ABSTRACT

Title of Dissertation:	ESSAYS ON THE EFFECTS OF PUBLIC POLICIES ON HOUSING, EMPLOYMENT, AND INCOME INEQUALITY
	Luke William Pardue Doctor of Philosophy, 2021
Dissertation directed by:	Professor Melissa Kearney Department of Economics

In this dissertation, I study the impact of public policies on three related but distinct economic outcomes: employment, housing stability, and income inequality.

In the second chapter, I examine the employment impact of the Paycheck Protection Program, a key element of the federal government's fiscal stimulus efforts during the 2020 coronavirus-induced recession. To assess the effect of this support on small business employment, I exploit differential timing in when firms rolled off headcount requirements needed to receive loan forgiveness. I find that as the PPP covered period expired, companies reduced active employment by a statistically significant 0.41% per week and 1.6% in the four weeks post-expiration. I estimate that, in aggregate, 907,200 jobs were lost within the four weeks after firms' covered periods expired, as companies no longer need to maintain pre-COVID-19 headcount levels to receive PPP loan forgiveness.

In the third chapter, I investigate the effectiveness of housing vouchers, the most common form of low-income rental assistance, in preventing households from facing eviction. I examine this question using newly-available public data on the universe of court-ordered evictions in the United States and exploiting plausibly exogenous variation in historical housing voucher allocation. I find that every four to six vouchers prevent one eviction in a given county, and that this effect is greater in counties with higher rent burdens and longer voucher waitlists. A simple back-of-the-envelope calculation suggests that, on average, one fourth of the cost of a housing voucher can be recovered through savings in eviction prevention alone.

In the final chapter, I conduct an empirical simulation exercise that gauges the plausible impact of increased rates of college attainment on a variety of measures of income inequality and economic insecurity. Using two different methodological approaches—a distributional approach and a causal parameter approach—I find that increased rates of bachelor's and associate degree attainment would meaningfully increase economic security for lower-income individuals, reduce poverty and near-poverty, and shrink gaps between the 90th and lower percentiles of the earnings distribution. However, increases in college attainment would not significantly reduce inequality at the very top of the distribution.

ESSAYS ON THE EFFECTS OF PUBLIC POLICIES ON HOUSING, EMPLOYMENT, AND INCOME INEQUALITY

by

Luke William Pardue

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:
Professor Melissa Kearney, Chair
Professor Katharine Abraham
Professor Erich Battistin
Professor John Haltiwanger
Professor Nolan Pope

Dedication

To my mom and dad.

Acknowledgements

I owe a deep debt of gratitude to Melissa Kearney for her guidance during graduate school. Her brilliance is matched only by her patience and generosity with her time. She has pushed me to be more creative in my research, more rigorous in my analysis, and more precise in my writing.

I also am thankful for the mentorship of John Haltiwanger, who has provided a model of the type of economist I would like to be since my first year of graduate school. My committee members Katharine Abraham and Nolan Pope have both offered valuable feedback on earlier work in class and in seminars, and Erich Battistin has kindly agreed to serve as the Dean's Representative.

Many faculty members in the Economics Department at the University of Maryland have contributed time and effort to my graduate studies. In particular, John Shea's comments on earlier drafts of these chapters greatly improved them, and Judy Hellerstein pushed me in class and in seminars to think harder about my research. I am immensely thankful for the work done by Vickie Fletcher, Terry Davis, Jessica Gray, Amanda Statland, and Mark Wilkerson – much of which I am sure I am not aware of – to facilitate my research and progress through graduate school. Outside of the department, I would like to thank Brad Hershbein and Phil Levine, with whom I have collaborated and whose efforts vastly improved our work. Michael Strain also offered constructive comments on the first draft of the second chapter of this dissertation.

I am also grateful to my friends and classmates, including Ed Olivares, Ben Layton, Chris Roudiez, Nathalie Gonzalez, Elena Ramirez, Lea Rendell, Macarena

Kutscher, Mrin Chatterjee, Jamie Grasing, Cody Tuttle, Matthew Staiger, George Zuo, and other fellow graduate students, whose feedback and support vastly improved my work and my graduate school experience.

My path to graduate school was made possible by the mentorship and support of many people. I am thankful to Bill Evans, Jim Sullivan, and Kasey Buckles at the University of Notre Dame, who offered me my first opportunity to contribute to and produce academic research. John Kandrac, at the Federal Reserve Board, nurtured my creativity as an economist and showed me what it means to turn good ideas into good research.

Finally, I am unspeakably grateful to my friends and to my family for their constant support and love. Gable Brady, Matthew Bruno, Kyle Craft, Catherine Porto, Nathan and Sarah Pigott, and Mary Clare Rigali have been with me through the ups and downs of graduate school. My siblings Bill, Maggie, Molly, Mariah, and Jake have all provided immense love and support while ensuring that I remain humble, and their partners Carlie, Adam, Matthieu, Brandon, and Grace have never failed to pick up their slack in this task. My nieces and nephews Caroline, Jack, Hunter, Liam, and Milana have been a source of joy and unconditional love. Most importantly, my parents, Bill and Mariah, made me who I am today. The memory of my dad drives my work every day and has brought me to this point.

Table of Contents

List of Tables

List of Figures

Disclaimer

The fourth chapter of this dissertation was previously published as: Hershbein, Brad, Melissa S. Kearney, and Luke W. Pardue. “College Attainment, Income Inequality, and Economic Security: A Simulation Exercise.” American Economic Association: Papers and Proceedings, 110: 352-355. It is reproduced here with the following notice.

Chapter 1: Introduction

Economists justify government interventions in competitive markets on the grounds of efficiency or equity – because, for instance, a market failure or externality exists that prevents a market from achieving a socially efficient outcome, or because a more equal distribution of economic resources will *per se* improve social welfare. In this dissertation, I study the effects of three government interventions intended to promote economic efficiency or equity.

The first two chapters examine programs focused on improving efficiency, through providing small businesses with resources to preserve worker-firm relationships during a global pandemic and through providing households resources to avoid housing instability and its many external costs. I begin with a chapter looking at a novel government program introduced during the COVID-19 pandemic and resulting recession: the Paycheck Protection Program (PPP). At the onset of the COVID-19 pandemic, Congress created the Paycheck Protection Program (PPP) to help small businesses stay afloat and maintain employment. Through the program's initial expiration on August 8, 2020, the federal government guaranteed $525 billion in loans to businesses through PPP (Treasury Department 2020).

PPP was designed as a tool to replace the revenue of small businesses, who do not typically hold large savings buffers or have easy access to capital markets, in order to preserve valuable worker-firm specific capital (e.g., knowledge of specific businesses procedures, or within-firm relationships) and to prevent costly small business closures during a temporary public health emergency.

Policymakers have historically relied on firms to build savings buffers and markets to reallocate resources efficiently during business downturns, but this program has been justified on the basis that these firms have limited ability to build savings buffers (and we would not expect them to hold savings for a global pandemic) and lack the access to liquid capital markets available to larger businesses. This temporary revenue replacement was thus designed to avoid costly layoffs or firm closures among a set of businesses for which moral hazard is not a concern and market-based solutions may not exist.

Small businesses were eligible to apply for loans covering up to eight weeks of payroll, and these loans were up to 100% forgivable, given companies maintained a specified average level of employment over the length of the "covered period," in addition to other conditions (SBA 2020b). I exploit detailed payroll microdata, with information on the date of loan disbursement and the length of time the company must maintain employment levels, to estimate the effect of the Paycheck Protection Program on employment. I find that, as the PPP Covered Period expired, companies reduced active employment by a statistically significant 0.41 percent per week. Event study estimates over the four weeks after expiration indicate that firms reduced headcount by 1.6% the month after expiration. I find smaller, insignificant results in counties that have managed to contain coronavirus case growth, suggesting that revenue replacement policies such as PPP can be effective at maintaining employment (and preserving the value of worker-firm relationships) given that external conditions improve as well.

This work has important implications for the design of future small business support programs during times of crisis. These PPP headcount requirements acted as a binding constraint on firm employment behavior, indicating that many businesses emerged from their loan periods in no better shape than when they entered -- largely because public health measures did not improve. Many firms rolled off these requirements as the pandemic continued, suggesting that flexibility in the loan terms, such as tying expiration dates to a measure of public health conditions, is particularly important.

While the first chapter examines a program that was developed and implemented in a matter of weeks during a nearly unprecedented public health crisis, the second chapter of this dissertation takes a look at a longer-established program: the Section 8 Housing Voucher Program, the nation's largest housing subsidy program. These vouchers allow tenants to lease units directly from landlords in the private rental market, with the government then paying a portion of the rent directly to the landlord. In 2018, the federal government spent $50 billion dollars on low-income housing assistance programs, an amount comparable to the budget for food stamps (SNAP) or the Supplemental Security Income program (OMB, 2019).

Housing vouchers are an in-kind subsidy meant to increase the consumption of housing and reduce housing instability. One source of housing instability is eviction, the forcible removal of tenants from their rental unit, most commonly for nonpayment of rent. Each year, over 2 million households, or 5 percent of the renting population, face the threat of eviction from their housing unit (Desmond et al., 2018b). Recent and ongoing work in economics finds that evictions cause a lasting increase in housing instability, as well as mental health hospitalizations, and emergency room utilization (Humphries et al., 2019; Collinson and Reed, 2018). Collinson and Reed (2018) estimate that, between increased hospitalizations and use of homeless shelters, the cost of an eviction totals roughly $8,000 within two years. One justification for housing assistance policies is thus the reduced burden imposed on institutions that bear these costs.

I explore the effect of housing vouchers in reducing a county's eviction rate by combining a newly-available dataset compiling court records of eviction cases across the country with data on the availability and use of housing vouchers in each area. To avoid potential endogeneity issues, I exploit plausibly exogenous historical variation in one part of the voucher allocation formula: the stock of a county's rental housing built before 1940. I use this data as of 2000, during the last substantial expansion of the housing voucher program. This instrument is strongly predictive of the county's level of housing vouchers in 2016, and data on the stock of rental units built immediately after 1940 has no predictive power on voucher levels. Furthermore, this instrument is uncorrelated with measures of community decline or attitudes towards housing within the area.

I find that, across a number of specifications, every four to six vouchers prevent one eviction in a county. These effects are larger in markets where renters pay a larger share of income on rent, suggesting that vouchers play an important role in reducing cost burdens, making rent affordable for low-income families. There is also a larger effect in areas with greater demand for these vouchers (as measured by time on the voucher waiting list), meaning a targeted expansion of the program could have a greater impact on preventing evictions. A back-of-the-envelope calculation suggests that as much as a quarter of the average annual cost of a voucher can be recovered from savings through eviction prevention alone.

Finally, I turn from interventions intended to improve efficiency to examining the effect of one set of public policies on income inequality, which has been increasing in the United States over the last several decades. Among prime-age, full-time year-round workers the ratio of earnings between the 90^{th} percentile and 10^{th} percentile increased from 4.63 to 5.45 between 1979 and 2018. This rise in inequality is in part due to rising dispersion between workers with different education levels. For example, in 1979, median earnings among high school graduate FTFY workers were approximately $38,300 (in 2018 dollars), as compared to about $53,400 among FTFY workers with a bachelor's degree (BA) or higher. In 2018, the comparable numbers were $40,000 and $70,000. The divergent economic outcomes of those with and without a college degree have led many observers to emphasize the need for increased skill attainment, in particular increased college attainment, to both boost individual economic security and to address rising income inequality.

In this chapter, I conduct a simulation exercise that gauges the plausible impact of increased rates of college attainment on a variety of measures of income inequality and economic insecurity. Although several channels for increasing college attainment have been proposed—including additional funding for higher education institutions, expanded access to free or reduced tuition for students, and behavioral or information interventions— this analysis does not examine any single policy intervention. Instead, it employs a simulation method, incorporating existing causal estimates, to focus on outcomes.

We simulate three counterfactual scenarios, (1) raising the share of the workers with at least a bachelor's degree to 50 percent, (2) raising the share of the sample with an associate degree to 15 percent *and* the BA share to 50 percent, and (3) raising the AA share to 20 percent and the BA share to 60 percent. For each scenario, we assign the "new" AA and BA holders simulated earnings in two ways. The *distribution method* assigns a random draw from the distribution of existing AA or BA (including those with higher than a BA), conditioning on one of 12 age-race-sex cells. The *causal parameter* method assigns a causal estimate of the marginal AA or BA returns using parameters from the existing literature: high school graduates who are assigned an AA receive a 29 percent annual earnings increase, and high school graduates who are treated with a BA a 68 percent annual earnings increase.

In both the distribution and causal parameter method, we further adjust earnings for the relative wage effect that is likely to result from an increase in the share of the population with a college degree. To incorporate this relative wage response into our simulation exercise, we follow the common paradigm in the academic literature, as described in Autor and Acemoglu (2011), and specify a two-factor CES production function model.

The results of this simulation exercise reveal that a sizable increase in rates of college attainment would meaningfully increase earnings of individuals near the bottom of the earnings distribution relative to those near the top. For instance, while the 90/10 ratio increased from 4.63 to 5.45 between 1979 and 2018, simulation 3 (increasing AA rates to 20 percent and BA rates to 60 percent) would bring that ratio down to 5.16 or 5.00, depending on the method, reversing from more than half of the actual increase over this period. This reduction stems from increases in the 10th percentile of FTFY earnings and smaller proportional change at the 90th percentile. However, increased college attainment will have minimal effects on reducing overall inequality back to the levels in 1979, as a greater share of the population with college degrees will not meaningfully affect earnings at the highest parts of the distribution, where much of the rise in inequality has taken place.

We also examine the effect of increasing educational attainment on economic security, as captured by several poverty measures. Simulations applying the distribution method imply that increasing the BA share to 60 percent and the AA share to 20 percent would reduce the poverty rate by 2.39 percentage points, from 11.3 to 8.91 percent in the sample using all civilian adults age 25 to 54. Reductions in the near-poverty or low-income rate (family income less than 150 and 200 percent of the poverty threshold, respectively) are larger, with the first falling from 18.5 to 14.2 percent, and the second falling from 26.5 to 20.4 percent. Both of these simulated rates are lower than their actual levels in 1979.

In this way, the policy prescription of increased educational attainment should appeal to those whose primary concern is the economic security of poorer individuals, but it will not satisfy the goals of those whose primary concern is the reduction of overall income inequality or income shares at the top of the distribution.

In all, this dissertation explores government interventions intended to improve economic efficiency or reduce inequality. First, I find that while the Paycheck

Protection Program was likely on-net effective at preventing job loss and maintaining efficient employee-firm relationships, but its time-limited nature mitigated the employment impact of this program. Second, looking at the federal housing voucher program, I find that between four and six additional vouchers prevent one eviction in a given city, with the savings from these avoided evictions covering about one quarter of the costs of the voucher program each year. Finally, in my fourth chapter, I find that increasing educational attainment can improve earnings at the bottom of the distribution, bringing up measures of income inequality between the 90th percentile and the bottom half, but will do little to meaningfully reduce overall income inequality because these trends are driven by the upper tail of the earnings distribution.

Chapter 2: Small Business Experiences as the Paycheck Protection Program Ends: Evidence from Covered Period Expiration

2.1 Introduction

As the COVID-19 pandemic unfolded in the spring of 2020, causing a significant drop in economic activity, Congress created the Paycheck Protection Program (PPP) as one part of a nearly unprecedented fiscal stimulus effort. This program was developed as a forgivable loan, backed by the federal government, to help small businesses stay afloat and maintain employment as the U.S economy was temporarily shut down in order to contain the novel coronavirus. Through the initial program's expiration on August 8, 2020, the federal government has guaranteed $525 billion in loans to businesses through PPP (Treasury Department 2020).

Given the scale of the program and its unique features, as compared to past stimulus efforts, identifying the effectiveness of PPP at maintaining firm employment is an important question that has received increasing attention. Previous economic research has used data on loan applications along with pre-determined business size cutoffs to estimate the effectiveness of this program among PPP-eligible and -ineligible firms (Hubbard and Strain 2020, Chetty et al. 2020, Autor et al. 2020). They find small, and often imprecise, positive effects of this program on employment, firm financial health, and survival.

This research differs in several respects. First, while prior work has examined business hiring patterns as PPP was rolled out in the spring and summer, this work examines these trends as PPP ends and companies roll off of the requirements necessary to receive loan forgiveness. Second, the dataset used here includes detailed information on the terms of businesses' PPP loans, including, most importantly, the loan disbursement date and the length of time the company was required to maintain payrolls. No research to date using either public or private data contains this detail. Finally, while much research focuses on the effect of PPP on small businesses near the 500-employee-size eligibility, I focus on the experiences of companies much smaller in size who may face even greater liquidity concerns than larger businesses.

I exploit detailed payroll microdata, with information on the date of loan disbursement and the length of time the company must maintain employment levels, to estimate the effect of the Paycheck Protection Program on employment. I find that, as the PPP Covered Period expired, companies reduced active employment by a statistically significant 0.41 percent per week. Event study estimates over the four weeks after expiration indicate that firms reduced headcount by 1.6% the month after expiration. I find smaller, insignificant results in counties that have managed to contain coronavirus case growth, suggesting that revenue replacement policies such as PPP can be effective at maintaining employment (and preserving the value of worker-firm relationships) given that external conditions improve as well.

This work has important implications for the design of future payroll support programs. These PPP headcount requirements were a binding constraint on firm employment behavior, indicating that many businesses emerged from their PPP experience in no better shape than when they received the loan -- largely because public health measures did not improve. Many firms rolled off these requirements as the pandemic continued and the economy was at a particularly fragile point, suggesting that flexibility in the loan terms, such as tying expiration dates to a measure of public health conditions, is warranted.

2.2 Background, Data, and Methodology

2.2.1 The Paycheck Protection Program as a Policy Response

From April 3, 2020 to August 8, 2020 all small businesses, generally defined as firms with fewer than 500 employees, were able to apply for loans through the Paycheck Protection Program.[1] Businesses were eligible to apply for loans covering up to eight weeks of payroll costs through private lenders, who then submitted the loan for approval to the Small Business Administration. These loans are up to 100% forgivable, given companies adhere to specific guidelines set out by the SBA. Companies must spend at least 60% of the loan on payroll costs, and the remainder may only be spent on eligible expenses including rent, mortgage interest, and utilities.

Additionally, and importantly for this piece, in order to receive full loan forgiveness, the loan recipient must maintain a specified average level of employment over the length of the "covered period" (SBA 2020b). The Covered Period is the 8- or 24-week period beginning on the PPP loan disbursement date.[2] With few exceptions, the required employment level the firm must maintain is either relative to

[1]Small businesses are generally those with fewer than 500 employees, though exact thresholds vary by industry. Additionally, these size thresholds for PPP eligibility purposes were extended for businesses with NAICS 2-digit industry code 72 (hotels and restaurants).

[2]A borrower may choose an 8-week covered period only if he or she received the loan before June 5, 2020.

employment over the period from February 15, 2019, to June 30, 2019 or any 12-week period between May and September 2019 for seasonal employers (ADP 2020).

This program was implemented as one piece in a broad suite of policies enacted by Congress in the spring to provide relief during the initial stages of an economic downturn that was both sharp and deep, but also forecast to be a temporary decline from which the economy would rapidly recover. Specifically, PPP was designed as a tool to replace the revenue of small businesses, who do not typically hold large savings buffers or have easy access to capital markets, in order to preserve valuable worker-firm specific capital and prevent costly small business closures during a temporary public health emergency.

Thus, PPP is a novel fiscal policy in the U.S. in at least two ways: first, worker relief has almost exclusively been provided through unemployment benefits, which replace a portion of a worker's wages after he or she has been terminated. This structure, however, allows for the destruction of any worker-firm specific capital (e.g. knowledge of specific businesses procedures, or within-firm relationships) when the employee is terminated. If, during a temporary downturn, employee-employer ties can be maintained through a revenue replacement program, it may be cost-beneficial for the government to provide such relief.

Additionally, PPP is a novel policy in that it has essentially provided grants to small businesses in order to replace lost revenue, when policymakers have historically relied on firms to build savings buffers and markets to reallocate resources efficiently during business downturns. Policymakers justified this program for small businesses on the basis that these firms have limited ability to build savings buffers (and we would not expect them to hold savings for a global pandemic[3]) and lack the access to liquid capital markets available to larger businesses. This temporary revenue replacement was thus designed to avoid costly firm closures among a set of businesses for which moral hazard is not a concern and market-based solutions may not exist.

Prior work has largely found that PPP loans were effective at increasing employment and improving other measures of firm health, although this remains an evolving area of research. Autor et al. (2020), in closely related work, use payroll microdata from private processor ADP to estimate that companies just below the size eligibility cutoffs boosted net employment by 3.25% relative to companies slightly above (and thus ineligible) as PPP rolled out in the spring.

Chetty et al. (2020) finds similar results using county-level data from private financial management applications: they estimate a modest, though imprecisely-estimated, employment effect of about 2 percentage points. Hubbard and Strain (2020a), using firm-level data from Dun & Bradstreet, finding more modest effects of PPP on employment (although they only have data for application of loans over $150,000). They also find PPP had a positive effect on a measure of firm financial health and reduced the likelihood of business closure.

Although PPP loans were effective at increasing employment and avoiding layoffs and firm closure, the program was not well-targeted to small businesses, due

[3]The absence of pandemic insurance markers is interesting, though perhaps not surprising given market failures in insurance markets for other low-probability high-cost events, such as flood/natural disaster insurance.

to information frictions, a delay in processing small business applications, and a lower likelihood of approval (Neilson et al. 2020). This is especially troubling given that survey data suggests that small businesses were in a particularly vulnerable position at the onset of the pandemic: the median small business had fewer than two weeks of cash on hand (Bartik et al. 2020). In this way, timely evidence on small business experiences as the PPP program ends and the economy is in an only modestly-improved position is particularly important.

2.2.2 Data

I use data from Gusto, a payroll and benefits platform for small businesses, to measure firm-level hiring and termination rates as companies cross the 8- or 24-week period during which they were required to maintain employment levels relevant for loan forgiveness.

In Gusto's Hiring & Termination dataset, a hire corresponds to an employer creating a new employment with a hiring date, and a termination corresponds to the entry of a termination date for a given employment. Layoffs corresponded to terminations where the employer listed the reason for the termination as a "layoff" (one of the choices from a standardized list in Gusto's termination flow). In order to capture a more time-sensitive view of employer activities, I record hires, terminations, and layoffs on the date that they were entered into the system, rather than their effective date.

I connect this payroll microdata with information entered by small business owners in Gusto's "PPP Forgiveness Tracker," an application offered to the small business who received PPP loans, which aids them in keeping track of funds. Importantly, this tracker includes the date each PPP loan was disbursed and the length of the company's covered period.

Table 2.1 presents summary statistics of firms tracking PPP loans through the Forgiveness Tracker as of the first week of March 2020. These firms are very small businesses with an average of 7.8 employees at this point in time, though there is somewhat significant variation. The 99th percentile of firm size in this dataset is 52 employees. The average loan size is just about $80,000, lower than the average loan size of $101,000 across all PPP loans (Treasury Department 2020). This average loan amount from the SBA is likely skewed by a small number of large companies not present in Gusto's database, where 68.6% of all PPP loans were below $50,000 and 81.7% were below $100,000.

Examining the experience of such very small businesses is crucial to any comprehensive assessment of the Paycheck Protection Program, and is a new contribution of this paper. Prior analysis of the employment effects looks largely across the 500-employee eligibility threshold, and these relatively larger firms are likely to have much different experiences with PPP. These small businesses also represent a large part of the program. Using loan-level data released by the SBA, I estimate that 83% of all recipients employed fewer than 20 workers, and these companies received roughly one-third (34%) of the total loan amount (SBA 2020c).[4]

[4]This data contains limited information on each loan recipient's self-reported industry, size, location, and loan amount. It does not include data on the length of the recipient's covered period.

The average loan size among firms with fewer than 20 employees was just over $45,800 -- smaller than Gusto's average amount of $80,000.

Appendix Table 2A.1 compares the sectoral composition of Gusto's PPP recipients (who use the Forgiveness Tracker) to the universe of PPP recipients in the loan-level dataset. Because many firms in Gusto's dataset received loans below $150,000, this table breaks out the composition for businesses receiving loans less than $150,000 (as identified in the SBA dataset) and then all businesses. The distribution of PPP recipients across industries in Gusto's data is largely in line with the universe of PPP loans (both <$150,000 and all loans), with one main exception. One fifth of PPP loans went to Gusto customers whose NAICS code is not classifiable. which is a result of how industries are assigned in Gusto's data. Customers self-assign themselves to a series of subsectors, including "Technology," which does not exactly align with a specific NAICS-classified sector.

2.2.3 Methodology

I formalize the empirical approach in a standard two-way fixed-effects difference-in-difference framework, where treatment timing is identified as the first full week after a company's covered period ends. The main estimating equation is thus of the form:

$$y_{it} = \alpha + \delta_t + \gamma_i + \beta * expired_{it} + \varepsilon_{it} \qquad (2.1)$$

Where y_{it} is the outcome of interest for company *i* in week *t*. I include a set of week fixed effects (δ_t) and company fixed effects (γ_i), and *expired* indicates the covered period has ended. The coefficient of interest, β, can be interpreted as the change in the outcome variable (hiring rate, for instance) after a company's covered period ends, relative to the weeks prior to expiration. A causal interpretation of this coefficient relies on the identifying assumption that no contemporaneous events occur that might affect a company's hiring and termination decisions (e.g. policies, changes in the economic outlook, etc).

An encouraging piece of evidence is the relatively broad distribution of covered period expiration dates. Appendix Figure A.1 plots the portion of companies whose covered period ends for each week along the year. There are two masses of companies, determined by the length of the covered period assigned. A large group of companies' covered period ended around the beginning of July, which would correspond to getting loans at the beginning of April and having an 8-week covered period. The covered period for another mass of PPP loans ends in October and November, again consistent with an April-May loan date but a 24-week covered period.[5] Were these companies' expiration dates concentrated in a single week or set of weeks, I may be concerned that seasonal employment patterns confound the variable of interest.

This analysis also relies on the assumption that the headcount requirement was enforced or at least that firms believed it would be enforced. Importantly, any non-compliance -- either because firms did not believe this rule would be enforced or because firms still could not afford to keep employees on payroll for eight or twenty-four weeks after the loan was disbursed -- will make it more difficult to find an effect

[5]I also examine how these effects differ across firms with 8- and 24-week covered periods, presented in Appendix Table 2A.2.

using these methods. This analysis measures firms' reactions in the five weeks immediately after the covered period expires, any actions the firms take outside this treatment window will either bias the effect downwards (if, say, firms terminate employees before the covered period ends) or will not be included (if firms terminate employees after five weeks post-covered period).

I also estimate effects separately by week, in an event-study design, estimating the following equation:

$$y_{it} = \alpha + \delta_t + \gamma_i + \sum_{d=-5}^{5} \beta_d * D(t = d) + \epsilon_{it} \qquad (2.2)$$

where D is an indicator variable for a given number of weeks relative to the expiration of the covered period. Figure 8 plots the estimated leads and lags coefficients, as well as their associated standard errors. Coefficients are estimated relative to "Week 0," the week containing the end of the covered period.

2.3 Results

2.3.1 Main Regression Results

I formally estimate the effect of the expiration of the covered period using the regression framework described above. Table 2.2 presents results of the regression pooled across all companies, and then separately for companies in industries hardest-hit by the pandemic, and all others.[6] Standard errors are clustered at the company

[6]"Hard-hit" industries include NAICS codes 44, 71, and 72, and 81 (Retail Trade, Arts & Entertainment, Leisure & Hospitality, and Other Personal Services).

level to allow for an arbitrary error variance-covariance matrix within each company (Bertrand, Duflo, and Mullainathan 2003).

The first column of results displays the estimates pooled across all sectors. The main coefficient of interest is presented in Panel A, estimating the effect of the expiration of the covered period on the net hiring rate. I estimate that as companies rolled off their covered period, they reduced hiring by a statistically significant 0.41 percentage points. Off of a pre-treatment baseline mean of 0.22%, this change represents a 200% decrease in the net hiring rate. This decline is due to a slight drop in hiring (of a statistically insignificant 0.05 points), but largely to a spike in terminations (of a statistically significant 0.25 percentage points) as companies do not need to maintain headcount levels to receive PPP loan forgiveness.

The second column presents estimates among firms in industries hardest-hit by the pandemic: those in retail, food and beverage establishments, and other sectors that traditionally rely on face-to-face interaction. While the reduction in sample size leads to less precision (only about 17% of the sample is in these hard-hit industries) the point estimate is slightly larger in magnitude at -0.53. On the other hand, as the covered period expires, firms in all other industries reduce headcount by a smaller amount, -0.39 percentage points, as presented in the final column. Furthermore, looking at the termination rates in Panel C, there is a larger, and statistically significant, spike in terminations among hard hit industries. These point estimates, while often imprecisely estimated, provide suggestive evidence headcount losses associated with the end of PPP are larger in firms within hard-hit industries that have continued to struggle.

Figure 1 plots the results for these three groups from the event study design, where these effects are estimated by week. The estimates plotted in this figure can be found in Appendix Table 2A.3. The omitted time period is Week -1, the week immediately prior to the first full week in which the covered period had expired. In the full sample, net hiring rates are near zero and statistically insignificant for the 5 weeks prior to expiration. These effects drop to a statistically significant -0.334 percentage points (relative to week -1) the week of expiration and continue to show net headcount decline in the following weeks. Taking the four weeks after expiration together (weeks 0 to 3), these results suggest that in the month after PPP recipients' covered periods expired, net employment fell by 1.6 percent.

The point estimates within the hart-hit industries sample suggests both a steeper and more prolonged employment drop among these companies. The first week after expiration, hiring falls by 0.89 percentage points, and remains at this low level for the remainder of the sample period, while net hiring among companies in other sectors drops by 0.21 percentage points in the first week and moves back toward zero by the end of the sample.

2.3.2 Effects Across More- and Less-Recovered Areas

One possible explanation for this drop in employment as forgiveness requirements expire is the simple fact that the pandemic, and resulting economic weakness, was ongoing as firms rolled off their covered period. PPP was designed as a support measure to help firms replace a very temporary drop in revenue, but

COVID-related economic weakness extended far longer than provided for by these loans.

I examine the role that the continued coronavirus outbreak played in driving firm behavior by exploiting differences in the evolution of local health conditions between the time a firm received a PPP loan and when their covered period expired. If local conditions played an important role in determining whether firms were able to re-open and return to prior revenue levels, I would expect firms located in areas that fared relatively better along this measure to have been able to maintain employment as PPP requirements expired. There is good reason to believe that local public health conditions play an important role in local economic health, particularly in contrast to legal restrictions such as business closures or stay-at-home orders. Goolsbee and Syverson (2021) find that public health conditions, as measured by county-level COVID death rates, explain a far greater share of consumer traffic patterns than legal restrictions alone.

To measure local health conditions, I use county daily COVID-19 cases, assembled by the New York Times (New York Times 2020). I calculate the percentage change in average weekly cases within each firm's county between the week that firm received its loan and the week its covered period expired. There are few counties where Coronavirus cases declined over the course of the 8- or 24-week periods in this dataset. Indeed, among the roughly 1,000 counties represented in this dataset, median case growth between loan disbursement and expiration was 643%. I examine firm experiences across three groups: the bottom 10%, where COVID-19 cases increased by less than 104%; the top 10%, in which cases great by 1700% or more; and the middle 80% of firms.

Table 2.3 presents regression estimates of the effect of covered period expiration on the net hiring rate for all firms in these three groups. There was slight attrition in assigning county identifiers to the firms in this sample, and the first column presents the analysis for the full (reduced) sample of all firms to which COVID case levels at loan disbursement and expiration were assigned. The overall effect across the full sample, of -0.359 is close to the full sample estimate in Table 2.1 of -0.414 percentage points.

There is a monotonic increase in the effect of covered period expiration as county case growth rises, from the second through fourth columns. Within the 10% of firms experiencing relatively little case growth, there is no significant effect of covered period expiration, perhaps because consumers became less sensitive to these (relatively low) cases levels, resumed economic activity, and firms were able to meet payroll expenses with business revenue. The results for firms in the middle 80% of case growth closely matched the overall sample. However, firms in counties with the highest case growth dramatically reduced their headcount as the covered period expired, by -0.640 percentage points, though only significant at the 10 percent level. These results lend credence to the hypotheses that much of the employment loss following the end of PPP-related headcount requirements is due to the continuation of the health emergency and ensuing economic weakness.

Figure 2 presents the results of the corresponding event-study estimates across these three mutually exclusive groups (lowest 10%, middle 80% and top 10% of case growth). These numbers are displayed in Appendix Table 2A.4. In the first full week of expiration, net hiring fell by 0.47 percentage points among companies in counties that experienced large increases in COVID-19 cases. Among companies where cases stayed relatively low, the point estimate in the first week is in fact positive. By the end of the sample period, the point estimate in the high-growth sample remained low at -0.77 percentage points, suggesting companies were still reducing headcount, whereas net hiring in fact rose by 0.17 percentage points in the best-performing counties during the first treatment week and generally remained near-zero through the end of the sample.

To put these findings in the context of the larger universe of PPP recipients, I can conduct a similar exercise for firms in the SBA loan-level dataset. For each firm included, I calculate the weekly average case growth eight or twenty-four weeks after the firm's loan approval data. This analysis differs from what I am able to perform using Gusto data because the SBA data does not provide the firm's county or the length of the covered period. I examine *state-level* case growth (instead of county) and I assign firms 8 or 24-week covered periods to match the observed distribution of covered period length in Gusto's data. First, I assign 24-week covered periods to all firms who received loans on or after the week of June 5, 2020 (in line with the program rules) and then randomly assign 55% of firms who received loans before then to to the 24-week group in order[7]. Results do not materially change as I vary the number of 8-week recipients who move to a 24-week period.

[7]This step is meant to simulate the number of firms that received loans before the longer covered period was introduced and "switched" into the 24-week covered period when they were able to do so.

Overall, COVID experiences among firms in this PPP dataset look roughly similar to those described above. The bottom ten percent of firms experienced a 116% rise in state-level weekly cases across the covered period, compared to 104% in Gusto's data, while the top 10% experienced a 2,434% rise in cases, higher than the 1,700% above. This pattern suggests that, solely based on how public health conditions evolved, termination patterns could have been even more dramatic among the universe of PPP recipients.

2.4 Discussion

2.4.1 Aggregate Effects

Finally, I use these findings to infer aggregate employment effects of the expiration of these PPP headcount requirements among businesses who received PPP loans and provide an updated estimate of PPP's cost-per-job-saved.

I employ this measure as developed by Autor et al. (2020), with the caveat that cost-per-job saved is by no means a sufficient statistic by which to evaluate PPP. First, current cost-per-job saved estimates include relatively short- and medium- term effects of this program, and the effects of this program depend on the long-term evolution of employment among these firms. Second, as pointed out in Hubbard and Strain (2020a), there are many costs and benefits of this program not included in such a measure, including but not limited to the value of worker-firm specific capital that has been preserved, the (social and economic) value of preventing widespread small

---

As noted in an earlier section, about 75% of firms in the Gusto dataset switched to the longer covered period.

business liquidation, and the value of this program in speeding up the economic recovery. Any holistic assessment of this program must take those factors into account.

I compute aggregate employment effects by taking the preferred treatment estimate across all firms and scaling it by an estimate of the number of workers who worked at firms taking PPP loans. I follow the formula

$$Employment\ Effects = \beta * Takeup\,\% * N_t \tag{2.3}$$

Where β is the treatment effect estimated above: firms' net headcount fell by 1.6 percent in the four weeks after the covered period expired. *Takeup %* is an estimate of the portion of small businesses who received PPP loans, and N_T is an estimate of the number of employees who work at small businesses.

Take-up is calculated from dividing the numbers of small businesses who received a PPP loan by the count of all such firms, obtained from the SBA loan level database and the Census Bureau's *Statistics of U.S. Businesses* (SUSB).[8] Take-up among small businesses was quite high, at an estimated 81%. Autor et. al (2020) estimate that at the beginning of 2020, 70 million employees worked for a firm which was eligible for PPP.

With these estimates in hand, I calculate that the end of the covered period reduced aggregate employment by (0.016*0.81*70,000,000 =) 907,200 jobs within the four weeks following the covered period expiration.

[8]To calculate this take-up rate, I focus on firms within the size range in which the treatment effect was estimated: firms with fewer than 50 employees. There are 6% of firms with missing size values in the dataset, which I attribute based on the distribution of non-missing size values.

Autor et al. (2020) estimate that PPP saved approximately 2.3 million jobs, at a cost of $518 billion,[9] for a cost of $224,000 per job supported. My findings suggest that as PPP ended, these firms eliminated 907,200 jobs. Thus, on net, PPP preserved 1.4 million jobs, at a net cost of roughly $370,000 per job saved.

Two caveats are warranted: First, I am applying a *local treatment effect* estimated across a subset of PPP recipients: small businesses generally employing fewer than 50 people. Larger firms facing fewer information barriers may be more aware of the exact terms of PPP loans and thus more strongly react to the covered period ending, leading to larger treatment effects for this group. Alternatively, it is possible that larger firms were better able to recover from the recession during the summer and spring and thus were not inclined to increase terminations as the covered period expired.

2.4.2 Has the Paycheck Protection Program Succeeded?

This work finds that small businesses receiving PPP loans quickly and significantly cut headcount, beginning the week that loan forgiveness requirements expire. These effects are largest within certain sectors hardest-hit by the ongoing pandemic and concentrated in counties that have fared the worst in containing the spread of the coronavirus. It is still difficult, however, to make broad conclusions about the success or failure of both the Paycheck Protection Program and payroll-side relief policy more generally for a few reasons. First, this pandemic and the relief

[9]After the latest draft of Autor et al. (2020) was released, Treasury and SBA released final data on the total PPP loan amount of $525 billion (US Treasury Department 2020). For consistency between estimates, I retain the $518 billion figure here.

policies enacted at its onset are still in the relatively early stages. As a policy designed to maintain the relationships between workers and firms, and to prevent wasteful firm closure, the Paycheck Protection Program ought to be judged on the long-run value of these worker-firm relationships it preserved and firm closures that it prevented. Research on both of these questions is in early stages, and it is too early to make confident conclusions.

Furthermore, these findings are greatly context-specific to the setting that the United States economy finds itself in currently. PPP was designed by policymakers as a temporary revenue-replacement program, providing payroll support under the assumptions that businesses could be "put on ice" until the novel coronavirus quickly receded and economic activity returned to normal levels. Clearly, neither of those assumptions have proven to be the case. Indeed, if one clear lesson can be drawn from this experience thus far, it is that data-dependent policies that can adapt as conditions evolve are preferable to hard-and-fast time limits based on initial predictions. For instance, economic relief during a public health emergency could be tied to public health indicators such as virus case counts or vaccine distribution. As this initial work on the expiration of PPP has found, the positive initial results of economic relief may be temporary if the assumptions under which this aid was designed do not materialize.[10]

[10]Glenn Hubbard and Michael Strain have made a similar point in arguing that the expiration of a second round of PPP should be tied to a measure of widespread vaccine distribution (Hubbard and Strain 2020b).

2.5 Conclusion

In this paper, I examine the effect of the expiration of one key component of the paycheck protection program: the requirement that firms maintain average headcount levels for an extended period of time. I find that firms significantly reduce their net hiring -- and increase terminations -- just as they are released from requirements as a part of PPP loan forgiveness. These effects are particularly pronounced (though sometimes less precisely estimated) for certain firms in subsectors hit particularly hard by the pandemic. I find that firms located in areas that fared relatively better in containing the coronavirus did not significantly reduce headcount as requirements expire, lending credence to the hypothesis that this behavior was driven by firms who continued to struggle as the pandemic dragged on.

Extrapolating these findings to a larger universe of PPP recipients yields an estimate that over 900,000 jobs were lost as firms rolled off their requirements for PPP forgiveness. These findings provide empirical evidence that struggling firms were quite responsive to the release of their constraints under the terms of PPP loan forgiveness, and that contrary to the aim of PPP, these loans were not able to provide enough financial support for them to come out of the assistance period in a stable economic position, largely because public health conditions had not materially improved as these loans expired.

2.6 Tables and Figures

Table 2.1: Summary Statistics of Gusto Payroll Database

	Mean	Std. Dev
Firm Size (No. employees)	7.8	11.0
PPP Loan Amount ($)	$79,999	$140,697
Portion with 8-week Covered Period	43.66%	
Portion with 24-week Covered Period	56.34%	
Number of Firms	37, 316	

Sources: Gusto, LLC.
Notes: Payroll data is calculated at the firm-level. PPP loan terms are self-reported in the PPP Forgiveness Tracker. Number of employees measures active employees on a firm's payroll.

Table 2.2: Impact of the Covered Period Expiration on Firm Employment (Percentage Points)

	All Companies	Hard-Hit Industries	All Others
A. Net Hiring Rate	-0.414***	-0.525	-0.389***
	(0.145)	(0.518)	(0.137)
Pre-Treatment Mean:	0.221%	0.226%	0.220%
B. Gross Hiring Rate	-0.049	0.131	-0.087
	(0.060)	(0.230)	(0.053)
Pre-Treatment Mean:	1.198%	1.790%	1.071%
C. Gross Termination Rate	0.246***	0.756***	0.135**
	(0.059)	(0.233)	(0.066)
Pre-Treatment Mean:	0.977%	1.564%	0.851%
Number of Observations	304,915	53,773	252,142

*** = p<0.01, ** = p<0.05, * = p<0.1
Sources: Gusto, LLC
Notes: Payroll data is calculated at the firm-level PPP loan terms are self-reported in the PPP Forgiveness Tracker. Number of employees measures active employees on a firm's payroll. Standard errors are clustered at the firm level. "Hard-Hit Industries" includes firms with NAICS codes 44, 71, and 72, and 81 (Retail Trade, Arts & Entertainment, Leisure & Hospitality, and Other Personal Services).

Table 2.3: The Impact of the Covered Period Expiration on Firm Employment, by Change in Local COVID-19 Cases

	All	Lowest 10%	Middle 80%	Highest 10%
Net Hiring Rate	-0.359**	-0.144	-0.363*	-0.640*
	(0.161)	(0.378)	(0.192)	(0.337)
Pre-Treatment Mean	0.138	0.501	0.192	0.080
Number of Observations	227,965	22,720	180,618	24,627

*** = p<0.01, ** = p<0.05, * = p<0.1
Sources: Gusto, LLC, COVID-19 case counts from The New York Times (2020)
Notes: Payroll data is calculated at the firm-level. PPP loan terms are self-reported in the PPP Forgiveness Tracker. Number of employees measures active employees on a firm's payroll. Pre-treatment means are calculated for the five weeks prior to expiration. Standard errors are clustered at the firm level. Local COVID-19 cases are calculated as change in weekly average cases in that firm's county between the week each firm received its PPP loan and the week the covered period expired.

Figure 2.1: Effect of Expiration of Covered Period on Net Hiring Rate: "Hard-Hit" Industries and Others

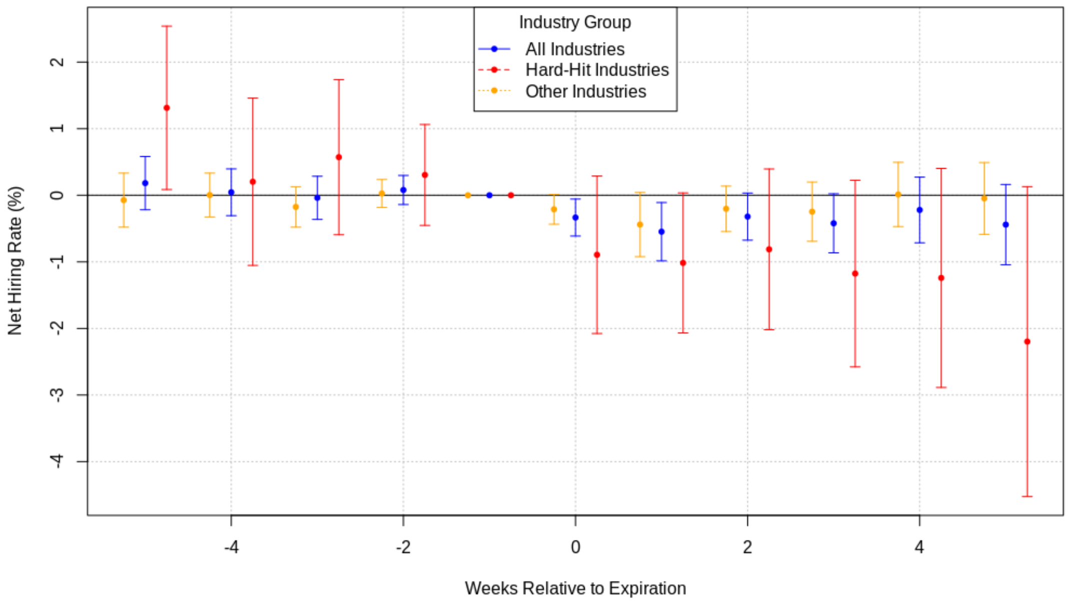

Source: Author's calculations of Gusto data.
This figure plots the weekly estimates of the effect of a firm's covered period expiration on it's net hiring rate. Estimates are calculated relative to the week prior to expiration (Week = -1). Standard errors are clustered at the firm level. "Hard-Hit Industries" includes firms with NAICS codes 44, 71, and 72, and 81 (Retail Trade, Arts & Entertainment, Leisure & Hospitality, and Other Personal Services).

Figure 2.2: Effect of Expiration of Covered Period on Net Hiring Rate: Local County COVID-19 Case Growth

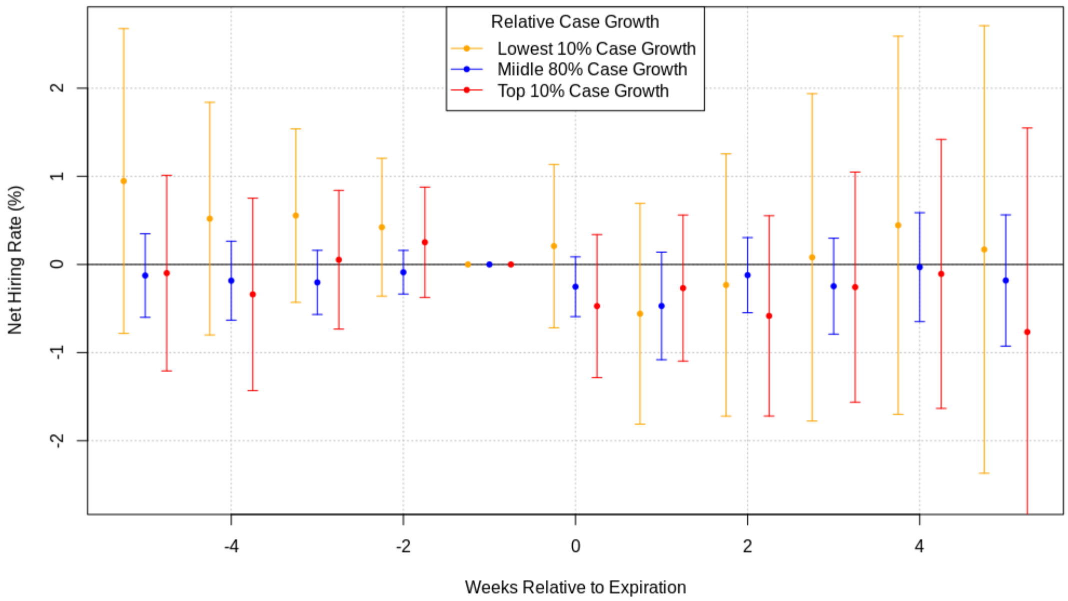

Source: Author's calculations of Gusto data.
This figure plots the weekly estimates of the effect of a firm's covered period expiration on it's net hiring rate. Estimates are calculated relative to the week prior to expiration (Week = -1). Standard errors are clustered at the firm level. Local COVID-19 cases are calculated as change in weekly average cases in that firm's county between the week each firm received its PPP loan and the week the covered period expired.

2.7 Appendix Materials

Table 2A.1: Distribution of PPP Funds Across Sectors

NAICS Code	Industry	Gusto	SBA Small Loans	SBA All Loans
11	Agriculture, Forestry, Fishing, and Hunting	0.23	2.34	1.54
21	Natural Resource Extraction	0.21	0.43	0.86
22	Utilities	0.27	0.18	0.28
23	Construction	3.82	10.18	12.48
31	Manufacturing	7.06	5.02	10.41
42	Wholesale Trade	2.82	3.72	5.34
44	Retail Trade	5.47	8.79	7.69
48	Transportation and Warehousing	1.18	3.30	3.35
51	Information	2.21	1.35	1.78
52	Finance and Insurance	3.76	3.46	2.29
53	Real Estate and Rental and Leasing	2.35	4.15	2.97
54	Professional Services	14.76	12.58	12.75
56	Administrative and Support and Waste Management Services	1.10	4.77	5.03
61	Educational Services	2.77	1.43	2.31
62	Health Care and Social Assistance	10.30	12.47	12.94
71	Arts, Entertainment, and Recreation	3.21	2.18	1.56
72	Accommodations and Food Services	4.24	9.53	8.03
81	Other services	4.77	9.72	5.99
92	Public Administration	0.40	0.29	0.34
99	Unknown/Unclassifiable	20.79	3.91	1.78

Sources: Gusto, LLC and Small Business Administration (2020c).
Notes: "SBA Small Loans" includes all loans under $150,000. NAICS codes are assigned to businesses using Gusto software by sector affiliation chosen by businesses as they join the platform. NAICS code 51 ("information") in Gusto's platform users who assigned themselves into the "Technology" subsector, which does not exactly line up with any particular NAICS industrial classification.

Table 2A.2: The Impact of the Covered Period Expiration on Firm Employment, by Length of Covered Period (Percentage Points)

	(1) 8-Week	(2) 24-Week
A. Net Hiring Rate	-0.429***	-0.267**
	(0.143)	(0.150)
Pre-Treatment Mean	0.153	-0.0680
B. Gross Hiring Rate	-0.054	-0.089
	(0.063)	(0.084)
Pre-Treatment Mean	1.198	1.130
C. Gross Termination Rate	0.216***	0.043
	(0.072)	(0.096)
Pre-Treatment Mean	0.977	1.124
Number of Observations	286,736	180,180

*** = p<0.01, ** = p<0.05, * = p<0.1
Sources: Gusto, LLC
Notes: Payroll data is calculated at the firm-level. PPP loan terms are self-reported in the PPP Forgiveness Tracker. Number of employees measures active employees on a firm's payroll. Pre-treatment means are calculated for the Standard errors are clustered at the firm level.

Table 2A.3: Event Study: Effects of Covered Period Expiration, by Change in Local COVID-19 Cases

Weeks Prior to (-) or After (+) Expiration	Lowest 10%	Middle 80%	Highest 10%
-5	0.947 (0.882)	0.125 (0.242)	-0.098 (0.566)
-4	0.520 (0.673)	-0.183 (0.228)	-0.340 (0.557)
-3	0.555 (0.502)	-0.204 (0.186)	0.053 (0.402)
-2	0.423 (0.399)	-0.088 (0.126)	0.252 (0.319)
0	0.209 (0.473)	-0.253 (0.173)	-0.472 (0.414)
1	-0.559 (0.639)	-0.471 (0.311)	-0.268 (0.423)
2	-0.233 (0.759)	-0.121 (0.218)	-0.584 (0.580)
3	0.081 (0.948)	-0.247 (0.278)	-0.257 (0.666)
4	0.445 (1.095)	-0.030 (0.315)	-0.107 (0.779)
5	0.170 (1.296)	-0.182 (0.380)	-0.766 (1.181)
Number of Observations	36,306	288,801	39,064

*** = p<0.01, ** = p<0.05, * = p<0.1
Sources: Gusto, LLC
Notes: Payroll data is calculated at the firm-level PPP loan terms are self-reported in the PPP Forgiveness Tracker. Number of employees measures active employees on a firm's payroll. Standard errors are clustered at the firm level. Local COVID-19 cases are calculated as change in weekly average cases in that firm's county between the week each firm received its PPP loan and the week the covered period expired.

Figure 2A.1: Distribution of Firms, by End Date of Covered

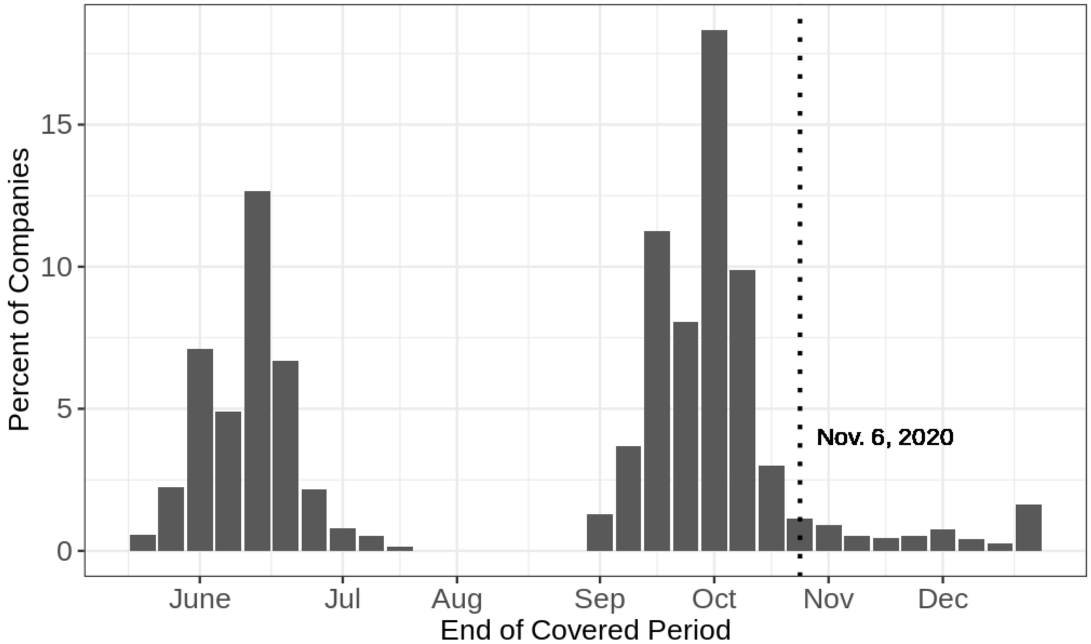

Source: Author's analysis of Gusto data.
"Covered Period" refers to the 8- or 24- week period (chosen by the firm) beginning at the date of loan disburseement over which firms were required to maintain pre-pandemic employment levels. Dates of covered period expiration are calculated from data on loan disbursement date and covered period length that firms entered in Gusto's PPP Forgiveness Tracker. This data tracks firms whose covered period expired on or before November 6, 2020.

Chapter 3: The Effect of Low-Income Rental Subsidies on Evictions: Evidence from the Housing Voucher Program

3.1 Introduction

In 2018, the federal government spent $50 billion dollars on low-income housing assistance programs, an amount comparable to the budget for food stamps (SNAP) or the Supplemental Security Income program (OMB, 2019). Research surrounding these rental assistance programs often focuses on their effects on participants' labor market outcomes and on their use as tools to explore "neighborhood effects" (Jacob and Ludwig, 2012; Kling et al., 2007; Chetty et al., 2016; Bergman et al., 2019). There has been limited work, however, on the impact of these programs on one important feature of a low-income renter's life: housing instability.

One particularly painful source of housing instability is eviction, the forcible removal of tenants from their rental unit. Each year, over 2 million households, or 5 percent of the renting population, face the threat of eviction from their housing unit (Desmond et al., 2018b). The causes and effects of evictions have been the subject of research among sociologists and public health researchers for a number of years, who find associations between these forced relocations and adverse economic outcomes, deterioration in physical and mental health, and longer-term residential instability (Desmond and Kimbro, 2015; Hartman and Robinson, 2003; Crane and Warnes, 2000). Recent economic research, with stronger causal identification, finds little effect of evictions on later economic outcomes, but does consistently find significant, long-lasting impacts on evictees' physical and mental health and subsequent likelihood of experiencing homelessness (Collinson and Reed, 2018; Humphries et al., 2019). Avoiding eviction thus delivers benefits both to affected individuals, in the form of improved health and housing stability, and public service providers who bear the cost of increased hospitalizations and use of homeless shelters.

I explore the effect of the largest form of low-income rental assistance, the federal Housing Voucher Choice program, on a county's eviction rate by combining a newly-available dataset compiling court records of eviction cases across the country with data on the avail- ability and use of housing vouchers in each area. To avoid potential endogeneity issues, I exploit plausibly exogenous historical variation in one part of the voucher allocation formula: the stock of a county's rental housing built before 1940. I use this data as of 2000, during the last substantial expansion of the housing voucher program.

This instrument is strongly predictive of the county's level of housing vouchers in 2016, and data on the stock of rental units built immediately after 1940 has no predictive power on voucher levels. Furthermore, this instrument is uncorrelated with measures of community decline or attitudes towards housing within the area.

I find that, across a number of specifications, every four to six vouchers prevent one eviction in a county. These effects are larger in markets where renters pay a larger share of income on rent, suggesting that vouchers play an important role in reducing cost burdens, making rent affordable for low-income families. There is also a larger effect in areas with greater demand for these vouchers (as measured by time on the voucher waiting list), meaning a targeted expansion of the program could have a greater impact on preventing evictions. A back-of-the-envelope calculation, using the estimated of cost of an eviction of $8,000 from Collinson and Reed (2018), suggests that as much as a quarter of the average annual cost of a voucher can be recovered from savings through eviction prevention alone.

3.2 Literature Review

Much of the research around the housing voucher program centers on its short- and long- term economic effects in terms of labor supply and earnings. While in the short-run, Jacob and Ludwig (2012) find that income limits and large subsidy phase-out rates cause voucher recipients to reduce their labor supply, recent long-term research of families randomly offered housing vouchers finds large inter-generational economic effects, with children of voucher recipients having significantly higher household incomes, higher rates of marriage, and improved educational outcomes 10-20 years later (Kling et al., 2007; Chetty et al., 2016). These long-term benefits intensified among a group of recipients who were restricted to move into low-poverty neighborhoods.

There has been little research, however, into whether these vouchers improve housing stability for low-income renters – a more direct aim of this program. One closely-related study evaluated HUD's "Welfare to Work" voucher program, which randomly assigned vouchers to families in six cities in 1999. This work found that vouchers significantly reduced the probability a family experienced homelessness or moved apartments in sub- sequent years (Wood et al., 2018). This work aims first to provide an updated estimate of the effectiveness of housing vouchers in improving housing outcomes, focusing on data from 2016, which will account for recent dynamics in low-income rental housing markets. Second, I use data both at the county and MSA level, which allows me to account for market-level spillover effects, which are one main justification for in-kind transfers such as housing vouchers. Finally, the outcome of focus here is an area's eviction rate - a measure of housing stability not yet examined in the voucher literature and one that is receiving a growing amount of national attention.

Work by sociologist Matthew Desmond suggests that evictions, while often caused by negative economic shocks, can themselves be the cause of spirals into economic distress, as renters struggle to find new housing while maintaining jobs and caring for family members (Desmond, 2016). A large body of work in sociology and public health find negative associations of evictions with a tenant's physical and mental health, increasing depression and the risk of suicide (Desmond and Kimbro, 2015; Fowler et al., 2015; Sandel et al., 2018). These findings are often descriptive, based on comparisons of renters who faced eviction and those who did not, and lack strong causal identification. Recent and ongoing work in economics, exploiting random assignment of judges to a tenant's eviction case, calls into question the causal link between evictions and material hardship (Humphries et al., 2019; Collinson and Reed, 2018). Both find that evictions cause a very slight decrease in earnings, particularly small in comparison to the large drop in income preceding eviction. Collinson et al. (2014), however, finds that evictions cause a lasting increase in housing instability, increasing the likelihood that individuals experience homelessness up to several years after the eviction judgment. Furthermore, they find that an eviction causes significant increases in mental health hospitalizations and emergency room utilization.

Research has also explored the effectiveness of policies meant to prevent evictions and the negative consequences that follow. Evans et al. (2016) finds that one-time transfers from Chicago's emergency cash assistance program to residents facing difficulty making a rent payment reduces the likelihood of applicants experiencing homelessness by 88%. Zewde et al. (2019) finds that the increased financial stability among low-income populations caused by the Affordable Care Act's Medicaid insurance expansion significantly decreased the evictions rate in states that took advantage of this expansion.

Taken together, such evidence suggests that housing instability is caused by acute negative economic shocks, such as job loss or an unexpected increase in medical costs, that prevent households from making rent payments. Housing vouchers are a likely candidate to prevent evictions, since they both directly reduce rent burdens through subsidies and also reduce the likelihood that negative economic shocks will make rent unaffordable through the scaling of tenant rental contributions to a set portion of income. At the market level, however, the incidence of such subsidies and the possibility of negative spillovers mean that it is not clear ex-ante whether vouchers will reduce an areas eviction rate. A brief conceptual model that highlights the main mechanisms through which vouchers may affect a market's eviction rate follows a review of the key institutional details of the housing voucher program and the eviction process.

3.3 Background and Institutional Features

3.3.1 Housing Vouchers

The program now known as Housing Choice Vouchers was created by Congress in Section 8 of the Housing and Community Development Act of 1974. Motivated by the high cost of constructing subsidized housing units directly, as well as the concentration of poverty in existing housing projects, Congress began to fund housing vouchers (Collinson et al., 2014). These vouchers allow tenants to lease units directly from landlords in the private rental market, with the government then paying a portion of the rent directly to the landlord. This program is funded by the federal government but administered by local authorities. The Department of Housing and Urban Development provides the funds to local Public Housing Authorities (PHAs), which then allocate the vouchers among the eligible population. As of 2018, the Housing Choice Voucher program was the largest rental assistance program in the U.S, serving roughly 2.3 million households at a cost of $20 billion dollars (HUD, 2019b,c).

In 1998, Congress established the current eligibility guidelines for the voucher program, which largely restricts assistance to families with incomes below half of the area's median income. Congress also stipulated that at least 75% of vouchers must be allocated to families earning less than 30% of the median income. In practice, the number of families eligible for vouchers vastly exceeds the number of vouchers available, and Public Housing Authorities maintain waitlists that can last for years. Fischer and Sard (2013) estimates that only one- quarter of eligible families receive

housing rental assistance, including vouchers. Currently, the average household receiving a voucher earns of 22% of the area median income (HUD, 2019c).
After receiving a voucher, households are able to lease any unit available on the private rental market, subject to minimum safety standards and the restriction that rent cannot initially exceed 40% of their income. Recipients pay 30% of rent and the government pays the remainder, up to a payment standard established at the metropolitan level – generally around half of an area's Fair Market Rent, as defined by HUD.

Importantly, the supply of housing vouchers last significantly expanded between 1998 and the 2003, increasing from about 1.5 million vouchers available to 2.1 million, as Congress funded an additional 600,000 vouchers (GAO, 2006). Since then, and through at least 2016, Congress has allocated no new vouchers for a community's general population, with the relatively small number of new vouchers funded from demolished public housing units or earmarked for homeless veterans and foster youths at risk of homelessness as they age out of the foster system (Couch, 2016).

3.3.1 Evictions

Eviction is a formal legal process, initiated by a property owner, in which a court orders the tenants of a rental unit to vacate that unit. The eviction process is determined at the local (usually county) level, and cases are heard in civil court (a small claims court or special housing court in large cities where these cases are more prevalent).

The process varies from county to county, but generally follows four steps. First, the landlord serves a tenant with an eviction notice. Legal justifications for eviction include a substantial violation of the lease or using the unit for illegal purposes, but the large majority of evictions are for non-payment of rent (Desmond et al., 2018a). Tenants in most cities then have a specified time period to resolve complaints listed on the eviction notice. For instance, in Chicago, non-payment of rent begins a five-day period in which the tenant may repay the sum owed (Humphries et al., 2019).

Should the period lapse without resolution, the landlord may then file an eviction case in court, and a hearing date is set. The landlord can file a single "cause of action," seeking solely to evict the tenant(s) from the unit, or multiple causes of action (a "joint action") in which she seeks a money judgment for unpaid rent or other damages (Desmond, 2016). On the date of the hearing, the judge assigned to the case or a court attorney typically attempts to work out an agreement between the tenant and landlord in order to avoid an eviction judgment and blemish the tenant's record. If no agreement is reach, the judge can either render an eviction judgment, ordering the tenant to vacate the unit, or she can dismiss the case. If a settlement is reached, the landlord and tenant must adhere to the terms of the agreement. If, at a later point in time, the judge finds that tenant is not meeting the terms, an eviction judgment is rendered; if the terms are met, the case is dismissed.

The final step of eviction is the execution of the eviction order. After an eviction judgment, the landlord files an eviction order with the Sheriff's office or City Marshal, who then removes the tenant's possessions and changes the locks.

3.3.2 Conceptual Model

A simple conceptual model of the low-income housing market can help inform the interpretation of the results presented. In this market, low-income renters demand units of a given quality and property owners supply these rental units. Housing vouchers are rental subsidies, which pay a portion of that household's rent (up to a limit). In this way, the main mechanism through which vouchers might affect evictions is by making rent more affordable for subsidized families, thereby reducing the likelihood of evictions. Vouchers are also in-kind subsidies, which under standard economic theory is meant to increase the consumption of housing because positive externalities lead the individually-optimal level of housing consumption to be below the socially-optimal level. These positive spillovers might also act to lower evictions within a market. For instance, low-income households might be doubling-up with families or friends, in violations of those tenants' lease. The voucher that allows the household to rent their own unit also reduces the probability the original tenants are evicted for violating the terms of their lease. Work by Currie and Yelowitz (2000) finds that vouchers reduce exactly this type of overcrowding.

The subsidies that vouchers provide, however, may also have incidence properties that spill over into non-voucher renters within the low-income rental market. Voucher recipients compete for rental units in the same market as non-voucher, low-income households. An introduction or expansion of vouchers in a market might raise rents – as individuals increase their consumption of housing or as landlords seek to capture these subsidies – making prices unaffordable for non-voucher households and increasing this population's probability of eviction. Recent research into an expansion of voucher generosity suggests that property owners of low-income rental units capture a large portion of these subsidies in the form of higher rents (Collinson and Ganong, 2018). An initial study by Susin (2002) found that these rent increases do in fact spill over into the non-voucher rental market, but more recent research with more credible identification finds that vouchers have no detectable effect on the price of low-income rental units within a city (Eriksen and Ross, 2015). Nevertheless, at least in theory, there are forces at work in even this simple conceptual model of a low- income housing market that may cause vouchers to increase or decrease total evictions within a city.

3.4 Data

To study the effect of a housing vouchers on eviction rates, I compile a cross-section of county- and metropolitan-level data on housing vouchers, eviction judgments, and demo- graphic variables from the year 2016, along with data on the area's level of housing built before 1940 as of 2000.

I use data on both counties and Metropolitan Statistical Areas (MSA's) as a complimentary approach in order to ensure enough variation in the data, but also examine effects close to the area encompassing a housing market. There are 2,197 counties in the sample for which evictions, voucher, and demographic data are available. Metropolitan Statistical Areas, on the other hand, define "socially and economically integrated areas" (US Census Bureau, 2000). These areas better define rental housing markets, the area in which individuals may move after receiving a voucher, and may better account for the spillover effects in the rental market due to vouchers. There are 207 Metropolitan areas in the sample, however, which may mask much of the variation in both the voucher and eviction data.

3.4.1 Housing Voucher Data

Data on housing vouchers is accessed from the Department of Housing and Urban Development's Picture of Subsidized Households (HUD, 2019c). This dataset contains information on both the total available and total occupied subsidized rental housing units, at varying levels of geographies and at several points in time over 1990 to 2018. Data is available for all of HUD's programs, including public housing and housing vouchers. I gather data on occupied housing units in the Housing Choice Voucher program (also known as Section 8 Vouchers and Certificates) at the county and MSA level for the year 2016, since that is the most recent year for which evictions data is available.

3.4.2 Eviction Data

Area-level eviction counts come from a publicly-available dataset assembled by the Eviction Lab at Princeton University (Desmond et al., 2018b). This dataset contains the universe of court-ordered evictions that occurred between 2000 and 2016, aggregated from individual court records up to various geographic levels. I download the county-level eviction count from 2016, the most recent year available. In this dataset, evictions are defined as having an eviction judgment rendered against the tenant (and the judge not later vacating that judgment). This method mirrors that in Humphries et al. (2019), but differs from Collinson and Reed (2018), which measures the effect of executed evictions – that is, the final step of forcible eviction by a City Marshal. An eviction order may not be executed if the tenant moves out willingly after the eviction order is rendered. Thus, the measure used here is a broader measure of forcible relocation than those only carried out by Sheriffs. I also aggregate this eviction data to the Metropolitan Statistical Area.

3.4.3 Demographic Data

I merge data on housing voucher utilization and eviction judgments with data on the age distribution of rental housing units from the 2000 Decennial Census and demographic data from the 2011-2016 American Community Survey 5-year Summary File, both gathered from the IPUMS National Historic Geographic Information Systems (Manson et al., 2019). I calculate the number of rental housing units built before 1940 as of the year 2000 from the Decennial Census data. Demographic and economic controls come from the 2011-2016 ACS five-year averages. I use the five-year ACS files, instead of the 2016 ACS one-year file, to take advantage of a larger sample size: all counties are available in the five-year file, while only the 352 largest counties are available in the 2016 one-year file. Appendix A presents results of this analysis using data from the one-year file, with the smaller sample size. The results are quantitatively similar and no significance levels are changed.

3.4.4 Data Summary

Table 3.1 presents summary statistics of major variables at both the county and MSA level. The statistics are generally similar across the geographic definitions, although there are fewer evictions and less aged rental housing when grouped by metropolitan areas. On average, in 2016, there was about one voucher for every one hundred people in a county or MSA, and one eviction for every 300 people.

I also include two measures of the relative need for vouchers in a given area: the percent of renter households cost-burdened (spending more than 30.4% of their income on rent) and the average time current voucher holders spent on the voucher waitlist. Both statistics paint a dim picture of the state of rental housing and rental assistance. On average, between 57% and 60% of renters spent more than a 30.4% of their income on rent in 2016, and the average length of time voucher holders spent on waitlists before receiving a voucher was between 31 months and over three years.

3.5 Empirical Strategy, Results, and Analysis

3.5.1 Empirical Strategy

Estimating the effects of vouchers on area eviction rates in the cross section requires variation in the level of housing vouchers that is unrelated to factors also influencing that area's level of tenant evictions. This requirement poses a challenge because a county's level of housing vouchers is likely strongly related to the economic circumstances of the renter population – that is, a county using a relatively large number of housing vouchers is also likely economically disadvantaged in other ways that contribute to higher eviction rates. Intuitively, and as Desmond and Gershenson (2016) confirms, the concentration of low-income renters is highly correlated with an area's eviction rate. Even after controlling for observable factors – such as an area's race composition and income distribution – there are likely factors unobservable to the researcher that influence both the level of housing vouchers and the prevalence of evictions within an area. If the effect of housing vouchers on evictions is negative, these endogeneity concerns will bias an OLS estimate towards, or potentially above, zero, leading to the conclusion that increases in housing vouchers increase evictions.

Table 3.2 presents results of such a regression. Column (1) simply regresses a county's level of housing vouchers per capita on its level of evictions per capita, producing a statistically significant, positive point estimate that each additional housing voucher increases a county's eviction rate. Column (2) includes controls for income, race, and age, and state fixed effects, which reduce this point estimate, but it remains positive and highly significant. Columns (3) and (4) present regressions at the MSA-level, but estimates are noisy and statistically insignificant from zero. Moreover, in these specifications, there still may be unobserved factors influencing both the level of housing vouchers and evictions per capita.

I argue that I can use one aspect of the process for allocating voucher funding across cities as an instrumental variable to isolate variation in housing vouchers that is orthogonal to factors also influencing a city's eviction rate. The Housing and Community Development Act of 1974 stipulated that newly-funded housing vouchers be allocated according to each area's low-income housing needs. These needs are determined by a weighted combination of six factors: (1) the renter population (20 percent); (2) the number of renter households with income below poverty (20 percent); (3) housing overcrowding, defined as the number of housing units with an occupancy ratio of 1.01 or more persons per room (10 percent); (4) housing vacancies (10 percent); (5) substandard housing, defined as the number of housing units built before 1940 and occupied by renter households with annual incomes at or below the poverty level (20 percent); (6) Other objectively measurable conditions (20 percent) (Code of Federal Regulations, 2019).

Of these factors, I use a portion of factor (5), meant to measure a city's degree of substandard housing, as an instrument for a city's level of housing vouchers. Although the level of pre-1940 housing units occupied by renters in poverty is likely endogenous to evictions, I find that the overall share of a city's housing built before 1940 is a relevant predictor of housing vouchers and is plausibly orthogonal to factors influencing the prevalence of evictions. Furthermore, I use the level of rental housing built before 1940 as of 2000 to predict the level of housing vouchers in 2016. I do this first to ease concerns that aged housing in 2016 may not be orthogonal to a city's eviction rate in 2016: while a city's current level of aged housing may be related to its attitudes toward rental housing (and in turn its eviction rate), the level of aged housing 16 years prior, when the large majority of vouchers were being allocated, is less plausibly related to current policies towards low- income renters. As noted earlier, by 2002, roughly 2.1 million of the 2.3 million vouchers in 2016 were already allocated, and the balance have not been made available to the general renter population. This aspect of the funding formula has been used once in the past as an instrument to examine the effect of housing vouchers on market-level outcomes: Sinai and Waldfogel (2005) use this instrument when gauging the targeting properties of subsidized housing.

One additional concern is that other government programs may use this statistic in determining funding levels for other assistance programs. If that is the case, this instrument might be picking up the collective effects of these programs instead of the voucher program alone. The Community Development Block Grant (CDBG) program does use the level of pre-1940 housing to determine funding for only a portion of recipient cities - so-called "Formula B" cities, who qualify for more funding than they would under the traditional "Formula A" (HUD, 2019a). Data on the level of CDBG funding is only available at the place level, which are smaller areas encompassing politically incorporated jurisdictions. In Appendix B, I conduct this same analysis at the place-level, and then include the level of CDBG funding as a control[11]. The results are unaffected by the inclusion of controls for this grant funding.

My empirical strategy estimates regressions of the form

$$EvictionsPerCapita_{c,s,2016} = \beta_0 + \beta_1 * VouchersPerCapita_{c,s,2016} + \beta_2 * X_{c,s} + \gamma_s + \epsilon_{c,s,2016} \tag{3.1}$$

Instrumenting for *VouchersPerCapita* with the level of pre-1940 rental housing as of 2000 in the first stage:

$$VouchersPerCapita_{c,s,2016} = \alpha_0 + \alpha_1 * Pre1940RentalUnitsPerCap_{c,s,2016} + \alpha_2 * X_{c,s} + \gamma_s + \mu_{c,s,2016} \tag{3.2}$$

Where c indexes the county in state s, and $X_{c,s}$ is a vector of observable demographic covariates - controlling for a county's race, income, and age distribution, the percent of individuals in an area who are married, and per capita public assistance income in each city. Furthermore, when this analysis is conducted at

[11] I am thankful to George Zuo for kindly providing data on the level of CDBG funding as of 2016.

the county level, I am able to include state-level fixed effects. I also conduct this analysis at the MSA level, which fully encompass a low-income rental housing market, and thus account for market-level spillovers discussed above.

To serve as a valid instrument, the level of pre-1940 housing must be both relevant and not associated with community factors that directly affect a city's level of evictions. I examine each requirement below.

3.5.2 Relevance

Table 3.3 presents the first stage regression of pre-1940 rental housing per capita on housing vouchers per capita, along with the controls described above, at the county and MSA level. An area's level of rental housing built before 1940 as of 2000 is strongly related to its level of housing vouchers. One additional pre-1940 rental unit per capita predicts a relatively precisely estimated 0.1 additional vouchers per capita. This coefficient is significant at the 99% confidence level. The first stage F-statistic in the preferred county-level specification with state fixed effects is 181.03, above the threshold of concern for a weak instrument (Stock and Yogo, 2012).

I also conduct a placebo test to ensure that this part of the formula, and not some other factor related to older housing, is driving the variation in voucher funding. I run this first stage analysis with the level of rental housing per capita build between 1940 and 1960 (as of 2000), the next age group available in Census data, as the left-hand-side variable, instead of pre-1940 housing per capita housing. These estimates are presented in columns (2) and (4) of Table 3.3. I find, while pre-1940 housing is a strong predictor of a city's level of vouchers, housing built between 1940 and 1960 is not significant at the 95% level at either the MSA- or county-level. While the point estimate at the county level is marginally significant, it is also an order of magnitude smaller than those in the “true” first stage and negative, meaning additional 1940-1960 housing decreases the county’s level of vouchers.

3.5.3 Exclusion

A greater concern with this instrumental variable approach is that this part of the formula, intended as a proxy for a city’s substandard rental housing does capture some factors of “community decline” or local attitudes towards housing, that directly affect eviction levels. Should this pre-1940 housing data be correlated with unobserved city factors that also cause evictions, I would be attributing spurious correlation to a causal effect in the second stage. While it is not possible to fully “prove” the exclusion restriction, I present results of several regression of pre-1940 housing on different proxies for these factors.

To test whether an increased share of older rental housing is related to observable measures of community decline or housing attitudes, Table 3.4 presents the main coefficient from separate regressions of pre-1940 housing per capita on two of variables that might serve as proxies for such a measure: housing units identified as having any of several substandard conditions[12], and new residential construction permits[13] in a county or the metropolitan area.

[12]Substandard housing conditions are defined in the ACS as units lacking complete kitchen or bathroom facilities (US Census Bureau, 2015).
[13]New building permits are merged in from the Census Bureau’s Building Permits Survey (US Census Bureau, 2019).

Table 3.4 presents the result of regressions of pre-1940 housing levels on these outcomes. No estimates are significant at the 90% confidence level: an area's level of pre-1940 rental housing has no significant relationship to the share of housing units with substandard conditions, As shown in columns (1) and (3). Columns (2) and (4) present estimates for the level of new building permits in an area, with similarly small and insignificant estimates. The lack of any identifiable relationships between these variables and the share of pre-1940 rental housing lends credence to the argument that this measure is plausibly exogenous to factors affecting a city's eviction rate and thus is a valid instrument.

3.5.4 Results

Table 3.5 presents results of the instrumental variable regressions. The results in columns (1) and (2) present coefficients estimated at the county level with and without state fixed effects. These estimates imply that roughly every four to six vouchers in a given city prevent one eviction (formally, increasing the number of vouchers-per-capita by one reduces the number of evictions per capita by 0.13 to 0.25). The county-level estimates are significant at the 99% confidence level.

The MSA-level estimates are presented in column (3). The point estimate is quite similar to the county-level estimates, but the large drop in the number of observations (from 2197 to 207) leads to a statistically insignificant coefficient.

Collinson and Reed (2018) make a back of the envelope calculation that the cost of an eviction totals roughly $8,000 within two years. In 2016, the average monthly rent subsidy of a housing voucher was $700, or $8,400 annually HUD

(2019c). Using the upper bound of the above estimates, that every four vouchers prevent one eviction, the cost savings from evictions alone can account for almost a quarter of the annual cost of the voucher program[14].

3.5.5 Analysis

To explore the mechanisms driving these results, I use two measures of the economic status of an area's renter population: the portion of cost-burdened renter households and the average time current voucher holders spent on the waitlist for their voucher[15].

First, I calculate the percent of cost-burdened renter households in each area, and divide the sample into those with rates of cost-burdened renters above the median and below the median (the median at the county-level is 57% and 58% at the MSA-level). Results in Table 3.6 present IV estimates for counties (column 1 and 2) and MSA's (columns 3 and 4). While the MSA results are again imprecisely estimated, in both cases, there is a large gap between these estimates in the low cost-burdened and high cost-burdened samples, with as few as two additional vouchers in high rent-burdened counties preventing every one eviction. These results suggest that the point estimates in the full sample are largely driven by areas which have a larger share of rent-burdened households, in turn providing evidence for the idea that vouchers can prevent evictions by simply making rent more affordable for these highly-burdened households.

---

[14]Four vouchers cost $700*12*4=$33,600 annually, while preventing one eviction costing a conservatively- estimated $8,000.

[15]Data on cost burdens is collected through the ACS, while data on average voucher wait length times is published alongside voucher utilization data in HUD's Picture of Subsidized Households (HUD, 2019c).

In Table 3.7, I examine differences in effects between cities that have longer and shorter waitlists for housing vouchers, as a measure of the demand for these resources relative to available supply. At both the county and MSA level there are larger and statistically significant effects in areas with longer waitlists (above the median of 25 months at the county level and 28 months at the MSA level), with the point estimate at the MSA level suggesting that in severely resource-constrained cities, voucher may prevent evictions on an almost one-for-one basis. In this way, expanding vouchers in areas with the most rent- burdened households, or in cities where there is the greatest need relative to current supply, may provide the largest eviction prevention effects.

3.6 Conclusion

In this paper, I examine the effect of housing vouchers on market-level eviction rates, exploiting a plausibly exogenous factor determining the current level of a city's housing vouchers. I find that, at the mean, between six and four additional vouchers prevent one eviction in a given city. Furthermore, the effectiveness of housing vouchers in preventing evictions is much larger in cities with a larger share of rent-burdened households and in cities where there is a larger demand for vouchers relative to supply. These results suggest that vouchers play an important role in preventing evictions by making rent more affordable, and targeting additional supply to areas most in need will have the largest effect.

3.7 Tables and Figures

Table 3.1: Summary Statistics

	County		MSA	
	Mean	Std. Dev.	Mean	Std. Dev.
Vouchers Per Capita	0.011	(0.007)	0.011	(0.004)
Evictions Per Capita	0.003	(0.003)	0.003	(0.003)
Percent Age 65+	0.144	(0.039)	0.142	(0.034)
Percent African American	0.130	(0.128)	0.123	(0.090)
Percent Native American	0.007	(0.024)	0.006	(0.008)
Percent Asian	0.052	(0.059)	0.062	(0.063)
Percent Other Race	0.079	(0.063)	0.086	(0.060)
Percent Married	0.367	(0.053)	0.364	(0.030)
Percent of Cost-Burdened Households	0.573	(0.065)	0.595	(0.051)
Per Capita Public Assistance Income	32.509	(19.567)	33.883	(18.936)
Percent of Cost-Burdened Households	0.573	(0.065)	0.595	(0.051)
Average Months Voucher Holder Spent on Waitlist	31.325	(97.929)	36.917	(104.486)
Observations	2197		207	

Notes: Statistics presented are for the year 2016, weighted by population. Evictions per capita measures eviction judgments, dated by time of case filing, per total individuals, made publicly available by the Eviction Lab (Desmond et al., 2018b). Vouchers per capita measures the number of vouchers in use, in a given area provided by (HUD, 2019c). All other data are from the 2011-2016 ACS 5-year summary file (Manson et al., 2019).

Table 3.2: OLS: Effects of Housing Vouchers Per Capita on Evictions Per Capita

	County		MSA	
	(1)	(2)	(3)	(4)
	No Controls	With Controls	No Controls	With Controls
Vouchers Per Capita	0.159***	0.088***	-0.070	0.023
	(0.010)	(0.012)	(0.057)	(0.060)
Percent African American		0.007***		0.015***
		(0.001)		(0.004)
Percent Native American		-0.009***		0.066***
		(0.002)		(0.021)
Percent Asian		0.002		0.008
		(0.002)		(0.009)
Percent Other Race		0.009***		-0.005
		(0.001)		(0.007)
Percent Married		-0.018***		-0.023
		(0.003)		(0.017)
Observations	2197	2197	207	207
Adjusted R^2	0.11	0.66	0.00	0.55
State F.E.	No	Yes	No	No
Other Controls	Yes	Yes	Yes	Yes
Observation Unit	County	County	MSA	MSA

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

Notes: Regressions are weighted by population. Other covariates include per-capita public assistance income and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories. "1940-1960 Rental Units" are rental housing units build between 1940 and 1960 as of 2000.

Table 3.3: First Stage: The Effect of Pre-1940 Rental Housing on Vouchers

	County		MSA	
	(1)	(2)	(3)	(4)
	First Stage	Placebo	First Stage	Placebo
Pre-1940 Rental Units Per Capita in 2000	0.112***		0.136***	
	(0.029)		(0.020)	
1940-1960 Rental Units Per Capita in 2000	-0.006*			0.026
	(0.003)			(0.320)
Percent African American	0.001	0.003	-0.006*	-0.007*
	(0.003)	(0.003)	(0.003)	(0.004)
Percent Native American	-0.025***	-0.025***	0.022	-0.016
	(0.003)	(0.003)	(0.026)	(0.035)
Percent Asian	0.001	0.005	-0.004	-0.018**
	(0.009)	(0.008)	(0.007)	(0.007)
Percent Other Race	0.022*	0.028**	-0.004	0.005
	(0.013)	(0.012)	(0.009)	(0.011)
Percent Married	-0.051***	-0.050***	-0.034	-0.021
	(0.009)	(0.009)	(0.022)	(0.022)
Observations	2197	2197	207	207
Adjusted R^2	0.75	0.73	0.68	0.57
State F.E.	Yes	Yes	No	No
Other Controls	Yes	Yes	Yes	Yes
Observation Unit	County	County	MSA	MSA
First Stage F-Statistic	181.03		56.85	

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

Notes: Regressions are weighted by population. Other covariates include per-capita public assistance income and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories. "1940-1960 Rental Units" are rental housing units build between 1940 and 1960 as of 2000.

Table 3.4: Exclusion Restriction: Pre-1940 Rental Housing and Measures of Community Decline

	County		MSA	
	(1) Substandard Condition	(2) Building Permits	(3) Substandard Condition	(4) Building Permits
Pre-1940 Rental Units Per Capita in 2000	0.097	-0.046	-0.025	0.003
	(0.070)	(0.030)	(0.103)	(0.056)
Percent African America	0.038***	-0.005	-0.031**	-0.025**
	(0.008)	(0.005)	(0.014)	(0.011)
Percent Native American	-0.045***	-0.009	-0.080	-0.074
	(0.010)	(0.008)	(0.073)	(0.062)
Percent Asian	-0.032*	-0.014	-0.071**	-0.051*
	(0.019)	(0.011)	(0.029)	(0.028)
Percent Other Race	0.099***	-0.015*	0.150***	-0.016
	(0.018)	(0.008)	(0.023)	(0.021)
Percent Married	-0.124***	0.019	-0.229***	0.041
	(0.027)	(0.019)	(0.082)	(0.057)
Observations	2197	2197	207	207
Adjusted R^2	0.82	0.28	0.85	0.25
State F.E.	Yes	Yes	No	No
Other Controls	Yes	Yes	Yes	Yes
Observation Unit	County	County	MSA	MSA

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

Notes: Regressions are weighted by population. Standard errors are in parentheses. Other covariates include public income assistance per capita, and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories. New residential construction is defined as multifamily housing permits, gathered from the Census Bureau's Building Permits Survey (US Census Bureau, 2019). Substandard conditions measures the lack of complete bath, or kitchen facilities in the unit (US Census Bureau, 2015).

Table 3.5: IV Estimates: Effect of Vouchers Per Capita on Evictions Per Capita

	County		MSA
	(1)	(2)	(3)
	No State F.E.	State F.E.	No State F.E.
Vouchers Per Capita	-0.133***	-0.243***	-0.170
	(0.030)	(0.049)	(0.118)
Percent African American	0.011***	0.008***	0.014***
	(0.001)	(0.001)	(0.003)
Percent Native American	-0.006**	-0.017***	0.063***
	(0.003)	(0.003)	(0.022)
Percent Asian	-0.012***	0.004**	0.005
	(0.002)	(0.002)	(0.008)
Percent Other Race	0.007***	0.017***	-0.004
	(0.002)	(0.002)	(0.006)
Percent Married	-0.015***	-0.035***	-0.027*
	(0.004)	(0.004)	(0.016)
Observations	2197	2197	207
R^2	0.39	0.55	0.60
State. F.E.	No	Yes	No
Other Controls	Yes	Yes	Yes
Observation Unit	County	County	MSA
First Stage F-Statistic	585.02	181.03	56.85

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

Notes: Regressions are weighted by population. Other covariates include per-capita public assistance income and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories.

Table 3.6: Effect of Vouchers, by Market's Median Rent Burden

	County		MSA	
	(1)	(2)	(3)	(4)
	Low Burden	High Burden	Low Burden	High Burden
Vouchers Per Capita	-0.157***	0.530***	0.126	-0.316
	-			
	(0.055)	(0.132)	(0.356)	(0.227)
Observations	1817	380	127	80
Adjusted R^2	0.53	0.49	0.53	0.50
State F.E.	Yes	Yes	No	No
Other Controls	Yes	Yes	Yes	Yes
Observation Unit	County	County	MSA	MSA

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

Notes: Regressions are weighted by population. Standard errors are in parentheses. Other covariates include race distribution, public income assistance per capita, and the percent of the population married, and percent in each of 4 race categories, 25 income categories, and 12 age categories. Rent burdened households spend more than 30.4% of household income on rent. Low and high rent-burdened areas are divided at the sample median percent of rent-burdened households. In the population-weighted median county (MSA), 57.4% (58.8%) percent of renting households are cost-burdened. High and low burdened groups within an area may not be equally sized since median is population-weighted.

Table 3.7: Effect of Vouchers, by Average Time on Voucher Waitlist

	County		MSA	
	(1)	(2)	(3)	(4)
	Short Waitlist	Long Waitlist	Short Waitlist	Long Waitlist
Vouchers Per Capita	-0.148***	-0.237**	0.130	-0.847**
	(0.049)	(0.095)	(0.145)	(0.330)
Observations	1674	523	140	67
Adjusted R^2	0.58	0.67	0.49	0.27
State F.E.	Yes	Yes	No	No
Other Controls	Yes	Yes	Yes	Yes
Observation Unit	County	County	MSA	MSA

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

Notes: Regressions are weighted by population. Standard errors are in parentheses. Other covariates include race distribution, public income assistance per capita, and the percent of the population married, and percent in each of 4 race categories, 25 income categories, and 12 age categories. Data on average waitlist lengths for current voucher holders at the county and MSA level are published in HUD's *Picture of Subsidized Households* (HUD, 2019c). Short and long waitlist areas are divided at the sample median. Voucher holders in the population-weighted median county (MSA) waited 25 (28) months on the voucher waitlist. Short and long waitlist groups may not be equally-sized since median is population-weighted.

3.8 Appendix A: Analysis with 2016 1-year ACS

Table 3A.1: First Stage: The Effect of Pre-1940 Rental Housing on Vouchers Per Capita

	(1) First Stage	(2) Placebo
Pre-1940 Rental Units Per Capita in 2000	0.102***	
	(0.039)	
1940-1960 Rental Units Per Capita in 2000		-0.009*
		(0.005)
Observations	352	352
Adjusted R^2	0.76	0.75
State F.E.	Yes	Yes
Other Controls	Yes	Yes
Observation Unit	County	County

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

Notes: Regressions are weighted by population. Standard errors are in parentheses. Other covariates include race distribution, public income assistance per capita, and the percent of the population married, and percent in each of 4 race categories, 25 income categories, and 12 age categories. “1940-1960 Rental Units” are rental housing units build between 1940 and 1960 as of 2000.

Table 3A.2: Exclusion Restriction: Pre-1940 Rental Housing and Measures of Community Decline

	(1) Substandard Condition	(2) Permits
Pre-1940 Rental Units Per Capita in 2000	-0.054	-0.058
	(0.103)	(0.051)
Percent African American	0.050***	0.005
	(0.013)	(0.010)
Percent Native American	-0.017	0.003
	(0.028)	(0.013)
Percent Asian	-0.059***	-0.009
	(0.021)	(0.015)
Percent Other Race	0.042**	-0.018
	(0.019)	(0.011)
Percent Married	-0.062	0.038
	(0.043)	(0.033)
Observations	352	352
Adjusted R^2	0.84	0.34
Other Controls	Yes	Yes
State F.E.	Yes	No
Observation Unit	County	County

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

Notes: Regressions are weighted by population. Standard errors are in parentheses. Other covariates include public income assistance per capita, and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories. New residential construction is defined as multifamily housing permits, gathered from the Census Bureau's Building Permits Survey (US Census Bureau, 2019). Substandard conditions measures the lack of complete bath, or kitchen facilities in the unit (US Census Bureau, 2015).

Table 3A.3: OLS and IV Estimates: Effect of Vouchers Per Capita on Evictions Per Capita

	(1)	(2)
	No State F.E.	State F.E.
Vouchers Per Capita	-0.182**	-0.328**
	(0.092)	(0.142)
Percent African American	0.012***	0.010***
	(0.003)	(0.003)
Percent Native American	-0.011	-0.026***
	(0.007)	(0.007)
Percent Asian	-0.009*	0.006
	(0.005)	(0.004)
Percent Other Race	0.003	0.012
	(0.008)	(0.012)
Percent Married	-0.013	-0.036***
	(0.013)	(0.013)
Observations	352	352
R^2	0.47	0.66
Other Controls	Yes	Yes
State F.E.	No	Yes
Observation Unit	County	County
First Stage F-Statistic	91.6	21.9

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

Notes: Regressions are weighted by population. Other covariates include per-capita public assistance income and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories.

3.9 Appendix B: Analysis at Place Level, Controlling for CDBG Funding

Table 3B.1: Effect of Vouchers Per Capita on Evictions Per Capita, Controlling for CDBG Funding

	(1) Without CDBG Funding Control	(2) With CDBG Funding Control
Vouchers Per Capita	-0.083***	-0.076***
	(0.018)	(0.016)
Percent African American	0.009***	0.009***
	(0.000)	(0.000)
Percent Native American	-0.008***	-0.007***
	(0.003)	(0.003)
Percent Asian	0.002***	0.002***
	(0.001)	(0.001)
Percent Other Race	-0.003***	-0.003***
	(0.001)	(0.001)
Percent Married	-0.017***	-0.017***
	(0.002)	(0.002)
Observations	7946	7946
Adjusted R^2	0.57	0.57
State F.E.	Yes	Yes
Other Controls	Yes	Yes
Observation Unit	County	County
First Stage F-Statistic	578.71	671.17

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

Notes: Regressions are weighted by population. Other covariates include per-capita public assistance income and the percent of the population in each of 4 race categories, 25 income categories, and 12 age categories. Data on the level CDBG funding was provided by a program representative at the Department of Housing and Urban Development.

Chapter 4: College Attainment, Income Inequality, and Economic Security: A Simulation Exercise

4.1 Introduction

College-educated workers today have much higher levels of earnings, income, and employment than those without college degrees, with especially large premiums awarded to those who hold a bachelor's degree or higher. As documented by numerous studies, the relative employment and earnings outcomes of individuals without a college degree have fared relatively poorly in the wake of advancements in technology, globalization, and trade, among other factors. Annual earnings of workers with a college degree or more have risen steadily over the past four or five decades, while the earnings of those with lower levels of education have stagnated or fallen (see for example, Autor 2014). Figure 1 shows that the college wage premium—which we initially define in this figure in accordance with previous literature as the difference in log annual earnings between those who have received a bachelor's degree and those who have not—increased steadily from the early 1980s through around 2000, at which point it flattened, but did not reverse. Today this college/high school wage premium remains at 90 percent and is similar for men (88 percent) and women (92 percent).[16]

[16]This wage premium calculation holds constant relative shares of sex-education-experience groups (two sexes, six education categories, and four potential experience categories), as relevant for the

Divergence in employment rates have exacerbated trends in relative earnings. Prime-age adults with no more than a high school degree have experienced a sizable decline in employment rates in recent decades, while employment rates among college-degree holders have fallen only slightly. For instance, among men age 25 to 34 with a high school degree but no college, employment rates fell from 89 to 82 percent between 1999 and 2018, as compared to a dip from 95 to 94 percent among their counterparts with at least a bachelor's degree (Abraham and Kearney, forthcoming). Not surprisingly, economic insecurity, as captured by the likelihood of living in or near poverty, is much higher among the non-college educated. In 2018, 4.4 percent of college graduates lived below the official federal poverty threshold, as compared to 12.7 percent of high school graduates and 25.9 percent of adults without a high school degree (Semega et al, 2019).

The divergent economic outcomes of those with and without a college degree have led many observers to emphasize the need for increased skill attainment, in particular increased college attainment, to boost individual economic security and address rising income inequality. The emphasis on increasing the supply of college graduates to the workforce as a response to the rise in earnings inequality is consistent with the arguments emphasized in the 2008 book by Goldin and Katz, *The Race Between Education and Technology*. The thesis of the book is based on the canonical supply and demand framework of wage determination. In highly simplified terms, the basic observation of Goldin and Katz (2008) is that during the 1980s and 1990s, the

populations of interest, and roughly follows the methodology of Autor, Katz, and Kearney (2008) and Acemoglu and Autor (2011). See appendix for further details.

demand for college-educated workers rose faster than the supply of college-educated workers, leading to a rise in their relative wage.[17]

In this paper, we conduct a simulation exercise that gauges the plausible impact of increased rates of college attainment on a variety of measures of income inequality and economic insecurity. Although several channels for increasing college attainment have been proposed—including additional funding for higher education institutions, expanded access to free or reduced tuition for students, and behavioral or information interventions—we set aside any consideration of the costs or effectiveness of these various approaches to focus on outcomes. The results of this simulation exercise reveal that a sizable increase in rates of college attainment would meaningfully increase economic security for individuals near the bottom of the earnings distribution. It would also shrink gaps between the 90th percentile and lower half of the earnings distribution, as well as between the median and bottom in most cases. However, increases in college attainment would not significantly reduce upper tail inequality or the amount of income going to earners in the top percentiles. The policy prescription of increased educational attainment should thus appeal to those whose primary concern is the economic security of poorer individuals, but it will not satisfy the goals of those whose primary concern is the reduction of overall income inequality or income shares at the top of the distribution.[18]

[17]This point has been suggested in related papers, including but not limited to Goldin and Margo (1992), Katz and Murphy (1992), and Card and Lemieux (2001).
[18]This paper builds on a 2015 policy memo that Hershbein and Kearney wrote with Larry Summers and posted on the Hamilton Project website (Hershbein, Kearney, and Summers, 2015). That memo described the results of simulating how the distribution of earnings would change if one of every ten men aged 25–64 without a bachelor's degree were to be assigned one, with a random draw from the earnings distribution of existing bachelor's-degree holders. In this paper, we expand on that earlier

4.2 Data and Methods

4.2.1 Data

Our primary data source for employment, earnings, income, and poverty status is the Annual Social and Economic Supplement of the Current Population Survey (March CPS), as provided by IPUMS (Flood et al. 2019). The March CPS provides detailed information on the composition of annual income for a relatively large, nationally representative sample of households and is released more quickly than other public datasets that contain earnings.[19] To illustrate changes in earnings and inequality over a longer horizon, we consider both the 1980 survey (covering earnings from 1979) and the 2019 survey (covering earnings from 2018).

We restrict our sample to adult civilians of prime age, 25−54, to minimize concerns about schooling and retirement decisions.[20] We define four mutually exclusive, exhaustive education categories: less than high school degree, high school graduate, associate degree, and bachelor's degree or higher. High school degree includes GED holders and those who attended college but did not get a degree. We measure employment as a binary variable that equals one if an individual worked a

---

analysis by including men and women, considering increased attainment of both associate and bachelor's degrees, using both a random distribution method and a causal parameter assignment method, examining multiple thresholds of increased educational attainment, and using current data.

[19]The March CPS microdata are released in the fall of the survey year and contain annual earnings data for the previous calendar year. American Community Survey (ACS) microdata constitute a larger sample but are released with a greater delay and contain less detailed earnings data that covers a longer time period due to the staggered nature of the survey throughout the year. We intend to repeat our simulation exercise with the decennial census and the ACS, for the sake of comparison.

[20]Previous literature has typically focused on the working-age population, 16–64, but since our simulation involves increasing educational attainment, we believe it makes more sense to focus on the population for whom the additional attainment is more reasonable and exclude those for whom further schooling is less likely.

positive number of weeks in the previous calendar year *and* had positive labor earnings; we define full-time, full-year workers (FTFY) as those usually working at least 35 hours per week and at least 40 weeks of the year. We define an individual's annual labor earnings as the sum of wages and salaries and non-negative business income over the same time period.[21] We adjust earnings for inflation to year 2018 dollars using the personal consumption expenditures (PCE) deflator from the Bureau of Economic Analysis. Because poverty status is based on family rather than individual income, we construct an individual's poverty threshold ratio by dividing that individual's total family income by the official poverty thresholds for the individual's family size and type.[22]

Table 4.1 presents summary statistics showing the earnings and income of adults in 1979 and 2018 for different samples. The first row of each panel reports selected percentiles of the real earnings distributions for all FTFY workers age 25 to 54 in 1979 and 2018, respectively. Subsequent rows show percentiles of the earnings distribution for men and women separately, and then the pooled and gender-specific earnings distributions for all individuals age 25 to 54, regardless of work status.

The rise in inequality over this period is evident from these numbers. Among men, unconditional earnings at the 10th, 25th, and 50th percentiles fell between 1979 and 2018, and FTFY earnings at these percentiles were generally stagnant or

[21]We exclude from the sample individuals for whom either of these components of earning is imputed. About 1% of our 1980 and 2019 samples have a component of earnings topcoded. We do not attempt to adjust for topcoding, but do implement a correction to use current topcoding methods for the 1979 sample, using historical income data generated by Larrimore et al. (2008).
[22]These thresholds are provided annual by the Census Bureau (see https://www.census.gov/data/tables/time-series/demo/income-poverty/historical-poverty-thresholds.html) and are already included in the IPUMS extracts we use.

increased only slightly. At the 75th, 90th, and 99th percentiles, however, earnings rose substantially, both unconditionally and for FTFY men. Among women, both unconditional and FTFY earnings increased at all highlighted percentiles, but the gains were much larger at the higher end of the distribution. Notably, earnings are zero at both the 10th and 25th percentiles of the unconditional sample for women in both 1979 and 2018, and although the 10th percentile of the unconditional earnings distribution is positive for men in 1979, it is zero in 2018. This sharp decline at the bottom reflects a lower likelihood of prime-age men having been employed at any point during the year; this likelihood fell from 92 to 85 percent, with the decline almost entirely concentrated among men without a college degree.[23]

Appendix Table 4A.1 and Appendix Figure 2 show the earnings distributions of FTFY workers in 1979 and 2018 by level of education. The table and figure show clearly how earnings gaps have increased between education groups. For example, in 1979, median earnings among high school graduate FTFY workers were approximately $38,300 (in 2018 dollars), as compared to about $53,400 among FTFY workers with a bachelor's degree (BA) or higher. In 2018, the comparable numbers were $40,000 and $70,000. The gap between the 90th percentile of earnings among high school graduates and BA holders grew by an even greater amount. In 1979, the 90th percentile of earnings among high school FTFY workers was roughly $72,600, as compared to $113,200 among BA holders, but the comparable numbers in 2018 were $80,000 and $155,000, a near doubling.

[23]Appendix Figures 1 and 2 plot kernel density estimates of the earnings distributions of FTFY workers in 1979 and 2018—pooled, and then separately for men and women.

These increases in wage inequality across education and time have occurred simultaneously with increases in educational attainment—although, as Goldin and Katz (2008) have argued, at a slower rate than previously. The first panel of Tables 4.2a and 4.2b shows the shares of the FTFY prime-age workforce (group A), FTFY male prime-age workforce (group B), FTFY female prime-age workforce (group C), and all prime-age men (group D) with different levels of education. Among the FTFY workforce, the share with at least a bachelor's degree has risen from about one-quarter in 1979 to 45 percent by 2018, with a much more modest increase in the associate degree share from about 9 to 11 percent. Because of the faster growth in educational attainment for women relative to men, the educational increases for men specifically are smaller, with BA-plus shares rising from 26 to 41 percent for FTFY men and from 25 to 36 percent for all prime-age men, unconditional on work status. Given the observed changes in earnings by education for different groups, our simulation exercise asks how earnings distributions would change were the education shares for these groups to be shifted.

4.2.2 Methods

We simulate three counterfactual scenarios. Simulation 1 raises the share of the sample—across the different samples described above—with at least a bachelor's degree (BA share) to 50 percent. Simulation 2 raises the share of the sample with an associate degree (AA share) to 15 percent *and* the BA share to 50 percent. Simulation 3 raises the AA share to 20 percent and the BA share to 60 percent. Both new AA holders and new BA holders are drawn from the existing high school graduate population. For each scenario, we assign the "new" AA and BA holders simulated earnings in two ways. The *distribution method* assigns a random draw from the distribution of existing AA or BA (including those with higher than a BA), conditioning on one of 12 cells: 10-year age category (25–34, 35–44, 45–54), race (white and other), and sex (male and female). The *causal parameter* method assigns a causal estimate of the marginal AA or BA returns using parameters from the existing literature, as described below.[24] One benefit of the distribution method is that it allows an individual who is currently out of the workforce to be assigned positive earnings if they are simulated to earn a college degree. The causal parameter method does not allow for employment responses at the extensive margin. The distribution method also allows for heterogeneity in treatment effects, whereas the causal parameter method uses a uniform percentage increase in earnings among the entire sample. On the other hand, the causal parameter method may come closer to capturing the "marginal" policy parameter of interest. We thus view the two methods as complements.

In the causal parameter approach, high school graduates who are assigned an AA receive a 29 percent annual earnings increase. This estimate is based on averaging the effects found for associate degree receipt in Bahr et al. (2014) and Stevens et al. (2015). These papers identify causal estimates using well-established individual fixed-effects methodologies. We assign the high school graduates who are

[24]While it would be desirable to use group-specific causal returns to different degree levels, the literature has not produced robust causal estimates for different demographic groups, and so we assign the same AA premium and BA premium to all individuals who have their college status shifted.

treated with a BA a 68 percent annual earnings increase. This is an approximation of the likely causal effect of BA attainment for a marginal student admitted to a less selective university, based on the findings of Zimmerman (2014). Zimmerman uses a regression discontinuity approach and estimates that individuals just admitted to a less-selective state university have a 22 percent increase in earnings 8 to 14 years after high school graduation relative to those just missing admission. To get an IV estimate of the effect of BA attainment, Zimmerman scales this earnings increase by the probability of attendance conditional on admissions (49 percent) and the probability of BA completion conditional on attendance (50 percent), yielding an IV estimate of a 90 percent earnings increase as compared to below-threshold earnings. This is almost surely an upper bound because, as Zimmerman acknowledges, admission to the university likely affects earnings through other channels, namely, credit completion without a degree. We thus adjust downward the 90 percent estimate. To do so, we assume that roughly a quarter (or 5 percentage points) of the 22 percent earnings increase associated with admission comes from the attendance without completion channel. We thus apply the scaling to a 17 percent earnings increase, obtaining a 68 percent "IV" estimate of BA attainment, rather than 90 percent.

This 68 percent estimate is likely a conservative measure of the earnings premium because it does not allow for the additional earnings premium that would be associated with a more selective institution.[25] Based on a regression discontinuity

[25]Additionally, the baseline earnings from which Zimmerman's estimates are drawn include some individuals who attend community colleges, whose earnings may be somewhat higher than those of high school graduates without any college attendance.

admissions cutoff at a more selective university than the one considered by Zimmerman (2014), Hoekstra (2009) estimates a 20 percent local average treatment effect on earnings of enrolling at a state flagship university, as compared to the likely counterfactual of attending a less selective institution. Thus, a reasonable extension to the assignment of a 68 percent causal parameter (which we do not incorporate) would be to assign some share of new BA holders an additional (multiplicative) 20 percent premium.

In both the distribution and causal parameter method, we further adjust earnings for the relative wage effect that is likely to result from an increase in the share of the population with a college degree. To incorporate this relative wage response into our simulation exercise, we follow the common paradigm in the academic literature, as described in Autor and Acemoglu (2011), and specify a two-factor CES production function model. In one case, the model includes BA and high school degree workers, and in the other case, the model includes AA and high school degree workers.

Appendix C describes our methodology for estimating relative wage effects and presents the resulting relative wage responses. We estimate that within our sample, a 1 percent increase in the relative supply of labor with a BA or more to non-BA high school graduates will narrow the relative wage premium by 0.25 percent; analogously, a 1 percent increase in the relative supply of AA-degree holders to high school graduates will decrease that relative wage premium by 0.18 percent.[26] For

---

[26]Although the relative wage parameter estimates from the regressions are defined for (i) BA (including BA-plus) and high school graduates, and (ii) AA and high school graduates, when applying

instance, the first simulation raises the BA completion rate from 45.1 to 50.3 percent for our FTFY sample (group A in Table 4.2a). In terms of relative supply effects in the labor market, considering all adults (not just FTFY or prime-age) and weighting each individual by their hours worked last year, this amounts to a change from 41.5 to 44.7 percent, which is roughly a 14 percent increase in the hours-weighted relative supply of BA to non-BA labor [0.415/(1–0.415) = 0.708; 0.447/(1–0.447) = 0.808; 0.808/0.708 = 1.141]. Thus, our simulation adjusts for a 0.14 * 0.25 ≈ 4 percent narrowing of the wage premium. This narrowing is assumed to fall equally on each group, raising non-BA earnings by 2 percent and lowering BA earnings (including BA-plus) by 2 percent. Because we draw from the pool of high school graduates to assign college degrees, the relative supply of associate degree holders and high school graduates also changes for the first simulation, with this ratio increasing by 7 percent, leading to a 0.07 * 0.18 ≈ 1 percent narrowing of that wage premium. When both AA and BA attainment is changed, as in Simulations 2 and 3, we narrow the wage premia sequentially: first adjusting the AA/high school wage premium, then narrowing the BA/high school wage premium. In Simulation 2, the AA/high school wage premium narrows by 7 percent and the BA/high school wage premium shrinks by 4 percent. In Simulation 3, the wage premia fall by 22 and 11 percent, respectively.

4.3 Results

Tables 4.2a and 4.2b show the practical impact of the three simulations on the numbers and shares of degree holders for four samples: all FTFY workers, FTFY

---

the adjusted wages to the population, we include individuals without a high school degree in the lower-skill group, implicitly treating them as perfect substitutes.

men, FTFY women, and all men unconditional on work status. We focus on these four samples because the unconditional sample of women includes a large share of non-workers. As shown in the top panel, in 2018 45.1 percent of FTFY prime-age workers held at least a BA and 10.9 percent held an AA. (Among all adults age 25 to 54, 39.9 percent held at least a BA and 10.7 percent held an AA; not shown in the table). Simulation 1 raises the BA share to 50 percent, which is a modest increase when the sample is limited to FTFY workers. For the full sample of prime-age individuals, this increase is more substantial, requiring that 11.1 million more adults hold a bachelor's degree (from 39.9 million to 51.0 million; not shown in the table). Simulation 2 maintains the bachelor's degree share increase to 50 percent and adds an increase in the share of the sample with associate degrees to 15 percent; while the latter is only a 4–5 percentage point bump from 2018 levels, it represents a relatively large proportional increase. Simulation 3 increases the respective shares to 60 and 20 percent. This requires an additional 21 million more prime-age adults to hold a bachelor's degree and 9.9 million more to hold an associate degree, which are ambitiously large gains; even among the FTFY sample, the respective increases are 10 million and 6.2 million (group A, Tables 4.2a and 4.2b). As described above, the simulation imparts new degrees to the current population of high school graduates, which in 2018 composed 40.1 percent (41.4 million) of prime-age adults and 37.7 percent (24.8 million adults) of prime-age FTFY workers.[27]

[27]Note that the simulations for the first three groups (FTFY samples) are based on raising education for *all* FTFY workers by the stated amounts, not men and women separately in the FTFY men and FTFY women samples. For the all-men group, education is raised for all (prime-age) men by the stated amount.

Table 4.3 illustrates how one of our counterfactual simulations affects the earnings distribution. It reports both observed earnings percentiles and simulated earnings percentiles, for each simulation, using the distribution method, for all FTFY workers, FTFY men, FTFY women, and all men. The simulations raise earnings in all four samples for roughly the lower three-quarters of the earnings distribution, with the strongest gains in the middle. The highest percentiles, however, show much smaller gains, or even losses among FTFY men, due to the general equilibrium effects that lower the college wage premium.

We are particularly interested, however, in how these changes affect distributional outcomes. Table 4.4a thus reports percentile earnings ratios for the sample of all prime-age FTFY workers, including changes based on all three simulations, according to both the distribution and causal parameter methods. As can be seen in the table, there were large increases in the 90/10, 90/25, and 90/50 percentile earnings ratios between 1979 and 2018, reflecting disproportionate growth at the top of the distribution (Table 4.1). However, there was actually a slight decrease in the 50/10 ratio over this period.

As the lower panel of Table 4.4a indicates, the simulation of a sizable increase in the rate of bachelor's degree attainment would lead to meaningful reductions in earnings ratios between the 90^{th} and lower percentiles among FTFY workers, and this is true for either simulation method, as both produce similar results. For example, the 90/10 ratio increased from 4.63 to 5.45 between 1979 and 2018. Simulation 3 (increasing AA rates to 20 percent and BA rates to 60 percent) would bring that ratio down to 5.16 (distribution method) or 5.00 (causal parameter method), reversing from more than half to all of the actual increase over this period.[28] As suggested by Table 4.3, the reduction stems from increases in the 10th percentile of FTFY earnings and smaller proportional change at the 90th percentile. The same simulation also substantially reduces the 90/50, 90/25, and 50/25 earnings ratios, although the reductions are less dramatic. Simulations 1 and 2, which involve smaller shifts in degree attainment, produce correspondingly smaller, but still sizable, reductions in these inequality measures. Interestingly, the causal parameter method produces slightly larger reductions in the percentile ratios than the distribution method, and the difference increases as the simulation becomes more extreme in the education shifts.[29]

The estimates reported in Table 4.4a incorporate relative wage effects estimated using data from 1979 to 2018. If we instead estimate relative wage effects using data from 1963 to 2018, consistent with previous literature, the depressive effect of increased BA attainment on relative wages would be larger and the depressive effect of increased AA attainment on relative wages would be smaller (as shown in Appendix Tables 4C.1 and 4C.2). Appendix Table 4A.3 reproduces the results from Table 4.4a using these relative wage effects instead. As can be seen in the table, the simulated reductions in the 90/10 and 90/25 wage ratios are even larger. Simulation 3 reduces the 90/10 ratio to 4.85 (distribution method) and 4.64 (causal

[28]Appendix Table 4 A.2 reports the analogous results when relative wage effects are not taken into account. The resulting reductions in inequality are, as expected, smaller, especially for the causal parameter method. For instance, the simulated 90/10 earnings ratio under simulation 3 becomes 5.20 according to both methods, as compared to 5.06 and 4.95 when relative wages are adjusted.

[29]This gap likely relates to the large earnings variance *among* college graduates; while the causal parameter method unambiguously increases earnings, the distribution method can result in some "treated" individuals having their earnings reduced, if the draw is sufficiently bad.

parameter method). It reduces the 90/25 ratio to 3.50 (distribution method) and 3.29 (causal parameter method).

Tables 4.4b and 4.4c report results separately for FTFY men and women, using our baseline approach. As in the pooled sample, the results from both the distribution and causal parameter methods show that for both men and women, a sizable increase in the rate of bachelor's degree attainment would lead to meaningful reductions in earnings ratios between the 90th and lower percentiles. For example, among FTFY men, the 90/10 ratio increased between 1979 and 2018 from 3.86 to 5.58; simulation 3 would bring that ratio down to 5.18 (distribution method) or 5.04 (causal parameter method), reducing the increase in inequality by up to one-third. Among FTFY women, the 90/10 ratio increased from 3.6 to 5.0; simulation 3 would bring that ratio down to 4.44 (distribution method) or 4.28 (causal parameter method), reducing the increase in inequality by about one half. Sizable reductions are also observed for the 90/25 and 50/25 ratios. Again, we see only small reductions (or for FTFY women, increases) in the 99/90 ratio, consistent with the rising dispersion in earnings among college graduates.[30] The causal parameter method produces slightly larger reductions in the percentile ratios than the distribution method, but these should be interpreted with caution as we do not have separate causal parameter estimates for men and women.

As discussed above, employment rates for prime-age men have fallen over time, especially for less-educated prime-age men. Thus, it is also illustrative to

[30]These reductions would likely be even smaller with better corrections for topcoded earnings.

examine how our simulations would affect earnings ratios and employment rates (proxied by positive earnings) for all prime-age men, regardless of work status. Table 4.4d reports observed and simulated earnings ratios for this latter sample.[31] As the 10th percentile of earnings for this sample is zero in both 1979 and 2018, we omit ratios with the 10th percentile in the denominator. The remaining ratios all experienced large increases over the nearly 40-year period, chiefly driven by reductions in earnings at the lower (and even middle) percentiles, which are in turn a symptom of the 7-point reduction in employment rates. The causal method is less useful for this sample, since it only increases earnings of those with positive earnings and does not allow for an extensive margin effect on employment. Not surprisingly, the simulated effects on income inequality for this sample are smaller at the lower end using the causal method than the distribution method. The distribution method results show that increasing the BA rate to 60 percent and the AA rate to 20 percent could lead to meaningful reductions in the 50/25 and 90/25 unconditional earnings ratios. The 50/25 ratio, which rose from 1.71 to 2.18, would fall to 1.88. The 90/25 ratio, which rose from 3.33 to 6.00, would fall to 4.83. As expected, the more intense simulations are associated with larger reductions. The distribution method simulation also suggests the employment rate would rise by 1.2–2.8 percentage points, suggesting gains below the 25th percentile not captured by the displayed ratios.[32]

[31]We do not report analogous results for the unconditional pooled sample of men and women or women separately, as 34.6 percent of women reported no earnings in 1979, making comparisons of unconditional earnings ratios over time less meaningful.

[32]We consider the employment rate to have increased when individuals switch from zero to positive earnings under the simulation. The causal parameter method affects only the intensive margin and thus the employment rate is unchanged by this method.

Table 4.5 reports the results of the simulated increase in college attainment on measures of individual level economic insecurity, as captured by four poverty measures: deep poverty (family income less than 50 percent of the federal poverty threshold), poverty (family income less than the poverty threshold), near poverty (family income less than 150 percent of the threshold), and low income (family income less than 200 percent of the federal poverty threshold). Here we follow official rules and define an individual's poverty status by whether that individual's *family income* is less than the corresponding Census poverty threshold, which varies by family size and composition.[33] As reported in the table, all four measures of poverty increased between 1979 and 2018 among adults age 25 to 54. The share of prime-age adults living below the poverty line increased from 8.2 to 11.3 percent, and the share living in deep poverty increased from 3.0 to 5.6 percent.

To simulate the effect on poverty of increased college attainment, we calculate simulated poverty rates by taking an individual's 2018 family income and adding any of their own additional earnings assigned by the simulation. Our calculation assumes family structure is fixed and there is no induced change in other family members' earnings; nor does it adjust income for any changes in taxes and transfers that would result from an increase in family earnings. Because this approach ignores any potential increase in taxes and reduction in transfer benefits, it likely overstates the increase in "true" household income and corresponding reduction in poverty. (An obvious exception is that some households might see an increase in their Earned

[33]For an explanation of how official poverty statistics are calculated and the 2018 federal poverty thresholds, see: https://www.census.gov/topics/income-poverty/poverty/guidance/poverty-measures.html.

Income Tax Credit.) However, because in-kind transfers and taxes are excluded from official poverty estimates, our approach is reasonable when using that measure as reference.[34]

Based on the results of applying the distribution method, the simulated effect of increasing the BA share to 60 percent and the AA share to 20 percent is to reduce the poverty rate by 2.39 percentage points, from 11.3 to 8.91 percent in the sample using all civilian adults age 25 to 54. Reductions in the near-poverty or low-income rate are larger, with the first falling from 18.5 to 14.2 percent, and the second falling from 26.5 to 20.4 percent. Both of these simulated rates are lower than their actual levels in 1979. The rate in deep poverty also falls, but only modestly, from 5.6 to 5 percent. This reflects the fact that very few people with a high school degree live in deep poverty (7.1 percent). To decrease rates of deep poverty, an intervention that targets high school dropouts (who have a deep-poverty rate of 12.7 percent) would likely be more effective.

The corresponding estimates from the causal parameter method imply smaller reductions of roughly half the magnitude of those from the distribution method. This, in large part, reflects that the former method does not allow for changes in the likelihood of employment and only increases earnings for those who have positive

[34]We experimented with using the NBER Taxsim model to adjust family income for taxes; however, since Taxsim calculates taxes owed and credits received, but does not include information about transfer benefits, the estimated numbers are still not an accurate measure of what would likely happen to household income net of taxes and transfers if earnings increased. In any case, the effects on poverty calculations are likely to be small because official poverty statistics do not adjust household income for taxes paid or tax credits received, nor do they include in-kind benefits such as SNAP. We have thus decided to report two benchmark estimates for poverty effects, one that simply adds earnings to existing household income and one that calculates poverty rates based only on earnings.

earnings, while the latter method allows for these changes, which are particularly likely to affect (near-) poverty measures.

Appendix Table 4A.4 reports the results from calculating poverty rates using only observed family earnings, ignoring other sources of income. These rates do not correspond to official poverty statistics, but they allow us to gauge rates of economic self-sufficiency, as captured by the share of prime-age adults in families who earn enough money to live above the federal poverty threshold, or multiples thereof. In 2018, 19.5 percent of individuals lived in families with earnings less than the federal poverty threshold, up from 17.2 percent in 1979. Using the distribution method, raising the BA share to 50 percent would reduce this poverty measure to 17.3 percent, back to its 1979 level; additionally raising the AA share to 15 percent would reduce this poverty rate to 16.8 percent, and raising the BA share to 60 percent and the AA share to 20 percent would further reduce the poverty rate to 14.8 percent. This total reduction of 4.65 percentage points would correspond to 6 million fewer prime-age adults in poverty (based on 2018 Census population counts). Similar declines would occur for the other poverty thresholds.

The reduction in poverty rates from simulations using the causal parameter method is much smaller, for the reasons discussed above. The difference in the simulated effects of poverty rates between the two methods highlights how important the effect of increased college attainment on the employment margin is for poverty avoidance and basic economic security.

We further consider what increased college attainment of prime-age individuals could do for child poverty rates. The results are presented in Appendix

Table 4A.5. The share of children in poverty (calculated using our data and the official federal poverty threshold) held steady between 1979 and 2018: 16.79 percent and 16.83 percent. These simulations suggest that the increased household earnings associated with increasing the BA attainment rate to 60 percent and the AA rate to 20 percent would reduce child poverty rates to 13.47 percent (distribution method) or 14.84 percent (causal parameter method.) There would be an even larger percentage point reduction in the share of children living in households with income less than 200 percent of the poverty rate. Increased economic security among children should be considered a key benefit that would result from more educated prime-age adults, many of whom are parents to young children.

4.4 Conclusion

In this analysis we have simulated the effects of increasing college attainment, both bachelor's and associate degrees, of men and women age 25 to 54 to gauge the likely effects on earnings and earnings inequality. We have conducted the simulation using two distinct approaches. The distribution method assigns individuals whose college status is randomly shifted a draw from the earnings distribution of college-educated workers. The causal parameter method assigns workers whose college status is randomly shifted a single earnings premium, based on existing studies in the literature. Both approaches suggest that increasing the educational attainment of adults without a college degree will increase their average earnings, with gains concentrated in the lower half of the earnings distribution. The distribution method further allows for an increase in the likelihood of work, which is particularly important for raising earnings at the bottom of the distribution. The results of the simulation also show meaningful reductions in rates of poverty and near-poverty (family income less than 200 percent of the federal poverty threshold). Increasing rates of college degree attainment will also moderately reduce inequality, mostly by raising the lower-middle part of the earnings distribution relative to the upper-middle. However, increased college attainment will have minimal effects on reducing overall inequality back to the levels in 1979, as a greater share of the population with college degrees will not meaningfully affect earnings at the highest parts of the distribution, where much of the rise in inequality has taken place.

In this paper we have provided a quantitative approximation to what could be achieved in terms of reduced income inequality and increased individual economic security through a meaningful, albeit feasible, increase in the share of prime-age adults with a college degree. We have not attempted to argue for or against any particular way of achieving that result, though obviously the question of *how* to achieve increased college attainment is of the utmost importance. Nor have we made the claim that increasing college attainment is *sufficient* to address the current degree of income inequality or income insecurity. We view the results of this analysis as suggesting that increasing college attainment is an important—and potentially necessary—policy response to the rise in income inequality experienced over recent decades, but one that belongs alongside a number of other policy responses aimed at increasing the economic security of low-income Americans.

4.5 Tables and Figures

Table 4.1: Summary Statistics by Year and Sample: Earnings Percentiles and Inequality Measures

Year	p10	p25	p50	p75	p90	p99	Gini
			Panel A: 1979				
Full-time, full-year workers							
All	17,482	27,579	41,206	58,062	80,868	174,185	0.332
Men	22,934	34,837	50,514	67,642	88,544	203,216	0.310
Women	14,515	21,007	29,031	40,170	52,256	87,093	0.275
All individuals							
All	0	2,032	25,257	46,667	69,674	145,155	0.537
Men	5,632	26,128	44,603	63,868	87,093	188,701	0.381
Women	0	0	8,709	26,128	40,643	71,039	0.614
			Panel B: 2018				
Full-time, full-year workers							
All	22,000	32,000	50,000	78,000	120,000	320,000	0.403
Men	24,000	35,000	55,000	85,000	134,000	400,000	0.409
Women	20,000	30,000	45,000	68,000	100,000	260,000	0.383
All individuals							
All	0	6,500	34,000	60,000	100,000	268,000	0.565
Men	0	20,000	43,614	75,000	120,000	310,000	0.517
Women	0	0	25,000	50,000	80,000	200,000	0.598

Note: Statistics are calculated for civilian men and women ages 25 to 54. Earnings are defined as the sum of annual wage, salary, and positive business income, adjusted for inflation (to 2018 dollars) using the personal consumption expenditures (PCE) deflator of the Bureau of Economic Analysis. Employment is defined as having positive earnings in the reference year, and full-time, full-year workers are those working at least 40 weeks in the previous calendar year and at least 35 hours usually worked per week.
Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.2a: Numbers (in millions) and Shares of Degree Holders: Full-Time Full-Year Workers

	Group A: FTFY			Group B: FTFY Men		
	High School Graduate	AA Holder	BA or greater	High School Graduate	AA Holder	BA or greater
			Panel A: Observed			
1979	21.2 49.1%	3.7 8.7%	10.7 24.8%	13.0 45.9%	2.5 8.9%	7.5 26.4%
2018	24.8 37.7%	7.2 10.9%	29.6 45.1%	15.4 41.4%	3.8 10.1%	15.1 40.6%
			Panel B: Simulations for 2018			
Raise BA share to 50%	21.3 32.4%	7.2 10.9%	33.1 50.3%	13.3 35.7%	3.8 10.1%	17.2 46.3%
+ Raise AA share to 15%	18.5 28.2%	10.0 15.2%	33.1 50.3%	11.5 31.0%	5.5 14.8%	17.2 46.3%
20% AA share, 60% BA share	8.7 13.3%	13.2 20.1%	39.6 60.2%	5.4 14.6%	7.6 20.3%	21.3 57.2%

Note: High school graduates are defined as those with a high school degree (or equivalent) or some college, but no degree. For 1979 data, we define associate degree (AA) holders as those with exactly two years of college education, and "BA or greater" as those with four or more years of college education.
Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.2b: Numbers (in millions) and Shares of Degree Holders: Full-Time Full-Year Workers

	Group C: FTFY Women			Group D: All Men		
	High School Graduate	AA Holder	BA or greater	High School Graduate	AA Holder	BA or greater
			Panel A: Observed			
1979	8.2 55.1%	1.2 8.3%	3.2 21.7%	15.4 45.0%	2.9 8.6%	8.4 24.6%
2018	9.3 32.7%	3.4 11.9%	14.5 50.9%	21.5 43.8%	4.7 9.6%	17.9 36.4%
			Panel B: Simulations for 2018			
Raise BA share to 50%	8.0 28.2%	3.4 11.9%	15.8 55.5%	15.7 31.9%	4.7 9.6%	23.8 48.3%
+ Raise AA share to 15%	7.0 24.5%	4.4 15.6%	15.8 55.5%	13.3 27.0%	7.1 14.4%	23.8 48.4%
20% AA share, 60% BA share	3.3 11.6%	5.7 19.8%	18.3 64.2%	5.4 10.9%	9.8 19.9%	29.0 59.0%

Note: High school graduates are defined as those with a high school degree (or equivalent) or some college, but no degree. For 1979 data, we define associate degree (AA) holders as those with exactly two years of college education, and "BA or greater" as those with four or more years of college education.
Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.3: Simulated Effects of Increasing College Shares on Annual Earnings Distributions: Using Distribution Approach

	Share with AA	Share with BA or greater	Annual Earnings (Thousands) p10	p25	p50	p75	p90	p99	Gini
Panel A: 2018 Baseline									
FTFY Men and Women	10.9%	45.1%	22.0	32.0	50.0	78.0	120.0	320.0	0.403
FTFY Men	10.1%	40.6%	24.0	35.0	55.0	85.0	134.0	400.0	0.409
FTFY Women	11.9%	50.9%	20.0	30.0	45.0	68.0	100.0	260.0	0.383
All Men	9.6%	36.4%	0	20.0	43.6	75.0	120.0	310.0	0.517
Panel B: Simulation 1									
FTFY Men and Women	10.9%	50.3%	23.3	33.8	51.3	79.9	122.5	343.0	0.401
FTFY Men	10.1%	46.3%	24.6	35.9	57.4	89.2	141.1	392.0	0.407
FTFY Women	11.9%	55.5%	20.5	30.8	46.1	68.6	100.9	272.7	0.379
All Men	9.6%	48.4%	0	24.0	48.0	81.6	128.1	355.2	0.503
Panel C: Simulation 2									
FTFY Men and Women	15.2%	50.3%	23.5	34.3	52.8	79.2	122.5	343.0	0.397
FTFY Men	14.8%	46.3%	25.3	37.0	58.0	89.7	141.1	392.0	0.402
FTFY Women	15.6%	55.5%	21.1	31.5	47.5	68.6	100.9	274.4	0.375
All Men	14.4%	48.4%	0	24.1	49.1	81.6	129.5	345.6	0.498
Panel D: Simulation 3									
FTFY Men and Women	20.1%	60.2%	24.6	36.4	55.8	84.5	126.8	351.3	0.387
FTFY Men	20.3%	57.2%	27.4	39.2	60.9	94.5	141.8	382.7	0.391
FTFY Women	19.8%	64.2%	23.4	33.1	47.3	70.9	104.0	283.5	0.363
All Men	19.9%	59.0%	0	27.3	51.3	84.6	132.0	364.0	0.480

Note: Results are presented for each simulation under the distributional assignment approach: 1) increasing the share of all individuals with a BA or more (in the FTFY or entire sample) to 50 percent; 2) increasing the share with a BA or more to 50 percent and the share with an AA to 15 percent; and 3) increasing these shares to 60 percent and 20 percent, respectively.
Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.4a: Observed and Simulated Percentile Earnings Ratios: Full-Time, Full-Year Workers

	p50/ p10	p90/ p10	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini
Panel A: Observed							
1979	2.36	4.63	1.49	2.93	1.96	2.15	0.332
2018	2.27	5.45	1.56	3.75	2.40	2.67	0.403
Panel B: 2018 Simulations							
Distribution Method							
1) Raise BA share to 50%	2.20	5.25	1.52	3.62	2.39	2.80	0.401
2) + Raise AA share to 15%	2.24	5.21	1.54	3.57	2.32	2.80	0.397
3) 60% BA share, 20% AA share	2.27	5.16	1.53	3.47	2.27	2.77	0.387
Causal Parameter Method							
1) Raise BA share to 50%	2.20	5.25	1.52	3.62	2.39	2.64	0.397
2) + Raise AA share to 15%	2.24	5.21	1.54	3.57	2.32	2.64	0.394
3) 60% BA share, 20% AA share	2.26	5.00	1.52	3.36	2.21	2.62	0.382

Note: The distribution method assigns "treated" individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.4b: Observed and Simulated Percentile Earnings Ratios: Full-Time, Full-Year Men

	p50/ p10	p90/ p10	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini
Panel A: Observed							
1979	2.20	3.86	1.45	2.54	1.75	2.30	0.310
2018	2.29	5.58	1.57	3.83	2.44	2.99	0.409
Panel B: 2018 Simulations							
Distribution Method							
1) Raise BA share to 50%	2.33	5.74	1.60	3.93	2.46	2.78	0.407
2) + Raise AA share to 15%	2.29	5.57	1.57	3.82	2.43	2.78	0.402
3) 60% BA share, 20% AA share	2.22	5.18	1.53	3.56	2.33	2.70	0.391
Causal Parameter Method							
1) Raise BA share to 50%	2.32	5.56	1.58	3.79	2.39	2.86	0.402
2) + Raise AA share to 15%	2.31	5.39	1.58	3.68	2.33	2.86	0.399
3) 60% BA share, 20% AA share	2.17	5.04	1.49	3.46	2.32	2.67	0.385

Note: The distribution method assigns "treated" individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.4c: Observed and Simulated Percentile Earnings Ratios: Full-Time, Full-Year Women

	p50/ p10	p90/ p10	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini
Panel A: Observed							
1979	2.20	3.86	1.45	2.54	1.75	2.30	0.310
2018	2.29	5.58	1.57	3.83	2.44	2.99	0.409
Panel B: 2018 Simulations							
Distribution Method							
1) Raise BA share to 50%	2.25	4.92	1.50	3.28	2.19	2.70	0.379
2) + Raise AA share to 15%	2.25	4.78	1.51	3.20	2.12	2.72	0.375
3) 60% BA share, 20% AA share	2.02	4.44	1.43	3.14	2.20	2.73	0.363
Causal Parameter Method							
1) Raise BA share to 50%	2.25	4.80	1.50	3.20	2.13	2.62	0.375
2) + Raise AA share to 15%	2.23	4.68	1.49	3.12	2.10	2.63	0.371
3) 60% BA share, 20% AA share	2.02	4.28	1.43	3.03	2.12	2.59	0.359

Note: The distribution method assigns “treated” individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors’ calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.4d: Observed and Simulated Percentile Earnings Ratios: All Men

	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini	Employment Rate
Panel A: Observed						
1979	1.71	3.33	1.95	2.17	0.381	92.48%
2018	2.18	6.00	2.75	2.58	0.517	85.41%
Panel B: 2018 Simulations						
Distribution Method						
4) Raise BA share to 50%	2.00	5.34	2.67	2.77	0.503	86.71%
5) + Raise AA share to 15%	2.04	5.36	2.63	2.67	0.497	86.90%
6) 60% BA share, 20% AA share	1.88	4.83	2.57	2.76	0.480	88.21%
Causal Parameter Method						
4) Raise BA share to 50%	2.25	5.87	2.61	2.56	0.506	85.41%
5) + Raise AA share to 15%	2.17	5.65	2.60	2.58	0.503	85.41%
6) 60% BA share, 20% AA share	2.08	5.11	2.46	2.59	0.492	85.41%

Note: The distribution method assigns "treated" individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4.5: Poverty Rates: All Prime-Age Individuals

	Deep Poverty (<50% FPL)	Poverty (FPL)	Near Poverty (<150% FPL)	Low Income (<200% FPL)
Panel A: Observed				
1979	2.97%	8.21%	14.91%	23.22%
2018	5.59%	11.30%	18.48%	26.46%
Panel B: 2018 Simulations				
Distribution Method				
1) Raise BA				
2) share to 50%	5.17%	10.17%	16.39%	23.24%
3) + Raise AA				
4) share to 15%	5.13%	9.95%	15.88%	22.59%
3) 60% BA share,				
20% AA share	4.95%	8.91%	14.19%	20.41%
Causal Parameter Method				
1) Raise BA				
2) share to 50%	5.42%	10.72%	17.26%	24.38%
3) + Raise AA				
4) share to 15%	5.34%	10.51%	16.81%	23.89%
3) 60% BA share,				
20% AA share	5.26%	10.13%	16.08%	22.76%

Note: "FPL" is the federal poverty threshold as calculated by the Census Bureau for different family structures. Simulated changes in poverty rates reflect changes to each household member's income through direct wage or general equilibrium relative wage effects. Any resulting changes to transfer payments are not reflected in this analysis.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Figure 4.1: Trends in Wage Premia for Bachelor's-Plus/Noncollege and Associate Degree/Noncollege

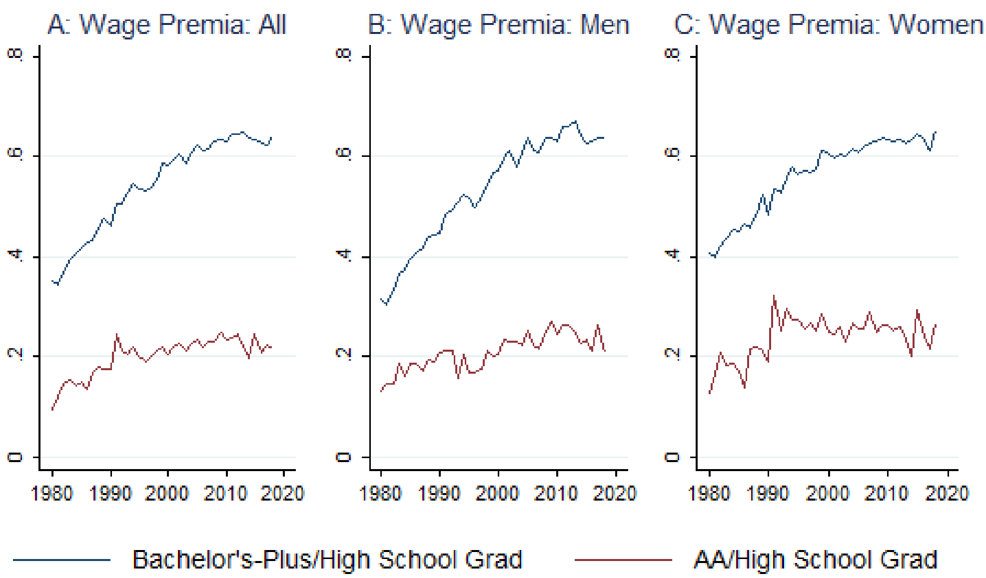

Note: These wage premia series depict a fix-weighted ratio of bachelor's-plus to high school or AA to high school wages for a composition-constant set of sex-education-experience groups (two sexes, six education categories, and four potential experience categories), similar in methodology to Autor, Katz,and Kearney (2008) and Acemoglu and Autor (2011). See appendix for further details.
Source: 1980-2019 March CPS and author's calculations.

4.6 Appendix A: Summary Statistics and Additional Simulations

Table 4A.1: Summary Statistics by Education: Earnings Distribution and Inequality Measures: FTFY Workers

Year	p10	p25	p50	p75	p90	p99	Gini
1979							
Overall	17,482	27,579	41,206	58,062	80,868	174,185	0.332
Less than HS	13,064	20,322	31,112	46,449	62,390	101,608	0.324
HS Degree	17,419	26,128	38,321	55,739	72,577	119,027	0.302
AA	21,773	30,482	43,546	60,965	81,287	145,157	0.293
BA or Greater	26,708	37,740	53,417	75,480	113,221	287,406	0.332
2018							
Overall	22,000	32,000	50,000	78,000	120,000	320,000	0.403
Less than HS	15,000	20,000	28,000	40,000	55,000	130,000	0.328
HS Degree	20,000	27,700	40,000	57,000	80,000	170,000	0.339
AA	24,000	31,200	47,000	65,000	90,000	175,000	0.327
BA or Greater	32,000	47,000	70,000	100,000	155,000	450,000	0.387

Note: Statistics are calculated for civilian men and women ages 25 to 54. Earnings are defined as the sum of annual wage, salary, and positive business income, adjusted for inflation (to 2018 dollars) using the personal consumption expenditures (PCE) deflator of the Bureau of Economic Analysis. Employment is defined as having positive earnings in the reference year, and full-time, full-year workers are those working at least 50 weeks in the previous calendar year and at least 35 hours usually worked per week. See text for description of education categories.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4A.2: Observed and Simulated earnings ratios: Full-Time, Full-Year Workers: No Relative Wage Effects

	p50/ p10	p90/ p10	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini
Panel A: Observed							
1979	2.36	4.63	1.49	2.93	1.96	2.15	0.332
2018	2.27	5.45	1.56	3.75	2.40	2.67	0.403
Panel B: 2018 Simulations							
Distribution Method							
4) Raise BA share to 50%	2.25	4.92	1.50	3.28	2.19	2.70	0.379
5) + Raise AA share to 15%	2.25	4.78	1.51	3.20	2.12	2.72	0.375
6) 60% BA share, 20% AA share	2.02	4.44	1.43	3.14	2.20	2.73	0.363
Causal Parameter Method							
4) Raise BA share to 50%	2.21	5.43	1.52	3.74	2.46	2.80	0.406
5) + Raise AA share to 15%	2.26	5.43	1.53	3.68	2.41	2.72	0.403
6) 60% BA share, 20% AA share	2.24	5.20	1.56	3.61	2.32	2.85	0.398

Note: The distribution method assigns "treated" individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4A.3: Observed and Simulated Earnings Ratios: Full-Time, Full-Year Workers, Using Relative Wage Effects Estimated over 1963–2018 Data

	p50/ p10	p90/ p10	p50/ p25	p90/ p25	p90/ p50	p99/ p90	Gini
Panel A: Observed							
1979	2.36	4.63	1.49	2.93	1.96	2.15	0.332
2018	2.27	5.45	1.56	3.75	2.40	2.67	0.403
Panel B: 2018 Simulations							
Distribution Method							
7) Raise BA share to 50%	2.25	5.21	1.54	3.57	2.32	2.80	0.398
8) + Raise AA share to 15%	2.24	5.15	1.56	3.58	2.30	2.80	0.395
9) 60% BA share, 20% AA share	2.08	4.85	1.50	3.50	2.33	2.71	0.378
Causal Parameter Method							
7) Raise BA share to 50%	2.25	5.21	1.54	3.57	2.32	2.64	0.394
8) + Raise AA share to 15%	2.22	5.10	1.54	3.53	2.30	2.64	0.392
9) 60% BA share, 20% AA share	2.08	4.64	1.47	3.29	2.23	2.60	0.375

Note: The distribution method assigns "treated" individuals a random draw from the earnings distribution of the assigned group (AA or BA). The causal parameter method increases the earnings of the treated by a factor consistent with existing literature (see text). We include general equilibrium effects on wages in these simulations, as explained in the text.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4A.4: Poverty Rates Relative to Earned Income for Prime-Age Individuals (All Prime-Age Individuals Simulation)

	Deep Poverty (<50% FPL)	Poverty (FPL)	Near Poverty (<150% FPL)	Low Income (<200% FPL)
Panel A: Observed				
1979	12.24%	17.15%	23.72%	31.89%
2018	14.42%	19.47%	26.20%	33.74%
Panel B: 2018 Simulations				
Distribution Method				
5) Raise BA				
6) share to 50%	12.89%	17.31%	23.29%	29.89%
7) + Raise AA				
8) share to 15%	12.46%	16.76%	22.62%	29.21%
3) 60% BA share,				
20% AA share	10.95%	14.82%	20.30%	26.71%
Causal Parameter Method				
5) Raise BA				
6) share to 50%	14.21%	18.63%	24.72%	31.45%
7) + Raise AA				
8) share to 15%	14.08%	18.37%	24.22%	30.91%
3) 60% BA share,				
20% AA share	13.95%	17.89%	23.34%	29.67%

Note: "FPL" is the federal poverty threshold as calculated by the Census Bureau for different family structures. Simulated changes in poverty rates reflect changes to each household member's income through direct wage or general equilibrium relative wage effects. Any resulting changes to transfer payments are not reflected in this analysis.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4A.5: Child Poverty Rates (All Prime-Age Individuals Simulation)

	Deep Poverty (<50% FPL)	Poverty (FPL)	Near Poverty (<150% FPL)	Low Income (<200% FPL)
Panel A: Observed				
1979	6.33%	16.79%	27.45%	39.36%
2018	7.54%	16.83%	27.92%	38.23%
Panel B: 2018 Simulations				
Distribution Method				
9) Raise BA 10) share to 50%	6.88%	15.24%	25.09%	34.90%
11) + Raise AA 12) share to 15%	6.78%	15.05%	24.67%	34.40%
3) 60% BA share, 20% AA share	6.50%	13.47%	22.62%	32.21%
Causal Parameter Method				
9) Raise BA 10) share to 50%	7.18%	15.87%	26.01%	36.21%
11) + Raise AA 12) share to 15%	7.07%	15.60%	25.55%	35.75%
3) 60% BA share, 20% AA share	6.89%	14.84%	24.39%	34.49%

Note: "FPL" is the federal poverty threshold as calculated by the Census Bureau for different family structures. Simulated changes in poverty rates reflect changes to each household member's income through direct wage or general equilibrium relative wage effects. Any resulting changes to transfer payments are not reflected in this analysis.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

Table 4A.6: Child Poverty Rates Relative to Earned Income: All Individuals Sample

	Deep Poverty (<50% FPL)	Poverty (FPL)	Near Poverty (<150% FPL)	Low Income (<200% FPL)
Panel A: Observed				
1979	25.69%	33.33%	42.27%	52.49%
2018	30.58%	38.73%	47.51%	55.22%
Panel B: 2018 Simulations				
Distribution Method				
13) Raise BA				
14) share to 50%	29.17%	36.37%	44.25%	51.79%
15) + Raise AA				
16) share to 15%	28.72%	35.95%	43.81%	51.44%
3) 60% BA share,				
20% AA share	27.38%	33.90%	41.55%	49.23%
Causal Parameter Method				
13) Raise BA				
14) share to 50%	30.21%	37.43%	45.40%	53.20%
15) + Raise AA				
16) share to 15%	30.01%	37.13%	44.87%	52.77%
3) 60% BA share,				
20% AA share	29.74%	36.30%	43.72%	51.48%

Note: "FPL" is the federal poverty threshold as calculated by the Census Bureau for different family structures. Simulated changes in poverty rates reflect changes to each household member's income through direct wage or general equilibrium relative wage effects. Any resulting changes to transfer payments are not reflected in this analysis.

Source: Authors' calculations of March Current Population Survey 1980 and 2019 (Flood et al., 2019).

4.7 Appendix B: Figures Plotting FTFY Earnings Distributions

Figure 4B.1: Earnings Distributions, FTFY, by Sex

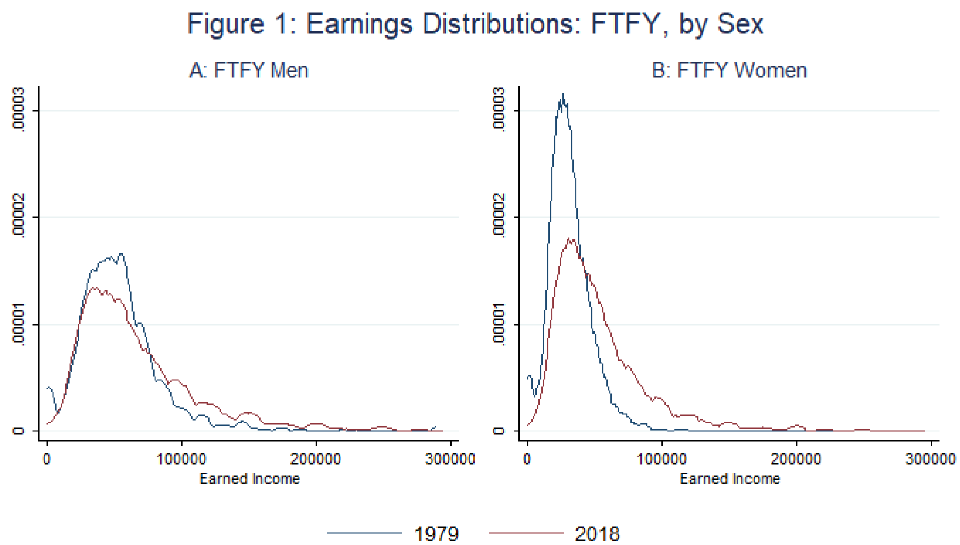

Figure 4B.2: Earnings Distributions, FTFY, by Education

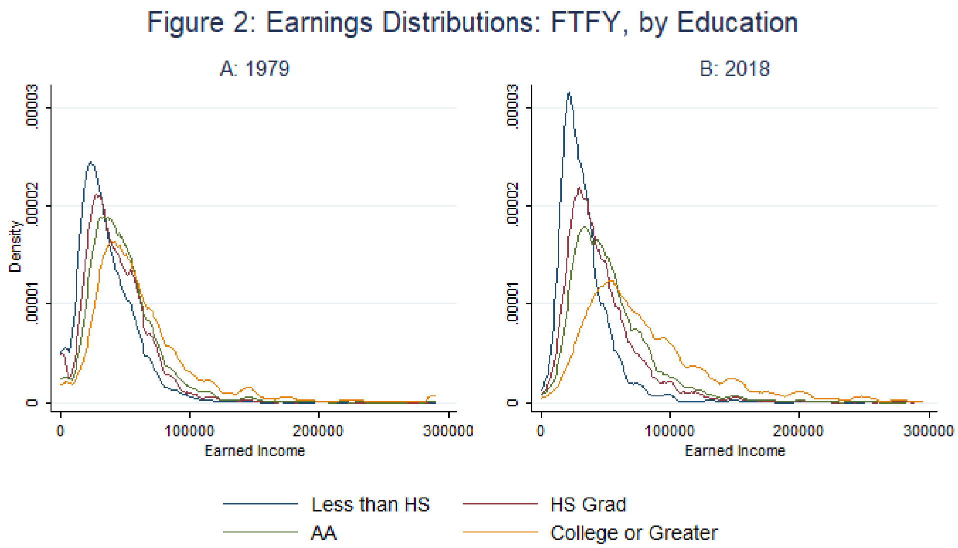

4.8 Appendix C: Estimation of Wage Premia and CES Substitution Elasticities

4.8.1 Wage Premia

We calculate composition-adjusted BA/high school and AA /high school relative wages overall and by age or experience using the March CPS sample. These data are sorted into sex-education-experience groups based on a breakdown of the data into two sexes, six education categories (high school dropout, high school graduate, some college, associate's degree, college plus, and greater than college), and four potential experience categories (0–9, 10–19, 20–29, and 30+ years). Log weekly wages of full-time, full-year workers are regressed in each year separately by sex on the dummy variables for four education categories, a quartic in experience, black and other race dummies, and interactions of the experience quartic with three broad education categories (high school graduate, some college, and college plus). The (composition-adjusted) mean log wage for each of the 48 groups in a given year is the predicted log wage from these regressions evaluated for whites at the relevant experience level (5, 15, 25, or 35 years depending on the experience group). Mean log wages for broader groups in each year represent weighted averages of the relevant (composition-adjusted) cell means using a fixed set of weights, equal to the mean share of total hours worked by each group over 1963 to 2018 (or 1979-2018, depending on the specification) from the March CPS.

We calculate BA/high school and AA/high school relative supply measures using the March CPS sample. We form a labor "quantity sample" equal to total hours worked by all employed workers (including those in self-employment) with 0 to 39 years of potential experience in 48 gender-education-potential experience cells: experience groups are ten-year categories of 0-9, 10-19, 20-29, and 30-39 years; education groups are high school dropout, high school graduate, some college, associate's degree holder, college graduate, and post-college. The quantity data are merged to a corresponding "price sample" containing real mean full-time weekly (March CPS) wages by year, gender, potential experience, and education. (Wage data used for the price sample correspond to the earnings samples described above.)

Wages in each of the 48 earnings cells in each year are normalized to a relative wage measure by dividing each by the wage of high school graduate males with ten years of potential experience in the contemporaneous year. We compute an "efficiency unit" measure for each gender-experience-education cell as the arithmetic mean of the relative wage measure in that cell over 1964 through 2018 (or 1979-2018). The quantity and price samples are combined to calculate relative log college/high school and log associate's degree/high school supplies. We define the efficiency units of labor supply of a gender-education-potential experience group in year *t* as the efficiency unit wage measure multiplied by the group's quantity of labor supply in year *t*. Following Autor, Katz, and Krueger (1998) and Card and Lemieux (2001), we calculate aggregate college-equivalent labor supply as the total efficiency units of labor supplied by college or college-plus workers plus half of the efficiency units of labor supplied by workers with some college. Similarly, aggregate high school–equivalent labor supply is the sum of efficiency units supplied by high school or lower workers, plus half of the efficiency units supplied by workers with some college. Our BA/high school (and AA/high school) log relative supply index is the natural logarithm of the ratio of BA-equivalent to non-BA-equivalent (or AA/non-AA) labor supply (in efficiency units) in each year. This measure is calculated overall for each year and by ten-year potential experience groupings.

4.8.2 Elasticity Estimates

We then use these measures of relative wages and relative supply to create estimates of how the wage premia will respond to changes in the relative supply from our simulations. Following Acemoglu and Autor (2011), we begin with a constant elasticity of supply production function with two inputs: high-high skill labor (proxied for by those with a BA or more) and low-skill labor (high school graduates),

$$Y = \left[(A_L L)^{\frac{\sigma-1}{\sigma}} + (A_H H)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \qquad (4A.1)$$

where σ is the elasticity of substitution between high- and low-skill labor, and A_L and A_H are factor-augmenting technology terms. We can express the log wage premium as a function of relative supply and technology,

$$\ln w_1 = \ln \frac{w_H}{w_L} = \frac{\sigma-1}{\sigma} ln\left(\frac{A_H}{A_L}\right) - \frac{1}{\sigma} ln\left(\frac{H}{L}\right) \qquad (4A.2)$$

owing for a log-linear time trend for demand of skills, we can then estimate the following equation:

$$\ln w_1 = \ln \frac{w_H}{w_L} = \frac{\sigma-1}{\sigma}\gamma_0 + \frac{\sigma-1}{\sigma}\gamma_1 t - \frac{1}{\sigma} ln\left(\frac{H}{L}\right) \qquad (4A.3)$$

The resulting coefficient on the relative supply term (from the above log-log specification) measures what percent the wage premium will fall for a given percent increase in the relative supply of BA holders.

To estimate the analogous relative supply effects for a change in the AA/HS relative supply, we amend the above two-factor production function to allow for a nest within the "low-skill" input: Associate Degree holders (M) and those with a high school degree (L).

$$Y = \left[\left(A_L \left(\alpha L^{\frac{\eta-1}{\eta}} + \beta M^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \right)^{\frac{\rho-1}{\rho}} + (A_H H)^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}} \qquad (4A.4)$$

where now ρ is the elasticity of substitution between high- and low-skill labor, η measures the elasticity with the low-skill labor nest, and α and β are also factor-augmenting technology terms. We can express the log wage premium as a function of relative supply and technology,

As above, we can express the AA/High School premium as

$$\ln w_2 = \ln \frac{w_M}{w_L} = \frac{\eta - 1}{\eta} ln\left(\frac{\beta}{\alpha}\right) - \frac{1}{\eta} ln\left(\frac{M}{L}\right) \qquad (4A.5)$$

Allowing for a log-linear time trend in demand for skills driven yields

$$\ln w_2 = \ln \frac{w_M}{w_L} = \frac{\eta - 1}{\eta} \delta_0 + \frac{\eta - 1}{\eta} \delta_1 t \; - \frac{1}{\eta} ln\left(\frac{M}{L}\right) \qquad (4A.6)$$

Table 4C.1 presents estimates of equation (1) above for several sample restrictions. The post-1992 interaction is included to allow for an evident trend change in the demand for skills around 1992. Using the same data and methodology as Acemoglu and Autor (2011) and data from 1963 to 2008 (as they do), we are able to replicate their coefficient estimate of -0.644 (reported in Table 8 of their handbook chapter). In this table we extend the data through 2018 and obtain an estimated coefficient on the relative BA/HS supply of -0.712. Because our simulations include only 25-54 year-olds and are restricted to the period from 1979 to 2018, it seems appropriate to restrict the estimating data to that age group and time period. Column (2) uses data back to 1963, but restricts the sample to 25-54 year-olds. Column (3) restricts the estimating data to our population sample and later time period. We incorporate the estimate from Column (3) in our main specifications; it implies that a one percent increase in the relative supply of college graduates (relative to high school graduates) will reduce the wage premium by 0.25 percent.

Table 4C.2 presents estimates of equation (2), the response of the AA/High School or wage premium to changes in the AA/HS-less relative supply. We include the same progression of sample restrictions as before. These estimates are more stable than the BA/HS data above. Column (3), with the preferred sample, suggests that a one percent increase in the AA/HS-less relative supply leads to a 0.18 percent decrease in that wage premium.

Table 4C.1: Bachelors-Plus/High School Relative Wage Response

	1963-2018		1979-2018
	16-64 y.o.	25-54 y.o.	25-54 y.o.
Relative Supply	-0.712***	-0.480***	-0.252***
	(0.0714)	(0.0407)	(0.0858)
Time	0.0302***	0.0213***	0.0198***
	(0.00249)	(0.00142)	(0.00161)
Time x Post-1992	-0.0128***	-0.00814***	-0.0105***
	(0.00158)	(0.00110)	(0.000106)
R^2	0.964	0.970	0.978

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

Table 4C.2: AA/High School Relative Wage Response

	1963-2018		1979-2018
	16-64 y.o.	25-54 y.o.	25-54 y.o.
Relative Supply	-0.0972***	-0.0697***	-0.183***
	(0.0303)	(0.0286)	(0.0496)
Time	0.00510***	0.00429***	0.00571***
	(0.000736)	(0.000698)	(0.000882)
R^2	0.837	0.801	0.733

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

References

Abraham, Katharine, and Melissa S. Kearney. 2020. “Explaining the Decline in the U.S. Employment-to-Population Ratio: A Review of the Evidence.” *Journal of Economic Literature*, 58(3): 585-643.

Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics, Volume 4B*. New York, NY: North Holland. pp. 1043–1171.

ADP Inc., 2020. “CARES SBA: PPP Loan Forgiveness Reports FAQs” Accessed 13 November 2020.https://www.adp.com/-/media/adp/redesign2018/images/customer-service/ppp/ppp-forgiveness-faqs-wfn.pdf?la=en&hash=75E39F54B804AB3B9494E2A977F94C9E07F33FEE

Autor, David. 2014. “Skills, Education, and the Rise of Earnings Inequality Among the `Other 99 Percent’.” *Science* 344(6186): 843–851.

Autor, David, David Cho, Leland D. Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B. Peterman, David Ratner, Daniel Villar, Ahu Yildirmaz (2020). “An Evaluation of the Paycheck Protection Program Using Administrative Payroll Microdata” Working Paper. http://economics.mit.edu/files/20094

Bahr, Peter R., Susan Dynarski, Brian Jacob, Daniel Kreisman, Alfredo Sosa, and Mark Wiederspan. 2015. “Labor Market Returns to Community College Awards: Evidence From Michigan.” CAPSEE Working Paper.

Bartik, Alexander, Marianne Bertrand, Zoe Cullen, Edward L. Glaeser, Michael Luca, Christopher Stanton (2020). “The impact of COVID-19 on small business outcomes and expectations,” *Proceedings of the National Academy of Sciences*, 117 (30): 17656-17666. 10.1073/pnas.2006991117

Bergman, P., R. Chetty, S. DeLuca, N. Hendren, L. Katz, and C. Palmer (2019). Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice. NBER Working Paper No. 26164. https://www.nber.org/papers/w26164.

Bertrand, Marianne, Ester Duflo, and Sendhil Mulainathan (2004). “How Much Should I Trust Differences-In-Differences.” *Quarterly Journal of Economics*. 119(1): 249-275. https://doi.org/10.1162/003355304772839588

Bartik, Alexander, Zoe Cullen, Edward Glaeser, Michael Luca, Christopher Stanton, and Adi Sunderam (2020). "The Targeting and Impact of Paycheck Protection Program Loans to Small Businesses," NBER Working Paper #27623. http://www.nber.org/papers/w27623

Card, David, and Thomas Lemiuex. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? ACohort-Based Analysis." *Quarterly Journal of Economics* 116(2): 705−746.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and the Opportunity Insights Team (2020). "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." Opportunity Insights Working Paper. https://opportunityinsights.org/wp-content/uploads/2020/05/tracker_paper.pdf

Chetty, R., N. Hendren, and L. Katz (2016). The effects of housing assistance on labor supply: Evidence from a voucher lottery. *American Economic Review 106* (4), 855–902.

Code of Federal Regulations (2019). US Code of Federal Regulations: Title 24-791: Housing and Urban Development: Allocations of Housing Assistance Funds. https://www.govinfo.gov/content/pkg/CFR-2019-title24-vol4/xml/CFR-2019-title24-vol4-part791.xml.

Collinson, R. and P. Ganong (2018). How do changes in voucher design affect rent and neighborhood quality. *American Economic Journal: Economic Policy 10* (2), 62–89.

Collinson, R. and D. Reed (2018). The effects of evictions on low-income households. Working Paper. Couch, L. (2016). Housing Choice Vouchers. https://nlihc.org/sites/default/files/ Sec4.12_Housing-Choice-Vouchers_2015.pdf.

Crane, M. and A. Warnes (2000). Evictions and Prolonged Homelessness. *Housing Stud- ies 15* (5), 757–773.

Currie, J. and A. Yelowitz (2000). Are public housing projects good for kids? *Jounal of Public Economics 75* (1), 99–124.

Davis, Stephen, R. Jason Faberman, and John Haltiwanger (2006). "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links." *Journal of Economic Perspectives*. 20(3): 3-26. *https://www.aeaweb.org/articles?id=10.1257/jep.20.3.3*

Deming, David J., and Christopher R. Walters. 2017. “The Impact of Price Caps and Spending Cuts on U.S. Postsecondary Attainment.” NBER Working Paper No. 23736.

Desmond, M. (2016). *Evicted: Poverty and Profit in the American City*. New York, New York: Crown.

Desmond, M. and C. Gershenson (2016). Who gets evicted? assessing individual, neigh- borhood, and network factors. *Social Science Research 64*, 1–16.

Desmond, M., A. Gromis, L. Edmonds, J. Hendrickson, K. Krywokulski, L. Leung, and A. Porton (2018a). *Eviction Lab Methodology Report*. Princeton.

Desmond, M., A. Gromis, L. Edmonds, J. Hendrickson, K. Krywokulski, L. Leung, and A. Porton (2018b). *Eviction Lab National Database: Version 1.0*. Princeton.

Desmond, M. and R. T. Kimbro (2015). Eviction’s fallout: Housing, hardship, and health. *Social Forces 94* (1), 1–30.

Eriksen, M. and A. Ross (2015). Housing vouchers and the price of rental housing. *American Economic Journal: Economic Policy 7* (3), 154–176.

Evans, W., J. Sullivan, and M. Wallskog (2016). The Impact of Homelessness Prevention Programs on Homelessness.

Fischer, W. and B. Sard (2013). *Chart Book: Federal Housing Spending is Poorly Matched to Need*. Washington, DC.

Fowler, K., M. Gladden, K. Vagi, J. Barnes, and L. Frazier (2015). Increase in suicides associated with home eviction and foreclosure during the us housing crisis: Findings from 16 national violent death reporting states, 2005-2010. *American Journal of Public Health 105* (2), 311–316.

GAO (2006). Rental Housing Assistance: Policy Decisions and market Factors Explain changes in the costs of the Section 8 Programs. https://www.gao.gov/products/ gao-06-405.

Hartman, C. and D. Robinson (2003). Evictions: The hidden housing problem. *Housing Policy Debate 14* (4), 461–501.

Flood, Sarah Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2019. *Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]*. Minneapolis, MN: IPUMS. https://doi.org/10.18128/D030.V6.0.

Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race Between Education and Technology*. Cambridge, MA: Harvard University Press.

Goldin, Claudia, and Robert A. Margo. 1992. "The Great Compression: The Wage Structure in the United States at Mid-Century." *Quarterly Journal of Economics* 107(1): 1–34.

Goolsbee, Austan and Chad Syverson (2021). "Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020" *Journal of Public Economics*. 193(104311):1-8. https://www.sciencedirect.com/science/article/pii/S0047272720301754

Hershbein, Brad, Melissa S. Kearney, and Lawrence H. Summers. 2015. "Increasing Education: What It Will and Will Not Do for Earnings and Inequality." *Brookings Up Front*. March 31. Available at https://www.brookings.edu/blog/up-front/2015/03/31/increasing-education-what-it-will-and-will-not-do-for-earnings-and-earnings-inequality/.

Hoekstra, Mark. 2009. "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach." *Review of Economics and Statistics* 91(4): 717–724.

Hubbard, Glenn and Michael Strain (2020a). "Has the Paycheck Protection Program Succeeded?" *Brookings Papers on Economic Activity*. https://www.brookings.edu/bpea-articles/has-the-paycheck-protection-program-succeeded/

Hubbard, Glenn and Michael Strain (2020b). "The Economy Needs a Little More PPP" *Wall Street Journal*. Editorial. Accessed 19 December 2020. https://www.wsj.com/articles/the-economy-needs-a-little-more-ppp-11607469632

HUD (2019a). CDBG Formula and Appropriation Process. https://www.hudexchange. info/onecpd/assets/File/CDBG-Formula-Appropriation-Process-Transcript. pdf.

HUD (2019b). Department of housing and urban development fiscal year 2018 budget. https://www.hud.gov/sites/documents/FY_18_CJS_COMBINED.PDF.

HUD (2019c). *A Picture of Subsidized Households*. Washington, D.C.

Humphries, J. E., N. Mader, D. Tannenbaun, and W. van Dijk (2019). Does Eviction Cause Poverty? Quasi-Experimental Evidence from Cook County, IL. Working Paper.

Jacob, B. and J. Ludwig (2012). The effects of housing assistance on labor supply: Evidence from a voucher lottery. *American Economic Review 102* (1), 272–304.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107(1): 35–78.

Kling, J., J. Leibman, and L. Katz (2007). Experimental analysis of neighborhood effects. *Econometrica 75* (1), 83–119.

Larrimore, Jeff, Richard V. Burkhauser, Shuaizhang Feng, and Laura Zayatz. 2008. "Consistent Cell Means for Topcoded Incomes in the Public Use March CPS (1976-2007)." *Journal of Economic and Social Measurement* 33(2/3): 89‒128.

Manson, S., J. Schroeder, D. V. Riper, and S. Ruggles (2019). IPUMS National Historical Geographic Information System: Version 14.0 [Database]. http://doi.org/10.18128/ D050.V14.0.

Neilson, Christopher, John Eric Humphries, and Gabriel Ulyssea Working (2020). "Information Frictions and Access to the Paycheck Protection Program," NBER Working Paper #27624. http://www.nber.org/papers/w27624

OMB (2019). Historical tables. https://www.whitehouse.gov/omb/ historical-tables/.

Semega, Jessica, Melissa Kollar, John Creamer, and Abinash Mohanty. 2019. "U.S. Census Bureau Income and Poverty in the United States: 2018." *U.S. Census Bureau Report Number p60-266*. Table B-1: https://www.census.gov/data/tables/2019/demo/income-poverty/p60-266.html

Small Business Administration. 2020a. "Paycheck Protection Program: Loan Forgiveness Application".https://www.sba.gov/sites/default/files/2020-05/3245-0407%20SBA%20Form%203508%20PPP%20Forgiveness%20Application%20FINAL_Fillable-508.pdf. Accessed 6 November 2020.

Small Business Administration. 2020b. "Paycheck Protection Program: FAQ's on PPP Loan Forgiveness Application." https://www.sba.gov/sites/default/files/2020-10/PPP%20%20Loan%20Forgiveness%20FAQs%20%28October%2013%2C%202020%29-508.pdf. Accessed 9 November 2020.

Small Business Administration. 2020c. "SBA Paycheck Protection Program Loan Level Data." https://home.treasury.gov/policy-issues/cares-act/assistance-for-small-businesses/sba-paycheck-protection-program-loan-level-data. Accessed 6 November 2020.

Stevens, Ann H., Michal Kurlaender, and Michel Grosz. 2019. "Career Technical Education and Labor Market Outcomes." *Journal of Human Resources* 54(4): 986–1036.

The New York Times. (2020). Coronavirus (Covid-19) Data in the United States. Accessed 18 December 2020, from https://github.com/nytimes/covid-19-data.

U.S. Census Bureau. 2017. Statistics of U.S. Businesses (SUSB). Accessed 9 November 2020. https://www.census.gov/programs-surveys/susb.html

U.S. Treasury Department. 2020. "Paycheck Protection Program Loan Report: 8/8/2020." Accessed 9 November 2020. https://home.treasury.gov/system/files/136/SBA-Paycheck-Protection-Program-Loan-Report-Round2.pdf

Zimmerman, Seth. 2014. "The Returns to College Admission for Academically Marginal Students." *Journal of Labor Economics* 32(4): 711–754.