month and that benefits do not last until the next issuance date (Wilde and Ranney (2000); Shapiro (2005); Castner and Henke (2011); Todd (2015)), a pattern termed the SNAP benefit cycle. Abundant evidence suggests that the SNAP benefit cycle has significant nutritional consequences for participants. In particular, the SNAP cycle has been associated with low caloric intake during the last week of the benefit month (Wilde and Ranney (2000)), a decline in the daily caloric intake and the daily dollar value of food consumed (Shapiro (2005)), and a decline in dietary quality throughout the month (Tarasuk et al. (2007); Kharmats et al. (2014)). Further, research indicates that the SNAP cycle might also exacerbate the risk of food insecurity. Hamrick and Andrews (2016) find that SNAP recipients are more likely to have a day without eating at the end of the SNAP benefit month and Weinstein et al. (2009) find that low-income households have a significantly higher probability to report food insecurity at the end of the month as at the beginning.

A wide range of factors have been examined to explain the SNAP benefit cycle. Among these factors are institutional features of the program such as benefit size, and the timing of benefit disbursal (see Todd (2015); Beatty et al. (2019)), behavioral factors such as the time preferences of SNAP households (see Shapiro (2005); Smith et al. (2016)), the fungibility of cash income and SNAP benefit income (see Smith et al. (2016)), the shopping frequency of SNAP households (see Wilde and Ranney (2000)) and finally household demographic characteristics such as female headship, residence in the South, distance to store, and use of cash welfare (see Wilde and Ranney (2000)).

While the existence of the SNAP benefit cycle is well documented, there is relatively less research focused on examining SNAP spending behavior at the micro level. Further, many of the existing studies rely on survey datasets that employ small sample sizes, use relatively short periods of analysis, and rely on self-reported SNAP participation and SNAP spending leading to potentially biased estimates due to measurement error. Finally, much of the existing research, models SNAP recipient spending behavior as originating from a single distribution, potentially ignoring heterogeneity in the spending behavior of these households, that might help explain the nature and causes of the cycle and uncover potential policy levers to target it.

In this study, I analyze a novel SNAP administrative panel dataset covering the universe of SNAP participating households in Georgia. This data contains every SNAP transaction for every participant over 5 years (2011 - 2015) and includes the addresses of SNAP authorized stores. Relative to previous studies, this dataset provides several advantages. First, SNAP participation and SNAP transactions are administratively recorded, freeing my estimates from issues caused by measurement error in these variables. Second, the extensive sample and panel nature of the data allows me to incorporate household specific variation in spending behavior in my analysis. Finally, the data offers great detail on the store choice of SNAP participants allowing for an estimation of the effect of these choices on the cycle.

Leveraging this data, I examine the SNAP benefit cycle, and investigate potential variables that correlate with the cycle. In particular, I investigate the relationship between the cycle and SNAP benefit size, store choice, shopping behavior, urban or rural residence, benefit saving behavior, and length of participation in SNAP. To examine heterogeneity in the cycle, I use a panel data finite mixture model[1] that allows different groups of SNAP recipients to display distinct patterns of spending throughout the benefit month. I use the empirical estimates to identify which group of SNAP households might be responsible for the SNAP benefit cycle, and

[1]Mixture models are probabilistic models that assume the data are generated by several different unobserved processes and probabilistically assign each observation to each process or group (Dempster et al. (1977)). These models have been widely used in a number of fields, including labor economics (Heckman et al. (1990); Gritz (1993)), marketing (Wedel et al. (1993)), industrial organization (Wang et al. (1998)), and health economics (Deb and Trivedi (1997)). For a survey of the statistical literature on these models, see McLachlan (2000).

further to identify household characteristics and behaviors that correlate with the cycle differentially among groups.

The mixture results produce several important insights. I discover a clear classification of SNAP households into two groups - designated as slow and fast spenders - that display starkly different patterns of spending of their monthly SNAP benefits. The overwhelming majority of SNAP households, 77%, are classified as fast spenders. The smaller group of slow spenders spends approximately a tenth of their benefit within the first 7 days, while fast spenders spend a little more than two thirds of their benefit. Additionally, I detect correlations between a number of household variables and the SNAP cycle. First, a higher household benefit leads to greater spending, within the first 7 days, with the effect being more pronounced for faster spending households. This finding is noteworthy given that fast spending households receive a larger benefit, on average, compared to slow households, and spend a larger percentage of their benefit throughout the month, as well. Second, completing additional monthly transactions during the benefit month, leads to greater spending, within the first 7 days, for slow households, while reducing spending, during the same period, for fast spending households. These results are consistent with the fact that slow households conduct significantly less transactions both during the first 7 days and the full benefit month, compared to fast households.

Third, bulk shopping[2] is negatively associated with spending, during the first 7 days of the benefit month, for slow households, while being positively associated with spending for fast spending households. Slow households are more likely to bulk shop compared to fast spending households, but given their lower spending throughout the first 7 days, it seems evident that slow households bulk shop later in the month. Fourth, household residence does not play a prominent role in the SNAP benefit cycle for slow households, but small differences in the cycle are detected for fast households. Fifth, both groups share similar store type preferences, on average, with the preferred store for shopping, by a wide margin, being supermarkets. Consistent with this behavior, the estimated effects of shopping at the remaining store categories on the SNAP benefit

[2]A bulk shopping episode occurs when a household completes at least one transaction, during a given benefit month, where it spends 40% or more of its benefit.

cycle, are statistically significant but quite small in magnitude, across both groups. Finally, length of participation in SNAP is positively correlated with the SNAP benefit cycle, with the average tenure of the fast spending households being significantly higher compared to slow households.

The remainder of this paper is organized as follows. Section 2 describes the background and relevant literature. Section 3 describes the data. Section 4 presents preliminary descriptive evidence on the SNAP cycle. Section 5 describes the methodology. Section 6 discusses the results. Section 7 concludes.

1.2 Background

1.2.1 SNAP Institutional Background

The Supplemental Nutrition Assistance Program (SNAP) stated goal is to "permit low-income households to obtain a more nutritious diet...by increasing their purchasing power" (United States Department of Agriculture, (2022)). To supplement a household's income, SNAP provides benefits via debit-like Electronic Benefits Transfer (EBT) cards that can be used to buy food in authorized retail food outlets.

To participate in SNAP, households must meet certain eligibility standards. First, a household's gross income generally cannot exceed 130 percent of the Federal poverty guidelines ($2,871 per month for a family of four in the United States in FY 2022). Additionally, a household's net income, gross income after certain deductions, cannot exceed 100 percent of the poverty guidelines ($2,209 per month for a family of four in the United States in FY 2022). Further, SNAP eligibility is also limited based on the household's resources. In FY 2022, households were permitted up to $2,500 in countable resources, whereas households with at least one elderly household member were allowed up to $3,750 in countable resources (United States Department of Agriculture, (2022)).[3]

Once a household is deemed eligible, it is assigned a certification period. Certification periods vary widely, ranging from 3 months to as much as 48 months. These periods are typically based on the likelihood of a change in the SNAP household's financial circumstances. Further,

[3] SNAP has a series of exceptions in eligibility requirements that apply to certain households. For a detailed discussion of SNAP eligibility requirements, see United States Department of Agriculture, (2022).

SNAP requires that most households be interviewed at least once every 12 months for recertification purposes. In FY 2022, the average certification period was 16 months (United States Department of Agriculture, (2022)).

SNAP maximum benefit amounts are defined using the Thrifty Food Plan (TFP), a low cost nutritionally adequate food plan that varies by household size and composition (Gundersen et al. (2019)).[4] Household SNAP benefit levels are determined by subtracting 30% of the household's net income from the value of the TFP. Thus, SNAP benefits do not cover a household's full food budget, except when a household has zero net income, in which case, the household receives the TFP maximum benefit. In FY 2022, the average SNAP household received a monthly benefit of $297 (United States Department of Agriculture, (2022)). It is estimated that food-at-home expenditures account for 74% of SNAP households total food expenditures. SNAP benefits account for over 60% of the average food-at-home expenditures of SNAP households. This figure rises to 80% for SNAP households in deep poverty (household income below 50% of the Federal Poverty Guidelines) (Tiehen et al. (2017)).

All states and territories provide SNAP benefits once monthly through EBT cards. In some states, SNAP benefits are issued on the same day for all participating households, typically on the first of the month. In contrast, other states stagger their benefit disbursal over multiple days. For example, in Georgia, benefit disbursal dates range from the 5th to the 23rd of the month.[5] There is considerable variation in these dates across states. Further, if unspent, SNAP benefit can carry over, for up to 9 months post disbursal. Any benefit that is not spent within 9 months of disbursal, is permanently removed from the account (United States Department of Agriculture, (2024)).

1.2.2 The SNAP Cycle: Conceptual Background

The life-cycle or permanent income hypothesis is a well established notion in current economic theory. This theory assumes that households' expectations of lifetime income will have a greater influence on their present decision-making than their current income (Friedman (1957); Hall (1978)). Given that consumers are rational, have perfect information and access to credit,

[4]For more details regarding the Thrifty Food Plan, see Thrifty food Plan.
[5]For more details regarding SNAP benefit disbursal dates, see Benefit Disbursal Dates.

they can always borrow against future income to conduct consumption decisions in the present. A testable prediction from this theory is that consumption should be consistent with average lifetime income since individuals can borrow against future income and save when income is highest to maintain this level of consumption. However, empirical evidence indicates that in many cases individuals do not smooth their consumption over their lifetimes, and even shorter periods. This departure could be due to market imperfections or a variety of other economic causes. In the case of the SNAP benefit cycle, several factors have been proposed to justify the lack of spending and consumption smoothing. I discuss some of these factors in turn.

Perhaps the most evident of these factors is the size of SNAP benefits. SNAP maximum benefit amounts are defined using the Thrifty Food Plan (TFP), a low cost nutritionally adequate food plan that varies by household size and composition (Gundersen et al. (2019)). Household SNAP benefit levels are determined by subtracting 30% of the household's net income from the value of the TFP. Thus, SNAP benefits do not cover a household's full food budget, except when a household has zero net income, in which case, the household receives the TFP maximum benefit. Evidence suggests however that SNAP benefits are generally not sufficient to meet the needs of SNAP recipients. The general argument posits that the TFP benefit design is not adequate as it fails to account for the time needed to prepare meals at home, does not properly account for the dietary needs of adolescent household members, and does not address geographic variations in the cost of goods (Yaktine and Caswell (2013), Ziliak (2016), Gundersen et al. (2019)). Relatedly, Todd (2015) examines the role of benefit size in the SNAP benefit cycle by exploiting an increase in SNAP benefits due to the American Recovery and Reinvestment Act (ARRA). Lending credence to this argument, Todd (2015) finds that before ARRA, SNAP households experienced a decline of 25% in caloric intake during the fourth week of the month, a decline which is no longer present after the implementation of ARRA. Todd (2015) concludes that an increase in SNAP benefits may help smooth food intake during the benefit month for SNAP recipient households.

Besides the generosity of SNAP benefits, paycheck cycles, and household receipt of other means-tested cash welfare or social insurance programs, have also been examined as potentially affecting the benefit cycle. Other sources of income could mitigate the SNAP cycle by increasing resources available to SNAP households. Further as wages increase, they would eventually reduce the value of the SNAP entitlement, therefore reducing SNAP spending in total and as a share of overall food spending, further reducing the magnitude of the SNAP cycle. On the other hand, the overlap between the timing of SNAP benefit issuance and receipt of social insurance and welfare payments, or paychecks, could lead to an overestimation of the SNAP cycle (Beatty et al. (2019)). Using FoodAPS data, Beatty et al. (2019) examine the interactions between the SNAP, paycheck and other benefit cycles. The authors find only modest evidence that the SNAP cycle is affected by the timing of other sources of income, be it wage income or income from other government programs. In particular, they find evidence that the cycle is more pronounced for workers who are paid on a weekly or monthly basis, but little evidence of a correlation between the SNAP cycle and cycles in other income streams. The authors conclude that other sources of income neither mitigate nor exacerbate the SNAP cycle.

There is evidence that the SNAP benefit cycle is partially due to the time inconsistent preferences of SNAP recipient households (Shapiro (2005), Mastrobuoni and Weinberg (2009)). Based on Laibson (1997) theory of hyperbolic discounting, time inconsistent households display short-run impatience, preferring present consumption to future consumption. These households are more likely to spend when resources are abundant, which leads to a fast depletion of financial resources and higher susceptibility to experiencing negative income shocks and food insecurity later in the month. Shapiro (2005) finds that SNAP recipients' spending and food intake behavior is consistent with hyperbolic discounting. SNAP households' caloric intake and the dollar value of food consumed declines throughout the benefit month, implying a shift from higher quality/cost to lower quality/cost foods at the end of the benefit month. Relatedly, Smith et al. (2016) find that SNAP households spend approximately 96 cents of every food dollar on food at home on the day that benefits are issued. By the end of the month, just over three quarters of the household food budget is spent on food at home.

An additional factor that has been linked to the SNAP benefit cycle is income fungibility, referring to the degree of substitutability between SNAP income and income from other sources. Previous literature finds that SNAP households have a higher marginal propensity (MPS) to spend on food from benefit income compared to cash income (Fraker et al. (1995); Breunig and Dasgupta (2002)). Laibson (1998) demonstrates that the monthly consumption path of an impatient household depends on its propensity to spend out of existing resources. Given that impatient SNAP households spend relatively more of their benefit upon receipt, the higher MPS on food out of benefit income could exacerbate this effect. Smith et al. (2016) examine the role of income fungibility and time inconsistent preferences on the SNAP benefit cycle and find that income fungibility and time-inconsistent preferences work in tandem with SNAP households spending 9.3% more on food at home out of SNAP than out of non-SNAP income on the day of benefit arrival.

Finally, transportation difficulties, time constraints, a restrictive local food environment, or stigma could also be responsible for the SNAP benefit cycle. In particular, each of these factors could lead to less frequent shopping by SNAP households, resulting in large spending at the beginning of the month and trouble with storing food for consumption in later weeks as a consequence. Examining this possibility, Wilde and Ranney (2000) propose a theory that households weigh the advantages of frequent shopping (less food spoilage, less need for smaller trips to closer, higher-priced stores toward the end of the month) against the disadvantages of frequent major grocery trips (stigma, loss of leisure time, and so on). Using an endogenous switching regression model where the consumer simultaneously chooses either a frequent or infrequent shopping regime and food intake levels in each half of the month, the authors find significant differences in the food intake for each regime. Specifically, for households that shop frequently, food intake is quite smooth over the food stamp month, while for households that shop infrequently, food energy intake declines significantly throughout the course of the food stamp month. Further, the authors find that total income, measured as food stamp benefits plus cash income, has no measurable marginal effect on food energy intake or the probability of shopping frequently. However, demographic factors and geographical factors affect both food energy intake and the probability of shopping frequently. In particular, participation in the Special Supplemental Nutrition Assistance Program for Women, Infants, and Children (WIC) is found to have a strong positive effect on food energy intake. Urban residence, receiving cash welfare, female headship, residence in the South, and distance to "major" grocery store all reduce the probability of shopping frequently.

The SNAP literature has extensively explored the nutritional consequences of the SNAP cycle and the myriad factors that might be responsible for this pattern of spending. A less explored area of research is the potential unobserved heterogeneity that might exist across SNAP households in response to SNAP receipt. To the best of my knowledge, the only study that directly addresses this question is Dorfman et al. (2019). Using a sample of 163 households from FoodAPS, and a finite mixture model, the authors uncover two groups of SNAP recipient households, which they denote as the patient and impatient groups. Out of 163 households, they find 63 households in the impatient group (39% of the sample) and 100 households in the patient group (61% of the sample). Additionally, the authors find that impatient households spend 4 times as much of their monthly benefit within the first 4 days of the month (67.2% of their benefit) as the patient households (17.7% of their benefit).

This paper provides several contributions to the literature. First, it makes use of a novel restricted administrative panel dataset of SNAP recipients which avoids common issues that affect past studies, which often rely on survey or commercially available data, such as misreported SNAP participation (Kreider et al. (2012)), misreported SNAP expenditure information, and undersampling of disadvantaged populations (National Academies of Sciences, Engineering and Medicine (2020)). Second, I re-examine the SNAP benefit cycle and investigate several variables that might correlate with the cycle such as: household benefit, household urban or rural residence, household store choice, household shopping behavior and finally household tenure in SNAP. I also explore observed heterogeneity in the cycle by household benefit size, tenure in SNAP and household benefit saving behavior. These results add to the extensive literature examining the myriad factors that correlate with the cycle. Finally I directly examine unobserved heterogeneity in the cycle. Using a panel data finite mixture model, that optimally separates SNAP households based on their spending behavior, I investigate the existence of different groups of SNAP households who might display significantly different spending patterns. Additionally, I use the mixture results to investigate household specific variables that might have differential effects on the spending behavior of these households. These results contribute to the more scarce literature examining heterogeneity in the SNAP benefit cycle. Further, to the best of my knowledge, this paper is the first in this literature to study heterogeneity in the cycle while explicitly accounting for household specific temporal variation in the cycle through a panel data finite mixture model.

1.3 Data

1.3.1 Households

I use administrative data containing the universe of SNAP transactions for all SNAP households in Georgia between January 2011 and December 2015. This data consists of approximately 400 million transactions from 1,848,137 unique participant households. For each transaction I observe the amount, the date, and the store name and address, but do not observe which items are purchased. For tractability, I collapse transaction-level data to a benefit month summary of transactions for each household. A benefit month for a given household is the period starting with the date of benefit receipt and ending the date before the next issuance. A typical benefit month lasts 31 days, as SNAP households typically receive benefits on the same calendar day during different months. Georgia SNAP households receive benefits at different dates during a given month, thus benefit month dates are household specific. The period of analysis ranges from January 2011 to December 2015, resulting in 60 benefit months in the data. Finally, I restrict the analysis sample to include only households whose benefit amounts received in any benefit month are greater or equal to \$15 and less or equal to \$1500.[6] The final analytical sample contains 34,770,457 benefit months generated by 1,555,674 households.

[6] These cutoffs are consistent with the minimum and maximum allotments provided by the food and nutrition service at Food and Nutrition Service.

This dataset has several advantages that make it well suited to answering my research questions. First, the data covers the universe of SNAP participating households in Georgia. Thus SNAP participation is confirmed by administrative records, freeing my results from potential bias due to misclassification of SNAP participation. Further, my results are generalizable to the population of SNAP participants in Georgia and similar populations in other states. The second important feature of this data is that SNAP spending behavior is recorded electronically through administrative records, preventing errors related to self-reported spending behavior, and allowing for accurate and consistent tracking of household spending behavior over time. The third important feature of this data is the detailed information on stores where SNAP benefits are spent. This data allows for accurate tracking of store choice by SNAP households. Finally, the panel nature of this data allows me to examine household spending behavior over time, and investigate the role of length of participation in SNAP on the benefit cycle.

While the administrative data offer rich information on the benefit redemption patterns of SNAP households in Georgia, they are limited in other ways. First, the data do not contain information on the specific items that households purchase. Therefore I cannot examine the types of food purchased or the nutritional value of food consumed. Second, I do not observe any demographic information related to the households. Therefore I cannot explore the role of certain household characteristics such as household address, family size, income, race, among other variables, on SNAP household benefit spending behavior. Finally, the data contain only SNAP redemptions. Given that SNAP transactions constitute approximately 50% of SNAP household total food spending on average (Tiehen et al. (2017)), the findings apply to a significant portion, but not the entirety, of household food spending.

1.3.2 Stores

One advantage of this dataset is its rich set of information regarding stores where households spend SNAP benefits. The data contain information about approximately 15,000 individual stores in Georgia where SNAP benefits are spent. For each store, I observe the name, street address and the dollar value of all SNAP transactions that occur at that store. However, I do not observe which foods households buy with their SNAP benefits. I rely on this information to examine the role of store choice on the SNAP benefit cycle. To accomplish this, I classify stores into 6 categories: supermarkets, grocery, convenience, combination, specialty foods and other. These categories are derived from the USDA Store Tracking and Redemption System (STARS).[7]

Supermarkets refers to stores that sell a wide variety of grocery and other store merchandise, and have multiple checkout lanes and registers. This collapsed category includes STARS supermarkets, superstores and military commissaries. Grocery stores refers to stores that carry a small, moderate, or wide selection of all four staple food categories, with their primary stock being food items. This collapsed category includes STARS small, medium and large grocery stores. Convenience refers to self-service stores that offer a limited line of convenience items and primarily sell a variety of canned goods, dairy products, pre-packaged meats and other grocery items in limited amounts. This collapsed category is identical to the STARS convenience category. Combination refers to stores where the primary business is the sale of general merchandise but also sell a variety of food products. This category includes dollar stores, drug stores, and general stores. This category is identical to the STARS combination grocery/other stores category. Specialty foods refers to stores that operate as a cooperative or specialize in the sale of specific products such as bread/cereal products, fruits and/or vegetable products, meat products, and seafood products. Finally, Other refers to all remaining stores in the data that do not fit neatly into any of the above categories.[8]

1.3.3 Summary Statistics

Table 1 reports summary statistics for the analytical sample. The sample consists of benefit months generated by Georgia SNAP households from January 2011 to December 2015. There are 34,770,457 benefit months generated by 1,555,674 households. On average, SNAP households spend approximately 60% of their SNAP benefits within the first week of the benefit month, a strong indicator of the benefit cycle. SNAP households spend on average approximately

[7]For more information regarding the STARS classification system, see SNAP Store Type Definitions.

[8]For a detailed, yet non-exhaustive, list of the individual stores that belong to the six collapsed categories, see Table A1.

Table 1. Summary Statistics 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households

	Mean	*Standard Deviation*
Amount Spent		
$ Within 7 days	167.47	150.23
% Within 7 days	59.72	33.54
$ Within Month	261.69	188.05
% Within Month	91.55	18.44
Transactions & Benefit Amount		
# of Transactions (7 days)	4.31	3.68
# of Transactions (Month)	8.76	6.87
Bulk Shopping (%)	49.26	49.99
Benefit Amount ($)	287.80	198.00
SNAP Tenure		
Tenure (Ben. Months)	35.80	15.80
Household Location		
Rural	16.87	37.45
Urban	75.88	42.78
Mixed	7.26	25.94
Store Choice ($)		
Supermarkets ($)	215.83	171.77
Grocery ($)	12.00	41.42
Convenience ($)	13.75	36.09
Combination ($)	13.40	32.80
Specialty Foods ($)	3.27	17.65
Other ($)	3.44	24.57
Store Choice (%)		
Supermarkets (%)	74.71	28.65
Grocery (%)	4.19	13.47
Convenience (%)	5.00	12.41
Combination (%)	5.37	12.86
Specialty Foods (%)	1.12	5.89
Other (%)	1.16	7.67
Households	1,555,674	
Benefit Months	34,770,457	

Notes: The sample consists of benefit months generated by the universe of SNAP participating households in Georgia during the period of January 2011 to December 2015. The entire sample contains 34,770,457 benefit months, generated by 1,555,674 households. Store choice refers to the $'s or % of benefit spent by households on average, at each store category, during a given benefit month.

92% of their benefit during a given benefit month.[9] On average, households complete 4.31 transactions within the first 7 days and approximately twice as many, 8.76 transactions, during the full benefit month. This pattern of behavior indicates that households shop quite frequently during the first 7 days of the benefit month, consistent with their significant spending during the same period. Regarding stores where households choose to shop, supermarkets are the preferred choice. On average, households spend 74.71% of their monthly benefit at supermarkets, 5.37% at combination stores, 5% at convenience stores, 4.19% at grocery stores, 1.12% at specialty food stores and 1.16% at other stores.

To supplement my analysis, I define a number of additional variables. Tenure represents the number of benefit months a household spends in SNAP. Bulk shopping is a binary indicator that takes value 1 if a household has at least one transaction that amounts to 40% or more of their benefit on a given benefit month. This variable aims to identify households who prefer to conduct few high value shopping transactions throughout the benefit month. Bulk shopping could serve as an indicator of several household circumstances, such as a preference for buying durable foods in bulk for storage, or an unwillingness to conduct multiple shopping trips throughout the benefit month due to, transportation difficulties, stigma associated with SNAP benefit use, among other factors. Further this variable allows me to investigate whether frequent shopping households have a less pronounced benefit cycle compared to infrequent (bulk shopping) households. To construct a proxy measure of household urban or rural residence, I rely on store zip codes. Using the Rural-Urban Commuting Area (RUCA) codes[10], I classify stores as either rural or urban using their zip codes. RUCA codes use measures of population density, daily commuting, and urbanization, to classify U.S. census tracts broadly as metropolitan, micropolitan, small town, and rural commuting areas. These classifications can also be applied at the zip code level. RUCA codes have been widely used by various health researchers and are also currently used as the basis

[9]SNAP benefits can be rolled over, and occasionally households do not spend the full benefit they receive on a given benefit month.

[10]For access to the documentation and RUCA codes datasets used in the analysis, see Economic Research Service RUCA Codes.

for eligibility for many federal programs.[11] I rely on standard classifications of RUCA codes to identify rural and urban zip codes.[12]

In a given benefit month, I observe whether households complete at least two thirds of their transactions at rural or urban stores. If a household completes at least two thirds of their transactions at urban stores, during a given benefit month, then it is considered an urban household during that month. If instead, it completes two thirds of its transactions at rural stores, it is considered a rural household during the particular benefit month. If the household does not display a preference for shopping at urban or rural stores, during a given benefit month, then it is classified as a mixed household during that month.

On average, households spend approximately 36 benefit months (3 years) in the program. In about half of all benefit months, households conduct at least one transaction where they spend 40% or more of their benefit. Finally, households are classified as urban households in 75.88% of benefit months, rural in 16.87% of benefit months, and mixed in 7.26% of benefit months.

1.4 SNAP Cycle by Benefit Size and Other Characteristics

I begin my analysis with simple descriptive tables to examine the SNAP cycle, identify potential variables that correlate with the cycle, and investigate heterogeneity by benefit size and tenure in the program.

Table 2 reports the results for the main (analytical) sample and by benefit size categories. I define three categories. Low benefit households which are households whose maximum benefit received in any benefit month is less than $100, moderate benefit households whose benefit received in any given benefit month is between $100 and $250 inclusive, and finally, high benefit households whose minimum benefit received in any benefit month is greater than $250.[13]

Perhaps, unsurprisingly, there is a positive association between the amount of dollars spent

[11] For additional information on the origin and use of RUCA codes see Rural Health Research Center.

[12] In particular, I use the rural and urban classification definition listed as **categorization c**, provided by the Rural Health Research Center at the University of Washington, which originally developed the RUCA classification system jointly with the U.S. Department of Agriculture Economic Research Service (ERS) and the Federal Office of Rural Health Policy (FORHP). For detailed information on the available urban and rural classification systems, including **categorization c**, see Rural Health Research Center.

[13] These categories are mutually exclusive but not collectively exhaustive.

Table 2. Summary Statistics 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main sample and by Benefit Size

	Main Sample	*BA*<$*100*	*BA*∈*[*$*100,*$*250]*	*BA*>$*250*
Amount Spent				
$ Within 7 days	167.47	21.97	109.96	295.14
	(150.23)	(23.03)	(66.71)	(187.85)
% Within 7 days	59.72	58.77	59.92	56.04
	(33.54)	(42.25)	(35.02)	(29.45)
$ Within Month	261.69	33.79	167.63	476.85
	(188.05)	(23.71)	(43.73)	(183.95)
% Within Month	91.55	93.38	91.47	90.69
	(18.44)	(15.99)	(19.21)	(18.51)
Transactions & Benefit Amount				
# of Transactions (7 days)	4.31	1.39	3.65	5.90
	(3.68)	(1.51)	(3.14)	(4.20)
# of Transactions (Month)	8.76	2.47	7.07	12.90
	(6.87)	(2.10)	(5.20)	(7.70)
Bulk Shopping (%)	49.26	85.42	55.82	29.20
	(49.99)	(35.29)	(49.66)	(45.47)
Benefit Amount ($)	287.80	36.35	183.30	527.68
	(198.00)	(24.36)	(27.80)	(176.14)
SNAP Tenure				
Tenure (Ben. Months)	35.80	32.66	27.43	36.68
	(15.80)	(17.42)	(17.75)	(16.54)
Households	1,555,674	104,109	332,444	312,193
Benefit Months	34,770,457	1,730,014	4,279,901	6,419,025

Note: Main sample and breakdown by benefit size. BA represents benefit amount. Three groups are displayed: households whose maximum benefit amount received in any benefit month is less than $100, minimum of $100 and maximum of $250 inclusive, and households whose minimum benefit amount received in any benefit month is greater than $250. The benefit categories are mutually exclusive, but are not collectively exhaustive. For example, households who receive a minimum of $100 and a maximum benefit greater than $250 are not represented in these categories.

Table 2 (Cont'd): Summary Statistics 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main sample and by Benefit Size

	Main Sample	***BA<$100***	***BA∈[$100,$250]***	***BA>$250***
Household Location				
Rural	16.87	24.73	14.83	14.88
	(37.45)	(43.14)	(35.54)	(35.59)
Urban	75.88	67.67	77.75	78.46
	(42.78)	(46.77)	(41.59)	(41.11)
Mixed	7.26	7.60	7.42	6.66
	(25.94)	(26.50)	(26.21)	(24.94)
Store Choice				
Supermarkets ($)	215.83	27.13	135.54	396.01
	(171.77)	(23.54)	(58.26)	(184.56)
Grocery ($)	12.00	1.73	7.40	23.41
	(41.42)	(7.73)	(25.20)	(64.59)
Convenience ($)	13.75	1.42	10.89	22.86
	(36.09)	(5.81)	(26.40)	(52.82)
Combination ($)	13.40	2.74	9.38	22.04
	(32.80)	(7.94)	(22.24)	(47.59)
Specialty Foods ($)	3.27	0.41	2.20	5.74
	(17.65)	(3.23)	(11.57)	(26.36)
Other ($)	3.44	0.37	2.22	6.80
	(24.57)	(3.71)	(15.31)	(38.87)
Households	1,555,674	104,109	332,444	312,193
Benefit Months	34,770,457	1,730,014	4,279,901	6,419,025

Note: Main sample and breakdown by benefit size. BA represents benefit amount. Three groups are displayed: households whose maximum benefit amount received in any benefit month is less than $100, minimum of $100 and maximum of $250 inclusive, and households whose minimum benefit amount received in any benefit month is greater than $250. The benefit categories are mutually exclusive, but are not collectively exhaustive. For example, households who receive a minimum of $100 and a maximum benefit greater than $250 are not represented in these categories.

within the first 7 days of the benefit month and the benefit received by households. Low benefit households spend \$21.97 within the first 7 days, followed by \$109.96 for moderate benefit households and \$295.14 for high benefit households. However, the results also reveal that the three groups spend approximately the same percentage of their benefit within the first 7 days. Low benefit households spend 58.77% of their benefit within the first 7 days, followed by 59.92% for moderate benefit households and 56.04% for high benefit households. These results seem to indicate that the SNAP benefit cycle stays fairly consistent across household groups who receive different benefit amounts. Regarding their spending behavior, during the full benefit month, low benefit households spend 93.38% of their benefit, followed by 91.47% for moderate benefit households and 90.69% for high benefit households.

A few additional differences are apparent across the three groups. There is a positive association between benefit size and the number of transactions a household completes, both during the first 7 days and throughout the full benefit month. There is a pronounced negative association between bulk shopping and benefit size. Low benefit households are more likely to reside in rural areas compared to moderate and high benefit receipt households. Finally store type preferences seem fairly consistent across all three household groups with supermarkets being the preferred store followed by either combination, grocery, or convenience stores and then the remaining store categories.[14]

In Table 3, I examine the SNAP benefit cycle for the main sample and by tenure in SNAP. I define three groups of SNAP households. Low tenure households with tenure less or equal to 12 benefit months (1 year or less), moderate tenure households with tenure between 12 and 24 benefit months (1 year to 2 years non-inclusive), and high tenure households with tenure greater or equal to 24 benefit months (2 years or more). There is a positive association between tenure and the amount of dollars spent within the first 7 days of the benefit month. Low tenure households spend \$119.65 within the first 7 days, followed by \$140.13 for moderate tenure households and \$180.02 for high tenure households. Further, the results indicate that there is a

[14] All differences between the benefit size category groups are statistically significant at the 1% level. For the results of significance tests of differences of means see Appendix Table A2.

positive association between tenure and the percent of SNAP benefit spent within the first 7 days. Low tenure households spend 51.53% within the first 7 days, followed by 56.49% for moderate tenure households and 61.58% for high tenure households. These results indicate that the SNAP benefit cycle might be more prominent for households who spend a longer period of time in the program. Regarding their spending behavior during the full benefit month, low tenure households spend 90.50% of their benefit during the full benefit month, followed by 91.39% for moderate tenure households and 91.74% for high tenure households. Additionally, low tenure households receive the lowest benefit, on average, \$246.06, followed by moderate tenure households, \$260.19, and finally high tenure households, \$299.49.

A few additional differences are apparent across the three groups. There is a positive association between tenure and the number of transactions a household completes, both during the first 7 days of the benefit month and throughout the full month. There is a slight negative association between bulk shopping and tenure. Low tenure households are more likely to reside in urban areas compared to moderate and high tenure households. Finally store type preferences seem fairly consistent across all three household groups with supermarkets being the preferred store followed by convenience, combination, grocery stores and finally the remaining store categories.[15] Next I systematically assess the SNAP benefit cycle through regression analyses as well as a panel data finite mixture model.

1.5 Empirical Strategy

1.5.1 Regression Analysis

My main research design models household i's log of SNAP dollars spent within the first 7 days as follows:

$$log7_{it} = \alpha_i + \boldsymbol{\beta}\mathbf{x}_{it} + \epsilon_{it} \tag{1}$$

$$\alpha_i = \psi + \boldsymbol{\gamma}\bar{\mathbf{x}}_i + \eta_i \tag{2}$$

[15] All differences between the tenure category groups are statistically significant at the 1% level. For the results of significance tests of differences of means see Appendix Table A3.

Table 3. Summary Statistics 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main sample and by Tenure

	Main Sample	***T≤12***	***T∈(12,24)***	***T≥24***
Amount Spent				
$ Within 7 days	167.47	119.65	140.13	180.02
	(150.23)	(115.03)	(123.11)	(157.43)
% Within 7 days	59.72	51.53	56.49	61.58
	(33.54)	(36.22)	(34.39)	(32.71)
$ Within Month	261.69	220.00	235.96	273.01
	(188.05)	(156.61)	(164.16)	(195.26)
% Within Month	91.55	90.50	91.39	91.74
	(18.44)	(19.61)	(18.60)	(18.23)
Transactions & Benefit Amount				
# of Transactions (7 days)	4.31	3.33	3.85	4.55
	(3.68)	(3.24)	(3.39)	(3.77)
# of Transactions (Month)	8.76	7.72	8.30	9.01
	(6.87)	(6.20)	(6.53)	(7.01)
Bulk Shopping (%)	49.26	51.70	50.28	48.69
	(49.99)	(49.97)	(50.00)	(49.98)
Benefit Amount ($)	287.80	246.06	260.19	299.49
	(198.00)	(167.97)	(173.42)	(205.19)
SNAP Tenure				
Tenure (Ben. Months)	35.80	7.93	18.28	43.42
	(15.80)	(3.02)	(3.08)	(9.90)
Households	1,555,674	636,521	287,162	631,991
Benefit Months	34,770,457	3,861,033	5,094,561	25,814,863

Note: Main sample and breakdown by length of participation in SNAP. Length of participation in SNAP is measured by tenure (T), the number of benefit months a households spends in the program. Three groups are displayed: households whose tenure in SNAP is less or equal to 12 benefit months (1 year or less), between 12 and 24 benefit months exclusive (more than a 1 year and less than 2 years), and households whose tenure is greater or equal to 24 benefit months (2 years or more).

Table 3 (Cont'd): Summary Statistics 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and by Tenure

	Main Sample	$T \leq 12$	$T \in (12,24)$	$T \geq 24$
Household Location				
Rural	16.87	13.24	14.57	17.86
	(37.45)	(33.89)	(35.28)	(38.30)
Urban	75.88	79.44	78.16	74.90
	(42.78)	(40.41)	(41.32)	(43.36)
Mixed	7.26	7.32	7.27	7.24
	(25.94)	(26.04)	(25.97)	(25.92)
Store Choice				
Supermarkets ($)	215.83	186.76	198.60	223.58
	(171.77)	(148.83)	(154.71)	(177.37)
Grocery ($)	12.00	7.32	8.64	13.36
	(41.42)	(30.64)	(33.39)	(44.08)
Convenience ($)	13.75	10.76	11.92	14.56
	(36.09)	(29.68)	(31.58)	(37.72)
Combination ($)	13.40	10.00	11.28	14.33
	(32.80)	(27.31)	(29.03)	(34.17)
Specialty Foods ($)	3.27	2.06	2.38	3.63
	(17.65)	(12.95)	(13.92)	(18.86)
Other ($)	3.44	3.11	3.14	3.55
	(24.57)	(23.21)	(23.14)	(25.03)
Households	1,555,674	636,521	287,162	631,991
Benefit Months	34,770,457	3,861,033	5,094,561	25,814,863

Note: Main sample and breakdown by length of participation in SNAP. Length of participation in SNAP is measured by tenure (T), the number of benefit months a households spends in the program. Three groups are displayed: households whose tenure in SNAP is less or equal to 12 benefit months (1 year or less), between 12 and 24 benefit months exclusive (more than a 1 year and less than 2 years), and households whose tenure is greater or equal to 24 benefit months (2 years or more).

In particular, equations 1 and 2 are combined to yield:

$$log7_{it} = \psi + \boldsymbol{\beta}\mathbf{x}_{it} + \boldsymbol{\gamma}\overline{\mathbf{x}}_i + \nu_{it} \tag{3}$$

where $log7_{it}$ represents the log of SNAP dollars spent within the first 7 days[16], for household i during benefit month t. The vector $\mathbf{x}_{it}$ includes the log of monthly benefit received by a given household; the number of transactions the household completes during a given benefit month; a binary indicator of whether the household completed at least one transaction where it spent 40% or more of its benefit during the benefit month; the log of dollars spent, during the benefit month, at grocery stores, convenience stores, combination stores, specialty food stores, or other stores; binary indicators for whether a household is a rural or mixed household; and a benefit month trend. The vector $\overline{\mathbf{x}}_i$ includes time averages of all the time varying covariates in $\mathbf{x}_{it}$. Regression errors are clustered at the household level.

Equation 3 illustrates the correlated random effects approach in panel data models (Mundlak (1978) ; Chamberlain (1982)). A competing approach, which is commonly used in panel data models, is the fixed effects approach. I argue that, in my setting, the correlated random effects approach is preferred due to three main reasons. First, due to the nature of my data, it is plausible that for a significant portion of SNAP households, key independent variables such as the log of total credit, household urban, rural or mixed residence, among others, will show little variation over extended periods of time. This lack of variation could be due to a lack of changes in households' financial circumstances, residence, or behaviors. A fixed effects approach would prevent me from estimating the effect of these variables which are key in my analysis. The second, and potentially more important reason rests on a now well-established technical result.

[16]In previous studies, the percentage of benefit spent within the the first few days (3 to 4) of the benefit month, has been explored as a dependent variable (for an example, see Dorfman et al. (2019)). In my analysis, an analogous counterpart to this variable is ratio, defined as the percentage of benefit spent within the first 7 days of the benefit month. However, using ratio as a dependent variable is complicated by the fact that ratio has significant point mass at zero and one, implying that multiple households often spend either none or all of their benefit within the first 7 days of the benefit month. Accommodating the nature of ratio in finite mixture models involves dealing with meaningful complications regarding estimation and identification of mixture subgroups, while not providing much added benefit, compared to a model that employs log of dollars spent within the first 7 days as the main dependent variable.

Specifically, given that I employ a linear regression model, the correlated random effects approach yields coefficients on the time varying variables that are identical to the ones recovered from a fixed effects regression employing the same time-varying variables.[17] Finally, one of the main goals of my analysis, is to examine unobserved heterogeneity in the SNAP benefit cycle, through a finite mixture model. As I will demonstrate in the next section, the correlated random effects approach leads to a simple and powerful adaptation of the finite mixture framework to panel datasets. A fixed effects approach however is not as easily implemented in the finite mixture framework.

1.5.2 Finite Mixture of Distributions

In order to explore unobserved heterogeneity in the SNAP benefit cycle, I apply a finite mixture model to the data. In this application, I use a normal finite mixture model which assumes that each group follows a normal linear regression model. My dependent variable is the log of SNAP benefit spent in the first 7 days of the benefit month. I expect to observe most of the variation in the spending patterns of SNAP recipients during this period, leading to accurate classification of these households into groups, conditional on their spending. In addition, I also control for the same set of variables as described in my main specification, given by Equation 3.

In the finite mixture model, the random variable of interest is assumed to be a draw from a population that is an additive mixture of distinct subpopulations or classes (c) in proportions p_c where $\sum_{c=1}^{C} p_c = 1, p_c > 0 \; \forall \; c = 1, 2, ..., C$. For the log of SNAP dollars spent within the first 7 days of the benefit month, $log7_{it}$, the mixture density can be described in general by

$$g(log7_{it}|\Theta) = p_1 g_1(log7_{it}|\Theta_1) + p_2 g_2(log7_{it}|\Theta_2) + \cdots + p_C g_C(log7_{it}|\Theta_C) \qquad (4)$$

where the class densities $g_c(log7_{it}|\Theta_c)$ are assumed to be normals, i.e.,

$$g_c(log7_{it}|\Theta_c) = \frac{1}{\sqrt{2\pi\sigma_c^2}} \exp\left(-\frac{1}{2\sigma_c^2}(log7_{it} - \psi - \boldsymbol{\beta}_c \mathbf{x}_{it} - \boldsymbol{\gamma}_c \bar{\mathbf{x}}_i)^2\right) \qquad (5)$$

[17] For discussion regarding this result, see Mundlak (1978); Wooldridge (2010).

The mixing probabilities, p_c, regression coefficients, γ_c and β_c, and the standard deviation parameters, σ_c, are allowed to vary across classes.

The model is estimated by maximum likelihood using the Expectation Maximization (EM) algorithm and the Newton-Raphson method. The mixing probabilities are jointly estimated with the class-specific regression coefficients and standard deviations.

1.6 Results

1.6.1 Regression Results

Main Sample. Table 4 presents the results of estimating Equation 3 for the main sample and by the benefit size categories used in the descriptive analysis in Section 1.4. In column 1, I report the results for the main sample. In columns 2, 3, and 4, I report the results for the three benefit size categories.

Table 4 provides several interesting results. Examining the results for the main sample only, I find that consistent with the descriptive results, there is a positive correlation between credit received by households and the amount of dollars spent within the first 7 days of the benefit month. In particular, a 1% increase in total household benefit leads to a 0.86% increase in the amount of benefit spent within the first 7 days of the benefit month. This result indicates that on average, for the main sample, increasing household benefit does not significantly reduce the magnitude of the SNAP benefit cycle.

Continuing with the main sample, completing an additional shopping transaction, during the benefit month, increases the amount of dollars spent within the first 7 days by 2.21%. This result is consistent with the descriptive finding that, on average, half of all benefit month transactions occur after the first 7 days of the benefit month. Bulk shopping, during a given benefit month, increases the amount of dollars spent within the first 7 days by 50.35%, indicating that potentially a significant amount of bulk shopping occurs during the first 7 days of the benefit month. Being a mixed household reduces the amount of dollars spent within the first 7 days by 1.17%, indicating that households who shop with similar frequency at urban and rural stores, spend slightly less of their benefit during the first 7 days of the benefit month, compared to households who shop primarily at urban stores.

Finally, a 1% increase in dollars spent at either a grocery, combination, specialty or other store, increases the amount of dollars spent within the first 7 days by 0.026%, 0.006%, 0.024% and 0.024% respectively. Conversely, a 1% increase in dollars spent at a convenience store reduces the amount of dollars spent within the first 7 days by 0.005%. The effects estimated across the store categories are quite small, unsurprisingly, as most of household spending occurs at the reference store category, supermarkets.

1.6.2 Heterogeneity by Benefit Size, Saving behavior & Tenure

In columns 2 to 4 in Table 4, I present the results for Equation 3 by benefit size category. Total household benefit remains positively correlated with spending within the first 7 days. In particular, a 1% increase in total household benefit leads to an increase in the amount of benefit spent within the first 7 days of the benefit month, of 0.83% for low benefit households, 1.49% for moderate benefit households, and 0.40% for high benefit households. The magnitude of these effects differ meaningfully among the three groups, potentially underlying differences in the SNAP cycle among them that warrant further exploration.

Completing an additional shopping transaction, during the benefit month, increases the amount of dollars spent within the first 7 days by 6.77%, 3.67%, and 2.01% for low, moderate and high benefit households respectively. Thus, additional monthly transactions have a more pronounced effect on spending within the first 7 days for lower benefit households. Bulk shopping during a given benefit month increases the amount of dollars spent within the first 7 days by 59.39%, 54.22%, and 43.06%, for low, moderate and high benefit households respectively, indicating that bulk shopping has a more pronounced effect on spending within the first 7 days for lower benefit households. Low benefit households, who reside in rural areas, spend 1.89% more dollars within the first 7 days, compared to low benefit urban resident households. Conversely, high benefit households, who reside in rural areas, spend 1.19% less dollars within the first 7 days, compared to high benefit urban resident households. Moderate and

Table 4. Regression Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Sample and Heterogeneity by Benefit Size

	(1) Main Sample	(2) BA<$100	(3) BA∈[$100,$250]	(4) BA>$250
Log Credit	0.8594***	0.8330***	1.4853***	0.4001***
	(0.0008)	(0.0054)	(0.0108)	(0.0039)
Transactions	0.0221***	0.0677***	0.0367***	0.0201***
	(0.0001)	(0.0010)	(0.0003)	(0.0001)
Bulk Shopping	0.5035***	0.5939***	0.5422***	0.4306***
	(0.0007)	(0.0038)	(0.0020)	(0.0015)
Log Grocery	0.0262***	0.0197***	0.0279***	0.0237***
	(0.0002)	(0.0020)	(0.0008)	(0.0004)
Log Convenience	-0.0046***	-0.0232***	-0.0088***	-0.0058***
	(0.0002)	(0.0021)	(0.0008)	(0.0005)
Log Combination	0.0058***	–0.0056***	0.0050***	0.0029***
	(0.0002)	(0.0014)	(0.0007)	(0.0004)
Log Specialty	0.0241***	0.0157***	0.0282***	0.0222***
	(0.0003)	(0.0028)	(0.0010)	(0.0005)
Log Other Stores	0.0242***	0.0180***	0.0302***	0.0249***
	(0.0004)	(0.0043)	(0.0015)	(0.0007)
Rural	-0.0033	0.0189**	0.0054	-0.0119**
	(0.0021)	(0.0085)	(0.0066)	(0.0051)
Mixed	-0.0117***	-0.0009	-0.0128***	-0.0214***
	(0.0014)	(0.0065)	(0.0044)	(0.0033)
Period	-0.0022***	-0.0023***	0.0008***	-0.0033***
	(0.0000)	(0.0001)	(0.0001)	(0.0000)
R-squared	0.3253	0.2110	0.0744	0.1377
Households	1,555,674	104,109	332,444	312,193
Benefit Months	34,770,457	1,730,014	4,279,901	6,419,025

Regressions use correlated random effects models. The dependent variable in all regressions is the log of dollars spent within the first 7 days of the benefit month. BA represents benefit amount. Four groups are displayed: (1) The main sample consisting of all households in the sample; (2) Households whose maximum benefit amount received in any benefit month is less than $100 ; (3) Households who receive a minimum benefit of $100 and maximum of $250 inclusive ; (4) Households whose minimum benefit amount received in any benefit month is greater than $250. The benefit categories are mutually exclusive, but are not collectively exhaustive. For example, households who receive a minimum of $100 and a maximum benefit greater than $250 are not represented in these categories. (*10%,**5%,***1%).

high benefit households who are classified as mixed area residents, spend 1.28% and 2.14% less dollars within the first 7 days, compared to similar benefit urban resident households. Finally, I examine the estimated effects across the store categories. These results are remarkably similar across low, moderate and high benefit households. Neither store category appears to play an outsized role in the SNAP benefit cycle for either of the three benefit groups. This result is consistent with the findings from the descriptive analysis where it was evident that the preferred store category among all three groups was supermarkets.

An important feature of the SNAP program is that SNAP benefits can rollover across different benefit months. Thus, households who spend a small amount of benefit during a given benefit month, could be choosing to accumulate benefits for spending in later benefit months. Since my main outcome measure only examines spending behavior within the first 7 days, it is unable to distinguish between households who spend nothing within the first 7 days, but spend most of their benefit in a given benefit month, and those who do not spend most or all of their benefit during the benefit month, and instead save it for future spending. Therefore, failing to account for household saving behavior could lead to an underestimation of the extent of the benefit cycle, as saving behavior could appear indistinguishable from budgeting behavior. To account for saving behavior, I re-estimate my main specification, Equation 3, using a sample that excludes any household who spends less than 75% of their benefit in any given benefit month. I present these results in Table 5.

Most findings discussed in Table 4 remain the same in Table 5. Average household benefit size remains positively correlated with spending within the first 7 days, and the magnitude of effects follow the same pattern as for the main sample benefit size categories. The results regarding transactions, bulk shopping, rural or mixed residence also remain qualitatively similar. Finally, no substantive differences are discovered regarding the choice of store where households shop and the effect of this choice on spending within the first 7 days. Overall, these results provide suggestive evidence that the saving behavior of households is not a major factor affecting the SNAP benefit cycle.

Table 5. Regression Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Sample and Heterogeneity by Benefit Size and Saving Behavior

	(1) Main Sample	(2) BA<$100	(3) BA∈[$100,$250]	(4) BA>$250
Log Credit	0.8594***	0.8495***	1.4725***	0.4460***
	(0.0008)	(0.0056)	(0.0114)	(0.0040)
Transactions	0.0221***	0.0595***	0.0289***	0.0144***
	(0.0001)	(0.0011)	(0.0003)	(0.0001)
Bulk Shopping	0.5035***	0.5684***	0.4887***	0.3829***
	(0.0007)	(0.0039)	(0.0021)	(0.0016)
Log Grocery	0.0262***	0.0181***	0.0246***	0.0205***
	(0.0002)	(0.0021)	(0.0009)	(0.0005)
Log Convenience	-0.0046***	-0.0219***	-0.0073***	-0.0055***
	(0.0002)	(0.0021)	(0.0009)	(0.0005)
Log Combination	0.0058***	-0.0054***	0.0027***	0.0002
	(0.0002)	(0.0015)	(0.0007)	(0.0004)
Log Specialty	0.0241***	0.0147***	0.0244***	0.0181***
	(0.0003)	(0.0029)	(0.0011)	(0.0005)
Log Other Stores	0.0242***	0.0164***	0.0259***	0.0212***
	(0.0004)	(0.0044)	(0.0016)	(0.0007)
Rural	-0.0033	0.0151*	0.0092	-0.0064
	(0.0021)	(0.0088)	(0.0070)	(0.0053)
Mixed	-0.0117***	-0.0049	-0.0158***	-0.0215***
	(0.0014)	(0.0067)	(0.0046)	(0.0034)
Period	-0.0022***	-0.0023***	0.0010***	-0.0028***
	(0.0000)	(0.0002)	(0.0001)	(0.0000)
R-squared	0.3253	0.2079	0.0590	0.1245
Households	1,555,674	93,336	278,554	251,610
Benefit Months	34,770,457	1,634,050	3,773,554	5,499,285

Regressions use correlated random effects models. The dependent variable in all regressions is the log of dollars spent within the first 7 days of the benefit month. BA represents benefit amount. Four groups are displayed: (1) The main sample consisting of all households in the sample; (2) Households whose maximum benefit amount received in any benefit month is less than $100 and who spend 75% or more of their benefit in any given benefit month ; (3) Households who receive a minimum benefit of $100 and maximum of $250 inclusive and who spend 75% or more of their benefit in any given benefit month ; (4) Households whose minimum benefit amount received in any benefit month is greater than $250 and who spend 75% or more of their benefit in any given benefit month. These categories are mutually exclusive, but are not collectively exhaustive. (*10%,**5%,***1%).

An additional question of interest, concerns the relationship between length of participation in SNAP and the SNAP benefit cycle. For example, one could argue that low tenure households might display a more pronounced SNAP benefit cycle due to less familiarity with the SNAP program, with eligible stores, eligible products, and the most effective way to use SNAP benefits. However, high tenure SNAP households might instead show a more pronounced SNAP benefit cycle, if one associates a longer tenure with greater need for SNAP assistance, and a greater willingness to remain in the program, given that households must continuously meet recertification requirements. To examine this relationship, I estimate Equation 3 for the main sample and by tenure in SNAP, using the same three tenure categories as discussed in the descriptive results, Section 1.4, and present my results in Table 6.

First, consistent with the descriptive results, tenure is positively correlated with greater spending during the first 7 days of the benefit month. In particular, a 1% increase in total household benefit leads to an increase in the amount of benefit spent, within the first 7 days of the benefit month, of 0.58% for low tenure households, 0.80% for moderate tenure households, and 0.91% for high tenure households. Overall, these results show that an increase in total household benefit has a more pronounced effect on spending within the first 7 days for the highest tenure households.

Completing an additional shopping transaction, during the benefit month, increases the amount of dollars spent within the first 7 days by 5.05%, 3.44%, and 1.70% for low, moderate, and high tenure households, respectively. Thus, I find a more pronounced effect of additional monthly shopping transactions, on spending within the first 7 days, for lower tenure households. Bulk shopping during a given benefit month is associated with a significant increase in the amount of dollars spent within the first 7 days of 50.04%, 53.45%, and 50.24%, for low, moderate, and high tenure households, respectively. Thus, bulk shopping has a more pronounced effect on spending, within the first 7 days, for moderate tenure households. In contrast, being classified as a mixed household reduces the amount of dollars spent within the first 7 days by 1.39%, 1.21%, and 1.17% for low, moderate and high tenure households, respectively.

Table 6. Regression Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Sample and Heterogeneity by Tenure in SNAP

	(1) Main Sample	(2) T≤12	(3) T∈(12,24)	(4) T≥24
Log Credit	0.8594***	0.5832***	0.8019***	0.9078***
	(0.0008)	(0.0034)	(0.0020)	(0.0009)
Transactions	0.0221***	0.0505***	0.0344***	0.0170***
	(0.0001)	(0.0003)	(0.0002)	(0.0001)
Bulk Shopping	0.5035***	0.5004***	0.5345***	0.5024***
	(0.0007)	(0.0024)	(0.0018)	(0.0008)
Log Grocery	0.0262***	0.0225***	0.0271***	0.0265***
	(0.0002)	(0.0012)	(0.0008)	(0.0002)
Log Convenience	-0.0046***	-0.0180***	-0.0114***	-0.0032***
	(0.0002)	(0.0011)	(0.0007)	(0.0003)
Log Combination	0.0058***	0.0009	0.0039***	0.0064***
	(0.0002)	(0.0009)	(0.0006)	(0.0002)
Log Specialty	0.0241***	0.0191***	0.0242***	0.0248***
	(0.0003)	(0.0015)	(0.0010)	(0.0003)
Log Other Stores	0.0242***	0.0308***	0.0258***	0.0233***
	(0.0004)	(0.0018)	(0.0012)	(0.0004)
Rural	-0.0033	-0.0127	-0.0116**	0.0010
	(0.0021)	(0.0086)	(0.0059)	(0.0022)
Mixed	-0.0117***	-0.0139**	-0.0121***	-0.0117***
	(0.0014)	(0.0055)	(0.0039)	(0.0015)
Period	-0.0022***	0.1085***	0.0152***	-0.0032***
	(0.0000)	(0.0004)	(0.0001)	(0.0000)
R-squared	0.3253	0.2087	0.2526	0.3599
Households	1,555,674	636,521	287,162	631,991
Benefit Months	34,770,457	3,861,033	5,094,561	25,814,863

Regressions use correlated random effects models. The dependent variable in all regressions is the log of dollars spent within the first 7 days of the benefit month. Length of participation in SNAP is measured by tenure (T), the number of benefit months a households spends in the program. Four groups are displayed: (1) The main sample consisting of all households; (2) Households whose tenure in SNAP is less or equal to 12 benefit months (1 year or less); (3) Households whose tenure in SNAP is between 12 and 24 benefit months exclusive (more than 1 year and less than 2 years); (4) Households whose tenure in SNAP is greater or equal to 24 benefit months (2 years or more). (*10%,**5%,***1%).

Additionally, moderate tenure rural households spend 1.16% less dollars within the first 7 days of the benefit month compared to moderate tenure urban households.

Finally, I examine the effects across the store categories. Similar to the main sample results and the heterogeneity analysis by benefit size, and saving behavior, no significant differences across the tenure groups are detected. Neither store category appears to play an outsized role in the SNAP benefit cycle for either of the three tenure groups. This result is consistent with the descriptive analysis, where it was evident that the preferred store category among all three tenure groups was Supermarkets.

Overall, the regression results provide several insights. First, a more generous benefit does not automatically lead to a less pronounced benefit cycle. In particular, a 1% increase in household benefit can lead to an increase in the amount of dollars spent within the first 7 days that ranges from 0.4% to 1.49%. The effect is most pronounced for moderate benefit households, followed by low and high benefit households. This pattern does not change once I account for household's saving behavior. Additionally, the elasticity of the amount of dollars spent within the first 7 days, with respect to household benefit, increases with tenure.

Second, additional monthly transactions are positively associated with the severity of the benefit cycle. In particular, an additional monthly transaction leads to an increase in the amount of dollars spent within the first 7 days, that ranges from 1.44% to 6.77%. The effect of additional monthly transactions on the SNAP benefit cycle decreases with benefit size, saving behavior, and tenure. Bulk shopping is positively associated with the benefit cycle. Specifically, bulk shopping during a given benefit month, leads to a large increase in dollars spent within the first 7 days ranging from 38% to 59%. The effect of bulk shopping on the SNAP benefit cycle decreases with benefit size, and saving behavior.

Neither rural nor mixed residence seem to have very prominent effects on the benefit cycle, compared to urban residence. Finally, spending additional dollars during the month at either grocery, convenience, combination, specialty and other stores, while decreasing the amount of benefit spent at supermarkets, leads to small changes in the SNAP benefit cycle. The estimated effects, for the main sample, of a 1% increase in the amount of dollars spent at any of these stores, on the amount of dollars spent within the first 7 days, range between -0.0046% and 0.0262%. Thus relative to the supermarkets, the remaining store categories, seem to play a very small role in the benefit cycle. The results for household location and store choice are robust across the benefit size, saving behavior, and tenure subgroups.

1.7 Finite Mixtures Results

To uncover unobserved heterogeneity in the SNAP benefit cycle, I estimate a series of two-class finite mixture models. I restrict my attention to two-class finite mixture models for three main reasons. First, Figure 1 shows compelling evidence that there are two subpopulations of SNAP households.

Figure 1. Kernel Density, Household Average Log of \$ Spent Within the First 7 days, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Sample

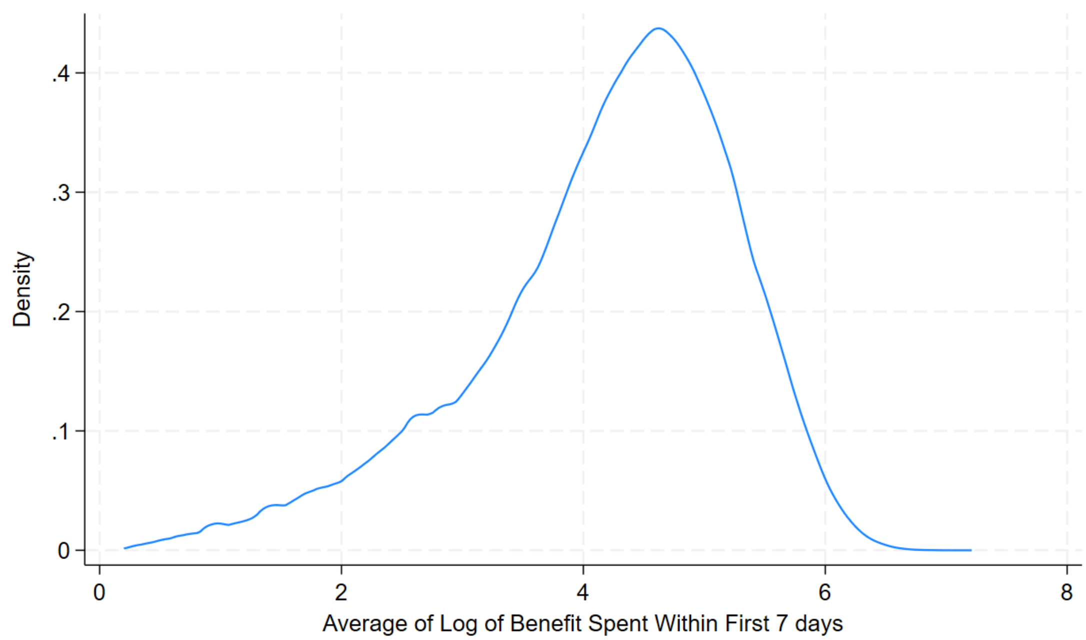

Note: The Main Sample consists of all Georgia SNAP households during 2011 - 2015. The entire sample contains 34,770,457 benefit months generated by 1,555,674 households.

Second, previous studies[18] examining the SNAP benefit cycle have provided evidence supporting the existence of only two groups of SNAP spenders. Third, I estimate three class mixtures and find that results remain qualitatively similar.[19] Thus, for ease of exposition, I discuss two-class mixture results.

I start my analysis by estimating a two-class mixture using a 10% random sample of all households in the data.[20] This subsample contains 3,474,589 benefit months generated by 155,567 households. This analysis aims to shed light on the unobserved heterogeneity in the SNAP benefit cycle without restricting the subsample by household characteristics or spending behavior. In other words, the results from this analysis are representative of the unobserved heterogeneity present in the main (analytical) sample and, as such, serve as the main results from the finite mixture analysis. Additionally, to test the robustness of my finite mixture results, to household tenure in SNAP and household saving behavior, I re-estimate my two-class finite mixture model on two additional subsamples.

First, I restrict my data to only include households who receive a minimum benefit of $100 dollars or more and spend at least 75% of their benefit in any given benefit month. This restriction yields 918,481 households who generate 19,795,513 benefit months. Subsequently, I extract a 10% random sample of households from this restricted sample. The resulting subsample contains 91,848 households who generate 1,986,510 benefit months. There are two main reasons for this analysis. First, unobserved heterogeneity in the cycle could depend on benefit size, as households who receive larger benefits might be more likely to smooth benefit spending. However, given that benefit size is tied to household size and potentially precarious financial circumstances, these households might instead display a more pronounced benefit cycle. Second, as previously mentioned, failing to account for household saving behavior could lead to an underestimation of the extent of the benefit cycle, as saving behavior could appear indistinguishable from budgeting behavior. This analysis aims to examine whether the

[18] See Dorfman et al. (2019).
[19] These results are reported in Appendix Table A4.
[20] To facilitate the feasibility of the finite mixture analysis, only subsamples are employed to conduct the analysis.

unobserved heterogeneity, potentially present in the SNAP benefit cycle, varies conditional on the size of household benefit and household saving behavior.

Finally, I impose a number of restrictions on the main sample to create my final subsample. In this most restrictive subsample, I include only households who receive a minimum benefit of \$100 dollars or more, spend at least 75% of their benefit in any given benefit month, spend a minimum of 24 benefit months in the program (2 years), and whose first recorded month of participation in the SNAP program is January 2012. These restrictions yield 38,059 households who generate 1,238,555 benefit months. Given the limited size of the restricted sample, I include all 38,059 households in the final subsample. This subsample allows me to examine variation in the benefit cycle among households who receive enough SNAP benefit to potentially smooth their consumption, display no benefit saving behavior, and have spent a considerable amount of time in the program allowing them to acclimate to program dynamics, familiarize themselves with SNAP eligible food items, SNAP stores, and recertification requirements. Further, by restricting household's first recorded benefit month to January 2012, my finite mixture estimation excludes households with tenure prior to the start date of observation in my data which is in January of 2011.

From here on out, I will be referring to the three subsamples discussed as follows. The 10% random sample of all households in the data will be called the main subsample. The 10% random sample of households who receive a minimum benefit of \$100 dollars or more and spend at least 75% of their benefit in any given benefit month, will be called restricted subsample 1. The remaining subsample, the most restrictive, will be called restricted subsample 2. In Table 7, I present summary statistics for the main sample and main subsample. In Table 8, I present summary statistics for the restricted sample and restricted subsample 1. In Table 9, I present summary statistics for the main sample, and for restricted subsample 2. Tables 7 and 8 indicate that the summary statistics from the random samples closely match the statistics from the samples they were extracted from.

Table 7. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and Main Subsample

	Main Sample	***Main Subsample***
Amount Spent		
$ Within 7 days	167.47	167.52
	(150.23)	(150.35)
% Within 7 days	59.72	59.79
	(33.54)	(33.57)
$ Within Month	261.69	261.60
	(188.05)	(188.56)
% Within Month	91.55	91.57
	(18.44)	(18.43)
Transactions & Benefit Amount		
# of Transactions (7 days)	4.31	4.31
	(3.68)	(3.68)
# of Transactions (Month)	8.76	8.75
	(6.87)	(6.88)
Bulk Shopping (%)	49.26	49.40
	(49.99)	(50.00)
Benefit Amount ($)	287.80	287.57
	(198.00)	(198.50)
SNAP Tenure		
Tenure (Ben. Months)	35.80	35.84
	(15.80)	(15.84)
Households	1,555,674	155,567
Benefit Months	34,770,457	3,474,589

Note: The main sample consists of benefit months generated by the Universe of SNAP participating households in Georgia. The main subsample refers to the benefit months generated by a 10% random sample of households from the main sample.

Table 7 (Cont'd): Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and Main Subsample

	Main Sample	***Main Subsample***
Household Location		
Rural	16.87	16.75
	(37.45)	(37.35)
Urban	75.88	76.01
	(42.78)	(42.70)
Mixed	7.26	7.23
	(25.94)	(25.90)
Store Choice		
Supermarkets ($)	215.83	215.74
	(171.77)	(172.16)
Grocery ($)	12.00	12.11
	(41.42)	(41.68)
Convenience ($)	13.75	13.66
	(36.09)	(36.02)
Combination ($)	13.40	13.34
	(32.80)	(32.78)
Specialty Foods ($)	3.27	3.28
	(17.65)	(17.77)
Other ($)	3.44	3.46
	(24.57)	(24.55)
Households	1,555,674	155,567
Benefit Months	34,770,457	3,474,589

Note: The main sample consists of benefit months generated by the Universe of SNAP participating households in Georgia. The main subsample refers to the benefit months generated by a 10% random sample of households from the main sample.

Table 8. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Restricted Sample and Restricted Subsample 1

	Restricted Sample	***Restricted Subsample 1***
Amount Spent		
$ Within 7 days	206.24	206.28
	156.43	156.47
% Within 7 days	59.85	59.82
	31.55	31.55
$ Within Month	326.18	326.44
	184.91	185.01
% Within Month	93.10	93.10
	16.04	16.04
Transactions & Benefit Amount		
# of Transactions (7 days)	4.92	4.93
	3.83	3.84
# of Transactions (Month)	10.28	10.30
	7.14	7.16
Bulk Shopping (%)	43.06	42.99
	49.52	49.51
Benefit Amount ($)	351.69	352.02
	191.24	191.49
SNAP Tenure		
Tenure (Ben. Months)	35.35	35.38
	16.08	16.02
Households	918,481	91,848
Benefit Months	19,795,513	1,986,510

Note: The restricted sample consists of benefit months generated by households who receive a minimum benefit of $100 dollars or more and spend at least 75% of their benefit in any given benefit month. Restricted subsample 1 refers to the benefit months generated by a 10% random sample of these households.

Table 8 (Cont'd): Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Restricted Sample and Restricted Subsample 1

	Restricted Sample	***Restricted Subsample 1***
Household Location		
Rural	15.12	15.14
	35.82	35.84
Urban	77.88	77.90
	41.50	41.49
Mixed	7.00	6.96
	25.52	25.44
Store Choice		
Supermarkets ($)	269.57	269.80
	173.94	174.36
Grocery ($)	15.40	15.54
	48.13	48.40
Convenience ($)	17.00	17.01
	41.11	41.01
Combination ($)	15.91	15.79
	36.97	36.86
Specialty Foods ($)	3.92	3.84
	19.73	19.57
Other ($)	4.39	4.45
	28.48	28.50
Households	918,481	91,848
Benefit Months	19,795,513	1,986,510

Note: The restricted sample consists of benefit months generated by households who receive a minimum benefit of $100 dollars or more and spend at least 75% of their benefit in any given benefit month. Restricted subsample 1 refers to the benefit months generated by a 10% random sample of these households.

Table 9. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and Restricted Subsample 2

	Main Sample	***Restricted Subsample 2***
Amount Spent		
$ Within 7 days	167.47	179.81
	150.23	133.14
% Within 7 days	59.72	58.66
	33.54	31.75
$ Within Month	261.69	291.82
	188.05	161.84
% Within Month	91.55	92.59
	18.44	16.68
Transactions & Benefit Amount		
# of Transactions (7 days)	4.31	4.41
	3.68	3.45
# of Transactions (Month)	8.76	9.30
	6.87	6.48
Bulk Shopping (%)	49.26	44.93
	49.99	49.74
Benefit Amount ($)	287.80	316.71
	198.00	168.49
SNAP Tenure		
Tenure (Ben. Months)	35.80	33.65
	15.80	6.13
Households	1,555,674	38,059
Benefit Months	34,770,457	1,238,555

Note: The main sample consists of benefit months generated by the Universe of SNAP participating households in Georgia. Restricted subsample 2 consists of benefit months generated by households who receive a minimum benefit of $100 or more in any given benefit month, spend at least 75% of their benefit in any given benefit month, spend 2 years or more in SNAP and their first recorded month of participation in the program is January 2012.

Table 9 (Cont'd): Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and Restricted Subsample 2

	Main Sample	***Restricted Subsample 2***
Household Location		
Rural	16.87	14.49
	37.45	35.20
Urban	75.88	78.07
	42.78	41.38
Mixed	7.26	7.44
	25.94	26.24
Store Choice		
Supermarkets ($)	215.83	243.26
	171.77	155.23
Grocery ($)	12.00	13.10
	41.42	41.58
Convenience ($)	13.75	13.25
	36.09	33.72
Combination ($)	13.40	14.53
	32.80	34.42
Specialty Foods ($)	3.27	2.95
	17.65	16.25
Other ($)	3.44	4.73
	24.57	28.90
Households	1,555,674	38,059
Benefit Months	34,770,457	1,238,555

Note: The main sample consists of benefit months generated by the Universe of SNAP participating households in Georgia. Restricted subsample 2 consists of benefit months generated by households who receive a minimum benefit of $100 or more in any given benefit month, spend at least 75% of their benefit in any given benefit month, spend 2 years or more in SNAP and their first recorded month of participation in the program is January 2012.

1.7.1 Main Subsample

Table 10 reports the results from the two-class finite mixture estimated using the main subsample, a 10% random sample of all households in the data. Two groups of SNAP recipients are found, which I denote as slow and fast spenders. Out of 155,567 households, I find 35,780 households in the slow group (23% of the sample) and 119,787 households in the fast group (77% of the sample). The fast group spends 12 times as many SNAP dollars ($147) within the first 7 days of the benefit month as the slow group ($12). The finite mixture results reveal a number of interesting patterns. The estimated coefficients indicate that a 1% increase in household SNAP benefit increases the amount of SNAP dollars spent within the first 7 days by 0.45% for slow households and 0.92% for fast households.

Completing an additional shopping transaction, during the benefit month, is associated with a 6.92% increase in the amount of SNAP dollars that slow spenders use in the first 7 days and a 1.04% reduction for fast spenders. Bulk shopping has a more pronounced effect on spending within the first 7 days for the slow spending group. In particular, I find bulk shopping during a given benefit month reduces the amount of SNAP dollars spent within the first 7 days by 53.64% for slow households while increasing the amount of SNAP dollars spent within the first 7 days by 27.43% for fast households. In addition, I also document that rural and mixed household residence have little or no effect on the amount of dollars spent within the first 7 days of the benefit month compared to urban residence. Finally, I turn to the relationship between store choice and the SNAP benefit cycle where supermarkets are the base category. Here, I find that a 1% increase in the amount of SNAP dollars spent at either of the 5 store categories leads to a small, and mostly statistically significant, 0.0035% to 0.0384%, increase in SNAP benefit spent in the first 7 days across both groups. These results indicate that none of the 5 store categories seem to have a substantial effect on the SNAP benefit cycle.

Next I use the finite mixture estimates to calculate the posterior probabilities of households belonging to the slow or fast spending group.[21] With these posterior probabilities, I

[21] The panel data finite mixture model generates multiple posterior probabilities for each household. Specifically, for a given household, each benefit month will be associated with a set of posterior probabilities, one for each group. I

Table 10. Two-Class Finite Mixture Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Subsample

	Main Subsample	2-Part Mixture	
		Group 1	Group 2
Log Credit	0.8602***	0.4517***	0.9235***
	(0.0025)	(0.0051)	(0.0007)
Transactions	0.0220***	0.0692***	-0.0104***
	(0.0002)	(0.0007)	(0.0001)
Bulk Shopping	0.5046***	-0.5364***	0.2743***
	(0.0021)	(0.0066)	(0.0010)
Log Grocery	0.0269***	0.0284***	0.0120***
	(0.0007)	(0.0027)	(0.0002)
Log Convenience	-0.0054***	0.0035	0.0086***
	(0.0008)	(0.0026)	(0.0002)
Log Combination	0.0062***	0.0384***	0.0023***
	(0.0007)	(0.0022)	(0.0002)
Log Specialty	0.0241***	0.0104***	0.0132***
	(0.0009)	(0.0037)	(0.0003)
Log Other Stores	0.0238***	0.0105**	0.0113***
	(0.0012)	(0.0043)	(0.0004)
Rural	0.0086	-0.0101	0.0036**
	(0.0065)	(0.0166)	(0.0018)
Mixed	0.0000	0.0095	-0.0039***
	(0.0044)	(0.0117)	(0.0012)
Period	-0.0023***	0.0033***	-0.0014***
	(0.0001)	(0.0002)	(0.0000)
R-squared	0.3271		
Mean (Y)	$167	$12	$147
Probability		23%	77%
Households	155,567		
Benefit Months	3,474,589		

Note: Regression and mixture employ correlated random effects models. The dependent variable in the regression and mixture is the log of dollars spent within the first 7 days of the benefit month. The main subsample column displays the results from a correlated random effects regression using a 10% random sample of all Georgia SNAP households during 2011 - 2015. Group 1 and 2 refer to the 2 part estimated classes from the mixture using the same 10% random sample. (*10%,**5%,***1%).

classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20).[22] The classified sample contains 1,839,745 benefit months generated by 75,482 households. Of the 75,482 households, 6,985 belong to the slow spending group and 68,497 belong to the fast spending group. For each of these groups, I report the respective means of the independent variables in Table 11.

Table 11 offers several interesting insights. First, it is evident that slow households spend a much smaller percentage of their benefit within the first 7 days. In particular, slow households spend 8.67% of their benefit within the first 7 days compared to 70.07% for fast spending households. Consistent with this pattern of spending, slow households complete significantly less transactions, within the first 7 days, 0.65 compared to 5.29 transactions for fast households. Examining the behavior of these households throughout the full benefit month, slow households continue to spend less both in absolute and percentage terms, and complete less transactions compared to fast households. In particular, on average, slow households spend $125.65, use 80.29% of their benefit, and complete 4.42 transactions within a benefit month, compared to $293.54, 93.12% of benefit, and 9.50 transactions for fast spending households. Additionally, slow households are more likely to bulk shop in any given benefit month, 62.36% chance, compared to fast spending households, 48.87% chance. However, given their lower spending throughout the first seven days, it is plausible that most of these bulk transactions occur after the first 7 days.

A priori, it is unclear whether a more generous benefit could lead to a less or more pronounced cycle. SNAP is a means-tested program, and presumably households who receive greater benefits are those in most need of these benefits. However, the larger benefit could allow these households to smooth their consumption, for instance, by stocking on durables and

average these probabilities to retrieve a set of 2 posterior probabilities for each household, across the whole period of analysis. These average posterior probabilities are used to classify households as belonging to either the slow or fast group.

[22] These conservative classification thresholds are chosen to ensure an accurate classification of households into the respective groups. However, a consequence is that some households are not classified as either type, given their values of the average posterior probabilities. These households are not included in the analysis comparing means by group.

Table 11. Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Subsample

	Classified Households	***Group 1 Slow***	***Group 2 Fast***	**Diff.**
Amount Spent				
$ Within 7 days	209.90	13.81	212.54	-198.72***
	(160.97)	(35.73)	(160.37)	
% Within 7 days	69.26	8.67	70.07	-61.40***
	(29.34)	(21.47)	(28.57)	
$ Within Month	291.31	125.65	293.54	-167.89***
	(196.92)	(133.18)	(196.69)	
% Within Month	92.95	80.29	93.12	-12.83***
	(16.72)	(29.32)	(16.42)	
Transactions & Benefit Amount				
# of Transactions (7 days)	5.23	0.65	5.29	-4.63***
	(3.91)	(1.41)	(3.89)	
# of Transactions (Month)	9.44	4.42	9.50	-5.08***
	(7.03)	(4.91)	(7.03)	
Bulk Shopping (%)	49.05	62.36	48.87	13.49***
	(49.99)	(48.45)	(49.99)	
Benefit Amount ($)	315.04	181.50	316.83	-135.32***
	(205.39)	(182.32)	(205.09)	
SNAP Tenure				
Tenure (Ben. Months)	38.13	15.67	38.43	-22.76***
	(15.39)	(17.24)	(15.13)	
Households	75,482	6,985	68,497	
Benefit Months	1,839,745	24,378	1,815,367	

Note: Using the finite mixture estimates, from the main subsample analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20). Out of 155,567 households, in the main subsample, only 75,482 meet the above criteria. The classified households generate 1,839,745 benefit months. Of the 75,482 households, 6,985 belong to the slow spenders group and 68,497 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

Table 11 (Cont'd): Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Main Subsample

	Classified Households	***Group 1 Slow***	***Group 2 Fast***	**Diff.**
Household Location				
Rural	17.86	16.76	17.88	-1.12***
	(38.30)	(37.35)	(38.32)	
Urban	75.04	74.67	75.04	-0.37
	(43.28)	(43.49)	(43.28)	
Mixed	7.10	8.57	7.08	1.49***
	(25.69)	(28.00)	(25.65)	
Store Choice				
Supermarkets ($)	239.63	104.31	241.45	-137.14***
	(180.62)	(124.05)	(180.57)	
Grocery ($)	13.81	4.72	13.94	-9.21***
	(44.91)	(23.57)	(45.11)	
Convenience ($)	15.75	6.37	15.88	-9.51***
	(39.47)	(23.99)	(39.62)	
Combination ($)	14.54	6.80	14.64	-7.85***
	(34.29)	(21.67)	(34.41)	
Specialty Foods ($)	3.93	1.10	3.96	-2.87***
	(20.03)	(9.51)	(20.13)	
Other ($)	3.65	2.35	3.67	-1.32***
	(25.52)	(19.53)	(25.59)	
Households	75,482	6,985	68,497	
Benefit Months	1,839,745	24,378	1,815,367	

Note: Using the finite mixture estimates, from the main subsample analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20). Out of 155,567 households, in the main subsample, only 75,482 meet the above criteria. The classified households generate 1,839,745 benefit months. Of the 75,482 households, 6,985 belong to the slow spenders group and 68,497 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

non-perishables, and planning their monthly food consumption by conducting multiple grocery trips. The finite mixture model estimates, and comparison of means post mixture estimation, indicate that the households who receive the highest benefit, present the more severe benefit cycle.

Further, it is also unclear whether longer tenure in SNAP should be associated with a more or less severe benefit cycle. Similar to a larger benefit, a longer tenure could also signal a greater need to participate in SNAP by a given household, especially considering that recertification periods depend on a household's financial circumstances. Further, longer tenure households might also be those who have chosen to continuously recertify further signaling their interest to stay in the program. In addition, a longer tenure may also provide a learning opportunity for households, to gain a better understanding of SNAP dynamics, SNAP benefits, and SNAP stores, among other factors, which may improve their SNAP benefits consumption patterns and reduce their probability of experiencing a severe benefit cycle. Here, the results also indicate a positive correlation between tenure and the SNAP benefit cycle. In particular, I find that slow households spend considerably less time on average in the program, 16 benefit months (1 year and 4 months) compared to fast spending households who stay in the SNAP program for 38 benefit months (3 years and 2 months).

Finally, regarding locations where households reside, I do not find major differences between slow and fast households. Most households in either group are classified as urban households (approximately 75%), followed by rural (approximately 17%) and finally mixed households (approximately 8%). Additionally, the choices of which store type to shop at, throughout the benefit month, are fairly similar between both groups as well. The preferred choice are supermarkets, followed by either convenience or combination stores, grocery stores and the other two remaining categories, specialty and other stores.

1.7.2 Unobserved Heterogeneity by Saving Behavior & Tenure

Next, I discuss the results from estimating two-class mixture models using restricted subsamples 1 and 2. In Table 12, I present results for restricted subsample 1, the subsample of households who receive a minimum benefit of $100 or more and spend at least 75% of their benefit in any given benefit month. The striking feature of these results is the robustness of the finite mixture model estimates to accounting for the benefit size and saving behavior restrictions. Many of the findings from the main subsample analysis remain qualitatively the same.

In particular, I again find evidence of two groups of spenders, slow and fast spenders. The fast spending group spends 13 times as many dollars ($200) within the first 7 days of the benefit month as the slow group ($15). A 1% increase in the household SNAP benefit increases the amount of dollars spent within the first 7 days by 0.23% for slow spenders and 0.90% for fast spenders. Completing an additional shopping transaction, during the benefit month, is associated with a 6.94% increase in the amount of SNAP dollars spent within the first 7 days by slow spenders, and a 1.10% reduction for fast spenders. Bulk shopping, during a given benefit month, reduces the amount of dollars spent within the first 7 days by 53.18% for slow spenders and increases the amount of dollars spent within the first 7 days by 27.21% for fast spenders. Rural and mixed household residence continue to have little or no effect on the amount of dollars spent within the first 7 days of the benefit month compared to urban residence. Finally, store choice patterns and their effect on the SNAP benefit cycle remain qualitatively the same across both groups. For either group, none of the 5 store categories seem to have much of an effect on the SNAP benefit cycle.

Similar to the previously estimated results, I compare the means of the key independent variables, across the two groups estimated by this mixture model, using the same assignment rule and threshold, conditional on the posterior probabilities. The classified sample contains 54,732 households. Of the 54,732 households, 3,329 belong to the slow spenders group and 51,403 belong to the fast spenders group. The results are shown in Table 13. These results again remain remarkably consistent with the findings from the main subsample. Slow spenders continue to spend a much lower percentage of their benefit within the first 7 days. In particular, slow spenders spend 5.51% of their benefit within the first 7 days compared to 66.88% for fast spenders. Slow spenders continue to complete significantly less transactions, within the first 7 days, 0.49 compared to 5.69 for fast spenders and complete less transactions, throughout the full benefit month as well, 6.59 compared to 10.86 for fast spenders. Slow spenders continue to be more likely to bulk shop, 65.42% chance compared to 42.43% for fast spenders, receive a lower benefit on average compared to fast spenders, \$248.77 and \$371.68, respectively, and spend considerably less time in the program on average, 6 benefit months (6 months) compared to 38 benefit months (3 years and 2 months) for fast spenders.

Regarding locations where households reside, and their preferred store type, the patterns again remain extremely similar to the main subsample. Most households in either group are classified as urban households 81.52% and 76.88% for slow and fast households respectively, followed by rural 10.61% and 16.28% respectively, and finally mixed households 7.87% and 6.84% respectively. Finally the choices of where to shop throughout the benefit month are fairly similar between both groups as well. The preferred choice are supermarkets, followed by convenience, grocery or combination and lastly, specialty and other stores.

Finally, I discuss the results from estimating a two-class mixture model using restricted subsample 2. In Table 14, I present the results for restricted subsample 2, the subsample of households receiving a minimum benefit of \$100 or more, who spend at least 75% of their benefit in any given benefit month, spend a minimum of 24 benefit months in the program (2 years), and whose first recorded month of SNAP participation in the data is January 2012. Despite the very restrictive nature of this subsample, the findings again remain remarkably consistent. Two groups of SNAP recipients are found, slow and fast spenders. Fast spenders spend 13 times as many dollars (\$179) within the first 7 days of the benefit month as slow spenders (\$13). A 1% increase in the household SNAP benefit increases the amount of dollars spent within the first 7 days by 0.27% for slow spenders and 0.87% for fast spenders.

Completing an additional shopping transaction, during the benefit month, is associated with a 7.74% increase in the amount of SNAP dollars spent within the first 7 days by slow spenders, and a 1.10% reduction in the amount of dollars spent within the first 7 days by fast spenders. Bulk shopping, during a given benefit month, reduces the amount of dollars spent within the first 7 days by 41.54% for slow spenders, and increases the amount of dollars spent

Table 12. Two-Class Finite Mixture Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 1

	Restricted Subsample 1	2-Part Mixture	
		Group 1	Group 2
Log Credit	0.7086***	0.2327***	0.9034***
	(0.0053)	(0.0149)	(0.0016)
Transactions	0.0173***	0.0694***	-0.0110***
	(0.0003)	(0.0009)	(0.0001)
Bulk Shopping	0.4343***	-0.5318***	0.2721***
	(0.0027)	(0.0099)	(0.0011)
Log Grocery	0.0235***	0.0223***	0.0116***
	(0.0009)	(0.0036)	(0.0003)
Log Convenience	-0.0046***	-0.0015	0.0086***
	(0.0010)	(0.0035)	(0.0003)
Log Combination	0.0009	0.0222***	0.0017***
	(0.0008)	(0.0031)	(0.0002)
Log Specialty	0.0185***	-0.0095*	0.0122***
	(0.0011)	(0.0051)	(0.0003)
Log Other Stores	0.0228***	0.0250***	0.0109***
	(0.0015)	(0.0058)	(0.0004)
Rural	0.0118	-0.0017	0.0053**
	(0.0087)	(0.0282)	(0.0025)
Mixed	-0.0155***	-0.0153	-0.0011
	(0.0059)	(0.0189)	(0.0016)
Period	-0.0017***	0.0053***	-0.0015***
	(0.0001)	(0.0004)	(0.0000)
R-squared	0.1780		
Mean (Y)	$123	$15	$200
Probability		18%	82%
Households	91,848		
Benefit Months	1,986,510		

Note: Regression and mixture employ correlated random effects models. The dependent variable in the regression and mixture is the log of dollars spent within the first 7 days of the benefit month. The restricted subsample 1 column displays the results from a correlated random effects regression using a 10% random sample of households who receive a minimum benefit of $100 dollars or more and spend at least 75% of their benefit in any given benefit month. Group 1 and 2 refer to the 2 part estimated classes from the mixture using the same 10% random sample. (*10%,**5%,***1%).

Table 13. Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 1

	Classified Households	*Group 1 Slow*	*Group 2 Fast*	**Diff.**
Amount Spent				
$ Within 7 days	240.04	12.79	241.05	-228.26***
	(161.43)	(36.42)	(161.06)	
% Within 7 days	66.61	5.51	66.88	-61.37***
	(28.64)	(15.88)	(28.39)	
$ Within Month	345.89	233.43	346.39	-112.96***
	(191.33)	(138.56)	(191.38)	
% Within Month	93.54	94.31	93.54	0.77***
	(15.55)	(12.37)	(15.57)	
Transactions & Benefit Amount				
# of Transactions (7 days)	5.67	0.49	5.69	-5.20***
	(3.98)	(1.16)	(3.97)	
# of Transactions (Month)	10.84	6.59	10.86	-4.28***
	(7.31)	(5.59)	(7.31)	
Bulk Shopping (%)	42.53	65.42	42.43	22.99***
	(49.44)	(47.57)	(49.42)	
Benefit Amount ($)	371.14	248.77	371.68	-122.91***
	(197.12)	(146.23)	(197.14)	
SNAP Tenure				
Tenure (Ben. Months)	37.60	5.96	37.74	-31.78***
	(15.45)	(10.19)	(15.32)	
Households	54,732	3,329	51,403	
Benefit Months	1,318,749	5,844	1,312,905	

Note: Using the finite mixture estimates, from the restricted subsample 1 analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20). Out of 91,848 households, in restricted subsample 1, only 54,732 meet the above criteria. The classified households generate 1,318,749 benefit months. Of the 54,732 households, 3,329 belong to the slow spenders group and 51,403 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

Table 13 (Cont'd): Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 1

	Classified Households	***Group 1 Slow***	***Group 2 Fast***	**Diff.**
Household Location				
Rural	16.26	10.61	16.28	-5.67***
	(36.90)	(30.80)	(36.92)	
Urban	76.90	81.52	76.88	4.64***
	(42.15)	(38.82)	(42.16)	
Mixed	6.85	7.87	6.84	1.03**
	(25.26)	(26.93)	(25.25)	
Store Choice				
Supermarkets ($)	284.89	192.83	285.30	-92.47***
	(180.48)	(136.21)	(180.55)	
Grocery ($)	16.83	8.12	16.87	-8.74***
	(50.96)	(35.35)	(51.01)	
Convenience ($)	18.74	13.15	18.76	-5.61***
	(43.60)	(38.84)	(43.62)	
Combination ($)	16.65	10.82	16.68	-5.86***
	(37.47)	(32.44)	(37.49)	
Specialty Foods ($)	4.34	1.87	4.35	-2.48***
	(21.18)	(12.94)	(21.21)	
Other ($)	4.44	6.64	4.43	2.20***
	(28.70)	(39.62)	(28.64)	
Households	54,732	3,329	51,403	
Benefit Months	1,318,749	5,844	1,312,905	

Note: Using the finite mixture estimates, from the restricted subsample 1 analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20). Out of 91,848 households, in restricted subsample 1, only 54,732 meet the above criteria. The classified households generate 1,318,749 benefit months. Of the 54,732 households, 3,329 belong to the slow spenders group and 51,403 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

within the first 7 days by 29.80% for fast spenders. Rural and mixed household residence continue to have little or no effect on the amount of dollars spent within the first 7 days of the benefit month compared to urban residence. Finally, store choice patterns and their effect on the SNAP benefit cycle remain qualitatively the same across both groups, with neither of the 5 store categories having a pronounced effect on the SNAP benefit cycle.

Comparing the means of the independent variables, across the two groups, many of the previous results remain qualitatively the same. The results are presented in Table 15. The classified sample contains 24,111 households. Of the 24,111 households, 29 belong to the slow spenders group and 24,082 belong to the fast spenders group. Slow spenders continue to spend a smaller percentage of their benefit within the first 7 days. In particular, slow spenders spend 11.06% of their benefit within the first 7 days compared to 65.60% for fast spenders. Slow spenders continue to complete significantly less transactions, within the first 7 days, 0.71 compared to 5.08 for fast spenders, and complete less transactions, throughout the full benefit month as well, 3.84 compared to 9.78 for fast spenders. Further, slow spenders continue to be more likely to bulk shop, 83.39% chance compared to 44.70% for fast spenders, and receive a lower benefit on average compared to fast spenders, $197.82 and $328.17, respectively.

Regarding locations where households reside, and their preferred store type to shop, the patterns again remain extremely similar to the main subsample. Most households in either group are classified as urban households, 89.17% and 77.23%, for slow and fast households respectively, followed by rural, 5.79% and 15.40% respectively, and finally mixed households, 5.04% and 7.37% respectively. Finally the choices of where to shop throughout the benefit month are fairly similar between both groups as well. The preferred choice are supermarkets, followed by combination, grocery or convenience and lastly, the other two remaining categories, specialty and other stores.

1.7.3 Covariates Affecting Probability of Assignment

The mixture results reveal the existence of two groups of SNAP households: slow and fast spenders. In this section, I estimate linear probability models that attempt to predict group

Table 14. Two-Class Finite Mixture Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 2

	Restricted Subsample 2	2-Part Mixture	
		Group 1	Group 2
Log Credit	0.6505***	0.2660***	0.8692***
	(0.0070)	(0.0186)	(0.0022)
Transactions	0.0229***	0.0774***	-0.0110***
	(0.0004)	(0.0012)	(0.0001)
Bulk Shopping	0.4907***	-0.4154***	0.2980***
	(0.0034)	(0.0117)	(0.0014)
Log Grocery	0.0256***	0.0256***	0.0114***
	(0.0012)	(0.0047)	(0.0004)
Log Convenience	-0.0136***	-0.0010	0.0077***
	(0.0013)	(0.0047)	(0.0004)
Log Combination	0.0027**	0.0351***	0.0031***
	(0.0011)	(0.0038)	(0.0003)
Log Specialty	0.0193***	-0.0073	0.0135***
	(0.0016)	(0.0068)	(0.0005)
Log Other Stores	0.0216***	0.0156**	0.0114***
	(0.0019)	(0.0068)	(0.0006)
Rural	0.0132	0.0103	0.0014
	(0.0116)	(0.0352)	(0.0033)
Mixed	-0.0136*	-0.0130	0.0002
	(0.0073)	(0.0230)	(0.0021)
Period	0.0102***	0.0369***	-0.0002***
	(0.0002)	(0.0005)	(0.0000)
R-squared	0.1412		
Mean (Y)	$109	$13	$179
Probability		19%	81%
Households	38,059		
Benefit Months	1,238,555		

Note: Regression and mixture employ correlated random effects models. The dependent variable in both is the log of dollars spent within the first 7 days of the benefit month. The restricted subsample 2 column displays the results of a regression using benefit months generated by households who receive a minimum benefit of $100 or more, spend at least 75% of their benefit in any given benefit month, spend 2 years or more in SNAP and their first recorded month of participation in the program is January 2012. Group 1 and 2 refer to the 2 part estimated classes from the mixture using the same subsample. (*10%,**5%,***1%).

Table 15. Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 2

	Classified Households	***Group 1 Slow***	***Group 2 Fast***	**Diff.**
Amount Spent				
$ Within 7 days	206.59	22.52	206.81	-184.29***
	(136.19)	(50.52)	(136.11)	
% Within 7 days	65.53	11.06	65.60	-54.54***
	(29.20)	(24.30)	(29.14)	
$ Within Month	303.89	182.93	304.03	-121.11***
	(166.73)	(66.98)	(166.76)	
% Within Month	93.04	93.52	93.04	0.48
	(16.26)	(15.62)	(16.26)	
Transactions & Benefit Amount				
# of Transactions (7 days)	5.08	0.71	5.08	-4.38***
	(3.60)	(1.19)	(3.60)	
# of Transactions (Month)	9.77	3.84	9.78	-5.94***
	(6.60)	(2.91)	(6.60)	
Bulk Shopping (%)	44.74	83.39	44.70	38.69***
	(49.72)	(37.24)	(49.72)	
Benefit Amount ($)	328.01	197.82	328.17	-130.34***
	(172.69)	(72.45)	(172.72)	
SNAP Tenure				
Tenure (Ben. Months)	34.03	33.25	34.03	-0.79***
	(6.24)	(6.08)	(6.24)	
Households	24,111	29	24,082	
Benefit Months	792,816	933	791,883	

Note: Using the finite mixture estimates, from the restricted subsample 2 analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")≥.80) or fast spender (Pr("Slow")<.20). Out of 38,059 households, in restricted subsample 2, only 24,111 meet the above criteria. The classified households generate 792,816 benefit months. Of the 24,111 households, 29 belong to the slow spenders group and 24,082 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

Table 15 (Cont'd): Comparing Characteristics Across the Groups, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 2

	Classified Households	***Group 1 Slow***	***Group 2 Fast***	**Diff.**
Household Location				
Rural	15.39	5.79	15.40	-9.61***
	(36.09)	(23.36)	(36.10)	
Urban	77.24	89.17	77.23	11.95***
	(41.93)	(31.09)	(41.94)	
Mixed	7.37	5.04	7.37	-2.33**
	(26.12)	(21.88)	(26.13)	
Store Choice				
Supermarkets ($)	252.32	163.25	252.43	-89.18***
	(160.04)	(71.72)	(160.09)	
Grocery ($)	13.85	4.88	13.86	-8.98***
	(43.24)	(23.19)	(43.25)	
Convenience ($)	14.69	3.31	14.70	-11.40***
	(35.64)	(12.49)	(35.65)	
Combination ($)	14.98	9.64	14.98	-5.34***
	(34.23)	(35.29)	(34.23)	
Specialty Foods ($)	3.28	0.32	3.28	-2.96***
	(17.34)	(3.96)	(17.35)	
Other ($)	4.78	1.53	4.78	-3.26***
	(29.19)	(16.17)	(29.20)	
Households	24,111	29	24,082	
Benefit Months	792,816	933	791,883	

Note: Using the finite mixture estimates, from the restricted subsample 2 analysis, I calculate the posterior probabilities of households belonging to group 1 (slow spenders) or group 2 (fast spenders). With these posterior probabilities, I classify a household either as a slow spender (Pr("Slow")$\geq$.80) or fast spender (Pr("Slow")<.20). Out of 38,059 households, in restricted subsample 2, only 24,111 meet the above criteria. The classified households generate 792,816 benefit months. Of the 24,111 households, 29 belong to the slow spenders group and 24,082 belong to the fast spenders group. Diff. refers to the difference in means between slow and fast spenders.

membership based on observable household characteristics. This analysis uses the classified sample from the main subsample mixture analysis. Group membership is determined using the same approach discussed in Section 1.7.1. Given that group membership relies on the average posterior probability of assignment, to identify household variables that correlate with group membership, I rely on household averages of some of the initial independent variables of interest and further, generate additional variables at the household level, that could provide relevant insights.

In particular, the covariates included are: the average monthly benefit received by a given household measured in hundreds of dollars; the average number of transactions the household completes during a given benefit month; bulk shopper, a binary indicator that identifies households whose probability of bulk shopping in any given benefit month is greater than 67%; rural, a dummy which identifies households who are classified as rural in at least 67% of their benefit months; urban, a dummy which identifies households who are classified as urban in at least 67% of their benefit months; mixed, a dummy which identifies the remaining households who do not meet the rural or urban criteria; finally, I generate a set of dummies which identify stores where households prefer to spend their benefits. In particular, a household is defined as having a preference for shopping at one of the six store categories, if it spends the highest monthly percentage of their benefit, on average, at that store.

In Table 16, I present the results from this analysis. The base group are slow spenders. All variables included are found to have a statistically significant effect on the probability of a household being a fast spender. However, these variables offer little predictive power ($R^2 = 0.0416$) and often are not economically significant. In particular, an increase in the average household benefit of \$100 reduces the probability of a household being a fast spender by 0.11 percentage points. Households who shop more frequently, on average, are 0.22 percentage points more likely to be fast spenders. Bulk shoppers are 0.41 percentage points less likely to be fast spenders. Rural households are 0.25 percentage points more likely to be fast spenders, while mixed households are 0.24 percentage points less likely to be fast spenders. Spending an additional month in the program increases the probability of a household being a fast spender by 0.12 percentage points.

Finally, regarding store preferences, households who prefer to shop at alternative store categories instead of supermarkets, have a lower probability of being fast spenders. The estimated reductions, in the probability of being a fast spender, range from 1.74 percentage points to 5.71 percentage points. Specifically, the estimated reductions due to preferring grocery, other, convenience, specialty and combination stores instead of supermarkets are 1.74, 2.04, 2.44, 5.70, and 5.71 percentage points respectively. These effects are also the largest among the variables employed, implying that shopping preferences might serve as a potential useful indicator of SNAP household spending type.

1.8 Discussion & Conclusion

Previous research found that SNAP recipients spend a significant amount of their monthly benefit early in the benefit month. Further, this research indicated that this behavior was associated with nutritional consequences for SNAP households, potentially diminishing the effectiveness of the SNAP program in reducing food insecurity. However, most of the existing studies model SNAP recipient spending behavior as coming from the same distribution, potentially ignoring heterogeneity in the spending behavior of SNAP households that could explain the SNAP benefit cycle and uncover potential policy levers to target it. The basic premise of this study is that this approach misspecifies the true model and omits potentially important information.

Using restricted administrative panel data containing the universe of SNAP participating households in Georgia, and a panel data finite mixture model, I find compelling evidence supporting my view. My results indicate that there are two groups of SNAP households, slow and fast spenders. Fast spenders are the larger group, comprising an estimated 77% of households, with the remaining 23% being slow spenders. Additionally, fast spenders display a significant benefit cycle, spending approximately 70% of their monthly benefit within the first 7 days, compared to 9% for slow households. Exploiting the panel nature of my data, I examine the role

Table 16. Linear Probability Model Estimates, Covariates Affecting Probability of Assignment, Classified Households (Main Subsample)

	(1)
Total Credit	-0.0011***
	(0.0001)
Transactions	0.0022***
	(0.0000)
Bulk Shopper	-0.0041***
	(0.0002)
Rural	-0.0025***
	(0.0002)
Mixed	-0.0024***
	(0.0003)
Grocery	-0.0174***
	(0.0006)
Convenience	-0.0244***
	(0.0009)
Combination	-0.0571***
	(0.0009)
Specialty	-0.0570***
	(0.0040)
Other Stores	-0.0204***
	(0.0013)
Tenure	0.0012***
	(0.0000)
R-squared	0.0416
Households	75,482
Benefit Months	1,839,745

The sample used is the classified sample obtained from the mixture analysis using the main subsample. Of the 75,482 households in the sample, 6,985 belong to the slow spenders group and 68,497 belong to the fast spenders group. The dependent variable is a binary indicator identifying households who are fast spenders. The classified households generate 1,839,745 benefit months. (*10%,**5%,***1%).

of two features of the SNAP program that could plausibly have important repercussions for the SNAP benefit cycle, namely, that SNAP benefits can rollover and further that SNAP households must recertify periodically to remain eligible for receiving SNAP benefits. I find that my results remain qualitatively the same, when I account for household saving behavior due to the rollover feature, or household tenure in the program.

Overall, my mixture results offer both positive and negative news. The positive news are that the slow group of SNAP spenders appear to budget their SNAP benefits throughout the month. Although, they only spend 9% of their benefit on average within the first 7 days, they spend approximately 80% throughout the full month. Further these households, receive a lower benefit, on average, compared to the fast spending households, \$182 and \$317 respectively. The negative news is that the overwhelming majority of SNAP households are fast spenders. These households spend a significant percentage of their benefit within the first 7 days, despite receiving a larger benefit on average, and further spend a longer period of time in the program, on average, 3 years and 2 months compared to 1 year and 4 months for the slow households.

Given the existing evidence linking the SNAP benefit cycle, to lower nutrient intake throughout the benefit month and potential food insecurity, these results provide strong justification for research focused on examining the SNAP cycle in greater detail and identifying potential policy levers to mitigate the cycle. Previous studies have suggested a number of potential avenues. First, creating educational programs focused on providing information to assist SNAP households in budgeting SNAP benefits might be effective. Potential tools that could encourage smoother spending and consumption include consumer education in food budgeting, nutrition, sources of affordable and healthy foods such as government food programs or local organizations, encouraging the use of grocery lists, among other factors. In recent years, the nutrition education component of SNAP (SNAP-Ed) has been significantly funded to provide SNAP households with guidance on obtaining healthy and affordable eating options. Future research could examine whether these efforts have yielded any gains in reducing the SNAP benefit cycle.

Another policy option would be to increase the frequency of SNAP benefit disbursal, or at the very least, allow households to opt into more frequent SNAP benefit disbursal. This policy has been suggested in a number of studies (Shapiro (2005); Hastings and Washington (2010); Todd (2015); Dorfman et al. (2019)). Further, there is some anedoctal evidence to indicate that some SNAP households might indeed prefer a more frequent disbursal schedule (Fraker et al. (1995)). While previous studies have examined the costs and benefits of a more frequent disbursal schedule, such as Shapiro (2005), the current SNAP program and its features might warrant a reexamination of this policy and its potential costs and benefits.

A third policy option would be to increase the size of SNAP benefits. SNAP benefits are determined using the TFP, a low cost nutritionally adequate food plan that varies by household size and composition (Gundersen et al. (2019)). However, until recently, researchers argued that the TFP was long overdue for an update as it failed to reflect current economic circumstances and meet SNAP households dietary and food needs (Yaktine and Caswell (2013); Ziliak (2016); Gundersen et al. (2019)). Todd (2015) provides results that support this reasoning. Specifically, examining the temporary 14% increase in benefits due to ARRA, Todd (2015) finds that this increase eliminated the end-of-month decline in caloric intake experienced across her sample of SNAP households. As a result of an update in the TFP in 2021, there has been a 21% increase in maximum SNAP benefits (U.S. Department of Agriculture (2021)). Future studies could evaluate the consequences of this benefit update to the SNAP benefit cycle.

Finally, it is important to recognize that the SNAP benefit cycle is not only due to the institutional features of the SNAP program, but likely also caused by household specific characteristics and circumstances. In fact, Wilde and Ranney (2000) find that demographic and geographic factors such as urban residence, residence in the South, female headship and distance to major grocery store all correlate positively with the cycle. Thus, it is evident that any policy efforts aimed at targeting the SNAP benefit cycle must consider the complex interplay of the myriad factors that underlie its existence. As the SNAP benefit cycle is closely tied to the specific circumstances of the households, be those financial, geographical or behavioral, ideal policies might be the ones, that in the absence of external interference, SNAP households would choose for themselves.

Chapter II: A Flexible Model of Food Security: Estimation and Implications for Prediction

Note: This chapter represents coauthored work with Will Davis, Rusty Tchernis and Christian Gregory. Funding for this project comes from a cooperative agreement with the United States Department of Agriculture's Economic Research Service.

2.1 Introduction

The United States Department of Agriculture (USDA) defines food security as "consistent, dependable access to enough food for active, healthy living" (Rabbitt et al. (2024)). In 2023, approximately 86.5% of U.S. households were food secure. The remaining 13.5%, roughly 18.0 million households, were food insecure, implying that they experienced difficulty during the past year providing nutritionally adequate food for all household members (Rabbitt et al. (2024)). In addition to the immediate nutritional consequences of food insecurity, the condition represents a vital issue for public health given its associations with poor chronic health outcomes like obesity, diabetes, and hypertension (Gundersen and Ziliak (2015)). Accurate measurement of food security allows researchers, policymakers, and other stakeholders to assess the prevalence of food insecurity, evaluate the evolving needs of food insecure households, and measure the effectiveness of related programs like the Supplemental Nutrition Assistance Program (SNAP) (Bickel et al. (2000)).

Household food security is a complex latent construct, making it challenging to fully capture with a single observable variable. To determine household-level food security, researchers use multiple indicators related to associated conditions, behaviors, and experiences (Barrett (2010) , Opsomer et al. (2002), National Research Council (2006)). To account for the domains of food security, the USDA uses a household-level survey instrument called the core Food Security Module (FSM). The FSM was first piloted in the 1995 Current Population Survey (CPS) Food Security Supplement (FSS) and is now the most commonly used survey instrument for U.S. food security measurement.[23] The FSM includes a 10-item adult food security questionnaire completed by every responding household and an additional 8-item questionnaire

[23]For a brief history of module's development, *see* Smith and Gregory (2023).

used to measure child food security for households with children. The experiential FSM questions cover a range of conditions and behaviors that demonstrate varying levels of severity. These include household members being worried about food running out before having money to purchase more, not being able to afford balanced meals, and going a full day without eating due to inadequate food resources.[24] When analyzing FSM responses, many researchers use the USDA's food security scale to assign each household to a food security category. These USDA categories are based on a household's number of affirmative FSM responses.[25] Each household is assigned to one of four possible adult food security categories using the 10 adult FSM questions. The scale is extended to include three child food security categories using the 8 child FSM questions.

There are two main benefits to measuring food security using the FSM and USDA scale categories. First, since the FSM includes at most 18 items, it can be added to new or existing surveys without substantially increasing respondent and enumerator burden. This burden can be further reduced by screening households. Specifically, households that respond negatively to questions about less severe conditions or behaviors at the beginning of the FSM questionnaire are not asked questions related to more severe conditions/behaviors. In some surveys like the CPS, the inclusion of an additional external screener prevents households with certain characteristics from participating in the FSM entirely.[26] Second, the nature of the FSM scale categories simplifies calculation, interpretation, and identification. More specifically, assigning households to categories simply involves the summation of affirmative FSM responses. This user-friendly approach makes the concept of food security accessible and convenient for researchers,

[24] For a full list of FSM questions, *see* Rabbitt et al. (2024).

[25] In this context, an affirmative response is one which indicates an undesirable outcome with regards to food security, implying that households with lower food security (greater food insecurity) are more likely to respond affirmatively to FSM questions. As discussed in Section 2.2.1, we code FSM responses in our model so that higher numerical values indicate higher food security.

[26] The CPS FSS contains external and internal screening protocols to reduce respondent survey burden and avoid questions that may be inappropriate given information provided earlier in the survey. Households that pass the external screen are exempt from taking the entire FSM and they are automatically assigned to the USDA's full food security category. Furthermore, the adult FSM questionnaire in the CPS includes two internal screeners. Households that register no food stress in their responses to a given set of adult internal screener questions are not asked the remaining adult FSM questions, and their responses to skipped questions are assumed to be negative. The National Health and Nutrition Examination Survey does not employ an external screener, implying that no households are exempt from taking the entire FSM questionnaire. It does, however, employ two internal screeners in the adult food security questionnaire.

policymakers, and other stakeholders. FSM scale categories also provide easily interpreted and relevant targets for policies focused on improving household food security.

While the USDA FSM scale categories are deterministic, the scale itself was developed using estimates from a probabilistic 1-parameter logistic Rasch model (1-PRM). The 1-PRM belongs to a class of methods from Item Response Theory (IRT) which are often used to measure latent constructs like test performance and psychological traits (Rasch (1960), Hambleton (1989), Baker (2001), Andrich and Marais (2019)). The IRT modeling approach has several strengths. Firstly, the 1-PRM has been widely employed in the field of education research, resulting in a comprehensive understanding of its statistical properties. Researchers were therefore able to consult IRT experts to confirm that measuring food security with the 1-PRM was consistent with IRT methods and produced reasonable results (Ohls et al. (2001)). Second, using the 1995-1997 CPS waves, the scale's original creators demonstrated that 1-PRM parameter estimates were robust across time and population subgroups (Ohls et al. (2001)). Finally, the 1-PRM can be used when households have incomplete FSM responses, whether the missing responses are due to forced skip patterns from screeners or random non-response (Opsomer et al. (2002)).

Despite these strengths, the 1-PRM has limitations. First, the dichotomous 1-PRM can only be used with binary "affirmative/negative" type response variables. Some FSM questions like "In the last 12 months, did you ever cut the size of any of the children's meals because there wasn't enough money for food?" can only be answered as "yes" or "no", but 9 of the 18 FSM questions have three ordered response options.[27] These three-response-option FSM questions must first be converted into binary variables for use in the 1-PRM, removing potentially valuable information about the severity of each household's underlying food security. The implications of only using binary variables for all FSM responses has been discussed often in the food security measurement literature (most notably in National Research Council (2006) and Nord (2012)), but to the best of our knowledge, Nord (2012), Tanaka et al. (2020), and Opsomer et al. (2002) are the only other studies to estimate food security without converting polytomous FSM questions into

[27] Throughout the paper, we refer to "yes/no" response option questions as binary variables/questions and use polytomous for FSM questions with three potential responses.

binary variables. Second, while the FSM scale's creators showed that the 1995 estimates were stable across time using the two following waves of CPS data, the food security environment of U.S. households may have changed since 1995. Acknowledging this point, the scale's original creators state "…we recommend continuing to estimate the IRT model item parameters each year, as data become available".[28] This recommendation implies that estimating the underlying measurement model once and then using parameter estimates to derive deterministic categories may be problematic if parameters need to be continually updated. This recommendation also highlights the need for ongoing refinement of food security measurement methods to ensure validity and relevance. Finally, as discussed in National Research Council (2006), the 1-PRM does not provide a direct measure of uncertainty for its food security estimates, ignoring the inherent uncertainty of latent trait estimation.

In light of the strengths and weaknesses of the USDA's FSM scale and 1-PRM, we propose an alternative food security measurement model which we estimate using responses to the 10 adult FSM questions.[29] Specifically, we estimate adult food security at the household level using a Bayesian Graded Response Model (BGRM).

Our BGRM has several attractive features applicable to food security measurement. First, the BGRM estimates latent food security as a continuous variable. A continuous measure allows for a more detailed understanding of each household's underlying food security level that can supplement valuable information from food security categories. Second, our standard BGRM specification accommodates any mix of binary and polytomous ordered response variables, capturing information regarding response severity present in the 5 polytomous adult FSM questions. Third, as a Bayesian model, our BGRM estimates latent food security as a set of draws from each household's conditional posterior distribution of food security. As specifically recommended by National Research Council (2006), estimating food security as draws from a posterior distribution of possible values provides an inherent way to measure uncertainty in food

[28]For a detailed discussion of the 1-PRM's estimation and stability across time, *see* Ohls et al. (2001).

[29]In this study, we use the 10 adult FSM questions to estimate households' adult food security. We do not include the 8 child FSM questions answered by households with children. All 18 FSM questions can be included in our model without modification, but adult and child food security are thought to represent distinct latent traits.

security estimates along with point estimates. This approach allows us to easily produce household-level measures of food security uncertainty, measures which can then be used in post-estimation analysis.

We demonstrate the performance and features of our BGRM by estimating household-level adult food security with data from the 2017 - 2018 waves of the National Health and Nutrition Examination Survey (NHANES). We use the NHANES for two primary reasons. The first is that the NHANES contains the full FSM. Second, the NHANES includes continuous Monthly Food Spending (MFS) at supermarkets and grocery stores for each household which is absent from other surveys. In a demonstration exercise, we add continuous MFS into our model along with the set of binary/polytomous adult FSM questions, highlighting how our model can be altered to allow for continuous variables and the accompanying changes to its estimation algorithm.

Our NHANES estimates highlight several key findings. First, as expected, our results show significant uncertainty associated with household-level food security estimates. We find higher uncertainty among adults assigned to the USDA's marginal food secure category and lower uncertainty for food insecure households. Second, we find overlap in BGRM food security point estimates across standard USDA-defined categories. These findings align with previous research suggesting that probabilistically assigning households to multiple categories may be appropriate (National Research Council (2006), Nord (2012)), something that can easily be done with our household-level measures of uncertainty. Third, we find significant variation in BGRM food security point estimates within USDA-defined food security categories. This within-category variation suggests differences in the range of food security levels for households assigned to the same USDA category. Finally, following (Gundersen (2008), Balistreri (2016)), we demonstrate one potential application of our continuous food security point estimates by calculating three Foster, Greer, and Thorbecke (FGT) (Foster et al. (1984)) indices, which capture the prevalence, depth, and severity of food insecurity. Our results indicate that the BGRM produces significantly lower values for all three FGT indices compared to those calculated using 1-PRM Rasch scores for the same sample. This suggests that our model identifies fewer households in the NHANES sample as food insecure and estimates lower levels of severity than the USDA's official measure based on the binary 1-PRM, further highlighting the impacts of alternative modeling assumptions.

While our primary results come from the NHANES, the USDA's national food security estimates are measured using the CPS FSS. Given that both the CPS and NHANES represent large samples of U.S. households, our BGRM should produce qualitatively similar estimates using both surveys. Additionally, noteworthy results found using CPS data may have more direct implications for government agencies and policymakers in the U.S. given their official uses. In Appendix E., we estimate our BGRM using data from the 2018 CPS FSS. Qualitatively, we find BGRM estimates for the CPS that are strikingly similar to those from the NHANES, highlighting the robustness of our model and the FSM questionnaire itself.

We also conduct an example exercise with the NHANES demonstrating our model's capacity to incorporate continuous variables alongside binary and polytomous ones. Specifically, we include MFS alongside the set of adult FSM questions to estimate a novel latent construct. This exercise produces several noteworthy findings. First, the distributional characteristics of our estimates with MFS are qualitatively similar to those derived from the FSM questions alone. However, while the BGRM point estimates for many households change minimally with the inclusion of MFS, some households see significant changes, reflecting the possible influence of adding a variable that captures economic access to food to the experiential FSM. Importantly, however, we note this exercise is intended solely as a demonstration of our model's flexibility rather than as a proposal for a validated new measure. Further theoretical exploration would be necessary to validate and clarify the interpretive meaning of any new latent construct, something outside the scope of this initial methodologically focused study. Nonetheless, the ability to incorporate any mix of binary, polytomous, and continuous variables is, to the best of our knowledge, a unique feature of our BGRM compared to alternative IRT models used in food security measurement, further widening the set of possible variables for future measure construction.

Overall, using our BGRM to estimate food security produces several important insights. First, our findings indicate that there is significant uncertainty in household food security estimates, something that has been recognized by many studies in the existing literature (e.g., Nord (2012), National Research Council (2006)). While previous studies have recognized this issue, existing measurement models do not produce measures of latent trait uncertainty as a natural part of model estimation. Our model's ability to provide household-level measures of uncertainty in food security estimation provides an attractive way to account for this uncertainty inherent to latent trait estimation. Second, the considerable overlap in our estimates across common USDA-defined food security categories highlights the importance of carefully considering the implications of deterministic, rather than probabilistic, category assignment for researchers and policymakers alike. Third, by exploiting the BGRM's ability to incorporate a mix of different variable types, we can fully integrate information from all FSM questions, both binary and polytomous. Furthermore, with our demonstration exercise using both adult FSM responses and MFS, we illustrate how our BGRM can effectively capture information from continuous variables adjacent to food security like food spending, something that current models cannot do. For future research, our BGRM represents a valuable tool for integrating related concepts like economic access to food into augmented food security measures, or for estimation of new latent traits like nutrition security that will most likely require some mix of binary, polytomous, and continuous variables (Seligman et al. (2023)), all without necessitating significant adjustments to the model's structure.

2.2 The Bayesian Graded Response Model of Food Security

Before discussing the model's full mathematical description, we begin with an intuitive overview of our BGRM.[30] For simplicity, we start with the case of binary response variables which shares the same logic as the polytomous case. Our BGRM assumes that each household's set of FSM responses is partially determined by their value of a continuous latent food security variable, δ. As a latent variable, δ is never directly observed. Given how we code FSM responses

[30]For readers generally interested in Graded Response Models and Bayesian Item Response Theory, *see* Samejima (1968), Fox (2010), and Johnson and Albert (2006).

(see Footnote 36), households with higher levels of δ are more likely to report higher numerical response values across all FSM questions, indicative of higher food security.

For the binary case, each household can only provide a response of “yes” or “no” to each FSM question. Similar to the latent variable specification of a standard probit model (Muthén (1979)), our BGRM assumes that a household’s observed binary response, y, to each FSM question is determined by an unobserved (latent) continuous response variable, y^*, which reflects underlying food security. In turn, the value of y^* is a function of the household’s value of δ, a question specific intercept parameter, μ, and an error term, e. Each household has a J length vector of y^* values where J is the total number of FSM questions. With a household’s value of y^*, they respond to the specific FSM question with a “yes” or a “no” depending on if their value of y^* is above or below some interior response threshold parameter, γ. When including polytomous questions, the number of interior γ threshold parameters to identify increases by one, allowing for three response ranges in the distribution of y^* rather than two.

As a Bayesian model, we use a set of conjugate uninformative prior distributions, presented in Section 2.2.1, and the model’s likelihood function to define each parameter’s conditional posterior distribution.[31] Each posterior distribution is conditional on the value of all other model parameters and the set of observed binary and polytomous FSM responses. Unlike the frequentist approach which estimates a single value for each parameter with an associated sampling distribution, the BGRM produces a series of draws from each parameter’s conditional posterior distribution. The draws of each parameter taken across the iterations therefore serves as the model’s final estimates. Point estimates for each parameter can then be calculated using the posterior distribution mean taken across all draws.

[31] A conjugate prior distribution is a prior that simplifies Bayesian updating by being mathematically compatible (or “conjugate”) with the model’s likelihood function. This approach allows each parameter’s conditional posterior distribution to have the same functional form as its prior distribution. An uninformative prior, also commonly called a non-informative or weakly informative prior, is chosen to reflect minimal prior knowledge about the parameter being estimated, so that the posterior distribution is mostly determined by the observed data rather than prior knowledge. For example, specifying a normal prior distribution with a very large variance would be considered an uninformative prior.

We use a Markov Chain Monte Carlo (MCMC) algorithm to estimate our model, specifically a Gibbs Sampler algorithm.[32] While the form of these conditional posterior distributions are defined prior to estimation, sampling from each first requires specifying initial values for the model's parameters. These pre-specified parameter values serve as starting points which are then continually updated across iterations until the model converges. Starting with these initial values, a series of draws from each parameter's posterior distribution is taken sequentially across many iterations.[33] In each iteration, a new draw is taken from the parameter's conditional posterior distribution. This new draw then serves as an updated value of the parameter in the following steps of the iteration until a new, updated draw of the parameter is taken in the next iteration. This sequential updating procedure is repeated until the desired number of total iterations is achieved and the set of conditional posterior distributions stabilize, known as convergence. We remove draws from early Gibbs Sampler iterations in a process known as "burn in" to ensure that the model's posterior parameter distributions have had adequate time to converge and stabilize. The set of post-burn-in draws from each parameter's conditional posterior distribution then serve as our estimates.[34] The means of these conditional posterior distribution draws are used as parameter point estimates, and their 95% credible intervals are used as a common measure of estimation uncertainty.[35]

[32] Casella and George (1992) provide a more detailed, yet still intuitive, overview of the Gibbs Sampler algorithm and the conditions required for model convergence.

[33] Assuming that the conditions for convergence are satisfied, the only notable implication of using different starting parameter values is the number of iterations needed for the model to converge. Intuitively, choosing starting parameter values closer to the parameters' post-convergence posterior means reduces the number of iterations required for convergence.

[34] Unlike frequentist methods, Bayesian estimation does not produce a single estimate for each parameter with an associated sampling distribution; rather, the entire process of iterative sampling is how the model is estimated. Thus, in the Bayesian context, estimation and sampling are inseparable.

[35] In Bayesian statistics, a credible interval represents a range of credible parameter values calculated using draws from the parameter's conditional posterior distribution. The 95% credible interval specifically represents the range that includes 95% of post-burn-in parameter draws. In other words, the parameter's value falls outside the 95% credible interval for only 5% of draws. Larger 95% credible intervals imply greater levels of uncertainty.

2.2.1 Model Specification

Moving to the full mathematical specification of our BGRM, we use households' adult FSM questionnaire responses to measure their latent food security. Our BGRM includes the following parameters and variables:

1. y_{ij} denotes the discrete response of household $i = 1, \dots, I$, to FSM question $j = 1, \dots, J$. Each question j has a set of discrete, and ordered, potential response options ranging from 1 to C_j, implying that $y_{ij} \in \{1, \dots, C_j\}$.

2. The variable δ_i represents the latent food security level of household i which is invariant across the set of FSM responses provided by the same household.[36]

3. The vector of question-specific intercept parameters $\boldsymbol{\mu} = \mu_1, \dots, \mu_J$.

4. The vector of response threshold parameters for each question j denoted by $\boldsymbol{\gamma^j} = \gamma_0^j,\ \gamma_1^j,\ \dots,\ \gamma_{C_j}^j$, such that $\gamma_{k-1}^j \leq \gamma_k^j$ for all $k = 1, \dots, C_j$.

5. Random variation in the elicited food security level of household i with respect to question j, e_{ij}, which captures error in each question's ability to reflect latent food security.

With these variables, the model of household-level FSM responses is defined as follows. Household i answers FSM question j with a response of $y_{ij} = k$ if and only if:

$$\gamma_{k-1}^j < \mu_j + \delta_i + e_{ij} \leq \gamma_k^j \tag{6}$$

[36]Affirmative responses to FSM questions (e.g., a "yes" response rather than a "no" response) are designed to be negatively related to food security. To interpret δ as a measure of latent food security rather than food insecurity, we code the numerical response options of all FSM questions such that higher numerical response values (i.e., higher values of y_{ij}) correspond to higher food security levels (higher δs). This approach aligns with the conceptualization of food security as a continuum. For USDA FSM categories, the continuum ranges from full food security (analogous to the highest values of δ in our model) to very low food security (lowest values of δ). Since this is only a matter of interpretation, however, response values could easily be coded such that higher values are indicative of lower food security rather than higher, creating a measure of food insecurity.

We assume that the distribution of each e_{ij} is Normal with mean 0 and variance σ_j^2. Furthermore, for the purpose of identification, we restrict the value of the first and last response threshold parameter for each question j such that $\gamma_0^j = -\infty$ and $\gamma_{C_j}^j = \infty$.

The model of FSM responses given by (6) implies that the probability household i provides a response k to FSM question j is such that:

$$P(y_{ij} = k) = \Phi\left(\frac{\gamma_k^j - (\mu_j + \delta_i)}{\sigma_j}\right) - \Phi\left(\frac{\gamma_{k-1}^j - (\mu_j + \delta_i)}{\sigma_j}\right) \tag{7}$$

where $\Phi()$ is the CDF of the standard Normal distribution.

The response probabilities given by (7) describe the likelihood function of a graded response model with a probit link function, written formally as:

$$L(\boldsymbol{\mu}, \boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2) = \prod_i \prod_{j \in \mathscr{C}_i} \left[\Phi\left(\frac{\gamma_{y_{ij}}^j - (\mu_j + \delta_i)}{\sigma_j}\right) - \Phi\left(\frac{\gamma_{y_{ij}-1}^j - (\mu_j + \delta_i)}{\sigma_j}\right)\right] \tag{8}$$

The outer product of (8) extends over all households in the sample, and the inner product extends over the set of questions j answered by household i, denoted by $\boldsymbol{\mathscr{C}_i}$.[37]

Examining the model's likelihood in (8), we see that not all model parameters are identified. As is the case with many ordinal response models, the value of the likelihood does not change with affine transformations of each question's response threshold parameters, intercepts, error variances, or latent food security levels. To address this, we restrict $\gamma_1^j = 0$ for all FSM questions j and make the normalizing assumption that $\delta_i \sim N(0, 1)$ which serves as the prior distribution for δ.

[37] As indicated by $\boldsymbol{\mathscr{C}_i}$, our model does not assume that all households provide responses to every FSM questions. This allows for partial non-responses which is a vital feature as many surveys screen respondents, allowing them to skip some or all FSM questions.

The remaining prior distributions used in our model for μ, σ^2, and γ are given as:

$$\mu_j \sim N(\bar{\mu}, \sigma_\mu^2),\ \forall j = 1, \dots, J \tag{9}$$

$$\sigma_j^2 \sim IG(\alpha, \beta),\ \forall j = 1, \dots, J \tag{10}$$

$$\gamma_k^j \sim U(a, b)1(\gamma_{k-1}^j < \gamma_k^j \leq \gamma_{k+1}^j),\ \forall k = 2, \dots, C_j - 1 \tag{11}$$

where $IG()$ denotes the inverse gamma distribution and $U()$ denotes the uniform distribution. Our set of prior distributions, including the normal prior placed on latent food security, represent conjugate prior distributions. In turn, conjugate prior distributions produce conditional posterior distributions that correspond to known probability distributions. This approach makes sampling from the posterior distributions substantially more tractable and computationally efficient.

Exploiting a unique feature of Bayesian models, we use a process known as data augmentation to define a new latent variable, y^*, such that for household i and FSM question j:

$$y_{ij}^* = \mu_j + \delta_i + e_{ij} \tag{12}$$

As discussed in our intuitive model overview, the value of y_{ij}^* determines the value of observed response, y_{ij}. Specifically, from (6) and (12), $y_{ij} = k$ if and only if:

$$\gamma_{k-1}^j < y_{ij}^* \leq \gamma_k^j \tag{13}$$

which gives the following augmented likelihood function over our model parameters and augmented response variable, y^*:

$$L(\boldsymbol{\mu}, \boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\sigma}^2, \boldsymbol{y}^*) = \prod_i \prod_{j \in \mathscr{C}_i} \phi\left(\frac{y_{ij}^* - (\mu_j + \delta_i)}{\sigma_j}\right) 1(\gamma_{y_{ij}-1}^j < y_{ij}^* \leq \gamma_{y_{ij}}^j) \tag{14}$$

where 1() represents an indicator function. Estimation of the model's structural parameters and latent variables is performed using the MCMC sampling algorithm outlined in the following sub-section.

2.2.2 Estimation Algorithm and Model Features

We now present the estimation algorithm used to estimate our model before briefly discussing results of an initial simulated data exercise and some of our model's key features/limitations. We use the following Gibbs Sampler algorithm to sample draws from the conditional posterior distributions of the model's parameters and latent variables. Each step of the algorithm outlined below is performed sequentially in every iteration. An iteration is completed after Step 5, at which point the next iteration begins starting at Step 1 using the parameter draws from the previous iteration, continually updating their values.

Step 1. Sample $\boldsymbol{\delta}$. For $\boldsymbol{\delta} = [\delta_1, \ldots, \delta_I]'$, let the estimation equation be given as:

$$\boldsymbol{y}^* - \boldsymbol{\mu} \otimes 1_{\boldsymbol{I}} = \boldsymbol{\Lambda}\boldsymbol{\delta} + \boldsymbol{e}$$

where $\boldsymbol{y}^* = [y_1^*, \ldots, y_J^*]'$ is a stacked vector of y^*'s, $1_{\boldsymbol{I}}$ is a $I \times 1$ vector of 1's, $\boldsymbol{\Lambda} = I_I \otimes 1_{\boldsymbol{J}}$, $\boldsymbol{I_I}$ is an $I \times I$ identity matrix, $1_{\boldsymbol{J}}$ is a $J \times 1$ vector of 1's, $\boldsymbol{e} = [e_1, \ldots, e_J]'$, and $\otimes$ denotes the Kronecker product.

All elements of $\boldsymbol{\delta}$ are drawn simultaneously from the following full conditional posterior distribution:

$$\boldsymbol{\delta}|\boldsymbol{\mu}, \boldsymbol{\sigma}^2, \boldsymbol{y}^*, \boldsymbol{Y} \sim N(\boldsymbol{d}, \boldsymbol{D}) \tag{15}$$

where $\boldsymbol{D} = \left[\boldsymbol{I_I} + \boldsymbol{\Lambda}'(\boldsymbol{I_I} \otimes \boldsymbol{\Sigma}^{-1})\boldsymbol{\Lambda}\right]^{-1}$, $\boldsymbol{d} = \boldsymbol{D}\left[\boldsymbol{\Lambda}'(\boldsymbol{I_I} \otimes \boldsymbol{\Sigma}^{-1})(\boldsymbol{y}^* - \boldsymbol{\mu} \otimes 1_{\boldsymbol{I}})\right]$, and $\boldsymbol{\Sigma}$ is a $J \times J$ variance-covariance matrix with diagonal elements $(\sigma_1^2, \ldots, \sigma_J^2)$ and zeros for the off diagonal elements.

Step 2. Sample Elements of $\boldsymbol{\mu}$. For each element of $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_J]'$, the estimation equation is given as:

$$\boldsymbol{y}_j^* - \boldsymbol{\delta} = 1_{\boldsymbol{I}}\mu_j + \boldsymbol{e}_j$$

where $\boldsymbol{y}_j^*$ is a stacked vector of length I containing all households' continuous response variables for question j.

Each μ_j is then drawn from the following full conditional posterior distribution:

$$\mu_j | \boldsymbol{\delta}, \boldsymbol{\sigma}^2, \boldsymbol{y}^*, \boldsymbol{Y} \sim N(v, V) \tag{16}$$

where $V = \left(\frac{1}{\sigma_\mu^2} + \frac{1_{\boldsymbol{I}}' 1_{\boldsymbol{I}}}{\sigma_j^2}\right)^{-1}$ and $v = V\left[\frac{\bar{\mu}}{\sigma_\mu^2} + \frac{1_{\boldsymbol{I}}'(\boldsymbol{y}_j^* - \boldsymbol{\delta})}{\sigma_j^2}\right]$.

Step 3. Sample Elements of $\boldsymbol{\gamma}$. Samples for each element γ_k^j of $\boldsymbol{\gamma^j} = [\gamma_0^j, \gamma_1^j, \ldots, \gamma_{C_j}^j]$, such that $k = 2, \ldots, C_j - 1$, are drawn from the following full conditional posterior distribution:

$$\gamma_k^j | \gamma_{-k}^j, \boldsymbol{y}^*, \boldsymbol{Y} \sim U(L, R) \tag{17}$$

where $L = max[max(\boldsymbol{y}_j^* | y_{ij} = k), \gamma_{k-1}^j]$, and $R = min[min(\boldsymbol{y}_j^* | y_{ij} = k + 1), \gamma_{k+1}^j]$.

Step 4. Sample Elements of $\boldsymbol{y}^*$. Samples for each element y_{ij}^* of $\boldsymbol{y}^* = [y_1^*, \ldots, y_J^*]'$ are drawn from the following full conditional posterior distribution:

$$y_{ij}^* | \boldsymbol{\mu}, \boldsymbol{\delta}, \boldsymbol{\sigma}^2, \boldsymbol{\gamma}, \boldsymbol{Y} \sim N\left(\mu_j + \delta_i, \sigma_j^2\right) \tag{18}$$

truncated to the interval $(\gamma_{y_{ij}-1}^j, \gamma_{y_{ij}}^j)$.

Step 5. Sample Elements of $\boldsymbol{\sigma}^2$. For each element of $\boldsymbol{\sigma}^2 = [\sigma_1^2, \ldots, \sigma_J^2]'$, the estimation equation is given as:

$$\boldsymbol{y}_j^* = 1_{\boldsymbol{I}}\mu_j + \boldsymbol{\delta} + \boldsymbol{e}_j$$

Each σ_j^2 is then drawn from the following full conditional posterior distribution:

$$\sigma_j^2 | \boldsymbol{\mu}, \boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{Y} \sim IG(a, b) \tag{19}$$

where $a = \alpha + \frac{I}{2}$, and $b = \left[\frac{1}{\beta} + \frac{(\boldsymbol{y}_j^* - 1_{\boldsymbol{I}}\mu_j - \boldsymbol{\delta})'(\boldsymbol{y}_j^* - 1_{\boldsymbol{I}}\mu_j - \boldsymbol{\delta})}{2}\right]^{-1}$.

With the model and Gibbs Sampler algorithm fully defined, we first verify model performance using simulated data.[38] Specifically, we first simulate a set of binary and polytomous response variables matching the properties of the FSM. We then verify if our model can adequately retrieve the known data generating parameter values. To simulate the data, we define values for each parameter in our BGRM and use the model to generate simulated responses to 10 questions (5 polytomous and 5 binary) for 5,000 units (households). We then estimate our model using the simulated data, running the Gibbs Sampler algorithm for 30,000 iterations with the first 10,000 iterations removed for burn-in. After estimation, we test to see if our model accurately converges to the known data-generating parameter values used to create our simulated data. Detailed information regarding the simulated data exercise and its results can be found in Appendix B., but to summarize, the results of our simulated data exercise suggest that the BGRM performs well under standard parameter retrieval tests used in the related literature (Zhu and Stone (2011), Kieftenbeld and Natesan (2012)).

Our BGRM has several attractive features. First, our model estimates food security as a continuous latent trait, potentially capturing more variation in severity than the USDA's FSM

[38] All model estimation is conducted in the MATLAB programming language. All data and code used in this study are freely available from the authors upon request.

scale categories alone. Having access to a continuous measure also allows us to calculate other measures/indices which we demonstrate in our FGT exercise. While a continuous measure of food security may provide more granular information than discrete categories, it comes at the cost of simplicity and, in some respects, interpretability. More specifically, the four USDA FSM scale food security categories provide a straightforward method for assigning households to easily interpreted, policy relevant, groups. This deterministic approach does, however, rely on the assumption that each category is mutually exclusive rather than "fuzzy", meaning that a household cannot be probabilistically assigned to two or more categories based on their level of uncertainty. In Figure 3, we find overlap in the distribution of BGRM posterior mean food security levels across three USDA categories. While we do not conduct this specific analysis in our study, our BGRM estimates could also be used to create or modify food security categories based on distributional features of the latent trait estimates. Furthermore, as recommended by National Research Council (2006), uncertainty estimates from our model can be utilized to calculate the probability that a household falls within a specific food security category. Probabilistic category assignment with frequentist IRT methods has been explored in Nord (2012). While not probabilistically assigning food security categories, Gregory (2020) uses measures of uncertainty from a dichotomous Bayesian 4-parameter IRT model to probabilistically assign household misreporting probabilities with informative prior distribution.

Second, by not needing to convert polytomous FSM questions into binary variables, we can directly include additional information regarding the severity of related experiences. This issue, among others, was raised by National Research Council (2006) in their report detailing potential modifications to improve USDA food security measurement. With respect to polytomous food security models, our work builds on Nord (2012), Tanaka et al. (2020), and Opsomer et al. (2002) who use the same frequentist polytomous IRT model, the Partial Credit Rasch Model (PCRM), to estimate household food security. Similar to our BGRM, their PCRM can be estimated with a mix of ordered polytomous and binary outcome variables. While Nord (2012) suggests that directly modeling the FSM's polytomous responses leads to little improvement compared to the 1-PRM, Tanaka et al. (2020) find that the PCRM does improve estimation precision, highlighting the potential benefits of moving beyond binary IRT methods. Unlike the PCRM, which is built using the logistic distribution, our BGRM can be interpreted as an extension of the probit model based on the normal distribution.[39] In our Bayesian framework, the normalizing assumption allows us to use conjugate non-informative prior distributions, providing analytically tractable conditional posterior distributions for our model parameters. While we use non-informative priors in our study, informative prior distributions can be used to incorporate *a priori* information similar to Gregory (2020). This unique feature of Bayesian models is especially useful for estimation with small sample sizes.

Third, drawing samples from each household's posterior distribution of δ provides a direct way to measure estimation uncertainty. This approach aligns with a recommendation from National Research Council (2006) which states that the large amounts of uncertainty inherent in food security estimates should not be ignored. Instead, National Research Council (2006) explicitly recommends modeling food security values as a posterior distribution of potential values rather than a single point estimate. This recommendation aligns exactly with our BGRM. Using 95% credible intervals as a measure of uncertainty for each household's food security posterior distribution, our model allows us to capture the general uncertainty present in latent trait estimation. This work builds on Nord (2012) who considers the potential role of uncertainty in his report responding to National Research Council (2006) . In addition to probabilistic category assignment discussed above, measures of uncertainty could also be used in multivariate analysis. As one example, households' 95% credible interval magnitudes could be used as weights in a regression, leading to coefficient estimates that are primarily informed by households with lower levels of uncertainty.

Finally, the purpose of our analysis is not ultimately to create a fixed food security scale or categories, but rather to estimate latent food security through full model estimation. By

[39] As proposed by Bock and Aitkin (1981), general IRT models beyond the 1-parameter model cannot directly calculate the unit-level latent trait parameter without additional assumptions. They propose using Marginal Maximum Likelihood estimation after first assuming a distribution for the parameter. This approach is commonly used in the IRT literature, with many models, including our BGRM, assuming a normal distribution for the latent trait parameter.

estimating the model in each application, we can capture potential changes in the underlying relationship between observable variables and latent food security occurring across time, groups, and data sets. This approach aligns with the suggestion from Ohls et al. (2001) that the underlying food security measurement model be estimated using different surveys and years to detect any notable changes.[40]

In addition to the features discussed above, our methodology shares two of the 1-PRM's greatest strengths. First, our BGRM only uses the set of FSM questions to which a household responds to estimate food security. Doing so allows us to include households with both complete and incomplete responses. This flexibility is valuable in capturing food security for a wider range of households rather than only for households with valid responses to all FSM questions. Second, similar to the 1-PRM, our BGRM determines the relative contribution of each manifest variable to the measurement of latent food security based on data rather than relying on expert opinion like many other scales/indices. This data-driven approach ensures that the food security measure is grounded in empirical evidence based on relationships present in the data.

2.3 Foster-Greer-Thorbecke Methodology

To contextualize the findings from our BGRM and exploit the continuous nature of our model's food security estimates, we calculate three Foster-Greer-Thorbecke (FGT) measures of food insecurity prevalence, depth and severity (Foster et al., 1984). We then compare the BGRM FGT measures to those created using the USDA's 1-PRM Rasch scores. This section contains a general description of the FGT methodology and how it is applied to our BGRM estimates.[41]

In cases where the threshold of a continuous variable holds policy relevance, such as the poverty line or food security classification cutoff, the FGT methodology provides notable benefits. Specifically, the FGT methodology allows for the measurement of not only prevalence (e.g., the percentage of a sample below some threshold), but also depth (average distance from a threshold) and severity (skewness among those under the threshold). The FGT methodology has

[40] Additionally, as discussed in a conference paper by Opsomer et al. (2002), the 1-PRM is overparametrized, implying that comparability of model parameter estimates across data sets requires normalization.

[41] Our FGT analysis is included as an illustration of the model's empirical flexibility, not as a core contribution of the study.

been applied in various contexts in the related literature, including food insecurity (Gundersen, 2008), food expenditure poverty (Yang et al., 2019) , and food stamp benefits and child poverty (Jolliffe et al., 2005).

Analyzing the FGT measures of food insecurity is attractive for several reasons. Many published reports and studies examining food insecurity rely on categorical measures of food insecurity which classify households as food insecure based on whether they fall below or above certain thresholds of food security and insecurity. However these categorical measures do not directly account for differences in food insecurity between food insecure households. By contrast, the FGT based indices of food insecurity utilize the full information, from both the USDA food security Rasch scores and our BGRM measure of food security, to examine variation in the extent and severity of food insecurity among food insecure households.

Let $N = \{1, \dots, n\}$ denote the set of all households under consideration, with n representing the total number of households. For all $i \in N$, s_i denotes the BGRM posterior mean for household i. In the Rasch methodology context, s_i represents a Rasch score for the household, with larger Rasch scores indicating higher food insecurity. However, in our context higher BGRM posterior means indicate higher food security. For consistency between the BGRM and Rasch approaches, we reflect our BGRM posterior means such that higher values are indicative of greater food insecurity when calculating the set of FGT indices.

We assume that for every $i \in N$, $s_i \in [-\infty, z]$, where z represents the BGRM posterior mean associated with the highest possible level of food insecurity, where a household answers affirmatively to all FSM questions. The Rasch model cannot calculate scale scores for extreme households – those that respond either entirely affirmatively or entirely negatively to all FSM questions (Nord (2012), Bickel et al. (2000)). Thus, for the households that answer all questions affirmatively, the standard methodology is to assign them the scale scores corresponding to affirming 9.5 items for households without children. We follow this methodology when assigning Rasch scores to the households that affirm all questions in the FSM in our sample. Further, we adopt a similar approach in the context of the BGRM. In particular, we calculate the difference in the average BGRM posterior means for households with 8 and 9 affirmative responses to the FSM. Subsequently, we add this difference to the maximum BGRM posterior mean of households with 9 affirmative responses and set z equal to this value.[42]

We define e as the benchmark level of food insecurity, and set it equal to the maximum BGRM posterior mean for households with two affirmative responses in the Food Security Module questionnaire. When using Rasch scores to construct FGT measures, e is set equal to the Rasch score for households with two affirmative responses in the Food Security Module questionnaire (Gundersen, 2008). Thus, a household is considered food insecure if and only if $s_i > e$. For every household i, we define the Food Insecurity Index (FII) as $(s_i - e)$ if $s_i > e$ and 0 otherwise. Consequently, households deemed as food secure, $s_i \leq e$, will have FII equal to 0 in all FGT measures. The FII of a household can be thought of as a measure of the degree of food insecurity of the household; in the poverty literature, an equivalent concept would be the individual's "shortfall" from the poverty line. Finally, we define the Normalized Food Insecurity Index (NFII), denoted by d_i, as

$$d_i = \frac{s_i - e}{z - e} \quad \text{if} \quad s_i > e \quad ; \quad d_i = 0 \quad \text{otherwise} \tag{20}$$

We use the NFII of all households to compute aggregate measures of food insecurity for all households. Denote d as the degree of food insecurity experienced by the group of all households. Consistent with previous studies (Gundersen, 2008), we use three common aggregation rules applied to the function d. Following Foster et al. (1984), these aggregation rules can be expressed by varying the value of α in the following equation:

$$d^{\alpha} = \frac{\sum_{i=1}^{n} (d_i)^{\alpha}}{n} \tag{21}$$

[42] The deterministic assignment of trait boundaries for extreme response households is a known limitation of the FGT methodology itself and is not specific to the BGRM framework.

When $\alpha = 0$, d is defined as the food insecurity rate, which reflects the proportion of households in our sample that are food insecure. When $\alpha = 1$, d is the defined as the food insecurity gap, a measure of the depth of food insecurity. Specifically, the food insecurity gap measures the average shortfall of households relative to the food security threshold value. The gap measure assigns equal weights to food secure and food insecure households, thus it can underestimate the extent of food insecurity among food insecure households when there is a large number of food secure households in the sample. When $\alpha = 2$, d is defined as the squared food insecurity gap, a measure of the severity of food insecurity. The squared food insecurity gap produces a measure similar to the average shortfall of households relative to the food security threshold, but assigns a larger weight to the most food insecure households with higher values of d_i.

2.4 Data

Our primary analysis uses data from the 2017 - 2018 waves of the National Health and Nutrition Examination Survey (NHANES).[43] The NHANES is one of the major programs conducted by the National Center for Health Statistics (NCHS), providing data from a sample of roughly 5,000 individuals from the U.S. in each wave and for 2017-18 includes oversamples of Hispanic persons (inclusive of all races), non-Hispanic black persons, non-Hispanic Asian persons, non-Hispanic whites with incomes less than 185% of the federal poverty line, and Non-Hispanic white persons and persons of other races and ethnicities aged 0–11 years or 80 years and over.[44] The NHANES includes questions regarding the demographic, dietary, socioeconomic, and health-related characteristics of adults and children. In addition to self-reported variables provided by NHANES respondents, trained NHANES medical personnel collect medical, dental, and physiological measurements and administer laboratory tests. The NHANES is used to assess the prevalence of major diseases, disease risk factors, nutritional status, and their associations with health promotion and disease prevention efforts.

[43] While more recent waves of the NHANES are available at the time of this study, we use the 2017-2018 waves to avoid potential data collection or response concerns resulting from the COVID-19 pandemic.

[44] For additional information regarding the 2017 - 2018 NHANES, see NHANES Documentation 2017 - 2018.

Most importantly for our study, the NHANES includes the full FSM questionnaire and household food spending, measured as the monthly dollar amount spent by households at supermarket or grocery stores, namely MFS.[45] MFS is part of the Consumer Behavior questionnaire in the NHANES which aims to provide interview data on food expenditures at the family level. This combination of data allows us to estimate our BGRM (with and without MFS) and assign food security categories based on the USDA's FSM scale for comparison purposes.

In line with the USDA's official 1-PRM, we first remove what are commonly known as "households with extreme responses", which we refer to as "extreme response households" in this study, to create our final sample (Nord, 2012). Extreme response households include households that respond either entirely affirmatively or entirely negatively to all FSM questions. Since IRT models rely on within-household response variation across questions to identify a latent trait, extreme response households cannot be included. As discussed in Nord (2012) for example, the USDA's 1-PRM omits households with extreme FSM responses from the sample prior to estimation as the model cannot identify their latent trait without within-household-variation across responses. Instead, only households with at least some level of response variation are used in estimation.[46]

Our final 2017-2018 NHANES sample consists of 1,659 adults, though there is variation in the number of observations across variables due to missingness. Table 17 provides descriptive statistics for our NHANES sample without extreme response households. At 40%, most adults in the sample are considered marginally food secure under the standard USDA FSM scale categories.[47] The remaining shares of low and very low food secure adults are 37% and 23%,

[45] Specifically, the variable asks: "During the past 30 days, how much money (did your family/did you) spend at supermarkets or grocery stores?"

[46] Alternatively, our model can produce point estimates for extreme response households that match *a priori* expectations (Davis et al. (2024)), but the credible intervals (measure of uncertainty) surrounding these point estimates are substantial given the limited amount of information available aside from which end of the distribution they fall in. This high level of uncertainty implies that the point estimates for extreme response households are largely uninformative even if they were used in estimation. In this initial study, we do not assign latent trait scores to extreme value households post-estimation unlike the USDA's 1-PRM. Instead, we focus on the estimates produced directly by our model for the sample of non-extreme response households. The only exception to this is our FGT analysis which assigns latent trait values to extreme value households in a similar manner to the USDA's 1-PRM.

[47] These summary statistics are calculated without using NHANES-provided weights to account for the survey's complex design, implying that the statistics are not nationally representative.

respectively. For health outcomes, 41% of adults are obese; 17% have been diagnosed with diabetes; and 36% have been diagnosed with high blood pressure. The mean monthly food spending of households (MFS) is $393 with a standard deviation of $297. Females make up 53% of the sample. Hispanic and white are the most commonly reported races in our sample at 30% and 29%, respectively. Blacks, Asians, and other races/ethnicities represent 27%, 8%, and 6% of our NHANES sample, respectively. 82% of adults in the sample are U.S. citizens. 10% of the sample holds a bachelors degree or higher. 32% have attended some college or hold a two-year associates degree. 30% of the sample have a high school degree or equivalent, and the remaining 29% have less than a high school degree. Finally, 39% of the adults in the sample are married. The second largest group by marital status are never married adults at 21%. The remaining respondents are either widowed (7%), divorced (13%), living apart (13%), or separated (6%).

Before discussing results, it is important to note that the NHANES is a complex survey, requiring the application of survey weights to produce nationally representative statistics. Because our BGRM is implemented in a Bayesian framework that does not easily accommodate design-based weights, we do not apply survey weights in our estimation of NHANES sample statistics. As a result, all findings beyond the household-level are illustrative and not nationally representative. Therefore, these results should in no way be interpreted as a replacement for the USDA's official population measures.

2.5 Results

2.5.1 Main Results

Using the full sample of NHANES households without extreme responses, we run our BGRM Gibbs Sampler algorithm for 50,000 total iterations. To ensure adequate time for convergence, we remove the first 25,000 draws from the final set of estimates in a process known as "burn-in", leaving us with 25,000 post-burn-in draws from each conditional posterior distribution. Figure 2 displays the BGRM posterior mean food security (FSEC) estimates for each

Table 17. Summary Statistics 2017 - 2018 National Health and Nutrition Examination Survey (NHANES) Sample, Excluding Extreme Response Households

	Years: 2017 - 2018		
	Mean	Standard Deviation	Count
Food Security			
Marginal Food Secure Households	0.40	0.49	1659
Low Food Secure Households	0.37	0.48	1659
Very Low Food Secure Households	0.23	0.42	1659
Health Outcomes			
Obese	0.41	0.49	1638
Diabetes	0.17	0.38	1619
High Blood Pressure	0.36	0.48	1653
Gender			
Female	0.53	0.50	1659
Food Spending			
Food Spending	393.15	296.69	1659
Race			
Hispanic	0.30	0.46	1659
White	0.29	0.45	1659
Black	0.27	0.44	1659
Asian	0.08	0.27	1659
Other	0.06	0.24	1659
Citizenship			
Citizen	0.82	0.38	1653
Education Level			
Less than High School	0.29	0.45	1573
High School	0.30	0.46	1573
Some College & Associate Degree	0.32	0.47	1573
College Graduate & Above	0.10	0.29	1573
Marital Status			
Married	0.39	0.49	1575
Widowed	0.07	0.26	1575
Divorced	0.13	0.34	1575
Separated	0.06	0.24	1575
Never Married	0.21	0.41	1575
Living Apart	0.13	0.33	1575

Notes: Estimates are unweighted and reflect the sample only. The sample consists of adults from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). The entire sample contains 1,659 adults. We exclude households who answered all questions affirmatively or negatively. Number of observations vary across variables due to missing values. Food security categories follow USDA definitions. Marginal Food Secure households have 1-2 affirmative responses to the adult FSM questionnaire, followed by 3-5, and 6-9 affirmative responses for Low and Very Low Food Secure households, respectively. Food Spending is a continuous variable measuring the monthly food spending of households at supermarket or grocery stores, in dollars. All other variables are indicator variables.

adult arranged in ascending order, along with colored horizontal lines representing the 95% credible interval range of each adult's posterior distribution of FSEC estimates.[48]

Figure 2 highlights several interesting findings regarding the distribution of uncertainty for our household-level latent food security estimates.[49] First, we find that there is non-trivial levels of uncertainty in each household's latent trait estimate. This finding aligns with those from other food security measurement studies in so far as estimation of a latent trait using a limited binary/polytomous response option instrument leading to noisy household-level estimates (e.g., Nord (2012), National Research Council (2006)). With that being said, the observed levels of uncertainty in our estimates are far from large enough to make household-level estimates uninformative. Perhaps more interesting than the magnitudes of uncertainty themselves is the distribution of uncertainty. We find the lowest levels of uncertainty for households in the left tail of the distribution (most food insecure) with higher levels of uncertainty for households with point estimates towards the right tail (most food secure). This finding implies that our model produces more precise estimates for more food insecure households (lower latent trait point estimates). Intuitively, the model can more confidently place food insecure households in the left tail of the distribution, but less so with households in the right tail, households that would typically be considered marginally food secure under the official FSM scale. This finding also aligns with prior literature, such as Coleman-Jensen (2010), which notes that marginally

[48] Estimating our BGRM produces 25,000 post-burn-in simulated FSEC draws from each household's conditional posterior distribution. We therefore use the posterior mean of each household's distribution as our point estimates. 95% credible intervals represent the range in which 95% of a household's 25,000 FSEC values are drawn, representing a measure of estimation uncertainty/variability.

[49] It is important to again note that we do not use survey weights to estimate our model, thereby not accounting for the NHANES' complex survey design. Our estimates are therefore not nationally representative. We did not include survey weights for two primary reasons. First, incorporating survey weights into Bayesian statistical models is still in its relative infancy (e.g., see Si et al. (2020), Gunawan et al. (2020)). There are potential ways to directly incorporate survey weights into the model, but it would involve fundamentally altering the model's likelihood function. Alternatively, researchers can use weights post-estimation to construct descriptive statistics like the mean of latent food security, but this approach does not account for the effects of uncertainty inherent to using weights, so the uncertainty estimates produced by our model would not either. Second, while nationally representative descriptive statistics are key for national monitoring, our primary goal is to describe the BGRM and how it can be used to generate household-level food security estimates. As explored in Solon et al. (2015), using survey weights is not vital for applications focused on individual/household-level analysis.

food-secure households often share characteristics with food-insecure households, complicating their classification.

Figure 2. Posterior Mean Values of BGRM Latent Food Security with 95% Credible Intervals Using NHANES 2017 - 2018

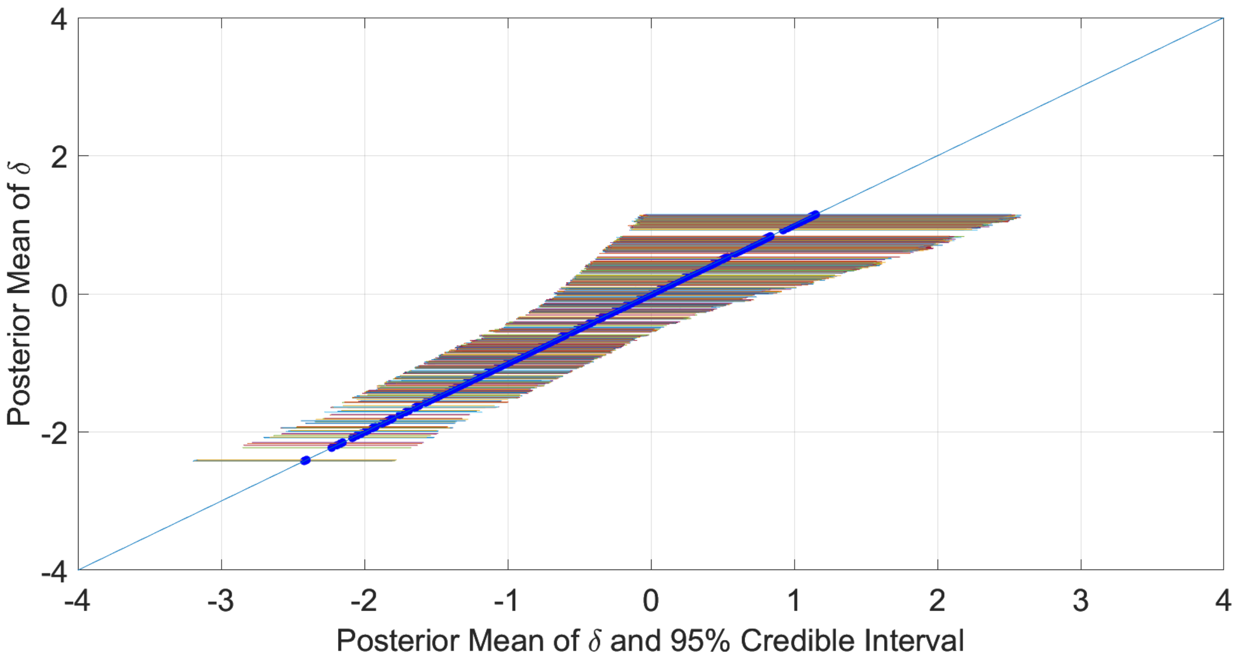

Note: Estimates are unweighted and reflect the sample only. The sample consists of households with responses to the 10 adult FSM questions from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 1,659 households. Each blue dot on the 45^o degree line represents the BGRM posterior mean latent food security value for a given household in the sample. The colored lines running through the blue dots represent the magnitude of the 95% credible intervals for each household.

While understanding the level of uncertainty in our household-level latent food security estimates is important in and of itself, measures of uncertainty can also be used directly or indirectly in multivariate analysis. As just one example, in Column 3 of Table F5, we calculate the probability that each household's latent food security falls in each of the four quartiles of δ's overall distribution. Households with 95% credible intervals that fall entirely within a single quartile range will have a probability equal to 1 of being in that quartile. Households with 95% credible intervals stretching across multiple quartile thresholds, however, will have some probability of falling within 2 or more quartiles. Using this probabilistic approach captures both the position of the household's food security point estimate and the magnitude of their uncertainty. Incorporating measures of uncertainty, like 95% credible intervals, into future multivariate analysis could provide richer insights and more accurate interpretations of food security or similar latent constructs. These posterior distributions allow researchers to directly quantify household-level uncertainty in food security estimates, allowing for applications such as precision-weighted regression analysis, probabilistic classification, and targeted policy interventions for households with high expected insecurity and high uncertainty.

While not explicitly done in our initial study, one specific recommendation for how uncertainty estimates could be used in future analysis is weighting regressions based on uncertainty. Researchers can use the magnitude of each household's 95% credible interval as a weight in regression analysis. Households with narrower credible intervals (i.e., lower uncertainty) would therefore have a larger influence on estimated coefficients, while those with wider credible intervals (more uncertainty) would contribute less. This approach can help control for variability based on estimate reliability, possibly allowing for more robust conclusions when analyzing the relationships between food security and related outcomes of interest. Regardless of the approach taken, our BGRM's ability to directly produce household-level estimates of uncertainty along with point estimates is an attractive feature under many circumstances.

To examine the relationship between our estimates and standard USDA food security categories, Figure 3 shows the distribution of BGRM posterior mean food security estimates separated by USDA-defined food security category. Each household in our NHANES sample is assigned to one of three categories following the USDA's methodology, specifically: marginal food secure, low food secure, and very low food secure.[50] We then plot histograms of BGRM posterior mean FSEC values for households in each category, showing how the distribution of BGRM point estimates varies within and across official USDA categories. As expected, the rank order is consistent across our BGRM point estimates and the USDA food security categories.

[50] The USDA food security categories are defined such that Marginal Food Secure households have between 1-2 affirmative responses to the adult FSM followed by 3-5 and 6-9 affirmative responses for Low and Very Low Food Secure households, respectively. Full Food Secure households, the most food secure USDA category, are not included in this analysis as we remove extreme response households with 0 affirmative responses from the sample. Similarly, Very Low Food Secure extreme response households with 10 affirmative responses are not included. For additional information about USDA food security categories, see Rabbitt et al. (2023).

Figure 3. Distribution of BGRM Posterior Mean Food Security by USDA FSM Scale Food Security Categories Using NHANES 2017 - 2018

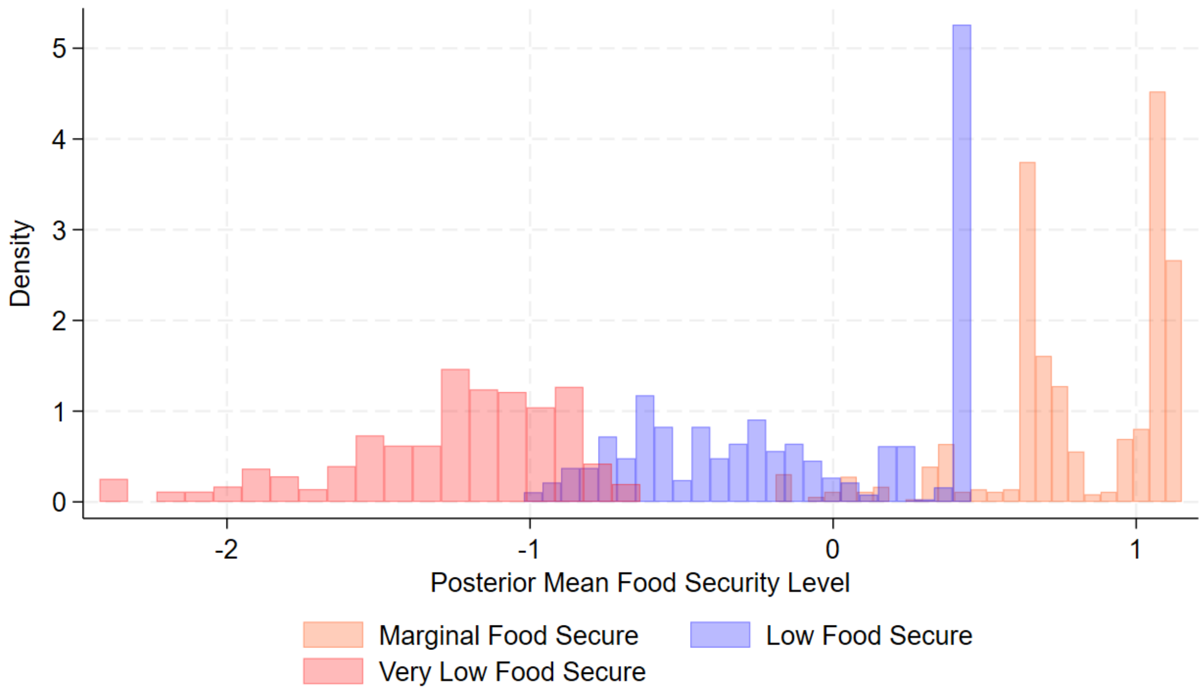

Note: Estimates are unweighted and reflect the sample only. The sample consists of households with responses to the 10 adult FSM questions from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 1,659 households. Each household in our sample has a BGRM posterior mean value of food security that is estimated by the model. Using the responses of the same households to the adult food security questionnaire questions, households are categorized into 1 of 3 food security categories based on the standard USDA food security category definitions. Marginal Food Secure households have 1-2 affirmative responses to the adult FSM questionnaire, followed by 3-5, and 6-9 affirmative responses for Low and Very Low Food Secure households, respectively. Histograms depicting the household distribution of BGRM posterior mean food security values for households belonging to the same FSM food security category are displayed using the 3 colors above.

Adults assigned to the marginal food secure category under the USDA scale have the highest BGRM posterior mean values on average followed by low and very low food secure households, respectively. While the general ordering of our posterior mean estimates agree with the USDA categories, we find evidence of significant variation in BGRM posterior mean values within USDA food security categories, highlighting potential differences in underlying severity levels for households assigned to the same deterministic category. Being able to identify the presence and magnitude of this within-category variation in severity is an attractive feature of our model's continuous food security point estimates.

We also find notable overlap in the distributions of our BGRM posterior means separated into the USDA's marginal, low, and very low food security categories. Traditional, mutually exclusive categories assume no overlap, but our results indicate that such separation may not hold when categories are solely defined by the number of affirmative FSM responses. Specifically, overlap among the lower three USDA categories suggests that relying on a strict category assignment could lead to over- or under-estimation of households' true food security status. For example, a significant portion of households in the left tail of the marginal food secure category have BGRM posterior means that fall below those of households in the right tail of the low food secure category. This finding implies that some households may be misclassified into higher or lower categories than their BGRM posterior means would suggest. This overlap further highlights the potential value of a probabilistic approach to food security categorization, particularly for households with posterior means near category thresholds. Given the uncertainty in each household's posterior distribution of latent trait estimates, households with 95% credible interval ranges stretching across category boundaries might be better represented by a probability of belonging to each category, rather than a single deterministic assignment.

Furthermore, Figure 3 shows some visually identifiable areas of high density in our BGRM estimates within existing USDA categories which could prove meaningful for future research. For example, the distribution of BGRM posterior mean estimates for households in the marginal food secure category appears qualitatively bimodal, implying that the point estimates of households in this category are largely separated into two areas. We also find a significant mass of households in the very right tail of the distribution for low food secure households. Examining these areas of high density may be a promising avenue for future research identifying new potential food security categories or refining existing ones. Additionally, as shown in Figure F4, we find a similar bimodal quality in the marginal food secure category using our CPS sample, implying that is not simply a product of the NHANES sample. This consistency is key as the

USDA uses the CPS to calculate national food security rates each year. While doing so is outside the scope of this study, careful consideration of these distributional features and how they relate to category definitions is valuable given their critical research and policy relevance.

Finally, as shown in Appendix C., we examine the degree to which our BGRM posterior mean FSEC estimates capture variation in health outcomes commonly associated with food security. Specifically, we estimate regressions of three binary variables, self-reported obesity, diagnosis of diabetes, and diagnosis of high blood pressure on our BGRM posterior mean point estimates and a set of variables controlling for gender, age, race, level of education, marital status, citizenship status, household size, and annual household income. We compare the resulting BGRM regression estimates to those from regressions of health outcomes on binary category indicators of low and very low food security defined using the USDA's FSM scale, leaving marginal food secure as the reference group; the number of raw affirmative FSM responses for each adult in the sample; and the respondent-level probabilities of falling within a given quartile of the overall BGRM FSEC posterior distribution. While this regression analysis is not of crucial importance for our initial methods-focused study, it does serve as a validity check of our model's results. For example, if we find that higher food security estimates produced by our model lead to higher rates of obesity, diabetes, or high blood pressure, that would serve as a cause for concern given the deep well of results from the food security literature suggesting the opposite relationship (e.g., Gundersen and Ziliak (2015)). We find no results to suggest counter intuitive associations between our measure and the outcomes of interest. We also find highly similar results for the BGRM and USDA derived measures in our regressions, implying that there is no notable loss in explanatory power from our highly flexible, but more complex, BGRM.

We use the NHANES to estimate our primary results as the survey includes both the health outcome variables and MFS (which we utilize in our Section 2.6 demonstration exercise below). As previously discussed, however, official USDA national food security estimates come from the CPS FSS. As a robustness check we estimate our BGRM using the 2018 CPS FSS and compare the results to those from the NHANES. We discuss this CPS robustness check in more detail in Appendix E..

From Figure F3, a CPS analog to Figure 2, we find significant uncertainty in our CPS estimates similar to the NHANES. The distribution of uncertainty is similar between our CPS and NHANES results, with the largest uncertainty for marginal food secure households. However, we do find that uncertainty is somewhat lower in the middle of the distribution than the left tail, implying that very low food secure households have greater uncertainty than low food secure households on average in the CPS. Perhaps more notably, however, are the distributional features of our CPS point estimates separated by USDA-defined food security categories shown in Figure F4. Overall, the distribution of BGRM point estimates by USDA FSEC category is very qualitatively similar to the NHANES estimates shown in Figure 3. This finding is expected as both surveys provide large samples of U.S. adults. Overall, we view these similarities as a demonstration of robustness, both for our model and the FSM questionnaire itself.

Finally to produce aggregate measures of food insecurity prevalence, depth, and severity in our NHANES sample, we apply the FGT methodology to our continuous BGRM estimates. We compare our results to FGT measures calculated using the USDA's household-level Rasch scores. These results are presented in Table 18 which illustrates several key findings. First, the food insecurity rate calculated using 1-PRM Rasch scores is 23%, compared to 15% with BGRM estimates. Additionally, the food insecurity gap and severity measures are significantly larger when calculated with Rasch scores compared to the BGRM. These differences likely reflect variations in methodological assumptions between both models. Rasch scores rely on binary responses which may overstate food insecurity by rigidly classifying households, particularly those near category thresholds, into more extreme groups. For example, responding to a question with "sometimes" is considered equivalent to a response of "often" when converting FSM questions into binary responses prior to estimating the 1-PRM. In contrast, our BGRM estimates latent food security using full information from the FSM's polytomous response variables, implying that the difference in underlying severity between someone with a "sometimes"

Table 18. FGT Food Insecurity Measures, 2017 - 2018 National Health and Nutrition Examination Survey (NHANES) Sample

	All Households	
Food Insecurity Measure	BGRM	RASCH
Food Insecurity Rate ($\alpha = 0$)	.1543	.2256
Food Insecurity Gap ($\alpha = 1$)	.0546	.0955
Squared Food Insecurity Gap ($\alpha = 2$)	.0307	.0567

Note: Estimates are unweighted and reflect the sample only. The sample consists of adults from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). The entire sample contains 4,738 adults. The BGRM column displays the FGT measures derived from the estimated BGRM Posterior Means of Food Security for all households. The RASCH column displays the FGT measures derived from the Rasch scores for all households in the sample.

response can be differentiated from an "often" response. In other words, our BGRM produces estimates that capture the relative position of households along the food security spectrum and incorporates uncertainty, resulting in potentially conservative FGT measures. Overall, these findings illustrate how different modeling approaches can influence the aggregate characterization of food insecurity, providing useful insights for researchers and policymakers seeking to understand the potential impacts of adopting alternative measurement frameworks. We do emphasize again, however, that these differences do not represent a notable empirical finding, but rather reflects modeling assumptions and classification thresholds. These BGRM estimates should not be interpreted as replacements for the USDA's official food insecurity prevalence. Additionally, while not explored in this initial study, measures of uncertainty from BGRM estimates could be used to identify uncertainty in FGT estimates, allowing for a more nuanced view of food security prevalence, depth, and severity.

2.6 Continuous Variable Estimation Demonstration Exercise: Monthly Food Spending

One of the main methodological advantages of our BGRM is its ability to include any mix of different observable variable types, even continuous variables which is a unique feature. For example, the 1-PRM restricts the set of potential variables to those with binary "yes"/"no" type responses. While extensions of the 1-PRM like the PCRM used in Nord (2012), Tanaka et al. (2020), and Opsomer et al. (2002) allow for combinations of binary and ordered polytomous variables, the model cannot be estimated with continuous variables. Alternatively, with minor modification, our BGRM can be used to estimate latent traits with any mix of binary, ordered polytomous, and continuous variables. We highlight this feature in a demonstration exercise where we estimate our model using both the set of binary/polytomous adult FSM questions and an additional continuous variable: MFS.[51] This application produces a new latent construct that combines experiential information regarding adult food security with a measure of economic access/affordability in the form of monthly household food spending at supermarkets/grocery stores. We emphasize from the beginning, however, that the inclusion of MFS is solely for illustrative purposes and is not intended as a validated latent construct. We do not claim that economic access and food security should be combined without theoretical or empirical justification that is beyond the scope of this study. Rather, this demonstration shows how the BGRM can be altered to flexibly incorporate continuous variables—something not possible with traditional Rasch models used to estimate food security.[52]

We estimate our BGRM using the 10 binary/polytomous adult FSM questions and continuous MFS. Since MFS is a continuous variable, including it in our BGRM requires minor alterations to the model and Gibbs Sampler algorithm. We describe these changes in detail in

[51] In an earlier working paper, we used household-level Healthy Eating Index (HEI) scores as the continuous variable in our demonstration (Davis et al. (2024)). While continuous HEI may prove useful for future measurement applications, MFS is used in this study as it more directly aligns with the theoretical domains of food security.

[52] While capturing information from MFS may allow for a better understanding of consumers' economic access and affordability of food, we do not propose our MFS/FSM latent construct as a valid measure. Far more theoretical and exploratory work would be needed to validate any new measure, something outside the scope of this initial methodologically focused study.

Appendix D.. Similar to our primary results, the BGRM Gibbs Sampler algorithm with FSM and MFS is run for 50,000 iterations, removing the first 25,000 draws for burn-in.

Figure 4. Distribution of BGRM Posterior Mean Latent Construct Using FSM and Continuous Monthly Food Spending (MFS) by USDA FSM Scale Food Security Category Using NHANES 2017 - 2018

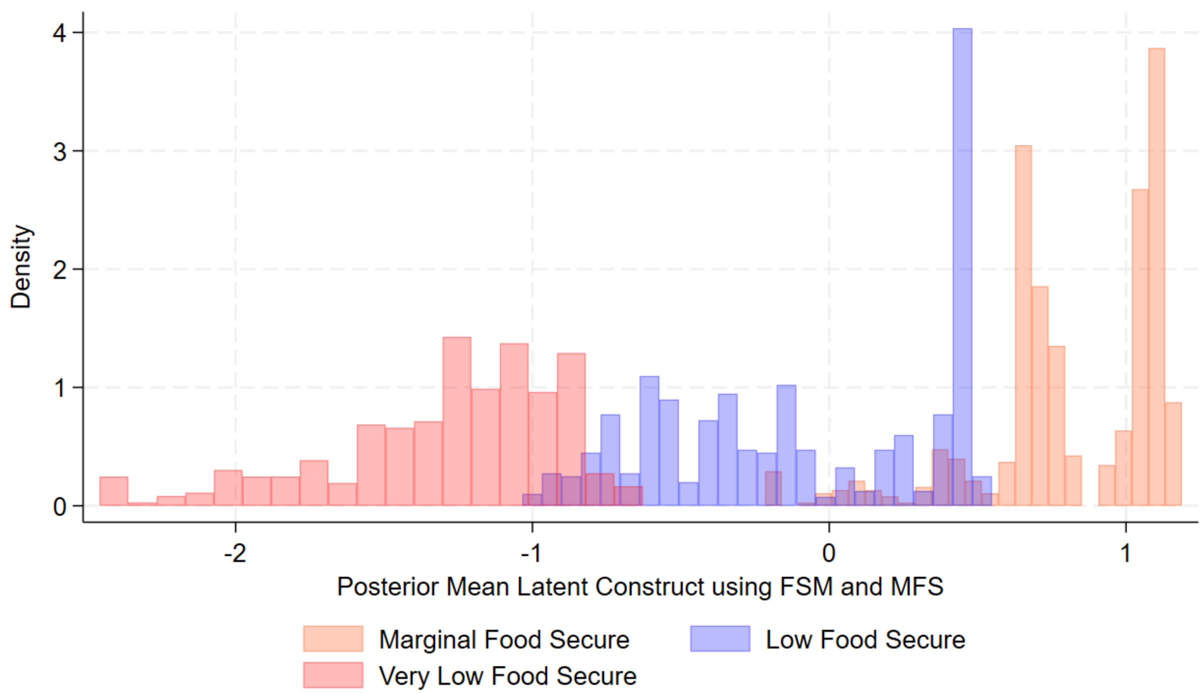

Note: Estimates are unweighted and reflect the sample only. The sample consists of households with responses to the 10 adult FSM questions from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 1,659 households. For each household in our sample a BGRM posterior mean latent construct using FSM and continuous monthly food spending is estimated by the model. Using the responses of the same households to the adult food security questionnaire, households are categorized into 1 of 3 food security categories based on the standard USDA food security category definitions. Marginal Food Secure households have 1-2 affirmative responses to the adult FSM questionnaire, followed by 3-5, and 6-9 affirmative responses for Low and Very Low Food Secure households, respectively. Histograms depicting the household distribution of BGRM posterior mean latent construct values for households belonging to the same FSM food security category are displayed using the 3 colors above.

Figure 4 shows the distribution of BGRM latent trait estimates for our novel FSM/MFS construct separated by USDA-defined FSM categories. Similar to Figure 3 made without MFS, Figure 4 shows overlap across USDA food security categories and significant within-category variation in our FSM/MFS point estimates. Additionally, the distribution of our new latent construct is qualitatively similar to the distribution of BGRM food security point estimates without MFS. For example, we find that the distribution of BGRM means for the new latent construct appears to remain bimodal for households in the USDA's marginal food secure group after adding MFS.[53]

To more directly compare the estimates produced by our BGRM using just the FSM to those of the new latent construct using FSM and MFS, we calculate the absolute difference in posterior mean latent trait values within-households. More specifically, we subtract the posterior mean estimate of latent food security using only the FSM questions from each household's posterior mean estimated with continuous MFS included. We then take the absolute value of the difference as some households will have negative differences while others have positive differences, but our primary concern is the magnitude of the difference rather than its direction. Higher absolute differences imply a greater level of disagreement between the two latent trait point estimates for the same household while smaller differences imply that the estimates are in greater agreement. We then plot the distribution of these absolute differences in Figure 5.

While adding continuous MFS to our BGRM does change each household's posterior mean latent trait estimate, the absolute difference between both is relatively small on average. The mean absolute difference between the two measures is 0.023 which is equivalent to roughly 2.5% of a standard deviation of the sample's distribution of posterior mean food security estimates (standard deviation is equal to 0.894 for the measure with MFS and 0.882 using just the FSM questions without food spending). That said, as shown in Figure 5, the absolute difference between estimates is not uniform, with far larger differences for some households than others. For example, the 95th percentile of absolute difference is 0.06 which is roughly 10% of one standard deviation of the posterior mean distribution. The highest absolute difference in our sample is 0.13, or roughly 15% of one standard deviation of the posterior mean. Therefore, while adding MFS leads to a relatively small impact on average, some households see non-trivial changes in their posterior mean food security level after adding our measure of economic access/affordability.

[53] We also compute a counterpart to Figure 2 using our estimates with MFS. The counterpart to Figure 2 is shown in Appendix D. as Figure F2. This figure does not change meaningfully after including MFS.

Figure 5. Distribution of Absolute Difference Between BGRM Latent Construct Posterior Means With and Without Continuous Food Spending Using NHANES 2017 - 2018

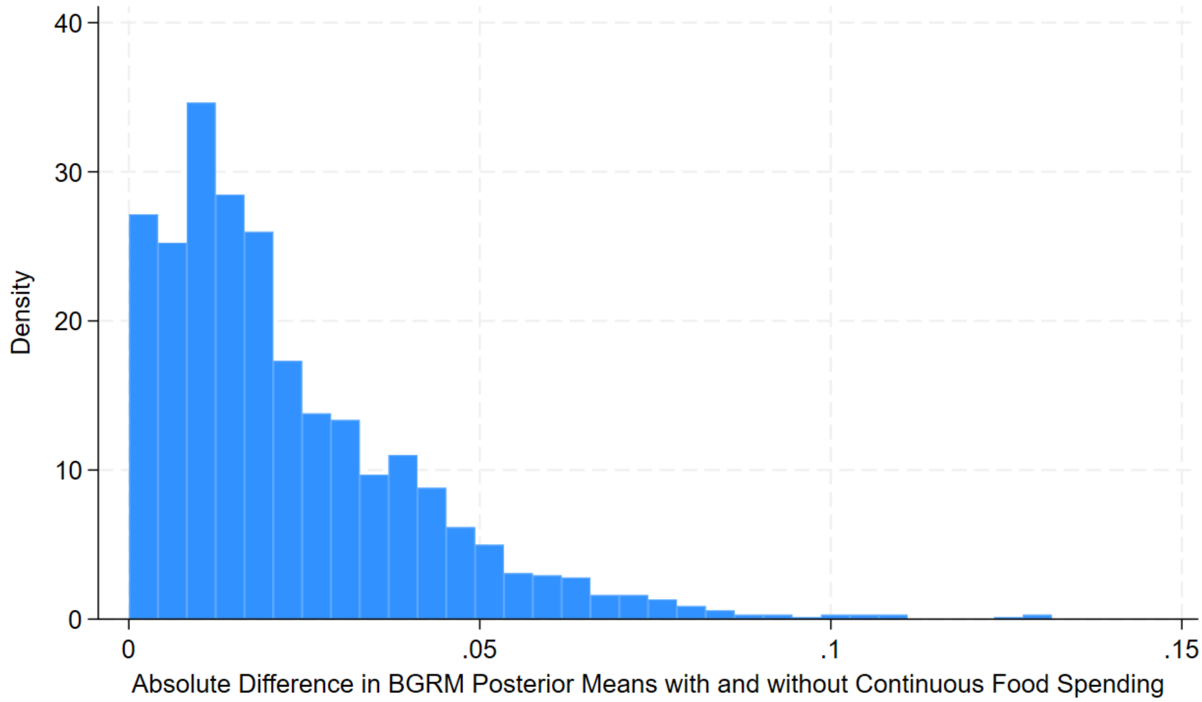

Note: Estimates are unweighted and reflect the sample only. The sample consists of households with responses to the 10 adult FSM questions from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 1,659 households. For each household in our sample a BGRM posterior mean value of food security is estimated by the model. Additionally, for each household, we also estimate a BGRM posterior mean latent construct using FSM and continuous monthly food spending. For each household, we calculate the absolute difference between the BGRM posterior mean value of food security and the BGRM posterior mean latent construct using FSM and continuous monthly food spending. The histogram depicts the household distribution of the absolute differences between the two latent constructs estimated by the model.

While the inclusion of MFS in our BGRM serves only to demonstrate model flexibility, there are theoretical grounds for future research to consider continuous economic access variables like MFS in food security research. Household food spending provides valuable insights into the economic access dimension of food security which the FSM may not fully capture. While the FSM assesses experiences and behaviors associated with food insecurity, it may not fully capture the financial constraints households face when securing food from retail sources. Incorporating MFS or similar continuous variables may help researchers better understand the economic pressures influencing food security status, complementing the experiential focus of the FSM and offering a more nuanced perspective on a multidimensional issue. This approach also aligns with broader definitions of food security from organizations like the USDA and World Bank (Rabbitt et al. (2024), World Bank (2024)), which emphasize both physical and economic access to adequate, safe, and nutritious food. Continuous measures such as MFS provide a proxy for each household's economic capacity to access food, addressing critiques in the literature that call for more comprehensive food security measures that include economic dimensions (Barrett and Lentz (2016), Manikas et al. (2023)). The BGRM's flexibility to integrate a mix of variable types offers a promising path forward for incorporating continuous variables like MFS into food security and related measures. However, we again emphasize that in this study, our MFS exercise is solely intended as a demonstration of the BGRM's capabilities rather than as a proposal for a validated new measure of food security or some other latent construct. Future work would be required to thoroughly explore and validate a combined approach, ensuring that it aligns with both theory and practical applications in food security measurement.

2.7 Conclusion

In this study, we introduce the Bayesian Graded Response Model (BGRM) as a novel method for measuring household food security and other latent constructs. The standard specification of our BGRM offers several advantages for food security measurement. First, it produces continuous point estimates of latent food security, providing detailed insights into households' relative food security levels. These estimates can complement existing discrete categories like those from the USDA's FSM scale. In an example exercise, we follow previous work in the food security literature (e.g., Gundersen (2008), Balistreri (2016)) and use our BGRM point estimates to calculate Foster, Greer, and Thorbecke (FGT) (Foster et al. (1984)) indices of food insecurity prevalence, depth, and severity for our NHANES sample, demonstrating the value of continuous estimates. Second, the BGRM accommodates binary and polytomous ordered response variables without requiring dichotomization, preserving information from the FSM's polytomous response option questions. Third, as a Bayesian model, the BGRM naturally produces household-level uncertainty estimates alongside point estimates, allowing researchers to incorporate precision into analyses, such as by weighting multivariate regression models based on estimation uncertainty.

Our application of the BGRM to the 2017–2018 NHANES data highlights its utility and flexibility. We find significant uncertainty in estimates for households in the USDA's marginally food-secure category, aligning with prior calls in the literature for measures that explicitly account for latent trait uncertainty (National Research Council (2006)). Moreover, we find significant within-category variation and our BGRM results reveal overlap in latent food security estimates across USDA categories, suggesting that some households may be misclassified under the current framework. This supports the potential for probabilistic category assignment as discussed in prior studies (e.g., National Research Council (2006); Nord (2012)).

While the PCRM used in previous studies also allows for a mix of binary and polytomous variables (Nord (2012); Tanaka et al. (2020); Opsomer et al. (2002)), one of the BGRM's most significant strengths is that it can be easily altered to allow for any mix of binary, polytomous, and continuous variables in model estimation. We demonstrate this feature by combining adult FSM response data with Monthly Food Spending (MFS) at grocery stores/supermarkets from the NHANES, estimating a novel latent construct. While the purpose of this exercise is to demonstrate model flexibility, it also highlights a potential method for combining continuous economic access measures with discreet experiential/behavioral data in a joint modeling framework. By capturing economic access to food, MFS complements the experiential focus of the FSM, offering a more holistic view of household food security. However, we reiterate that far more theoretical and empirical work would be needed to validate any new latent construct, something that is outside the scope of this initial study.

Future research can expand on the BGRM in several ways. Improving the computational efficiency of the model would enhance accessibility, particularly for researchers with limited computing resources. While all MATLAB code used to estimate the BGRM is freely available from the authors on request, developing software packages for more common programs with

Bayesian functionality like R, Python, or Stata could further broaden approachability. Beyond food security, the model's ability to incorporate continuous variables positions it as a promising tool for measuring related constructs, such as nutrition security, which will most likely require distilling diverse variables—both discreet and continuous—into a single latent trait (Seligman et al. (2023)). We see nutrition security measurement as a notable future application of our BGRM if a mix of discreet and continuous variable types are required as proposed by Seligman et al. (2023).

While this study is primarily methodological, it is grounded in the applied challenges of food security measurement. Offering a framework that incorporates uncertainty, preserves polytomous responses, and allows for flexible variable inclusion, the BGRM aligns with several of the longstanding recommendations from National Research Council (2006). We view this model as a step forward for both food security research and broader latent trait measurement. Our BGRM complements existing food security measurement approaches by offering flexibility and precision that may be especially valuable in specific research contexts, for example, when household-level measures of uncertainty or continuous measurement variables are required. While the BGRM is not intended to replace current methods used for food security measurement, it provides researchers with a powerful tool for addressing new and complex questions in food security and beyond.

Chapter III: Economic Conditions and Substance Use Among Teenagers in the United States

3.1 Introduction

The United States spends more than any other country on healthcare but ranks 27th in life expectancy, a difference which is largely due to substance misuse and associated physical and mental health problems (U.S. Department of Health and Human Services, 2016). Further, substance misuse and substance use disorders in the United States carry a high financial cost, costing more than $400 billion a year in crime, health and lost productivity (U.S. Department of Health and Human Services, 2016). Alcohol, marijuana and tobacco are the most commonly used substances by youth (Johnston et al. (2020)). Despite declines in use in the last years, among 8th to 12th graders, 15.6% report current (past 30 days) marijuana use, 18.2% report current alcohol use and 3.7% report current cigarette use (Johnston et al. (2020)). Since individuals generally start using and abusing substances during adolescence (U.S. Department of Health and Human Services, 2016), it is important to understand triggering factors, to guide policy to prevent use and abuse of substances among this population.

The economic literature has a well documented body of evidence on the relationship between economic conditions and substance use including alcohol, marijuana, tobacco and illicit drugs. While informative, most of these studies have focused on adults with relatively little attention devoted to adolescents. This study contributes to this literature by examining the relationship between economic conditions and the use of alcohol, marijuana and tobacco among adolescents in the United States. I use state-level aggregated data from the Youth Risk Behavior Surveillance System (YRBSS) (Centers for Disease Control & Prevention) for the sample period of 1991 to 2019. The YRBSS is a biennial national school survey of adolescents designed to monitor health risk behaviors (Pesko et al. (2016)). I use the adult unemployment rate as a measure of macroeconomic conditions. This is a standard approach in the literature (Arkes (2007); Carpenter et al. (2017); Tekin et al. (2013)).

The theoretical literature studying the relationship between economic conditions and substance use identifies three causal pathways: an income effect, economic stress effect and an opportunity cost effect (Ayllón and Ferreira-Batista (2018)). The income effect asserts that if drugs and alcohol are normal goods, hard economic times would lead to reduction in consumption due to lower incomes. Evidence of this effect can be found in Neumayer (2004), Granados (2005), Ruhm and Black (2002), Ruhm (2005), and Xu (2013). The economic stress effect links economic conditions and substance use through psychological factors. In particular, it asserts that individuals might resort to drugs and alcohol as a form of self-medication to deal with uncertainty about employment or lack of job opportunities in the labor market. In turn, the opportunity cost effect asserts that economic recessions lower the opportunity cost of drug and alcohol use due to lower wages or less job opportunities available in the labor market. Evidence consistent with this effect is found in Arkes (2011), Arkes (2012), Dee (2001), Bradford and Lastrapes (2014), and Currie and Tekin (2015). Thus, economic conditions and substance use would display a procyclical relationship if the income effect is dominant and a countercyclical relationship otherwise.

Given the existing evidence on the relationship between economic conditions and substance use among adults, why would one expect different results for adolescents? First, different or additional factors may play a role in the relationship between teenage substance use and economic conditions. The amount of unsupervised time available for teenagers, which plausibly could depend on the state of the economy, could be a significant factor. Income effects could vary across substances, drugs versus alcohol for example, given different costs for access. Further, a weaker economy could lead to increases in black market activity for drugs, and affect the ease of obtaining said drugs (Arkes (2007)). Thus, the effect of economic conditions on substance use among teenagers might differ from adults which motivates this analysis.

Besides providing insight on the effects of the business cycle on teenage substance use, this research could also be useful in identifying locations where teenage drug use might be on the rise and lead to a better understanding of the timing of the demand for drug treatment services (Arkes (2007)). There is evidence that drug use patterns occur in similar stages to an epidemic, with periods of incubation, expansion, plateau and decline. And further that these epidemics tend to be local (see Arkes (2007), Golub and Johnson (1993)). Despite existing evidence that early stage prevention programs can have a significantly negative influence on the duration and intensity of an epidemic, there is a dearth of systems designed to detect the onset of an epidemic (Caulkins (1998)). This issue is further complicated, by the fact that most of the known determinants of drug use - e.g: parent's education and family composition, prices of illicit drugs - display minimal change in the aggregate over time or are elusive to measurement in the population (see Ellickson et al. (1999), Gaviria and Raphael (2001), Cook and Moore (2001), Caulkins (1999), Pacula et al. (2001), Saffer and Chaloupka (1999), Bachman et al. (1998)). This study presents an easily measured predictor of teenage drug use that changes regularly, the adult unemployment rate.

Finally, in regards to the timing of the demand for drug treatment services, this research could provide guidance on funding decisions of drug treatment services during weaker economic periods. During state budget crises, often fueled by weak economic conditions, drug treatment services are prime candidates for receiving funding cuts. However, if evidence were to arise from this analysis that weak economic conditions lead to higher teenage substance use and abuse, and consequently higher need and demand for drug treatment services, funding cuts of these services, would not only be ill-advised but could potentially lead to severe long-term consequences for teenage drug users (Arkes (2007)).

Estimating state and year fixed effects models, I find evidence of a relationship between macroeconomic conditions and teenage substance use for females and 12 to 14 year olds. Females are more likely to drink alcohol and 12 to 14 year olds are more likely to smoke during periods of higher unemployment. The result for females is consistent with the findings of Pabilonia (2014). The remainder of this paper proceeds as follows: section 2 presents a theoretical framework and literature review. Section 3 describes the data. Section 4 describes the empirical approach. Section 5 presents the results and section 6 concludes.

3.2 Conceptual Framework & Literature Review

3.2.1 Conceptual Framework

Local economic conditions could affect teenage substance use through several mechanisms. These mechanisms could plausibly be divided into four main categories: (1) effects of parental job loss on teenage substance use, (2) effects of community-level job losses on teenage substance use, (3) effects of teenage job loss on teenage substance use, and (4) social networks and peer effects. Discussed below is each class of mechanism.

Effects of Parental Job Loss on Teenage Substance Use. Firstly, parental unemployment could reduce families' wealth. The length of the unemployment spell, as well as the level of family wealth, would determine the severity of financial deprivation and consequently the extent of financial stress placed on the teenager and the family. The effect would be more pronounced for teenagers whose parents have lost their jobs due to industry downturns, since these adults often must start over in new firms or industries (Jacobson et al. (1993); Stevens (1997)). Further, family income and material resources can affect children's wellbeing, as well as their achievement test scores (Morris and Gennetian (2003); Dahl and Lochner (2012)). Evidence from longitudinal studies using child fixed effects showed that parental unemployment decreases grade point averages (Rege et al. (2011)), increases grade repetition (Kalil and DeLeire (2002); Stevens and Schaller (2011)), and increases school-related behavioral problems (Hill et al. (2011)). Thus, in addition to lower financial security, poorer academic performance could amplify stress that motivates teenage substance use as a coping mechanism.

Secondly, parental unemployment could cause changes in families' physical and mental health, as well as parenting behaviors. Parental unemployment could affect teenage substance use by altering parents' physical and mental health and parent-teenager interactions. There are evidence that psychological and mental health declines for unemployed individuals when compared to employed individuals (Ananat et al. (2011); McKee-Ryan et al. (2005)). Additionally, studies employing longitudinal data observing families before and after an unemployment spell found evidence of debilitating family functions, as well as poorer child-to-parent interactions (Elder Jr et al. (1995); McLoyd (1998)). Conversely, parental job loss could lead to parents spending more time with children, which could lead to improved parent-to-teenager interaction thereby discouraging teenage substance use. However, evidence shows that unemployed parents do not necessarily spend more time with their children than employed parents (Edwards (2011); Kalil and Ziol-Guest (2013)).

Despite all the evidence discussed, positive life changes from parental unemployment could occur. Parents could choose to change careers due to a lengthy unemployment spell and find higher levels of life satisfaction from the new career choice. This change could lead to higher levels of family wealth and potentially better parent-to-teenager interaction. Additionally, parents could choose to migrate to a different region where better employment prospects are available. Migration could afford teenagers an opportunity to escape factors related to their environment that motivated the use of substances. Aggregately unclear is that parental unemployment may lead to higher teenage substance use.

Effects of Community-Level Job Losses on Teenage Substance Use. The effect of higher local unemployment rates on teenage substance use is not only manifested through a family member experiencing unemployment but could also be manifested through community-level effects. Some evidence suggests that firm layoffs as well as shutdowns impact not only job losers, but also residents of the impacted community, regardless of their employment status. Black, McKinnish, and Sanders (2003; 2005b; 2005a) examined booms and busts in the steel and coal industries in the 1970s and 1980s and found that industry downturns cause higher unemployment both within and outside of initially affected sectors. Large job losses at the community level also lower earnings for those who remain employed (Blanchflower and Oswald (1994); Ananat et al. (2011)) and individuals entering the labor market during this period experience a lifelong decline in earnings (Oreopoulos et al. (2012)). Furthermore, the effect of industry downturns is not restricted to reduced employment and earnings. Individuals residing in affected communities might also experience increased anxiety and stress, even if they are not among the job losers. Clark et al. (2010) and Luechinger et al. (2010) use longitudinal studies with individual fixed effects to show that an increase in the regional unemployment rate decreases employed individuals' reported life satisfaction. In addition, time series analyses by Dooley and Catalano, find a positive association between the local unemployment rate and psychological distress among employed individuals (Dooley and Catalano (1984); Dooley and Catalano (1984)).

Thus, even among teenagers whose family members might not experience unemployment, the effects described above could operate through their parents to create a host of issues that ultimately lead to teenage substance use. These issues could range from family financial distress due to lower parental earnings, to lower parental life satisfaction causing poorer teenage-to-parent interactions, to poorer academic performance for teenagers as a result of these stresses. Further, as parents attempt to respond to poorer regional economic conditions, this too could lead to migratory decisions that may result in teenagers being removed from existing social networks and school environment. Teenagers adjusting to their migration changes could have a positive or negative response in terms of their mental health and life satisfaction. The added stress of these migratory changes could also contribute to the incidence of or increase in teenage substance use.

Effects of Teenage Job Loss on Teenage Substance Use. This category of mechanisms only applies to teenagers who are legally allowed to work or are active members of the labor force. These teenagers could be directly affected by higher local unemployment rates, through actual job loss or poorer job prospects. Given that these teenagers are engaged in the labor force, the economic effects through which the unemployment rate affects teenage substance use for these individuals are synonymous to the ones that undergird the relationship for adults. Namely: (1) the income effect asserts that, if drugs and alcohol are normal goods, one would observe lower consumption during stressful economic times as a result of lower incomes; (2) the economic stress effect argues that drugs and alcohol might be used as a form of self-medication to deal with uncertainty about the availability of potential job opportunities in the labor market; and, (3) the opportunity cost effect argues that hard economic times might lead to higher substance use by lowering the opportunity cost of the former due to fewer job opportunities in the market. As a result, through these channels, economic conditions and teenage substance use would display a procyclical relationship if the income effect is dominant and a countercyclical relationship otherwise.

Social Networks and Peer Effects. Another category of mechanisms through which favorable or unfavorable local economic conditions could affect teenage substance use would be through teenagers' interactions with classmates, siblings, friends and work colleagues. Peer substance use has been shown to be one of the strongest predictors of teenage substance use (Barnes and Welte (1986); Huizinga et al. (1995); Scull et al. (2010)). Furthermore, social networks and peer effects constitute channels that could motivate substance use for teenagers from higher socioeconomic backgrounds, who might have been relatively insulated from unfavorable economic conditions. A plausible chain of causality could unfold where poor local or state level economic conditions could lead to increased teenage substance use for vulnerable teenagers. This group would include individuals from lower socio-economic backgrounds, or whose parents were directly affected by worsening economic conditions. Teenagers, whether directly affected or unaffected by these poor economic conditions, could frequent the same school, work and social environments with teenagers who engage in substance use. Through these interactions, the mechanisms of peer influence on teenage substance use would be exerted.

Peer influence could be exerted through several pathways. Peers could influence teenagers to engage in substance use through positive and negative peer pressure. Studies of middle school youth have documented a positive association between smoking and drinking and direct peer pressure (Kung and Farrell (2000); Simons-Morton et al. (2001)). Teenagers could engage in substance use through an observational learning process by associating with valued peers who use substances. Furthermore, beliefs about social norms regarding substance use could be transmitted through relationships with peers who use substances (Oetting and Beauvais (1986); Scull et al. (2010)).

Some studies find that sibling influence is a stronger predictor of teenage substance use than parental influence (Brook et al. (1999); Windle (2000)), and an equal or greater influence than peer influence (Brook et al. (1990); Needle et al. (1986); Vink et al. (2003)). Older siblings could influence teenage substance use by modeling (Bandura and Walters (1977)) or by creating accessibility (Forster et al. (2003); Needle et al. (1986)). Additionally, older siblings could also moderate the effect of peer substance use on younger siblings' substance use. Needle et al. (1986) documents that younger siblings are more likely to use substances when both older siblings and peers are using them. Brook et al. (1990) argue that this effect could be taken as evidence that older siblings, who do not use substances or use them infrequently, could deter the influence of peers' substance use on younger siblings' substance use.

In sum, the four main categories of mechanisms through which parental unemployment could affect teenage substance use are (1) effects of parental job loss on teenage substance use, (2) effects of community-level job losses on teenage substance use, (3) effects of teenage job loss on teenage substance use, and (4) social networks and peer effects. Under these categories, the mitigating and exacerbating factors as well as mechanisms of action are firm and industry spillover effects, family wealth, family physical and mental health, teenager physical and mental health, teenager academic performance, teenager labor force participation, the nature of drugs and alcohol as goods, social networks, peer interactions, and siblings' effects. Although extensive, this list of factors is by no means exhaustive. However, it provides a robust conceptual framework through which one can consider the potential mechanisms of action that underscore the relationship between economic conditions and teenage substance use.

At this stage, it would be pure speculation to argue which effects dominate the relationship between economic conditions and teenage substance use, and furthermore whether these effects differ significantly for adults. Since teenagers are typically more financially constrained than adults, perhaps the income effect would be larger for them. Teenagers' employment-based social networks could have less stigma surrounding drug and substance use, while adults might be more concerned with stigma in their employment networks, thus the social network effect could be larger for teenagers. Since teenagers who are involved in the labor force have both academic and job responsibilities, the opportunity cost effect could be more pronounced for them when compared to adults (Arkes (2007)).

Unraveling these mechanisms could be complicated, as many of the intermediate effects (income, hours of work, worse family interaction) could be endogenous to teenage drug use. However, accounting for these factors would control for part of the total effect of the economy. Since the main goal of this analysis is to estimate the total effect of the economy rather than a partial effect, these factors must be excluded from the model to capture the total effect of the economy on teenage substance use (Arkes (2007)).

3.2.2 Literature & Contribution

As previously mentioned, there is a voluminous body of evidence focused on the relationship between economic conditions and substance use. In this section, I discuss the findings from a selective list of studies, as this literature has been well summarized elsewhere (Ruhm and Jones (2012); Cawley and Ruhm (2012)). Research focusing on the effect of macroeconomic conditions on substance use has spanned a range of substances, from alcohol, marijuana and tobacco to harder drugs such as cocaine, amphetamines, heroin, among others.

Among the studies focused on alcohol use and macroeconomic conditions, are methodologically sophisticated studies which control for both time and geography fixed effects and identify the effects of the business cycle on substance use through within area (commonly within-state) changes in the unemployment rate (Carpenter et al. (2017)). These studies have mostly focused on time periods before the 2007 Great recession, and they find lower alcohol consumption among heavy drinkers and greater consumption among light drinkers during hard economic times (Dee (2001); Ruhm (1995); Ruhm and Black (2002); Pacula (2011)).

Various studies have also investigated the relationship between alcohol consumption and the business cycle using data covering the Great recession. Bor et al. (2013) use data from the Behavioral Risk Factor Surveillance System (BRFSS) and document increased abstinence from alcohol among U.S. adults but an increase in total consumption caused by a rise in the number of moderate and heavy drinkers. Tekin et al. (2013) also using the BRFSS, document higher binge drinking during the great recession, but argue this increase was economically insignificant. Cotti et al. (2015) use household scanner data, and find that among households who purchase alcohol, demand for packaged alcohol demonstrated a procyclical relationship with the business cycle during the Great Recession.

Research focused on the use of illicit drugs has been less widespread mostly due to lack of data (Carpenter et al. (2017)). Hollingsworth et al. (2017) use CDC Mortality data from the Multiple Cause of Death Files (MCOD) and AHRQ microdata from the State Emergency Department Databases (SEDD) covering 1999 - 2014. They find strong evidence that opioid-related deaths and emergency department (ED) visits increase during times of economic weakness. Pabilonia (2014) uses Youth and Risk Behavior Surveillance System data on 15-17 year olds covering 2003-2011. She finds weak evidence of a relationship between the business cycle and marijuana use among white respondents, but documents a significant counter-cyclical relationship among black males. Arkes (2007, 2011) uses data from the NLSY97 to study drug use among adolescents and young adults. He documents a strongly counter-cyclical pattern for marijuana and illicit drug use among teenagers. Among 20-24 year olds, he finds suggestive evidence of a countercyclical pattern for marijuana use, but weak evidence of a relationship between the business cycle and illicit drug use.

Neither Pabilonia (2014) nor Arkes (2007, 2011) are able to isolate the effect of the business cycle (i.e.: economic conditions), on the use of specific illicit drugs aside from marijuana. Carpenter et al. (2017) contribute to this literature by using the National Survey on Drug Use and Health (NSDUH) to study the effect of economic conditions on illicit drugs other than marijuana. Using data spanning 2002-2015, they find mixed evidence on the cyclicality of illicit drug use, but robust evidence on the counter-cyclicality of the intensity of prescription pain reliever use as well as clinically relevant substance use disorders involving opioids.

Among the studies discussed above, the one that most closely matches my analysis is Pabilonia (2014), as she uses the YRBSS and studies the same outcomes: marijuana, alcohol and cigarettes. However our analyses differ in important aspects. Pabilonia (2014)'s analysis focuses on the effects of the great recession on teenage substance use as well as other activities and her sample period spans 2003 - 2011. My analysis focuses on the broader effect of macroeconomic conditions on teenage substance use, and covers a larger sample period 1991 - 2019. In addition, I focus on a wider age group, 12 to 18 year olds, in contrast to Pabilonia (2014) who focuses on 15 to 17 year olds. Thus my analysis complements the one conducted by Pabilonia (2014) and provides updated evidence on the relationship between economic conditions and substance among teenagers in the United States.

3.3 Data

3.3.1 YRBSS

I use individual level data from the Youth Risk Behavior Surveillance System (YRBSS). The YRBSS is a biennial national school-based survey of adolescents conducted both at the national and state level. The national survey is conducted by the Centers for Disease Control and Prevention (CDC), while state surveys are conducted by state health and education agencies. Its objective is to monitor risky health behaviors among adolescents which include drug, alcohol and tobacco use. It relies on a two-stage, cluster sample design, where the data are weighted by school and student nonresponse to produce a representative sample of students in grades 9-12 in its jurisdiction. The typical sampling period is between February and May of each year (Pesko et al. (2016); Brener et al. (2013)). My period of analysis is 1991 to 2019. Given the biennial nature of the survey, this period of analysis yields 15 years of data.

YRBSS respondents are asked a series of questions regarding alcohol, drugs, and tobacco use. These questions range from use since last interview, age at first use, the number of days they have used marijuana, alcohol or tobacco in the past 30 days, the number of cigarettes they have had in the last 30 days, among many others. I use 4 measures of cigarette, alcohol and marijuana use: (1) heavy Cigarette use (every day over the past 30 days), (2) regular cigarette use (at least 20 days over the past 30 days), (3) current Marijuana use (at least 1 day over the past 30 days), (4) current alcohol use (at least 1 day over the past 30 days).

3.3.2 Macroeconomic Conditions

There are a few potential economic variables to use in this analysis. The economic variable needs to be measurable or available monthly or as an annual average at the state level.

Teenage employment statistics, such as the teenage unemployment rate or teenage labor force participation rate, are obvious candidates, but there are certain disadvantages associated with using these variables. First, these data are usually obtained from the Current Population Survey, for which there are very few teenagers for most states. This could potentially lead to high sampling error for all but the largest states (Arkes (2007)). Second, the inclusion of a teenage employment variable in the model, might cause model endogeneity if exogenous changes in teenage drug use have an effect on teenage labor supply (Arkes (2007)).

These problems would be significantly reduced by relying on employment statistics for individuals aged 16 years and over. Larger sample sizes would lead to less sampling error, and state employment statistics would be minimally affected by changes in teenage labor supply resulting from exogenous changes in teenage drug use (Arkes (2007)). Thus, I use the state adult unemployment rate[54] for individuals aged 16 and over which is considered the standard measure of the strength of the labor market. Specifically, I use the 12 month average of the monthly state unemployment rate extracted from the Bureau of Labor Statistics (BLS). This is a standard approach in the literature (Arkes (2007); Carpenter et al. (2017) ; Tekin et al. (2013)).

3.3.3 Descriptive Statistics

In Tables 19 and 20, I present descriptive statistics on substance use among teenagers in the United States, over the sample period 1991-2019. On average, more than two thirds of the sample (41.95%) report consuming alcohol on at least one day over the past 30 days. The corresponding figures for marijuana and cigarettes are approximately the same at 22%. The percentage of teenagers who report regular and heavy cigarette use is considerably smaller. 9.50% report smoking cigarettes on at least 20 days over the past 30 days, and 7.10% report smoking every day over the past 30 days.

There is moderate variation in the observations per outcome due to different rates of non-response accross the years. In regards to macroeconomic conditions, the average unemployment rate and labor force participation rate over the sample period are 5.95% and

[54] I explore models which include both the adult unemployment rate and adult labor force participation rate as explanatory variables. Results based on these models are presented in the appendix.

Table 19. Weighted Means for Outcomes and Explanatory Variables

	Years: 1991-2019	
	Mean	Observations (N)
Outcomes		
Heavy Cigarette Use (every day over the past 30 days) (%)	7.10	198,573
Regular Cigarette Use (at least 20 days over the past 30 days) (%)	9.50	198,573
Current Cigarette Use (at least 1 day over the past 30 days) (%)	22.33	198,573
Current Marijuana Use (at least 1 or 2 times over the past 30 days) (%)	21.78	204,119
Current Alcohol Use (at least 1 day over the past 30 days) (%)	41.95	193,125
Explanatory Variables		
Unemployment Rate (%)	5.95	208,244
Adult Labor Force Participation Rate (%)	65.20	208,244

Notes: The sample consists of U.S. teens between ages 12 - 18, for years 1991 - 2019, who responded to the YRBSS.

65.20% respectively. The average state real beer tax and average state real cigarette excise tax are \$0.29 and \$0.97, respectively. In regards to demographics, Whites are the largest racial group in the sample at 61.17%, followed by Blacks (13.73%), Other (11.47%), Hispanics (9.68%), and Asian or Pacific Islanders (3.96%). 15, 16 and 17 year olds are the largest age groups at 24.52%, 25.92% and 24.47% respectively. They are followed by 18 year olds or older teenagers and 14 years olds at 14.34% and 10.56% respectively. 12 and 13 year olds have an almost negligible representation in the sample, at 0.10% and 0.08% respectively. Finally, the sample is approximately evenly distributed accross the 4 grade groups, 9th, 10th, 11th and 12th grade with the sample percentages being 27.04%, 25.49%, 23.98%, and 23.48% respectively.

3.4 Methods

To investigate the effect of economic conditions on substance use among teenagers, I estimate a standard state and year fixed effects model, that relies on plausibly exogenous variation in the timing and extent of macroeconomic fluctuations accross states (Carpenter et al. (2017)).

Table 20. Weighted Means for Outcomes and Explanatory Variables

	Years: 1991-2019	
	Mean	Observations (N)
Taxes		
Average Real Beer Taxes (State) ($ per gallon)	0.29	208,244
Average Real Cigarette Excise Taxes (State) ($ per pack)	0.97	208,244
Demographics & School Characteristics		
Race		
White (%)	61.17	208,244
Black (%)	13.73	208,244
Hispanic (%)	9.68	208,244
Asian or Pacific Islander (%)	3.96	208,244
Other (%)	11.47	208,244
Age		
12 years old or younger (%)	0.10	208,244
13 years old (%)	0.08	208,244
14 years old (%)	10.56	208,244
15 years old (%)	24.52	208,244
16 years old (%)	25.92	208,244
17 years old (%)	24.47	208,244
18 years old or older (%)	14.34	208,244
Grade		
9th Grade (%)	27.04	208,244
10th Grade (%)	25.49	208,244
11th Grade (%)	23.98	208,244
12th Grade (%)	23.48	208,244

Notes: The sample consists of U.S. teens between ages 12 - 18, for years 1991 - 2019, who responded to the YRBSS.

The specification used takes the form:

$$Y_{ist} = \beta_0 + \beta_1 UR_{st} + \beta_2 X_{ist} + \beta_3 Z_{st} + \beta_4 \gamma_s T_t + \gamma_s + \delta_t + \varepsilon_{ist} \quad (22)$$

where Y_{ist} are the teenage substance use outcomes, for individual i residing in state s at year t. X_{ist} is a vector of individual specific characteristics such as gender, race, age, and grade; Z_{st}

includes state cigarette and beer excise taxes; cigarette indoor use laws; medical and recreational marijuana laws; indoor vaping restrictions; minimum legal purchase laws; e-cigarette taxes; and state ACA Medicaid expansion status; UR_{st} represents the 12 month average of the monthly adult unemployment rate for state s in year t.

The coefficient of interest in Eq.(22) is β_1, which conditional on the use of state and year fixed effects (γ_s and δ_t, respectively) is identified from within-state changes in substance use outcomes concurrent with within-state changes in the unemployment rate (Carpenter et al. (2017)). Permanent differences accross states in unobserved determinants of substance use outcomes are controlled for by the state dummies. Additionally, the year dummies control for secular year-specific effects common to all states (Carpenter et al. (2017)). Finally, ($\gamma_s * \delta_t$) are state specific linear time trends to control for omitted factors that vary over time within states. Standard errors are clustered at the state level.

3.5 Results

3.5.1 Primary Results

Table 21 shows 3 state level specifications for my dependent variables of interest: heavy cigarette use, regular cigarette use, current cigarette use, current alcohol use and current marijuana use. Panel 1 shows the specification which includes all controls, state and year fixed effects and state specific linear time trends. Panel 2 excludes all controls as well as state specific linear time trends and only uses state and year fixed effects. Panel 3 excludes all controls, state and year fixed effects and state specific linear time trends (i.e.:pooled ols). My preferred specification is in panel 1 as it better controls for possible confounding factors.

Turning to the findings from panel 1, for cigarette use, the coefficients on the unemployment rate are positive but statistically insignificant. A one-percentage-point increase in the unemployment rate is estimated to increase heavy cigarette use by 0.4 percentage points, which corresponds to 5.63% of the mean, and regular cigarette use by 0.4 percentage points, which corresponds to 4.21% of the mean. A one percentage point increase in the unemployment rate is estimated to increase current cigarette use by 0.4 percentage points, which corresponds to

1.79% of the mean. These results are consistent with Pabilonia (2014) which mostly documents positive but statistically insignificant effects of the unemployment rate on smoking.

For marijuana use, the coefficient on the unemployment rate is also positive but statistically insignificant. A one-percentage-point increase in the unemployment rate is estimated to increase current marijuana use by 0.2 percentage points, which corresponds to 0.92% of the mean. This result stands in contrast to Pabilonia (2014) and Arkes (2007) who find evidence of counter-cyclical marijuana use for teenagers.

Finally, for alcohol use, the coefficient on the unemployment rate is positive but statistically insignificant as well. A one-percentage-point increase in the unemployment rate is estimated to increase current alcohol use by 0.7 percentage points, which corresponds to 1.67% percent of the mean. This evidence is consistent with the results of Arkes (2007) and Pabilonia (2014). They stand in contrast to the findings of Ruhm (1995), Ruhm and Black (2002) and Dee (2001) which document pro-cyclical overall adult alcohol use. Thus, I find no evidence of a relationship between economic conditions and teenage cigarette, marijuana or alcohol use. However, it would be premature to conclude that these results indicate the absence of a relationship between macroeconomic conditions and teenage substance use. The null effects reported could arise if the pro-cyclical and counter-cyclical effects previously discussed nullify one another. Further, given that this analysis aggregates different racial, gender and age groups, it fails to account for differential effects accross these dimensions. In the next section, I analyze this relationship by gender, race and age in order to investigate the differential effects of economic conditions on teenage substance use accross these dimensions.

3.5.2 Comparisons Across Gender, Race and Age

In Table 22, I present estimates from my preferred specification by gender, age, and race. Panel 1 repeats the full sample results (from panel 1 of Table 21). The remaining panels separately present the findings for males, females, whites, blacks, 12 - 14 year olds and 15 - 18 year olds.

Females are more likely to currently drink alcohol during weaker economic times. For a one percentage point increase in the unemployment rate, females are 1 percentage point (2.38% of the mean) more likely to have drunk any alcohol in the last 30 days. This result is consistent with Pabilonia (2014). Teenagers aged 12 - 14 are more likely to currently smoke cigarettes during

Table 21. Coefficient Estimates on the Unemployment Rate in Teenage Substance-Use Equations (Preferred Specification vs Additional Models)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 1.Preferred Specification: All controls, state and year fixed effects & state specific linear time trends					
Unemployment Rate	0.004 (0.003) [-0.002,0.010]	0.004 (0.003) [-0.003,0.011]	0.004 (0.004) [-0.004,0.012]	0.002 (0.002) [-0.003,0.006]	0.007 (0.004) [-0.000,0.015]
Adjusted R^2	0.053	0.065	0.084	0.031	0.060
Panel 2. Only state and year fixed effects. Exclude controls. Exclude state specific linear time trends					
Unemployment Rate	0.001 (0.002) [-0.004,0.006]	0.000 (0.003) [-0.005,0.006]	-0.002 (0.004) [-0.009,0.006]	0.001 (0.003) [-0.006,0.007]	0.004 (0.004) [-0.004,0.012]
Adjusted R^2	0.030	0.038	0.058	0.010	0.026
Panel 3. No controls. No fixed effects. No state specific linear time trends					
Unemployment Rate	-0.003* (0.001) [-0.006,-0.001]	-0.004* (0.002) [-0.008,-0.000]	-0.004 (0.003) [-0.009,0.001]	-0.001 (0.002) [-0.005,0.003]	-0.000 (0.002) [-0.005,0.004]
Adjusted R^2	0.001	0.001	0.000	0.000	-0.000
Mean Rate (%)	7.10	9.50	22.33	21.78	41.95
Observations	198573	198573	198573	204119	193125

Notes: Linear Probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. The only controls used in the regression are listed in panel's title. All controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

weaker economic times. For a one percentage point increase in the unemployment rate, the probability of having smoked cigarettes increases by 0.9 percentage points (4.03% of the mean). Most of the remaining estimates accross all subsamples and outcomes are positive but statistically insignificant. Therefore, accross most outcomes and subsamples, I find no evidence of a relationship between macroeconomic conditions and teenage substance use. The overall lack of statistically significant estimates accross the outcomes and subsamples differs from results by Arkes (2007), which documents positive and statistically significant effects of macroeconomic conditions on white and Black teenage marijuana use and alcohol use. Additionally, Arkes (2007) documents statistically significant effects of macroeconomic conditions on male and female teenage marijuana use.

Table 22. Coefficient Estimates on the Unemployment Rate by Gender, Race, and Age

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
			Panel 1. All		
Unemployment Rate	0.004 (0.003) [-0.002,0.010]	0.004 (0.003) [-0.003,0.011]	0.004 (0.004) [-0.004,0.012]	0.002 (0.002) [-0.003,0.006]	0.007 (0.004) [-0.000,0.015]
Observations	198573	198573	198573	204119	193125
			Panel 2. Male		
Unemployment Rate	0.003 (0.003) [-0.004,0.009]	0.002 (0.004) [-0.005,0.009]	0.000 (0.004) [-0.007,0.007]	0.001 (0.002) [-0.004,0.006]	0.005 (0.004) [-0.004,0.013]
Observations	97004	97004	97004	99991	94258
			Panel 3. Female		
Unemployment Rate	0.005 (0.003) [-0.001,0.012]	0.006 (0.004) [-0.002,0.013]	0.007 (0.005) [-0.003,0.017]	0.002 (0.004) [-0.005,0.009]	0.010* (0.005) [0.000,0.019]
Observations	101569	101569	101569	104128	98867
			Panel 4. 12 - 14 years old		
Unemployment Rate	0.006 (0.004) [-0.002,0.014]	0.004 (0.003) [-0.003,0.011]	0.009* (0.003) [0.002,0.015]	-0.001 (0.004) [-0.009,0.007]	0.006 (0.006) [-0.005,0.018]
Observations	19940	19940	19940	20435	19436

Notes: Table continues on next page

3.5.3 Robustness Tests

In this section, I conduct additional tests on the robustness of my main results to changes in samples or specifications. The results of these analyses are presented in Table 23. In panel 1, I repeat the preferred specification for comparison. In panel 2, I estimate the original model but exclude state specific linear time trends. These trends were included to control for omitted factors

Table 22 Cont'd: Coefficient Estimates on the Unemployment Rate by Gender, Race, and Age

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
		Panel 5. 15 - 18 years old			
Unemployment Rate	0.003 (0.003) [-0.002,0.009]	0.003 (0.003) [-0.003,0.010]	0.002 (0.004) [-0.007,0.011]	0.002 (0.002) [-0.003,0.007]	0.007 (0.004) [-0.001,0.015]
Observations	178633	178633	178633	183684	173689
		Panel 6. White			
Unemployment Rate	0.004 (0.004) [-0.004,0.012]	0.004 (0.004) [-0.005,0.013]	0.004 (0.004) [-0.005,0.012]	-0.002 (0.003) [-0.008,0.004]	0.005 (0.005) [-0.004,0.015]
Observations	86492	86492	86492	88381	84968
		Panel 7. Black			
Unemployment Rate	-0.001 (0.002) [-0.005,0.002]	-0.003 (0.002) [-0.007,0.002]	0.002 (0.005) [-0.008,0.013]	0.010 (0.005) [-0.000,0.020]	0.010 (0.007) [-0.004,0.023]
Observations	42248	42248	42248	43698	40812

Notes: Linear probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. Controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

that vary over time within states such as unobserved preferences and tastes among teenagers (Hansen et al. (2017)). In panel 3, I investigate the existence of non-parallel pre-trends as well as a lagged effect of the unemployment rate on teenage substance use by including a 2-year lead and lag of the unemployment rate to the preferred specification. In panel 4, to account for the possibility that economic measures from small states could be affected by sampling error, I exclude observations from the 7 least populous states - Delaware, District of Columbia, Hawaii, Montana, Rhode Island, South Dakota, and Vermont. Finally, in panel 5, to account for the high degree of variation in economic conditions within large states, I exclude observations from the four most populous states - California, Texas, New York and Florida.

The models that exclude state specific linear time trends produce generally smaller estimates than the original model, but the coefficients on the unemployment rate remain statistically insignificant. The coefficient on current cigarette use experiences a sign reversal. Higher unemployment is now estimated to decrease current cigarette use by 0.2 percentage points (0.89% of the mean). In the models that include the 2 year lead and lag of the unemployment rate, the coefficients on the 2 year and lag for all outcomes are statistically insignificant. Thus I find no evidence that past unemployment influences current cigarette, alcohol or marijuana use. I also find no evidence that my results are due to the existence of non-parallel pre-trends observed 2 years prior. In these models, the coefficients on the unemployment rate appear similar in magnitude when compared to the same coefficients in the preferred specification (panel 1). However, they are all generally larger and the coefficient on current alcohol use now becomes statistically significant. A one percentage point increase in the unemployment rate increases current alcohol use by 1 percentage point (2.38% of the mean).

Similar to the results from the main specification, the estimates from the models that exclude people from the 7 least-populous states display no statistically significant evidence of a relationship between economic conditions and substance use among teenagers. The coefficient estimates on the unemployment rate are extremely similar in magnitude and retain the same signs as in the preferred specification. These results are potentially due to the subsample and full sample

Table 23. Coefficient Estimates on the Unemployment Rate in Teenage Substance-Use Equations (All Panels Display a Modified Version of the Preferred Specification or Sample Used. Each Panel Title Lists the Change)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 1. Preferred Specification: All controls, state and year fixed effects & state specific linear time trends					
Unemployment Rate	0.004 (0.003) [-0.002,0.010]	0.004 (0.003) [-0.003,0.011]	0.004 (0.004) [-0.004,0.012]	0.002 (0.002) [-0.003,0.006]	0.007 (0.004) [-0.000,0.015]
Adjusted R^2	0.053	0.065	0.084	0.031	0.060
Panel 2. Exclude State Specific Linear Time Trends					
Unemployment Rate	0.002 (0.002) [-0.003,0.006]	0.001 (0.002) [-0.003,0.006]	-0.002 (0.003) [-0.009,0.004]	0.001 (0.002) [-0.004,0.006]	0.002 (0.003) [-0.005,0.009]
Adjusted R^2	0.051	0.063	0.082	0.030	0.059
Panel 3. Include 2-year lag and lead. Include State Specific Linear Time Trends					
Unemployment Rate	0.005 (0.003) [-0.001,0.012]	0.005 (0.004) [-0.002,0.013]	0.005 (0.004) [-0.003,0.014]	0.004 (0.003) [-0.002,0.010]	0.010* (0.004) [0.002,0.019]
2-year lagged unemployment rate	0.001 (0.002) [-0.003,0.005]	0.001 (0.003) [-0.005,0.006]	-0.002 (0.004) [-0.011,0.007]	-0.004 (0.003) [-0.011,0.003]	0.001 (0.004) [-0.007,0.009]
2-year lead unemployment rate	-0.004 (0.002) [-0.009,0.001]	-0.004 (0.003) [-0.009,0.002]	-0.002 (0.003) [-0.008,0.004]	-0.002 (0.003) [-0.008,0.004]	-0.007 (0.003) [-0.014,0.000]

Notes: Table continues in next page

being similarly sized. The estimates from the models that exclude observations from the 4 most populous states in the sample also display no statistically significant evidence of a relationship between economic conditions and teenage substance use. The coefficients on the unemployment rate are also very similar in magnitude to the ones in the preferred specification, although for these models the sample sizes for different outcomes differ by a meaningful amount of observations.

Table 23 Cont'd: Coefficient Estimates on the Unemployment Rate in Teenage Substance-Use Equations (All Panels Display a Modified Version of the Preferred Specification or Sample Used. Each Panel Title Lists the Change)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 4. Exclude 7 least populous states					
Unemployment Rate	0.004 (0.003) [-0.002,0.010]	0.004 (0.003) [-0.003,0.011]	0.004 (0.004) [-0.004,0.012]	0.002 (0.002) [-0.003,0.006]	0.007 (0.004) [-0.000,0.015]
Observations	194422	194422	194422	199855	189237
Panel 5. Exclude 4 most populous states					
Unemployment Rate	0.004 (0.003) [-0.003,0.011]	0.004 (0.004) [-0.004,0.012]	0.004 (0.005) [-0.006,0.014]	0.001 (0.002) [-0.004,0.005]	0.007 (0.005) [-0.002,0.017]
Observations	128310	128310	128310	131296	124263
Mean Use Rate (%)	7.10	9.50	22.33	21.78	41.95

Notes: Linear Probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. All controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.6 Discussion

This paper analyzes the relationship between economic conditions and teenage substance use, using state-level data from the YRBSS. Results from state-year fixed effects models provide mixed evidence for the relationship between the adult unemployment rate and substance use among teenagers. No evidence is found of a relationship between macroeconomic conditions and substance use when analyzing the full sample. This result is mostly consistent with the findings of Pabilonia (2014). It stands in contrast to the findings of Arkes (2007), Ruhm (1995), Ruhm and Black (2002), and Dee (2001). There is evidence of a relationship between macroeconomic conditions and teenage substance use for females and 12 to 14 year olds. Females are more likely to drink alcohol and 12 to 14 year olds are more likely to smoke during periods of higher unemployment. The results for females are consistent with the findings of Pabilonia (2014).

The discrepancy between my findings and the research discussed above could be due to several reasons. Arkes (2007) and Pabilonia (2014) analyze 16 to 18 year olds and 15 to 17 year olds respectively. My sample includes a wider age group namely 12 to 18 year olds. The differences in our full sample results could potentially be due to differential effects among 12 to 14 year olds which are captured in my analysis. On the other hand, the discrepancy between my results and those of Ruhm (1995), Ruhm and Black (2002) and Dee (2001) could be due to the different age groups analyzed, as the authors focus only on adults. In particular, the mechanisms through which the economy affects substance use, could have differing effects on teenagers compared to adults, leading to different patterns of use in response to the same economic conditions.

There are some caveats and limitations of this study that are worth mentioning. First, there might be concerns about the reliability of self-reported data on illegal activities such as marijuana use. Any correlation between systemic misreporting and economic conditions could affect the estimates (Arkes (2007)). A second issue could be that the state might be too broad a geographic region to measure the economic statistic reflecting economic conditions. However economic statistics measured at a more reduced level such as Metropolitan statistical areas, might suffer from more measurement error given smaller sample sizes (Arkes (2007)).

Despite these caveats, the results do offer some insight, as I document evidence of counter- cyclical alcohol and cigarette use among females and 12 to 14 year olds. As previously discussed, there are several potential mechanisms that could explain these findings. Hard economic conditions could decrease the opportunity cost of alcohol and cigarette use, and cause the onset of mental health issues such as depression due to joblessness or family discord caused by financial difficulties. Furthermore, these conditions could lead to worse academic performance, higher levels of stress and use of substances as a coping strategy. However, weaker economic times could also contribute to lower alcohol and cigarette use by lowering teenagers' incomes or increasing their amount of supervised time. For females as well as 12 to 14 year olds, my results indicate that the counter- cyclical mechanisms dominate.

Finally, my results imply that economic conditions must be considered, when deciding the extent and future of public funding for substance use programs aimed at teenagers. Worse economic conditions often lead to state budget cuts (Carpenter et al. (2017)), placing these programs in a vulnerable position, and program funding cuts could potentially exacerbate the effects of hard economic times on teenage substance use.

Appendix A. Additional Tables and Figures for Chapter 1

Table A1. Constructed Store Categories, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program

Supermarkets	Grocery	Convenience	Combination	Specialty	Other
Walmart	Sai Fresh Foods	Quiktrip	Dollar General	Peacock Meats	Twenterprise
Kroger	Soonoco Friends	Quick Mart	Race Trac	Sweet Thangs Bakery	Country Fresh Foods
Publix	King Food Mart	Citgo Quick Service	Dollar Tree	Crab King II	Dyans Food
Aldi	Family Food Market	Chevron Food Mart	Walgreens	Butcher boyz	USA Best Choice
Target	Twin City Grocery	Stop N Save	Rite Aid	Rocky Creek Seafood	Temple Seafood
Piggly Wiggly	Harvest Time Market	Exxon Food Mart	K.D. Asian Market	Tucker Farmers Market	Plantation Meats
Food Depot	Barkers Market	Quick Stop	Fred's Store	Anytime Bread	Massey Foods
Food Lion	Snipes Fine Foods	Petro Gas Station	CVS	Cajun Meat Company	The Steak Man
Sam's Club	Gaskins Grocery	Shell Food Mart	Ideal Tropical Foods	Red Barn Produce	Simpson Meats
Ingles	Dublin Foods 30	Circle K Store	Family Dollar Store	Fatboys Produce	Anthonys Store

Notes: Non-exhaustive list detailing some of the individual stores that belong to each constructed store category. These categories are derived from the USDA Store Tracking and Redemption System (STARS). Supermarkets refers to stores that sell a wide variety of grocery and other store merchandise, and have multiple checkout lanes and registers. This category includes STARS supermarkets, superstores and military commissaries. Grocery stores refers to stores that carry a small, moderate, or wide selection of all four staple food categories, with their primary stock being food items. This category includes STARS small, medium and large grocery stores. Convenience refers to self-service stores that offer a limited line of convenience items and primarily sell a variety of canned goods, dairy products, pre-packaged meats and other grocery items in limited amounts. This category is identical to the STARS convenience category. Combination refers to stores where the primary business is the sale of general merchandise but also sell a variety of food products. This category includes dollar stores, drug stores, and general stores. This category is identical to the STARS combination grocery/other stores category. Specialty foods refers to stores that operate as a cooperative or specialize in the sale of specific products such as bread/cereal products, fruits and/or vegetable products, meat products, and seafood products. Finally, Other refers to all remaining stores in the data that do not fit neatly into any of the above categories.

Table A2. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and by Benefit Size

	Main Sample	***BA<$100***	***BA∈[$100,$250]***	**BA>$250**	***P-value***	***1 vs 2***	***2 vs 3***	***1 vs 3***
Amount Spent								
$ Within 7 days	167.47 (150.23)	21.97 (23.03)	109.96 (66.71)	295.14 (187.85)	0.0000	-87.99***	-185.18***	-273.17***
% Within 7 days	59.72 (33.54)	58.77 (42.25)	59.92 (35.02)	56.04 (29.45)	0.0000	-1.15***	3.88***	2.74***
$ Within Month	261.69 (188.05)	33.79 (23.71)	167.63 (43.73)	476.85 (183.95)	0.0000	-133.84***	-309.22***	-443.05***
% Within Month	91.55 (18.44)	93.38 (15.99)	91.47 (19.21)	90.69 (18.51)	0.0000	1.91***	0.78***	2.69***
Transactions & Benefit Amount								
# of Transactions (7 days)	4.31 (3.68)	1.39 (1.51)	3.65 (3.14)	5.90 (4.20)	0.0000	-2.26***	-2.25***	-4.51***
# of Transactions (Month)	8.76 (6.87)	2.47 (2.10)	7.07 (5.20)	12.90 (7.70)	0.0000	-4.61***	-5.83***	-10.43***
Bulk Shopping (%)	49.26 (49.99)	85.42 (35.29)	55.82 (49.66)	29.20 (45.47)	0.0000	29.60***	26.62***	56.22***
Benefit Amount ($)	287.80 (198.00)	36.35 (24.36)	183.30 (27.80)	527.68 (176.14)	0.0000	-146.95***	-344.39***	-491.33***
SNAP Tenure								
Tenure (Ben. Months)	35.80 (15.80)	32.66 (17.42)	27.43 (17.75)	36.68 (16.54)	0.0000	5.23***	-9.25***	-4.02***
Households	1,555,674	104,109	332,444	312,193				
Benefit Months	34,770,457	1,730,014	4,279,901	6,419,025				

Note: Main sample and breakdown by benefit size. BA represents benefit amount. Three groups are displayed: households whose maximum benefit amount received in any benefit month is less than $100, households who receive a minimum benefit of $100 and a maximum benefit of $250 inclusive, and households whose minimum benefit amount received in any benefit month is greater than $250. P-value for F test of joint difference in means across all three groups is displayed. Pairwise t-tests for difference of means are also displayed. 1 vs 2 refers to the first group means compared to the second group means (BA<$100 vs BA∈[$100,$250]). The remaining comparisons are interpreted analogously. The benefit categories are mutually exclusive, but are not collectively exhaustive. For example, households who receive a minimum of $100 and a maximum benefit greater than $250 are not represented in these categories.

Table A2. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and by Benefit Size (continued)

	Main Sample	*BA<$100*	*BA∈[$100,$250]*	BA>$250	*P-value*	*1 vs 2*	*2 vs 3*	*1 vs 3*
Household Location								
Rural	16.87 (37.45)	24.73 (43.14)	14.83 (35.54)	14.88 (35.59)	0.0000	9.90***	-0.05*	9.85***
Urban	75.88 (42.78)	67.67 (46.77)	77.75 (41.59)	78.46 (41.11)	0.0000	-10.08***	-0.70***	-10.79***
Mixed	7.26 (25.94)	7.60 (26.50)	7.42 (26.21)	6.66 (24.94)	0.0000	0.18***	0.76***	0.94***
Store Choice								
Supermarkets ($)	215.83 (171.77)	27.13 (23.54)	135.54 (58.26)	396.01 (184.56)	0.0000	-108.41***	-260.47***	-368.88***
Grocery ($)	12.00 (41.42)	1.73 (7.73)	7.40 (25.20)	23.41 (64.59)	0.0000	-5.67***	-16.01***	-21.68***
Convenience ($)	13.75 (36.09)	1.42 (5.81)	10.89 (26.40)	22.86 (52.82)	0.0000	-9.47***	-11.97***	-21.44***
Combination ($)	13.40 (32.80)	2.74 (7.94)	9.38 (22.24)	22.04 (47.59)	0.0000	-6.64***	-12.66***	-19.30***
Specialty Foods ($)	3.27 (17.65)	0.41 (3.23)	2.20 (11.57)	5.74 (26.36)	0.0000	-1.79***	-3.54***	-5.34***
Other ($)	3.44 (24.57)	0.37 (3.71)	2.22 (15.31)	6.80 (38.87)	0.0000	-1.85***	-4.58***	-6.43***
Households	1,555,674	104,109	332,444	312,193				
Benefit Months	34,770,457	1,730,014	4,279,901	6,419,025				

Note: Main sample and breakdown by benefit size. BA represents benefit amount. Three groups are displayed: households whose maximum benefit amount received in any benefit month is less than $100, households who receive a minimum benefit of $100 and a maximum of $250 inclusive, and households whose minimum benefit amount received in any benefit month is greater than $250. P-value for F test of joint difference in means across all three groups is displayed. Pairwise t-tests for difference of means are also displayed. 1 vs 2 refers to the first group means compared to the second group means (BA<$100 vs BA∈[$100,$250]). The remaining comparisons are interpreted analogously. The benefit categories are mutually exclusive, but are not collectively exhaustive. For example, households who receive a minimum of $100 and a maximum benefit greater than $250 are not represented in these categories.

Table A3. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and by Tenure

	Main Sample	*T≤12*	*T∈(12,24)*	T≥24	*P-value*	*1 vs 2*	*2 vs 3*	*1 vs 3*
Amount Spent								
$ Within 7 days	167.47 (150.23)	119.65 (115.03)	140.13 (123.11)	180.02 (157.43)	0.0000	-20.48***	-39.89***	-60.37***
% Within 7 days	59.72 (33.54)	51.53 (36.22)	56.49 (34.39)	61.58 (32.71)	0.0000	-4.96***	-5.09***	-10.04***
$ Within Month	261.69 (188.05)	220.00 (156.61)	235.96 (164.16)	273.01 (195.26)	0.0000	-15.96***	-37.05***	-53.01***
% Within Month	91.55 (18.44)	90.50 (19.61)	91.39 (18.60)	91.74 (18.23)	0.0000	-0.90***	-0.34***	-1.24***
Transactions & Benefit Amount								
# of Transactions (7 days)	4.31 (3.68)	3.33 (3.24)	3.85 (3.39)	4.55 (3.77)	0.0000	-0.52***	-0.70***	-1.22***
# of Transactions (Month)	8.76 (6.87)	7.72 (6.20)	8.30 (6.53)	9.01 (7.01)	0.0000	-0.58***	-0.71***	-1.29***
Bulk Shopping (%)	49.26 (49.99)	51.70 (49.97)	50.28 (50.00)	48.69 (49.98)	0.0000	1.42***	1.59***	3.01***
Benefit Amount ($)	287.80 (198.00)	246.06 (167.97)	260.19 (173.42)	299.49 (205.19)	0.0000	-14.13***	-39.30***	-53.43***
SNAP Tenure								
Tenure (Ben. Months)	35.80 (15.80)	7.93 (3.02)	18.28 (3.08)	43.42 (9.90)	0.0000	-10.35***	-25.14***	-35.49***
Households	1,555,674	636,521	287,162	631,991				
Benefit Months	34,770,457	3,861,033	5,094,561	25,814,863				

Note: Main sample and breakdown by length of participation in SNAP. Length of participation in SNAP is measured by tenure (T), the number of benefit months a households spends in the program. Three groups are displayed: households whose tenure in SNAP is less or equal to 12 benefit months (a year or less), between 12 and 24 benefit months exclusive (more than a 1 year and less than 2 years), and greater or equal to 24 benefit months (2 years or more). P-value for F test of joint difference in means across all three groups is displayed. Pairwise t-tests for difference of means are also displayed. 1 vs 2 refers to the first group means compared to the second group means (T<12 vs T∈(12,24)). The remaining comparisons are interpreted analogously. The tenure categories are mutually exclusive and collectively exhaustive.

Table A3. Summary Statistics, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program (SNAP) Households, Main Sample and by Tenure (continued)

	Main Sample	*T≤12*	*T∈(12,24)*	T≥24	*P-value*	*1 vs 2*	*2 vs 3*	*1 vs 3*
Household Location								
Rural	16.87	13.24	14.57	17.86	0.0000	-1.33***	-3.29***	-4.62***
	(37.45)	(33.89)	(35.28)	(38.30)				
Urban	75.88	79.44	78.16	74.90	0.0000	1.29***	3.26***	4.55***
	(42.78)	(40.41)	(41.32)	(43.36)				
Mixed	7.26	7.32	7.27	7.24	0.0000	0.05**	0.03*	0.08***
	(25.94)	(26.04)	(25.97)	(25.92)				
Store Choice								
Supermarkets ($)	215.83	186.76	198.60	223.58	0.0000	-11.84***	-24.97***	-36.82***
	(171.77)	(148.83)	(154.71)	(177.37)				
Grocery ($)	12.00	7.32	8.64	13.36	0.0000	-1.33***	-4.72***	-6.05***
	(41.42)	(30.64)	(33.39)	(44.08)				
Convenience ($)	13.75	10.76	11.92	14.56	0.0000	-1.16***	-2.64***	-3.80***
	(36.09)	(29.68)	(31.58)	(37.72)				
Combination ($)	13.40	10.00	11.28	14.33	0.0000	-1.27***	-3.06***	-4.33***
	(32.80)	(27.31)	(29.03)	(34.17)				
Specialty Foods ($)	3.27	2.06	2.38	3.63	0.0000	-0.32***	-1.25***	-1.57***
	(17.65)	(12.95)	(13.92)	(18.86)				
Other ($)	3.44	3.11	3.14	3.55	0.0000	-0.03	-0.42***	-0.45***
	(24.57)	(23.21)	(23.14)	(25.03)				
Households	1,555,674	636,521	287,162	631,991				
Benefit Months	34,770,457	3,861,033	5,094,561	25,814,863				

Note: Main sample and breakdown by length of participation in SNAP. Length of participation in SNAP is measured by tenure (T), the number of benefit months a households spends in the program. Three groups are displayed: households whose tenure in SNAP is less or equal to 12 benefit months (a year or less), between 12 and 24 benefit months exclusive (more than a 1 year and less than 2 years), and greater or equal to 24 benefit months (2 years or more). P-value for F test of joint difference in means across all three groups is displayed. Pairwise t-tests for difference of means are also displayed. 1 vs 2 refers to the first group means compared to the second group means (T<12 vs T∈(12,24)). The remaining comparisons are interpreted analogously. The tenure categories are mutually exclusive and collectively exhaustive.

Table A4. Three-Class Finite Mixture Estimates, 2011 - 2015 Georgia Supplemental Nutrition Assistance Program Households, Restricted Subsample 1

	Restricted Subsample 1	3-Part Mixture		
		Group 1	Group 2	Group 3
Log Credit	0.7086***	0.2090***	0.8976***	0.9960***
	(0.0053)	(0.0151)	(0.0018)	(0.0000)
Transactions	0.0173***	0.0743***	-0.0071***	-0.0000***
	(0.0003)	(0.0009)	(0.0001)	(0.0000)
Bulk Shopping	0.4343***	-0.5217***	0.2865***	-0.0001***
	(0.0027)	(0.0099)	(0.0010)	(0.0000)
Log Grocery	0.0235***	0.0201***	0.0119***	-0.0000
	(0.0009)	(0.0037)	(0.0003)	(0.0000)
Log Convenience	-0.0046***	0.0027	0.0092***	-0.0000
	(0.0010)	(0.0036)	(0.0003)	(0.0000)
Log Combination	0.0009	0.0287***	0.0037***	-0.0000
	(0.0008)	(0.0032)	(0.0003)	(0.0000)
Log Specialty	0.0185***	-0.0134**	0.0124***	0.0000
	(0.0011)	(0.0053)	(0.0004)	(0.0000)
Log Other Stores	0.0228***	0.0241***	0.0115***	0.0000
	(0.0015)	(0.0060)	(0.0005)	(0.0000)
Rural	0.0118	-0.0159	0.0010	0.0000
	(0.0087)	(0.0287)	(0.0029)	(0.0000)
Mixed	-0.0155***	-0.0252	-0.0066***	-0.0000
	(0.0059)	(0.0191)	(0.0018)	(0.0000)
Period	-0.0017***	0.0059***	-0.0016***	0.0000***
	(0.0001)	(0.0004)	(0.0000)	(0.0000)
R-squared	0.1780			
Mean (Y)	$123	$13	$191	$308
Probability		17%	74%	9%
Households	91,848			
Benefit Months	1,986,510			

Regression and mixture employ correlated random effects models. The dependent variable in the regression and mixture is the log of dollars spent within the first 7 days of the benefit month. The restricted subsample 1 column displays the results from a correlated random effects regression using a 10% random sample of households who receive a minimum benefit of $100 dollars or more and spend at least 75% of their benefit in any given benefit month. Group 1, 2, and 3 refer to the 3 part estimated classes from the mixture using the same 10% random sample. (*10%,**5%,***1%).

Appendix B. Simulated Data Exercise

Prior to estimating the BGRM with real data, we test model performance in a simulated data exercise. The goal of this exercise is to ensure that our model can accurately retrieve known data generating parameters used to simulate a set of household responses. Data are simulated for a set of $I = 5,000$ households answering a set of $J = 10$ questions.[55] The potential responses for each question is set to mimic those of the household food security component of the FSM, implying that $C = [C_1, C_2, ..., C_{10}]$ is defined as $C = [3, 3, 3, 2, 3, 2, 2, 2, 2, 3]$. Our BGRM is well suited to this mix of binary and polytomous ordered response data as opposed to strictly binary or polytomous alternative IRT methods.

The set of data generating structural parameters is given in Table F1. The food security level of each household i is constructed such that $\delta_i \sim N(0, 1)$. The error term for each household i's response to question j is constructed such that $e_{ij} \sim N(0, \sigma_j^2)$, and each y_{ij}^* and y_{ij} are constructed using equations (12) and (6), respectively. The Gibbs Sampler algorithm outlined in Section 2.2.2 was run for 30,000 iterations with the first 10,000 draws removed for burn-in.

Table F2 compares the data generating values of our parameters to their estimated posterior mean and 95% credible intervals.[56] Beginning with our estimates of μ, we see that the posterior mean value is qualitatively similar to the data generating value in all cases. Furthermore, the true data generating value of each μ_j falls within the posterior 95% credible interval in all cases. Looking to our estimates of σ^2, we again see that the posterior mean and data generating values are reasonably close to one another. Alternatively, the 95% credible interval of σ^2 does not cover the true value in the single case of question 7. Finally, moving to the estimates of γ_2, we find that while qualitatively similar to their data generating values, the 95% credible intervals of γ only cover the true parameter value in two out of the five total cases. This again

[55] In this study, we only consider the set of 10 FSM questions answered by all households, ignoring the set of 8 additional child food security questions for households with children. In the future, the set of 8 child food security questions can be added to the model, though adult and child food security are thought to be separate latent constructs.

[56] The values in parentheses in Table F2 show 95% credible intervals given by the post-burn-in posterior draw of each parameter corresponding to the specified percentile value. In this specific case, the 20,000 post-burn-in iterations are first ordered from smallest to largest. The left credible interval value is then given by the 500th draw and the right interval as the 19,500th draw. The "NE" designation implies that the specified parameter is fixed and therefore not estimated by the model.

may be due to not having enough iterations to achieve convergence, but it may also be the result of utilizing limited response data.

As an alternative to checking whether the true value falls within the estimated 95% credible interval of each parameter, similar studies often rely on various forms of the Root Mean Squared Error (RMSE) calculated using the data generating and estimated parameter values (Zhu and Stone (2011), Kieftenbeld and Natesan (2012), Broomell and Bhatia (2014)). Taking this approach, we calculate the RMSE of our estimates for each parameter using the full set of post-burn-in draws from the posterior distribution. The RMSE of each parameter type is then averaged, giving us the average RMSE of μ, σ^2, and γ_2. We find that the average RMSE of μ, σ^2, and γ_2 are 0.029, 0.0432, and 0.0437, respectively. In line with other common statistical tests, there is no single RMSE value that signifies adequate parameter retrieval. However, we do find that our RMSE's fall well below the thresholds for convergence used in similar studies (Zhu and Stone (2011), Kieftenbeld and Natesan (2012)).

We now discuss our model's ability to accurately predict each household's latent food security variable. Figure F1 shows the posterior mean value of δ for each household along the x-axis and true data generating values of δ on the y-axis, along with a 45 degree line representing perfect agreement between the two measures. The results shown in Figure F1 provide two key insights. First, our estimated values of δ correspond well with their true values. Households with higher posterior mean values of δ are associated with higher data generating food security levels. Second, Figure F1 shows a substantial amount of binning in our estimates of δ. Specifically, the posterior mean values of δ fall into visually separable groups with discrete jumps across most groups. This binning is a byproduct of relying on discrete ordinal data. For a given set of households with identical responses in the observed data, the model can only partially distinguish between the relative food security level of households within the set. This limitation leads to households with the same responses having similar posterior mean values of δ, forming bins.

Furthermore, in this specific exercise, values of μ, σ^2, and γ are constructed so that they are the same across all questions. With equal parameter values across questions, questions are also treated as equal in so far as their relative importance in determining δ. For example, an individual who answers "3" to question 1 and "1" to all other questions will be placed into the same bin of posterior mean δ's as another individual who answers "3" to question 2 and "1" to all others. Alternatively, variation in parameter values across questions increases both the predictive power of δ and the number of final bins since questions are now distinct from one another in how they relate to δ. Given that assigning identical parameters across questions can therefore be seen as a particularly challenging scenario, we still find that our model produces estimates that generally converge to their true value.

Finally, we compare the relative abilities of both the BGRM and official USDA FSM scale to accurately predict household food security status. First, we designate the 20th percentile of our data generating δ's as the cutoff separating food security and food insecurity, implying that households with values of δ below the 20th percentile are categorized as "food insecure" while households with δ's above the 20th percentile are "food secure". Next we categorize households as either food secure or food insecure using the posterior mean estimates of δ from our BGRM. To then categorize households using the FSM scale, we use the standard approach where households are considered food insecure if they have 3 or more responses indicative of food insecurity.[57] Finally, we redefine the food insecurity cutoff percentile to the 50th and 5th percentiles to evaluate the sensitivity of both measures to changes in underlying food insecurity rates. This involves changing the percentile used for both the data generating δ's and posterior mean δ's while keeping the FSM scale cutoff values fixed.

We measure classification accuracy using the proper match rate and mismatch rate. To define these rates, note that each household i has some true food security status F_i such that $F_i = 1$ if the household is food secure and 0 otherwise based on the previously described threshold percentiles. Given each household's estimated food security status, $\hat{F}_i$, from either the BGRM or FSM scale cutoffs, we then define a proper match $P_i = 1$ as the case where $F_i = \hat{F}_i$ and

[57] Since responses to each of the 10 questions are positively related to food security, responses indicative of food insecurity include a response of 1 or 2 to questions 1, 2, 3, 5, and 10, and a response of 1 to questions 4, 6, 7, 8, and 9.

a mismatch $P_i = 0$ as the case where $F_i \neq \hat{F}_i$. The proper match rate is then calculated as $(\sum_{i=1}^{I} P_i)/I$ and the mismatch rate is given by 1 minus the proper match rate. Table F3 shows the proper match rates and mismatch rates for both the BGRM and FSM scale under our three percentile thresholds of food security. Beginning with the 20th percentile definition in Panel A of Table F3, we see that the proper match rate for the BGRM and FSM scale are 0.94 and 0.52, respectively. This implies that while food security categories estimated using the BGRM were correct 94% of the time, categorizations from the FSM scale were only correct for roughly 52% of households. We find that the BGRM similarly outperforms the FSM scale when the food security cutoff is set to the 50th and 5th percentile in Panel B and Panel C, respectively.

While the BGRM's assignments of food security categories outperform those produced using the USDA's FSM scale with the 3 response rule, we can also adjust the FSM scale such that both measures more closely match the true underlying data. Specifically, we adjust the FSM scale's cutoff of "food insecurity indicative responses" needed to classify a household as food insecure until the share of food secure and food insecure households produced by the FSM scale most closely matches the true data. For example, with the food security threshold set to the 20th percentile, adjusting the FSM scale to require 7 indicative responses produces shares of food secure and food insecure households equal to roughly 77% and 23% of the sample, respectively. This same process is repeated for our alternative thresholds and the match/mismatch rates are then given in Table F4. As expected, Table F4 shows that adjusting the number of responses needed in the FSM scale to mirror the data increases its performance considerably. While the gap in performance between the BGRM and FSM scale categorizations is smaller, the BGRM still outperforms the FSM scale in all cases. Therefore, while the vast majority of studies rely on the traditional 3 response FSM scale to define food security/insecurity regardless of the sample, the BGRM is better able to assign households food security status in our simulated data exercise even in the case where the FSM scale is adjusted to most closely match a known threshold.

Appendix C. Health Outcomes Regressions

We examine the ability of our BGRM posterior mean FSEC measure to capture variation in health outcomes commonly associated with food security. Specifically, we use data from the 2017-18 NHANES to estimate regressions of three binary health outcomes: self-reported obesity, diagnosis of diabetes, and diagnosis of high blood pressure on our BGRM posterior mean estimates. In each regression, we employ a vector of controls which capture gender, age, race, level of education, marital status, citizenship status, household size, and annual household income. We compare the resulting BGRM regression coefficients to regressions of our health outcomes of interest on binary indicators of low, and very low food security defined using the USDA's FSM scale, leaving marginal food security as the reference group; the number of affirmative FSM responses for each adult in the sample; and the respondent-level probabilities of falling within a given quartile of the BGRM posterior distribution.

Given this set of explanatory variables, the estimated regressions for each binary health outcome are represented through the linear model,

$$H_i = \beta_0 + \boldsymbol{\beta}_1 \boldsymbol{Z}_i + \boldsymbol{\beta}_2 \boldsymbol{X}_i + \epsilon_i \tag{C1}$$

where H_i represents a given binary health outcome for adult i, Z_i represents the vector of explanatory variables, X_i represents the vector of control variables, and ϵ_i is a stochastic error term. The parameters of interest are represented by $\boldsymbol{\beta}_1$.

Results from these regressions are shown in Table F5. Table F5 provides several insights. The BGRM posterior means have a statistically significant association with a diagnosis of high blood pressure. A one standard deviation increase in adult BGRM posterior mean FSEC value decreases the probability of being diagnosed with high blood pressure by 2.7 percentage points. This result aligns with *a priori* expectations as higher BGRM posterior mean food security values correspond to higher levels of food security. With each additional affirmative response to the set of FSM questions, adults are 1.1 percentage points more likely to have high blood pressure.

Further, the BGRM posterior mean values explain more of the variation in obesity than the set of USDA-defined FSM food security category indicator variables.

We find little evidence of a statistically significant relationship between the BGRM quartile probabilities and obesity or being diagnosed with diabetes. In regards to being diagnosed with high blood pressure, a one standard deviation increase in the probability of an individual falling within the first, second, or third quartile of the BGRM posterior distribution increases the probability of that individual being diagnosed with high blood pressure by 5.7, 4.6, and 4.6 percentage points, respectively. Conversely, adults that are low food secure according to the FSM scale are 6.5 percentage points more likely to be diagnosed with diabetes, and 7.7 percentage points more likely to be diagnosed with high blood pressure.

Taken together, the results of our regression analyses suggest that there are associations between our BGRM posterior mean estimates, FSM derived measures, and being diagnosed with high blood pressure. We find mostly statistically insignificant effects for obesity and being dianosed with diabetes across BGRM and FSM measures. The adjusted R-squared of our regressions using BGRM posterior mean are similar to those of regressions using affirmative FSM count. Given the similar health outcome regression results produced using both BGRM and FSM scale indicators, the flexibility, ability to measure uncertainty in estimation, and other attractive features of our BGRM provide additional capabilities to those of more common methods with relatively little loss in predictive power and gains for certain outcomes.

Appendix D. Estimation Including Monthly Food Spending

In the primary BGRM specification, our estimation procedure solely relied on respondents' answers to the adult FSM food security questions. For each question in the survey, we estimated up to three parameters: μ_j, σ_j, γ_j.[58] For these questions, we could not simultaneously estimate σ_j and λ_j, as the y^*_{ij} were not observed and therefore σ_j and λ_j were not identified simultaneously.

We now add the monthly food spending (MFS) associated with each participant/household in our sample. The estimation procedure for the parameters corresponding to the original 10 FSM questions remains unchanged. However, for the food spending variable, we can now estimate μ_j, σ_j, and λ_j simultaneously. Estimation of σ_j and λ_j is possible because the value of food spending for the i'th individual is treated as the true value of y^*_{ij} for that individual. Thus for the food spending variable, the conditional posterior distributions of μ_j, σ_j, and λ_j, differ slightly from those for FSM-related parameters. We display these conditional posterior distributions below. These conditional posterior distributions are then sampled from directly in the Gibbs Sampler algorithm. While we use MFS as our continuous variable of interest in this study, the same approach allows for any additional continuous variable(s), highlighting the BGRM's flexibility.

Conditional Posterior of μ_j: The estimation equation for the MFS variable j is given as:

$$y^*_j - \lambda_j \delta = 1_I \mu_j + e_j$$

The μ_j is then drawn from the following full conditional posterior distribution:

$$\mu_j | \delta, \sigma^2, y^*, Y \sim N(v, V) \tag{D1}$$

where $V = \left(\frac{1}{\sigma^2_\mu} + \frac{1'_I 1_I}{\sigma^2_j} \right)^{-1}$ and $v = V \left[\frac{\bar{\mu}}{\sigma^2_\mu} + \frac{1'_I (y^*_j - \lambda_j \delta)}{\sigma^2_j} \right]$.

[58]Gammas were only estimated for FSM questions with more than two potential responses.

Conditional Posterior of σ_j^2: The estimation equation is given as:

$$y_j^* = 1_I\mu_j + \lambda_j\delta + e_j$$

The σ_j^2 is then drawn from the following full conditional posterior distribution:

$$\sigma_j^2|\mu, \delta, \gamma, Y \sim IG(a, b) \tag{D2}$$

where $a = \alpha + \frac{I}{2}$, and $b = \left[\frac{1}{\beta} + \frac{(y_j^* - 1_I\mu_j - \lambda_j\delta)'(y_j^* - 1_I\mu_j - \lambda_j\delta)}{2}\right]^{-1}$.

Conditional Posterior of λ_j: The estimation equation is given as:

$$y_j^* - 1_I\mu_j = \lambda_j\delta + e_j$$

The λ_j is then drawn from the following full conditional posterior distribution:

$$\lambda_j|\delta, \sigma^2, y^*, Y \sim \mathcal{N}(G, g) \tag{D3}$$

where $g = \left(\frac{\delta'\delta}{\sigma_j^2} + \frac{1}{\sigma_\lambda^2}\right)^{-1}$, and $G = g\left[\frac{\delta'(y_j^* - 1_I\mu_j)}{\sigma_j^2} + \frac{\lambda}{\sigma_\lambda^2}\right]$.

Appendix E. Estimation Using Current Population Survey (2018)

The FSM was first piloted in the 1995 Current Population Survey (CPS) Food Security Supplement (FSS). Ever since, the USDA has used annual CPS data to calculate official national food security rates in the United States annually. Consequently, to examine the robustness of our BGRM and identify potential implications for national food security rates, we estimate latent food security using FSM responses in the 2018 CPS FSS.[59] The 2018 CPS FSS includes 37,300 interviewed households with 89,665 person records. The CPS FSS contains external and internal screening protocols to reduce respondent survey burden and avoid questions that may be inappropriate for some respondents given information provided earlier in the survey. Households that pass the external screen are exempt from taking the entire FSM and they are automatically assigned to the USDA's full food security category. Furthermore, the adult FSM questionnaire in the CPS includes two internal screeners. Households that register no food stress in a given set of adult internal screener questions are not asked the remaining adult FSM questions, and their responses to skipped questions are assumed to be negative.[60]

Unlike the CPS, the NHANES does not use an external screener for the FSM, but it does use internal screeners. To make our CPS results more comparable to the NHANES, we exclude all households from our CPS sample that were screened out of completing the full survey by the external screener. This restriction reduces our CPS sample size from 37,300 households to 14,326 households. Finally, for consistency with the NHANES BGRM estimation exercises, we also remove all extreme response households. This restriction yields our final sample of 6,065 households. Using this reduced sample, we estimate our model and produce Figures F3 and F4 shown below, counterparts to Figures 2 and 3 in the main text produced using the NHANES sample.

[59] We use the 2018 FSS for consistency with our 2017-18 NHANES results.

[60] For more information regarding screening procedures for the 2018 CPS FSS, see 2018 CPS FSS Technical Documentation.

Appendix F. Additional Tables and Figures for Chapter 2

Table F1. Simulated Data Exercise Data Generating Parameters

j	μ	σ^2	γ_2
1	1	0.5	1.5
2	1	0.5	1.5
3	1	0.5	1.5
4	1	0.5	∞
5	1	0.5	1.5
6	1	0.5	∞
7	1	0.5	∞
8	1	0.5	∞
9	1	0.5	∞
10	1	0.5	1.5

Note: Each j represents a FSM response question that households without children must answer to be assigned a domestic household food security category. Each μ represents question j specific intercept parameter. Each σ^2 represents question j's specific variance associated with the elicited food security level of household i. Each γ_2 represents a response threshold for question j.

Table F2. Simulated Data Exercise Data Generating Parameters vs. Posterior Means and 95% Credible Intervals

j	μ	$\hat{\mu}$	σ^2	$\hat{\sigma}^2$	γ_2	$\hat{\gamma}_2$
1	1	0.9972	0.5	0.4736	1.5	1.4572
		(0.9470,1.0468)		*(0.4268,0.5224)*		*(1.4050,1.5021)*
2	1	0.9922	0.5	0.5048	1.5	1.4921
		(0.9499,1.0360)		*(0.4613,0.5508)*		*(1.4553,1.5304)*
3	1	0.9754	0.5	0.4751	1.5	1.4396
		(0.9335,1.0178)		*(0.4339,0.5175)*		*(1.3986,1.4745)*
4	1	1.0029	0.5	0.4585	∞	NE
		(0.9505,1.0579)		*(0.3980,0.5252)*		
5	1	0.9970	0.5	0.5315	1.5	1.4632
		(0.9504,1.0441)		*(0.4841,0.5825)*		*(1.4226,1.4998)*
6	1	0.9684	0.5	0.4502	∞	NE
		(0.9169,1.0209)		*(0.3877,0.5181)*		
7	1	0.9820	0.5	0.4186	∞	NE
		(0.9312,1.0354)		*(0.3608,0.4854)*		
8	1	1.0096	0.5	0.4918	∞	NE
		(0.9557,1.0634)		*(0.4244,0.5651)*		
9	1	1.0090	0.5	0.5125	∞	NE
		(0.9577,1.0619)		*(0.4451,0.5885)*		
10	1	1.0247	0.5	0.5080	1.5	1.5368
		(0.9825,1.0670)		*(0.4654,0.5530)*		*(1.5062,1.5669)*

Note: Posterior mean value of μ is qualitatively close to the data generating value in all cases and the true data generating value of μ falls within the posterior 95% credible interval in all cases. Posterior mean and data generating values of σ^2 are reasonably close to one another. The 95% credible intervals of σ^2 cover the true value in all but one case. Posterior mean and data generating values of γ_2 are qualitatively similar. The 95% credible intervals of γ^2 cover the true value in 2 out of 5 cases.

Figure F1. Data Generating Values of δ vs. Posterior Mean Draws of δ

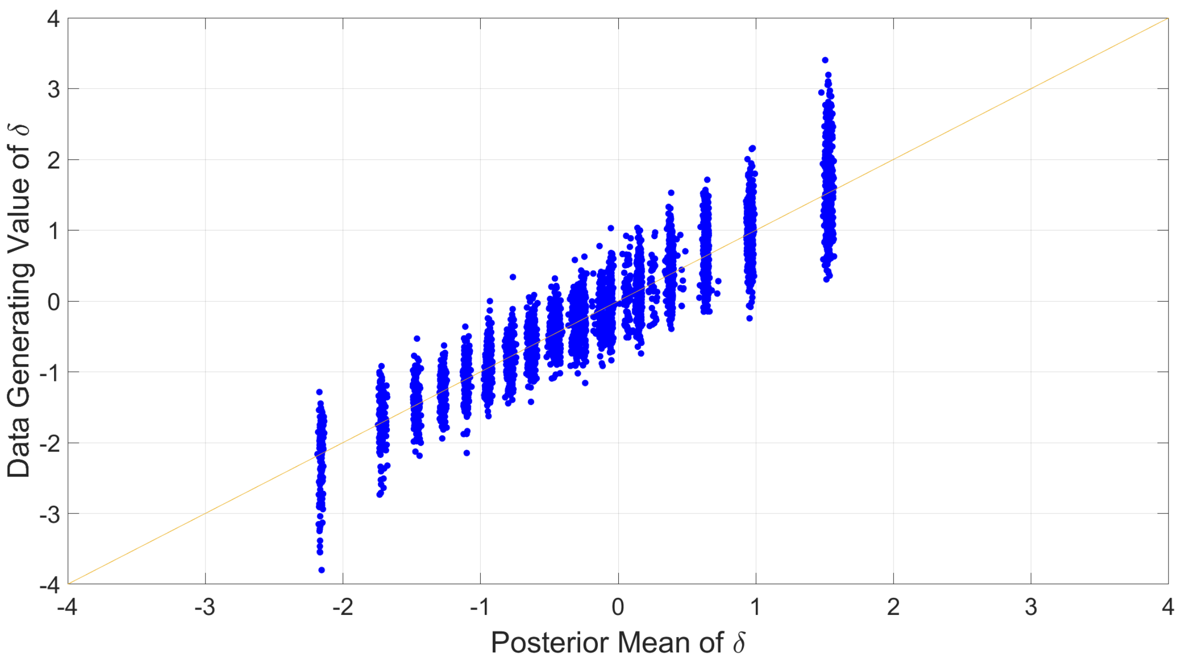

Note: Estimated values of δ correspond well with their true values. Households with higher posterior mean values of δ are associated with higher data generating food security levels. There is considerable binning in our estimates of δ. This binning is a byproduct of the fact that we rely on ordinal polytomous data.

Table F3. Food Security Categorization Accuracy of BGRM and FSM Scale

	Proper Match Rate	Mismatch Rate
Panel A: Food Insecurity Defined as δ below the 20th Percentile		
	Proper Match Rate	Mismatch Rate
BGRM	0.9416	0.0584
FSM Scale	0.5194	0.4806
Panel B: Food Insecurity Defined as δ below the 50th Percentile		
	Proper Match Rate	Mismatch Rate
BGRM	0.9080	0.0920
FSM Scale	0.8078	0.1922
Panel C: Food Insecurity Defined as δ below the 5th Percentile		
	Proper Match Rate	Mismatch Rate
BGRM	0.9720	0.0280
FSM Scale	0.3694	0.6306

Note: In panel A, we detect a large difference between the BGRM and FSM scale proper match rates. This result implies that food security categories estimated using the BGRM were correct 94% of the time, compared to only 52% of the time for the FSM scale. The scale of the difference between the BGRM scale and the FSM scale proper match rate stays relatively the same in both panel B and C where the security cutoffs are set to the 50th and 5th percentile respectively. Thus, generally, food security cutoffs estimated using the BGRM outperform the ones estimated using the FSM scale.

Table F4. Food Security Categorization Accuracy of BGRM and Adjusted FSM Scale

	Proper Match Rate	Mismatch Rate
Panel A: Food Insecurity Defined as δ below the 20th Percentile		
BGRM	0.9416	0.0584
FSM Scale	0.9158	0.0842
Panel B: Food Insecurity Defined as δ below the 50th Percentile		
	Proper Match Rate	Mismatch Rate
BGRM	0.9080	0.0920
FSM Scale	0.8810	0.1190
Panel C: Food Insecurity Defined as δ below the 5th Percentile		
	Proper Match Rate	Mismatch Rate
BGRM	0.9720	0.0280
FSM Scale	0.9654	0.0346

Note: We adjust the FSM scale's number of "food insecurity indicative responses" needed to classify a household as food secure until the share of food secure and food insecure households produced by the FSM scale most closely matches the true data. The performance of the FSM scale improves considerably, however the BGRM still outperforms the FSM scale in all cases.

Table F5. Coefficient Estimates of the BGRM Posterior Mean of Food Security vs. FSM Scale Food Security Categories on Adult Households Health Outcomes

	Obesity				Diabetes				**High Blood Pressure**			
	(1)	**(2)**	**(3)**	**(4)**	(5)	(6)	(7)	(8)	**(9)**	**(10)**	**(11)**	**(12)**
BGRM Posterior Mean	0.004 (0.015)				-0.007 (0.011)				-0.030** (0.014)			
Affirmative FSM Count		-0.002 (0.006)				0.003 (0.004)				0.011** (0.005)		
1st Quartile			0.054 (0.070)				0.025 (0.052)				0.154** (0.063)	
2nd Quartile			0.060 (0.063)				0.064 (0.046)				0.186*** (0.056)	
3rd Quartile			0.175 (0.128)				0.040 (0.095)				0.226** (0.115)	
Very Low Food Secure				-0.022 (0.035)				0.020 (0.026)				0.043 (0.032)
Low Food Secure				0.005 (0.030)				0.065*** (0.022)				0.077*** (0.027)
Adj. R-squared	0.0463	0.0463	0.0467	0.0460	0.1467	0.1467	0.1464	0.1514	0.1961	0.1961	0.1994	0.1977
Observations (N)	1,394	1,394	1,394	1,394	1,378	1,378	1,378	1,378	1,407	1,407	1,407	1,407
Mean DV	0.42	0.42	0.42	0.42	0.18	0.18	0.18	0.18	0.37	0.37	0.37	0.37

Notes: There are 3 binary outcomes employed: self-reported obesity, diagnosis of diabetes, and diagnosis of high blood pressure. For each of these outcomes four regressions are displayed respectively: (1) A regression where the outcome is regressed on the BGRM Posterior Means of Food Security for the sample; (2) A regression where the outcome is regressed on the number of FSM questions adults answered affirmatively; (3) A regression where the outcome is regressed on the probabilities of individuals' draws of latent food security falling within a given quartile ; (4) A regression where the outcome is regressed on the FSM Scale Categories of Food Security. The same controls are used accross all regressions: gender, age, education, marital status, citizenship, race, household size and annual household income. Mean DV stands for mean of dependent variable. The sample consists of the 2017 - 2018 adult food security data of the National Health and Nutrition Examination Survey (NHANES), excluding households who answered all questions affirmatively or negatively. (* 10%, ** 5%, and *** 1%).

Figure F2. Posterior Mean Values of BGRM Latent Construct Using FSM and Continuous Monthly Food Spending (MFS) with 95% Credible Intervals Using NHANES 2017 - 2018

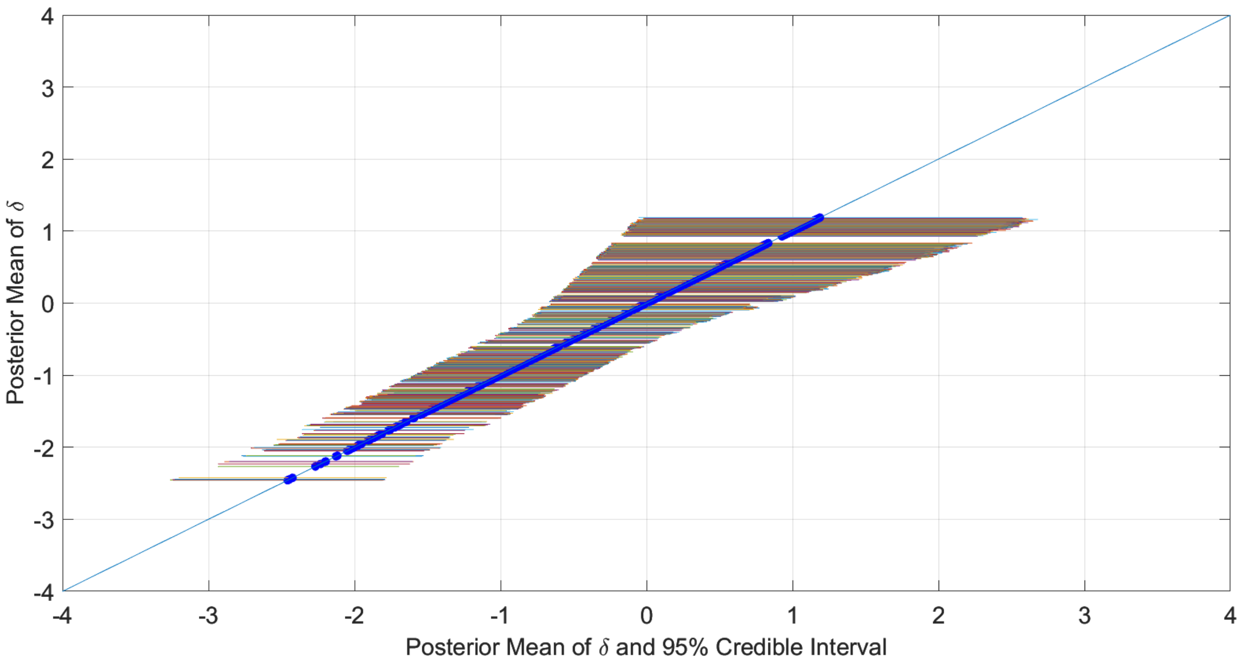

Note: The sample consists of households with responses to the 10 adult FSM questions from the 2017 - 2018 National Health and Nutrition Examination Survey (NHANES). We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 1,659 households. Each blue dot on the 45^o degree line represents the BGRM posterior mean latent construct using FSM and continuous Monthly Food Spending (MFS) value for a given household in the sample. The colored lines running through the blue dots represent the magnitude of the 95% credible intervals for each household.

Figure F3. Posterior Mean Values of BGRM Latent Food Security with 95% Credible Intervals Using CPS 2018 Food Security Supplement

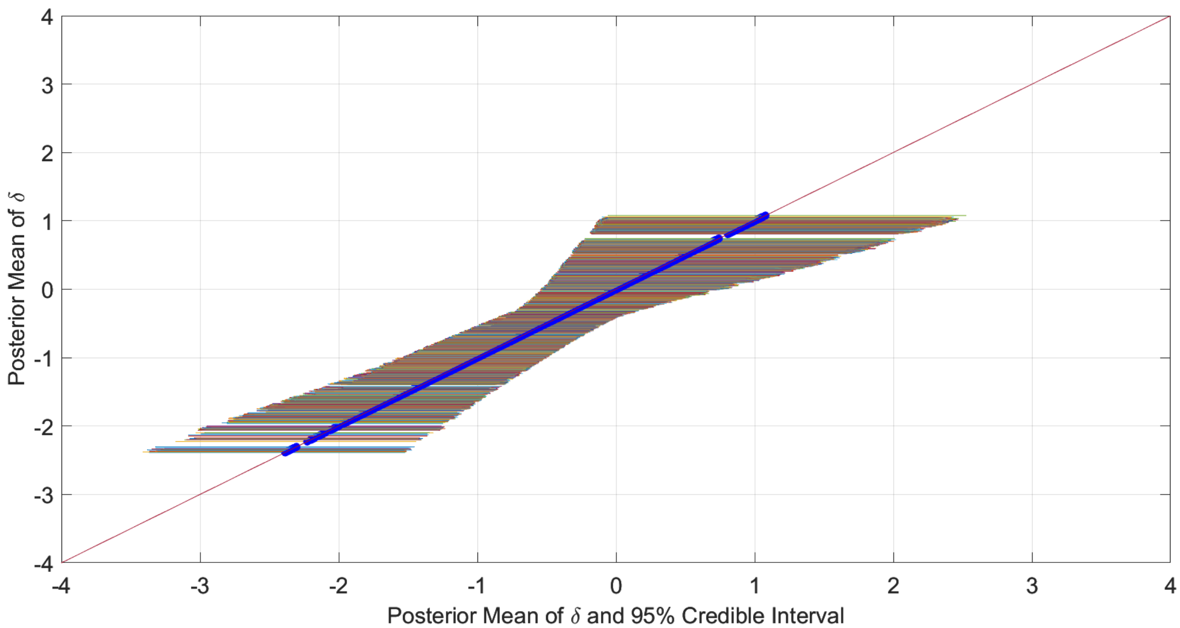

Note: The sample consists of households with responses to the 10 adult FSM questions from the 2018 Current Population Survey (CPS) Food Security Supplement. We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 6,065 households. Each blue dot on the 45^o degree line represents the BGRM posterior mean latent food security value for a given household in the sample. The colored lines running through the blue dots represent the magnitude of the 95% credible intervals for each household.

Figure F4. Distribution of BGRM Posterior Mean Food Security by USDA FSM Scale Food Security Categories Using CPS 2018 Food Security Supplement

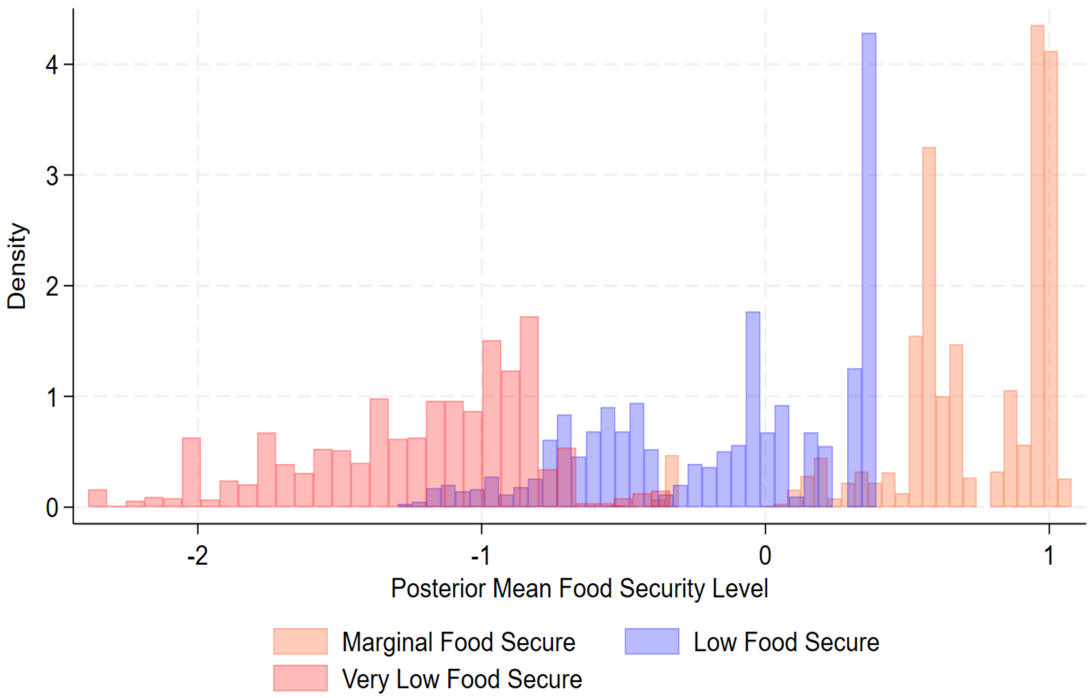

Note: The sample consists of households with responses to the 10 adult FSM questions from the 2018 Current Population Survey (CPS) Food Security Supplement. We exclude households who answered all food security questions affirmatively or negatively. The remaining analytical sample contains 6,065 households. Each household in our sample has a BGRM posterior mean value of food security that is estimated by the model. Using the responses of the same households to the adult food security questionnaire questions, households are categorized into 1 of 3 food security categories based on the standard USDA food security category definitions. Marginal Food Secure households have 1-2 affirmative responses to the adult FSM questionnaire, followed by 3-5, and 6-10 affirmative responses for Low and Very Low Food Secure households, respectively. Histograms depicting the household distribution of BGRM posterior mean food security values for households belonging to the same FSM food security category are displayed using the 3 colors above.

Appendix G. Additional Tables and Figures for Chapter 3

Table G1. Coefficient Estimates on the Unemployment Rate and Adult Labor Force Participation Rate in Teenage Substance-Use Equations (Preferred Specification vs Additional Models)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 1. Preferred Specification: All controls, state and year fixed effects & state specific linear time trends					
Unemployment	0.004	0.004	0.004	0.002	0.007
Rate	(0.003)	(0.003)	(0.004)	(0.002)	(0.004)
	[-0.001,0.009]	[-0.002,0.010]	[-0.004,0.011]	[-0.003,0.006]	[-0.000,0.015]
ALFP	-0.003	-0.002	0.000	-0.000	-0.000
	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)
	[-0.008,0.002]	[-0.009,0.005]	[-0.011,0.011]	[-0.008,0.008]	[-0.009,0.008]
Adjusted R^2	0.053	0.065	0.084	0.031	0.060
Panel 2. Only State and year fixed effects. Exclude state specific linear time trends					
Unemployment	0.000	-0.000	-0.002	0.000	0.004
Rate	(0.002)	(0.003)	(0.004)	(0.003)	(0.004)
	[-0.005,0.005]	[-0.005,0.005]	[-0.009,0.005]	[-0.006,0.007]	[-0.004,0.012]
ALFP	-0.002	-0.002	-0.003	-0.001	0.000
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
	[-0.005,0.002]	[-0.006,0.002]	[-0.009,0.004]	[-0.007,0.004]	[-0.006,0.006]
Adjusted R^2	0.030	0.038	0.058	0.010	0.026
Panel 3. No controls. No fixed effects. No state specific linear time trends					
Unemployment	-0.002	-0.002	-0.000	-0.002	0.003
Rate	(0.001)	(0.002)	(0.003)	(0.002)	(0.002)
	[-0.005,0.001]	[-0.006,0.002]	[-0.005,0.005]	[-0.005,0.002]	[-0.000,0.007]
ALFP	0.003*	0.004*	0.009**	-0.002	0.008**
	(0.001)	(0.002)	(0.003)	(0.001)	(0.003)
	[0.000,0.006]	[0.001,0.008]	[0.003,0.016]	[-0.004,0.000]	[0.003,0.014]
Adjusted R^2	0.003	0.004	0.006	0.000	0.003
Mean Use Rate (%)	7.10	9.50	22.33	21.78	41.95
Observations	198573	198573	198573	204119	193125

Notes: Linear Probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. ALFP represents adult labor force participation rate. All controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table G2. Coefficient Estimates on the Unemployment Rate and Adult Labor Force Participation Rate by Gender, Race, and Age (All Panels Use the Preferred Specification and Vary the Subgroup Analyzed: All, Male, Female, 12 - 14 Years Old, 15 - 18 Years Old, White, Black)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 1. All					
Unemployment	0.004	0.004	0.004	0.002	0.007
Rate	(0.003)	(0.003)	(0.004)	(0.002)	(0.004)
	[-0.001,0.009]	[-0.002,0.010]	[-0.004,0.011]	[-0.003,0.006]	[-0.000,0.015]
ALFP	-0.003	-0.002	0.000	-0.000	-0.000
	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)
	[-0.008,0.002]	[-0.009,0.005]	[-0.011,0.011]	[-0.008,0.008]	[-0.009,0.008]
Observations	198573	198573	198573	204119	193125
Panel 2. Male					
Unemployment	0.002	0.002	0.001	0.001	0.005
Rate	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)
	[-0.004,0.008]	[-0.005,0.009]	[-0.007,0.008]	[-0.004,0.006]	[-0.004,0.014]
ALFP	-0.004	-0.002	0.005	0.002	0.004
	(0.003)	(0.004)	(0.006)	(0.005)	(0.005)
	[-0.011,0.003]	[-0.010,0.006]	[-0.007,0.017]	[-0.009,0.012]	[-0.006,0.015]
Observations	97004	97004	97004	99991	94258
Panel 3. Female					
Unemployment	0.005	0.005	0.006	0.002	0.009*
Rate	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
	[-0.001,0.011]	[-0.002,0.013]	[-0.002,0.015]	[-0.006,0.009]	[0.001,0.018]
ALFP	-0.002	-0.002	-0.006	-0.002	-0.005
	(0.003)	(0.004)	(0.006)	(0.004)	(0.005)
	[-0.007,0.003]	[-0.010,0.006]	[-0.017,0.006]	[-0.011,0.007]	[-0.014,0.004]
Observations	101569	101569	101569	104128	98867
Panel 4. 12 - 14 years old					
Unemployment	0.006	0.004	0.009*	-0.002	0.006
Rate	(0.004)	(0.003)	(0.003)	(0.004)	(0.006)
	[-0.002,0.013]	[-0.003,0.011]	[0.002,0.016]	[-0.011,0.007]	[-0.006,0.017]
ALFP	-0.003	-0.002	0.004	-0.006	-0.006
	(0.003)	(0.003)	(0.005)	(0.007)	(0.007)
	[-0.009,0.003]	[-0.008,0.005]	[-0.007,0.014]	[-0.020,0.009]	[-0.020,0.007]
Observations	19940	19940	19940	20435	19436

Note: Table continues in next page

Table G2. Coefficient Estimates on the Unemployment Rate and Adult Labor Force Participation Rate by Gender, Race, and Age (All Panels use the Preferred Specification and Vary the Subgroup Analyzed: All, Male, Female, 12 - 14 Years Old, 15 - 18 Years Old, White, Black) (continued)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 5. 15 - 18 years old					
Unemployment Rate	0.003	0.003	0.002	0.002	0.007
	(0.003)	(0.003)	(0.004)	(0.002)	(0.004)
	[-0.002,0.008]	[-0.003,0.010]	[-0.006,0.011]	[-0.003,0.007]	[-0.001,0.015]
ALFP	-0.003	-0.002	-0.001	0.000	0.000
	(0.003)	(0.004)	(0.006)	(0.004)	(0.004)
	[-0.008,0.003]	[-0.010,0.006]	[-0.012,0.011]	[-0.008,0.009]	[-0.009,0.009]
Observations	178633	178633	178633	183684	173689
Panel 6. White					
Unemployment Rate	0.003	0.003	0.004	-0.002	0.006
	(0.004)	(0.004)	(0.004)	(0.003)	(0.005)
	[-0.004,0.010]	[-0.005,0.011]	[-0.004,0.011]	[-0.008,0.003]	[-0.004,0.015]
ALFP	-0.006*	-0.005	-0.002	-0.001	0.003
	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
	[-0.013,-0.000]	[-0.014,0.004]	[-0.012,0.009]	[-0.011,0.009]	[-0.007,0.012]
Observations	86492	86492	86492	88381	84968
Panel 7. Black					
Unemployment Rate	-0.001	-0.002	0.003	0.010*	0.010
	(0.002)	(0.002)	(0.005)	(0.005)	(0.007)
	[-0.005,0.003]	[-0.007,0.002]	[-0.007,0.013]	[0.000,0.020]	[-0.004,0.024]
ALFP	0.003	0.004	0.003	0.005	0.006
	(0.002)	(0.002)	(0.005)	(0.006)	(0.006)
	[-0.001,0.007]	[-0.001,0.008]	[-0.006,0.013]	[-0.007,0.017]	[-0.007,0.019]
Observations	42248	42248	42248	43698	40812
Mean Use Rate (%)	7.10	9.50	22.33	21.78	41.95

Note: Linear probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. Controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. All models include state specific linear time trends. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table G3. Coefficient Estimates on the Unemployment Rate and Adult Labor Force Participation Rate in Teenage Substance-Use Equations (All Panels Display a Modified Version of the Preferred Specification or Sample Used. Each Panel Title Lists the Change)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 1. Exclude State Specific Linear Time Trends					
Unemployment	0.002	0.001	-0.002	0.001	0.002
Rate	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
	[-0.003,0.006]	[-0.003,0.006]	[-0.009,0.004]	[-0.004,0.006]	[-0.005,0.009]
ALFP	-0.000	0.001	-0.001	-0.001	-0.001
	(0.002)	(0.003)	(0.004)	(0.003)	(0.004)
	[-0.005,0.005]	[-0.004,0.006]	[-0.008,0.007]	[-0.008,0.006]	[-0.008,0.007]
Panel 2. Include 2-year lag and lead. Include State Specific Linear Time Trends					
Unemployment	0.005	0.005	0.005	0.004	0.010*
Rate	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)
	[-0.001,0.011]	[-0.002,0.013]	[-0.003,0.014]	[-0.002,0.010]	[0.002,0.019]
ALFP	-0.002	-0.001	-0.001	-0.002	0.001
	(0.003)	(0.004)	(0.006)	(0.004)	(0.005)
	[-0.008,0.003]	[-0.009,0.007]	[-0.012,0.011]	[-0.011,0.007]	[-0.008,0.010]
2-year Lagged	-0.000	0.000	-0.002	-0.005	0.001
unemployment rate	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)
	[-0.004,0.004]	[-0.006,0.006]	[-0.011,0.007]	[-0.012,0.002]	[-0.008,0.010]
2-year Lead	-0.003	-0.004	-0.002	-0.002	-0.007
unemployment rate	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
	[-0.008,0.001]	[-0.009,0.002]	[-0.008,0.004]	[-0.008,0.004]	[-0.014,0.000]
Panel 3. Include 2-year lag and lead. Exclude State Specific Linear Time Trends					
Unemployment	0.005	0.005	0.002	0.003	0.007
Rate	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)
	[-0.001,0.011]	[-0.002,0.012]	[-0.006,0.010]	[-0.003,0.010]	[-0.002,0.015]
ALFP	0.000	0.001	-0.001	-0.002	0.000
	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)
	[-0.005,0.005]	[-0.005,0.007]	[-0.010,0.007]	[-0.009,0.005]	[-0.008,0.008]
2-year lagged	-0.001	-0.001	-0.004	-0.003	-0.000
unemployment rate	(0.002)	(0.003)	(0.004)	(0.003)	(0.004)
	[-0.005,0.003]	[-0.006,0.005]	[-0.012,0.004]	[-0.010,0.003]	[-0.008,0.008]
2-year lead	-0.005	-0.005	-0.004	-0.002	-0.008*
unemployment rate	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
	[-0.010,0.000]	[-0.011,0.000]	[-0.011,0.003]	[-0.009,0.006]	[-0.016,-0.001]
Mean Use Rate (%)	7.10	9.50	22.33	21.78	41.95
Observations	198573	198573	198573	204119	193125

Note: Table continues in next page

Table G3. Coefficient Estimates on the Unemployment Rate and Adult Labor Force Participation Rate in Teenage Substance-Use Equations (All Panels Display a Modified Version of the Preferred Specification or Sample Used. Each Panel Title Lists the Change) (continued)

Explanatory Var.	Heavy Cigarette Use	Regular Cigarette Use	Current Cigarette Use	Current Marijuana Use	Current Alcohol Use
Panel 4. Exclude 7 least populous states					
Unemployment Rate	0.004	0.004	0.004	0.001	0.007
	(0.003)	(0.003)	(0.004)	(0.002)	(0.004)
	[-0.002,0.009]	[-0.002,0.010]	[-0.004,0.011]	[-0.003,0.006]	[-0.000,0.015]
ALFP	-0.003	-0.002	-0.000	-0.000	0.001
	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)
	[-0.008,0.002]	[-0.009,0.005]	[-0.011,0.011]	[-0.009,0.008]	[-0.008,0.009]
Observations	194422	194422	194422	199855	189237
Panel 5. Exclude 4 most populous states					
Unemployment Rate	0.004	0.004	0.004	0.001	0.007
	(0.003)	(0.004)	(0.005)	(0.002)	(0.005)
	[-0.003,0.010]	[-0.004,0.011]	[-0.006,0.013]	[-0.004,0.006]	[-0.002,0.017]
ALFP	-0.002	-0.000	0.000	0.001	0.000
	(0.003)	(0.004)	(0.006)	(0.005)	(0.005)
	[-0.008,0.003]	[-0.009,0.008]	[-0.012,0.013]	[-0.009,0.011]	[-0.010,0.010]
Observations	128310	128310	128310	131296	124263
Mean Use Rate (%)	7.10	9.50	22.33	21.78	41.95

Notes: Linear Probability models were estimated. Observations stand for the number of non-missing values for the dependent variable in the sample. YRBSS weights used. All controls include race, gender, age, grade; cigarette and beer excise taxes, cigarette indoor use laws, medical and recreational marijuana laws, state ACA Medicaid expansion status, indoor vaping restrictions, minimum legal purchase laws, e-cigarette taxes, state fixed effects, and year fixed effects. Standard errors clustered at the state level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

Ananat, Elizabeth Oltmans, Anna Gassman-Pines, Dania V Francis, and Christina M Gibson-Davis, "Children left behind: The effects of statewide job loss on student achievement," Technical Report, National Bureau of Economic Research 2011.

Andrich, David, *Rasch models for measurement*, SAGE, 1988.

_ and Ida Marais, *A course in Rasch measurement theory*, Springer, 2019.

Arkes, Jeremy, "Does the economy affect teenage substance use?," *Health economics*, 2007, *16* (1), 19–36.

_ , "Recessions and the participation of youth in the selling and use of illicit drugs," *International Journal of Drug Policy*, 2011, *22* (5), 335–340.

_ , "How does youth cigarette use respond to weak economic periods? Implications for the current economic crisis," *Substance use & misuse*, 2012, *47* (4), 375–382.

Ayllón, Sara and Natalia N Ferreira-Batista, "Unemployment, drugs and attitudes among European youth," *Journal of health economics*, 2018, *57*, 236–248.

Bachman, Jerald G, Lloyd D Johnson, and Patrick M O'Malley, "Explaining recent increases in students' marijuana use: impacts of perceived risks and disapproval, 1976 through 1996.," *American journal of public health*, 1998, *88* (6), 887–892.

Baker, Frank B, *The basics of item response theory*, ERIC, 2001.

Balistreri, Kelly Stamper, "A decade of change: measuring the extent, depth and severity of food insecurity," *Journal of family and economic issues*, 2016, *37*, 373–382.

Bandura, Albert and Richard H Walters, *Social learning theory*, Vol. 1, Prentice hall Englewood Cliffs, NJ, 1977.

Barnes, Grace M and John W Welte, "Patterns and predictors of alcohol use among 7-12th grade students in New York State.," *Journal of studies on alcohol*, 1986, *47* (1), 53–62.

Barrett, Christopher and Erin C. Lentz, "Hunger and Food Insecurity," in "The Oxford Handbook of the Social Science of Poverty," Oxford University Press, 05 2016.

Barrett, Christopher B, "Measuring food insecurity," *Science*, 2010, *327* (5967), 825–828.

Beatty, Timothy KM, Marianne P Bitler, Xinzhe Huang Cheng, and Cynthia Van der Werf, "SNAP and paycheck cycles," *Southern Economic Journal*, 2019, *86* (1), 18–48.

Bickel, Gary, Mark Nord, Cristofer Price, William Hamilton, John Cook et al., "Guide to measuring household food security," 2000.

Bickel, Gary W, William L Hamilton et al., "Household food security in the United States in 1995: technical report of the food security measurement project," 1997.

Black, Dan A, Terra G McKinnish, and Seth G Sanders, "Does the availability of high-wage jobs for low-skilled men affect welfare expenditures? Evidence from shocks to the steel and coal industries," *Journal of Public Economics*, 2003, *87* (9-10), 1921–1942.

_ , _ , and _ , "Tight labor markets and the demand for education: Evidence from the coal boom and bust," *ILR Review*, 2005, *59* (1), 3–16.

Black, Dan, Terra McKinnish, and Seth Sanders, "The economic impact of the coal boom and bust," *The Economic Journal*, 2005, *115* (503), 449–476.

Blanchflower, David G and Andrew J Oswald, *The wage curve*, MIT press, 1994.

Bock, R Darrell and Murray Aitkin, "Marginal maximum likelihood estimation of item parameters: Application of an EM algorithm," *Psychometrika*, 1981, *46* (4), 443–459.

Bor, Jacob, Sanjay Basu, Adam Coutts, Martin McKee, and David Stuckler, "Alcohol use during the Great Recession of 2008–2009," *Alcohol and alcoholism*, 2013, *48* (3), 343–348.

Bradford, W David and William D Lastrapes, "A prescription for unemployment? Recessions and the demand for mental health drugs," *Health economics*, 2014, *23* (11), 1301–1325.

Brener, Nancy D, Laura Kann, Shari Shanklin, Steve Kinchen, Danice K Eaton, Joseph Hawkins, and Katherine H Flint, "Methodology of the youth risk behavior surveillance system—2013," *Morbidity and Mortality Weekly Report: Recommendations and Reports*, 2013, *62* (1), 1–20.

Breunig, Robert V and Indraneel Dasgupta, "A Theoretical and Empirical Evaluation of The Functional Forms Used to Estimate The Food Expenditure Equation of Food Stamp Recipients: Comment," *American Journal of Agricultural Economics*, 2002, *84* (4), 1156–1160.

Brook, Judith S, David W Brook, and Martin Whiteman, "Older sibling correlates of younger sibling drug use in the context of parent-child relations.," *Genetic, Social, and General Psychology Monographs*, 1999, *125* (4), 451–468.

_ , Martin Whiteman, Ann Scovell Gordon, and David W Brook, "The role of older brothers in younger brothers' drug use viewed in the context of parent and peer influences," *The Journal of genetic psychology*, 1990, *151* (1), 59–75.

Broomell, Stephen B and Sudeep Bhatia, "Parameter recovery for decision modeling using choice data.," *Decision*, 2014, *1* (4), 252.

Carpenter, Christopher S, Chandler B McClellan, and Daniel I Rees, "Economic conditions, illicit drug use, and substance use disorders in the United States," *Journal of Health Economics*, 2017, *52*, 63–73.

Casella, George and Edward I George, "Explaining the Gibbs sampler," *The American Statistician*, 1992, *46* (3), 167–174.

Castner, Laura and Juliette Henke, "Benefit redemption patterns in the supplemental nutrition assistance program," Technical Report, Mathematica Policy Research 2011.

Caulkins, Jonathan P, "Can supply factors suppress marijuana use by youth," *Drug Policy Analysis Bulletin*, 1999, *7*, 3–5.

Caulkins, JP, "Drug prevention: The paradox of timing," *Federation of American Scientists' Drug Policy Analysis Bulletin*, 1998, *5*, 1–3.

Cawley, John and Christopher J Ruhm, "The economics of risky health behaviors, Vol. 2," 2012.

Chamberlain, Gary, "Multivariate Regression Models for Panel Data," *Journal of econometrics*, 1982, *18* (1), 5–46.

Clark, Andrew, Andreas Knabe, and Steffen Rätzel, "Boon or bane? Others' unemployment, well-being and job insecurity," *Labour economics*, 2010, *17* (1), 52–61.

Coleman-Jensen, Alisha Judith, "U.S. food insecurity status: toward a refined definition," *Social Indicators Research*, 2010, *95*, 215–230.

Cook, Philip J and Michael J Moore, "Environment and persistence in youthful drinking patterns," in "Risky behavior among youths: An economic analysis," University of Chicago Press, 2001, pp. 375–438.

Cotti, Chad D, John M Gordanier, and Orgul D Ozturk, "Does Distributing SNAP Benefits Later in The Month Smooth Expenditures?," *Food Policy*, 2021, *104*, 102123.

Cotti, Chad, Richard A Dunn, and Nathan Tefft, "The great recession and consumer demand for alcohol: a dynamic panel-data analysis of US households," *American Journal of Health Economics*, 2015, *1* (3), 297–325.

Council, NRC National Research, "Food Insecurity and Hunger in the United States: An Assessment of the Measure. Washington, DC. The National Academies Press.," `https://doi.org/10.17226/11578` 2006.

Currie, Janet and Erdal Tekin, "Is there a link between foreclosure and health?," *American Economic Journal: Economic Policy*, 2015, *7* (1), 63–94.

Dahl, Gordon B and Lance Lochner, "The impact of family income on child achievement: Evidence from the earned income tax credit," *American Economic Review*, 2012, *102* (5), 1927–1956.

Davis, Will, Jose O Xilau, Rusty Tchernis, and Christian A Gregory, "A Flexible Model of Food Security: Estimation and Implications for Prediction," Technical Report, National Bureau of Economic Research 2024.

Deb, Partha and Pravin K Trivedi, "Demand for medical care by the elderly: a finite mixture approach," *Journal of applied Econometrics*, 1997, *12* (3), 313–336.

_ and _ , "Finite Mixture for Panels with Fixed Effects," *Journal of Econometric Methods*, 2013, *2* (1), 35–51.

Dee, Thomas S, “Alcohol abuse and economic conditions: evidence from repeated cross-sections of individual-level data,” *Health economics*, 2001, *10* (3), 257–270.

Dempster, Arthur P, Nan M Laird, and Donald B Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 1977, *39* (1), 1–22.

Dooley, David and Ralph Catalano, “The epidemiology of economic stress,” *American journal of community psychology*, 1984, *12* (4), 387–409.

Dorfman, Jeffrey H, Christian Gregory, Zhongyuan Liu, and Ran Huo, “Re-examining the SNAP benefit cycle allowing for heterogeneity,” *Applied Economic Perspectives and Policy*, 2019, *41* (3), 404–433.

Dutta, Indranil, Craig Gundersen, and Prasanta K Pattanaik, “Measures of food insecurity at the household level,” *Food security indicators, measurement, and the impact of trade openness: series: WIDER studies in development economics*, 2007, pp. 42–61.

Edwards, Ryan, “American time use over the business cycle,” *New York City: Mimeo, Queens College, CUNY*, 2011.

Ellickson, Phyllis L, Rebecca L Collins, and Robert M Bell, “Adolescent use of illicit drugs other than marijuana: how important is social bonding and for which ethnic groups?,” *Substance use & misuse*, 1999, *34* (3), 317–346.

Forster, J, V Chen, T Blaine, C Perry, and T Toomey, “Social exchange of cigarettes by youth,” *Tobacco control*, 2003, *12* (2), 148–154.

Foster, James, Joel Greer, and Erik Thorbecke, “A class of decomposable poverty measures,” *Econometrica: journal of the econometric society*, 1984, pp. 761–766.

_ , _ , and _ , “The Foster–Greer–Thorbecke (FGT) poverty measures: 25 years later,” *The Journal of Economic Inequality*, 2010, *8*, 491–524.

Fox, Jean-Paul, *Bayesian item response modeling: Theory and applications*, Springer, 2010.

Fraker, Thomas M, Alberto P Martini, and James C Ohls, “The Effect of Food Stamp Cashout on Food Expenditures: An Assessment of The Findings from Four Demonstrations,” *Journal of human resources*, 1995, pp. 633–649.

Friedman, Milton, “A Theory of the consumption function,” 1957.

Gaviria, Alejandro and Steven Raphael, “School-based peer effects and juvenile behavior,” *Review of Economics and Statistics*, 2001, *83* (2), 257–268.

Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel, “Issuance and Incidence: SNAP Benefit Cycles and Grocery Prices,” *American Economic Journal: Economic Policy*, 2022, *14* (1), 152–178.

Golub, Andrew Lang and Bruce D Johnson, *Crack's decline: Some surprises across US cities*, US Department of Justice, Office of Justice Programs, National Institute of …, 1993.

Granados, José A Tapia, "Recessions and mortality in Spain, 1980–1997," *European Journal of Population/Revue européenne de Démographie*, 2005, *21* (4), 393–422.

Gregory, Christian A, "Are we underestimating food insecurity? Partial identification with a Bayesian 4-parameter IRT model," *Journal of Classification*, 2020, *37*, 632–655.

Gritz, R Mark, "The impact of training on the frequency and duration of employment," *Journal of econometrics*, 1993, *57* (1-3), 21–51.

Gunawan, David, Anastasios Panagiotelis, William Griffiths, and Duangkamon Chotikapanich, "Bayesian weighted inference from surveys," *Australian & New Zealand Journal of Statistics*, 2020, *62* (1), 71–94.

Gundersen, Craig, "Measuring the extent, depth, and severity of food insecurity: an application to American Indians in the USA," *Journal of Population Economics*, 2008, *21*, 191–215.

_ and James P Ziliak, "Food insecurity and health outcomes," *Health affairs*, 2015, *34* (11), 1830–1839.

_ , Brent Kreider, and John Pepper, "The economics of food insecurity in the United States," *Applied Economic Perspectives and Policy*, 2011, *33* (3), 281–303.

_ , Elaine Waxman, and Amy S Crumbaugh, "An examination of the adequacy of Supplemental Nutrition Assistance Program (SNAP) benefit levels: impacts on food insecurity," *Agricultural and Resource Economics Review*, 2019, *48* (3), 433–447.

Hall, Robert E, "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence," *Journal of political economy*, 1978, *86* (6), 971–987.

Hambleton, Ronald K, "Principles and selected applications of item response theory.," *Educational Measurement Third Edition.*, 1989.

Hamrick, Karen S and Margaret Andrews, "SNAP participants' eating patterns over the benefit month: a time use perspective," *PloS one*, 2016, *11* (7), e0158422.

Hansen, Benjamin, Joseph J Sabia, and Daniel I Rees, "Have cigarette taxes lost their bite? New estimates of the relationship between cigarette taxes and youth smoking," *American Journal of Health Economics*, 2017, *3* (1), 60–75.

Hastings, Justine and Ebonya Washington, "The first of the month effect: consumer behavior and store responses," *American economic Journal: economic policy*, 2010, *2* (2), 142–62.

_ and Jesse M Shapiro, "How are SNAP benefits spent? Evidence from a retail panel," *American Economic Review*, 2018, *108* (12), 3493–3540.

Heckman, James J, Richard Robb, and James R Walker, "Testing the mixture of exponentials hypothesis and estimating the mixing distribution by the method of moments," *Journal of the American Statistical Association*, 1990, *85* (410), 582–589.

Hill, Heather D, Pamela A Morris, Nina Castells, and Jessica Thornton Walker, "Getting a job is only half the battle: Maternal job loss and child classroom behavior in low-income families," *Journal of Policy Analysis and Management*, 2011, *30* (2), 310–333.

Hollingsworth, Alex, Christopher J Ruhm, and Kosali Simon, "Macroeconomic conditions and opioid abuse," *Journal of health economics*, 2017, *56*, 222–233.

Huizinga, David, Rolf Loeber, and Terence P Thornberry, *Recent findings from the program of research on the causes and correlates of delinquency*, The Office, 1995.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan, "Earnings losses of displaced workers," *The American economic review*, 1993, pp. 685–709.

Johnson, Valen E and James H Albert, *Ordinal data modeling*, Springer Science & Business Media, 2006.

Johnston, Lloyd, Richard Miech, Patrick O'Malley, Jerald Bachman, John Schulenberg, and Megan Patrick, "Monitoring the Future national survey results on drug use, 1975-2019: Overview, key findings on adolescent drug use," 2020.

Jolliffe, Dean, Craig Gundersen, Laura Tiehen, and Joshua Winicki, "Food Stamp Benefits and Child Poverty," *American Journal of Agricultural Economics*, 2005, *87* (3), 569–581.

Jr, Glen H Elder, Jacquelynne S Eccles, Monika Ardelt, and Sarah Lord, "Inner-city parents under economic pressure: Perspectives on the strategies of parenting," *Journal of Marriage and the Family*, 1995, pp. 771–784.

Kalil, Ariel and K Ziol-Guest, "The Great Recession and married parents' use of time," *Russell Sage Foundation. https://www. russellsage. org/sites/default/files/The% 20Great% 20Recession% 20and% 20Married% 20Parents% 20Use% 20of% 20Time% 20September*, 2013, *203*, 202013.

_ and Thomas DeLeire, "Parental Job Loss and Early Adolescent Adjustment in Black and White Families. JCPR Working Paper.," 2002.

Kharmats, Anna Y, Jessica C Jones-Smith, Yun Sang Cheah, Nadine Budd, Laura Flamm, Alison Cuccia, Yeeli Mui, Angela Trude, and Joel Gittelsohn, "Relation between the Supplemental Nutritional Assistance Program cycle and dietary quality in low-income African Americans in Baltimore, Maryland," *The American journal of clinical nutrition*, 2014, *99* (5), 1006–1014.

Kieftenbeld, Vincent and Prathiba Natesan, "Recovery of graded response model parameters: A comparison of marginal maximum likelihood and Markov chain Monte Carlo estimation," *Applied Psychological Measurement*, 2012, *36* (5), 399–419.

Kreider, Brent, John V Pepper, Craig Gundersen, and Dean Jolliffe, "Identifying the effects of SNAP (food stamps) on child health outcomes when participation is endogenous and misreported," *Journal of the American Statistical Association*, 2012, *107* (499), 958–975.

Kung, Eva M and Albert D Farrell, "The role of parents and peers in early adolescent substance use: An examination of mediating and moderating effects," *Journal of Child and Family Studies*, 2000, *9*, 509–528.

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

_ , "Life-Cycle Consumption and Hyperbolic Discount Functions," *European economic review*, 1998, *42* (3-5), 861–871.

Leftin, Joshua, Mia Monkovic, Francisco Wang, Nima Rahimi, Andrew Wen, and Alma Vigil, "Technical Documentation for the Fiscal Year 2022 Supplemental Nutrition Assistance Program Quality Control Database and the QC Minimodel," 2024.

Luechinger, Simon, Stephan Meier, and Alois Stutzer, "Why does unemployment hurt the employed?: Evidence from the life satisfaction gap between the public and the private sector," *Journal of Human Resources*, 2010, *45* (4), 998–1045.

Manikas, Ioannis, Beshir M Ali, and Balan Sundarakani, "A systematic literature review of indicators measuring food security," *Agriculture & food security*, 2023, *12* (1), 10.

Mastrobuoni, Giovanni and Matthew Weinberg, "Heterogeneity in intra-monthly consumption patterns, self-control, and savings at retirement," *American Economic Journal: Economic Policy*, 2009, *1* (2), 163–189.

McKee-Ryan, Frances, Zhaoli Song, Connie R Wanberg, and Angelo J Kinicki, "Psychological and physical well-being during unemployment: a meta-analytic study.," *Journal of applied psychology*, 2005, *90* (1), 53.

McLachlan, Geoffrey, "Finite mixture models," *A wiley-interscience publication*, 2000.

McLoyd, Vonnie C, "Socioeconomic disadvantage and child development.," *American psychologist*, 1998, *53* (2), 185.

Moffitt, Robert A and David C Ribar, "Rasch analyses of very low food security among households and children in the three city study," *Southern Economic Journal*, 2016, *82* (4), 1123–1146.

Morris, Pamela A and Lisa A Gennetian, "Identifying the effects of income on children's development using experimental data," *Journal of Marriage and Family*, 2003, *65* (3), 716–729.

Mundlak, Yair, "On The Pooling of Time Series and Cross Section Data," *Econometrica: journal of the Econometric Society*, 1978, pp. 69–85.

Muthén, Bengt, “A structural probit model with latent variables,” *Journal of the American Statistical Association*, 1979, *74* (368), 807–811.

National Academies of Sciences, Engineering and Medicine, “A Consumer Food Data System For 2030 and Beyond,” 2020.

Needle, Richard, Hamilton McCubbin, Marc Wilson, Robert Reineck, Amnon Lazar, and Helen Mederer, “Interpersonal influences in adolescent drug use—the role of older siblings, parents, and peers,” *International Journal of the Addictions*, 1986, *21* (7), 739–766.

Neumayer, Eric, “Recessions lower (some) mortality rates:: evidence from Germany,” *Social science & medicine*, 2004, *58* (6), 1037–1047.

Nord, Mark, “Assessing Potential Technical Enhancements to the U.S. Household Food Security Measure,” Technical Report, United States Department of Agriculture 2012.

Oetting, Eugene R and Fred Beauvais, “Peer cluster theory: Drugs and the adolescent,” *Journal of Counseling & Development*, 1986, *65* (1), 17–22.

Ohls, James, Larry Radbill, and Allen Schirm, “Household Food Security in the United States, 1995-1997: Technical Issues and Statistical Report,” Technical Report, Mathematica Policy Research 2001.

Opsomer, Jean D, Helen H Jensen, Sarah M Nusser, Dorin Drignei, and Yasuo Amemiya, “Statistical considerations for the USDA food insecurity index,” 2002.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz, “The short-and long-term career effects of graduating in a recession,” *American Economic Journal: Applied Economics*, 2012, *4* (1), 1–29.

Pabilonia, Sabrina Wulff, “The Effects of the Great Recession on Teenagers’ Risky Health Behaviors and Time Use,” Technical Report, IZA Discussion Papers 2014.

Pacula, Rosalie Liccardo, “Substance use and recessions: What can be learned from economic analyses of alcohol?,” *International Journal of Drug Policy*, 2011, *22* (5), 326–334.

_ , Michael Grossman, Frank J Chaloupka, Patrick M O’Malley, Lloyd D Johnston, and Matthew C Farrelly, “Marijuana and youth,” in “Risky behavior among youths: An economic analysis,” University of Chicago Press, 2001, pp. 271–326.

Pesko, Michael F, Jenna M Hughes, and Fatima S Faisal, “The influence of electronic cigarette age purchasing restrictions on adolescent tobacco and marijuana use,” *Preventive medicine*, 2016, *87*, 207–212.

Rabbitt, Matthew P, “Measuring the Effect of Supplemental Nutrition Assistance Program Participation on Food Insecurity Using a Behavioral Rasch Selection Model,” *Unpublished paper, University of North Carolina at Greensboro*, 2014.

— and Alisha Coleman-Jensen, "Rasch analyses of the standardized Spanish translation of the U.S. household food security survey module," *Journal of Economic and Social Measurement*, 2017, *42* (2), 171–187.

— , Laura J Hales, Michael P Burke, and Alisha Coleman-Jensen, "Household food security in the United States in 2022," 2023.

— , M Reed-Jones, Laura J Hales, and Michael P Burke, "Household food security in the United States in 2023," 2024.

Rasch, Georg, *Probabilistic models for some intelligence and attainment tests*, Danmarks pædagogiske Institut, 1960.

Rege, Mari, Kjetil Telle, and Mark Votruba, "Parental job loss and children's school performance," *The Review of Economic Studies*, 2011, *78* (4), 1462–1489.

Ruhm, Christopher J, "Economic conditions and alcohol problems," *Journal of health economics*, 1995, *14* (5), 583–603.

— , "Healthy living in hard times," *Journal of health economics*, 2005, *24* (2), 341–363.

— and AM Jones, "Understanding the relationship between macroeconomic conditions and health," *The Elgar Companion to Health Economics, Second Edition, Cheltenham, Edward Elgar*, 2012, pp. 5–14.

— and William E Black, "Does drinking really decrease in bad times?," *Journal of health economics*, 2002, *21* (4), 659–678.

Saffer, Henry and Frank Chaloupka, "The demand for illicit drugs," *Economic inquiry*, 1999, *37* (3), 401–411.

Samejima, Fumi, "Estimation of latent ability using a response pattern of graded scores," *ETS Research Report Series*, 1968, *1968* (1).

Scull, Tracy M, Janis B Kupersmidt, Alison E Parker, Kristen C Elmore, and Jessica W Benson, "Adolescents' media-related cognitions and substance use in the context of parental and peer influences," *Journal of youth and adolescence*, 2010, *39*, 981–998.

Seligman, Hilary K, Ronli Levi, Victoria O Adebiyi, Alisha Coleman-Jensen, Joanne F Guthrie, and Edward A Frongillo, "Assessing and monitoring nutrition security to promote healthy dietary intake and outcomes in the United States," *Annual Review of Nutrition*, 2023, *43*, 409–429.

Shapiro, Jesse M, "Is there a daily discount rate? Evidence from the food stamp nutrition cycle," *Journal of public Economics*, 2005, *89* (2-3), 303–325.

Si, Yajuan, Rob Trangucci, Jonah Sol Gabry, and Andrew Gelman, "Bayesian hierarchical weighting adjustment and survey inference," *Survey Methodology*, 2020, p. 181.

Simons-Morton, Bruce, Denise L Haynie, Aria D Crump, Patricia Eitel, and Keith E Saylor, "Peer and parent influences on smoking and drinking among early adolescents," *Health education & behavior*, 2001, *28* (1), 95–107.

Smith, Travis A. and Christian A. Gregory, "Food Insecurity in the United States: Measurement, Economic Modeling, and Food Assistance Effectiveness," *Annual Review of Resource Economics*, 2023, *15* (Volume 15, 2023), 279–303. Publisher: Annual Reviews Type: Journal Article.

Smith, Travis A, Joshua P Berning, Xiaosi Yang, Gregory Colson, and Jeffrey H Dorfman, "The effects of benefit timing and income fungibility on food purchasing decisions among supplemental nutrition assistance program households," *American Journal of Agricultural Economics*, 2016, *98* (2), 564–580.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge, "What are we weighting for?," *Journal of Human resources*, 2015, *50* (2), 301–316.

Stevens, Ann Huff, "Persistent effects of job displacement: The importance of multiple job losses," *Journal of Labor Economics*, 1997, *15* (1, Part 1), 165–188.

_ and Jessamyn Schaller, "Short-run effects of parental job loss on children's academic achievement," *Economics of Education Review*, 2011, *30* (2), 289–299.

Substance, Abuse, Health Services Administration US Mental, Office of the Surgeon General (US et al., "Facing addiction in America: The surgeon general's report on alcohol, drugs, and health," 2016.

Tanaka, Victoria T, G Engelhard, and Matthew P Rabbitt, "Examining differential item functioning in the Household Food Insecurity Scale: Does participation in SNAP affect measurement invariance," *Journal of Applied Measurement*, 2019, *20* (1), 100–111.

_ , George Engelhard, and Matthew P Rabbitt, "Modeling Household Food Insecurity with a Polytomous Rasch Model," in "Quantitative Psychology: 84th Annual Meeting of the Psychometric Society, Santiago, Chile, 2019 84" Springer 2020, pp. 319–331.

Tarasuk, Valerie, Lynn McIntyre, and Jinguang Li, "Low-income women's dietary intakes are sensitive to the depletion of household resources in one month," *The Journal of nutrition*, 2007, *137* (8), 1980–1987.

Tekin, Erdal, Chandler McClellan, and Karen Jean Minyard, "Health and health behaviors during the worst of times," Technical Report, National Bureau of Economic Research 2013.

Tiehen, Laura, Constance Newman, and John A Kirlin, "The Food-Spending Patterns of Households Participating in The Supplemental Nutrition Assistance Program: Findings from USDA's FoodAPS," 2017.

Todd, Jessica E, "Revisiting the Supplemental Nutrition Assistance Program cycle of food intake: Investigating heterogeneity, diet quality, and a large boost in benefit amounts," *Applied Economic Perspectives and Policy*, 2015, *37* (3), 437–458.

United States Department of Agriculture,, "Do Supplemental Nutrition Assistance Program Benefits Expire," 2024.

United States Department of Agriculture, (2022), Food and Nutrition Service, Office of Policy Support, "Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2022," 2022.

U.S. Department of Agriculture, "Thrifty Food Plan, 2021," 2021.

Vink, Jacqueline M, Gonneke Willemsen, and Dorret I Boomsma, "The association of current smoking behavior with the smoking behavior of parents, siblings, friends and spouses," *Addiction*, 2003, *98* (7), 923–931.

Wang, Peiming, lain M Cockburn, and Martin L Puterman, "Analysis of patent data—a mixed-Poisson-regression-model approach," *Journal of Business & Economic Statistics*, 1998, *16* (1), 27–41.

Wedel, Michel, Wayne S DeSarbo, Jan Roelf Bult, and Venkatram Ramaswamy, "A latent class Poisson regression model for heterogeneous count data," *Journal of applied Econometrics*, 1993, *8* (4), 397–411.

Weinstein, James L, Katie S Martin, and Ann M Ferris, "Household food security varies within month and is related to childhood anemia," *Journal of Hunger & Environmental Nutrition*, 2009, *4* (1), 48–61.

Wilde, Parke E and Christine K Ranney, "The monthly food stamp cycle: shopping frequency and food intake decisions in an endogenous switching regression framework," *American Journal of Agricultural Economics*, 2000, *82* (1), 200–213.

Windle, Michael, "Parental, sibling, and peer influences on adolescent substance use and alcohol problems," *Applied Developmental Science*, 2000, *4* (2), 98–110.

Wooldridge, Jeffrey M, *Econometric Analysis of Cross Section and Panel Data*, MIT press, 2010.

World Bank, "What is Food Security?," https://www.worldbank.org/en/topic/agriculture/brief/food-security-update/what-is-food-security 2024. Online; accessed 25 November 2024.

Xu, Xin, "The business cycle and health behaviors," *Social Science & Medicine*, 2013, *77*, 126–136.

Yaktine, Ann L and Julie A Caswell, "Supplemental nutrition assistance program: examining the evidence to define benefit adequacy," 2013.

Yang, Yanliang, George C. Davis, and Wen You, "Measuring Food Expenditure Poverty in SNAP Populations: Some Extensions with an Application to the American Recovery and Reinvestment Act," *Applied Economic Perspectives and Policy*, 2019, *41* (1), 133–152.

Zhu, Xiaowen and Clement A Stone, "Assessing fit of unidimensional graded response models using Bayesian methods," *Journal of Educational Measurement*, 2011, *48* (1), 81–97.

Ziliak, James, “Modernizing SNAP benefits,” *The Hamilton Project Policy Proposal*, 2016, *6.*

Vita

José Xilau was born in Benguela, Angola on October 18th, 1992. José received his M.A. and Ph.D. in Economics from the Andrew Young School of Policy Studies of Georgia State University in 2022 and 2025, respectively. His primary field was health economics and his secondary field was labor economics. During his time at Georgia State University, José was a recipient of the Andrew Young School Dean's Fellowship, E.D. (Jack) Dunn Fellowship and Research Fund and George Malanos Economics Scholarship. He regularly attended and presented at multiple conferences and externally funded his research through a cooperative agreement with the USDA Economic Research Service. Before joining Georgia State, José earned a bachelor's and master's degree in economics from Boston University.