"Oh, SNAP!" The Impact of Nutritional Assistance on Grade Progression Rates for K-12 Students

Author: Liam Dietrich
Advisor: Prof. Paul Cichello

Presented to the Department of Economics
in partial fulfillment of the requirements
for a Bachelor of Arts Degree with Honors

Abstract

The COVID-19 pandemic impacted nearly every aspect of societies around the world when it struck in 2020. Food insecurity increased in almost all countries, even those with well-developed economies and safety nets, and education for all ages was drastically affected by social distancing guidelines and a concern for the safety of students and faculty alike. Using data from the 2019 and 2020 cross-sections of the Survey of Income and Program Participation, I evaluated the impact of an automatic increase in food assistance benefits in the United States during March 2020 on grade progression rates for students who were affected by the benefit increase. I find statistically insignificant results regarding the effect of this increase on grade progression rates, but a significant positive effect of being in the post-Covid period on these rates. These results could reflect the fact that school districts around the country broadly loosened the academic requirements for grade progression. Future research could evaluate the effect of food stamp benefits on the quality of education, and seek to overcome the limitations of the model used for this analysis.

Acknowledgements

I would first like to thank my advisor, Professor Paul Cichello. Without your help, I truly do not know how I would have completed such a monumental task as writing this thesis. Your classes have given me a strong fundamental understanding of statistics and econometrics that I know I will put to extensive use in my future, and I truly appreciate your patience and help throughout my time at Boston College as both a teacher and a mentor. I would also like to thank Professors Tracy Regan, Matthew Rutledge, and Geoffrey Sanzenbacher. Each of you have been instrumental in my growth as an economics major and a student at BC. Finally, I would like to thank my friends and family; in particular, my girlfriend, Elizabeth Mortati, who was an amazing source of support and encouragement throughout the process of writing this thesis, and my roommates, who put up with me ranting about random data issues or econometrics topics for the better part of our last two semesters here. With all of you behind me, I have had a formational experience at BC that has helped me become the person I am today.

I. Introduction

Food insecurity—a period when the availability of nutritional, safe food and socially acceptable means by which to acquire it are uncertain (Anderson 1990)—is a pervasive issue in developed and developing countries worldwide. In 2019, nearly two billion people were estimated to have experienced food insecurity globally, according to the United Nations State of Food Security and Nutrition in the World (SOFI) Report, with nearly 750 million of these people experiencing severe food insecurity. Upon the onset of the COVID-19 pandemic, though, global food insecurity increased dramatically—in 2021, nearly 2.4 billion people were food insecure, and as many as 820 million people were severely so.

The initial impacts of COVID-19 were far from limited to worsened food security, though. Education at all levels and for all ages also suffered a major blow, particularly in the early stages of the pandemic, when social distancing was generally regarded as the most effective method by which to slow the spread of disease. School closures were widespread, but differentially implemented, for inconsistent periods of time, and with varying degrees of effectiveness. Without in-person schooling, children and education faculty were faced with unprecedented challenges, leading to worsened mental health for both groups (Jalongo 2021). Children regressed in reading and math skills, undoing decades of progress in each area. Low-income students and students belonging to already-vulnerable populations found these problems exacerbated. Despite all of these negative impacts on education quality from COVID-19, individual districts were afforded greater freedom in determining students' eligibility for grade progression (advancing to the next grade) and retention (remaining in the same grade for an additional year), with all states receiving federal waivers for standardized testing requirements in the 2019-2020 academic year (Green 2020).

Previous research on the topics of food insecurity, nutritional assistance, government programs, and educational outcomes for primary and secondary school students is easy to find, but few pieces link each of these topics to each other, and fewer do so against the backdrop of the COVID-19 pandemic. The literature supports the notion that food insecurity is associated with worsened educational outcomes for students, but lacks clarity on whether food security is associated with improved educational outcomes (Shankar et al. 2017, Argaw et al. 2023). There is also research examining educational losses for children during the pandemic, but analysis of government intervention in the realm of education mostly focuses on education-specific policies enacted to stem these losses (Duflo 2001, Angrist & Krueger 1991).

My project characterizes the effect of a uniform increase in Supplemental Nutrition Assistance Program (SNAP) payments on grade progression rates for K-12 students receiving the benefits. To do so, I use data from two months of the Survey of Income and Program Participation (SIPP), a longitudinal, nationally representative dataset of the non-institutionalized United States civilian population. I find statistically insignificant evidence of any effect of this increase, but a statistically significant increase in progression rates after the onset of COVID-19 for all students. It is likely that the lack of impact from this increase is due to a general loosening or complete elimination of standardized benchmarks to determine a student's eligibility for advancing to the next grade; with higher across-the-board progression, it is unsurprising that grade progression is unaffected by most factors besides time.

This paper contributes to the existing literature in two primary ways: (1) by quantifying the impact of a program predominantly designed to address food insecurity on grade progression, described in this paper as a metric for educational development, and (2) by analyzing the effectiveness of an increase in transfer payments in a period of widespread upheaval, such as COVID-19. Current research has underexplored the effect of transfer programs on areas to which they are not directly related despite evidence of an intersection between the underlying problems these programs address and other areas of recipients' lives. With regards to SNAP, few studies have examined the effectiveness of the increase in benefits on alleviating pandemic-related instability, and those that have look primarily at its effect on food insecurity.

The remainder of this paper proceeds as follows. First, I present background information on food insecurity and K-12 education in the United States, the particular impacts of COVID-19 in these areas, and the specific policy change to be evaluated. Then, a review of existing literature provides context for this study and explores the research that has already been conducted in these areas. I describe the data source and outline the methodology used to analyze it. I then present my findings from this analysis and interpret the results. I conclude with a discussion of limitations of my project and an evaluation of potential areas of focus for future researchers.

II. Background

Food Insecurity & Assistance in the United States

Despite the United States' status as a developed country with the highest GDP in the world, its citizens are far from immune to the upward trend in food insecurity, with 33.8 million people experiencing food insecurity in 2021. To combat food insecurity in the U.S., the United States Department of Agriculture (USDA) administers the Supplemental Nutrition Assistance Program (SNAP), formerly known as the Food Stamp Program, to eligible citizens. SNAP is the largest source of nutrition assistance nationwide, serving 46 million Americans in fiscal year 2011 with a cost of $75 billion (Nestle 2019). While initially requiring beneficiaries to pay for a portion of the benefits themselves, the program has evolved into a means-tested supplemental income measure that is free for eligible households. As SNAP is a federal assistance program, all benefits are paid for by the federal government, with administrative costs also being partly shared by individual states. SNAP's federal nature also guarantees uniform eligibility requirements for all states except Alaska and Hawaii (which have slightly different income levels for eligibility). To be eligible for SNAP, an individual must pass two income and two asset tests, having: (1) gross income of less than 135% of the federal poverty level; (2) net income of less than 100% of the federal poverty level; (3) liquid asset worth of less than $2,000; (4) vehicle asset worth of less than $4,250. Individuals can also be categorically eligible for SNAP by receiving benefits from certain other means-tested government assistance programs.

In an attempt to help its most vulnerable citizens weather the pandemic, the U.S. government immediately expanded access to SNAP upon the onset of COVID-19 in March 2020 through the Coronavirus Aid, Relief, and Economic Security (CARES) Act, allowing eligible households to receive their maximum possible benefit from the program. Since many people—particularly low-income people—experienced significant income and/or employment shocks during the pandemic, this policy was intended to supplement people's income to help them afford food. The Trump administration left out the most vulnerable people from its expansion of benefits—if a household was already receiving the maximum amount of food assistance, they experienced no change in benefits. When President Biden assumed office, though, his administration changed this policy and began extending federal assistance to these people, granting all recipients at least $95 in benefits as well as a 15% benefit increase across the board. A timeline of the benefits received by three typical SNAP recipients—one who received the maximum benefit prior to the pandemic, one whose income was 100% of the Federal Poverty Level (FPL), and one whose income was 50% of the FPL—is included in the appendix (Figure D1). The timeline assumes that there were no income fluctuations over the time period depicted.

State Responses to Educational Losses During COVID-19

In March of 2020, schools nationwide began closing due to a lack of methods besides social distancing by which to slow the spread of COVID-19 in public spaces. Without clear federal guidance on the issue of school closures, states' policies in this regard were not identical: many states instituted mandatory school closures, while others only recommended that districts close schools, and left the ultimate decision up to individual districts (Slavin 2020). After the initial two-three week period of closures, it became apparent that longer closures would be necessary, and states were thus faced with the difficult decision of how to best service their students in light of drastically different circumstances than in the past. Most states agreed in many areas: a need to quickly pivot to remote learning, the existence of equity issues in doing so, and waiving standardized assessments for the initial year of the pandemic, for instance. States diverged, however, in many aspects, including decisions about graduation and grading policies (Reich et al. 2020).

A key point of contention for school districts was how to equitably evaluate their students' success in the initial phase of the pandemic. Since nearly all states were granted federal waivers of standardized testing requirements, the question of whether or not to allow struggling students to progress to the next grade became particularly difficult to answer. Without the usual benchmark of these assessments to help administrators understand a student's preparedness for the next grade, policies regarding grade promotion (progressing a student to the next grade/graduation) and retention (holding a student back to repeat a grade) varied widely at the state and district levels. In some states, no guidance on grade promotion or retention was given; in many others, states explicitly gave districts permission to adjust their policies in these regards as they saw fit. The result in all states was a highly individualized approach to educational policies at the district level.

As an example of the freedom afforded to each district, consider Michigan, where Gov. Gretchen Whitmer almost immediately passed the comprehensive Executive Order 2020-35 in response to COVID-19. The order gave school districts large amounts of discretion when choosing how to grant credit for graduation, including the abilities to grant full semester credit for coursework up through March 11, 2020, allow students to graduate with alternative assessment methods (e.g. a resumé evaluation), or provide an optional final exam/interdisciplinary learning approach to gauge a students' readiness for graduation (Whitmer 2020). The order also mandated that schools would close for in-person education until the end of the 2020 academic year, and that districts were to create and administer a plan for the virtual continuation of education, grade promotion, and grade retention within 4 weeks of the order's passing.

III. Literature Review

The following section discusses existing literature on the topics of food insecurity, SNAP benefits, educational attainment, and the effect of COVID-19 on each of these areas.

Food Insecurity: Measures & Impacts

Many researchers have sought to characterize the prevalence of food insecurity and the extent to which it impacts various outcomes for affected people. For the U.S., a standard measure of food insecurity came with the addition of the U.S. Household Food Security Survey Module (HFSSM) to the U.S. Census Bureau's Current Population Survey (CPS) in the early 1990s. The HFSSM includes a set of 18 questions (10 for households with no children) that assess the level of food insecurity and hunger a household experiences in a given year. There are four broad categories into which households with children can be qualified: *food secure* (fewer than three affirmative answers), *food insecure without hunger* (between three and seven affirmative responses), *moderately food insecure with hunger* (between seven and twelve positive answers) and *severely food insecure with hunger* (more than 12 affirmative positive responses) (Nord et al. 1999). Another category, *food insecure with hunger among children*, reflects positive responses to five of the eight questions specifically concerning children. While not all surveys quantify food insecurity with the 18 questions used in the CPS, most nutrition-related surveys include some measure of food insecurity, classified on a similar scale of low, moderate, and severe levels.

A significant body of research focuses on the causes of food insecurity for once-food-secure households. Heflin (2016) aggregates these causes broadly into what she terms *instability*, consisting of "employment shocks, household formation shocks, residential changes, income changes, household size changes, and disability shocks." She finds that increases in instability have significant effects on various measures of material hardship; however, decreases in instability generally do not, implying that efforts to increase benefits to disadvantaged populations may be impactful only in the sense that they prevent instability from worsening. It is also important to note that food insecurity does not afflict all households at equal rates. Rather, researchers observe differential severity and prevalence of food insecurity among many demographics (Alaimo 2005). Households with children, female-headed households, households with seniors, and Hispanic/Black households were more likely to experience food insecurity than the applicable comparison group for each category. It is likely that this phenomenon occurs because instability, as measured by Heflin, is more common in these groups than it is for the average American household.

Food insecurity is linked to a host of negative outcomes for people of all ages. Examining recent research in this field, Gundersen and Ziliak (2015) show that after controlling for other health-related risk factors, health outcomes for food-insecure people are drastically worse than their food-secure counterparts. Food insecurity is associated with worsened mental, physical, and oral health for nonsenior adults; for children and seniors, these effects are magnified. For instance, food-insecure children are over twice as likely to consider themselves in fair or poor health and between 1.4 and 2.6 times as likely to suffer from asthma than food-secure children, and food-insecure seniors experience livelihood deterioration comparable to food-secure seniors that are 14 years older. In all children—from infancy to adolescence—food insecurity is associated with worse behavioral outcomes, emotional development, and academic success (Shankar et al. 2017).

Economics of Education, Grade Retention, and Grade Advancement

It is this final metric, academic success, with which this paper is concerned. Educational outcomes for children are of particular importance, as these outcomes tend to have a snowball effect: children who fall behind early on in their academic careers tend to have worse educational outcomes than those who keep pace with their peers during this time (Choi et al. 2018). Education in general has a clear-cut association with future success in employment, higher returns on wage, and better health (Blanden et al. 2022); as a result, an indirect effect of food insecurity on children is to diminish the likelihood that they will experience better outcomes in these areas later in life due to an adverse effect on their education.

Duflo (2001) provides evidence of one of the aforementioned effects of increases in education. Duflo uses an exogenous school-building initiative in Indonesia to determine the real returns to primary education on wages. She finds that each school built per 1000 children leads to an average increase of 0.12-0.19 years of education and 1.5-2.7 percent increase in wages, implying economic returns to education ranging from 6.8-10.6 percent. While the magnitude of these results are specific to Indonesia, which experienced rapid GDP growth at the same time as its school-building program, the general relationship remains externally valid—increased access to primary school education has a positive effect on real wages later in life. Angrist and Krueger (1991) take advantage of compulsory schooling laws in the United States and instrument for increases in education with quarter of birth, as the interaction of these laws with birthdates creates exogenous variation in education. They find that increased education has a causal impact on earnings for those who are required to attend an additional year of secondary school, with nearly a 7.5 percent increase in later earnings for each additional year of education. Their research also supported the notion of "negative selection" in education, or the idea that students who struggle in school may actually gain more from additional education than those who already have the ability to succeed in school.

It is because of these relationships that in many countries, students who starkly underperform relative to their peers in school may be faced with involuntary (retention) or voluntary grade repetition. It is estimated that roughly 5 percent of students in grades 1-3 in the U.S. face grade retention, and between 7 and 15 percent of all K-12 students have been retained at some point in their academic career (Tingle et al. 2012). Yet the consensus from independent researchers and educational non-profit organizations is clear: grade retention has strongly negative consequences for students' future success. Grade retention is rated by sixth graders as the most stressful life event they could experience (compared to death of a parent and going blind, among others), and is linked to a host of motivational and mental health issues that have serious implications on children's psychosocial development (Anderson et al. 2005). Retention also has no effect on the social behaviors it is supposed to correct (misbehaving, acting out, social withdrawal), and instead leads to reported decreases in self-esteem and engagement in relationships of all kinds. Academically, retained students have worse academic outcomes and are less likely to complete high school than comparable students who are promoted (Mathys 2017).

An opposite approach to helping struggling students learn material is to advance them to the next grade regardless of their success in their current grade. This practice of *social promotion* initially gained some favor as a reaction to the inadequacy of grade retention in addressing students' needs. However, due to backlash against social promotion, and the public perception that students who were wrongly promoted would struggle in future grades (Lynch 2014), the federal government encouraged educators to find alternative methods to help struggling students. With the Clinton administration explicitly calling for an end to the practice in the 1990s, and the passage of the No Child Left Behind (NCLB) Act in 2001, social promotion fell out of favor as a remedial technique; instead, despite its known shortcomings, grade retention remains a more common method by which teachers hope to improve struggling students' academic abilities (Peterson & Hughes 2011).

The Impact of COVID-19

When the COVID-19 pandemic struck the United States in early 2020, a critical concern of policymakers was the pandemic's impact on food security for the millions of people who suddenly found themselves unemployed amidst supply chain disruptions that raised the prices of everyday necessities. As income and employment shocks are key drivers of a fall into food insecurity, it is unsurprising that during the pandemic, there were large year-over-year increases in food insecurity at much higher rates globally than in the years prior to COVID-19. In the U.S., increases in food insecurity strike various groups differently; for example, worsened food access at the beginning of the pandemic was partly offset by mobile grocery ordering, but this problem remained for people without access to reliable internet. Black, Brown, and Hispanic people were more susceptible to pandemic-induced food insecurity in the U.S., due to higher incidence of COVID-19 among these groups, overall lower incomes, and pre-existing inequities in food access (O'Hara and Toussaint 2021). In general, poorer people—but especially women, children, migrants—experienced worsened nutrition and food security as a result of the pandemic than their richer counterparts (Swinnen and McDermott 2020).

Education systems and outcomes were also dramatically impacted by the shift to online learning necessitated by COVID-19. While this may have led to better health outcomes, at least during the early stages of the pandemic, a significant body of research has emerged documenting the detriments of online learning to educational attainment, which were both far-reaching and severe. At the onset of the pandemic, educators were concerned about academic, social, and financial effects for children who suddenly found their learning experience drastically disrupted; unfortunately, many of these predictions have come true as more data surrounding the impact of the pandemic has been released. For students in primary and secondary school, COVID-19 was associated with worse reading and math scores (the 2022 release of *The Nation's Report Card* in the U.S. provides particularly shocking information about the scope of academic regression during the pandemic), higher rates of mental illness, lower social engagement, and disconnection from peers (Jalongo 2021). Low-income students faced a unique set of challenges with regards to education during the pandemic. Online learning necessitates access to reliable internet and a source of both daytime childcare and home schooling assistance; each of these is less certain to be available to low-income families than it is for comparatively high-income families. For lower-income families who participate in free and reduced-price lunch programs, online learning also represents an additional loss of food security through a loss of in-person education (Ambrose 2020).

The U.S. and state governments sought to address food insecurity at the onset of the pandemic, rightly fearing major decreases in food availability and affordability for houses impacted by COVID-19. The primary method by which the federal government sought to alleviate food insecurity was through an emergency allotment of SNAP benefits to cover the difference between the amount a household received and the maximum possible benefit for a household of that size. In 2021, this change was extended to include a temporary 15% increase in benefits for all beneficiaries. Based on national estimates of food insecurity, which stayed relatively stagnant throughout the pandemic annually after sharp increases in the months immediately following widespread quarantines in 2020, it is apparent that this program was successful in alleviating increases in food insecurity felt by countries around the world. Using a Bayesian structural time series analysis to construct counterfactual estimates of food insecurity for people in the United States, Bryant and Follett (2022) find that the 15% increase in SNAP benefits prevented 850,000 cases of food insecurity each week during the pandemic after it was enacted, providing convincing evidence of the program's success. In this model, they assume that (1) the additional benefits did not affect the underlying food insecurity of recipients, and (2) their model structure is unchanged by the introduction of these benefits. The authors also perform two robustness checks: first, they use a local dataset detailing food pantry visits for the Des Moines network of pantries and find that overall numbers of visits decreased by a comparable amount during the same time period of their first analysis; second, they perform this same analysis with random dates pre-2019 and find that 20 out of 24 analyses exhibit null results, indicating that this effect was due to SNAP interventions.

Contributions to Existing Literature

Despite the proven link between food insecurity and numerous other outcomes, including educational attainment and quality, current literature does not focus on the specific impact that SNAP benefits may have on these outcomes. This paper seeks to begin filling that gap and to highlight the fact that these benefits may have indirect effects in other areas. Specifically, I examine the impact of an increase in SNAP benefits on a previously unexplored outcome, progression to the next grade. Even though grade progression rates increased dramatically during the pandemic, by controlling for being in the post-Covid period and other variables which are known to impact a person's educational success, I characterize a separate effect of the benefit increase on grade progression. My paper also examines the efficacy of SNAP benefits in the early stages of the COVID-19 pandemic, a time when the benefits were most needed; my results are thus relevant to policymakers considering similar increases in benefits at the beginning of a large-scale crisis.

IV. Data

Introduction to the SIPP

This paper analyzes data from the 2018-2021 Survey of Income and Program Participation (SIPP). The SIPP is a longitudinal, nationally representative survey of the non-institutionalized civilian population administered by the U.S. Census Bureau, meant to cover individuals of all socioeconomic backgrounds in order to analyze the impact of various government programs, including SNAP. It contains comprehensive information on demographic characteristics, employment, education, assets and liabilities, health and well-being, and program participation for individuals in each panel. Each panel covers participants in four one-year interview waves, with reference periods that vary from time of interview to participants' entire lives, depending on the question. Interviews for the 2018 panel began in 2018 and concluded in 2021, allowing for analysis of short-term effects of the COVID-19 pandemic.

The structure of the SIPP allows for linkage between individuals into families and households, both at the monthly level and at the time of interview. Interviewers begin the survey by identifying a household reference person, who answers basic demographic questions about members of the household, including a detailed relationship between each household member, and provides answers to household-level questions. These responses are further divided into two categories where relevant: one which includes "Type 2 people," who live in the residence but do not belong to the family residing in the household, and are therefore not directly interviewed; and one which does not include Type 2 people. Since certain transfer programs allow benefits to be extended to these people, inclusion of Type 2 people can impact analysis of these programs; as such, income and poverty variables which include Type 2 people are used in this study.

Sampling Design & Weighting

The SIPP begins sampling by assigning at least one contiguous county to a primary sampling unit (PSU). If a single county does not have a large enough population (7,500 residents) to be considered its own PSU, it is combined with one or more adjacent counties until it can be. PSUs with over 100,000 housing units are automatically included in the SIPP; smaller PSUs are stratified based on certain poverty measures, then chosen with a probability proportional to their size. Within PSUs, addresses are separated into two strata with low and high concentrations of low-income households, and the higher-concentration stratum is oversampled. In this way, the SIPP attempts to include enough low-income citizens in the sample to allow for causal inference regarding these people, whereas a true random sample might not.

The SIPP provides weights to address the sampling structure and to allow information from multiple panels to be combined in the same reference period. The final person weight is made up of a base weight (to account for the probability of being selected for a sample unit), a subsampling adjustment, a mover adjustment (for people who move to different households in Waves 2+), a nonresponse adjustment, a cross-section adjustment (when combining data from multiple panels into one period), and an adjustment to account for known differences from population values. Since data from two distinct months is included, different weights are applied for each month to create an overall representative sample.

To generate descriptive statistics and regression results with correct point estimates and standard errors, the technique of balanced repeated replication (BRR) is employed. Like other replication methods, BRR randomly creates subsamples of the overall sample, computes parameter estimates within each sample, and uses the variability between subsamples to calculate an unbiased sampling variance estimate for the statistic (Zinn 2016). Specifically, in each replication, one of the two PSUs per stratum is removed, and the sampling weight of the remaining PSU is doubled. Then, the subsample is constructed using the remaining PSUs from each stratum. Each replication is weighted to be representative of the full sample, so that individual replications are representative samples for the same population. Another variation of the BRR approach is Fay's method, in which replicates are weighted with a less extreme value so as to create more stable and precise variance estimates (Judkins 1990). In the SIPP, the included PSUs are weighted at an additional 50% of their original weight, and the excluded PSUs are weighted at 50% less than their original weight (as opposed to doubling/completely removing the included/excluded PSUs, respectively) (U.S. Census Bureau).

Key Variables & Data Editing

The primary dependent variable for this study is an indicator variable for whether a person progressed to the next grade in the given year. To construct this variable, I first forward-filled enrollment data for missing months where appropriate—if a student was unenrolled for a small period of time, it was typically due to them being on summer vacation or in a different turnover between grades. For example, a 4th grader on summer vacation from June through August would initially be coded to have missing enrollment data in those months; after my adjustment, they were coded to be in 4th grade until August. A student was considered to have progressed to the next grade in a year if their grade level at the beginning of the year was lower than their grade level at the end of the year.

Another relevant variable for this study is a recoded food insecurity score. During the SIPP interviews, respondents are asked six questions related to food insecurity: (1) "The food you bought did not last?"; (2) "Could not afford balanced meals?"; (3) "In [reference year], did you ever cut the size of your meals or skip meals because there wasn't enough money for food?"; (4) "How often did [respondent] cut the size of his/her meals?"; (5) "In [reference year], did you ever eat less than you felt you should because there wasn't enough money to buy food?"; and (6) "In [reference year], were you ever hungry but didn't eat because there wasn't enough money for food?" Based on the number of "yes" answers to this question, households are classified as (1) High/marginal food secure (0-2 affirmative responses), (2) low food secure (3-4 affirmative responses), or (3) very low food secure (5+ affirmative responses). For the purposes of this analysis, a household was considered food insecure if it fell into either of the two latter categories.

The SIPP provides extensive information on program participation in many government transfer programs, including SNAP, Supplemental Security Income (SSI), General Assistance (GA), Temporary Assistance for Needy Families (TANF), and the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC), among others. For each of these programs, coverage, ownership, and dollar amount of benefits received by respondents is recorded. Within a spell (continuous period of time) of coverage, all information is constant, except for the value of benefits received, which is collected at the monthly level. This allows for comparison of benefit levels at different points in time and across multiple spells of coverage, if multiple spells exist.

Other variables include household demographics, such as race & sex of household reference person & respondents, household size, metropolitan status, primary language spoken at home, whether a member of the household served in the Armed Forces, number of children in the household, and number of seniors in the household. Income dynamics, such as total earned income from all members of the household older than 15 and household net worth, are tracked at the monthly level. A self-reported health outcome variable is also available for all respondents. Summary statistics for each relevant variable are presented in the following subsection.

It is important to note that certain assumptions were made in the process of creating a final dataset for my analysis, some of which were due to measurement error in the data. Education was topcoded at 22 years for Ph.D. & medical school graduates, and was considered to be zero years until a person had completed the first grade. Education was also filled in to reflect that a person's education could not decrease or become missing if it had already been defined. Finally, if a person was coded as being in the same grade for multiple academic years, but not as having repeated a grade in that time, they were considered to have progressed to the next grade over that year. Since repeating a grade is a memorable event, and since a person cannot logically remain in the same grade for multiple years without repeating a grade, I assumed this error was due to mistakes on the part of interviewers. After correcting these mistakes, which applied to 298 people (35.06%) in the sample, I find grade progression rates that are near the national average, although still slightly lower; this likely reflects that my sample is entirely made up of low-income students, who are less likely to advance to the next grade than their higher-income peers (Corman 2003).

Sample Creation & Treatment Assignment

When evaluating the impact of a natural experiment, in which some random exogenous change impacts only one group or multiple groups differentially, it is not always straightforward to assign people to a treatment group. In this case, since the federal government universally offered the SNAP increase to everyone eligible for benefits who was not already receiving the maximum benefit for their household size, a person was considered to have received the "treatment" of increased SNAP benefits if they were receiving benefits below their maximum amount in March of 2020. A person in my sample who did not receive SNAP benefits in March of 2020, or who was receiving maximum benefits already (so that they did not receive an increase in benefits) was considered to be in the control group. While the people in my control group are not necessarily eligible to receive benefits since program eligibility is not captured by the SIPP, I attempted to match the groups in my sample so that they would still be as similar as possible.

Table 1. *Sample Size Determination*

	Sample Size	Loss
Initial sample	54,276	-
Non-missing data	31,358	22,918
K-12 age	5,040	26,318
Enrolled 2019/2020	4,946	94
High school non-grad	4,770	176
Income/poverty <1.85	850	3,920

To construct my treatment and control groups in such a way that I was confident in their similarity, I restricted my sample in several ways. I began with only respondents with non-missing data in September 2019 and December 2020, a pre- and post-SNAP-increase month (N=31,358). I then restricted my sample to only include those who were the correct age for a K-12 student (maximum age in 2020 between five and 19 years old), because educational policies were highly specific to individual colleges and universities (N=5040). To ensure I did not include people who had already dropped out or graduated from high school, I dropped any person who was unenrolled or had missing enrollment data in both months of the sample (N=4946), and anyone who had graduated high school by September 2019 (N=4770). Finally, the sample was constrained to respondents earning incomes under 185% of the federal poverty level (FPL) for their household size. This threshold was chosen to reflect that many states allow people to be eligible for SNAP if their gross income is between 130% and 185% of the FPL, and regardless of whether a person actually receives SNAP benefits, they are fundamentally similar to other people who are eligible for the program. Some states use a threshold of 200% for determining SNAP eligibility, but because many more states set the income-to-poverty threshold at 200% for only households with elderly or disabled residents, choosing this as the cutoff for the sample may have resulted in over-inclusion of households with these groups. After restricting the sample in this way, the final number of respondents decreased to 850 people.

Table 2. *Descriptive Statistics, September 2019*

	Control	Treatment	Total
N	429 (50.5%)	421 (49.5%)	850 (100.0%)
Sex			
Female	215 (50.1%)	220 (52.3%)	435 (51.2%)
Male	214 (49.9%)	201 (47.7%)	415 (48.8%)
Race			
White alone	315 (73.4%)	263 (62.5%)	578 (68.0%)
Black alone	69 (16.1%)	108 (25.7%)	177 (20.8%)
Asian alone	19 (4.4%)	8 (1.9%)	27 (3.2%)
Other	26 (6.1%)	42 (10.0%)	68 (8.0%)
Grade level of spell of enrollment.	5.40 (3.53)	5.29 (3.53)	5.3 (3.53)
Age	10.28 (3.92)	10.31 (3.86)	10.30 (3.89)
Black head of household?			
Not Black	359 (83.7%)	310 (73.6%)	669 (78.7%)
Black	70 (16.3%)	111 (26.4%)	181 (21.3%)
Female head of household?			
Male	99 (23.1%)	57 (13.5%)	156 (18.4%)
Female	330 (76.9%)	364 (86.5%)	694 (81.6%)
Hispanic head of household?			
Not Hispanic	247 (57.6%)	264 (62.7%)	511 (60.1%)
Hispanic	182 (42.4%)	157 (37.3%)	339 (39.9%)
Parent 1 high school graduation status			
Did not graduate high school	139 (32.4%)	140 (33.3%)	279 (32.8%)
Graduated high school	290 (67.6%)	281 (66.7%)	571 (67.2%)
Food security			
High/marginal food security	340 (79.3%)	269 (63.9%)	609 (71.6%)
Low food security	61 (14.2%)	94 (22.3%)	155 (18.2%)
Very low food security	28 (6.5%)	58 (13.8%)	86 (10.1%)
Number of persons in household this month	5.24 (1.93)	4.44 (1.61)	4.84 (1.82)
Household income this month ($1000s)	2.92 (2.07)	2.09 (1.46)	2.51 (1.84)

Table 3. *Descriptive Statistics, December 2020*			
	Control	Treatment	Total
N	429 (50.5%)	421 (49.5%)	850 (100.0%)
Sex			
Female	214 (49.9%)	220 (52.3%)	434 (51.1%)
Male	215 (50.1%)	201 (47.7%)	416 (48.9%)
Race			
White alone	311 (72.5%)	266 (63.2%)	577 (67.9%)
Black alone	62 (14.5%)	106 (25.2%)	168 (19.8%)
Asian alone	24 (5.6%)	8 (1.9%)	32 (3.8%)
Other	32 (7.5%)	41 (9.7%)	73 (8.6%)
Grade level of spell of enrollment.	5.93 (3.69)	6.03 (3.64)	5.98 (3.66)
Age	11.56 (3.92)	11.51 (3.82)	11.54 (3.87)
Black head of household?			
Not Black	360 (83.9%)	310 (73.6%)	670 (78.8%)
Black	69 (16.1%)	111 (26.4%)	180 (21.2%)
Female head of household?			
Male	106 (24.7%)	62 (14.7%)	168 (19.8%)
Female	323 (75.3%)	359 (85.3%)	682 (80.2%)
Hispanic head of household?			
Not Hispanic	249 (58.0%)	264 (62.7%)	513 (60.4%)
Hispanic	180 (42.0%)	157 (37.3%)	337 (39.6%)
Parent 1 high school graduation status			
Did not graduate high school	142 (33.1%)	140 (33.3%)	282 (33.2%)
Graduated high school	287 (66.9%)	281 (66.7%)	568 (66.8%)
Food security			
High/marginal food security	346 (80.7%)	303 (72.0%)	649 (76.4%)
Low food security	59 (13.8%)	72 (17.1%)	131 (15.4%)
Very low food security	24 (5.6%)	46 (10.9%)	70 (8.2%)
Number of persons in household this month	5.29 (2.00)	4.47 (1.66)	4.88 (1.88)
Household income this month ($1000s)	2.93 (2.81)	2.12 (1.66)	2.53 (2.35)

Tables 2 & 3 present descriptive statistics of each variable used in this analysis, separated by treatment (received SNAP benefits in March 2020) and control group (did not receive SNAP benefits in March 2020) status, and for the entire sample overall. The sample comprises 850 individuals—429 in the control group, 421 in the treatment group—in two periods, September 2019 and December 2020. Due to small changes in demographic information between groups (e.g. people changing their reported sex or race), the following discussion of demographic statistics focuses on the information given in December 2020.

The treatment group is slightly more female (52.3% compared to 49.9%), less white (63.2% vs. 72.5%), more Black (25.2% vs. 14.5%), and less Asian (1.9% vs. 5.6%) than the control group. People in both groups are similar ages (11.51 vs. 11.56 mean years of age for treatment and control groups, respectively). In terms of household head demographics, the household heads of the treatment group are also more female (85.3% vs. 75.3%), more Black (26.4% vs. 16.1%), and less Hispanic (37.3% vs. 42.0%). This may be partly due to a general lower use of means-tested programs by Hispanic people (Bitler 2021). Parents of the two groups have nearly identical rates of high school graduation, 66.7% for the treatment group and 66.9% for the control group.

Monthly household income is almost $1000 different between the treatment and control groups (2.12 vs. 2.93 thousand dollars) and does not change appreciably between periods. A similar trend exists for household size: the control group has an average of 5.29 people in the household compared to 4.47 people in the treatment group in December 2020, and this is roughly the same as in September 2019. The largest difference, both across groups and across periods, is in the proportion of people experiencing food insecurity at any level. In September 2019, 79.3% of the control group is full or marginally food secure, compared to 63.9% of the treatment group; in December 2020, food security for both groups increases (80.7% and 72.0%, respectively), but much more so for the treatment group. While not the specific focus of this paper, this could indicate the success of the SNAP benefit increase in alleviating food insecurity for vulnerable populations, as SNAP recipients saw a larger decrease in food insecurity after the policy was put into effect; the findings from Bryant & Follett (2022) corroborate this idea.

V. Methodology

Theoretical Model

To evaluate the effect of a natural experiment, such as the March 2020 increase in SNAP benefits for all recipients, I needed to overcome the issue of *selection bias*, the issue that arises when observations in data are not chosen at random for some treatment or sample. In this case, there are several theoretical reasons that recipients of the SNAP increase would be systematically different from non-recipients even in its absence. Comparing SNAP recipients to the general population, it is easy to see that the two groups would be different. Being eligible for SNAP requires a lower income and asset level than most of the U.S. population, and people with lower incomes generally have worse educational outcomes—but even among people eligible for the program, there could be a gap between recipients and non-recipients. Since SNAP benefits are given only to those who apply, for instance, the typical SNAP recipient may be more motivated to succeed or somewhat better educated than another qualifying person by the fact that they successfully underwent the application process. The idea that SNAP recipients may have better motivation than eligible non-recipients is particularly troubling for comparison of the two groups, since there is no way to directly control for unobservable characteristics like this based on observed data. Difference-in-differences (DiD) models are designed specifically to overcome this type of bias.

To use the DiD approach, several conditions must be met. Within two groups, a treatment of some kind must have been administered to one group, but not to the other group; importantly, there cannot be any *leakage*, in which a member of the control group is affected by the treatment. There must be data on the outcome of interest for each observation in a period before and a period after the treatment. The most important assumption for DiD estimation, though, is that the treatment and control groups experience the same trend over time in the outcome of interest and would continue to experience this same trend even without the treatment. The existence of these *parallel trends* between groups is critical to the DiD method, because the approach hinges on the counterfactual assumption that in the absence of the treatment, the treatment group would have followed the same trend over time as the control group. If these criteria are met, then the standard DiD approach yields an unbiased estimate of the treatment's impact, controlling for initial differences between the treatment and control group and the change in outcome for the control group over time.

Since my dependent variable is at the person-year level, and the SIPP only includes data on individuals over a four-year period, verifying parallel trends in grade progression prior to the treatment was not possible. To provide some evidence that this assumption holds true for the treatment and control groups, I plotted monthly enrollment rates for each group from January 2019 until December 2020 (Figure 1). From this figure, it appears that the assumption is valid: the trends follow roughly the same pattern until April 2020, then diverge more than in previous months.

Figure 1. *Monthly Enrollment Rate 2019-2020, K-12 Students*

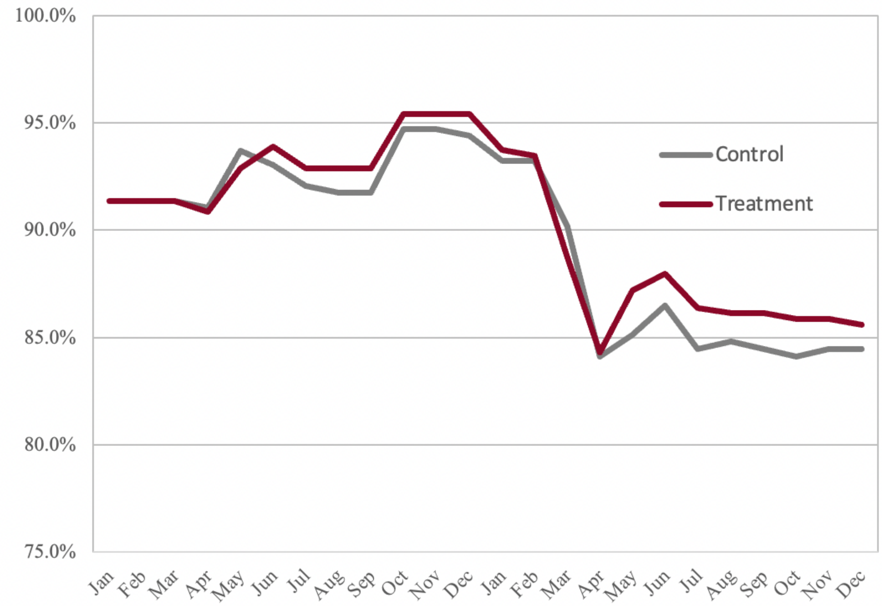

While it is encouraging to see relatively parallel trends before the intervention in March of 2020, and a divergence in the trends afterwards, the figure could suggest other interpretations. Most importantly, at the same time that the SNAP benefit increase took effect, nearly every facet of society changed as COVID-19 spread around the U.S.; it is possible, then, that this break in trend is the result of other noise due to this simultaneous change. By restricting my sample in such a way that the control and treatment groups are as similar as possible, I hoped to address this issue, since COVID-19 would presumably impact the treatment and control groups similarly. However, as discussed in the previous section, the demographic makeups of the treatment and control group are different, so the pandemic could have impacted the groups differently, resulting in a violation of the parallel trends assumption. It is also worth noting that there are some significant increases in enrollment in multiple months over this period: in September of 2019, this increase is likely attributable to the start of a new school year, which naturally suggests an increase in enrollment rates; in April of 2020, the increase in enrollment reflects students beginning virtual classes after being "unenrolled" when the pandemic first struck in March.

The below figure presents a simplified version of the rationale for the difference-in-differences approach. In the figure, Y_{dt} is an observed outcome for a person in group d (T = treatment, C = control) at time $t = \{0, 1\}$, with $t = 0$ representing pre-treatment and $t = 1$ representing post-treatment.

Table 4. *Difference in Differences Model*

	Pre-treatment	Post-treatment	Difference
Treatment group	Y_{T0}	Y_{T1}	$Y_{T1} - Y_{T0}$
Control group	Y_{C0}	Y_{C1}	$Y_{C1} - Y_{C0}$
Difference	$Y_{T0} - Y_{C0}$	$Y_{T1} - Y_{C1}$	$(Y_{T1} - Y_{T0}) - (Y_{C1} - Y_{C0})$

The naive difference between the treatment group in the pre- and post-treatment periods, $Y_{T1} - Y_{T0}$, does not account for initial differences between the treatment and control groups, which leads to biased estimates of the treatment's effectiveness. By taking the difference between each group's ending value and its initial value, then subtracting the control group's difference over time from the treatment group's difference over time, the DiD model accounts for this. Assuming that the control group provides a valid counterfactual for the treatment group, the "difference in differences" in the bottom right of the table will yield the average effect of the treatment on the treated (ATET).

Even though the DiD approach is theoretically unbiased when the above assumptions hold true, adding other relevant independent variables to a model can help increase the precision of the treatment effect estimate. For this reason, in addition to the standard DiD indicators, I included the following control variables: indicators for a person having a Black, female, and/or Hispanic head of household (*black_hoh, female_hoh,* and *hispanic_hoh*, respectively), as well as variables for the natural logarithm of a person's household size (*lhhsize*), monthly income in 1000s of dollars (*hhinck*), food insecurity status (*food_insecure*), and an indicator for whether their first listed parent graduated from high school (*par1_hsgrad*). Each of these covariates were chosen because of evidence from existing analyses of SNAP and known associations between them and educational attainment in general (see Section III). The logarithm of household size reflects that changes in household size are often non-linearly related to other aspects of the household, with each additional household member having a diminishing impact on household resources.

Limited Dependent Variable Models

Note: this section draws heavily from Wooldridge (2019).

Analyzing a binary dependent variable necessitates making a decision between using an ordinary least squares linear probability model (OLS, LPM) or maximum likelihood estimation (MLE) approach. Each method brings with it a distinct set of advantages and disadvantages.

The primary advantage of an LPM is that it is easier to interpret, since the displayed results are the marginal effects of each coefficient on the dependent variable. There are multiple disadvantages to using the LPM approach, however. Since the dependent variable is dichotomous (either taking a value of zero or one), its error variance is given by formula

$$Var(u) = P_{it}(1-P_{it})$$

where *u* is the idiosyncratic error term in the model, and P is a person's predicted probability of progressing to the next grade in a given year. It is clear from this notation that the assumption of homoscedasticity must be violated in this case, because the variance changes depending on a person's predicted probability of realizing the outcome. The error term can also only take on two values —since a person either has probability one or zero of realizing the binary outcome—so it is distributed binomially, rather than normally, and statistical inference is theoretically invalid. Finally, an LPM may predict probabilities greater than one or less than zero, because there is no assumption in OLS that the values fall in a particular range.

Often, a better method to analyze a binary dependent variable is logistic regression. Logistic regression uses an MLE approach, which maximizes the likelihood of observing a set of data given a probability specification (in this case, the logistic curve or sigmoid function) and its parameters. A brief technical description of MLE and logistic regression are included in the Appendix. This method ensures that the predicted probabilities are between zero and one, and allows for a nonlinear trend (which is more likely the functional form of the outcome variable). Typically, the main disadvantage of logistic regression is that the results are *log odds*, the natural logarithm of the ratio of an outcome's probability of being observed to its probability of not being observed, so the results do not have the easy interpretation of OLS results. There are two primary approaches to find a marginal result from logistic regression: (1) to take the derivative of the function with respect to each independent variable at the mean value of each variable, which yields the *marginal effect at the average* (MEA), and (2) to take the derivative of the function with respect to each independent variable for each observation at its observed values, then average the effects, which gives the *average marginal effect* (AME). Using the AME in this model provides a more clear interpretation than MEA—if using MEA, the marginal effects would be calculated at the sample proportion for any binary covariates, when they instead take on only values of zero and one.

With more advantages at face value, logistic regression would seem to be the obvious choice for a model like the one discussed in this paper. However, using logistic regression with repeated cross-sectional data, and for difference-in-differences estimation in general, brings with it a host of issues. A key assumption for logistic regression is independence of observations; when observations in a dataset are repeated measures of one individual, as they are in the SIPP, this assumption is violated (Stoltzfus 2011). With regards to difference-in-differences specifically, Lechner (2011) proves that the linear common trends assumption cannot be valid when using a nonlinear specification, unless the group-specific differences are zero, because the unobserved outcome term does not difference out. This is due to the fact that under the standard DiD assumption, the common trend between groups is additive (i.e. an increase in the control group's pre-treatment outcome is assumed to correspond to the same absolute increase in the treatment group's unobserved pre-treatment outcome), which is not the case for non-linear DiD specifications (Imbens and Athey 2006). Thus, using logistic regression for difference-in-differences analysis has multiple theoretical drawbacks that invalidate standard inference from such a model.

Fortunately, despite the shortcomings of the model, LPMs are still used in practice because its theoretical problems often result in small practical issues that are quickly solved, and present easily-interpreted results. In this analysis, standard errors are robust to heteroscedasticity because the method of balanced repeated replication ensures variance estimates that are unbiased estimates of the population variance (Rao 1996). The issue of error normality, while theoretically a problem for hypothesis testing and statistical analysis, is practically not an issue due to the sample size of my dataset—with a large enough sample size, the normality of errors does not significantly impact standard error estimates (Schmidt 2018). Even if some predicted probabilities may be greater than one or less than zero, the LPM is still useful for evaluating the marginal effect of independent variables not at their extreme values; in this analysis, less than 8% of observations fall outside of the correct range for each of the models with controls, and none do for the initial model. Finally, it is worth noting that LPMs perform better when more of the independent variables are binary or have few values—this study primarily includes variables which meet this criteria.

Due to the infeasibility of logistic regression with difference-in-differences, and the presence of solutions to common theoretical problems with the LPM in this dataset, results from the LPM are primarily shown and discussed. AMEs from a logistic regression are reported as a brief robustness check on the sign of the relevant independent variables, although the results are not discussed in detail because of the aforementioned issues with using logistic regression for difference-in-differences models.

VI. Results & Discussion

Table 5. *Difference in Differences, Yearly Grade Progression Rate*			
	Pre-treatment	**Post-treatment**	**Difference**
Treatment group	0.9272	0.9697	0.0425
Control group	0.9581	0.9819	0.0238
Difference	-0.0309	-0.0122	**0.0187**

The simple difference-in-differences table (Table 5) reveals three key pieces of information. First, being in the treatment group is associated with a 3.09 percentage point decrease in the likelihood that a person will progress to the next grade in a given year. Between 2019 and 2020, both groups see an increase in this likelihood. For the control group, the increase is 2.38 percentage points, and for the treatment group, the increase is 4.25 percentage points. Under the parallel trends assumption, the time trend would have been a 2.38 point increase for the treatment group as well; instead, the treatment group had an additional increase of 1.87 percentage points in their likelihood of progressing to the next grade.

$$progressed\_this\_year_{it} = \delta_0 treatment_{it} + \delta_1 post_{it} + \delta_2 treatment_{it} * post_{it} + u_{it} \quad (1)$$
$$progressed\_this\_year_{it} = \delta_0 treatment_{it} + \delta_1 post_{it} + \delta_2 treatment_{it} * post_{it} + \mathbf{X}\boldsymbol{\beta} + u_{it} \quad (2)$$

Results from three OLS difference-in-differences models are displayed in Table 6. Model (1) is a simple DiD model, with no other covariates; the regression exactly matches the difference-in-differences table above. Model (2) includes additional covariates, represented by the vector **X.** The first iteration of model (2) includes demographic information about the household head and household size. The second iteration includes each of the previously mentioned covariates as well as non-demographic information about a person: household income (in $1000s), and indicators for their first parent graduating high school and food insecurity status.

Table 6. *OLS Regression Results*			
	(1)	(2)	(3)
VARIABLES	Initial Model	Incl. Demographics	Full Model
1.treatment	-0.0309	-0.0287	-0.0302
	(0.0205)	(0.0215)	(0.0221)
1.post	0.0238**	0.0240**	0.0241**
	(0.00991)	(0.00998)	(0.00997)
1.te	0.0188	0.0189	0.0196
	(0.0191)	(0.0192)	(0.0191)
1.black_hoh		-0.0249	-0.0240
		(0.0191)	(0.0187)
1.hispanic_hoh		0.0160	0.0196
		(0.0133)	(0.0135)
1.female_hoh		-0.0264**	-0.0289***
		(0.0103)	(0.0106)
lhhsize		-0.0207	-0.0181
		(0.0202)	(0.0220)
hhinck			-0.000194
			(0.00232)
1.par1_hsgrad			0.0126
			(0.0172)
1.food_insecure			0.0133
			(0.0139)
Constant	0.958***	1.009***	0.994***
	(0.00977)	(0.0366)	(0.0458)
Observations	1,700	1,700	1,700
R-squared	0.011	0.020	0.022

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In each model, only the indicator for being in the post-treatment period is significant, and the treatment effect is statistically insignificant. Adding demographic covariates does not change the sign or significance of these variables. The additional controls also do not alter the general results, although the estimates and standard errors change very slightly. Among the controls, only the indicator for a female head of household is significant at the 5% level. Besides this, no conclusions can be definitively drawn about the exact relationship between the dependent and independent variables, due to high standard errors on the coefficient estimates.

The results indicate no statistically significant underlying difference between SNAP recipients and non-recipients with regards to the ability to advance to the next grade. This may be because for similarly low-income students, after controlling for other characteristics, there could truly be no major difference in their grade progression rates. However, the standard errors of the estimate result in a wide 95% confidence interval [(-0.0737, 0.0132)], so the true effect could be as large in magnitude as a 7.37 point decrease, or potentially even positive.

The indicator for being in the post-treatment period, as expected, is significant at the 5% level in all models. It is unsurprising that this variable is significant, because looser regulation regarding grade progression at the beginning of COVID-19 resulted in higher across-the-board progression rates, which indiscriminately affected students in each school which changed their regulations. In the final model, being in the post-benefit-increase period is associated with a 2.41 percentage point increase in a student's probability of progressing to the next grade, after controlling for other relevant variables [95% CI: (0.0045, 0.044)]. This effect is similar in each model, even without the addition of controls.

The lack of a statistically significant treatment effect in each model is not surprising, and corroborated by existing literature. Specifically, Heflin's findings that mitigating instability does not have a significant impact on a person's material hardship are relevant: by increasing SNAP benefits, the federal government may have prevented a source of instability from worsening for SNAP recipients, but their educational attainment would not have necessarily improved as a result. Like the initial difference between the treatment and control group, the estimate has a large 95% confidence interval [(-0.0181, 0.0572)], so the true effect could be much higher or lower than the point estimate suggests.

Interestingly, the only demographic control to have a significant impact on grade progression rates was the indicator for having a female head of household. Having a female head of household is associated with a 2.89 percentage point decrease [95% CI: (-0.0498, -0.0081)] in the probability of a student progressing to the next grade on average, *ceteris paribus*. While this is the expected sign of the coefficient, I was surprised that this was the only significant demographic control, given that existing literature supports a negative association between having a female, Black, or Hispanic head of household on educational outcomes. It could be that this association simply does not carry over to grade progression, rather than general educational outcomes (quantified through test scores or years of education); further research could explore this area.

The insignificance of most other coefficient estimates is consistent across each iteration of the three models, and was not necessarily surprising to find. Since grade progression was made much easier to attain through a variety of methods at the start of the COVID-19 pandemic, it makes sense that very little would impact grade progression rates besides individual district actions (for which I did not have data) and being in the post-COVID period. The high standard errors could also have been the result of low power in my model; this limitation is discussed in the following subsection.

The point estimates of the coefficients on the control variables in each model, while not the focus of this analysis, generally align with the theoretical direction they should take, with some notable exceptions. Based on existing literature, I expected the coefficients on *black_hoh*, *female_hoh*, *hispanic_hoh, lhhsize*, and *food_insecure* to all be negative; *hispanic_hoh* and *food_insecure* are both positive. It is possible that neither of these estimates are correct, though, as the 90% confidence intervals for both estimates include negative numbers [90% CI: (-0.008, 0.045); (-0.014, 0.041), respectively]. The coefficient estimates on the other variables—*par1_hsgrad* and *hhinck*—are both positive, as expected; again, however, the 90% confidence intervals for each estimate range widely over negative and positive numbers [90% CI: (-0.022, 0.047); (-0.005, 0.004), respectively].

It is worth noting that although the point estimates for each coefficient may seem small in magnitude, they represent large changes in the grade progression rate. The constant in the first regression is 0.958, meaning that for a member of the control group in 2019, the probability of progressing to the next grade in that year is 95.8%. The increase from this number to the progression rate in 2020, 2.38%, represents a very large increase relative to the initial rate, and the same idea holds for every estimate in the model. While the coefficient estimates relating to the treatment are not statistically significant from zero, the ranges of values in the 95% confidence intervals indicate that the benefit increase could have had a very economically significant impact on grade progression rates for people who received it.

Limitations

The main limitation of my analysis is low power, or, equivalently, a high probability of falsely failing to reject my null hypotheses. The low power is primarily related to the low range of predicted probabilities from my model. Due to the laws surrounding educational participation in the United States, grade progression does not vary much in my sample. While this is not an issue for the unbiasedness or validity of the OLS model, it does mean that the predicted probabilities of individuals progressing to the next grade in a given year are generally high and do not vary particularly much, being between 0.87 and 1.04. Since the predicted probabilities are a function of the independent variables, this implies that there is low variation in the independent variables, which results in high standard errors. The fact that this range is very close to one boundary of the dependent variable (which only takes on values of 0 and 1) also presents issues for the power of the LPM: it indicates that the model does not accurately predict when people do not progress to the next grade. Moreover, the process of OLS fits a linear relationship between the dependent and independent variables, but this relationship is inherently nonlinear when the dependent variable is binary. When a person already has a high probability of progressing to the next grade, the impact of other factors on this probability should be smaller than if the person had a lower probability of progressing; with such high predicted probabilities, the model may overestimate the true impact of the increase in SNAP benefits on grade progression.

It could be possible that high multicollinearity—linear correlation between the regressors of a model—in the independent variables may also increase the standard errors of the estimated coefficients. To test for multicollinearity, it is common to use a variance inflation factor (VIF), a measure of how much a coefficient's variance is inflated due to collinearity. The R-squared in the VIF calculation for an independent variable is the R-squared of the regression of that variable on all other independent variables. Thus, a high R-squared (a marker of high explanatory power of regressors on a regressand) translates into a higher VIF, which indicates higher collinearity in that variable. Table 7 presents results from VIF analysis computed after the final regression. Based on these results, it is not clear that there is high collinearity between all independent variables.

Table 7. *VIF Results*

Variable	VIF
1.treatment	2.21
1.post	1.98
1.te	3.07
1.black_hoh	1.18
1.hispanic_hoh	1.29
1.female_hoh	1.07
lhhsize	1.2
hhinck	1.15
1.par1_hsgrad	1.13
1.food_insecure	1.05
Mean VIF	1.53

The size of my sample could also present a concern. While 1700 observations is enough to conduct statistical analysis, increasing this sample size would help to increase the total variation in my independent variables, which would in turn improve the precision of my estimates.

VII. Conclusion

The United States offers many means-tested social services to help aid struggling members of its society; however, the impact of these programs on problems outside of their exact scope has not been extensively studied. This paper attempted to address this gap, using data from the Survey of Income and Program Participation to determine the impact of an exogenous increase in SNAP benefits on grade progression rates for recipients of the increase. Through a linear probability model difference-in-differences approach, I find inconclusive results regarding this effect, as well as on most of the relevant controls I included; however, being in the post-COVID period has a statistically significant positive effect on grade progression, and having a female head of household has a statistically significant negative effect.

Previous literature suggests that benefit increases may not have an impact on markers of stability, but rather act through a prevention of worsened instability (Heflin 2016). Assuming this is true, it makes sense that my results are generally inconclusive and statistically insignificant. If the increase in benefits only prevented grade progression rates from worsening, this would not show up as a significant effect in my regression. However, no definitive conclusion about the treatment effect can be drawn based on the lack of statistical significance. Future analyses may attempt to overcome the limitations of this paper to increase the precision of these estimates, potentially by using a larger sample or one that included more students who did not progress to the next grade.

My results also corroborate the idea that, as a response to COVID-19, educators were given great discretion in choosing the standards for grade progression. This is likely why there were so few factors that impacted grade progression in my sample, and why one of the only two factors which did was the indicator for being in the post-COVID period. It is possible, though, that the simultaneous changes from COVID-19, which drastically altered nearly every aspect of U.S. society in some way, obscured the effect of the benefit increase on grade progression or caused some bias in my estimates.

My model highlights the potential for future research in the area of educational effects of COVID-19. This paper focused on grade progression, a measure of education *quantity*, whereas education *quality* was also severely impacted by COVID-19. In the initial pandemic phase, teachers were forced into online learning, often with less than a month's preparation, which could have led to drastically worse quality of education. Trends over time for primary school math and reading scores support this notion, with both sets of scores decreasing dramatically after 2020 compared to earlier years. Thus, a future analysis could replicate the work in this paper, but evaluating the effect of a benefit increase on standardized test scores or a different marker of educational success. In this way, the impact of these benefits on educational outcomes could be made more clear.

APPENDIX

A. Logistic Regression and Maximum Likelihood Estimation Overview

Note: this section draws heavily from Wooldridge (2019)

To overcome the issues associated with the linear probability model, we assume that the probability of the outcome variable occurring (the *response probability*) is of the form

$$P(y=1 \mid \mathbf{X}) = G(\mathbf{X}\boldsymbol{\beta}), \qquad \text{(A1)}$$

where, for a model with n observations and k independent variables plus one constant, G is an undetermined nonlinear function which takes on values between 0 and 1 inclusive, $\mathbf{X}$ is an nxk matrix comprising a column of ones and a column for each explanatory variable, and $\boldsymbol{\beta}$ is a kx1 vector of coefficients. For logistic regression, or *logit* regression, G follows the logistic distribution or sigmoid function

$$G(z) = \exp(z)/[1+\exp(z)] \qquad \text{(A2)}$$

Unlike in OLS regression, logit models are based on the idea of an underlying *latent variable*, y*. This variable is unobserved, and is of the form

$$y^* = \mathbf{X}\boldsymbol{\beta} + e, \qquad y=1[y^*>0] \qquad \text{(A3)}$$

where the *indicator function* 1[·] equals 1 if the condition inside the brackets is true, and zero otherwise. The error term e is assumed to be independent of $\mathbf{X}$ and symmetrically distributed according to the standard logistic function around zero. This notation of (A4) implies that

$$\begin{aligned} P(y = 1 \mid \mathbf{X}) &= P(y^* > 0 \mid \mathbf{X}) = P(e > -(\mathbf{X}\boldsymbol{\beta}) \mid \mathbf{X}) = \\ &= 1 - P(e \le -(\mathbf{X}\boldsymbol{\beta}) \mid \mathbf{X}) = 1 - G[-\mathbf{X}\boldsymbol{\beta}] = G(\mathbf{X}\boldsymbol{\beta}), \qquad \text{(A4)} \end{aligned}$$

the original response probability for y.

Estimation of this equation is well suited to the technique of *maximum likelihood estimation* (MLE). MLE maximizes the probability of observing a given set of data based on a particular set of parameters; here, these parameters are represented by $\boldsymbol{\beta}$. The likelihood function for (y | $\mathbf{X}$) for observation i is given by

$$f(y \mid \mathbf{X}_i; \boldsymbol{\beta}) = [G(\mathbf{X}_i\boldsymbol{\beta})]^y[1 - G(\mathbf{X}_i\boldsymbol{\beta})]^{1-y}, \ y \in [0, 1] \qquad \text{(A5)}$$

The *log-likelihood* function for $(y \mid \mathbf{X})$ for observation i is found by taking the natural log of (A5), and is given by

$$l_i(\beta) = y_i \log[G(\mathbf{X}_i\beta)] + (1 - y_i)\log[1 - G(\mathbf{X}_i\beta)] \qquad \text{(A6)}$$

The log-likelihood for a random sample with n observations is found by summing the individual log-likelihoods over all n, and the maximum likelihood estimator finds the maximum of this log-likelihood. The MLE is consistent and both asymptotically normal and asymptotically efficient.

Importantly, MLE accounts for the heteroscedasticity in $\mathrm{Var}(y \mid \mathbf{X})$ since it fits the data to the particular distribution of $(y \mid \mathbf{X})$(in this case, the logistic distribution).

B. Logistic Regression AME Results

VARIABLES	(1) Logit Model	(2) Demographics	(3) Full
1.treatment	-0.0219	-0.0200	-0.0215
	(0.0146)	(0.0161)	(0.0168)
1.post	0.0316**	0.0314**	0.0319**
	(0.0140)	(0.0140)	(0.0141)
1.te	0.00227	0.00292	0.00340
	(0.0184)	(0.0183)	(0.0182)
1.black_hoh		-0.0222	-0.0204
		(0.0183)	(0.0174)
1.hispanic_hoh		0.0157	0.0200
		(0.0140)	(0.0143)
1.female_hoh		-0.0276***	-0.0295***
		(0.0106)	(0.0104)
hhsize		-0.00397	-0.00341
		(0.00398)	(0.00438)
hhinck			-0.000634
			(0.00256)
1.par1_hsgrad			0.0139
			(0.0189)
1.food_insecure			0.0132
			(0.0123)
Observations	1,700	1,700	1,700

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

The interpretation of the average marginal effects computed after logistic regression should be done with caution, for the reasons discussed in Section V. However, the statistical significance of each variable is the same as in the OLS regression model, and the sign of the significant variables is the same, providing some robustness to the OLS results.

C. *Survey Weights for the SIPP*

The SIPP provides a set of replicate and probability weights in order to account for the complex sampling design of the survey. Many statistical softwares, including STATA, have a "survey" option through which these weights are automatically applied; for ease of use, the SIPP documentation includes the specific code for commonly used softwares. In STATA, the code is as follows:

```
svyset [pw=WPFINWGT], brrweight(repwgt1-repwgt240) fay(.5) vce(brr) mse,
```

where WPFINWGT is an individual's probability weight for each month and [repwgt1, repwgt2,..., repwgt240] are the 240 replicate weights used for the process of balanced repeated replication.

D. Figures

Figure D1: Benefits received from 2020 Q1 – 2022 Q2 for three typical SNAP recipients. "Maximum" is the maximum possible SNAP benefit for a recipient, "New Max." is the new maximum possible SNAP benefit after a 15% increase at the beginning of 2021, and "Max. + $95" reflects the additional $95 people who were already at the maximum benefit level became eligible for in April 2021.

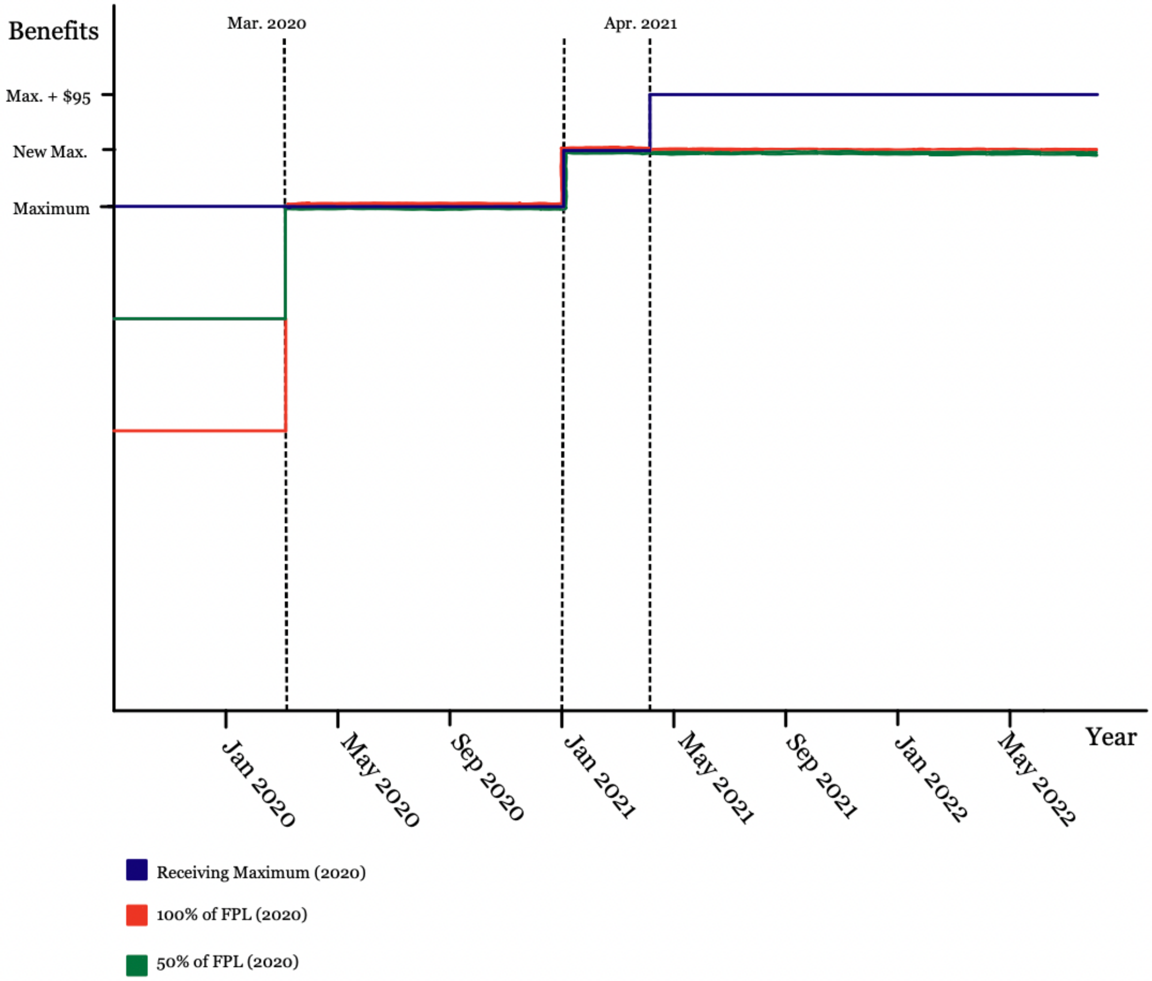

References

Alaimo, Katherine. “Food Insecurity in the United States: An Overview.” *Topics in Clinical Nutrition* 20, no. 4 (December 2005): 281–98.

Ambrose, Adrian Jacques H. “Inequities During COVID-19.” *Pediatrics* 146, no. 2 (August 1, 2020): e20201501. https://doi.org/10.1542/peds.2020-1501.

Anderson. “Core Indicators of Nutritional State for Difficult-to-Sample Populations.” *The Journal of Nutrition* 120 Suppl 11 (November 1990): 1559–1600. https://doi.org/10.1093/jn/120.suppl_11.1555.

Anderson, Gabrielle E., Shane R. Jimerson, and Angela D. Whipple. “Student Ratings of Stressful Experiences at Home and School.” *Journal of Applied School Psychology* 21, no. 1 (July 6, 2005): 1–20. https://doi.org/10.1300/J370v21n01_01.

Angrist, Joshua D., and Alan B. Krueger. “Does Compulsory School Attendance Affect Schooling and Earnings?” *The Quarterly Journal of Economics* 106, no. 4 (1991): 979–1014. https://doi.org/10.2307/2937954.

Argaw, Thomas Lemma, Jasmine Fledderjohann, Elisabetta Aurino, Sukumar Vellakkal. *Children’s Educational Outcomes and Persistence and Severity of Household Food Insecurity in India: Longitudinal Evidence from Young Lives*. The Journal of Nutrition, Volume 153, Issue 4, 2023, Pages 1101-1110. https://doi.org/10.1016/j.tjnut.2023.02.008.

Athey, Susan, and Guido W. Imbens. “Identification and Inference in Nonlinear Difference-in-Differences Models.” *Econometrica* 74, no. 2 (2006): 431–97. https://doi.org/10.1111/j.1468-0262.2006.00668.x.

Blanden, Jo, Matthias Doepke, and Jan Stuhler. “Educational Inequality*.” Working Paper. Working Paper Series. National Bureau of Economic Research, April 2022. https://doi.org/10.3386/w29979.

Bryant, Andrew, and Lendie Follett. “Hunger Relief: A Natural Experiment from Additional SNAP Benefits during the COVID-19 Pandemic.” *The Lancet Regional Health - Americas* 10 (June 1, 2022): 100224. https://doi.org/10.1016/j.lana.2022.100224.

Choi de Mendizábal, A., Gil Izquierdo, M., Mediavilla Bordalejo, M., & Valbuena Gómez, J. (2018). Predictors and Effects of Grade Repetition in Spain. *Revista De Economía Mundial*, (48). https://doi.org/10.33776/rem.v0i48.3882

Corman, H. “The Effects of State Policies, Individual Characteristics, Family Characteristics, and Neighbourhood Characteristics on Grade Repetition in the United States.” *Economics of Education Review* 22, no. 4 (August 1, 2003): 409–20. https://doi.org/10.1016/S0272-7757(02)00070-5.

Duflo, Esther. “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment.” *American Economic Review* 91, no. 4 (September 2001): 795–813. https://doi.org/10.1257/aer.91.4.795.

Green, Erica L. “Relief Offered from Testing and Student Loans as Virus Roils Education.” *The New York Times*, 20 Mar. 2020, www.nytimes.com/2020/03/20/us/politics/coronavirus-student-loans-education-testing.html.

Gundersen, Craig, Brent Kreider, and John Pepper. "The Economics of Food Insecurity in the United States." *Applied Economic Perspectives and Policy* 33, no. 3 (2011): 281–303. https://doi.org/10.1093/aepp/ppr022.

Gundersen, Craig, and Victor Oliveira. "The Food Stamp Program and Food Insufficiency." *American Journal of Agricultural Economics* 83, no. 4 (2001): 875–87. https://doi.org/10.1111/0002-9092.00216.

Gundersen, Craig, and James P. Ziliak. "Food Insecurity And Health Outcomes." *Health Affairs* 34, no. 11 (November 2015): 1830–39. https://doi.org/10.1377/hlthaff.2015.0645.

Heflin, Colleen. "Family Instability and Material Hardship: Results from the 2008 Survey of Income and Program Participation." *Journal of Family and Economic Issues* 37, no. 3 (September 1, 2016): 359–72. https://doi.org/10.1007/s10834-016-9503-6.

Jalongo, Mary Renck. "The Effects of COVID-19 on Early Childhood Education and Care: Research and Resources for Children, Families, Teachers, and Teacher Educators." *Early Childhood Education Journal* 49, no. 5 (September 1, 2021): 763–74. https://doi.org/10.1007/s10643-021-01208-y.

Judkins, David R. "Fay's Method for Variance Estimation." *Journal of Official Statistics* 6, no. 3 (1990): 223–39.

Lynch, Matthew. "Alternatives to Social Promotion and Retention." *Interchange* 44, no. 3 (December 1, 2013): 291–309. https://doi.org/10.1007/s10780-014-9213-7.

Lechner, Michael. "The Estimation of Causal Effects by Difference-in-Difference Methods." *Foundations and Trends in Econometrics* 4, no. 3 (2011): 165–224. https://doi.org/10.1561/0800000014.

Mathys, Cécile, Marie-Hélène Véronneau, and Aurélie Lecocq. "Grade Retention at the Transition to Secondary School: Using Propensity Score Matching to Identify Consequences on Psychosocial Adjustment." *The Journal of Early Adolescence* 39, no. 1 (January 1, 2019): 97–133. https://doi.org/10.1177/0272431617735651.

Nestle, Marion. "The Supplemental Nutrition Assistance Program (SNAP): History, Politics, and Public Health Implications." *American Journal of Public Health* 109, no. 12 (December 2019): 1631–35. https://doi.org/10.2105/AJPH.2019.305361.

Nord, Mark, and Gary Bickel. "Estimating the prevalence of children's hunger from the Current Population Survey Food Security Supplement." *Second Food Security Measurement and Research Conference*. Vol. 2. 1999.

O'Hara, Sabine, and Etienne C. Toussaint. "Food Access in Crisis: Food Security and COVID-19." *Ecological Economics* 180 (February 1, 2021): 106859. https://doi.org/10.1016/j.ecolecon.2020.106859.

Peterson, Lisa S., and Jan N. Hughes. "The Differences between Retained and Promoted Children in Educational Services Received." *Psychology in the Schools* 48, no. 2 (February 1, 2011): 156–65. https://doi.org/10.1002/pits.20534.

Rao JN, Rust KF. "Variance estimation for complex surveys using replication techniques." Stat Methods Med Res. 1996 Sep;5(3):283-310. doi: 10.1177/096228029600500305. PMID: 8931197.

Reich, Justin, Christopher J. Buttimer, Alison Fang, Garron Hillaire, Kelley Hirsch, Laura R. Larke, Joshua Littenberg-Tobias, et al. "Remote Learning Guidance From State Education Agencies During the COVID-19 Pandemic: A First Look." EdArXiv, April 2, 2020. https://doi.org/10.35542/osf.io/437e2.

Shankar, Priya, Rainjade Chung, and Deborah A. Frank. "Association of Food Insecurity with Children's Behavioral, Emotional, and Academic Outcomes: A Systematic Review." *Journal of Developmental & Behavioral Pediatrics* 38, no. 2 (February 1, 2017): 135–50. https://doi.org/10.1097/DBP.0000000000000383.

Schmidt AF, Finan C. "Linear regression and the normality assumption." J Clin Epidemiol. 2018 Jun;98:146-151. doi: 10.1016/j.jclinepi.2017.12.006. Epub 2017 Dec 16. PMID: 29258908.

Slavin, Robert E., and Nathan Storey. "The US Educational Response to the COVID-19 Pandemic." SSRN Scholarly Paper. Rochester, NY, July 15, 2020. https://papers.ssrn.com/abstract=3652585.

Swinnen, Johan, and John McDermott. "Covid-19 and Global Food Security." *EuroChoices* 19, no. 3 (2020): 26–33. https://doi.org/10.1111/1746-692X.12288.

Tingle, Lynne R., Jason Schoeneberger, and Bob Algozzine. "Full Article: Does Grade Retention Make a Difference?" *The Clearing House: A Journal of Educational Strategies, Issues and Ideas* 85, no. 5 (July 11, 2012): 179–85.

"THE STATE OF FOOD SECURITY AND NUTRITION IN THE WORLD 2021." United Nations. Accessed November 14, 2022. https://doi.org/10.4060/CB4474EN.

U.S. Census Bureau. Survey of Income and Program Participation [2018-2021 Panels]. U.S. Department of Commerce. https://www.census.gov/programs-surveys/sipp/data/datasets/2021-data/2021.html

Whitmer, Gretchen. "Executive Order 2020-35: Provision of K-12 Education during the Remainder of the 2019-2020 School Year - RESCINDED." Accessed April 28, 2023. https://www.michigan.gov/whitmer/news/state-orders-and-directives/2020/04/02/executive-order-2020-35.

Wooldridge, Jeffrey. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. MIT Press, 2010.

———. *Introductory Econometrics: A Modern Interpretation*. 9th ed. Mason, Ohio: Cengage Learning, 2019.

Zinn, Sabine. "Variance Estimation with Balanced Repeated Replication: An Application to the Fifth and Ninth Grader Cohort Samples of the National Educational Panel Study." *Methodological Issues of Longitudinal Surveys: The Example of the National Educational Panel Study*. Accessed April 25, 2023. https://doi.org/10.1007/978-3-658-11994-2_4.