Essays in Household Finance

Danial Salman

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington
2026

Reading Committee:
Léa Stern, Co-Chair
Stephan Siegel, Co-Chair
Philip Bond
Jarrad Harford

Program Authorized to Offer Degree:
Foster School of Business

University of Washington

Abstract

Essays in Household Finance

Danial Salman

Co-Chairs of the Supervisory Committee:

Léa Stern

Stephan Siegel

Department of Finance and Business Economics

This dissertation consists of two chapters on households' decision-making in anticipation of uncertain financial outcomes. The first chapter analyzes how households, facing uncertainty about future displacements and the value of their human capital, respond to the threat of automation. I link workers' subjective displacement expectations to their direct and social exposure to a disruptive technology: autonomous vehicles (AVs). I find that commercial driver licensing and truck-driving employment fall disproportionately in more AV-exposed areas. Remaining drivers extend their work hours and reduce mortgage market participation relative to less-exposed neighboring drivers. Changes in household spending on alcohol and tobacco products are consistent with heightened automation-induced anxiety. The results indicate that perceived displacement risk affects households' labor supply, credit behavior, and health, informing welfare assessments and policy responses to automation.

The second chapter turns to liabilities on household balance sheets, specifically student loans. We evaluate the immediate consumption response to President Biden's 2022 loan forgiveness announcement, which promised meaningful but uncertain debt relief for approximately 42 million borrowers. Retail stores in counties with a 1 percentage point higher share of eligible student loan borrowers saw a persistent 0.1% increase in weekly sales. This spending response was absent in counties with high shares of financially delinquent households, consistent with delinquent borrowers

being liquidity constrained and unable to smooth consumption. Novel data on debt relief eligibility and applications suggest that student loan borrowers anticipated relief they ultimately did not receive and adjusted their spending in response.

Acknowledgements

I thank my supervisory committee for their guidance and encouragement during my doctoral studies: Léa Stern (co-chair), Stephan Siegel (co-chair), Philip Bond, and Jarrad Harford. I particularly want to thank Léa for the time and energy she spent advising and advocating for me over the past five years. I also want to thank several supportive faculty members and colleagues who overlapped with me at the University of Washington, including but not limited to Mark Westerfield, Jordan Nickerson, Jonathan Karpoff, Fahad Khalil, Jorge Rivero, George Nurisso, Karim Farroukh, Alejandro González, Yvonne Ng, Xuan Xie, Kourosh Ghobadi, Robin Luo, and Hossein Poorvasei.

I also thank Marshall Van Alstyne, Guillermo Lagarda, and Seth Benzell, with whom I first discovered the joy of research at Boston University. I am grateful to Robert King for introducing me to financial economics and for guiding me through the PhD application process. Lastly, I thank Liliana Rojas-Suarez for mentoring me during my time as her research assistant at the Center for Global Development.

Dedication

To my wife, Arianna, my parents, Salman and Nazia, and my siblings, Rumsha and Hamaél.

Contents

Disclaimer

Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business.

The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Chapter 1

End of the Road? Autonomous Vehicles and Displacement Risk

1.1 Introduction

The classic debate between technology doomsayers and optimists has regained relevance. While some commentators argue that artificial intelligence (AI) will broadly displace labor, others maintain that the current wave of innovation resembles earlier episodes of creative destruction, where job creation ultimately outpaced job loss.[1] Although much of this discussion centers on aggregate labor markets, less is known about how perceived displacement risk translates into changes in household behavior. Whether households update their beliefs about impending displacement and take precautionary measures remains unclear, particularly when myopic or wishful thinking might lead them to downplay emerging risks (Beshears et al., 2018; Engelmann et al., 2024) or when their adaptive capacity is limited (Manning and Aguirre, 2025).

My paper analyzes how adversely affected households, facing uncertainty about future displacements and the value of their human capital, respond to the threat of automation. The dynamics of automation-induced labor adjustment are not fully understood. The modern academic literature often employs a task-based framework to study how new technologies affect labor demand, focusing on task displacement, productivity, and reinstatement effects (Acemoglu and Restrepo,

[1]Historically, the creative destruction wrought by the adoption of new technologies favored job creation over job destruction (Autor, 2015). Yet academic perspectives are shifting, with growing concern over the pace and scope of automation (Autor, 2022; Bond and Kremens, 2023; Jones, 2024; Korinek and Suh, 2024; Restrepo, 2025).

2018, 2019; Feigenbaum and Gross, 2024; Hampole et al., 2025; Ager, Goñi and Salvanes, 2026). By contrast, less attention has been paid to the labor supply side of this phenomenon. In principle, rational, forward-looking households may anticipate the looming obsolescence of an occupation and preemptively reduce their labor supply in response. While standard economic theory implies such an anticipatory effect, empirical demonstrations have been scarce and inconclusive.

I provide direct empirical evidence that households respond preemptively to automation threats, thereby highlighting a labor supply channel that has been underexplored in the literature. Focusing on truck drivers—a group of workers for whom displacement risk has been highly salient—I demonstrate that there have already been initial labor market reallocations before actual displacements are realized, i.e., before the widespread deployment of self-driving cars and trucks.[2] The approximately two million heavy and tractor-trailer truck drivers in the United States, specifically those with long-distance driving jobs with few specialized tasks, are most at risk of displacement given that automation of highway driving is a central objective of many developers of autonomous vehicles (AVs) (Viscelli, 2018; Gittleman and Monaco, 2020; Mohan and Vaishnav, 2022).

I rely on two main data sources. First, I obtain the universe of driver licenses in California and New York State. The de-identified data allow me to distinguish commercial truck drivers, including heavy-duty truck operators, using the listed license classes. Second, I hand-collect a panel dataset on large-scale AV testing across zip codes in the United States. Using residential information from the licenses, I compute each driver's direct or indirect exposure to these areas, which uniquely had ubiquitous exposure to self-driving cars. Robotaxis, while not the same as autonomous trucks, are informative about the progress of self-driving technology since information about their deployment and effectiveness may shift truck drivers' perceptions of displacement risk.

As a first step, I document a sharp fall in commercial driver licensing in zip codes exposed to AV testing in California relative to other zip codes in the same county that did not have AV deployments. The point estimates correspond to a 0.6 to 1 percentage point decline in the commercial driver license (CDL) share following AV-exposure, indicating a strong reaction in treated zip codes. I further show that zip codes that had a greater number of AV trips had disproportionately fewer commercial drivers by the end of the sample period.

[2] Fredrick Kunkle, "From road cowboys to robots: Truckers are wary of autonomous rigs," *The Washington Post*, May 23, 2017.

Next, given concerns about endogenous deployments of self-driving cars in specific locations in California and richer data in New York State, I study the role of social networks in influencing licensing behavior beyond AV's direct geographic footprint. Specifically, I link New York drivers, using their license zip codes, to Facebook's Social Connectedness Index, which captures the probability that Facebook users in a pair of locations are friends with each other (Bailey et al., 2018). Leveraging the quasi-exogenous geographical rollout of AV testing, I relate changes in commercial driving to changes in friend exposure to AV, as self-driving cars were deployed in increasingly more areas. My primary result is that the CDL share falls disproportionately in zip codes that are more socially connected to areas with AV testing. Quantitatively, a one standard deviation higher friend exposure induces a reduction in the CDL share by 23 basis points, corresponding to 9% relative to the mean.

In addition, I report heterogeneous effects of friend exposure. I find larger point estimates among younger drivers, for whom the present value implications of displacement risk are greater. These workers are less likely to enter a 'risky' occupation, and conditional on entry, they are more likely to incur the switching costs to exit. I estimate that the CDL share among the young falls by 41 basis points following a one standard deviation increase in friend exposure to AV. This result suggests that observed declines in employment for young workers in occupations exposed to AI may have an important supply component (Brynjolfsson, Chandar and Chen, 2025). I also show that the results are specific to long-haul, heavy truck drivers, not finding similar results for local delivery and bus drivers, who have a lower share of their tasks exposed to AV technology.

I find evidence that labor market adjustments to disruptive technologies take place faster if information frictions are reduced. Comparing the response to AV testing in the metro Phoenix area, I find sharper effects in New York following AV testing in San Francisco. The results align with greater knowledge diffusion across more socially connected regions (Diemer and Regan, 2022), given New York's relatively stronger ties to San Francisco. In placebo tests, I show that friend exposure to metropolitan areas without large-scale AV testing is not associated with statistically significant changes in the CDL share. I also find support for the baseline estimate in an out-of-sample analysis in that the share of employment in truck driving falls disproportionately in areas with higher social exposure to AV. Cumulatively, the evidence indicates that individuals with friends in early AV hot spots reduced their supply of labor to trucking more, relative to neighbors who had weaker ties to these locations.

Using the American Community Survey (ACS), I document complementary evidence on the intensive margin in terms of hours worked per driver and on real earnings, both of which are positively associated with exposure to AV testing. The results are consistent with a precautionary saving motive in that workers who revise their subjective displacement expectations upward work more to accumulate a buffer stock of savings (Deaton, 1991; Carroll, 1997; Low, Meghir and Pistaferri, 2010).[3] Higher real earnings per hour is also consistent with workers demanding increased compensation for displacement risk.

I then show that increased displacement risk may have reduced drivers' willingness to take on a mortgage. Life cycle portfolio choice theory suggests that less wealth would be allocated to risky assets such as stocks and real estate when the value of human capital, a relatively safe asset, falls (Gomes, Jansson and Karabulut, 2024). In this context, drivers may perceive a diminishment in the value of their human capital, not from the normal course of aging, but rather due to heightened displacement risk from the rollout of autonomous vehicles and artificial intelligence more broadly. Alternatively, the result may reflect a shift in households' long-term liquidity management behavior, as they reduce exposure to a fixed payment schedule.

Finally, I document evidence for the proposed mechanism: drivers, receiving information about self-driving cars through their social networks, revise their beliefs about their susceptibility to automation. Higher displacement expectations may be accompanied by an increase in anticipatory anxiety. Using the NielsenIQ Consumer Panel, I show that relatively more exposed households—those who have more friends in AV testing locations and where the head of household has a driving-related occupation—increased spending on alcohol and tobacco products, consumption of which has been linked to anxiety disorders and self-medication of anxiety symptoms (Kushner, Abrams and Borchardt, 2000; Dee, 2001; Morissette et al., 2007; Dávalos, Fang and French, 2012). These results are complementary to other studies that rely on different methods in indicating that many drivers have negative sentiment towards and concerns about automation (Shoag, Strain and Veuger, 2021; Orii et al., 2021; Van Fossen et al., 2023). While some workers may have engaged in wishful thinking, the licensing and employment results suggest that others acted upon their revised forecasts about

[3] Jiang et al. (2025) also find evidence of increased working hours in AI-exposed occupations which they attribute to enhanced productivity and monitoring of workers.

their occupation's automatability. In other words, I find evidence of effects of social networks on both workers' beliefs and behavior.

My findings contribute to the growing literature on the economics of automation. Most closely related to my paper, Cavounidis et al. (2023) develop an overlapping generations model that shows that young workers receive an obsolescence rent for entering a risky occupation, where risk refers to facing a negative labor demand shock with an uncertain date. They document stylized facts consistent with their model for teamsters, dressmakers, and milliners in the early 20$^{\text{th}}$ century, finding mixed evidence regarding the current state of trucking employment.[4] Using a more empirical approach, I provide causal estimates for the anticipatory effects of automation on the supply of truckers in the modern era with heightened concerns about new technologies.

My focus on displacement risk is closely associated with a collection of financial economics papers on labor income risk. One branch of the literature studies wage premia in highly levered firms (Titman, 1984; Berk, Stanton and Zechner, 2010; Agrawal and Matsa, 2013; Graham et al., 2023; Dore and Zarutskie, 2023). In equilibrium, employees receive a compensating differential for unemployment risk associated with working for a financially distressed firm. A second branch links innovative activities to increased displacement risk and falls in workers' human capital (Gârleanu, Kogan and Panageas, 2012; Kogan et al., 2020, 2023). For firms, perceived displacement risk may affect capital structure decisions (Matsa, 2018). Building on induced innovation and technology adoption logic in which labor scarcity raises automation (Zeira, 1998; Acemoglu and Restrepo, 2019; Hémous et al., 2025), the subjective displacement beliefs highlighted here may also generate a novel self-fulfilling feedback loop: as households reduce labor supply to occupations they view as high risk, labor costs in those occupations rise, thereby intensifying firms' incentives to automate.

More broadly, my results add to a literature examining the role of social networks in shaping households' decision making. Researchers estimate network and peer effects in the context of social distancing during the pandemic (Bailey et al., 2024), mortgage refinancing (Maturana and Nickerson, 2019), job finding (Gee, Jones and Burke, 2017), stock market and saving participation (Cannon, Hirshleifer and Thornton, 2024), insurance and tax credit take-up (Hu, 2022; Wilson, 2022), and trading behavior (Hirshleifer, Peng and Wang, 2024).

The remainder of this paper is structured as follows: Section 1.2 outlines the conceptual frame-

[4] Castex, Chow and Dechter (2024) illustrate similar patterns for a broader set of occupations.

work; Section 1.3 discusses the institutional background; Section 1.4 presents the data; Section 1.5 describes the direct effects of AV testing in California; Section 1.6 reports social spillovers in New York State; Section 1.7 provides evidence on the underlying mechanism; and Section 1.8 concludes.

1.2 Conceptual Framework

In this section, I set up a stylized equation, illustrating the key trade-off underpinning the empirical analysis.

Suppose individuals are infinitely lived. Each worker computes the present value of her labor income stream from her current occupation i relative to that from an alternative occupation j. Workers remain in their current occupation if the following inequality holds:

$$\frac{w_i}{r+\tilde{\delta}} \geq \frac{w_j}{r} - c \quad \forall j \neq i. \tag{1.1}$$

Let r denote the risk-free rate. The wage term w can equivalently be thought of as the dividend or yield on human capital, a nontradable asset (Viceira, 2001). I assume that labor income is uninsurable, and I abstract away from a hedging demand given the low share of financial market participation among individuals in my sample (see Table 1.1).

Subjective displacement risk $\tilde{\delta}$ reduces the present value of each worker's current occupation relative to her outside option. The trade-off is starker for entrants because they face no switching costs, denoted by c. Displacement risk in an occupation reflects the uncertainty of an upcoming labor demand adjustment, in terms of capital substituting for labor which would induce a reduction in employment and wages in that occupation. The uncertainty could be in the form of whether or when such an adjustment will occur.

Equation 1.1 demonstrates that increased subjective displacement expectations can generate a leftward shift in labor supply in anticipation, in partial equilibrium. Although a new automation technology could shift displacement risk across several occupations at once, I focus on a single focal occupation for simplicity, treating risk in all other occupations as fixed. Equivalently, $\tilde{\delta}$ can represent relative or incremental displacement risk in the focal occupation.

1.3 Institutional Background

This section provides a brief institutional background relevant for the research designs I implement. For a more comprehensive survey charting the history and technical research on self-driving cars, see Badue et al. (2021). See Viscelli (2016) and Levy (2023) for detailed descriptions of the modern trucking industry. For characteristics of the industry before and after the Carter deregulation specifically, refer to Zingales (1998).

1.3.1 Autonomous Vehicle Testing

For subjective displacement expectations to affect households' decision making, the new labor displacing technology must be salient. Autonomous vehicles (AVs) are one such technology. Self-driving cars and trucks have received much commercial and media attention due to their large potential labor market impacts.[5] A National Academies report predicted that self-driving vehicles would have become widespread by now, with significant effects on the employment of long-haul truckers (National Academies of Sciences, Engineering, and Medicine, 2017). Hype surrounding the diffusion of AV has intensified and waned over the years; however, rapid occupation loss in the transportation sector remains a concern.

Although news and academic articles are sources of public signals that can shift displacement beliefs, I rely on a setting that produces cross-sectional variation in private signals and, subsequently, in displacement beliefs in order to estimate causal effects. While AV testing has a rich history, going back many decades, I focus on the large-scale testing and deployment of robotaxi services in Arizona and California. Waymo (formerly known as the Google Self-Driving Car Project) was the first to offer rides in autonomous vehicles to the general public. Waymo initiated public trials of autonomous ride-hailing in the metro Phoenix area via its Early Rider Program in 2017, launching a commercial service dubbed Waymo One in the following year. In 2021, Waymo launched a similar pilot program in San Francisco, with Cruise (a subsidiary of General Motors) following soon after.[6] In 2024, Waymo began operations in Los Angeles, starting with a restricted number of testers before opening up its ride-hailing service to all residents.

[5] Conor Dougherty, "Self-Driving Trucks May Be Closer Than They Appear," *New York Times*, November 13, 2017.

[6] The California Department of Motor Vehicles suspended Cruise's AV permits on Oct 24, 2023. In December 2024, General Motors shut down Cruise's robotaxi operations.

1.3.2 The Trucking Industry

If you got it, a truck brought it to you! If you got your food, your clothing, your medicine; if you got fuel for your homes, fuel for your industries, a truck brought it to you. The day our trucks stop, America stops!

– Al Pacino as Jimmy Hoffa, former President of the Teamsters Union, in *The Irishman*

The trucking industry is large. Echoing the epigraph to this section, 63% of total freight (by value) was transported by truck in the United States in 2021, including 75% of domestic shipments.[7] The Department of Transportation (DOT) valued the trucking industry as a whole at $389.3 billion, corresponding to 1.7% of gross domestic product in 2021. The Bureau of Labor Statistics (BLS) estimated that there were upwards of 3.5 million individuals in the “driver/sales workers and truck drivers" occupation category in 2022, 86% of whom were truck drivers (BLS, Occupational Employment and Wage Statistics). Heavy and tractor-trailer truck drivers specifically composed almost two million workers, which may be an underestimate considering that a number of self-employed truckers are not counted in the BLS data (Viscelli, 2018). Not only is trucking a large source of employment, it is also geographically widespread, representing the most common occupation in 29 states.[8]

Phares and Balthrop (2022) document that the non-employed compose the largest category of workers who enter the trucking profession. The other top occupations contributing to inflows of new truckers include transportation and material moving workers as well as those in construction and extraction roles. The welfare impact of AVs could be sizable if fewer trucking jobs exist to cushion blue-collar workers against job losses and the accompanying large, persistent effects on earnings and mortality (Sullivan and von Wachter, 2009; Braxton and Taska, 2023; Fallick et al., 2025; Cockriel, 2025).[9]

[7]See *Freight Facts and Figures*, developed by the Bureau of Transportation Statistics, part of the Department of Transportation.

[8]Quoctrung Bui, “Map: The Most Common Job In Every State," *NPR*, February 5, 2015.

[9]At the same time, self-driving technology could lead to a large reduction in the number of lives lost in motor vehicle accidents (Kalra and Groves, 2017).

1.4 Data

1.4.1 Waymo Introduction

Using publicly available maps of Waymo's service areas, I manually identify all zip codes with autonomous ride-hailing accessible to the general public from 2017 to 2024. The data cover 136 zip codes in four metropolitan areas (Figure 1.1). In 2017, self-driving car coverage was limited to zip codes in Chandler, Tempe, and Mesa in the Phoenix metro area. In 2021, coverage expanded to parts of San Francisco, and again in 2022, to downtown Phoenix. By 2024, Waymo's commercial operations had expanded further to Los Angeles and Austin, while the original Phoenix and San Francisco locations grew larger in size.

Communities living within these service areas had ubiquitous exposure to self-driving cars, unlike virtually anywhere else. Indeed, an examination of Google Trends shows that even within metropolitan areas, normalized search volume for "Waymo" concentrates in exactly the places where self-driving cars were first deployed (Appendix Figure B.1).

1.4.2 Commercial Driver Licenses

Given that there is no federal repository of commercial driver licenses (CDLs) in the United States, I obtain records from two state driver licensing agency systems. I rely on Waymo deployments in San Francisco and Los Angeles to estimate direct effects of AV-exposure in California, and deployments in suburban and downtown Phoenix and San Francisco to measure social spillovers in New York State. I omit Los Angeles and Austin from the latter tests due to limited data availability.

The California Department of Motor Vehicles maintains a card history record, consisting of salient details about each license or identification card issued to a given individual. Through a Freedom of Information Act (FOIA) request, I obtained annual counts of active driver licenses at the zip code level from January 1, 2012 to June 9, 2025, with a breakdown by class and type.[10]

These data measure the stock of outstanding driver licenses in each zip code and year. As of 2024, I observe approximately 28 million driver licenses in California, of which approximately 550 thousand are Class A CDLs. Class A license holders are permitted to operate vehicles such

[10] Only one license is considered active for an individual at any given time. A license is considered to be 'active' from its issue date to the day prior to the next license's issue date or until it has been expired for 90 days, whichever comes first, and is counted as the active card for a given year if it was active as of December 31 of that year.

as tractor-trailers with a gross combination weight rating (GCWR) of more than 26,000 pounds, where the towed vehicle's Gross Vehicle Weight Rating (GVWR) is over 10,000 pounds.

In contrast, the New York State data capture new driver license issuances, a flow measure. I acquired a snapshot provided by the New York Department of Motor Vehicles containing 16,406,138 de-identified records from all non-expired driver licenses, permits, or non-driver identification cards as of June 16, 2023.[11] The data include demographics such as the license holder's age, sex, and zip code, and details about the license including class, status, and expiration year. Since licenses are generally valid for 8 years in New York State, I can infer the year of issuance from the expiration year.

The sample is composed of 211,147 truckers who obtained a Class A CDL between 2015 and 2023. The data include all new issuances and do not allow classification of first-time CDL holders versus those who renewed their CDLs. Table 1.1a presents descriptive statistics for Class A CDL holders in New York State who tend to be middle-aged men on average.

In both datasets, I restrict subsequent analyses to zip-years that include at least 5 Class A CDL holders, and I omit zip codes that fail this criterion in any given year. Zip codes can be considered accurate to a point, based on the timeliness with which individuals reported their change of address after a move. I restrict the samples to California and New York State zip codes only. In New York, I omit licenses without full driving privileges, namely those that serve as identification cards only.

1.4.3 Population Survey

The American Community Survey (ACS) is representative of the U.S. population, and it is richer than the licensing data in terms of information collected. Respondents self-report their demographics, education, housing status, occupation, and earnings among other characteristics. I use the 1% samples from the 2012-2021 ACS 1-year files obtained from the Integrated Public Use Microdata Series (IPUMS) database.[12] The ACS microdata are only available at a coarser geography: public use microdata areas (PUMAs). I build a zip code to PUMA crosswalk using Census tracts.[13]

[11] See https://data.ny.gov/.

[12] Due to the effects of the COVID-19 pandemic on data collection, the Census Bureau advises against comparing the ACS 2020 microdata 1-year file to other ACS microdata sample years. I therefore exclude it from the analysis. After 2021, the ACS uses new PUMA definitions, rendering comparisons with earlier sample years highly imprecise.

[13] Census tracts are fully contained within PUMA boundaries and form the fundamental units of PUMA geography. Using the Census provided zip code to tract relationship file, I assign each tract to the zip code where the tract has the largest population share, winsorizing at the left tail at the 1% level. I then use the Census provided tract to

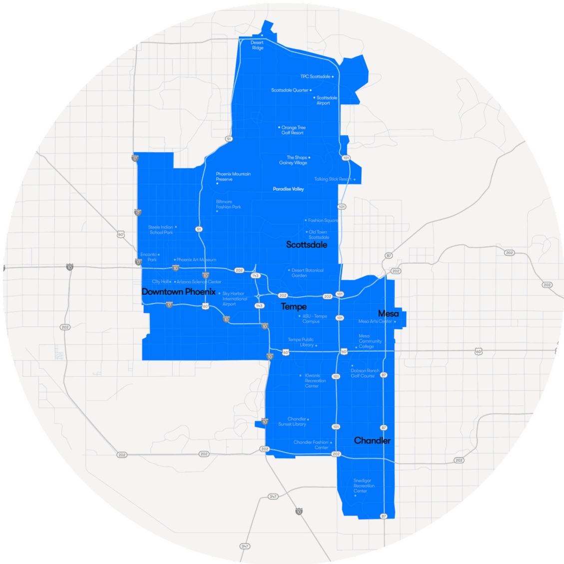

(a) Phoenix

(b) San Francisco

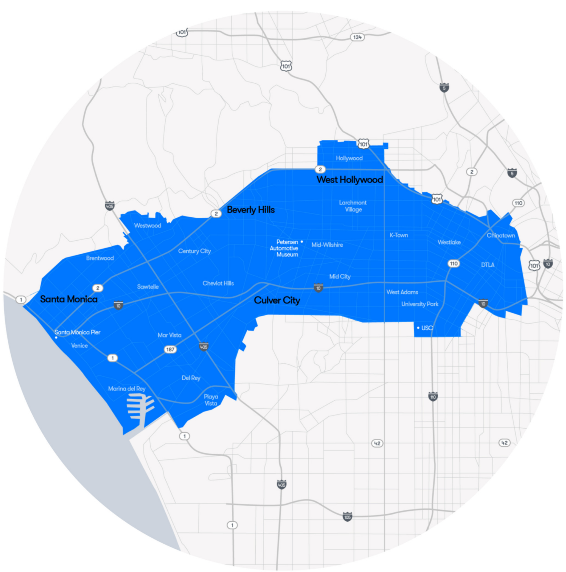

(c) Los Angeles

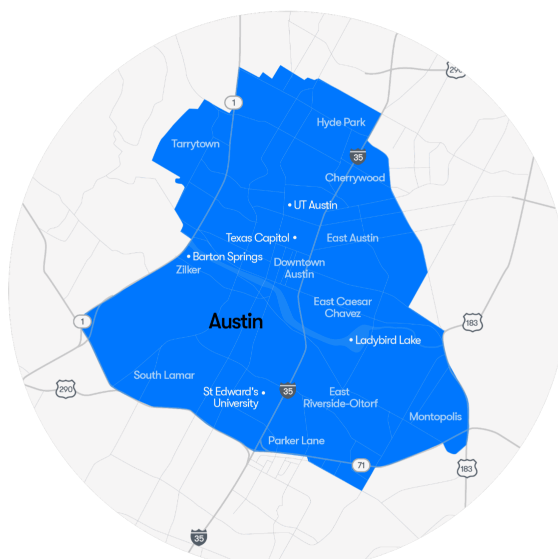

(d) Austin

Figure 1.1: Waymo Service Areas. The figures depict Waymo's service areas as of December 2024.

	Mean	SD	p10	p25	p50	p75	p90	N
Age	46.81	14.62	27.00	36.00	48.00	58.00	65.00	211,147
Female	0.02	0.14	0.00	0.00	0.00	0.00	0.00	211,147

(a) New York State

	Mean	SD	p10	p25	p50	p75	p90	N
Age	47.17	13.91	27.00	37.00	49.00	58.00	64.00	307,059
Female	0.06	0.24	0.00	0.00	0.00	0.00	0.00	307,059
Non-white	0.20	0.40	0.00	0.00	0.00	0.00	1.00	307,059
College	0.12	0.33	0.00	0.00	0.00	0.00	1.00	307,059
Hours worked	44.69	14.44	26.00	40.00	40.00	50.00	60.00	301,929
Real earnings (2010)	36.95	34.02	6.82	17.79	32.29	47.49	65.40	307,059
Mortgage	0.68	0.47	0.00	0.00	1.00	1.00	1.00	219,975
Investment income	0.07	0.25	0.00	0.00	0.00	0.00	0.00	307,059

(b) American Community Survey

	Mean	SD	p10	p25	p50	p75	p90	N
New York CDL share (%)	2.69	2.13	0.73	1.18	2.05	3.69	5.45	8,105
California CDL share (%)	2.81	2.27	0.46	1.12	2.33	3.90	5.72	28,295
ACS Driver share (%)	2.53	1.19	1.14	1.63	2.36	3.24	4.12	19,338

(c) Driver shares

Table 1.1: Summary Statistics. Panel (a) shows summary statistics for Class A commercial driver license holders in New York State. Panel (b) shows summary statistics for driver/sales workers and truck drivers in the ACS. Real earnings are reported in thousands of dollars using 2010 as the base year. Panel (c) shows summary statistics for Class A commercial license issuances in New York State as a share of total driver license issuances at the zip code level, active Class A commercial licenses in California as a share of total active licenses at the zip code level, and employment in the driver/sales workers and truck drivers category in the ACS as a share of the labor force at the PUMA level.

Table 1.1b presents descriptive statistics for driver/sales workers and truck drivers. Relative to heavy truck and tractor-trailer operators in New York State, there are more women represented in the ACS data. The vast majority of the workers in the latter dataset are employed as truck drivers (BLS, Occupational Employment and Wage Statistics). In addition, most of them are high school not college educated. A common misconception is that truckers are highly compensated. The ACS data indicate that this is not the case on average. However, hours worked, and hence

PUMA relationship file to establish the crosswalk.

earnings, may be higher than documented in Table 1.1 given that truckers have strong incentives to underreport them (Viscelli, 2016). Most of them have mortgages outstanding, and few report receiving any income in the form of an estate or trust, interest, dividends, royalties, and rents. It is unlikely that drivers hedge displacement risk by buying stocks negatively correlated with their income prospects, given the low share of those with investment income.

Despite differences in measurement of local trucking employment across datasets, Table 1.1c shows that driver shares range from 2-3% on average. Both the New York and California data measure Class A CDL employment at the zip code level, but the former tracks issuances and the latter captures the stock of active licenses. The ACS, in contrast, measures employment in truck driving as a share of the labor force at the PUMA level.

1.4.4 Social Networks

Since individuals' actual friendship networks are unobservable, I rely on an aggregated measure of social connectedness. Following Bailey et al. (2018), I use the Social Connectedness Index (SCI), computed as of October 2021:

$$\text{SCI}_{i,j} = \frac{FB \text{ Connections}_{i,j}}{FB \text{ Users}_i * FB \text{ Users}_j}, \tag{1.2}$$

where FB Connections$_{i,j}$ is the total number of Facebook friendship links between Facebook users living in location i and Facebook users living in location j. Dividing by the product of Facebook users in each location, the SCI measures the probability that Facebook users in a pair of locations are friends with each other.[14] Previous research has shown that Facebook data are representative of the U.S. population and that they reflect real, offline friendships between people (Bailey et al., 2018). The network is relatively constant over time, capturing long-term, historic connections between people in different geographies (Bailey et al., 2024; Hirshleifer, Peng and Wang, 2024).

[14] The publicly available measure of the SCI that I use here is scaled to have a maximum value of 1,000,000,000 and a minimum value of 1; therefore, it measures relative, not actual, probabilities.

1.4.5 Spending

My final source of data is the NielsenIQ Consumer Panel.[15] This dataset tracks high-frequency household expenditures on nondurable goods such as groceries, cosmetics, and general merchandise. Households use in-home scanners to record all of their purchases intended for personal, in-home use.

The Consumer Panel also contains households' demographic information, including zip code of residence and occupation categories. I focus on respondents where the head of household is in a driving-related occupation as of the 2021 panel year and whose data are non-missing for all months through December 2022. Demographics are collected in the fall prior to each panel year; accordingly, the occupation and zip codes correspond to fall 2020.

The occupation categories are coarse. The one capturing drivers also includes closely related blue-collar occupations, specifically factory/transportation workers and factory machine operators. However, approximately half of the workers within this occupation category are motor vehicle operators based on corresponding entries in the BLS Occupational Employment and Wage Statistics. I further limit the dataset to male respondents within this category to better match drivers in my main sample.

1.5 Direct Effects of AV on Licensing

I first document the direct effects of AV testing on driver-licensing behavior in California. I then turn to an analysis of the social transmission of displacement risk in New York State, where the data are richer and concerns about the endogenous entry of self-driving cars are mitigated. I begin by estimating difference-in-differences specifications of the form:

$$Y_{i,t+1} = \alpha_i + \sum_{t \neq -1} \beta_t \left(\mathrm{AV}_i \times \text{Event Year}_t\right) + \gamma_t' X_{i,t} + \varepsilon_{i,t+1}. \tag{1.3}$$

where $Y_{i,t+1}$ is the number of active Class A CDLs as a share of total active licenses within zip code i as of year $t+1$. α_i represents zip code fixed effects. AV_i is an indicator equal to 1 if Waymo deployed self-driving cars in zip code i over the sample period, and Event Year$_t$ is an indicator

[15] The data is provided by the Kilts Center at the University of Chicago Booth School of Business.

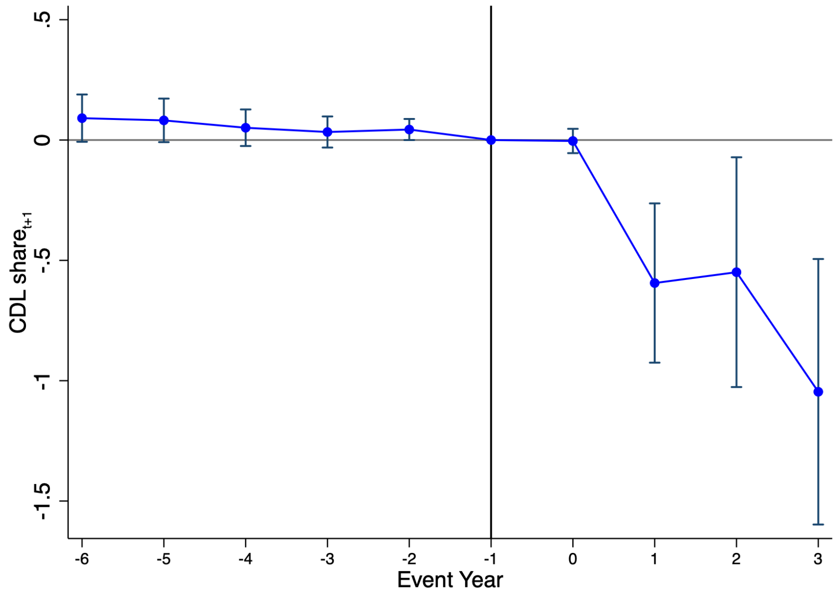

Figure 1.2: AV Testing in California. This figure plots coefficients estimated from a stacked difference-in-differences analysis based on whether AV testing took place within a zip code (treated) as of the sample period. The control units are the never-treated zip codes. The outcome variable is the share of Class A commercial driver licenses in a zip code. The specifications include zip code by stack and year by stack fixed effects and fixed effects for the following groups, interacted with dummies for each year: county, median age, female share, non-white share, bachelor's degree share, population density, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

for the event year of the outcome. The coefficient for the year preceding an event is omitted for reference.

$X_{i,t}$ includes various covariates from the 5-year ACS (2017-2021) for zip code i at year t interacted with year fixed effects. These include the median age, female share, bachelor's degree share, non-white share, population density, median earnings, and friend-weighted density and earnings. The latter two measures are constructed in a similar fashion to equation 1.5 described below. In order to isolate an anticipatory labor supply channel, I also include county-year fixed effects, which absorb variation in local conditions, in particular demand for truck drivers. Effectively, with these fixed effects, I am comparing drivers living in the same county but in different zip codes which have varying levels of exposure to AV.

Given potential issues with staggered difference-in-differences designs when there is treatment effect heterogeneity, Figure 1.2 reports results from a stacked regression where treated zip codes are only compared to never-treated units (Baker, Larcker and Wang, 2022). Estimated treatment effects are near zero in the pre-period and they are negative and statistically significant following the launch of AV testing. The interpretation is that fewer people are renewing their Class A CDLs and that there are fewer entrants into this license class following AV exposure. The point estimates in the post-period correspond to a 0.6 to 1 percentage point decline in the CDL share, suggesting a strong reaction in zip codes that experienced a rollout of AV. I find that estimated dynamic effects are similar under alternative designs and with log outcomes (e.g., Appendix Figures B.2 and B.3).

In Appendix Table B.1, I show evidence that the intensity of AV exposure is associated with larger effects. Using publicly available data from the California Public Utilities Commission (CPUC), I show that among treated zip codes, those with a greater number of AV trips in 2022-2024 had disproportionately lower CDL shares in 2025.[16] Column (1) reports that a 1 standard deviation increase in AV trips is associated with a 20 bp lower CDL share. The correlation persists within county and with additional controls.

A potential concern with the preceding tests is that Waymo's strategic entry decisions in California may be correlated with unobserved changes in local economic conditions, even within counties. For example, the income prospects of residents in particular zip codes may evolve in ways that affect both their labor market opportunities and Waymo's expected profitability. To address this concern, I adopt an alternative empirical strategy. Even in labor markets without active AV trials, workers may receive private signals about automation that can reshape their displacement expectations and, in turn, their occupational choices. Exploiting variation in Facebook-based social connectedness to Phoenix and San Francisco (the first AV testbeds), I trace how AV news propagates through networks to influence licensing behavior in New York State. This social-network-driven approach isolates the impact of subjective displacement risk and extends the analysis beyond AV's direct geographic footprint.

[16]The data can be obtained from https://www.cpuc.ca.gov/regulatory-services/licensing/transportation-licensing-and-analysis-branch/autonomous-vehicle-programs/quarterly-reporting.

1.6 Effects of Friend Exposure to AV on Licensing

1.6.1 Empirical Strategy

The empirical analysis in this section aims to estimate the causal link between occupation and location-specific labor supply and the associated risk of automation-induced displacement.[17] Specifically, I relate the labor supply of truckers within a zip code Y_i to the level of subjective displacement risk in a given period $Disp\tilde{R}isk$:

$$Y_i = \alpha + \delta Disp\tilde{R}isk_i + \gamma' X_i + \epsilon_i. \tag{1.4}$$

Unfortunately, identifying δ from equation 1.4 is challenging for two reasons. First, directly eliciting displacement beliefs is difficult due to selection issues in who chooses to become or remain a truck driver. At any given time, one would need to survey a representative sample of all workers who have considered or dismissed truck driving as a viable occupation. Second, beliefs are typically equilibrium objects; thus, one cannot assume that they are uncorrelated with the levels of the outcome.

To overcome these challenges, an ideal experiment would involve having workers randomly update their displacement expectations (Golin and Rauh, 2022). As this is likely infeasible outside of a laboratory, I instead rely on the quasi-exogenous geographical rollout of AV testing, news of which diffuses differentially along social networks, to identify the impact of increased subjective displacement risk. I construct my social network-based exposure measure, as an instrument for subjective displacement expectations, as:

$$\text{Friend Exposure to AV}_{i,t} = \sum_j \frac{\text{SCI}_{i,j}}{\sum_h \text{SCI}_{i,h}} * \mathbb{1}\left\{\text{AV}_{j,t}\right\}. \tag{1.5}$$

For an individual living in zip code i, her friend exposure to AV at time t is the sum over all zip codes j in the U.S., with an indicator for whether AV testing is occurring, multiplied by her normalized social exposure to those zip codes.[18] The measure is a friend-weighted average of exposure to AV testing.

[17] The exposition of my baseline empirical strategy is similar to Xu (2022).

[18] While the SCI is defined at the ZIP Code Tabulation Areas (ZCTA) level, I refer to zip codes instead here and in all subsequent analyses for simplicity. I also standardize friend exposure for interpretability.

The heat maps in Figure 1.3 illustrate the changing geography of friend exposure in New York State, as increasingly more zip codes were exposed to AV testing via their social networks. In 2017, early hot spots included New York City, Buffalo, Rochester, and Watertown. By 2021, communities in Long Island and Ithaca experienced the largest changes in friend exposure. In 2022, new hot spots included Syracuse and Albany. In aggregate, the largest growth in friend exposure occurred following AV testing in San Francisco and downtown Phoenix, reflecting the high social connectedness of urban areas.

Importantly, the sizable temporal and geographic variation in the maps indicates that different communities in New York State received news about AV from their friends at different times. In other words, in each year, different individuals happen to be most exposed to AV testing. Table B.2 shows the demographic correlates of changes in friend exposure. The magnitudes and the correlations are time-varying. For instance, in 2017, zip codes with higher median earnings had more exposure to AV testing; however, this relationship flipped in later years as testing expanded to San Francisco and downtown Phoenix.

1.6.2 Friend Exposure and Licensing Behavior

I study the relationship between friend exposure to AV and subsequent labor market adjustments by estimating the following reduced-form equation using comprehensive licensing data from New York State:

$$Y_{i,t+1} = \alpha_i + \beta \text{ Friend Exposure to AV}_{i,t} + \gamma_t' X_{i,t} + \epsilon_{i,t+1}. \tag{1.6}$$

I define $Y_{i,t+1}$ as Class A CDL issuances as a share of total license issuances within zip code i as of year $t+1$. α_i represents zip code fixed effects. $X_{i,t}$ captures a range of characteristics for zip code i at year t: time-varying ones such as the mean age and female share from the licensing data, and various covariates interacted with year fixed effects from the ACS. As before, I include county-year fixed effects which absorb variation in local demand conditions.

For β to receive causal interpretation, the identifying assumption is that the time-varying effects of unobservables are not systematically correlated with changes in friend exposure. Since AV deployment in Phoenix and San Francisco over time is likely exogenous with regard to workers'

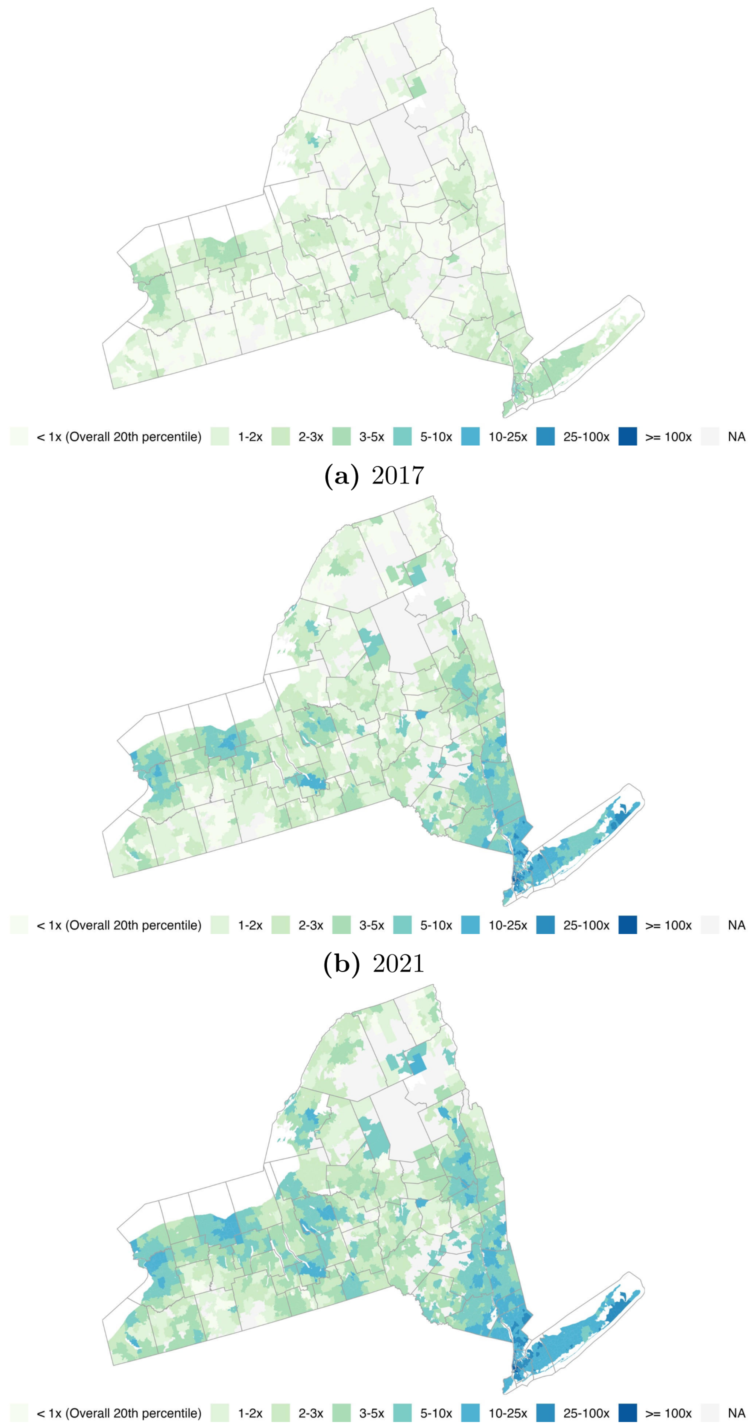

(a) 2017

(b) 2021

(c) 2022

Figure 1.3: Friend Exposure to Autonomous Vehicles. This figure plots zip code level friend exposure to autonomous vehicles (AVs) as defined in equation 1.5 for New York State. Panel (a) depicts friend exposure as of 2017. Panel (b) and (c) do so for 2021 and 2022 respectively. Darker colors correspond to higher friend exposure.

occupational decisions, particularly for those living in distant labor markets, the assumption is a plausible one.

Given the identifying assumption, Table 1.2 presents the results of the exposure research design. In Panel A, column (1), a 1 standard deviation increase in friend exposure leads to a 38 basis point (bp) decline in the Class A CDL share. With the full set of controls in column (2), the point estimate remains approximately the same. Columns (3) and (4) incorporate the zip code specific fixed effects to account for unobservable time-invariant factors in the error term. These specifications are akin to two-way fixed effects (TWFE) regressions, where friend exposure captures the interaction of a post-treatment period indicator and a measure of treatment intensity. The final point estimate is -23 bps, which represents a 9% reduction relative to the mean CDL share.

The effects also appear specific to Class A drivers. In Panel B, after controlling for zip code fixed effects, I do not find analogous results for Class B license holders, who tend to work locally and operate buses or single-unit trucks rather than tractor-trailers. Within a task-based framework, these drivers spend a larger share of their time on activities such as customer service, vehicle inspection, and freight loading and unloading—tasks that are largely unaffected by AV technology and therefore less susceptible to automation. The attenuated point estimates in columns (3) and (4) suggest that Class B license holders responded less to news about AV, consistent with higher reported concerns for long-haul truckers (Viscelli, 2018; Gittleman and Monaco, 2020; Mohan and Vaishnav, 2022).

A reduction in trucking employment could result from either fewer new drivers entering the profession or more existing drivers leaving. In other words, people may be less likely to become truckers, or current drivers might switch to other occupations (or retire). Over a short time frame, changes in the CDL share—and thus coefficients in Table 1.2—are more likely to reflect new entries rather than exits, as commercial drivers typically would not cancel their licenses when leaving the profession, simply allowing them to expire instead. Although approximately 10,500 (5%) drivers voluntarily surrendered their CDLs over the sample period, the results remain robust to removing them (Appendix Table B.3). Additionally, concerns about measurement error in the form of inaccurate zip code information are minimized, since the Department of Motor Vehicles verifies addresses at the time of license issuance.

Importantly, I find that the impact of friend exposure to AV on the CDL share is stronger among

	CDL share$_{t+1}$			
	(1)	(2)	(3)	(4)
Panel A. Class A Licensing				
Friend Exposure$_t$	-0.38***	-0.40**	-0.17**	-0.23**
	(0.14)	(0.16)	(0.070)	(0.089)
Zip Code FE	NO	NO	YES	YES
County-Year FE	YES	YES	YES	YES
License Controls	NO	YES	NO	YES
Census Controls-Year FE	NO	YES	NO	YES
Adj. R-squared	0.16	0.29	0.75	0.75
Observations	2160	2160	2160	2160
Panel B. Class B Licensing				
Friend Exposure$_t$	-0.34***	-0.15***	-0.033**	0.026
	(0.048)	(0.048)	(0.015)	(0.025)
Zip Code FE	NO	NO	YES	YES
County-Year FE	YES	YES	YES	YES
License Controls	NO	YES	NO	YES
Census Controls-Year FE	NO	YES	NO	YES
Adj. R-squared	0.37	0.63	0.81	0.84
Observations	2024	2024	2024	2024

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

Table 1.2: Commercial Driver Licensing in New York State. This table reports the results of regression 1.6. In Panel A, the outcome variable is the share of Class A commercial driver licenses in a zip code. In Panel B, the outcome variable is the share of Class B commercial driver licenses in a zip code. Friend exposure is defined as in equation 1.5. The license controls include mean age and female share. The Census controls by year fixed effects include population density, bachelor share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

younger workers. Specifically, the magnitudes of the estimated coefficients of friend exposure are almost twice as large for the sample of drivers below the median age in 2017 (Table 1.3).[19] Column

[19]The median age of 42 is computed in the unrestricted license dataset. I split the sample to focus on the share of young Class A CDL holders relative to total young drivers which cannot be conveyed with interaction terms in Table 1.3.

	CDL share$_{t+1}$			
	(1)	(2)	(3)	(4)
Friend Exposure$_t$	-0.58*	-0.84**	-0.29*	-0.41**
	(0.34)	(0.35)	(0.17)	(0.19)
Zip Code FE	NO	NO	YES	YES
County-Year FE	YES	YES	YES	YES
License Controls	NO	YES	NO	YES
Census Controls-Year FE	NO	YES	NO	YES
Adj. R-squared	0.10	0.22	0.80	0.80
Observations	1240	1240	1240	1240

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

Table 1.3: Young Drivers in New York State. This table reports the results of regression 1.6. The outcome variable is the share of Class A commercial driver licenses in a zip code among drivers 42 years old and younger. Friend exposure is defined as in equation 1.5. The license controls include mean age and female share. The Census controls by year fixed effects include: population density, bachelor share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

(4) shows that the CDL share among the young falls by 41 bps following a one standard deviation increase in friend exposure to AV. This suggests that younger individuals are more sensitive to perceived displacement risk: they are less likely to enter an occupation when displacement risk rises, and those who do enter are more likely to incur the switching costs to exit. This result is consistent with the predictions of the overlapping generations model in Cavounidis et al. (2023) and may inform interpretations of reduced employment of young workers in AI-exposed occupations as in Brynjolfsson, Chandar and Chen (2025). Alternatively, since younger people are more likely to observe content about self-driving cars online, the coefficients may instead reflect higher effective exposure.

Instead of isolating particular zip codes where AV testing is occurring, an alternative approach uses a broader definition. It is reasonable to consider all zip codes within the Phoenix and San Francisco metropolitan areas as AV testing locations, given that all residents of these cities can potentially interact with self-driving cars. My results are strongly robust to incorporating this broader definition of AV testing in equation 1.5 (Appendix Table B.4). However, I maintain the more granular measure of friend exposure in my main tests because it has more temporal variation (expansion of AV testing in the Phoenix metro area from 2017 to 2022). There may also be a salience factor in that individuals who live within AV testing areas encounter and ride in the vehicles more frequently, making them more likely to communicate news about AV to their social networks.

1.6.3 Dynamics of Licensing Behavior

To get a better sense of how average treatment effects evolve over time, I employ a difference-in-differences methodology within the following TWFE event study specification:

$$Y_{i,t+1} = \alpha_i + \sum_{t \neq -1} \beta_t \left(\text{HighExp}_i \times \text{Year}_t\right) + \gamma_t' X_{i,t} + \varepsilon_{i,t+1}. \tag{1.7}$$

$Y_{i,t+1}$, X_{it}, and α_i are defined as before. HighExp_i is an indicator equal to 1 if zip code i has friend exposure greater than the median friend exposure within its county as of 2017 or 2021, and Year_t is an indicator for the year of the outcome. The coefficient for the year preceding an event is omitted for reference.

I isolate the response to AV testing in Chandler, Tempe, and Mesa in 2017 and in San Francisco in 2021, excluding the downtown Phoenix expansion which may have been anticipated. While estimated coefficients for treatment effects fall following AV testing in metro Phoenix (Figure 1.4a), the decline in estimated coefficients is sharper following AV testing in San Francisco (Figure 1.4b).[20]

This exercise illustrates that labor market adjustments to disruptive technologies can take place faster if information frictions are relaxed. On average, high exposure zip codes experience a relative fall in their CDL shares by approximately 30 bps after the launch of large-scale robotaxi services in San Francisco. The results in Figure 1.4 align with higher social connectedness—and hence faster information diffusion—between urban areas such as New York and San Francisco, compared to New York and suburban communities in Phoenix's East Valley (Diemer and Regan, 2022).

A concern with the previous tests is that the friend exposure measure captures general information gleaned from social networks, such as the state of the economy or advances in AI more broadly,

[20]These tests do not suffer from potential issues related to staggered difference-in-differences designs because all units experience simultaneous shifts in social exposure to AV that differ only in magnitude not timing.

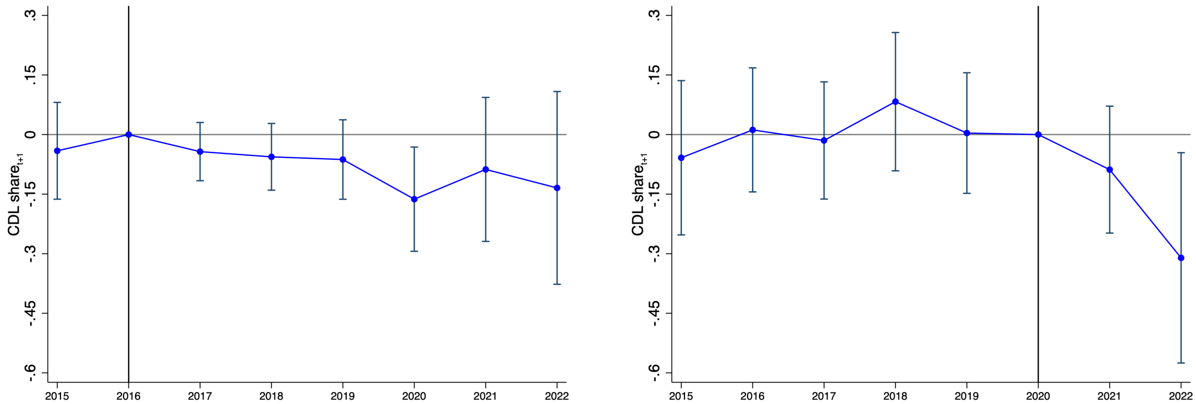

(a) Event: AV Testing in Phoenix Metro Area

(b) Event: AV Testing in San Francisco

Figure 1.4: Difference-in-Differences. The figures plot coefficients estimated from a difference-in-differences analysis based on whether a zip code was above (treated) or below (control) the median friend exposure to AV within its county as of 2017 (Panel a) or 2021 (Panel b). The outcome variable is the share of Class A commercial driver licenses in a zip code. The specifications include age and gender controls, fixed effects for each zip code, and fixed effects for the following groups, interacted with dummies for each year: county, population density, bachelor's degree share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

that can affect licensing behavior, rather than specific news about AV deployments. To address this concern, I conduct placebo tests which compute friend exposure to counterfactual AV testing zip codes. I use nearest neighbor matching without replacement to assign each test location to another within the contiguous United States but outside of Maricopa and San Francisco counties. Appendix Figure B.4 shows that matched zip codes, which did not have large-scale AV testing, are highly similar on observables to zip codes in metro Phoenix and San Francisco which did have AV testing. Given the similarity (and in some cases proximity) of the matched zip codes to the true testing locations, indirect information transmission about AV testing could negatively bias the placebo estimates.

I repeat the difference-in-differences methodology specified in equation 1.7 using the alternative measure of friend exposure. Appendix Figure B.5 presents analogous estimates to those in Figure 1.4. The placebo analyses suggest that specific news about AV matters. Splitting zip codes into treatment and control units on the basis of friend exposure to a "synthetic" Phoenix or San Francisco does not yield statistically significant estimates, despite the potential of negative bias.

Another related concern is that social connectedness may be correlated with unobservable differences between individuals that independently affect licensing behavior following the rollout of AV. Political affiliation may be one example, insofar that it drives friendship links with Phoenix and San Francisco and differential time-varying expectations about the trucking industry (Mian, Sufi and Khoshkhou, 2023). Indeed, while high and low friend exposure zip codes are similar along a number of dimensions, they also differ in terms of the bachelor's degree share and median earnings (Appendix Table B.5). While the placebo tests partially address this concern, I perform further tests below. Following the approach in Bailey et al. (2024), I relate changes in friend exposure to changes in licensing behavior each year as AV testing expanded geographically:

$$\Delta Y_{i,t+1} = \sigma_0 + \sigma_1 \Delta \text{ Friend Exposure to AV}_{i,t} + \sigma'_{2,t} X_{i,t} + \epsilon_{i,t+1}. \tag{1.8}$$

For σ_1 to receive causal interpretation, the identifying assumption is the same as before: the time-varying effects of unobservables are not systematically correlated with changes in friend exposure. The first difference estimator encapsulates the weaker condition of contemporaneous exogeneity rather than strict exogeneity which underlies the previous TWFE specifications. Moreover, this dynamic formulation allows me to focus on the effects of changes in friend exposure, controlling for previous changes in friend exposure.

Column (1) of Table 1.4 shows that a 1 standard deviation increase in friend exposure to AV, induces a 10 bp fall in the CDL share. In columns (2-4), I split the sample to focus on pivotal years and include prior changes in friend exposure as controls. I find that reductions in the CDL share are driven by the most recent increase in friend exposure, primarily following AV testing expansion in San Francisco and downtown Phoenix. For these results to be explained by unobserved factors rather than reflecting a causal impact of friend exposure to AV on the CDL share, one would have to argue that, year after year, zip codes with social ties to regions where AV testing was introduced happened to cut back on their trucking workforce for reasons unrelated to social exposure to labor displacing technologies.

1.6.4 Out-of-Sample Analysis

I find support for the previous findings in an out-of-sample analysis using the American Community Survey (ACS). I first aggregate friend exposure to the PUMA level by summing across relevant zip

	Δ CDL share$_{t+1}$	Δ CDL share$_{2018}$	Δ CDL share$_{2022}$	Δ CDL share$_{2023}$
Δ Friend Exposure$_t$	-0.10**			
	(0.047)			
Δ Friend Exposure$_{2017}$		-0.038	0.089	1.31*
		(0.19)	(0.48)	(0.75)
Δ Friend Exposure$_{2021}$			-0.071**	-0.066
			(0.033)	(0.041)
Δ Friend Exposure$_{2022}$				-0.88**
				(0.44)
County FE	YES x YEAR	YES	YES	YES
Census Controls FE	YES x YEAR	YES	YES	YES
Adj. R-squared	0.16	0.12	0.11	0.14
Observations	1890	270	270	270

Standard errors in parentheses

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

Table 1.4: Changes in Licensing in New York State. This table reports the results of regression 1.8. The outcome variable is the yearly change in the share of Class A commercial driver licenses in a zip code. Friend exposure is defined as in equation 1.5. The Census controls fixed effects include: population density, bachelor share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

codes j using the crosswalk described in Section 1.4:

$$\text{Friend Exposure to AV}_{k,t} = \sum_{j \in k} \text{Friend Exposure to AV}_{j,t}. \tag{1.9}$$

This measure captures how much signal each PUMA k receives collectively about AV; however, my results are robust to taking averages instead. Next, I estimate equation 1.6 in the broader ACS sample, covering all states and the District of Columbia.[21]

As before, I include county-year fixed effects to absorb variation in local conditions and to allow for a comparison between truck drivers living in nearby but different PUMAs with differential friend exposure.[22] Column (1) of Table 1.5 shows that employment in truck driving, as a share of the labor force, falls disproportionately in areas with higher social exposure to AV deployment in Phoenix and San Francisco. Quantitatively, a 1 standard deviation higher friend exposure is

[21]I require each PUMA-year observation to include at least 5 drivers, and I exclude PUMAs that fail this criterion in any given year. All variables are weighted by the corresponding Census provided person and household survey weights.

[22]Geographically large PUMAs, mainly in rural, low-density areas, are omitted automatically because they are not contained within a county.

	Driver share$_t$	Hours worked$_t$ (log)	Real earnings$_t$ (IHS)	Mortgage$_t$
Friend Exposure$_t$	-0.033***			-0.26*
	(0.011)			(0.15)
Friend Exposure$_{t-1}$		0.0071***	0.026***	
		(0.0018)	(0.0061)	
PUMA FE	YES	YES	YES	YES
County-Year FE	YES	YES	YES	YES
Controls	YES	YES	YES	YES
Controls-Year FE	YES	YES	YES	YES
Adj. R-squared	0.48	0.06	0.10	0.36
Observations	7200	6400	6400	7200

Standard errors in parentheses

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

Table 1.5: Labor and Financial Outcomes of Drivers. This table reports the results of regression 1.6. The outcome variables are driver/sales workers and truck drivers as a share of the PUMA labor force, the natural log of the usual hours worked per week in the preceding 12 months, the inverse hyperbolic sine (IHS) of the real earnings over the preceding 12 months, and the fraction of outstanding mortgages. Friend exposure is defined as in equation 1.5. The controls include mean age, female share, non-white share, and college share. The controls by year fixed effects include: population density, median earnings, and network-weighted density and earnings. Standard errors are clustered at the state level.

associated with a fall in the driver share by 3 bps, corresponding to 1.2% relative to the mean driver share in 2016.

Two important caveats about this analysis are that the sample ends in 2021, only one year into San Francisco's large-scale AV testing experiment, which motivated my choice of contemporaneous outcomes instead of leads, and secondly, drivers here belong to a larger occupational group. Although not entirely comparable, these two factors can explain the lower magnitude of the point estimate, relative to those in Table 1.2.

Leveraging the richness of the ACS, I analyze the relationship between friend exposure to AV and additional labor and financial outcomes. Column (2) of Table 1.5 documents a 1% increase in weekly work duration associated with 1 standard deviation higher friend exposure while column (3) shows a corresponding 3% increase in real earnings. Given that a non-trivial fraction of truck drivers are self-employed, I focus on real earnings as opposed to wages.[23] I use the lag of friend

[23] Approximately 30% of workers in the trucking industry are self-employed while 88% of long-distance trucking businesses are nonemployers (Day and Hait, 2019). The Bureau of Transportation Statistics, using Federal Motor Carrier Safety Administration registration data, estimates owner-operators were 11.1% of all truck drivers as of November 2023.

exposure as drivers reported their usual hours worked and real earnings over the preceding 12 months.[24] Finally, column (4) illustrates a negative association between friend exposure and the fraction of drivers with outstanding mortgages.

In combination, the ACS results suggest that the labor supply of truckers responded negatively to the deployment of AV on the extensive margin but increased slightly on the intensive margin. The latter is consistent with a precautionary saving motive in that workers who revise their subjective displacement expectations upward work more to accumulate a buffer stock of savings (Deaton, 1991; Carroll, 1997; Low, Meghir and Pistaferri, 2010). Higher real earnings per hour is also consistent with workers demanding increased compensation for displacement risk (Cavounidis et al., 2023).

Heightened expectations of displacement, as a form of background risk, may also induce a desire to rebalance one's exposure to risky assets (Gomes, Jansson and Karabulut, 2024). The mortgage share result indicates that drivers may have had reduced willingness to invest in real estate and carry a positive mortgage balance following the rollout of AV. At the same time, the result may indicate that households prefer to hold less debt following an increase in labor income risk for liquidity reasons.

1.7 Mechanism

Based on the conceptual framework in Section 1.2, workers switch occupations when subjective displacement risk rises. The preceding analyses use friend exposure to automation technologies as an instrument, in the reduced-form, for such risk to show that workers adjust their labor market behavior in anticipation of future displacements. In this section, I provide support for the first-stage in that workers revise their displacement expectations when exposed to news about automation via their social networks.

Using the NielsenIQ Consumer Panel, I first compute monthly total spending for 2,365 households where the head of household had a driving-related occupation as of the 2021 panel year and whose data are non-missing for all months through December 2022 (the results are robust to this choice). I also compute spending on alcohol and tobacco products specifically, consumption of which has been linked to anxiety disorders and self-medication of anxiety symptoms (Kush-

[24] As a consequence, columns (2) and (3) only reflect the impact of friend exposure to AV testing in Phoenix.

ner, Abrams and Borchardt, 2000; Dee, 2001; Morissette et al., 2007; Dávalos, Fang and French, 2012). Changes in substance use among exposed households following friend exposure to AV may reflect automation-induced anxiety and a revision in households' beliefs about their susceptibility to automation.

I estimate the following equation:

$$Y_{i,k,t+1} = \mu_i + \beta \left(\text{HighExp}_k^{2021} \times \text{Post}_t\right) + \gamma_t' \left(X_i \times \text{Month}_t\right) + \varepsilon_{i,k,t+1}, \tag{1.10}$$

where the outcome variable captures spending by household i located in zip code k in month $t+1$. HighExp_k is defined as before, an indicator equal to 1 if zip code k has friend exposure greater than the median friend exposure within its county as of 2021, while Post_t is an indicator for whether a month falls on or after August 2021, when Waymo launched public AV trials of autonomous ride-hailing in San Francisco. Unlike all the previous estimations which relied on geographical aggregates, I can include household-level fixed effects μ_i in this specification due to the longitudinal nature of the consumer data. X_i includes fixed effects for the county of residence, household size, household income, race, age, education, and network-weighted density and earnings. These are interacted with month fixed effects to absorb time-varying effects of location, household, and network characteristics.

Table 1.6 presents estimates of β, conveying whether high-exposure households have a differential spending response following the initiation of AV testing. Column (1) shows that the coefficient of interest is statistically indistinguishable from zero for total household expenditures on nondurable goods. However, column (2) shows a precisely estimated increase in spending on alcohol and tobacco products. Following the launch of AV testing in 2021, high-exposure households increase spending on these legal substances by $2.6, which is 15% relative to the mean in July 2021.[25] Columns (3) and (4) show that this increase is driven by higher expenditure on both alcoholic beverages and tobacco products separately.

The statistically significant, positive spending response is not present among workers in any other occupation category or in other broad spending categories (Figure 1.5).[26] The result is

[25] The majority of households do not report purchasing alcoholic beverages and tobacco products in a given month. The results are robust to dropping households that never report alcohol or tobacco consumption over the sample period.

[26] I omit occupation group 11 due to limited observations. I further omit occupation groups 7, 9, 10, and 12 which

	Total Spending$_{t+1}$	Alcohol and Tobacco$_{t+1}$	Alcohol Only$_{t+1}$	Tobacco Only$_{t+1}$
High Exposure x Post	6.72	2.60***	1.45**	1.16**
	(14.3)	(0.94)	(0.73)	(0.52)
Household FE	YES	YES	YES	YES
County-Month FE	YES	YES	YES	YES
Controls-Month FE	YES	YES	YES	YES
Mean of Outcome	856.3	17.4	13.5	3.88
Adj. R-squared	0.66	0.64	0.58	0.74
Observations	42918	42918	42918	42918

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

Table 1.6: Spending response to displacement risk. This table reports the results of regression 1.10. The outcome variables are monthly total spending and monthly spending on alcohol and tobacco. High exposure indicates whether a zip code was above (treated) or below (control) the median friend exposure to AV within its county as of 2021. The controls by month fixed effects include: household size, household income, race, age, education, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

thus both occupation- and spending-category-specific: changing either the occupation group or the spending category eliminates the effect.

Complementary to other studies, these results indicate that friend exposure to automation induced an increase in anticipatory anxiety among households in driving-related occupations (Shoag, Strain and Veuger, 2021; Orii et al., 2021; Van Fossen et al., 2023). The observed increases in alcohol and tobacco spending are reminiscent of behavioral responses of the kind emphasized in the broader “deaths of despair" literature (Case and Deaton, 2015).

1.8 Conclusion

The consequences of technological change hinge on how adversely affected households fare. My paper highlights that, far from being passive recipients of technology shocks, these households proactively adjust their labor supply, occupational choices, and financial behavior before jobs are automated. Social networks facilitate the diffusion of information about emerging risks, shaping the timing and magnitude of adjustments across places.

These anticipatory responses may themselves lead to a self-fulfilling prophecy. As households reduce labor supply to occupations they perceive as vulnerable, labor costs in those occupations rise, further strengthening firms’ incentives to automate (Zeira, 1998; Acemoglu and Restrepo,

refer to members of the armed forces, farmers, students, and those outside the labor force respectively.

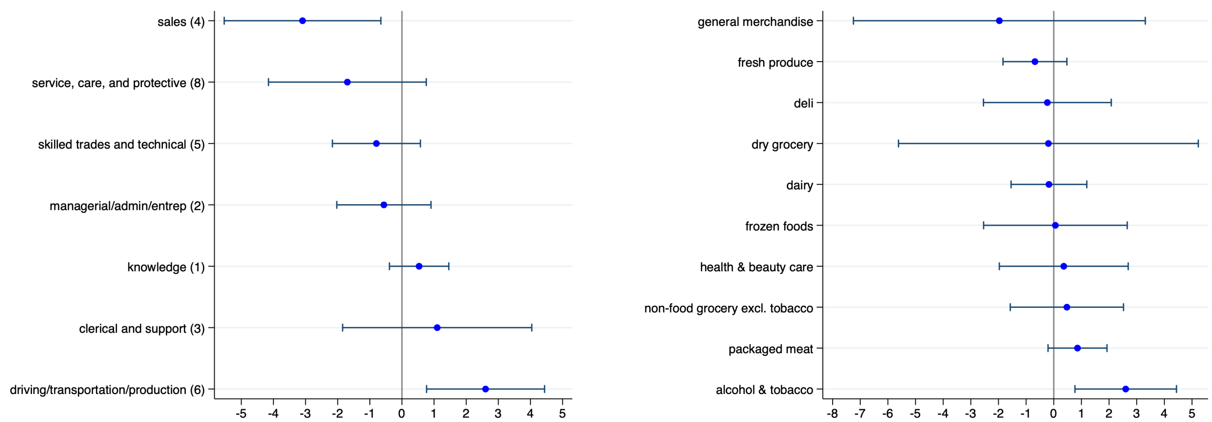

(a) Alcohol and tobacco spending by occupation

(b) Drivers' total spending by category

Figure 1.5: Spending Responses to AV Exposure. Panel A reports estimated coefficients and 95% confidence intervals from regression 1.10 among households by male head occupation category; NielsenIQ Consumer Panel occupation codes are in parentheses. The outcome variable is monthly total spending on alcohol and tobacco products. Panel B reports estimated coefficients and 95% confidence intervals from regression 1.10 on monthly total spending by category among households where the head of household is in a driving-related occupation. Standard errors are clustered at the zip code level.

2019; Hémous et al., 2025). In this way, subjective displacement beliefs not only forecast but also accelerate the pace of innovation and adoption of automation technologies.

Although there are a number of advantages to my empirical approach, it is worth repeating certain limitations. First, individuals' actual friendship networks are unobservable, so I rely on an aggregated zip code level measure, namely the SCI. Second, my measure of friend exposure to AV lends itself to estimating the marginal effect of private signals about AV, not aggregate effects from public signals which may be quantitatively larger.

My results suggest that welfare assessments evaluating the impact of automation may be inaccurate if they fail to account for an anticipatory channel. Moreover, the findings may also have implications for the design and timing of retraining programs and information treatments to better enable households to adapt to automation. An important avenue for further research is to examine the effects of displacement risk on a broader range of occupations, especially those affected by recent advances in artificial intelligence and robotics.

Appendix 1.A Description of Variables

Variables	Description
American Community Survey	
Age	The respondent's age in years as of the last birthday.
Female	Dummy equal to one if the respondent is female.
Non-white	Dummy equal to one if the respondent has not reported "White" race.
College	Dummy equal to one if the respondent has an Associate's, Bachelor's, Master's, Professional, or Doctoral degree.
Hours worked	The number of hours per week that the respondent usually worked in the previous 12 months.
Real earnings	Inflation-adjusted income earned from wages or a person's own business or farm in the previous 12 months.
Mortgage	Dummy equal to one if an owner-occupied housing unit was encumbered by a mortgage, loan, or other type of debt.
Investment income	Dummy equal to one if the respondent reports non-zero pre-tax income from an estate or trust, interest, dividends, royalties, and rents received during the previous year.
Population density	Persons per square mile.
Median earnings	Inflation-adjusted median earnings in the past 12 months for the population 16 years and over with earnings in the past 12 months.
New York State Department of Motor Vehicles	
Age	The license holder's age in years as of license issuance.
Female	Dummy equal to one if the license holder is female.

Appendix 1.B Additional Figures and Tables

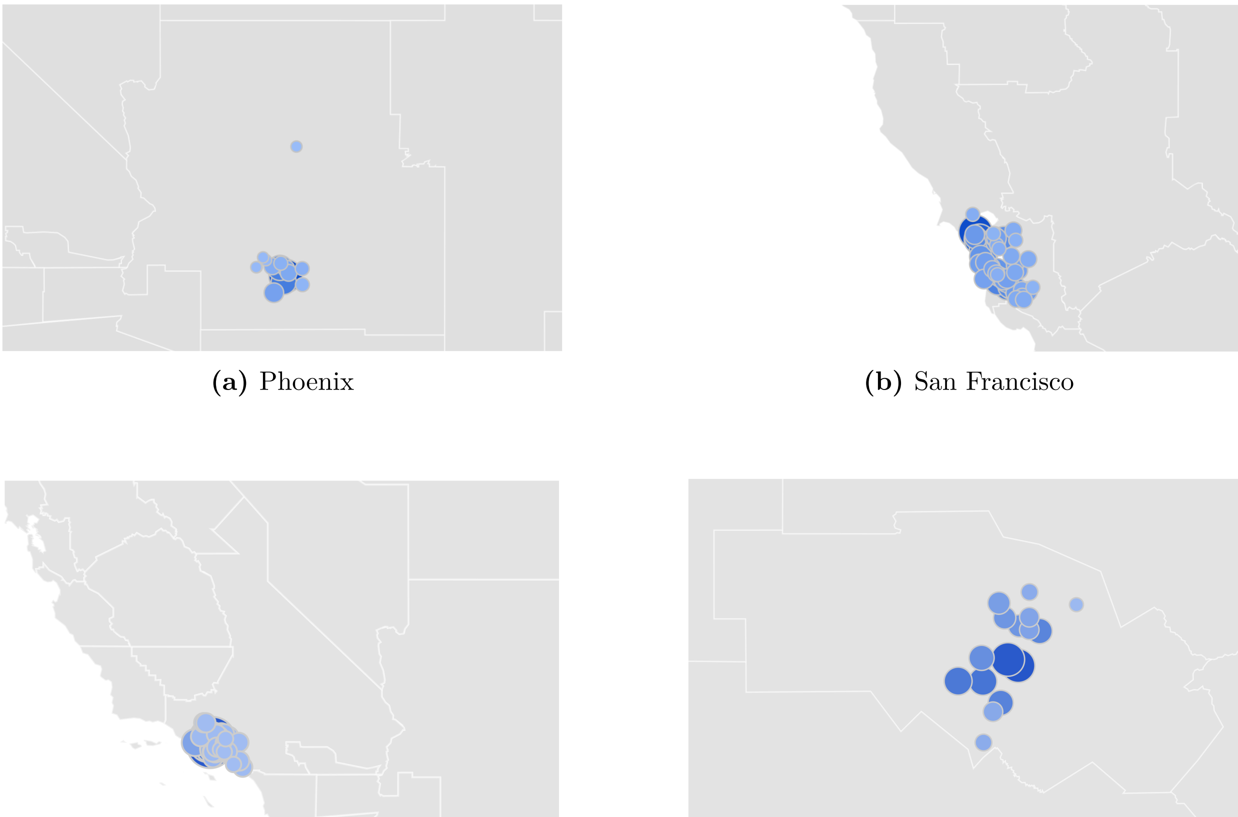

(a) Phoenix (b) San Francisco

(c) Los Angeles (d) Austin

Figure B.1: Google Trends. This figure plots aggregate Google searches for 'Waymo' in the Phoenix, San Francisco, Los Angeles, and Austin metro areas. The data cover the calendar years of 2017 (Panel a), 2021 (Panel b), and 2024 (Panels c & d). Darker colors correspond to higher normalized search intensity.

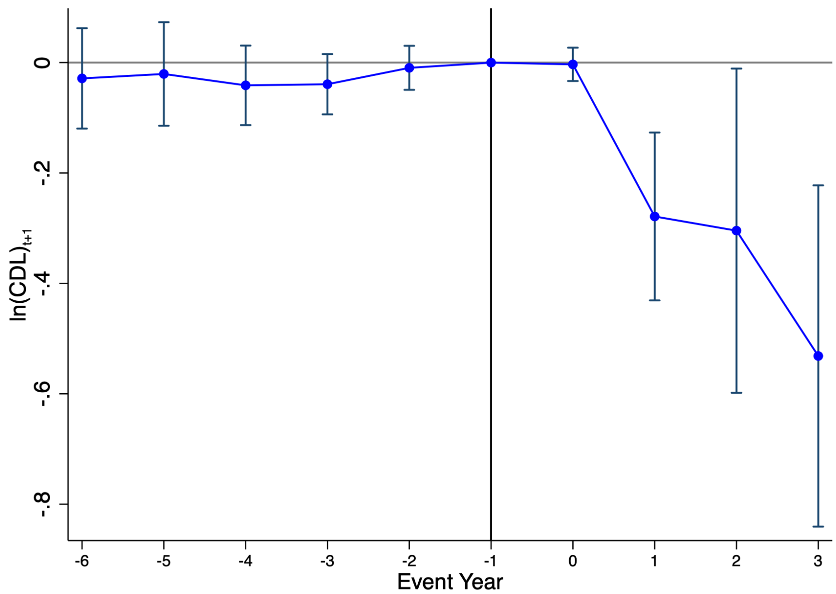

Figure B.2: Log Difference-in-Differences. This figure plots coefficients estimated from a stacked difference-in-differences analysis based on whether AV testing took place within a zip code (treated) as of the sample period. The control units are the never-treated zip codes. The outcome variable is the natural log of the number of Class A commercial driver licenses in a zip code. The specifications include zip code by stack and year by stack fixed effects and fixed effects for the following groups, interacted with dummies for each year: county, median age, female share, non-white share, bachelor's degree share, population density, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

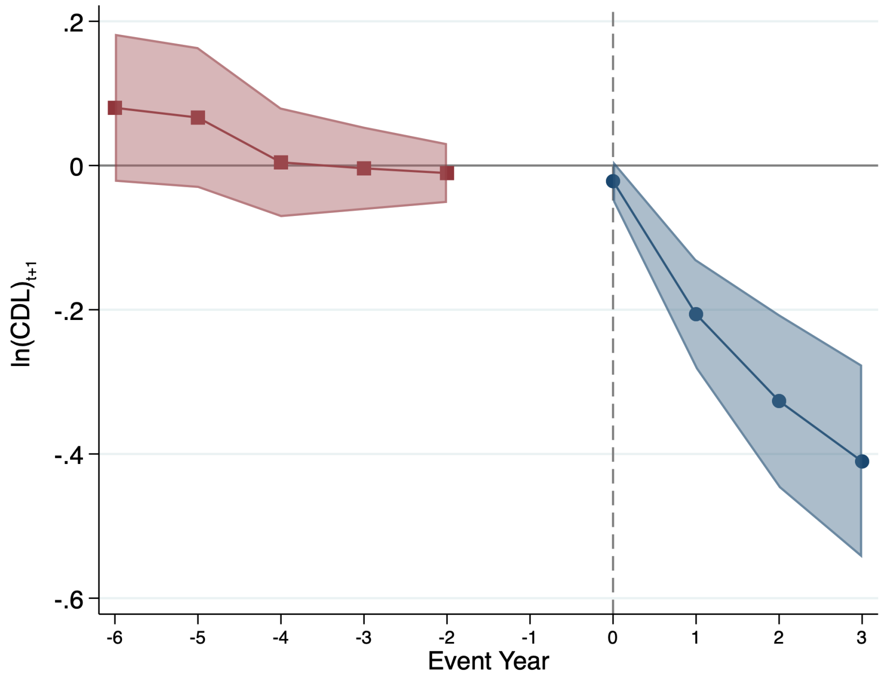

Figure B.3: Alternative Estimator. This figure plots the estimated coefficients and the associated 95% confidence intervals for dynamic treatment effects using the Sun and Abraham (2021) estimator. The treated units consist of zip codes which had AV testing during the sample period. The control units are the never-treated zip codes. The outcome variable is the natural log of the number of Class A commercial driver licenses in a zip code. The specifications include fixed effects for each zip code and fixed effects for the following groups, interacted with dummies for each year: county, median age, female share, non-white share, bachelor's degree share, population density, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

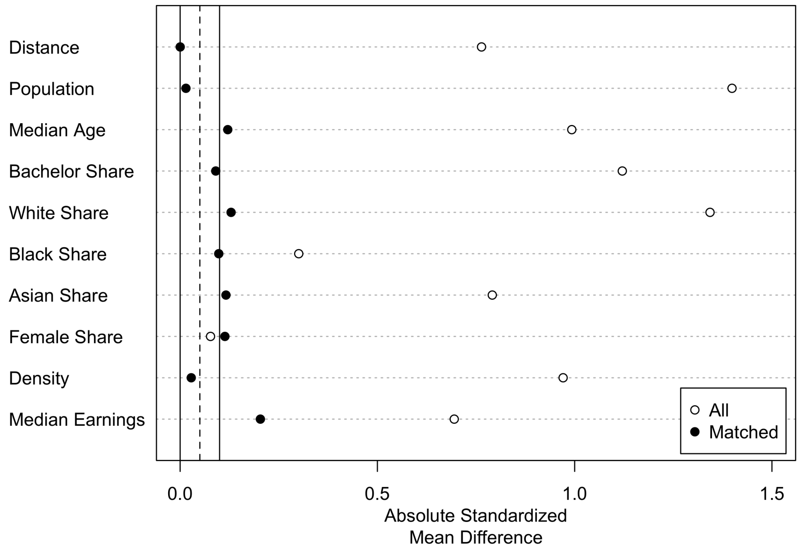

Figure B.4: Balance Plot. This figure reports absolute standardized mean differences in covariates between AV testing locations and matched zip codes as well as between AV testing locations and all zip codes. The matching pool includes all zip codes in the contiguous United States, excluding those in Maricopa and San Francisco counties. The covariates are total population, median age, bachelor share, white share, black share, Asian share, female share, population density, and median earnings. The matching algorithm was 1:1 nearest neighbor matching without replacement on the propensity score. Vertical lines indicate thresholds of 0, .05, and .1.

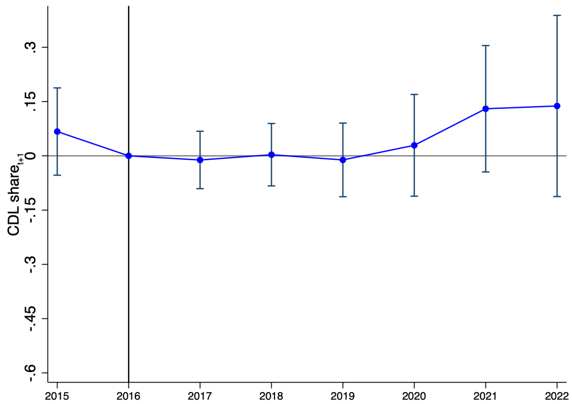

(a) Event: AV Testing in "Synthetic" Phoenix

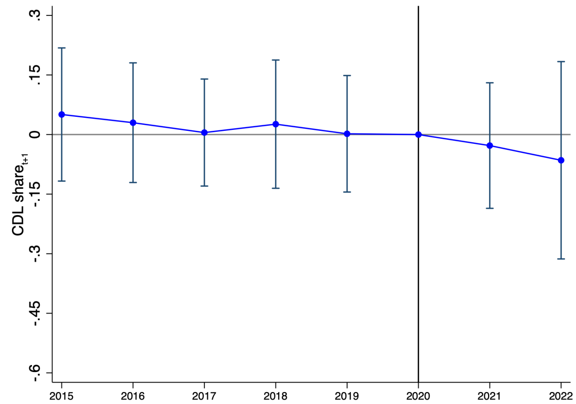

(b) Event: AV Testing in "Synthetic" San Francisco

Figure B.5: Placebo Tests. The figures plot coefficients estimated from a difference-in-differences analysis based on whether a zip code was above (treated) or below (control) the median friend exposure to AV within its county as of 2017 (Panel a) or 2021 (Panel b). Friend exposure is calculated to propensity score matched zip codes relative to those with actual AV testing. The outcome variable is the share of Class A commercial driver licenses in a zip code. The specifications include age and gender controls, fixed effects for each zip code, and fixed effects for the following groups, interacted with dummies for each year: county, population density, bachelor's degree share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

	CDL share$_{2025}$			
	(1)	(2)	(3)	(4)
AV Trips$_{2022-2024}$	-0.20***	-0.16***	-0.16***	-0.11*
	(0.042)	(0.056)	(0.051)	(0.061)
County FE	NO	NO	YES	YES
Census Controls	NO	YES	NO	YES
Adj. R-squared	0.12	0.72	0.10	0.72
Observations	68	68	68	68

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

Table B.1: AV Trips. This table reports cross-sectional regressions at the zip code level of the share of Class A commercial driver licenses in 2025 on the number of Waymo self-driving car trips from 2022 to 2024. AV trips is a standardized measure of the total number of trips ending in each zip code. The controls include median age, female share, non-white share, college share, population density, and median earnings.

	Δ Friend Exposure$_{2017}$	Δ Friend Exposure$_{2021}$	Δ Friend Exposure$_{2022}$
Mean Age	-0.0084***	-0.032***	-0.0093***
	(0.00090)	(0.0080)	(0.0023)
Female Share	-0.46*	3.98	0.16
	(0.26)	(6.38)	(0.43)
Density	0.0038***	0.056***	0.016***
	(0.00045)	(0.0098)	(0.0013)
College	1.57***	21.9***	4.14***
	(0.30)	(3.81)	(0.72)
Non-white	0.15***	-0.46	0.51***
	(0.034)	(0.40)	(0.087)
Real earnings (2021)	0.0045***	-0.027***	-0.0017
	(0.0014)	(0.011)	(0.0026)
Adj. R-squared	0.72	0.69	0.87
Observations	284	284	284

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

Table B.2: Determinants of Change in Friend Exposure. This table reports the demographic correlates of change in friend exposure. Friend exposure is defined as in equation 1.5. Standard errors are clustered at the zip code level.

	CDL share$_{t+1}$			
	(1)	(2)	(3)	(4)
Friend Exposure$_t$	-0.38**	-0.36**	-0.17*	-0.19*
	(0.16)	(0.18)	(0.085)	(0.10)
Zip Code FE	NO	NO	YES	YES
County-Year FE	YES	YES	YES	YES
License Controls	NO	YES	NO	YES
Census Controls-Year FE	NO	YES	NO	YES
Adj. R-squared	0.18	0.30	0.74	0.75
Observations	1960	1960	1960	1960

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

Table B.3: Licensing in New York State, excluding exits. This table reports the results of regression 1.6. The outcome variable is the share of Class A commercial driver licenses in a zip code. I exclude drivers who voluntarily surrender their CDLs during the sample period. Friend exposure is defined as in equation 1.5. The license controls include mean age and female share. The Census controls by year fixed effects include: population density, bachelor share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

	CDL share$_{t+1}$			
	(1)	(2)	(3)	(4)
Friend Exposure$_t$	-0.49***	-0.57***	-0.19**	-0.27**
	(0.19)	(0.21)	(0.086)	(0.11)
Zip Code FE	NO	NO	YES	YES
County-Year FE	YES	YES	YES	YES
License Controls	NO	YES	NO	YES
Census Controls-Year FE	NO	YES	NO	YES
Adj. R-squared	0.16	0.30	0.74	0.75
Observations	2160	2160	2160	2160

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

Table B.4: Licensing in New York State (broad AV exposure). This table reports the results of regression 1.6. The outcome variable is the share of Class A commercial driver licenses in a zip code. Friend exposure is defined as in equation 1.5. The license controls include mean age and female share. The Census controls by year fixed effects include: population density, bachelor share, non-white share, median earnings, and network-weighted density and earnings. Standard errors are clustered at the zip code level.

	Low Exposure			High Exposure			
	Mean	SD	Count	Mean	SD	Count	Diff
Mean Age	35.68	3.37	154	35.84	3.53	130	-0.16
Female Share	0.02	0.03	154	0.03	0.05	130	-0.01
College	0.12	0.04	154	0.15	0.04	130	-0.03
Non-white	0.43	0.31	154	0.42	0.26	130	0.00
Density	16241.45	23538.29	154	18516.83	23515.30	130	-2275.38
Real earnings (2021)	39833.75	9294.35	154	43377.83	8994.27	130	-3544.08
SCI weighted density	6637.42	5633.28	154	7100.82	4995.53	130	-463.40
SCI weighted earnings	43078.77	5540.23	154	44873.63	4870.30	130	-1794.87

(a) 2017

	Low Exposure			High Exposure			
	Mean	SD	Count	Mean	SD	Count	Diff
Mean Age	35.59	3.45	154	35.95	3.43	130	-0.35
Female Share	0.03	0.04	154	0.02	0.04	130	0.00
College	0.12	0.04	154	0.15	0.04	130	-0.03
Non-white	0.43	0.32	154	0.42	0.25	130	0.01
Density	15659.30	22517.95	154	19206.45	24589.81	130	-3547.14
Real earnings (2021)	39903.07	8880.68	154	43295.72	9507.27	130	-3392.64
SCI weighted density	6592.24	5630.58	154	7154.34	4993.61	130	-562.09
SCI weighted earnings	42929.99	5433.53	154	45049.88	4940.97	130	-2119.89

(b) 2021

Table B.5: Summary Statistics of Zip Codes with High and Low Friend-Exposure to AV in 2016. High exposure zip codes have friend exposure to AV greater than the median of their county of 2017 (Panel a) or 2021 (Panel b).

Chapter 2

Relief Beliefs: Effects of Anticipated Student Loan Forgiveness

with Xuan Xie

2.1 Introduction

Student loan forgiveness has been a subject of considerable debate in recent years.[1] Despite growing political support, particularly among those on the left, empirical evidence on the effects of student loan relief is limited. More broadly, whether and to what extent debt relief programs benefit debtors, including by stimulating consumption, remains an open question, especially relative to policies that provide short-term liquidity relief.

In this paper, we evaluate the immediate consumption response to President Biden's August 2022 loan forgiveness announcement which promised debt relief of \$10,000 to \$20,000 for approximately 42 million borrowers through an executive order. The policy announcement represented a meaningful, yet uncertain, wealth shock of 10-20% of median household wealth (Dinerstein, Yannelis and Chen, 2024). In theory, rational consumers would immediately incorporate news about changes in their permanent income. However, if many student loan borrowers were liquidity constrained or hand-to-mouth, they would have been unable to smooth consumption from the wealth shock. Moreover, given the legal challenges to the policy, many borrowers may have disbelieved in the likelihood of actually receiving a reduction in their student debt balances. The Supreme Court eventually blocked the debt relief plan in June 2023.

Although borrowers' liabilities were ultimately unaffected under the plan, we argue that the announcement itself is of independent interest given its scope and the attention it received at the time. Figure 2.1 presents normalized

[1] See, for example, Harris, K. and Trump, D. 2024, 'Presidential debate', ABC, 10 September.

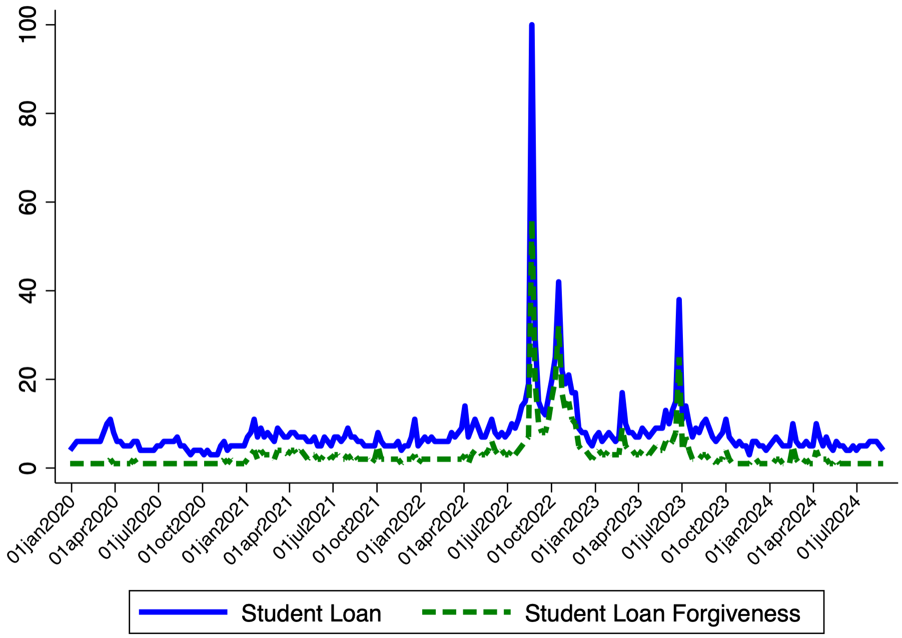

Figure 2.1: Google Trends. This figure plots weekly aggregate Google searches for 'student loan' and 'student loan forgiveness'. The data cover the weeks of December 29, 2019 to August 25, 2024. The counts are normalized and presented on a scale from 0-100, where 100 is the maximum search interest.

search volume from Google Trends for 'student loan' and 'student loan forgiveness'. Biden's announcement stands out as the most noteworthy event, with the latter two spikes occurring when the debt relief application website launched and when the Supreme Court blocked the plan. In contrast, the student loan payment moratorium of March 2020 barely registers despite borrowers having received immediate liquidity relief.

Compared to other studies on the relative merits of debt relief, our setting enables us to isolate the impact of a change in wealth, as the moratorium remained in place until October 2023. This stands in contrast to most other debt relief interventions, such as consumer bankruptcy protection, which typically involve simultaneous payment reductions. Furthermore, unlike other wealth shocks—such as fluctuations in housing or stock market wealth, which may have both permanent and transitory components—student loan forgiveness is entirely permanent, simplifying its interpretation (Dinerstein et al., 2025).

Using an exposure research design, exploiting the geographical variation of student loan borrowers, we show that retail stores located in counties with a 1 percentage point higher share of student loan borrowers, who were eligible for debt relief, saw a persistent 0.1% increase in weekly sales following the announcement. We find that the spending response was absent in counties with high shares of financially delinquent households. Our results indicate that student loan borrowers, who were not liquidity constrained, increased spending immediately upon receiving news about an upcoming, large principal reduction.

Despite the announced debt relief being a small fraction of lifetime permanent income, we find that borrowers were strongly responsive to it. The implied average quarterly marginal propensity to consume (MPC) is 4.5%, which is of the same order of magnitude but larger than typical estimates in the literature for MPCs derived from increases in illiquid wealth (Kaplan and Violante, 2022). Consistent with Baugh et al. (2021), an asymmetric consumption pattern, wherein the response to changes in net worth driven by liability reduction differs from those driven by asset appreciation, suggests a behavioral explanation. Relatedly, the economic magnitude of our estimate may reflect a specific aversion individuals have toward student debt (Gopalan et al., 2024).

To study the effects of student loan forgiveness, we obtained comprehensive data on the number of borrowers eligible for debt relief and their subsequent applications from the Department of Education via Freedom of Information Act (FOIA) requests. We combined this zip code-level dataset with retail scanner data from NielsenIQ to compute each store's differential exposure to President Biden's loan forgiveness announcement. The high-frequency and granular nature of the scanner data allows us to control for location-specific economic trends by employing high-dimensional fixed effects. Effectively, we compare stores belonging to the same chain located in the same metropolitan area at the same time but in different counties with varying levels of exposure to eligible student loan borrowers. Moreover, given the granularity of the scanner data, we document heterogeneous effects across broad product categories, showing positive treatment effects among a wide variety of goods.

Since debt relief eligibility may correlate with unobservable differences between individuals that independently affect spending behavior, we conduct placebo tests using counterfactual debt relief announcements. We do not find statistically significant differential spending responses among retail stores more exposed to eligible borrowers following placebo announcements in the same calendar week of different years. These results suggest that the August 2022 announcement was a pivotal event that shifted borrowers' spending behavior.

Next, we evaluate the immediate consumption response to the Supreme Court decision striking down the Biden administration's debt relief plan in June 2023. In contrast to the previous policy announcement, the Court's decision represented a negative shock to borrowers' net worth, attenuated by the degree to which borrowers expected the Supreme Court justices to rule the plan unconstitutional. We find evidence that borrowers reduced spending immediately following the Supreme Court reversal.

Our results align with a beliefs channel. If borrowers did not expect to eventually receive loan forgiveness—perhaps because they accurately anticipated the Supreme Court's decision—they would not have reacted to the plan's announcement or reversal, nor would they have submitted debt relief applications. However, we find that application counts strongly correlate with student loan forgiveness eligibility. Furthermore, we show that debt relief take-up rates were higher among women, high-income households, and in zip codes that lean more toward the Democratic Party.

Our results are robust to using an alternative measure of student loan borrowers. Leveraging representative credit bureau data on student loan debt, we repeat our empirical methodology and find similar point estimates. In these specifications, we estimate retail stores' exposure to all borrowers with outstanding student loan debt, not just those eligible for the debt relief program. Additionally, our findings remain strongly robust to more stringent cuts of the scanner data, which exclude stores that do not report sales consistently each week.

We add to a literature examining the effects of debt relief interventions. There is an ongoing debate about the merits of balance reduction policies relative to those providing liquidity relief. For instance, Ganong and Noel (2020) document that short-term payment reductions with no change in long-term obligations have significant effects while principal reductions without changes in payments have no effect. Indarte (2023) finds similar evidence for liquidity effects being stronger in consumer bankruptcy filings. Dobbie and Song (2020) find the opposite in terms of positive effects for interest write-downs but not for immediate payment reductions. Cespedes, Parra and Sialm (2026) argue that large principal reductions on mortgage loans have substantial effects even in the short run. We contribute to this literature by showing that balance reductions have non-trivial effects for households that are not liquidity constrained. Relative to the aforementioned papers, our results pertain to government debt relief interventions for which there is a paucity of evidence.

More directly, our findings contribute to the growing literature on student loans.[2] Closely related to our paper, Dinerstein, Yannelis and Chen (2024) study the student loan payment moratorium and the 2022 Biden debt relief announcement, finding effects on credit outcomes during the moratorium but not on borrowers' credit use around the forgiveness announcement. Relative to that paper, we study nondurable retail spending rather than debt-financed consumption on durable goods and use high-frequency scanner data linked to the universe of eligible borrowers.[3] More recently, Dinerstein et al. (2025) examine realized federal student loan discharge, while Di Maggio, Kalda and Yao (2026) study realized private student debt discharge. Relative to this work, we isolate the immediate spending response to a broad, salient, but ultimately unrealized forgiveness announcement.

Finally, we show that consumer spending is responsive to news and to beliefs about future household balance

[2] See Yannelis and Tracey (2022) and Looney and Yannelis (2024) for recent overviews.

[3] An advantage of the Department of Education eligibility data is that they capture all borrowers deemed eligible for relief, including relatively young cohorts who may have been especially exposed to the proposed cancellation. The Department estimated that 65% of eligible borrowers were under age 40 (White House, 2022).

sheets. In this regard, our work is most similar to Garmaise, Levi and Lustig (2024), who demonstrate that household consumption displays excess sensitivity to macroeconomic news, and to Chopra, Roth and Wohlfart (2025), who show that experimentally induced changes in home-price expectations affect household spending. Focusing on political news about household liabilities instead, we document that presidential announcements can shift households' behavior even amid legal uncertainty surrounding forthcoming executive actions. By tracking how spending evolves around both the initial loan forgiveness announcement and the Supreme Court decision, our findings also contribute to the broader literature on economic policy uncertainty and underscore the role of policymakers in shaping expectations (Baker, Bloom and Davis, 2016).

The remainder of this paper is structured as follows. In Section 2.2, we discuss the student loan market and the debt relief plan. Section 2.3 describes the data that we use in our analysis. Section 2.4 presents our empirical strategy and the main findings. Finally, Section 2.5 concludes.

2.2 Institutional Background

The student loan market has experienced tremendous growth over the past few decades. Looney and Yannelis (2024) document that since 2000 total debt balances have more than quadrupled and the number of borrowers has more than doubled. They show that as of 2023, approximately 46 million borrowers held $1.6 trillion in outstanding student loans which represents the largest source of household debt after mortgages. Federal student loans, that is, loans that are disbursed or guaranteed by the federal government, compose the vast majority of outstanding student loan debt. As of August 2023, the share of individuals with a credit bureau record who have any student loan debt is 15%; these borrowers have a median balance of $20,625 and a median payment of $166 (Urban Institute, 2024).

In recent years, several policy and legal changes have impacted the student loan market. We summarize these in Figure 2.2. In this paper, we focus on President Biden's August 2022 loan forgiveness announcement which promised debt relief of $20,000 to borrowers who received a Pell Grant in college with loans held by the Department of Education and $10,000 in debt cancellation to non-Pell Grant recipients.

Eligibility under the program was based on 2020 or 2021 income: less than $125,000 for individuals and less than $250,000 for households.[4] As Pell Grant recipients formed more than 60% of the borrower population, the Department of Education estimated that around 27 million borrowers would be eligible for the full $20,000 in debt relief (White House, 2022). Moreover, the Department estimated that the president's policy would cancel the full

[4]Dependent student loan borrowers' eligibility was based on parental income.

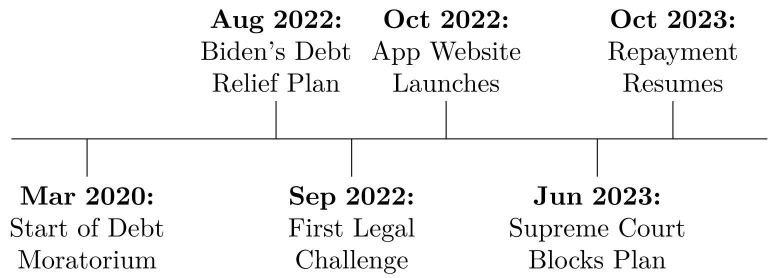

Figure 2.2: Timeline. This figure plots key events impacting the student loan market between March 2020 and October 2023.

remaining balance for roughly 20 million borrowers.

The announcement itself took place during the pandemic-era debt moratorium which froze required student loan payments. The Department of Education launched a website for processing debt relief applications and began verifying individual borrowers' eligibility in October 2022. However, the plan faced immediate political opposition and legal challenges, culminating in the Supreme Court blocking the Biden-Harris Administration from implementing the debt relief measures. Student loan payments resumed in October 2023.

2.3 Data

This section describes the data used in this paper. We construct our main dataset by aggregating the universe of eligible student loan borrowers to the county level and merging it with NielsenIQ's Retail Scanner data. Parts of the analysis use credit bureau data via the Urban Institute. We also construct demographic covariates from a variety of sources outlined below.

2.3.1 Department of Education

The U.S. Department of Education maintained zip code-level counts of eligible student loan borrowers under the proposed Biden-Harris debt relief plan, which we obtained through a Freedom of Information Act (FOIA) request. These counts are rounded, and zip codes with fewer than 10 borrowers are excluded. In total, the data represent about 42 million borrowers, 99% of whom are associated with specific zip codes. The Department did not include data on Pell Grant recipients in its analysis of eligible borrowers. Therefore, we cannot distinguish between those eligible for $10,000 or $20,000 in debt forgiveness at the zip code level.

The Department of Education launched a website on October 14$^{\text{th}}$, 2022 and received approximately 25 million debt relief applications until the website closed down under legal pressure on November 11$^{\text{th}}$, 2022.[5] By this time, the Department had approved two-thirds of the applications, sending them to loan servicers for further processing. From a second FOIA request, we obtained zip code-level counts of the debt relief applications. These counts are rounded, and zip codes with fewer than 100 borrowers are excluded. We observe all 25 million applications, 94% of whom are associated with specific zip codes.

The choropleth maps in Figure 2.3 plot eligible borrowers (Figure 2.3a) and debt relief applications (Figure 2.3b) as a share of the total adult population in each county. The maps convey a high degree of cross-sectional variation. In particular, counties in the South, Appalachia, and parts of the Midwest tend to have high shares of eligible student loan borrowers and applicants. In contrast, much of the Western United States and certain parts of the Northeast exhibit lower shares. Importantly, there is substantial within-state variation in the shares.

2.3.2 Retail Scanner

We measure weekly consumer spending using the NielsenIQ Retail Scanner dataset.[6] It captures a substantial part of retail consumption in the United States and has been used extensively in the economics and marketing literatures with growing use in finance (Dubois, Griffith and O'Connell, 2022).

The scanner data are at the store level. We observe more than 50,000 de-identified retail stores in various categories (food, drug, mass merchandiser, convenience, and liquor) with collective sales of greater than $500 billion in 2022. Appendix Table A.1 reports summary statistics on weekly spending in each retail channel. The dataset contains each store's county of location which we use to link to the counts of eligible student loan borrowers.

2.3.3 Credit Bureau

We source aggregated, representative data on household debt from the Urban Institute.[7] We use a county-level snapshot as of February 2022, 6 months before President Biden's debt relief announcement. Nationally, the share of people with a credit bureau record who have any student loan debt is 15% while their median balance is $20,108 and their median payment is $160. Approximately, one-quarter of individuals with a credit bureau record are financially delinquent in the sense that they have some debt (including student loans and other liabilities like mortgage and credit card debt) in collections.

[5] Borrowers submitted a short online application on https://studentaid.gov/debt-relief/application, by inputting their names, contact information, and Social Security Numbers and by affirming that they satisfied the plan's income eligibility thresholds.

[6] The data is provided by the Kilts Center at the University of Chicago Booth School of Business.

[7] The latest data can be obtained from https://apps.urban.org/features/debt-interactive-map/.

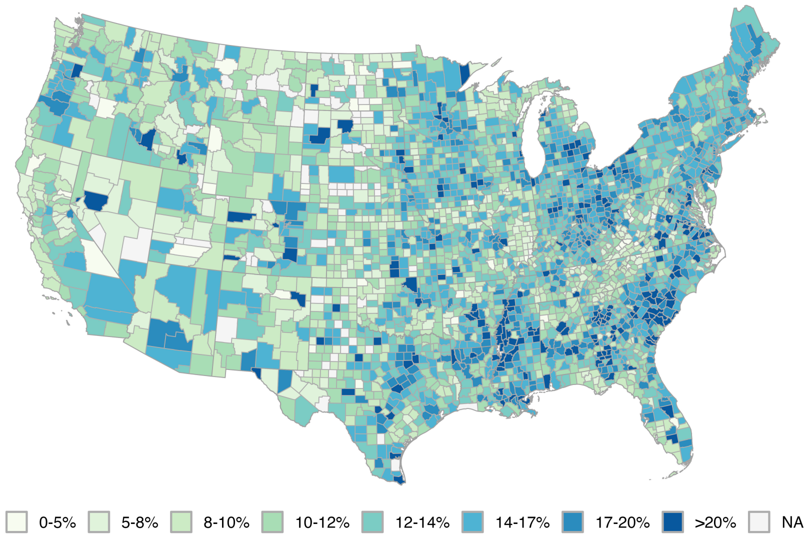

(a) Eligible Share

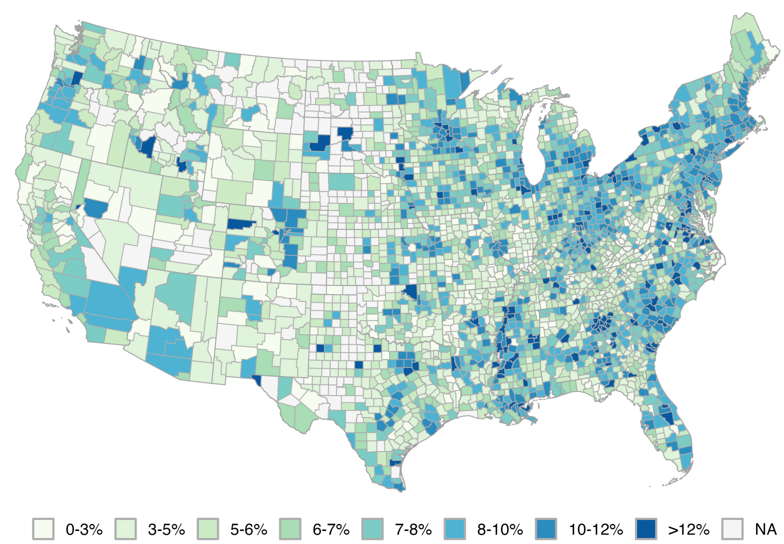

(b) Application Share

Figure 2.3: Distribution of Student Loan Forgiveness. This figure plots the county-level distribution of eligible borrowers and their debt relief applications. Panel (a) shows the share of adults in each county who were eligible for forgiveness under the plan. Panel (b) shows the number of applications as a share of total adults in each county. Darker colors correspond to higher shares.

2.3.4 Demographics

Finally, we supplement our FOIA data with demographic characteristics at the zip code and county level from the American Community Survey 5-Year estimates. We extract college enrollment data from the Integrated Postsecondary Education Data System (IPEDS). We also obtain partisanship information from the MIT Election Lab and the Federal Election Commission (FEC) to identify blue counties and individual campaign donations to the Democratic Party respectively during the 2020 election cycle.[8]

2.4 Results

2.4.1 Consumer Spending

Leveraging geographical variation in eligibility for the debt relief program, we first study the effect of President Biden's loan forgiveness announcement on consumer spending by estimating the following equation:

$$Y_{ikt} = \beta_{DiD} \textit{Eligible Share}_k \times Post_t + \mu_i + \tau_t + \epsilon_{ikt}. \tag{2.1}$$

The outcome is ln(total sales) for store i in county k in week t. We define *Eligible Share*$_k$ as the share of adults in county k who were eligible for forgiveness under the plan, while $Post_t$ is an indicator for whether a week falls around or after August 24, 2022, the date of the announcement. μ_i and τ_t are store and week fixed effects respectively. We cluster standard errors at the county level.

Eligible Shares are equilibrium objects; therefore, we cannot assume that they are uncorrelated with the levels of the outcome. However, identification rests on differential *changes* in store-level outcomes across counties with different levels of exposure. Because all counties were exposed to the announcement and treatment varies only in intensity, our design is a continuous treatment difference-in-differences. Under a standard parallel trends assumption, β_{DiD} captures the differential post-announcement change in log sales across stores facing different county-level exposure to eligible borrowers.[9]

Given this identifying assumption, Table 2.1 presents the results of the exposure research design. Column (1) shows that retail stores located in counties with a 1 percentage point higher share of student loan borrowers eligible for debt relief experienced a 0.2% increase in total sales following the announcement. In column (2), we incorporate

[8] We follow and update Meeuwis et al. (2022) in creating our zip code-level partisanship measure.

[9] A stronger assumption, referred to as strong parallel trends, is required to interpret these cross-dose comparisons causally as a treatment effect of higher versus lower exposure, since treatment-effect heterogeneity across counties with different eligible shares may also contribute to the estimated coefficient (Callaway, Goodman-Bacon and Sant'Anna, 2024).

	(1)	(2)	(3)	(4)
Eligible Share x Post	0.24***	0.22***	0.23***	0.10***
	(0.034)	(0.045)	(0.044)	(0.038)
Store FE	YES	YES	YES	YES
Week FE	YES	NO	NO	NO
Market-Week FE	NO	YES	YES	YES
Parent-Week FE	NO	NO	YES	YES
Channel-Week FE	NO	NO	YES	YES
Demographics-Week FE	NO	NO	NO	YES
Adj. R-squared	0.99	0.99	0.99	0.99
Observations	798240	798030	798015	798015

Standard errors in parentheses

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

Table 2.1: Spending at Retail Stores Post-Announcement. This table reports the results of regression 2.1. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is $ln(sales)$. Eligible Share is the share of adults in a county who were eligible for loan forgiveness, and Post is an indicator for whether a week falls around or after August 24, 2022. The demographic controls include: median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

designated market area-week fixed effects, absorbing variation in local conditions. In column (3), we further include parent company and retail channel fixed effects, each interacted with week dummies. With these high-dimensional fixed effects, we are effectively comparing stores belonging to the same chain, located in the same metropolitan area, and observed in the same week, but operating in different counties with different levels of exposure to eligible student loan borrowers.[10] We find that these coefficients are similar to those in column (1).

In column (4), we introduce county-level demographic controls, each interacted with week fixed effects to allow their effects to vary over time. These include median household income, total population, and the under-18 and college enrollment shares of the population. The latter two measures alleviate concerns related to back-to-school spending which arguably coincides with the announcement. We also include a dummy for blue counties, those who gave Biden a majority of the vote in the 2020 presidential election, to distinguish our estimates from a partisan sentiment effect. Our preferred estimate therefore indicates that retail stores located in counties with a 1 percentage point higher share of eligible student loan borrowers experienced a 0.1% increase in weekly sales following the announcement.

Our results are robust to using an alternative measure of student loan borrowers. Defining retail stores' exposure instead to all borrowers holding any student loan debt based on representative credit bureau data, we reestimate

[10]The NielsenIQ Retail Scanner data is anonymized, and we make no attempt to infer actual store names.

equation 2.1 and find similar results. We recover an identical point estimate in column (1) of Appendix Table A.2 relative to that in our baseline table. With additional fixed effects, the point estimates are somewhat stronger using the broader measure of student loan borrowers, suggesting that more borrowers anticipated relief than were actually eligible for forgiveness as calculated by the Department of Education. Our results are also strongly robust to a more stringent cut of the scanner data, where we omit stores that do not report sales each week (Appendix Table A.3).

Next, we calculate the implied MPC from the spending response to place our findings in context of the broader literature on MPCs out of illiquid wealth, such as housing or stock market equity (Kaplan and Violante, 2022). We start with personal consumption expenditure on nondurable goods of $3.9 trillion among 131.2 million households in 2022.[11] Using the coefficient in column (4) of Table 2.1, we estimate the change in nondurable consumption for the average household as $742 over a quarter. Given that the Department of Education estimated that roughly 27 million borrowers are Pell Grant recipients, eligible for the full $20,000 in loan forgiveness, we calculate an average increase in wealth of $16,429 (White House, 2022). Therefore, normalizing the change in consumption by the change in wealth, our estimate of the average quarterly MPC is 4.5%.[12]

Relative to typical estimates in the literature, which rely on appreciations in asset valuations, our estimate based instead on a reduction in a financial liability is similar but larger. An asymmetric consumption response to different types of changes in net worth suggests a behavioral explanation (Baugh et al., 2021) or a specific aversion to student debt (Gopalan et al., 2024). Permanent liability reductions may boost confidence in a way that asset gains, which may be prone to reversal, do not.

In Table 2.2, we show that the post-announcement spending response is stronger in counties with higher eligible share exposure. Relative to the lowest-exposure counties, stores in counties in the top half of the eligible share distribution saw meaningfully larger post-announcement sales increases. Column (4) reports that stores in the third and fourth quartiles of county eligibility exposure experience 1.2–1.4% larger post-announcement sales increases relative to the lowest quartile. By contrast, the estimate for the second quartile is smaller and statistically insignificant, suggesting that the relationship is nonlinear and that the spending response is concentrated in the upper half of the exposure distribution.

[11] See U.S. Bureau of Economic Analysis and Census Bureau series on https://fred.stlouisfed.org/series/PCND and https://fred.stlouisfed.org/series/TTLHH

[12] We abstract away from legal uncertainty and a potential fiscal multiplier for this exercise.

	(1)	(2)	(3)	(4)
Eligible Share Q2 x Post	0.0064	0.0085	0.0096*	0.0051
	(0.0057)	(0.0052)	(0.0049)	(0.0052)
Eligible Share Q3 x Post	0.012***	0.015***	0.018***	0.012***
	(0.0042)	(0.0041)	(0.0039)	(0.0040)
Eligible Share Q4 x Post	0.023***	0.023***	0.026***	0.014***
	(0.0038)	(0.0044)	(0.0043)	(0.0047)
Store FE	YES	YES	YES	YES
Week FE	YES	NO	NO	NO
Market-Week FE	NO	YES	YES	YES
Parent-Week FE	NO	NO	YES	YES
Channel-Week FE	NO	NO	YES	YES
Demographics-Week FE	NO	NO	NO	YES
Adj. R-squared	0.992	0.992	0.993	0.993
Observations	799028	798818	798803	798803
Mean eligible share: Q1	0.097	0.097	0.097	0.097
Mean eligible share: Q2	0.128	0.128	0.128	0.128
Mean eligible share: Q3	0.153	0.153	0.153	0.153
Mean eligible share: Q4	0.197	0.197	0.197	0.197

Standard errors in parentheses

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

Table 2.2: Spending at Retail Stores Post-Announcement by County Exposure Quartile. This table reports quartile-based difference-in-differences estimates of the effect of President Biden's August 24, 2022 student loan forgiveness announcement on retail sales. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is ln(sales). Counties are assigned to quartiles based on their pre-announcement eligible share, defined as the share of adults in the county eligible for debt relief; quartile 1 (Q1) is omitted, so the coefficients on Eligible Share Q $q \times$ Post for $q \in \{2, 3, 4\}$ are interpreted relative to the lowest-exposure quartile. Post is an indicator for whether a week falls around or after August 24, 2022. Mean eligible shares for each quartile are reported at the bottom of the table. The demographic controls include median household income, total population, partisanship, and the under-18 and college-enrollment shares of the population. Standard errors are clustered at the county level.

2.4.2 Dynamics of Consumer Spending

To get a better sense of how average treatment effects evolve over time, we employ a dynamic difference-in-differences methodology within the following event study specification:

$$Y_{ikt} = \mu_i + \sum_{t=1}^{T} \beta_t \left(\text{Eligible Share}_k \times \text{week}_t\right) + \sum_{t=1}^{T} \delta_t' \left(X_i \times \text{week}_t\right) + \varepsilon_{ikt}. \tag{2.2}$$

Y_{ikt}, *Eligible Share*$_k$, and μ_i are defined as before. $X_i \times \text{week}_t$ reflects all of our fixed effects from the last specification

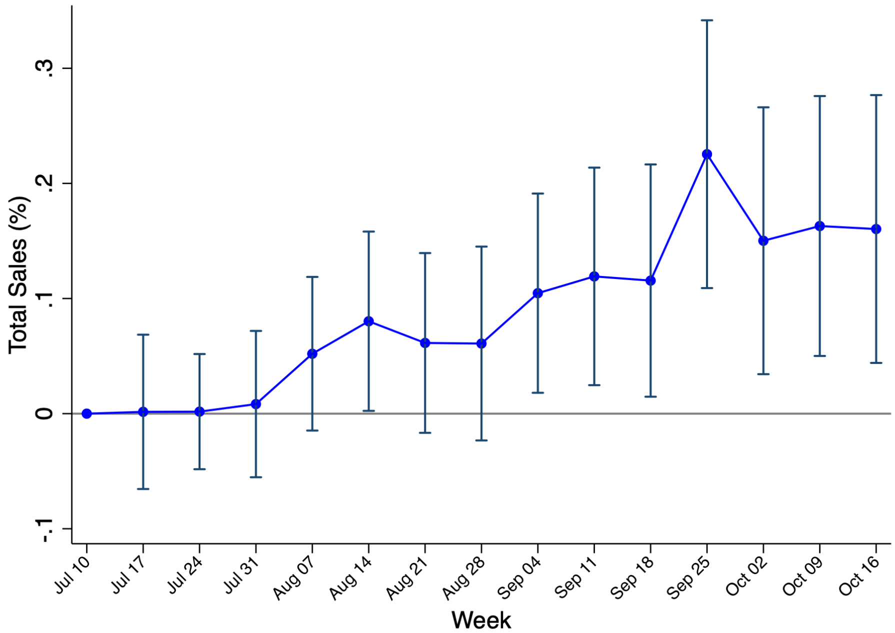

Figure 2.4: Dynamics of Spending Response. This figure reports estimated coefficients and 95% confidence intervals from regression 2.2. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The specification includes fixed effects for the following groups, interacted with dummies for each week: designated market area, parent company, retail channel, median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

in Table 2.1. We include data for the week of July 10 as $t = 0$, but omit a coefficient for this reference time period. Figure 2.4 shows that estimated treatment effects are near zero in the pre-period and that they are positive and statistically significant following the announcement, staying elevated around 0.1-0.2%.

Relative to estimates from July, the point estimates for the weeks of August 7 and August 14 suggest anticipation or policy leakage. A closer look at Google Trends in the days and weeks leading up to the announcement reveals clear spikes in search activity (Appendix Figure A.1). At the same time, news articles speculated about upcoming student debt relief, a key promise of President Biden's election campaign.[13] Notably, the spike in searches on August 16 coincides with the Department of Education discharging debt for approximately 200 thousand students of ITT Technical Institute which may have signaled the much broader executive action we study that involved debt cancellation for over 42 million borrowers.[14]

[13] Ann Carrns, "What to Know as the Pause on Student Loans Is Set to Expire," *New York Times*, August 12, 2022.

[14] Stacy Cowley, "Education Department wipes out $4 billion in ITT Tech student loans," *New York Times*, August 16, 2022.

Appendix Figure A.2 shows analogous event-study estimates by quartile of county eligible share exposure. The dynamic pattern is consistent with the static quartile results: post-announcement spending responses are concentrated in the upper half of the exposure distribution, while differences between the third and fourth quartiles are modest.

A concern with the previous tests is that eligibility may be correlated with unobservable differences between individuals that independently affect spending behavior following the announcement. For example, consumer types may not be uniformly distributed across counties even within a metropolitan area, with more urban counties having higher exposure to eligible borrowers and exhibiting different seasonal spending patterns. Indeed, while high and low exposure counties are similar along a number of dimensions, they also differ in terms of median household income, education, and partisanship (Appendix Table A.4).

To address this concern, we conduct placebo tests which convey differential spending responses to counterfactual debt relief announcements. We repeat the difference-in-differences methodology specified in equation 2.2, with the event window now reflecting announcements in the same calendar week but in different years. Appendix Figure A.3 presents analogous estimates to those in Figure 2.4. The placebo tests generally do not yield a similar pattern nor statistically significant estimates, suggesting that the August 2022 announcement was a pivotal event that shifted borrowers' spending behavior.

2.4.3 Liquidity Constraints

For debt relief expectations to affect households' decision-making, it is crucial that borrowers have sufficient financial flexibility to smooth consumption across time. To investigate the role of liquidity constraints, we conduct a triple difference-in-differences exercise, estimating specifications of the form:

$$\begin{aligned} Y_{ikt} = {} & \beta_{DDD} \textit{Eligible Share}_k \times \textit{Post}_t \times \textit{High Delinquency}_k \\ & + \beta_2 \textit{High Delinquency}_k \times \textit{Post}_t \\ & + \beta_3 \textit{Eligible Share}_k \times \textit{Post}_t + \mu_i + \tau_t + \epsilon_{ikt}. \end{aligned} \tag{2.3}$$

We introduce $\textit{High Delinquency}_k$ which is an indicator for whether counties have a share with any debt in collections greater than 26%, the national rate. Our hypothesis is that due to liquidity constraints, financially delinquent households would be less able to smooth consumption following President Biden's debt relief announcement. Therefore, our parameter of interest is β_{DDD} which captures the differential effect of exposure to eligible borrowers on retail stores in high-delinquency counties relative to those in low-delinquency counties following Biden's loan forgiveness announcement.

In line with our interpretation, column (3) of Table 2.3 shows that the positive spending response is absent in high-delinquency counties as β_{DDD} completely offsets β_3. For retail stores located in low-delinquency counties, we estimate that a 1 percentage point higher share of eligible student loan borrowers led to a 0.3% increase in weekly sales following the announcement.[15] This pattern is also consistent with Dinerstein, Yannelis and Chen (2024), who find stronger responses to student loan relief among borrowers without prior delinquencies. Together with Eaglin et al. (2025), our estimates suggest that the effects of debt relief depend importantly on borrowers' underlying financial slack.

For comparison purposes, we conduct similar exercises for high student loan debt and high student loan payment counties in columns (1) and (2), where the thresholds are defined relative to national medians. We find that the spending response was higher in counties with relatively lower outstanding student loan balances, where a greater number of borrowers stood to receive cancellation of their entire debt balances. In contrast, we do not observe statistically significant, differential spending responses between counties with higher rather than lower student loan payment amounts, consistent with the policy affecting expected debt balances rather than contemporaneous liquidity.

2.4.4 Spending by Broad Product Category

The preceding analyses document the effects of the debt relief announcement on nondurable consumption as a whole. To better understand the composition of this response, we next exploit the granularity of the scanner data and estimate our baseline specification separately across broad product categories.

Figure 2.5 displays the resulting category-level treatment effects. The pattern is mixed but generally tilted toward higher spending in the post-announcement period. Several categories exhibit positive estimates of comparable magnitude to the baseline effect, including grocery, non-food grocery, health and beauty care, non-alcoholic beverages, and baby products. The largest positive response appears in alcoholic beverages and mixers. By contrast, general merchandise shows a negative estimate.

Taken together, the figure suggests that the overall increase in retail spending is not driven by a single narrow class of goods. Instead, positive responses appear across several everyday consumption categories, consistent with an increase in day-to-day household spending following the announcement. At the same time, the dispersion in estimates

[15] We also observe a small but statistically significant time trend among high-delinquency counties relative to low delinquency ones. This could be due to different seasonal spending patterns between the two groups.

	(1)	(2)	(3)
Eligible Share x Post	0.17***	0.075	0.32***
	(0.050)	(0.046)	(0.085)
High SL Debt x Post	0.022*		
	(0.013)		
Eligible Share x High SL Debt x Post	-0.12*		
	(0.068)		
High SL Payment x Post		-0.021	
		(0.016)	
Eligible Share x High SL Payment x Post		0.100	
		(0.087)	
High Delinquency x Post			0.057***
			(0.016)
Eligible Share x High Delinquency x Post			-0.33***
			(0.088)
Store FE	YES	YES	YES
Week FE	NO	NO	NO
Market-Week FE	YES	YES	YES
Parent-Week FE	YES	YES	YES
Channel-Week FE	YES	YES	YES
Demographics-Week FE	YES	YES	YES
Adj. R-squared	0.99	0.99	0.99
Observations	781447	770410	797893

Standard errors in parentheses

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

Table 2.3: Triple Differences. This table reports the results of regression 2.3. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is $ln(sales)$. Eligible Share is the share of adults in a county who were eligible for loan forgiveness, and Post is an indicator for whether a week falls around or after August 24, 2022. High Delinquency is an indicator for whether the share with any debt in collections in a county is greater than 26%. High SL Debt is an indicator for whether the median student loan debt balance in a county is greater than \$20,108. High SL Payment is an indicator for whether the median student loan debt payment in a county is greater than \$160. The demographic controls include: median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

indicates meaningful heterogeneity in how households adjusted purchases across types of goods. Because our outcome is sales, the estimated increases may reflect changes in quantities purchased, shifts toward higher-priced varieties, or both.

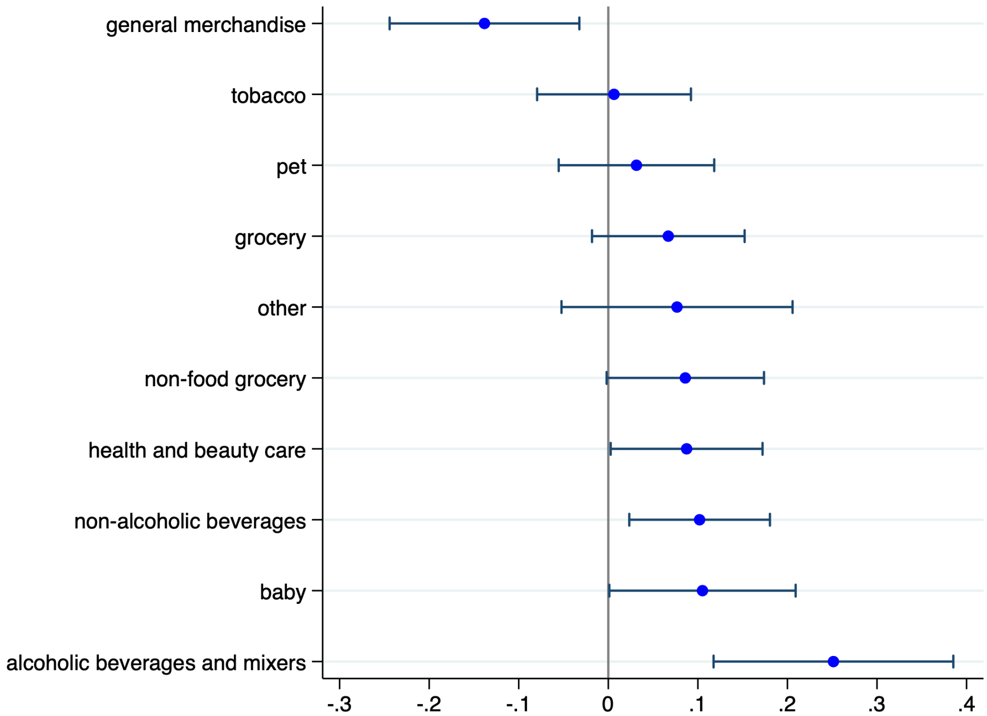

Figure 2.5: Heterogeneous Effects by Spending Category. This figure plots estimated coefficients and 95% confidence intervals from regression 2.1, estimated separately for each product category type. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is $ln(sales)$. The specification includes fixed effects for the following groups, interacted with dummies for each week: designated market area, parent company, retail channel, median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

2.4.5 Supreme Court Decision

In this subsection, we study the impact of the Supreme Court decision on consumption. In June 2023, the Supreme Court ruled the Biden administration's debt relief plan unconstitutional.[16] Ensuring that no borrowers received loan forgiveness under the plan, the Supreme Court decision represented a negative wealth shock to the extent that legal uncertainty about the plan's future remained. The justices voted 6-3 in *Biden v. Nebraska* holding that the administration overstepped its authority under the Higher Education Relief Opportunities for Students Act of 2003.

Figure 2.6 documents that retail stores located in counties with a 1 percentage point higher share of eligible student loan borrowers experienced a sharp drop in weekly sales of around 0.05% at the time of the Supreme Court decision. The negative impact is statistically significant at the 10% level. The lower economic magnitude, relative to

[16]For legal details, see Amy Howe, "Supreme Court strikes down Biden student-loan forgiveness program," *SCOTUSblog*, June 30, 2023.

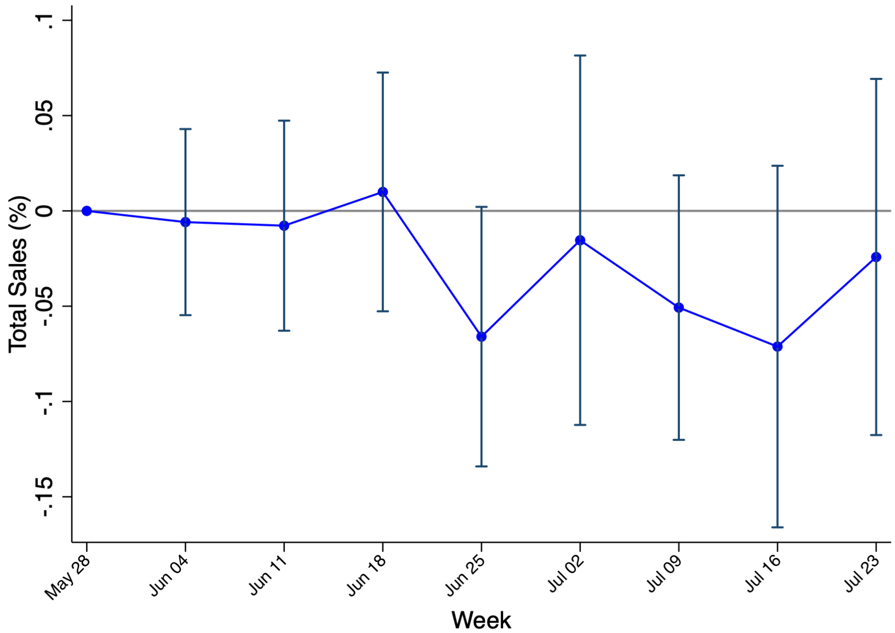

Figure 2.6: Dynamics of Spending Response. This figure reports estimated coefficients and 95% confidence intervals from regression 2.2. The data are at the store-week level and cover the weeks of May 28, 2023 to July 23, 2023. The specification includes fixed effects for the following groups, interacted with dummies for each week: designated market area, parent company, retail channel, median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

our main tests, suggests that borrowers had already partially incorporated news about the legality of the debt relief plan and updated their beliefs about forthcoming changes in their permanent income.[17]

2.4.6 Debt Relief Applications

Finally, we present further evidence that our results align with a beliefs channel and document demographic differences in belief formation. Using zip code-level counts of debt relief applications that borrowers submitted via `https://studentaid.gov/debt-relief/application`, we first estimate the following equation:

$$Application\ Share_k = \beta Eligible\ Share_k + \delta_k' X_k + \epsilon_k, \tag{2.4}$$

where $Application\ Share_k$ represents the number of debt relief applications in a given zip code as a share of the adult

[17]Contemporaneous prediction-market pricing is consistent with this interpretation. On June 24, 2023, the Kalshi contract on whether the Supreme Court would permit student loan forgiveness implied a 37% probability of forgiveness. We view this as suggestive evidence of public skepticism about the plan's survival, rather than a direct measure of borrower beliefs.

	(1)	(2)	(3)
Eligible Share	0.60***	0.59***	0.63***
	(0.0045)	(0.0058)	(0.0041)
County FE	NO	YES	YES
Demographics	NO	NO	YES
Adj. R-squared	0.86	0.92	0.96
Observations	18802	18094	15817

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

Table 2.4: Student Loan Borrowers Anticipated Debt Relief. This table reports the results of regression 2.4. The outcome variable is debt relief applications as a share of the adults in a zip code. Eligible Share is the share of adults in a zip code who were eligible for loan forgiveness. We control for gender, race, education, income categories, population groups, and partisanship. Standard errors are clustered at the county level.

population. The intuition behind this test is straightforward: if eligible borrowers were either unaware of the loan forgiveness plan or assigned a low subjective probability to receiving debt relief, we would expect $\beta = 0$. For instance, student loan borrowers would not have incurred the minimal temporal cost required to complete the application if they accurately anticipated the Supreme Court decision. In contrast, we find that application counts strongly correlate with student loan eligibility (Table 2.4). Column (3) shows that a 1 percentage point increase in the eligible share within a zip code is associated with a 0.63 percentage point increase in the application share.

We interpret these results primarily through a beliefs channel rather than an attention mechanism. This is because news about the debt relief plan was widely disseminated, and the Department of Education notified individual borrowers via email. One important caveat is that we cannot account for procrastination. Some borrowers, despite being well-informed and optimistic about the likelihood of loan forgiveness, may have delayed submitting their applications. Since the Department of Education was forced to suspend processing debt relief applications one month into the launch of the website, we do not observe the universe of eventual debt relief applicants.

Keeping this caveat in mind, we compute the application take-up rate within each zip code, as our proxy for beliefs, and document demographic differences via cross-sectional regressions using Census Bureau data. The third column of Table 2.5 shows that zip codes with a higher proportion of women and high-income households especially have higher debt relief take-up rates. Notably, we also observe a positive relationship between the share of donors to the Democratic Party and take-up rates.

For comparison purposes, we report demographic associations with the eligible share in column (1) and the application share in column (2). We observe that eligibility is increasing in income, which suggests that the debt relief policy was regressive in nature.[18] Interestingly, zip codes with a higher proportion of Asian borrowers and those aged 18-34 are associated with lower eligible and application shares, yet higher take-up rates.

2.5 Conclusion

In this paper, we provide empirical evidence on borrowers' consumption response to President Biden's August 2022 announcement, informing the contentious political debate on student loan forgiveness. Using comprehensive data on borrowers eligible for debt relief linked to high-frequency spending data, we find that borrowers reacted positively to the news about an upcoming reduction in their financial liabilities. Moreover, we document a negative reaction to the Supreme Court decision striking down the debt relief plan. Finally, we show that the majority of eligible borrowers submitted debt relief applications, suggesting that they anticipated receiving loan forgiveness.

Although there are a number of advantages to our empirical approach, it is worth repeating a key limitation. As no borrowers received forgiveness under the plan, our estimates do not capture the long-term effects of debt cancellation; rather, they measure initial reactions to the broad-based debt relief announcement and its subsequent reversal. However, these reactions have implications for policymakers, as they highlight the role of public communication in shaping expectations and influencing economic behavior.

While our results indicate that student loan forgiveness can have stimulatory effects, assessing the fiscal cost-effectiveness of such policies lies beyond the scope of this paper. An additional, important avenue for further research is to evaluate the full welfare effects of student loan forgiveness, taking distributional considerations seriously as in Catherine and Yannelis (2023).

[18] However, many borrowers living in low-income zip codes may have been eligible for the full $20,000 in loan forgiveness.

	Eligible Share	App Share	Take-up Rate
Biden's 2020 donation share	0.019***	0.018***	0.060***
	(0.0035)	(0.0023)	(0.0045)
Female share	0.15***	0.12***	0.25***
	(0.013)	(0.0085)	(0.019)
Bachelor degree or higher share	0.081***	0.070***	0.15***
	(0.0068)	(0.0050)	(0.011)
Population [18-34] years old share	-0.055***	-0.015*	0.061***
	(0.012)	(0.0087)	(0.020)
Population [35-50] years old share	0.026	0.038***	0.052*
	(0.021)	(0.014)	(0.028)
Population 50+ years old share	-0.14***	-0.071***	0.0077
	(0.011)	(0.0076)	(0.016)
Black alone share	0.15***	0.086***	0.023***
	(0.0046)	(0.0027)	(0.0050)
Asian alone share	-0.090***	-0.047***	0.058***
	(0.0082)	(0.0056)	(0.017)
Income <25k share	0.16***	0.090***	-0.014
	(0.013)	(0.0094)	(0.021)
Income [25-50k) share	0.17***	0.11***	0.073***
	(0.013)	(0.0098)	(0.022)
Income [50-100k) share	0.22***	0.16***	0.20***
	(0.013)	(0.0099)	(0.021)
Income [100-200k) share	0.27***	0.21***	0.38***
	(0.011)	(0.0086)	(0.020)
Mean of Outcome	0.15	0.088	0.57
Adj. R-squared	0.48	0.47	0.51
Observations	16660	16660	16660

Standard errors in parentheses

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

Table 2.5: Demographic Correlates of Debt Relief. This table reports cross-sectional regressions of debt relief eligibility, applications, and take-up on zip code-level demographic information. The omitted category for age is the fraction of the population that is < 18. The omitted category for income is the fraction of the population with income $>$ \$200k. Eligible Share is the share of adults in a zip code who were eligible for loan forgiveness. App Share is the share of adults in a zip code who applied for student debt relief. Take-up Rate is the number of applications as a share of the number of eligible borrowers in a zip code. Standard errors are clustered at the county level.

Appendix 2.A

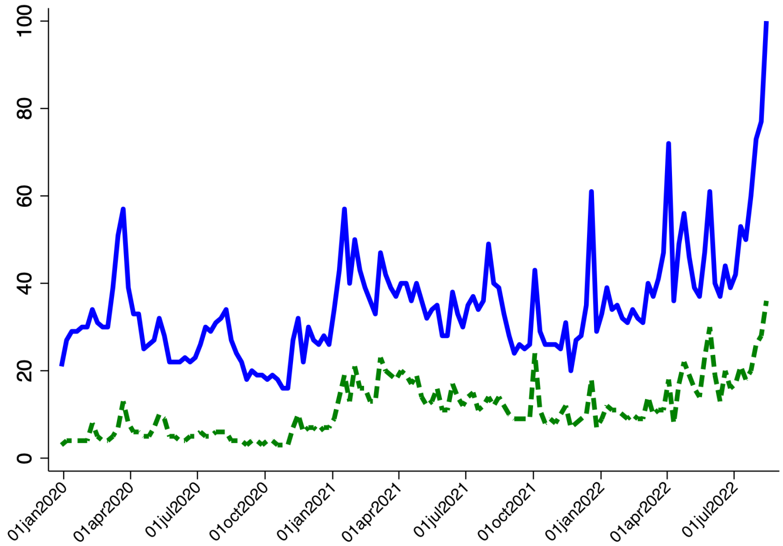

(a)

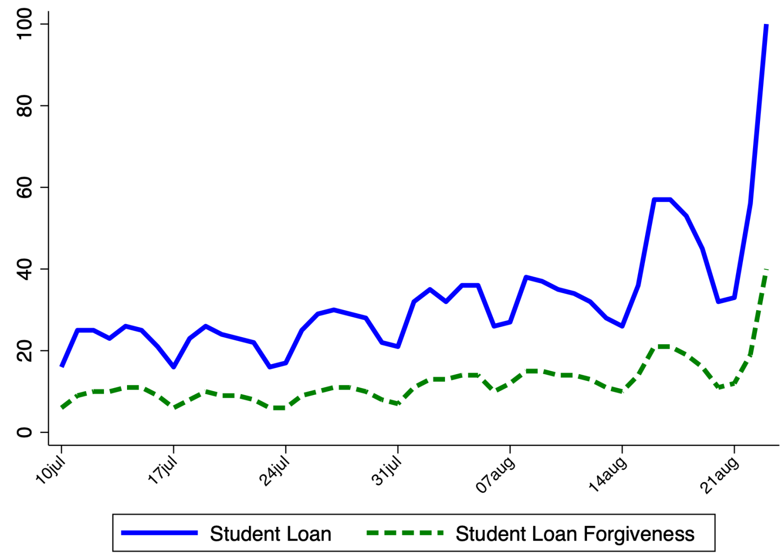

(b)

Figure A.1: Google Trends (before). This figure plots weekly aggregate Google searches for 'student loan' and 'student loan forgiveness'. Panel (a) covers the weeks of December 29, 2019 to August 14, 2022. Panel (b) covers the days of July 10, 2019 to August 23, 2022. The counts are normalized and presented on a scale from 0-100, where 100 is the maximum search interest.

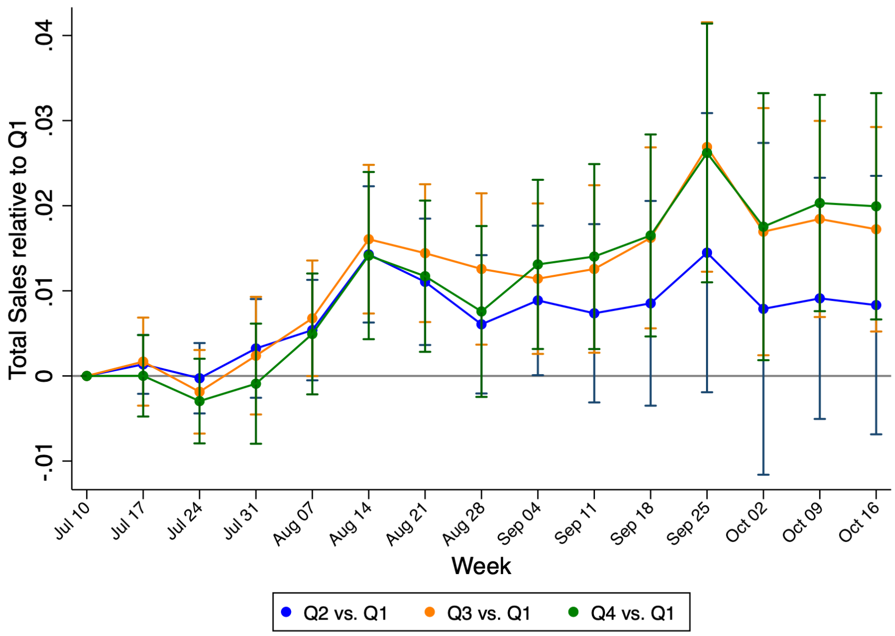

Figure A.2: Dynamics of Spending Response by County Exposure Quartile. This figure reports event-study coefficients and 95% confidence intervals from a quartile-based difference-in-differences specification estimating the effect of President Biden's August 24, 2022 student loan forgiveness announcement on retail sales. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is ln(sales). Counties are assigned to quartiles based on their pre-announcement eligible share, defined as the share of adults in the county eligible for debt relief. Quartile 1 (Q1) is omitted, so the reported coefficients trace week-by-week differences in sales for stores located in counties in quartiles 2–4 relative to stores in the lowest-exposure quartile. The specification includes fixed effects for the following groups, each interacted with week dummies: designated market area, parent company, retail channel, median household income, total population, partisanship, and the under-18 and college-enrollment shares of the population. Standard errors are clustered at the county level.

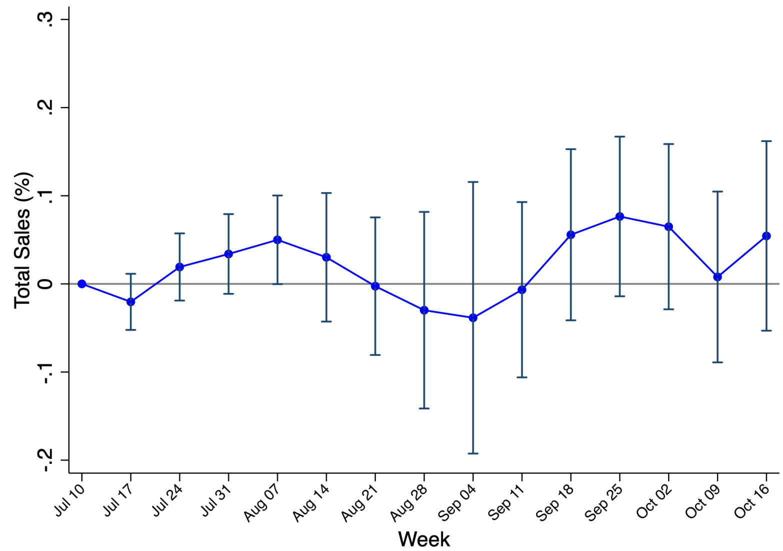

(a) 2021

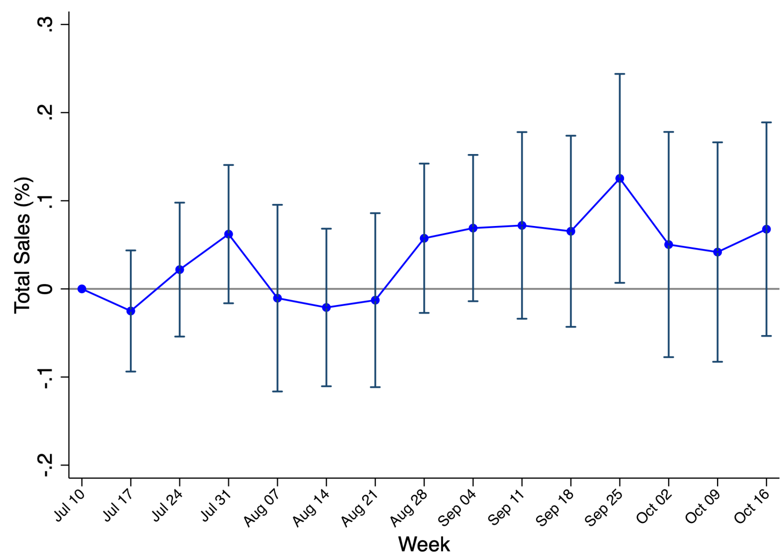

(b) 2023

Figure A.3: Placebo Tests. This figure reports estimated coefficients and 95% confidence intervals from regression 2.2. The data are at the store-week level and cover weeks in 2021 (Panel a) and 2023 (Panel b). The specification includes fixed effects for the following groups, interacted with dummies for each week: designated market area, parent company, retail channel, median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

	Mean	SD	p10	p25	p50	p75	p90	N
Drug	52.71	29.90	22.04	32.79	48.07	65.86	87.49	268,288
Food	603.13	362.51	237.06	353.66	528.25	761.59	1,069.90	166,431
Convenience	41.27	20.21	19.94	27.58	37.68	51.02	67.14	176,768
Liquor	84.05	62.09	22.82	41.37	72.70	103.83	153.40	4,274
Mass Merchandiser	136.41	245.11	15.43	20.98	30.71	91.70	534.10	183,288
Total	184.19	298.36	20.05	30.00	50.23	137.26	613.19	799,049

Table A.1: Summary statistics. This table reports summary statistics on weekly spending by store type. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. Sales are reported in thousands of dollars.

	(1)	(2)	(3)	(4)
SL Debt Share x Post	0.24***	0.31***	0.36***	0.24***
	(0.035)	(0.062)	(0.062)	(0.063)
Store FE	YES	YES	YES	YES
Week FE	YES	NO	NO	NO
Market-Week FE	NO	YES	YES	YES
Parent-Week FE	NO	NO	YES	YES
Channel-Week FE	NO	NO	YES	YES
Demographics-Week FE	NO	NO	NO	YES
Adj. R-squared	0.99	0.99	0.99	0.99
Observations	798906	798696	798681	797893

Standard errors in parentheses

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

Table A.2: Spending at Retail Stores Post-Announcement. This table reports the results of regression 2.1. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. The outcome variable is $ln(sales)$. SL Debt Share is the share of people in a county with a credit bureau record who have any student loan debt, and Post is an indicator for whether a week falls around or after August 24, 2022. The demographic controls include: median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

	(1)	(2)	(3)	(4)
Eligible Share x Post	0.23***	0.22***	0.23***	0.092**
	(0.034)	(0.044)	(0.044)	(0.036)
Store FE	YES	YES	YES	YES
Week FE	YES	NO	NO	NO
Market-Week FE	NO	YES	YES	YES
Parent-Week FE	NO	NO	YES	YES
Channel-Week FE	NO	NO	YES	YES
Demographics-Week FE	NO	NO	NO	YES
Adj. R-squared	0.99	0.99	0.99	0.99
Observations	785550	785340	785325	785325

Standard errors in parentheses

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

Table A.3: Spending at Retail Stores Post-Announcement. This table reports the results of regression 2.1. The data are at the store-week level and cover the weeks of July 10, 2022 to October 16, 2022. We omit stores that do not report sales information in each week of the year. The outcome variable is $ln(sales)$. Eligible Share is the share of adults in a county who were eligible for loan forgiveness, and Post is an indicator for whether a week falls around or after August 24, 2022. The demographic controls include: median household income, total population, partisanship, under-18 and college enrollment shares of the population. Standard errors are clustered at the county level.

	Low Exposure			High Exposure			
	Mean	SD	Count	Mean	SD	Count	Diff
Female share	0.49	0.03	1616	0.50	0.02	1586	-0.01
Non-adult share	0.22	0.04	1616	0.23	0.03	1586	-0.01
College enrollment share	0.03	0.09	1616	0.09	0.13	1586	-0.07
Democrat	0.10	0.30	1563	0.22	0.41	1533	-0.12
Median household income	52.85	14.08	1615	61.91	17.29	1586	-9.06
Total Population (log)	9.36	1.08	1616	11.21	1.27	1586	-1.85

Table A.4: Covariate balance. This table reports county-level summary statistics. We define counties as high (low) exposure if the share of eligible student loan borrowers is above (below) the median share within their state. Income is reported in thousands of dollars. We classify Democratic counties as those where the Democratic candidate received the majority of votes in the 2020 presidential election.

Bibliography

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2): 3–30.

Ager, Philipp, Marc Goñi, and Kjell G. Salvanes. 2026. "Gender-Biased Technological Change: Milking Machines and the Exodus of Women from Farming." *American Economic Review*, 116(1): 246–286.

Agrawal, Ashwini K., and David A. Matsa. 2013. "Labor unemployment risk and corporate financing decisions." *Journal of Financial Economics*, 108(2): 449–470.

Autor, David. 2022. "The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty."

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3–30.

Badue, Claudine, Rânik Guidolini, Raphael Vivacqua Carneiro, Pedro Azevedo, Vinicius B. Cardoso, Avelino Forechi, Luan Jesus, Rodrigo Berriel, Thiago M. Paixão, Filipe Mutz, Lucas de Paula Veronese, Thiago Oliveira-Santos, and Alberto F. De Souza. 2021. "Self-driving cars: A survey." *Expert Systems with Applications*, 165: 113816.

Bailey, Michael, Drew Johnston, Martin Koenen, Theresa Kuchler, Dominic Russel, and Johannes Stroebel. 2024. "Social Networks Shape Beliefs and Behaviors: Evidence from Social Distancing during the COVID-19 Pandemic." *Journal of Political Economy Microeconomics*. Publisher: The University of Chicago Press.

Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2018. "Social Connectedness: Measurement, Determinants, and Effects." *Journal of Economic Perspectives*, 32(3): 259–280.

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics*, 144(2): 370–395.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *The Quarterly Journal of Economics*, 131(4): 1593–1636.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. "Asymmetric Consumption Smoothing." *American Economic Review*, 111(1): 192–230.

Berk, Jonathan B., Richard Stanton, and Josef Zechner. 2010. "Human Capital, Bankruptcy, and Capital Structure." *The Journal of Finance*, 65(3): 891–926.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2018. "Behavioral Household Finance." In *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 1 of *Handbook of Behavioral Economics - Foundations and Applications 1*, , ed. B. Douglas Bernheim, Stefano DellaVigna and David Laibson, 177–276. North-Holland.

Bond, Philip, and Lukas Kremens. 2023. "Income and Inequality Under Asymptotically Full Automation."

Braxton, J. Carter, and Bledi Taska. 2023. "Technological Change and the Consequences of Job Loss." *American Economic Review*, 113(2): 279–316.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence."

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2024. "Difference-in-differences with a continuous treatment." National Bureau of Economic Research.

Cannon, Brad, David Hirshleifer, and Joshua Thornton. 2024. "Friends with Benefits: Social Capital and Household Financial Behavior."

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life-Cycle/Permanent-Income Hypothesis." *Quarterly Journal of Economics*.

Case, Anne, and Angus Deaton. 2015. "Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century." *Proceedings of the National Academy of Sciences*, 112(49): 15078–15083.

Castex, Gonzalo, Emma Chow, and Evgenia Dechter. 2024. "The Risk of Automation and Occupational Choice."

Catherine, Sylvain, and Constantine Yannelis. 2023. "The distributional effects of student loan forgiveness." *Journal of Financial Economics*, 147(2): 297–316.

Cavounidis, Costas, Qingyuan Chai, Kevin Lang, and Raghav Malhotra. 2023. "Obsolescence Rents: Teamsters, Truckers, and Impending Innovations."

Cespedes, Jacelly, Carlos Parra, and Clemens Sialm. 2026. "The Effect of Principal Reduction on Household Distress: Evidence from Mortgage Cramdown." *The Review of Financial Studies*, 39(2): 518–561.

Chopra, Felix, Christopher Roth, and Johannes Wohlfart. 2025. "Home Price Expectations and Spending: Evidence from a Field Experiment." *American Economic Review*, 115(7): 2267–2305.

Cockriel, William. 2025. "Machines Eating Men: Shoemakers and Their Children After the McKay Stitcher."

Dávalos, María E., Hai Fang, and Michael T. French. 2012. "Easing the Pain of an Economic Downturn: Macroeconomic Conditions and Excessive Alcohol Consumption." *Health Economics*, 21(11): 1318–1335.

Day, Jennifer, and Andrew Hait. 2019. "America Keeps on Truckin." *U.S. Census Bureau: America Counts*, Accessed: 2024-12-28.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*.

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

Diemer, Andreas, and Tanner Regan. 2022. "No inventor is an island: Social connectedness and the geography of knowledge flows in the US." *Research Policy*, 51(2): 104416.

Di Maggio, Marco, Ankit Kalda, and Vincent Yao. 2026. "Second Chance: Life with Less Student Debt." *The Journal of Finance*, 81(1): 507–550.

Dinerstein, Michael, Constantine Yannelis, and Ching-Tse Chen. 2024. "Debt Moratoria: Evidence from Student Loan Forbearance." *American Economic Review: Insights*, 6(2): 196–213.

Dinerstein, Michael, Samuel Earnest, Dmitri K. Koustas, and Constantine Yannelis. 2025. "Student Loan Forgiveness."

Dobbie, Will, and Jae Song. 2020. "Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers." *American Economic Review*, 110(4): 984–1018.

Dore, Timothy E, and Rebecca Zarutskie. 2023. "When Does Higher Firm Leverage Lead to Higher Employee Pay?" *The Review of Corporate Finance Studies*, 12(1): 36–77.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell. 2022. "The Use of Scanner Data for Economics Research." *Annual Review of Economics*, 14(Volume 14, 2022): 723–745.

Eaglin, Christopher, Apoorv Gupta, Filippo Mezzanotti, and Jonathan Zinman. 2025. "Working it out: Randomized Modification and Entrepreneurial Effort in a Collateralized Debt Market."

Engelmann, Jan B., Maël Lebreton, Nahuel A. Salem-Garcia, Peter Schwardmann, and Joël J. van der Weele. 2024. "Anticipatory Anxiety and Wishful Thinking." *American Economic Review*, 114(4): 926–960.

Fallick, Bruce, John Haltiwanger, Erika McEntarfer, and Matthew Staiger. 2025. "Job Displacement and Earnings Losses: The Role of Joblessness." *American Economic Journal: Macroeconomics*, 17(2): 177–205.

Feigenbaum, James, and Daniel P Gross. 2024. "Answering the Call of Automation: How the Labor Market Adjusted to Mechanizing Telephone Operation." *The Quarterly Journal of Economics*.

Ganong, Peter, and Pascal Noel. 2020. "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." *American Economic Review*, 110(10): 3100–3138.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas. 2012. "Displacement risk and asset returns." *Journal of Financial Economics*, 105(3): 491–510.

Garmaise, Mark J., Yaron Levi, and Hanno Lustig. 2024. "Spending Less after (Seemingly) Bad News." *The Journal of Finance*, 79(4): 2429–2471.

Gee, Laura K., Jason Jones, and Moira Burke. 2017. "Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network." *Journal of Labor Economics*, 35(2): 485–518. Publisher: The University of Chicago Press.

Gittleman, Maury, and Kristen Monaco. 2020. "Truck-Driving Jobs: Are They Headed for Rapid Elimination?" *ILR Review*, 73(1): 3–24. Publisher: SAGE Publications Inc.

Golin, M., and C. Rauh. 2022. "The Impact of Fear of Automation."

Gomes, Francisco, Thomas Jansson, and Yigitcan Karabulut. 2024. "Do Robots Increase Wealth Dispersion?" *The Review of Financial Studies*, 37(1): 119–160.

Gopalan, Radhakrishnan, Barton H. Hamilton, Jorge Sabat, and David Sovich. 2024. "Aversion to Student Debt? Evidence from Low-Wage Workers." *The Journal of Finance*, 79(2): 1249–1295.

Graham, John R., Hyunseob Kim, Si Li, and Jiaping Qiu. 2023. "Employee Costs of Corporate Bankruptcy." *The Journal of Finance*, 78(4): 2087–2137.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market."

Hémous, David, Morten Olsen, Carlo Zanella, and Antoine Dechezleprêtre. 2025. "Induced Automation Innovation: Evidence from Firm-Level Patent Data." *Journal of Political Economy*, 133(6): 1975–2028. Publisher: The University of Chicago Press.

Hirshleifer, David, Lin Peng, and Qiguang Wang. 2024. "News Diffusion in Social Networks and Stock Market Reactions." *The Review of Financial Studies.*

Hu, Zhongchen. 2022. "Social interactions and households' flood insurance decisions." *Journal of Financial Economics*, 144(2): 414–432.

Indarte, Sasha. 2023. "Moral Hazard versus Liquidity in Household Bankruptcy." *The Journal of Finance*, 78(5): 2421–2464.

Jiang, Wei, Junyoung Park, Rachel (Jiqiu) Xiao, and Shen Zhang. 2025. "AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents."

Jones, Charles I. 2024. "The AI Dilemma: Growth versus Existential Risk." *American Economic Review: Insights*, 6(4): 575–590.

Kalra, Nidhi, and David G Groves. 2017. *The enemy of good: Estimating the cost of waiting for nearly perfect automated vehicles.* Rand Corporation.

Kaplan, Greg, and Giovanni L. Violante. 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models." *Annual Review of Economics*, 14(Volume 14, 2022): 747–775.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2023. "Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data."

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Jae Song. 2020. "Technological Innovation and Labor Income Risk."

Korinek, Anton, and Donghyun Suh. 2024. "Scenarios for the Transition to AGI."

Kushner, Matt G, Kenneth Abrams, and Carrie Borchardt. 2000. "The relationship between anxiety disorders and alcohol use disorders: A review of major perspectives and findings." *Clinical Psychology Review*, 20(2): 149–171.

Levy, Karen. 2023. *Data Driven: Truckers, Technology, and the New Workplace Surveillance.* Princeton:Princeton University Press.

Looney, Adam, and Constantine Yannelis. 2024. "What Went Wrong with Federal Student Loans?" *Journal of Economic Perspectives*, 38(3): 209–236.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. "Wage risk and employment risk over the life cycle." *American Economic Review*, 100(4): 1432–1467.

Manning, Sam, and Tomás Aguirre. 2025. "How Adaptable Are American Workers to AI-Induced Job Displacement?" National Bureau of Economic Research.

Matsa, David A. 2018. "Capital Structure and a Firm's Workforce." *Annual Review of Financial Economics*, 10(Volume 10, 2018): 387–412. Publisher: Annual Reviews.

Maturana, Gonzalo, and Jordan Nickerson. 2019. "Teachers Teaching Teachers: The Role of Workplace Peer Effects in Financial Decisions." *The Review of Financial Studies*, 32(10): 3920–3957.

Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Duncan Simester. 2022. "Belief Disagreement and Portfolio Choice." *The Journal of Finance*, 77(6): 3191–3247.

Mian, Atif, Amir Sufi, and Nasim Khoshkhou. 2023. "Partisan Bias, Economic Expectations, and Household Spending." *The Review of Economics and Statistics*, 105(3): 493–510.

Mohan, Aniruddh, and Parth Vaishnav. 2022. "Impact of automation on long haul trucking operator-hours in the United States." *Humanities and Social Sciences Communications*, 9(1): 1–10. Publisher: Palgrave.

Morissette, Sandra Baker, Matthew T. Tull, Suzy Bird Gulliver, Barbara Wolfsdorf Kamholz, and Rose T. Zimering. 2007. "Anxiety, anxiety disorders, tobacco use, and nicotine: A critical review of interrelationships." *Psychological Bulletin*, 133(2): 245–272. Place: US Publisher: American Psychological Association.

National Academies of Sciences, Engineering, and Medicine. 2017. *Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?* Washington, D.C.:National Academies Press.

Orii, Lisa, Diana Tosca, Andrew L Kun, and Orit Shaer. 2021. "Perceptions on the Future of Automation in r/Truckers." *CHI EA '21*, 1–6. New York, NY, USA:Association for Computing Machinery.

Phares, Jonathan, and Andrew Balthrop. 2022. "Investigating the role of competing wage opportunities in truck driver occupational choice." *Journal of Business Logistics*, 43(2): 265–289.

Restrepo, Pascual. 2025. "We Won't Be Missed: Work and Growth in the AGI World." In *The Economics of Transformative AI.* University of Chicago Press.

Shoag, Daniel, Michael R. Strain, and Stan Veuger. 2021. "How Do Workers Perceive the Risks from Automation and the Opportunities to Retrain? Evidence from a Survey of Truck Drivers."

Sullivan, Daniel, and Till von Wachter. 2009. "Job Displacement and Mortality: An Analysis Using Administrative Data*." *The Quarterly Journal of Economics*, 124(3): 1265–1306.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199.

Titman, Sheridan. 1984. "The effect of capital structure on a firm's liquidation decision." *Journal of Financial Economics*, 13(1): 137–151.

Urban Institute. 2024. "Debt in America: An Interactive Map." Accessed: 2024-12-24.

Van Fossen, Jenna A, Amy M Schuster, Danielle Sperry, Shelia R Cotten, and Chu-Hsiang Chang. 2023. "Concerns, Career Decisions, and Role Changes: A Qualitative Study of Perceptions of Autonomous Vehicles in the Trucking Industry." *Work, Aging and Retirement*, 9(4): 399–406.

Viceira, Luis M. 2001. "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income." *The Journal of Finance*, 56(2): 433–470.

Viscelli, Steve. 2016. *The Big Rig: Trucking and the Decline of the American Dream.* Univ of California Press.

Viscelli, Steve. 2018. "Driverless? Autonomous Trucks and the Future of the American Trucker." Center for Labor Research and Education, University of California, Berkeley, and Working Partnerships USA.

White House. 2022. "Fact Sheet: President Biden Announces Student Loan Relief for Borrowers Who Need It Most." Accessed: 2024-12-24.

Wilson, Riley. 2022. "The Impact of Social Networks on EITC Claiming Behavior." *The Review of Economics and Statistics*, 104(5): 929–945.

Xu, Chenzi. 2022. "Reshaping Global Trade: The Immediate and Long-Run Effects of Bank Failures." *The Quarterly Journal of Economics*, 137(4): 2107–2161.

Yannelis, Constantine, and Greg Tracey. 2022. "Student Loans and Borrower Outcomes." *Annual Review of Financial Economics*, 14(Volume 14, 2022): 167–186.

Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *The Quarterly Journal of Economics*, 113(4): 1091–1117.

Zingales, Luigi. 1998. "Survival of the Fittest or the Fattest? Exit and Financing in the Trucking Industry." *The Journal of Finance*, 53(3): 905–938.