Empirical Essays on Firm Operations and Labor Economics

by

Samuel A. Stern

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2024

Doctoral Committee:

Professor John Leahy, Co-Chair
Associate Professor Pablo Ottonello, Co-Chair, University of Maryland
Associate Professor Peter Bahr
Professor Linda Tesar

Samuel A. Stern
sternsa@umich.edu
ORCID iD: 0009-0008-0403-4917

To my parents, Sally and Jeffrey,
and to my partner, Naomi

ACKNOWLEDGEMENTS

I am incredibly grateful to the encouraging faculty, caring friends, and loving family who have supported me over the past eight years. This dissertation was only possible with their help at each step of my graduate school journey.

I am extremely indebted to Pablo Ottonello for his thoughtful guidance and unwavering encouragement from the beginning of my time at Michigan. He included me in his own research, provided access to critical data, and spent countless meetings discussing the research contained in this dissertation. John Leahy offered invaluable insights into my research and helped navigate the publication process for the second chapter of this dissertation. John's macro workshop for graduate students shaped my understanding of macroeconomics literature and helped hone my presentation skills. Peter Bahr and Linda Tesar have been essential in my last chapter of graduate school. Their encouragement and new perspective on my research were essential to the final version of this dissertation.

Many others in the Department of Economics at University of Michigan were central to my development as a researcher. Wenting Song and Erin Markiewitz were amazing collaborators who know many of these papers as well as I do. Lutz Kilian and Andrés Blanco taught me econometric and numerical methods that were central to my graduate training. John Bound and Charlie Brown provided an insight narrative on the evolution of labor economics in their courses. Before attending graduate school, my understanding of economics was shaped by Gara Afonso and Nicola Cetorelli at the Federal Reserve Bank of New York, Christopher Snyder at Dartmouth College, and Jason Majewski at Bethlehem Central High School. Lastly, I am grateful to Jack Carter and other staff at the U.S. Census Bureau and Michigan's Federal Statistical Research Data Center who reviewed the statistics in the first chapter of this dissertation and ensured that their release satisfied the Census Bureau's criteria for disclosure avoidance.[1]

My life in Ann Arbor was filled with close friends and classmates who enriched my time as a graduate student. I will remember many barbecues and game nights spent with Paul

[1] The research for Chapter 1 was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1605. Resulting estimates were disclosed under disclosure clearance number, CBDRB-FY24-P1605-R10838. Any views expressed are my own and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product.

Kindsgrab, Shawn Martin, and Marco Rojas. Sarah Cissell and Gabe Monett offered friendship and conversation beyond economics. These years were colored the most by the love and support of my partner, Naomi Rawitz, who built a life with me in Ann Arbor.

Finally, I would like to thank my family for supporting me in more ways than I could ever detail here. My parents, Sally and Jeffrey, instilled the values that inspired my academic pursuits. My older sister, Rebecca, has remained a steadfast role model for as long as I can remember and my younger brother, Joshua, is the best companion that an older brother could ask for.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

ABSTRACT

This dissertation includes three papers that study firm operations and labor market adjustment using an empirical focus. The first chapter investigates the effect of workplace safety regulation on manufacturing operations. It matches a sample of random worksite inspections by OSHA to annual manufacturing data from the US Census Bureau and then estimates how these inspections affected establishment operations using a local projections approach to difference-in-differences. The results show that increased regulatory enforcement and safety compliance primarily affected establishment size rather than productivity or capital-intensity. Employment growth among inspected establishments was 15.6 percentage points lower over a decade, while productivity growth fell by 1.9 percentage points and the capital-labor ratio grew by 4.2 percentage points more before reverting to non-inspection levels. Over a five-year horizon, the direct cost of OSHA's random inspection program is estimated to be between \$5.5 and \$39.5 billion of manufacturing output in 2018 dollars. The same event study methodology is then applied to workplace incidents that triggered OSHA inspections. Workplace accidents or reports of unsafe conditions were followed by even larger declines in employment growth and more persistent changes to productivity and factor allocations. Complaints about unsafe working conditions preceded particularly large cuts to average worker compensation.

The second chapter, co-authored with Wenting Song, provides empirical evidence of the importance of firm attention to macroeconomic dynamics. We construct a text-based measure of attention to macroeconomic news and document that attention is polarized across firms and countercyclical. Differences in attention lead to asymmetric responses to monetary policy: expansionary monetary shocks raise market values of attentive firms more than those of inattentive firms, and contractionary shocks lower values of attentive firms by less. Attention also mitigates the effects of macroeconomic uncertainty on firm performance. In a quantitative rational inattention model that is calibrated with this new text-based measure, inattention drives monetary non-neutrality. As average attention varies over the business cycle, so does the efficacy of monetary policy.

The third chapter estimates the effect of access to higher education on local employment recoveries following the Great Recession. It begins by describing the historical magnitude

and composition of countercyclical enrollment between 1987 and 2019 to motivate the potential importance of higher education as a means of labor market adjustment. The chapter then presents an empirical strategy for estimating the effect of new enrollment during the Great Recession on local employment rates in subsequent years. The strategy uses a scaled Bartik instrument to isolate enrollment that is plausibly exogenous to local labor demand shocks and conditions on local educational attainment. The main results show that a one percentage point increase in the enrollment rate between 2007 and 2010 is associated with a 2.0 percentage point larger increase in the employment rate between 2007 and 2018. The effect is strongest among adults ages 25-34 and is largely consistent with the timing and age composition of new enrollment during the Great Recession.

CHAPTER 1

Labor Regulation and Manufacturing Operations: Evidence from OSHA Inspections

1.1. Introduction

The economic costs of workplace safety regulation are typically discussed in terms of productivity and competitiveness. From a neoclassical perspective, regulatory compliance raises the cost of labor, shifts the capital-labor ratio away from its optimum, and reduces a firm's productivity relative to competitors that are not under the same regulatory scrutiny. Existing research has used industry-wide measures of productivity and regulatory enforcement to estimate the consequences of safety regulation, but this approach does not account for how changes to productivity may affect safety nor can it tell us about individual firm responses. This paper uses random workplace inspections by the Occupational and Safety Health Administration (OSHA) to estimate the effect of safety regulation on manufacturing operations.

Since 1972, OSHA has conducted between 50,000 and 140,000 inspections each year to ensure safe working conditions. Between 1979 and 1997, 1.1 million of these inspections were randomly assigned within hazardous industries and initiated without warning. Most firms were found in violation of at least one safety standard but received only minor fines for initial violations. Despite facing nominal financial consequences, nearly all violators addressed their safety hazards to avoid steep penalties for repeat violations. In effect, these inspections were idiosyncratic shocks to regulatory compliance and workplace safety that were unrelated to firm operations or policy changes.

Since inspections were unanticipated and randomly assigned within-industry, their effect on firm operations can be estimated with a difference-in-differences (DiD) empirical design. Using the local projections approach to DiD for staggered treatment timing, this paper compares inspected firm sites (or establishments) to a control group that is inspected in later years. Each inspection is matched to administrative records on manufacturing operations from the US Census Bureau, and the treatment effect for a variety of outcomes is reported over the decade following inspection.

Resulting estimates show that regulatory enforcement and safety compliance primarily affected the scale of manufacturing operations. On average, employment growth among inspected establishments was 15.6 percentage points lower over a ten year horizon. These sites largely stagnated as competitors continued to grow unencumbered, resulting in a 5.2 percentage point gap in employment growth only two years after inspection. The severity of the employment effect is strongest among smaller, more labor-intensive, and more hazardous worksites.

Random inspections had a comparatively minor effect on establishment productivity. Average revenue-based total factor productivity (TFPR) growth among inspected plants was 1.9 percentage points lower after four years but eventually converged to the average TFPR growth rate among uninspected plants. This estimate is larger than industry-wide productivity effects found in Gray (1987) but comparable to establishment-level productivity losses due to environmental regulation from Greenstone et al. (2012). The temporary slowdown in observed productivity growth alongside a large and persistent decrease in employment growth is consistent with a total factor productivity (TFP) shock under monopolistic competition as in Hsieh and Klenow (2009).

Factor reallocation was also not a persistent consequence of random inspection. The average capital-labor ratio among inspected establishments grew by 4.2 percentage points more over five years but quickly reverted to pre-inspection levels. This temporary increase is consistent with slower capital adjustment, which is documented for both structures and equipment capital. Once establishments were able to adjust capital, it appears that the cost of safety compliance was largely factor neutral.

The same diff-in-diff methodology is then applied to OSHA inspections that were triggered by a report of unsafe working conditions or a workplace accident. Compared to random inspections, these safety incidents were followed by larger declines in employment growth and more persistent changes to establishment productivity and capital-intensity. Furthermore, establishments that were reported for unsafe work conditions exhibited sharper declines in average worker compensation.

The next section examines the contemporaneous relationship between workplace safety and other establishment characteristics. It finds that larger and less productive establishments received more safety violations and larger penalties from OSHA on average. However, inspection penalties per employee decreased with establishment size, which suggests that larger establishments are actually safer despite higher total penalties. These findings are important for understanding the types of establishments that are safe for workers and how increased regulatory enforcement may shift the distribution of establishment characteristics.

The final section approximates the direct cost of OSHA's enforcement program using estimated responses to random inspections. It considers two scenarios that differ in whether reduced output growth among inspected establishments is replaced by unaffected competitors. The direct cost of one year of random inspections over a five-year horizon is estimated to be \$5.5 billion of manufacturing output in 2018 dollars assuming elastic supply and \$39.5 billion under inelastic supply. This approach to approximating the cost of random inspections has some important limitations that are outlined at the end of the section.

Related Literature Workplace safety regulation gained attention in the 1980s as researchers sought to explain slower US productivity growth during the 1970s. Gray (1987) uses a panel of 450 manufacturing industries between 1958 and 1978 to estimate the productivity effects of both safety and environmental regulations. Increased OSHA enforcement – measured as the share of industry employees subject to inspection – is found to slow industry TFP growth by 0.27 percentage points per year. Bartel and Thomas (1987) focuses on heterogeneous effects of OSHA regulation on manufacturing in the late 1970s and finds that less concentrated industries and those with greater import competition exhibited a stronger negative relationship between profit margins and OSHA penalties. Dufour et al. (1998) finds mixed effects of safety regulation on productivity among Quebec manufacturers in the late 1980s: a policy that allowed workers to avoid hazardous tasks lowered productivity while other preventative policies and regulatory fines raised productivity. More recently, Levine et al. (2012) examines a sample of 409 random OSHA inspections in California and finds no discernible effect on employment, sales, or firm survival despite significant declines in injury rates.

A larger empirical literature on industrial regulation and manufacturing operations has focused on environmental policy, though many of the findings are relevant to regulatory compliance more broadly. Greenstone (2002) and Becker and Henderson (2000) use county-level variation in compliance with the 1970 Clean Air Act to show that increased regulatory scrutiny caused manufacturers to reduce employment and investment, and to relocate operations to areas under weaker regulation. Greenstone et al. (2012) employ the same geographic variation and find that county-wide manufacturing productivity declined by 2.6% in areas under greater regulatory scrutiny. He et al. (2020) documents an extreme outcome among manufacturers in China, where TFP declined by 24% on average among worksites under greater water quality scrutiny.

Other research has found neutral or even positive effects of environmental regulation on productivity, which is consistent with Porter and Linde (1995)'s hypothesis that well-designed environmental regulation may catalyze innovation. Alpay et al. (2002) finds faster productivity growth coincided with rising environmental standards among Mexican food manufacturers. Berman and Bui (2001) shows that stricter air quality standards imposed in Southern California during the 1980s caused only temporary productivity declines at oil refineries and that productivity recovered to the national average by 1992. Albrizio et al. (2017) finds that stricter regulation boosted short-run productivity growth within a panel of European manufacturing firms and that the effect was strongest among the most productive firms. Similar heterogeneity is documented by Gray and Shadbegian (2003), which shows that pollution abatement costs were lower among more productive pulp and paper mills.

1.2. Background on OSHA Programmed Inspections

Unannounced worksite inspections known as *programmed inspections* have been central to OSHA's enforcement strategy since the OSH Act became effective in 1971. They comprise about half of all worksite inspections and have historically targeted construction and manufacturing establishments.[1] The 1978 Supreme Court case Marshall v. Barlow's Inc. established that OSHA must schedule inspections according to "neutral" selection criteria to legally access worksites when an employer refuses entry (Siskind, 1993).[2] OSHA complied by randomly scheduling inspections in high-hazard industries and maintained this policy until 1998 when a new data initiative enabled more targeted—albeit neutral—scheduling. Between 1979 and 1997, OSHA randomly inspected 3.8% of manufacturing worksites annually on average, which totalled approximately 240,000 unique inspections.[3]

Inspection Scheduling The vast majority of programmed inspections were scheduled in industries with high rates of lost workdays due to injury or illness according to the Bureau of Labor Statistics (BLS). Establishments from the 200 most hazardous 4-digit SIC industries were randomly assigned a rank that determined inspection priority; those in the top 100 hazardous industries were assigned the higher of two randomized ranks to boost their chance of inspection. All remaining "low hazard" industries were still subject to oversight but

[1]See Appendix Tables A.1 and A.2 for a detailed summary of OSHA inspections by reason for referral and targeted industry.

[2]MacLaury (1988) reports that OSHA's policy of targeting high-hazard industries began in 1977 before the Marshall v. Barlow's decision and argues that the Supreme Court decision had little effect on OSHA policy. Regardless, the sample begins after the Supreme Court case to ensure quasi-random sampling.

[3]Estimates for the number of US manufacturing establishments are from the US Census Bureau's Business Dynamics Statistics.

received only 5% of random inspections.[4] Small employers, those that had received a comprehensive inspection within the last several years, and establishments that participated in one of OSHA's voluntary cooperative programs were exempt from programmed inspections altogether.[5]

Inspections were also scheduled through an ongoing series of federal and regional "emphasis programs" that targeted specific worksite hazards such as asbestos. To ensure access to worksites, these programs largely followed the same quasi-random procedure of inspecting establishments from high-risk industries for a given hazard (OSHA, 1995). About 10% of programmed inspections between 1979 and 1997 can be linked to an emphasis program.

Evolving Enforcement Policy Ongoing revisions to OSHA's regulatory authority and enforcement policy in the late 1970s and 1980s generated substantial variation in the frequency and rigor of inspections. The largest changes occurred in the 1980s when OSHA became a primary target of President Reagan's Regulatory Relief Task Force and later Congressional pushback. The number of programmed inspections initially increased in the early 1980s as inspections became less comprehensive, punitive, and responsive to external complaints (MacLaury, 1988). Between 1982 and 1987, compliance officers skipped comprehensive inspections at worksites whose injury and illness rates were below the national average injury rate. Congressional scrutiny in the mid-1980s led to a change in OSHA leadership, a reinstatement of the 1980 budget, and a return to stricter inspection standards that are best exemplified by OSHA's 1987 "egregious case" policy (Vike, 2007; Siskind, 1993). Enforcement policy for the remainder of the 1980s and 1990s was largely stable aside from an increase in violation fines in 1990 (Vike, 2007).

Inspection Procedure During a typical inspection, an OSHA compliance officer conducted a "walkaround" through all potentially hazardous areas of a worksite. Compliance officers were required to issue citations for all observed violations including those that the employer could rectify immediately. Once the officer completed their walkaround and reviewed worksite injury records, they presented their findings to the employer and an employee representative. Employers received information on all violations, fines, and deadlines for correcting hazards. They could formally contest citations with an independent federal review commission or discuss any violation with their regional OSHA office.[6]

[4]See OSHA (1995) for OSHA's last description of its scheduling procedure before switching to more targeted scheduling. This directive was preceded by earlier previous versions in 1981 and 1990.

[5]According to OSHA (1995), small employers were defined as establishments with 10 or fewer employees during the previous 12 months and the exemption window for previously inspected establishments ranged from two to five years.

[6]See OSHA's Field Operations Manual for more details.

OSHA officers and directors had some discretion in determining the penalty for violations, which complicates their use as a continuous measure of workplace safety. OSHA officers considered four factors when assessing a penalty: the gravity of the violation, employer size, good faith efforts of the employer, and the employer's history of previous violations. Penalties were commonly reduced for small employers, those who made good faith efforts to improve workplace conditions, or those with a history of safety compliance. Serious, willful, or repeat violations were typically exempt from any penalty reduction and often resulted in increased fines up to a statutory maximum (OSHA, 2023).

Outcomes and Efficacy Table 1.1 summarizes the number of violations per inspection, total penalty, and penalty per violation for random inspections at manufacturing establishments.[7] Worksites rarely passed programmed inspections without any violations: about 83% of inspections resulted in at least one citation, and worksites received 6 violations per inspection on average. Compliance officers identified 38% of all violations as serious, meaning that they posed "a substantial probability that death or serious physical harm could result" (OSHA, 2023). Despite the frequency and severity of violations, the median total fine assessed against worksites with at least one violation was only $400 in 2012 dollars, which is consistent with OSHA's stated intent to not punish initial violations with large financial penalties (OSHA, 2023).

Table 1.2 lists the 10 most commonly violated federal OSHA standards.[8] These categories span most features of manufacturing operations including equipment, structures, material, and energy. The most cited category is *Machinery and Machine Guarding*, which accounts for 36% of serious violations. A survey by Weil (1996) found that the two primary costs for machine-related violations were one-time capital expenditures that bring the machinery into compliance and training workers on safe operating procedures.

Following an inspection, managers were typically given about one month to rectify any safety violations identified by OSHA compliance officers. Failure to abate a hazard could result in additional citations, follow-up inspections, and fines for each day past the abatement deadline.[9] Manufacturers maintained an impressive record of hazard abatement over the sample period, leaving only 0.7% of all infractions unaddressed by OSHA's deadline. Such high compliance is likely explained by the threat of steep future penalties and increased oversight for missing OSHA deadlines.[10]

[7] Inspection data is published through the Department of Labor's Enforcement Data Catalog and described in greater detail in Section 1.3.

[8] States maintain their own sets of standards. Violations of state standards are included in the main analysis below.

[9] Extreme cases of non-compliance were referred to OSHA's Assistant Secretary for further action.

[10] See Weil (1996) for further discussion on reasons for high compliance despite low initial penalties.

Early research into OSHA's first years as an agency found little impact on workplace safety, yet later studies have consistently shown that programmed inspections reduce injury rates and increase safety compliance.[11] Gray and Scholz (1993) finds that manufacturers who received inspection violations between 1979 and 1985 had 22% lower injury rates over the next three years.[12] More recent work by Haviland et al. (2012) and Levine et al. (2012) estimate similarly strong effects in Pennsylvania and California between the mid-1990s and mid-2000s. Ko et al. (2010) uses a large panel of inspections between 1972 and 2006 to show that safety violations declined by 28-48% between a worksite's first and second inspection.

1.3. Manufacturing Responses to Safety Inspections

This section investigates the effect of increased regulatory oversight following an OSHA inspection on manufacturer choices and outcomes. Results are estimated using the local projections approach to difference-in-differences from Dube et al. (2022) and show that inspections primarily affect establishment size rather than productivity or factor allocations. This methodology is then extended to non-random inspections that were triggered by reports of unsafe work conditions or a workplace accident. Non-random inspections precede slightly larger declines in establishment size and more persistent effects on productivity and capital-intensity.

1.3.1. Data

OSHA maintains detailed records on 5.0 million inspections and 12.9 million safety violations dating back to 1972. Their original record-keeping system covered 29 states whose operations were conducted federally, while remaining states operated independently and only began reporting data to the federal agency in 1991.[13] Each inspection record includes information on the inspection's purpose, scope, target worksite, resulting violations, and abatement efforts.

[11] Viscusi (1979) and Bartel and Thomas (1985) conclude that OSHA had little impact on workplace safety using industry-level data between 1972 and 1979. Another series of papers compares establishments inspected early and late in the same calendar year, and finds no immediate difference in injury rates during that year (Smith, 1979; McCaffrey, 1983; Ruser and Smith, 1991).

[12] Gray and Mendeloff (2005) conduct a follow-up study where they extend the sample to 1998. They find that this initial effect of 22% moderates between 1987 and 1991, and then disappears between 1992 and 1998. However, a subsequent study by Haviland et al. (2012) attribute this trend to the fact that later sample years were heavily skewed towards larger worksites that were typically less responsive to inspections.

[13] Inspection schedules for each state were still created at the federal level and disseminated to all OSHA offices regardless of their control. Some independently operating states submitted historical records going back to the mid-1980s when they began reporting in 1991, though comprehensive national coverage does not begin until 1991.

The main analysis below focuses on a sample of random inspections at manufacturing establishments that were conducted between 1979 and 1997. The sample excludes inspections that occurred within five years of a previous inspection to avoid confounding treatments, and it excludes inspections at establishments with fewer than 20 employees which were often exempt from comprehensive inspections or treated more leniently by OSHA. The sample also excludes "records inspections" that were common in the mid-1980s and did not necessarily include an inspection of workspaces. Finally, the small fraction of non-compliant establishments that did not address safety hazards by OSHA's deadline are removed.

Additional analysis below considers non-random inspections that were triggered by unsafe work conditions. The first set of analysis examines inspections that followed an employee complaint or external referral about alleged safety hazards, and the second set uses inspections after a workplace accident that resulted in a fatality or hospitalization. The sample period for this analysis is extended to 1977-2013 since the random scheduling procedure that was in effect for 1979-1997 is not relevant to these inspections.[14] All other sample restrictions are the same as those for the sample of random inspections described above.

OSHA inspection data is combined with annual data on manufacturing operations from the US Census Bureau. Establishments are matched between the two sources using a Jaccard index, which measures the similarity in business name and address.[15] Establishment pairs that are located in the same state and with the most similar names are identified as potential matches, and those that score at least 85% for name similarity or 75% for both name and address similarity are selected as matches. Of the 240,000 programmed inspections in the main sample, 147,000 are successfully matched to US Census Bureau records and 122,000 appear in the Census datasets described below.

Census records should cover the entire population of private establishments with employees, so mismatches likely arise from discrepancies between the establishment name or address reported to an OSHA compliance officer and those on record with the US Census Bureau. If unmatched establishments differ systematically from the matched sample, then the findings below may not be valid for all inspected establishments. For instance, younger or smaller establishments could be underrepresented in the analysis below if they were more likely to have errors in their business information. Appendix Table A.3 shows that inspections in the matched subsample are fairly similar to those in the full sample, which suggests that

[14] The sample period begins in 1977 and ends in 2013 due to the availability of US Census Bureau data described below. Results that compare complaint/referral inspections to randomized programmed inspections use a restricted sample period from 1979 to 1997.

[15] Specifically, the Jaccard index measures share of common bigrams (or character pairs) between two strings. Business names and addresses are cleaned before calculating the Jaccard index to improve accuracy. They are converted to lower case, stripped of punctuation, and stripped of common "stop" words such as *incorporated* or *company* for business names, and *street* or *avenue* for addresses.

selection bias is not a major threat to the external validity of the findings below.[16]

Employment and payroll data are from the US Census Bureau's Longitudinal Business Database (LBD), which is an annual census of all private non-farm establishments with employees. The Census Bureau continuously updates the LBD using business surveys and tax records, and ensures longitudinal consistency by tracking organizational changes such as mergers and acquisitions (Chow et al., 2021). A separate effort by Fort et al. (2016) matches each establishment to a longitudinally consistent NAICS industry code, which enables industry-level analysis over the entire sample period. Average annual compensation is defined as total payroll divided by the number of employees, and it is converted to 2012 dollars using the Consumer Price Index for all urban consumers published by the Bureau of Labor Statistics.

The Census Bureau collected detailed information on manufacturing operations in the Annual Survey of Manufactures (ASM) until its discontinuation in 2021. The ASM covered a rotating panel of approximately 50,000 establishments that was resampled every five years, and the survey was extended into a census of all manufacturing establishments every five years as well.[17] Sufficiently large or important establishments were sampled every year, while all remaining establishments were assigned a sampling probability based on their relative industry importance and breadth of manufactured products.[18]

Establishment output and capital measures are constructed following Cunningham et al. (2022). Output is measured as the total value of shipments adjusted for resales and changes in inventories. Total capital is defined as the sum of structures and equipment capital, which are both constructed using the perpetual inventory method. Each capital type is initialized using the first reported value for capital in a census year of the ASM. Capital in subsequent years is then defined as the sum of depreciated capital from the prior year and capital expenditures from the current year.[19] Output and capital variables are deflated to 1997 dollars using industry-specific investment (*piinv*) and shipments (*piship*) deflators published by the NBER-CES Manufacturing Industry Database.

TFPR is estimated following Olley and Pakes (1996) and implemented with ASM data as in Kehrig (2015). Establishments are assumed to have Cobb-Douglas production technology

[16]Inspections in the matched subsample have slightly more violations and higher total penalties. Most matched inspections also appear to have occurred at larger establishments despite lower average employment. Overall, these statistics do not raise concerns over selection into the matched sample.

[17]The Census of Manufacturers (CMF) was conducted in years ending in 2 and 7, and collects an expanded set of establishment variables including the value of the establishment's capital stock.

[18]Between 1979 and 1997, all establishments with at least 250 employees were sampled each year. See Census ASM Methodology for more information on sampling. https://www.census.gov/programs-surveys/asm/technical-documentation/methodology.html.

[19]Capital is similarly backfilled in cases where an establishment reports annual capital expenditures before it first reports capital. Industry-specific investment price deflators and depreciation rates are from the BLS.

that takes total capital, labor, materials, and energy as inputs,

$$y_{i(j)t} = a_{i(j)t} + \beta_j^k k_{i(j)t} + \beta_j^l l_{i(j)t} + \beta_j^m m_{i(j)t} + \beta_j^e e_{i(j)t},$$

where production function elasticities, β_j^x, are estimated for each 6-digit NAICS industry, j, over the entire sample period.

Labor input, l, is defined as total hours worked. The ASM only reports hours worked for production workers, so total hours is approximated as production-worker hours adjusted by the ratio of total payroll to production-worker payroll,

$$TH_{it} = PH_{it} \times \left(\frac{SW_{it}}{WW_{it}} \right).$$

Since average wages likely differ between production and non-production workers, total hours is effectively measured in units of production worker hours. If payroll data is not available, then total hours is defined as production-worker hours. Materials, m, is defined as the cost of materials, resales, and contract work done for the establishment. Energy, e, is defined as the cost of electricity and fuels. As with output and capital, nominal values for materials and energy are deflated to 1997 dollars using their respective NBER-CES deflators: *pimat* for materials and *pien* for energy.

The control function approach of Olley and Pakes (1996) uses investment as a proxy variable for unobserved, time-varying productivity to avoid endogeneity bias when estimating production function elasticities.[20] Investment is defined as the sum of capital expenditures on structures and equipment and deflated using the NBER-CES investment deflator, *piinv*.

Table 1.3 summarizes inspection outcomes and establishment characteristics in the final sample of random inspections and compares them to a representative sample of all manufacturers over the same period. The final sample skews towards larger, more capital-intensive establishments that typically pay a higher annual wage. Despite these differences, revenue-based productivity is very similar between the two samples. Table 1.4 conducts a similar comparison using the two samples of non-random inspections: those following a complaint or referral, and those following a workplace accident. These establishments skew even larger than those in the sample of random inspections and received substantially larger fines despite fewer average violations.

[20]See Wooldridge (2009) and Ackerberg et al. (2015) for more detailed discussion.

1.3.2. Event Study Design

The average response to an OSHA inspection is estimated using a difference-in-differences empirical design. Establishments that are inspected in the current period serve as the treatment group, and those inspected in some future year compose the control group. Treatment timing is staggered since new cohorts of establishments are inspected each year and removed from the "not-yet-treated" control group. The main empirical design treats inspections as binary treatment events, though additional analysis uses the inspection penalty as a continuous treatment variable.

Average treatment effects are estimated using the local projection difference-in-differences estimator (LP-DiD) proposed in Dube et al. (2022). This method addresses "negative weights" bias that occurs with staggered treatment in a Two-Way Fixed Effects (TWFE) model (see Goodman-Bacon (2021) for further discussion), and it easily handles various departures from the standard TWFE design such as heterogeneous treatment effects. Treatment status for establishment i in year t is represented by the binary indicator, D_{it}, and takes a value of 1 in all years following treatment. The first difference in treatment status, ΔD_{it}, equals 1 in the treatment year and 0 in all other years. For a given logged outcome variable, y_{it}, the response to an inspection after k years is estimated with the following equation,

$$y_{i,t+k} - y_{i,t-1} = \alpha^k + \beta^k \Delta D_{it} + \eta_t^k + \nu_j^k + \varepsilon_{it}^k, \tag{1.1}$$

where η_t^k captures time fixed effects and ν_j^k controls for 6-digit NAICS industry fixed effects. Under the necessary DiD assumptions, the coefficient, β^k, recovers the average treatment effect of inspection on the outcome, y_{it}. Since outcomes are measured in log differences, the coefficient β^k can be interpreted as the percent change in the outcome variable over k years as a result of inspection. Standard errors are clustered by inspection year and establishment, and observations are weighted by inverse sampling probability for survey data.

The two central assumptions underlying a DiD empirical design are (i) *no anticipation* and (ii) *parallel trends*. The no anticipation assumption states that treated units do not anticipate and respond to treatment before it has occurred. This condition is satisfied for random inspections since OSHA did not warn establishments of upcoming inspections and managers could not adjust operations beforehand.[21] The same may not be true for non-random inspections examined below. For instance, a manager who learns that a local news station has reported on their unsafe working conditions may adjust operations in anticipation of an OSHA inspection.

[21] OSHA provided advanced notice for certain types of inspections but these inspections are identified in the data and removed from the sample.

The parallel trends assumption states that average outcomes between the treated and control groups would have followed parallel paths had treatment not occurred. In settings with multiple treatment groups and periods, the standard parallel trends assumption for a TWFE model is complicated by the choice of control group and strength of assumptions about pre-treatment trends (see Marcus and Sant'Anna (2021) for further discussion). For this design, I follow the parallel trends assumption for a control group of "not-yet-treated" units as outlined in Callaway and Sant'Anna (2021). In a potential outcomes framework where $y_{it}(0)$ denotes the outcome for establishment i in period t in the absence of treatment, this parallel trends assumption states that,

$$E[y_{i,t}(0) - y_{i,t-1}(0)|\Delta D_{it} = 1] = E[y_{i,t}(0) - y_{i,t-1}(0)|D_{it} = 0],$$

where all units receive treatment by the end of sample period, $D_{iT} = 1$.

This assumption is satisfied if the path of an outcome variable is not affected by any differences—other than treatment itself—between treated and not-yet-treated units. Such differences could arise if OSHA's inspection scheduling system changed over the sample people so that establishments inspected after the change were different from those inspected earlier on. They may also arise from the selection effects inherent to using not-yet-treated establishments in the control group. On average, not-yet-treated establishments will be younger, smaller, and survive longer than treated establishments.

Potential parallel trends violations are mitigated by controlling for industry fixed effects and excluding small establishments. Industry fixed effects absorb differences between treated and not-yet-treated establishments that emerge from shifts in the industries targeted by OSHA. The estimating equation would also ideally condition on establishment age to mitigate selection on age and survival, but this is not possible with available Census data. Some of these effects are mitigated by excluding small establishments whose outcomes are most sensitive to age.

The analysis below assesses for parallel trends violations by plotting the pre-treatment effect over five years. Most results that are estimated using random OSHA inspections appear to satisfy the parallel trends assumption since they do not exhibit pre-trends. In contrast, there are noticeable pre-trends when using non-random inspection samples. Panel (b) in Figure A.1 plots the employment effect of programmed inspections that occurred after OSHA began targeting hazardous worksites in 1998. The plot shows that employment declined shortly before OSHA's inspection, suggesting that a confounding factor—most likely the workplace injuries or illnesses that caught OSHA's attention—may influence employment growth after inspection as well.

The absence of pre-trends does not necessarily guarantee that average potential outcomes would be the same after treatment.[22] One remaining concern is that worsening selection bias in the control sample may contaminate the estimated treatment effect. Control establishments that remained "not-yet-treated" for longer were likely younger and smaller at the start of the event study. Over a longer horizon, the set of control units becomes more restricted and the estimated average treatment effect may become increasingly biased. Panels (c) and (d) in Appendix Figure A.1 assess for this bias in the main results presented below. Panel (c) restricts the control sample to establishments that received inspection within five years of the treatment group, and Panel (d) uses a balanced panel of establishments. Estimates are compared over a five year horizon and show that the effect of a random inspection is robust to both sample restrictions.[23]

Heterogeneous Treatment Effects Heterogeneous effects are a key feature of models that link distortive policies to aggregate productivity losses (see Guner et al. (2008) for an example). Characteristics such as size and capital-intensity likely affect an establishment's safety compliance, treatment by OSHA, and response to safety violations. Heterogeneous treatment effects can be estimated by interacting the first-differenced treatment indicator, ΔD_{it}, with a given establishment characteristic, x_{it}. The resulting estimating equation is a simple modification of Equation 1.1,

$$y_{i,t+k} - y_{i,t-1} = \alpha^k + \beta^k \Delta D_{it} + \rho^k \Delta D_{it} \cdot x_{it} + \gamma^k x_{it} + \eta_t^k + \nu_j^k + \varepsilon_{it}^k, \tag{1.2}$$

where the interaction coefficient, ρ^k, captures how the treatment response varies linearly with x_{it}. Although these estimates cannot be interpreted causally, they offer insight into how changes to workplace safety regulation and enforcement may impact manufacturers differently.

Continuous Treatment Variable Manufacturer responses to inspection likely vary with the severity of their safety violations. The magnitude of an establishment's compliance shock can be approximated using the total fine assessed by OSHA. As discussed in Section 1.2, penalties depend on several factors that are unrelated to worksite safety such as the number of employees and history of safety compliance. Nonetheless, an establishment's total fine contains useful information about a worksite's distance from safety compliance after

[22] See Roth et al. (2023) for further discussion of issues with testing for pre-trends.

[23] The horizon in Panel (c) cannot be extended beyond five years because all control units are inspected within five years. The balanced panel used in Panel (d) could be extended to the full horizon of 10 years, but estimates becomes substantially less precise.

conditioning on establishment size.

Letting p_{it} represent the log total penalty resulting from an inspection, the continuous effect of safety compliance can be estimated with the following model,

$$y_{i,t+k} - y_{i,t-1} = \alpha^k + \beta^k \Delta D_{it} + \delta^k \Delta D_{it} \cdot p_{it} + \rho^k \Delta D_{it} \cdot x_{it} + X_{it}\Gamma^k + \eta_t^k + \nu_j^k + \varepsilon_{it}^k, \quad (1.3)$$

where δ^k captures how establishment responses vary with the inspection penalty, x_{it} represents log total employment, and $X_{it} = [p_{it}\ x_{it}]$. Since p_{it} is defined as log total penalty, this analysis excludes inspections that did not result in a penalty.

An establishment's distance from safety compliance is endogenous to other choices and characteristics. Issues of "selection into treatment" are common for difference-in-differences designs with continuous treatment, and Callaway et al. (2021) shows that the average treatment effect is biased under standard parallel trends assumptions. The average causal response to a larger penalty conflates the consequences of worse safety conditions with the behavior of establishments that receive larger penalties. Put simply, a serious safety hazard might be handled differently by an establishment that rarely has such hazards than by an establishment that has them often.

A stronger parallel trends assumption is necessary to interpret $\hat{\delta}^k$ as the average causal response to a larger inspection penalty. Callaway et al. (2021) defines one such assumption as independence between treatment dose and the path of potential outcomes. Since there is no reason to believe that such an assumption is satisfied for random OSHA inspections, $\hat{\delta}^k$ can only be interpreted as the combined effect of increasing treatment dose and selection into receiving that dose. For this reason—alongside complications to measuring safety compliance with total penalty—the preferred specification treats inspections as binary events.

1.3.3. Manufacturing Operations following Random Inspections

Figure 1.1 plots the estimated average treatment effect of a random OSHA inspection on establishment employment, capital-labor ratio, and TFPR. Panel (a) shows that average employment growth at inspected establishments was 15.6 percentage points lower over the decade following inspection. Employment growth fell fastest immediately after inspection and continued to steadily decline over the observed horizon. Since employment grew by a similar magnitude in the control sample, the net effect implies that inspected establishments stalled as their competitors expanded.

Panel (b) shows that capital intensity increased for five years following inspection before ultimately reverting to pre-inspection levels. At its peak, the capital-labor ratio of inspected establishments grew by 4.6 percentage points more on average than at control establishments.

Survey evidence from Weil (1996) suggests that greater capital intensity could be the result of increased investment to improve hazardous equipment or structures. However, Figure 1.2 shows that both structures and equipment capital fell each year after inspection, which suggests that capital-intensity increased because establishments were simply slower to adjust capital than labor. Five years after inspection, average total capital growth had fallen by only 1.6 percentage points while employment growth fell by 10.1 percentage points.

Random inspections had an even smaller impact on establishment productivity growth as plotted in Panel (c). Average TFPR growth fell by 1.9 percentage points relative to uninspected establishments over four years. As with capital intensity, this decline reverted by the end of the sample horizon and suggests that increased safety compliance did not have a permanent effect on TFPR. In an environment with monopolistic competition and Cobb-Douglas production technology as in Hsieh and Klenow (2009), this temporary decline in TFPR alongside a persistent fall in employment could be the result of an idiosyncratic shock to TFP. Establishments experiencing these shocks would raise prices and reduce scale until their revenue product matches that of competitors.

Under the assumption that more hazardous establishments receive larger TFP shocks and scale back operations more, OSHA inspections effectively reallocate production to safer establishments. Panel (a) of Figure 1.3 appears to support this claim by showing that employment growth fell more at establishments that received larger inspection penalties. The analysis treats an inspection's penalty as a proxy for safety compliance and plots the average causal response, $\hat{\delta}^k$, from Equation 1.3. After conditioning on establishment size, the plot shows that average employment growth fell by 6.0 percentage points more when inspection penalties were twice as large. As discussed in Section 1.3.2, this estimate confounds the effect of worse safety violations with the type of establishment that has such violations. Regardless of the mechanism, the results suggest that more hazardous establishments shrink by more following inspection.

Panels (b) and (c) in Figure 1.3 show that an establishment's employment response also depends on its size and capital-intensity. Using estimates of ρ^k from Equation 1.2, these plots show that employment growth declines less among larger and more capital-intensive establishments: average employment growth following inspection is 6.3 percentage points higher at an establishment an twice as many employees and 4.6 percentage points higher at an establishment with twice as much capital to labor. These estimates offer insight into how OSHA's enforcement strategy through worksite inspections may affect the distribution of manufacturing establishments. For instance, reducing the number of inspections at small establishments may encourage their growth and lower industry concentration.

Comparison to existing evidence The decline in employment growth found in Panel A of Figure 1.1 contradicts estimates from Levine et al. (2012), which finds that randomly inspected establishments had 2.7 percent (s.d. 1.6) higher employment than control establishments within one year of inspection.[24] While the precise reason for this disparity is unclear, there are some noteworthy differences in methodology that may have contributed. Levine et al. (2012) uses a sample of 409 inspected establishments that are drawn from all private industries in California between 1996 and 2006, whereas the current sample is composed of manufacturing establishments from across the United States that were inspected between 1979 and 1997. Control units in Levine et al. (2012) fit the criteria for inspection but are not known to be inspected in the future, and a single control unit is matched to each inspected establishment based on their 2-digit SIC industry code and region within California. The current methodology compares establishments within narrower 4-digit SIC codes but does not consider geographic proximity. Finally, the employment effects estimated above emerge steadily over the decade following inspection, whereas the results in Levine et al. (2012) are based on employment within one year of inspection. It would be interesting to revisit their analysis using a longer sample horizon and check whether employment at inspected establishments eventually declines.

1.3.4. Manufacturing Operations following Safety Incidents

This section investigates how manufacturers adjust operations after a safety incident that resulted in inspection by OSHA. Unlike Section 1.3.3, these estimates cannot be interpreted causally because inspections are endogenous to workplace safety conditions.[25] Nonetheless, these results offer insight into how workplace hazards affect manufacturing operations, and how establishment responses differ from those following random inspection.

Figure 1.4 plots the average change in employment, capital-labor ratio, and TFPR after OSHA investigated an employee complaint or external referral about unsafe working conditions. Panel (a) shows that these reports had a larger average effect on employment growth than random inspection: average employment growth fell by 20.1 percentage points over the decade following inspection. The pre-treatment period exhibits a significant pre-trend, with employment growing rapidly in the years prior to inspection. This trend suggests that

[24] In comparison, Panel A of Figure 1.1 reports that average employment growth among inspected establishments was 1.1(0.2) percentage points lower in the inspection year and 3.6(0.6) percentage points lower in the year following inspection.

[25] This endogeneity of safe working conditions is less concerning when control units are not-yet-treated establishments than establishments that are never inspected by OSHA. Yet, inspected establishments may behave differently from not-yet-inspected establishments in the years immediately prior to inspection, which would violate the parallel trends assumption.

concern over unsafe work conditions often coincides with expanding operations, and that expansion without consideration for safety may backfire.

Panel (b) shows that capital-intensity increased substantially following a safety complaint or referral. Unlike the average capital-labor response to random inspection, this increase persisted over the entire event horizon and suggests a permanent change in capital intensity. Panel (c) shows a similarly persistent effect on TFPR, with average TFPR growth remaining below that of not-yet-inspected establishments. Interestingly, the capital-labor ratio and TFPR do not exhibit a clear pre-trend as found in Panel (a).

Figure 1.5 plots establishment changes after a workplace fatality or hospitalization. These accidents precede an even larger average decline in employment growth, which fell 22.7 percentage points relative to control establishment over the event horizon. As with complaint inspections, accident inspections were followed by an increase in the capital-labor ratio and decline in TFPR. Given the comparatively small sample of accidents inspections, these estimates are less precise than those based on other inspections but nonetheless suggest a persistent shift in capital intensity and productivity.

1.3.5. Wage Response to Random Inspection and Safety Incidents

In addition to factor allocations and productivity, workplace safety likely affects workers' wages. Figure 1.6 shows that random inspections and safety incidents were both followed by sharp declines in average wage growth. Random inspections in Panel (a) led to a 3.5 percentage point deficit in wage growth after five years. Unlike Figure 1.1, these estimates contain a pre-trend and may not satisfy the parallel trends assumption. Panel (b) shows that average annual wages declined fastest after a complaint or referral about unsafe work conditions. Compared to control units, average wage growth fell by 4.3 percentage points one year after inspection before steadily recovering. Workplace accidents in Panel (c) also resulted in persistently lower average compensation, though the effect is not as sharp as that following a report of unsafe work conditions.

These results show that increased safety compliance and safety incidents are associated with lower average wages, yet the underlying mechanism for their decline is unclear. Under the theory of compensating differentials, safer working conditions could be viewed as a form of non-pecuniary compensation that justifies lower wages. Managers might also cut wages if increased regulatory enforcement reduces profitability through lower productivity or increased safety expenses. In the case of complaints about unsafe work conditions, managers might retaliate against employees by lowering compensation.

1.4. Characterizing Safe Establishments

Random inspections offer insight into the contemporaneous relationship between workplace safety and establishment characteristics without concern over measurement error or selection bias.[26] Estimates show that more productive and larger establishments are safer when measured on a per employee basis. Since inspection responses strengthen with the severity of safety violations, increased regulatory enforcement may disproportionately affect smaller and lower productivity establishments.

These findings are based on analysis that relates workplace safety to the three main establishment characteristics considered above: size, capital-intensity, and productivity. Ex ante, each variable has an plausibly ambiguous relationship with safety. For instance, productivity and safety could be positively related if better managers simultaneously improved operations and compliance. Alternatively, productivity may be negatively correlated with safety if managers cut corners on compliance to boost productivity. The analysis below is intended to resolve some of this ambiguity about the relationship between safety and operations.

Workplace safety is measured using both the number of violations that an establishment receives and the total penalty incurred for those violations. Each outcome is related to establishment characteristics with following regression model,

$$s_{it} = \alpha + \boldsymbol{X_{it}\beta} + \eta_t + \nu_j + \varepsilon_{it}, \tag{1.4}$$

where s_{it} represents a given safety outcome, $\boldsymbol{X_{it}}$ is the vector of establishment characteristics, η_t captures time effects, and ν_j controls for 6-digit NAICS industry fixed effects. Establishment characteristics include employment, the capital-labor ratio, and TFPR. Each covariate is logged, so resulting estimates can be interpreted as the percent change in the safety outcome that is associated with doubling a given covariate. Standard errors are two-way clustered by year and industry.

Resulting estimates are presented in Table 1.5. The first row shows that larger establishments violate more safety standards and receive larger penalties from OSHA. However, the coefficient in column 2 suggest that safety violations decrease with establishment size: doubling the number of employees at an establishment is associated with a 32.7% increase in the fine assessed by OSHA on average. The second row shows a more ambiguous relationship between capital intensity and safety compliance. The first column finds that more

[26] Measurement error is a concern in alternative measures of workplace safety such as self-reported injury and illness records; selection bias is a concern in samples of non-random, anticipated inspections. For instance, estimates that are based on the sample of accident inspections might not accurately represent safer establishments.

capital-intensive firms receive fewer violations, yet the second column finds essentially no association between the capital-labor ratio and total fine assessed by OSHA. Finally, the third column finds a negative association between safety and productivity: column 2 estimates that doubling establishment TFPR is associated with a 5.6% decline in the resulting inspection fine. According to Table 1.3, the interquartile range for logged TFPR is 0.7, suggesting that establishment productivity has a fairly small effect on workplace safety.

1.5. Approximating the Cost of Random Inspections

The direct cost of OSHA's random inspection program depends on how easily supply shifted from inspected establishments to unaffected competitors. In an extreme case of inelastic supply where slower output growth at inspected establishments was not replaced by increased production among competitors, the direct cost can be approximated using the decline in real output growth found in Panel A of Appendix Figure A.2. In a more conservative case where output growth shifted to competitors but remaining production at inspected establishments occurred with lower productivity growth shown in Figure 1.1, the cost can be approximated using this decline in productivity growth.[27] Each of these cases suggests an average percent loss in real output associated with inspection. We can estimate the dollar cost of OSHA's random inspection program by applying this percent loss to the total output of all inspected establishments.

The total annual output of inspected establishments can be approximated using NBER-CES manufacturing data and employment information collected by OSHA. For each 4-digit SIC code, the share of industry production that is subject to random inspection is assumed to equal the share of employment at inspected establishments. Each industry's employment share is then multiplied by its real value of shipments and aggregated across all industries. The resulting estimate approximates the annual real value of output among inspected establishments, which averaged $505 billion in 2018 dollars between 1979 and 1997. For context, the average annual value of all manufacturing shipments over the same period was $11.2 trillion in 2018 dollars, and the estimated average share of output subject to inspection was 5.7%.

Figure 1.7 plots the approximate dollar cost of one year of random inspections over a 10 year horizon. These costs are constructed as the product of real output among inspected establishments and the percent decline in output growth associated with the inelastic and elastic supply scenarios described above. The red line plots the cost of foregone output

[27]Since employment at inspected establishments remained fairly stable over the event study horizon, this approach abstracts from changes in establishment size.

growth when supply does not shift to competitors and the blue line plots the cost of slower productivity growth when supply does shift. As expected, the direct cost under inelastic supply is substantially larger and more persistent than that under elastic supply. Aside from a large jump in the last year of the event study, the cost under inelastic supply rises fastest in the first few years after inspection and converges to approximately $50 billion in lost manufacturing output towards the end of the event horizon. Under elastic supply, the cost of slower productivity growth peaks after four years at $9.5 billion but eventually reverts as productivity growth at inspected establishments catches up to that of control establishments. Manufacturing production should become increasingly elastic over time and the direct cost of inspections likely approaches the costs estimated under elastic supply. That said, these results illustrate that cost estimates for regulatory enforcement are highly sensitive to assumptions about supply elasticity.

This approach to approximating inspection costs has important limitations that are worth outlining. First, it does not account for other ways that OSHA shapes workplace safety. Inspections raise the expected cost of safety violations and induce greater compliance, but even the possibility of an OSHA inspection likely deters establishments from violating safety standards. This approach also ignores heterogeneous effects of inspections by solely relying on average responses. The heterogeneous treatment effects found in Figure 1.3 suggest that larger establishments had weaker responses to inspection, which implies that using unconditional averages for all establishments may overestimate the total impact of inspections. Finally, the above approximations do not consider the elasticity of substitution between manufactured goods, related price responses, and other general equilibrium effects that result from random inspections. Nonetheless, this analysis provides a rough estimate for the direct cost of OSHA's random inspection enforcement strategy.

1.6. Conclusion

This paper uses a sample of random OSHA inspections and a local projection approach to difference-in-differences to study the relationship between manufacturing operations and workplace safety. It finds that increased regulatory enforcement and safety compliance greatly reduced establishment size but only temporarily impacted capital-intensity and productivity. The same sample of random inspections shows that larger and more productive firms are typically safer. Further event study analysis also found large declines in employment following a serious workplace accident or a report of unsafe working conditions. These safety incidents had more persistent effects on capital-intensity and productivity than random inspections. Both random inspection and safety incidents also led to lower average annual wages for workers.

These results suggest that within-establishment factor misallocation is not a major economic cost of safety regulation, though industry-level costs may arise due to compositional effects. When considering the effect of safety regulation on aggregate productivity, policy-makers should consider the disproportionate effect of inspection on certain types of establishments such as small employers. The negative relationship between establishment safety and productivity documented in this paper suggests site-specific targeting could increase aggregate productivity by shifting production away from hazardous, less productive manufacturers.

1.7. Tables and Figures

Table 1.1: Summary of Inspection Violations and Penalties

	N	Mean	S.D.	25th	Median	75th
Violations	240,465	6	7	1	4	8
Serious Violations	240,465	2	4	0	1	3
Total Penalty	240,465	1,809	21,462	0	178	1,168
Penalty per Violation	1,392,909	312	4,510	0	0	294
Penalty per Serious Violation	535,094	636	1,085	0	392	852

Notes: This table summarizes the number of violations and penalty incurred per inspection as well as the penalty incurred per violation for all OSHA programmed inspections at manufacturing worksites between 1979 and 1997. *Serious Violations* are defined as violations with "a substantial probability that death or serious physical harm could result." Penalties are deflated to 2012 dollars using the BEA investment price deflator. Data on OSHA inspection violations are available through the Department of Labor's Enforcement Data Catalog.

Table 1.2: Top 10 Most Violated Federal OSHA Standards and Average Penalty

	All Violations			Serious Violations		
	N	Mean	S.D.	N	Mean	S.D.
Machinery and Machine Guarding	311,909	374	1,752	192,753	525	886
Electrical	172,893	267	2,382	57,295	708	908
Toxic and Hazardous Substances	129,157	283	1,752	49,255	561	1,454
Walking-Working Surfaces	64,887	317	4,070	20,491	789	1,119
Hazardous Materials	80,672	274	1,352	32,576	582	1,201
Materials Handling and Storage	55,575	202	860	14,170	676	1,041
Personal Protective Equipment	51,564	277	1,531	18,223	636	1,459
General Environmental Controls	39,694	699	2,168	23,315	1,016	1,339
Fire Protection	31,899	104	410	4,366	628	818
Exit Routes and Emergency Planning	34,694	224	1,349	6,987	902	1,643

Notes: This table reports the number of violations and average penalty for the top 10 most commonly violated federal OSHA standards. Violations are from the sample of programmed inspections at manufacturing worksites between 1979 and 1997. Penalties are deflated to 2012 dollars using the BEA investment price deflator. *Serious Violations* are violations with "a substantial probability that death or serious physical harm could result." More information on Federal OSHA standards can be found here: https://www.osha.gov/laws-regs/regulations/standardnumber/1910. Data on OSHA inspection violations are available through the Department of Labor's Enforcement Data Catalog.

Table 1.3: Summary of Random Inspection Sample

Sample	Variable	N	Mean	S.D.	25th	Median	75th
Inspection Sample	Violations	48,000	7.0	7.5	2.0	5.0	10.0
	Total Penalty ($K)	48,000	2.4	12.8	0.0	0.4	1.9
	Employment	48,000	100.6	203.7	30.0	50.0	102.0
	Capital-Labor Ratio	18,000	59.4	81.0	21.5	40.4	69.0
	Ln(TFPR)	18,000	2.0	0.6	1.7	2.0	2.4
	Total Capital ($K)	18,000	7.8	38.5	1.0	2.4	5.8
	Structures Capital ($K)	18,000	3.5	22.8	0.4	1.0	2.6
	Equipment Capital ($K)	18,000	4.2	21.2	0.5	1.2	3.0
	Output ($K)	18,000	17.8	64.3	3.2	6.5	15.2
	Avg Annual Wage	48,000	21.6	10.8	14.3	20.0	26.9
Representative Sample	Employment	6,279,000	51.7	257.6	4.0	10.0	33.0
	Capital-Labor Ratio	915,000	67.9	399.9	16.9	33.3	64.5
	Ln(TFPR)	915,000	2.1	0.6	1.7	2.0	2.4
	Total Capital ($K)	915,000	8.4	60.4	0.3	1.0	3.4
	Structures Capital ($K)	915,000	3.7	31.5	0.1	0.4	1.5
	Equipment Capital ($K)	915,000	4.7	36.1	0.2	0.5	1.8
	Output ($K)	915,000	17.7	111.2	1.0	2.8	9.4
	Avg Annual Wage	6,279,000	114.5	2,599.0	10.0	16.6	25.1

Notes: This table summarizes inspection outcomes and establishment characteristics for the final sample of random inspections at manufacturing establishments between 1979 and 1997 (*Inspection Sample*). The table also summarizes a representative sample of manufacturing establishments over the same period (*Representative Sample*). *Violations* reports the number of violations discovered during the inspection and *Total Penalty* reports the total fine assessed against an inspected establishment (deflated to 2012 dollars using the BEA investment price deflator). All other variables definitions can be found in Section 1.3. Data on inspection outcomes are from the Department of Labor's Enforcement Data Catalog and establishments characteristics are from the US Census Bureau.

Table 1.4: Summary of Non-Random Inspection Samples

Sample	Variable	N	Mean	S.D.	25th	Median	75th
Complaint/Referral Sample	Violations	54,000	5.8	7.9	1.0	3.0	8.0
	Total Penalty ($K)	54,000	5.1	27.9	0.0	0.9	4.3
	Employment	54,000	201.7	572.6	40.0	80.0	185.0
	Capital-Labor Ratio	27,000	82.9	132.2	24.0	45.7	87.3
	Ln(TFPR)	27,000	2.1	0.6	1.7	2.1	2.4
	Total Capital ($K)	27,000	24.1	110.8	1.6	4.5	13.5
	Output ($K)	27,000	57.8	298.5	5.0	12.3	35.2
	Avg Annual Wage	54,000	28.1	16.9	15.7	24.9	36.6
Accident Sample	Violations	7,000	4.3	6.7	1.0	2.0	5.0
	Total Penalty ($K)	7,000	12.1	69.0	0.3	2.6	11.2
	Employment	7,000	216.9	564.0	45.0	91.5	202.7
	Capital-Labor Ratio	4,000	100.0	148.0	28.6	55.6	108.9
	Ln(TFPR)	4,000	2.1	0.7	1.8	2.1	2.5
	Total Capital ($K)	4,000	28.6	131.3	2.1	5.7	15.9
	Output ($K)	4,000	61.4	230.0	6.6	15.7	42.6
	Avg Annual Wage	7,000	32.5	17.8	20.4	29.3	40.8
Representative Sample	Employment	11,830,000	49.0	234.0	4.0	10.0	32.0
	Capital-Labor Ratio	2,142,000	83.5	482.1	19.7	39.8	78.8
	Ln(TFPR)	2,142,000	2.1	0.6	1.7	2.0	2.4
	Total Capital ($K)	2,142,000	9.2	64.6	0.3	1.1	4.0
	Output ($K)	2,142,000	22.6	385.5	1.0	3.0	10.5
	Avg Annual Wage	11,830,000	97.1	4,968.0	11.8	20.4	32.8

Notes: This table summarizes inspection outcomes and establishment characteristics for non-random OSHA inspections at manufacturing establishments between 1977 and 2013. *Complaint/Referral Sample* covers inspections that resulted from employee complaints or external referrals about alleged safety hazards and *Accident Sample* covers inspections after workplace accidents that resulted in fatalities or hospitalizations. The table also summarizes a representative sample of manufacturing establishments over the same period (*Representative Sample*). *Violations* reports the number of violations discovered during the inspection and *Total Penalty* reports the total fine assessed against an inspected establishment (deflated to 2012 dollars using the BEA investment price deflator). All other variables definitions can be found in Section 1.3. Data on inspection outcomes are from the Department of Labor's Enforcement Data Catalog and establishments characteristics are from the US Census Bureau.

Figure 1.1: Effect of Random OSHA Inspections on Manufacturing Operations

(a) Employment

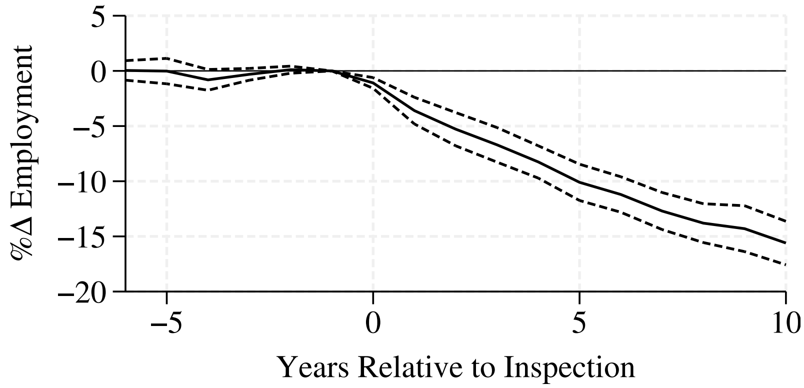

(b) Capital-Labor Ratio

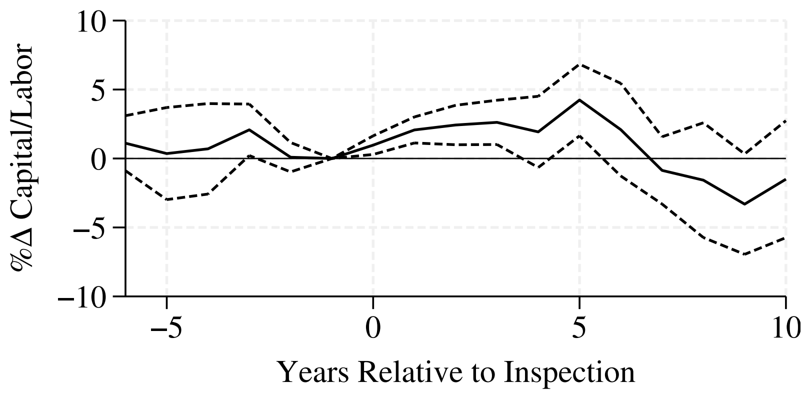

(c) TFPR

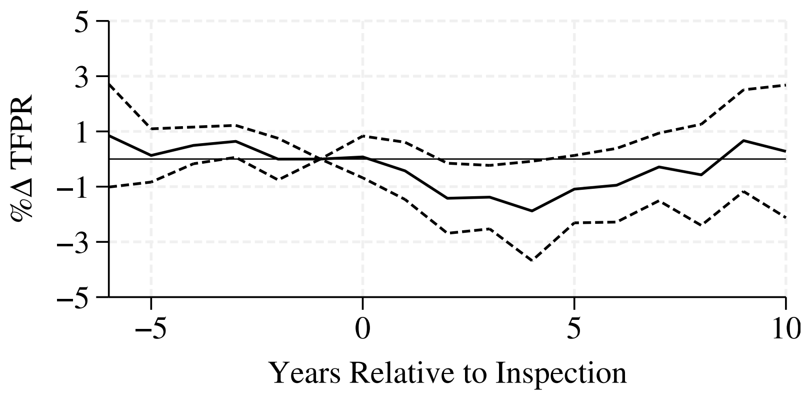

Notes: This figure plots the average response to randomly scheduled OSHA inspections at manufacturing establishments between 1979 and 1997. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in outcome due to inspection since the year preceding inspection. *Employment* in Panel (a) is defined as the number of employees; *Capital-Labor Ratio* in Panel (b) is the ratio of real total capital to the number of employees; and *TFPR* in Panel (c) is revenue-based total factor productivity (see Section 1.3 for further details). Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure 1.2: Effect of Random OSHA Inspections on Establishment Capital

(a) Total Capital

(b) Structures Capital

(c) Equipment Capital

Notes: This figure plots the average response to randomly scheduled OSHA inspections at manufacturing establishments between 1979 and 1997. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in capital due to inspection since the year preceding inspection. *Total Capital* in Panel (a) is defined as the sum of real structures and equipment capital; *Structures Capital* in Panel (b) is the real value of buildings and other structures; and *Equipment Capital* in Panel (c) is the real value of machinery and other equipment. Structures and equipment capital are both constructed using the perpetual inventory method as described in Section 1.3. Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure 1.3: Heterogeneous Treatment Effect on Employment

(a) Log Total Penalty

(b) Employment

(c) Capital-Labor Ratio

Notes: This figure plots the heterogeneous effect of random OSHA inspections on employment as a function of three establishment characteristics. The inspection sample covers programmed inspections at manufacturing establishments between 1979 and 1997. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in employment since the year preceding inspection that is associated with doubling a given establishment characteristic. *Log Total Penalty* in Panel (c) is logged total fine resulting from safety violations; *Employment* in Panel (b) is defined as the number of employees; and *Capital-Labor Ratio* in Panel (c) is the ratio of real total capital to the number of employees. Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure 1.4: Effect of Complaint/Referral Inspections on Manufacturing Operations

(a) Employment

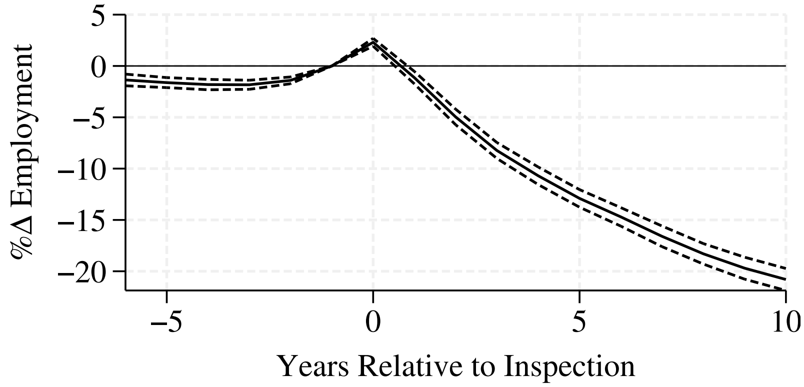

(b) Capital-Labor Ratio

(c) TFPR

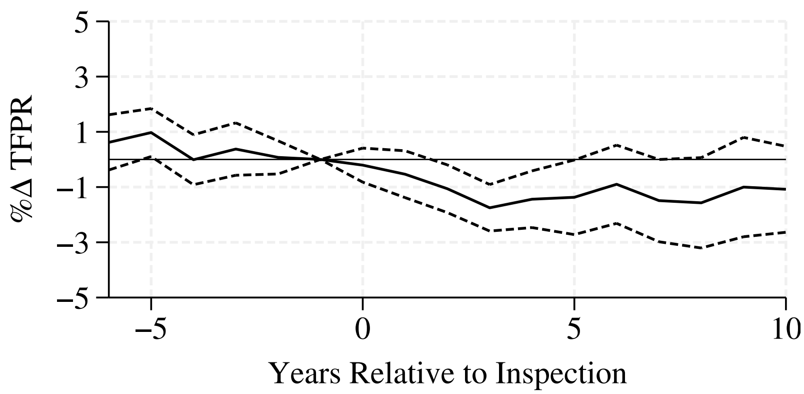

Notes: This figure plots the average response to OSHA inspections following an employee complaint or external referral at manufacturing establishments between 1977 and 2013. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in outcome associated with inspection since the year preceding inspection. *Employment* in Panel (a) is defined as the number of employees; *Capital-Labor Ratio* in Panel (b) is the ratio of real total capital to the number of employees; and *TFPR* in Panel (c) is revenue-based total factor productivity (see Section 1.3 for further details). Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure 1.5: Effect of Accident Inspections on Manufacturing Operations

(a) Employment

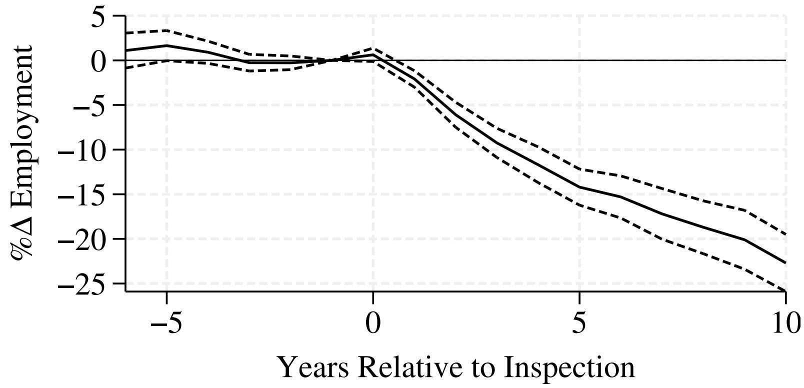

(b) Capital-Labor Ratio

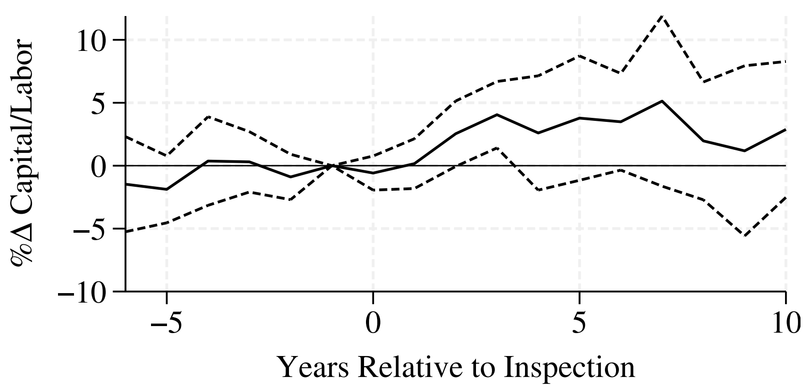

(c) TFPR

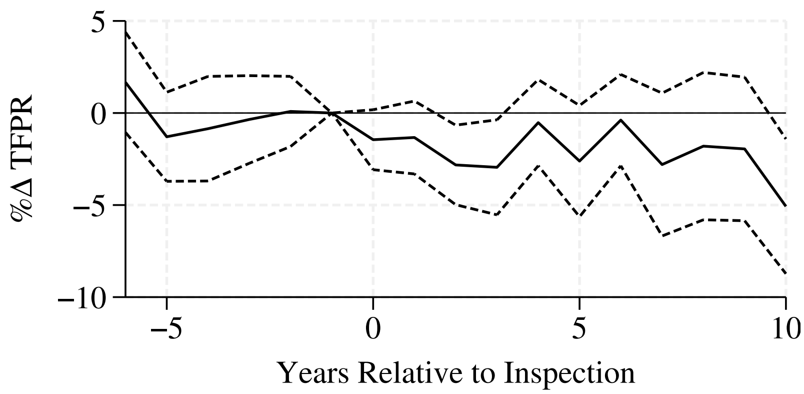

Notes: This figure plots the average response to OSHA inspections following a workplace accident at manufacturing establishments between 1977 and 2013. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in outcome associated with inspection since the year preceding inspection. *Employment* in Panel (a) is defined as the number of employees; *Capital-Labor Ratio* in Panel (b) is the ratio of real total capital to the number of employees; and *TFPR* in Panel (c) is revenue-based total factor productivity (see Section 1.3 for further details). Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure 1.6: Effect of OSHA Inspections on Average Annual Wages

(a) Random Inspections

(b) Complaint/Referral Inspections

(c) Accident Inspections

Notes: This figure plots the average wage response to three samples of OSHA inspections at manufacturing establishments. An establishment's average annual wage is defined as total payroll divided by the number of employees. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in wages associated with inspection since the year preceding inspection. *Random Inspections* in Panel (a) covers randomly scheduled programmed inspections from 1979-1997; *Complaint/Referral Inspections* in Panel (b) covers inspections following an employee complaint or external referral from 1977-2013; and *Accident Inspections* in Panel (c) covers inspections following a workplace accident from 1977-2013. Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Table 1.5: Safety Measures and Establishment Characteristics

	(1) Violations	(2) ln(Fine)
ln(Employment)	0.686	0.327
	(0.0695)	(0.0184)
ln(Capital-Labor)	-0.219	0.00956
	(0.0721)	(0.0112)
ln(TFPR)	-0.469	-0.0565
	(0.108)	(0.0247)
Constant	5.284	5.961
	(0.500)	(0.123)
Observations	49,000	29,000
R^2	0.091	0.338
within-R^2	0.00826	0.0596

Notes: This table estimates the relationship between establishment safety and other characteristics using a sample of randomly scheduled OSHA inspections at manufacturing establishments between 1979 and 1997. Estimates are from the model, $s_{it} = \alpha + \boldsymbol{X_{it}\beta} + \eta_t + \nu_j + \varepsilon_{it}$, where s_{it} represents an inspection safety outcome, $\boldsymbol{X_{it}}$ is a vector of establishment characteristics, η_t captures time effects, and ν_j controls for 6-digit NAICS industry fixed effects. The safety outcome in column 1 is the number of violations found during inspection and in column 2 is the total penalty assessed for those violations. Establishment characteristics include employment, the capital-labor ratio, and TFPR. Each covariate is logged, so resulting estimates can be interpreted as the percent change in the safety outcome that is associated with doubling a given covariate. The standard errors in parenthesis are two-way clustered by year and industry. Data on inspection outcomes are from the Department of Labor's Enforcement Data Catalog and establishments characteristics are from the US Census Bureau.

Figure 1.7: Direct Cost of OSHA's Random Inspection Program

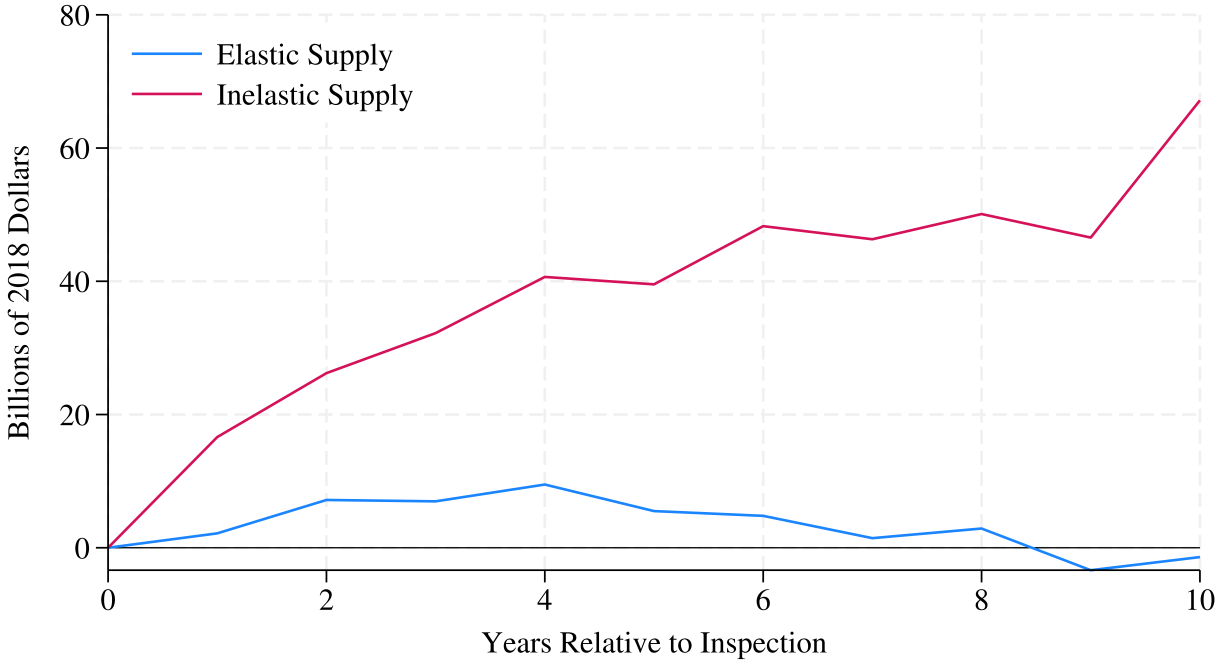

Notes: This figure plots two estimates for the direct cost of OSHA's random inspection program based either inelastic or elastic manufacturing supply. Each plot represents the cumulative cost of one year of random inspections over a ten year horizon. Costs are constructed as the product of real output from inspected establishments and the percent decline in output growth associated with either inelastic or elastic supply. The *Inelastic Supply* plot in red approximates the percent decrease in output using the average decrease in output growth among inspected establishments, and the *Elastic Supply* plot in blue approximates the output decrease using the decline in productivity growth following inspection. Total real output of inspected establishments is approximated as $505 billion in 2018 dollars based on NBER-CES manufacturing data and employment data from OSHA.

CHAPTER 2

Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach with Wenting Song

2.1. Introduction

Public information often goes unused because attention is scarce. Rational inattention models pioneered by Sims (2003) and a broader set of incomplete-information models (Mankiw and Reis, 2002; Woodford, 2009) consider firm managers who gather information to maximize value while facing cognitive costs of processing information. Inattention provides an intuitive microfoundation for monetary non-neutrality, yet empirically assessing the importance of attention is challenging because neither a firm's allocation of attention nor information-processing costs are readily observable.

This paper provides some of the first empirical evidence of the importance of firm attention to macroeconomic dynamics using a novel text-based measure. We document countercyclical firm attention and uncover substantial heterogeneity in attention across firms. Moreover, our measure is consistent with the asymmetric prediction of inattention models that attentive firms exhibit higher profit semi-elasticities in response to expansionary monetary shocks and lower semi-elasticities following contractionary shocks. We then use this measure to study macroeconomic implications of firm attention. Empirically, we find that attention mitigates the effects of macroeconomic uncertainty on firms' long-term performance. Quantitatively, we use this measure to calibrate information costs in a general equilibrium model with rationally inattentive firms. Firm inattention generates monetary non-neutrality and is a source of state dependence in monetary policy.

To construct our attention measure, we compile a corpus based on approximately 200,000 annual SEC filings of US publicly traded firms and search each document for macroeconomic keywords. We define two measures of attention: "prevalence," whether firm managers discuss macro conditions at all, and "intensity," the frequency with which managers discuss macro conditions. Our text-based classification of firm attention passes a number of sense checks: topic-specific attention is concentrated by industry; firms in more attentive industries adjust prices faster in response to monetary shocks; and attentive firms predict macro and firm-specific variables better in surveys.

We document two stylized facts about firm attention. First, firm attention is polarized. The majority of firms in our sample either mention macroeconomic conditions in every filing or in none of their filings. Second, attention is countercyclical. Among the remaining firms with time-varying attention, the number of firms that referenced macroeconomic news rose notably during recessions. We also study potential drivers of firm attention which include firm characteristics and macroeconomic uncertainty.

Our main empirical result validates our text-based measure by testing for asymmetry in firm performance that is predicted by inattention models: following a macroeconomic shock, the responses of firms with greater information-processing capacity should be closer to the optimal response regardless of the shock's direction. Therefore, attentive firms should exhibit higher profit elasticities in response to positive shocks and lower elasticities in response to negative shocks as they make decisions more accurately than inattentive competitors. We test for this asymmetry using an event-study design that exploits high-frequency variation in firms' market values around FOMC announcements. This test requires combining our prevalence attention measure with daily CRSP stock prices, quarterly Compustat firm financials, and high-frequency monetary shocks (constructed as in Gürkaynak et al., 2005; Gorodnichenko and Weber, 2016; Nakamura and Steinsson, 2018).

Consistent with the theoretical prediction, expansionary monetary shocks raise stock returns of attentive firms by 2% more than those of their inattentive peers, whereas contractionary shocks lower returns of attentive firms by 6% less. The suboptimal responses to monetary shocks by inattentive firms are direct evidence of the cost of inattentive behavior. Moreover, the asymmetry is inconsistent with several common concerns about measuring firm attention with text analysis: concern that filings contain macroeconomic buzzwords as a form of cheap talk to appease investors would imply a zero effect; concern that firms mention keywords solely as a function of exposure to monetary policy would imply symmetric responses to monetary shocks; and concern that stock returns vary with investor attention rather than firm attention would also fail to produce an asymmetric response.

We then examine how attention affects firm performance under varying degrees of aggregate uncertainty. We construct an uncertainty index based on the Survey of Professional Forecasters and measure performance in three dimensions: profitability, financial performance, and survival. The resulting estimates show that attentive firms outperform their inattentive competitors under increased uncertainty. Interestingly, attention appears to weakly reduce performance in low-uncertainty environments, which may hint at the cost of attention.

Finally, we present a quantitative rational inattention model calibrated using our new measure to study the aggregate implications of incomplete firm attention. In the model, firms with heterogeneous information costs optimally trade off between the precision of their signals of aggregate demand and the cost of acquiring and processing information. Consistent with our empirical findings, attentive firms have higher output semi-elasticities to expansionary monetary shocks and lower semi-elasticities to contractionary shocks. We incorporate observed countercyclicality of firm attention to show that the efficacy of monetary policy declines as the share of attentive firms rises and more firms set prices closer to the optimum. This new interpretation of attention-dependent monetary policy implies that central banks should expect policy interventions to be weaker when an aggregate shock has already drawn firm attention to macroeconomic policy.

Related Literature Our paper contributes to four strands of literature. First, we contribute to the empirical literature on macroeconomic expectations by developing an ongoing, broad-based measure of firm attention that extends back to the mid-1990s. Recent literature has highlighted the importance of expectations for macroeconomic policy[1] and consequently the need for empirical measures.[2] Existing research has successfully measured attention in lab experiments (Reutskaja et al., 2011), field experiments (Bartoš et al., 2016; Fuster et al., 2018), and for individual consumers and banks (Macaulay, 2020; Weitzner and Howes, 2021). Our methodology complements those measures as well as survey-based evidence on firm expectations by Tanaka et al. (2019), Coibion et al. (2018), and Candia et al. (2021), and enables researchers to explore questions that lie outside the coverage of existing surveys.

Second, our findings lend empirical support to a broad body of theoretical work on incomplete information as a source of monetary non-neutrality (Sims, 2003; Mankiw and Reis, 2002; Woodford, 2009). Microfoundations proposed in rational inattention and sticky information models are successful in explaining firm pricing (Maćkowiak and Wiederholt, 2009; Afrouzi, 2020; Yang, 2022), asset prices (Van Nieuwerburgh and Veldkamp, 2009), discrete choices (Matějka and McKay, 2015; Caplin et al., 2019), aggregate dynamics (Maćkowiak and Wiederholt, 2015; Afrouzi and Yang, 2021a), and reconciling micro and macro evidence (Auclert et al., 2020). Our results estimate a substantial cost of information frictions in the US data, providing empirical support for these theories.

Our findings on the relationship between countercyclical attention and monetary policy

[1] See, e.g., Coibion and Gorodnichenko (2015); Coibion et al. (2020); Malmendier and Nagel (2016).

[2] See Gabaix (2019) and Maćkowiak et al. (Forthcoming) for comprehensive surveys of existing measures of attention.

efficacy relates to existing literature on state dependencies of monetary policy. Tenreyro and Thwaites (2016) estimate non-linear responses in monetary policy that are weaker in recessions than in expansions. Vavra (2014), McKay and Wieland (2021), and Ottonello and Winberry (2020) consider volatility, durable consumption, and default risk as other channels through which state dependency arises. This paper suggests that attention may be an important source of state dependency of monetary policy.

Fourth, our paper relates to a broader and emerging literature that brings natural language processing techniques to economics. The seminal work of Loughran and McDonald (2011) applies the "bag of words" method to firm filings and develops word lists specific to economic and financial texts. Recent work has used textual analysis to study financial constraints (Buehlmaier and Whited, 2018), central bank communication (Hansen et al., 2018), firm-level political risk (Hassan et al., 2019), inflation expectation formation (Larsen et al., 2021), and uncertainty (Handley and Li, 2020). We contribute to this literature by constructing a set of keyword dictionaries based on macroeconomic news releases that correspond to nine macroeconomic topics. While this paper focuses on attention to monetary policy, our method for measuring attention and its effects can be generalized to other macroeconomic topics.

In a related paper, Flynn and Sastry (2022) independently and contemporaneously develop a text-based measure of firm attention to macroeconomic topics. They show, like we do, that their measure is countercyclical. Whereas we show that the stock prices of more attentive firms rise relative to less attentive firms in response to both positive and negative monetary shocks, they compare firms' labor market choices to those of a neoclassical model with full information and show the gap between model and firm behavior is negatively correlated with firm attention both over time and across firms. Together our papers present compelling evidence that our text-based measures contain information that is useful in predicting economic outcomes and that these predictions are consistent with interpreting these measures as measures of attention.

Road map The rest of the paper proceeds as follows: in Section 2.2 we describe our methodology for measuring attention and present evidence of the stylized facts listed above; in Section 2.3 we present a theoretical framework that incorporates attention and exposure to macro shocks and derive the predicted asymmetry; in Section 2.4 we outline an empirical strategy for testing the effects of attention on expected returns and present our results; in Section 2.5 we present the mitigating effects of attention on uncertainty; in Section 2.6 we construct a quantitative model of rational inattention and conduct policy counterfactuals; Section 2.7 concludes.

2.2. Textual Measure of Attention

This section presents our measure of firm attention, conducts preliminary validation exercises, and documents stylized facts about attention. We show that cross-industry patterns of our proposed measure and its correlation with price adjustment are consistent with predictions about attention behavior. We then highlight two key stylized facts: aggregate attention moves countercyclically over the sample period, and the majority of firms remain polarized between never and always paying attention. The section concludes with some reflections on the limitations and promise of textual analysis as a tool for measuring attention.

2.2.1. Data and methodology

10-K filings Our analysis uses all electronically available 10-K filings by publicly listed US companies between 1994 and 2019.[3] Under Regulation S-K, the US Securities and Exchange Commission (SEC) requires all public companies to disclose audited financial statements and a description of business conditions in these filings each year. Companies were phased into mandatory electronic filing between 1993 and 1996, meaning that our sample covers the universe of filers since 1996.[4] The final sample contains 201,751 documents submitted by 35,655 unique firms. Table 2.1 summarizes the length of these documents and unique vocabulary used by filers.

Table 2.1: 10-K filing length and vocabulary size

	N	Mean	Median	SD	Min	Max
Total word count	201,751	30,647	26,133	23,031	152	199,520
excl. stopwords	201,751	18,912	16,128	14,232	98	164,734
Unique word count	201,751	2,433	2,496	1,039	74	7,937
excl. stopwords	201,751	2,337	2,395	1,026	68	7,822

A discussion of economic conditions in an SEC filing typically appears in two contexts: (i) recent or future firm performance and (ii) the risk factors that shareholders face by investing in the company. The former context usually appears in Item 7, which requires managers to discuss the firm's financial conditions and results of operations. This section is written as a narrative and its length varies widely across firms (for instance, Item 7 of Alphabet's 2020

[3]Our methodology is also well-suited for quarterly 10-Q filings. However, we exclude these filings because they are less descriptive and do not require audited financial statements. We start our sample in 1994, since fewer than ten 10-K filings are available electronically in 1993 at the beginning of the phase-in process, and end our sample in 2019 before the onset of the COVID-19 pandemic.

[4]See SEC Release No. 33-7427 for more information about the phase-in process.

10-K filing is 17 pages long). Economic conditions as a source of risk appear in Items 1A and 7A, which detail general firm risks and near-term market risks, respectively.

Textual measure of firm attention To measure firm attention, we employ dictionary-based frequency counts that identify when firms discuss any of the following nine macroeconomic topics: general economic conditions, output, inflation, labor market, consumption, investment, monetary policy, housing, and oil. Each topic is matched with a keyword dictionary that consists of names of major macroeconomic releases from Econoday (the data provider behind Bloomberg's economic calendar) as well as words and phrases that commonly appear in popular articles on each topic. Any words or phrases that might apply to both aggregate- and firm-specific conditions are removed to avoid misidentification. For example, the phrase "interest rates" is excluded from the monetary policy dictionary because firms may mention interest rates in the context of their own liabilities. The dictionary of topics and associated keywords appears in Table B.1.

We then construct two measures of attention based on these keywords. Attention *prevalence*, d_{it}^k, indicates whether firm i mentioned any keyword related to a given topic k in period t:

$$d_{it}^k = \mathbb{1}(\text{Total topic-k words}_{it} > 0). \tag{prevalence}$$

Attention *intensity*, s_{it}^k, records the rate at which keywords are mentioned as a share of total words in the filing. We interpret this measure as the average intensity with which firms pay attention to economic conditions:

$$s_{it}^k = \frac{\text{Total topic-k words}_{it}}{\text{Total words}_{it}}. \tag{intensity}$$

The total word count is generated by following the parsing strategy in Loughran and McDonald (2011): each text is stripped of all numbers and "stop words," such as articles, and then mapped onto a dictionary of all words that appear in our sample of 10-K filings.

We treat *prevalence* as our baseline measure of firm attention in the majority of the paper. Since both measures are closely related, this avoids presenting duplicate results. The prevalence measure is also less susceptible to contamination by changes unrelated to firm attention. For instance, *intensity* will decrease if a firm adds a new appendix to its filing despite no change in its discussion of any topics listed above. Nonetheless, *intensity* is essential for understanding the intensive margin of attention and countercyclical variation documented below.

Table B.2 in Appendix B.1 reports summary statistics of the firms that are classified as attentive or inattentive according to our "general" topic. Firms that mention macro keywords

Figure 2.1: Average industry attention by topic

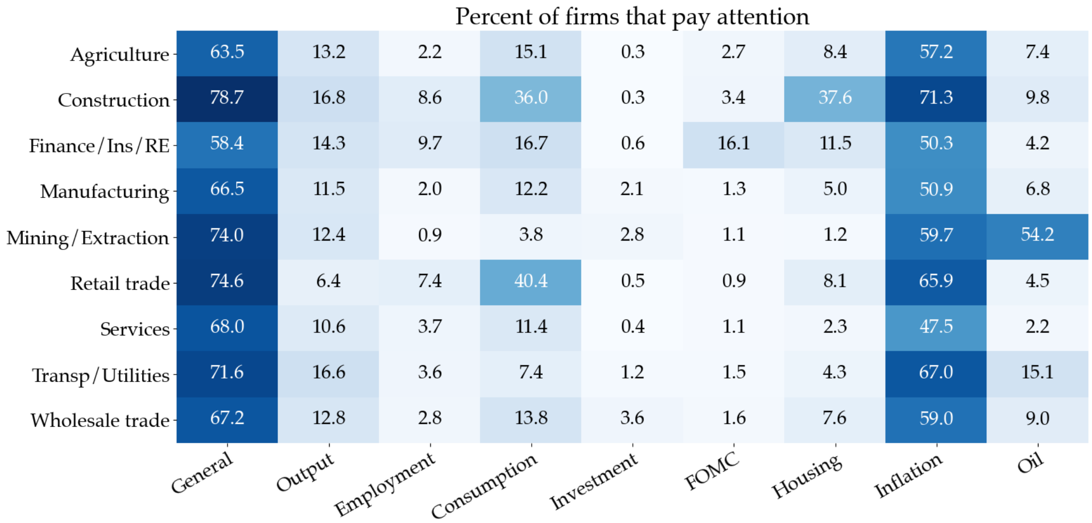

Notes: Heat map of the fraction of firms in an industry that discuss each macroeconomic topic. Industry is defined as a 2-digit NAICS. Darker color represents a higher fraction of firms that mention macro keywords.

tend to be larger—averaging \$7.3 billion in assets compared to \$2.9 billion among firms with no macro discussions—and older by just under four years on average. In contrast, average and median leverage appear fairly similar between the two groups of firms.

2.2.2. Sense check of the textual measure

This section uses cross-industry variation to test whether our prevalence measure is consistent with predictions of incomplete information models and then assesses how the measure relates to firms' forecast accuracy. We interpret the results as preliminary evidence that the prevalence measure may capture firm attention before presenting firm-level evidence in Section 2.4.

Cross-industry patterns of prevalence measure We first check whether the prevalence measure for the nine topics listed above is concentrated among commonly associated industries. Figure 2.1 reports the share of firms in each industry that discuss each topic in their 10-K filings, where industry is defined using 2-digit NAICS codes from Compustat. Since each topic uses a different set of keywords, differences in the prevalence measure across topics may reflect the relative popularity of keywords. Therefore, results should be interpreted across industries rather than across topics.

By and large, the prevalence measure for each topic is highest within related industries:

mining/extraction has the highest share of firms that discuss oil prices; retail trade has the highest share of firms that discuss consumption; and the financial sector has the highest prevalence on monetary policy (FOMC). While this cross-industry pattern is not unique to firm attention (for instance, a measure of profit exposure to each topic would produce the same pattern), it serves as a common sense check for our prevalence measure and suggests that textual analysis methods are capable of identifying firm attention.

Price adjustment following monetary shocks We next test whether industries with higher average prevalence adjust prices faster following monetary policy shocks, as predicted for attention in models of incomplete information (Mankiw and Reis, 2002; Maćkowiak and Wiederholt, 2009; Woodford, 2009). The association between prevalence and price response is estimated using the interaction between high-frequency monetary shocks—constructed as in Gorodnichenko and Weber (2016)[5]—and average industry prevalence in a local projection model (Jordà, 2005). Over an h-month horizon, our model takes the form

$$\log P_{s,t+h} - \log P_{s,t} = \alpha_s + \alpha_t + \beta_\nu^h \nu_t^{\mathrm{M}} + \beta_d^h d_{st} + \beta_{d\nu}^h d_{st} \nu_t^{\mathrm{M}} + \Gamma' Z_{st} + \varepsilon_{st}, \tag{2.1}$$

where $P_{s,t}$ is the BLS Producer Price Index (PPI) for industry s (4-digit NAICS) in month t, ν_t^{M} denotes the monetary shock in month t, d_{st} denotes average industry prevalence, and Z_{st} is a vector of controls including industrial production, a recession indicator, and industry size. We include industry and time fixed effects, $\{\alpha_s, \alpha_t\}$, and cluster standard errors by both industry and year. For ease of interpretation, monetary shocks are normalized so that positive values correspond to expansionary shocks. We exclude finance and utility industries as is common for estimating firm responses to monetary shocks.[6]

Figure 2.2 Panel (a) plots the estimated average price response, β_ν^h, and shows that prices rise in a hump-shaped manner following an expansionary monetary shock. At its peak, an unanticipated 25 basis point rate cut causes prices to rise by 1.8%. Panel (b) plots the marginal effect of average prevalence on an industry's price response, $\beta_{d\nu}^h$. Industries with a higher fraction of firms mentioning macro keywords raise prices faster in the first 10 months after a monetary shock, though the effect begins to decline after about seven months as the other industries catch up. This result is consistent with imperfect information models that predict faster price adjustment by attentive firms (e.g., Maćkowiak and Wiederholt, 2009).

Survey forecast accuracy The most direct test of our prevalence measure is whether it can predict a firm's forecasting accuracy. This can be implemented for a very limited subset

[5]See Section 2.4.1 for a detailed description of their methodology.

[6]See, for example, Ottonello and Winberry (2020), Acharya et al. (2020), and Cloyne et al. (2023).

Figure 2.2: Prevalence measure and price adjustment

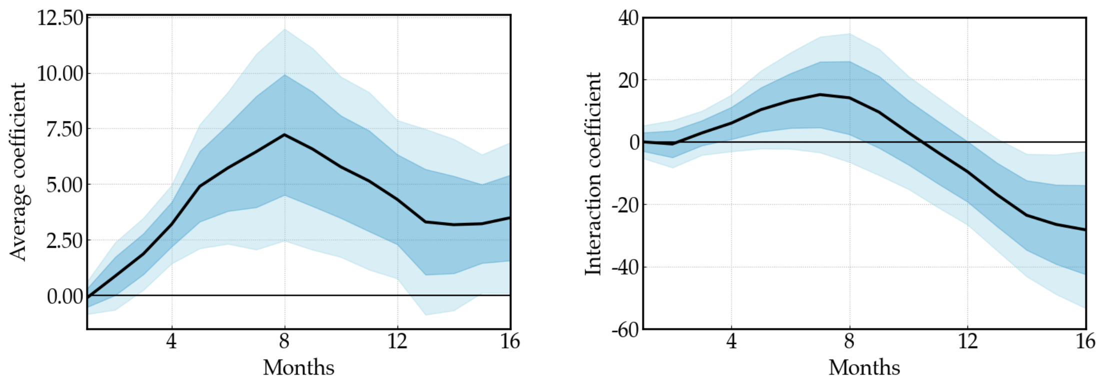

(a) Average responses to monetary policy **(b)** Marginal effects of attention

Notes: Panels (a) and (b) report the average and marginal coefficients, β_{ν}^{h} and $\beta_{d\nu}^{h}$, respectively, from estimating Equation (2.1) over months $h = 1, \cdots, 16$. We exclude finance and utility industries. Standard errors are double clustered by industry and year. Confidence intervals of 65% and 90% are reported. We have normalized the sign of monetary shocks so that positive shocks are expansionary.

of our sample that overlaps with a quarterly survey conducted by the Bank of Canada. The Business Outlook Survey (BOS) began in 1997 and interviews senior managers about macroeconomic and firm-specific conditions. It covers 100 firms every quarter based on quota sampling by industry, region, and size (Amirault et al., 2020).

Table 2.2: Attention and survey forecast accuracy

Panel A: Macro forecasts (2-year ahead inflation)

	General		**Inflation**		**Monetary**		
	Attn	Inattn	Attn	Inattn	Attn	Inattn	All
Avg accuracy	42%	33%	46%	33%	50%	40%	41%
N	125	12	87	50	12	125	137

Panel B: Micro forecasts (1-year ahead sales growth)

	General		**Consumption**		**Output**		
	Attn	Inattn	Attn	Inattn	Attn	Inattn	All
Avg accuracy	28%	14%	32%	24%	29%	26%	27%
N	116	7	47	76	38	85	123

Notes: This table reports average forecast accuracy by firm attention. Panel A reports the forecast accuracy of 2-year ahead inflation based on firm responses to BOS question 6.14. "General," "inflation," and "monetary" denote firm attention to the respective topics. Panel B reports the forecast accuracy of 1-year ahead sales growth based on firm responses to BOS question 2.6. "General," "consumption," and "output" denote firm attention to the respective topics.

Over the course of the survey, 137 firms in our sample have appeared in the BOS because they are either a US multinational with a presence in Canada or a Canadian firm listed in the US. Although these firms skew larger and are more likely to mention macro keywords than those in our full sample, they exhibit substantial variation in the prevalence measure across relevant topics.

The BOS includes two questions that pertain to firm forecasts. Question 6.14 asks about inflation expectations over the next two years, and question 2.6 asks about a firm's expected sales growth. The text for each is reproduced here:

> Question 6.14: Over the next 2 years what do you expect the annual rate of inflation to be based on the Canadian Consumer Price Index? (a) between 1–2%, (b) between 2–3%, (c) above 3%, (d) below 1%, (e) NA.
>
> Question 2.6: Over the next 12 months, is your firm's sales volume expected to increase (a) at a lesser rate, (b) the same rate, or (c) a greater rate, as over the past 12 months?

Since responses are multiple choice, we calculate the share of firms whose responses match the realized data and compare these shares between firms that are classified as attentive or inattentive using the prevalence measure. Responses to question 6.14 are compared to annual inflation over the next two years, from the OECD, and responses to question 2.6 are compared to sales volume in the next year, according to Compustat. We report forecast accuracy across all relevant economic topics: attention to general, inflation, and monetary news for forecasting inflation; and attention to general, output, and consumption news for forecasting sales.

Panel A in Table 2.2 shows that firms that are classified as attentive have more accurate inflation forecasts, and that the accuracy gap is highest for the inflation-specific prevalence measure: the accuracy rate of attentive firms was 13 percentage points higher than that of inattentive firms.

Panel B shows a similar pattern of accuracy when firms predict their own sales growth. Firms that are classified as attentive to aggregate demand (e.g., general, consumption, and output topics) forecast firm-specific demand better, which suggests that these firms translate macro information into better firm planning.

2.2.3. Stylized facts about firm attention

This section builds upon preliminary evidence that our text-based measures capture firm attention by summarizing how these measures vary over time and between firms, and then

Figure 2.3: Time series of attention to "economic conditions"

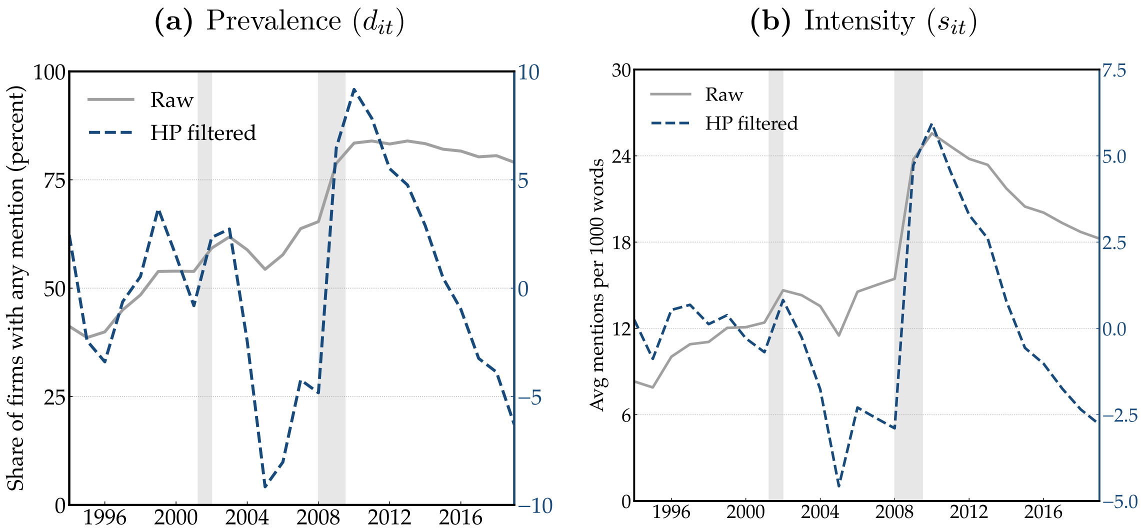

Notes: Time series of firm attention to the keyword "economic conditions." The left panel plots the prevalence measure and reports the share of firms that mention the keyword. The right panel plots the intensity measure and reports the average mentions of the keyword per 1,000 words. "Raw" refers to the unfiltered series, and "HP filtered" refers to the cyclical components of the HP-filtered series with smoothing factor $\lambda = 400$. Shares are reported in percent.

exploring potential drivers of firm attention. We document two key stylized facts: aggregate attention moves countercyclically over the sample period and the majority of firms remain polarized between never and always paying attention.

Countercyclical attention to economic conditions Both the share of firms that mention economic keywords and the intensity with which they are discussed vary countercyclically over our sample period. This is illustrated in Figure 2.3, which plots the annual average *prevalence* and *intensity* measures for the phrase "economic conditions," as well as detrended versions using an HP-filter ($\lambda = 400$).

Panel (a) shows that the share of firms mentioning "economic conditions" has steadily increased since 1994, with particularly rapid growth during the 2001 Recession and the Great Recession. Between 2008 and 2010, aggregate attention jumped by about 15 percentage points and remained elevated for the rest of the sample period. Average intensity in Panel (b) similarly spiked during the Great Recession but declined faster in subsequent years.

We point to these sharp dynamics around the 2001 Recession and the Great Recession as evidence of countercyclical attention, but we also acknowledge that our sample is limited and a longer time series is needed to confirm this pattern. In light of this, future sections use fluctuations in GDP growth and macroeconomic uncertainty—continuous measures related to business cycles—to provide further evidence of countercyclical attention and explore why attention increases in downturns.

Some models with endogenous firm attention predict the countercyclicality displayed in Figure 2.3. Maćkowiak and Wiederholt (2009) consider imperfect information firms that allocate attention between aggregate and idiosyncratic state variables. Increased aggregate uncertainty (itself countercyclical) induces these firms to shift attention toward aggregate conditions. Chiang (2021) decomposes the impact of lower expected productivity on attention into income and substitution effects. Countercyclical attention emerges among goods-producing agents when their marginal utility rises faster than the returns to attention falls under lower productivity.

Polarization in firm attention Despite the countercyclical dynamics documented above, most firms in our sample are polarized between either always or never discussing economic conditions in their 10-K filings. Figure 2.4 illustrates this by plotting each firm's share of filings that mention the same key phrase, "economic conditions."[7] The resulting distribution in Panel (a) is heavily concentrated at each extreme, with about three quarters of firms taking values of either 0 or 1. This suggests that most variation in attention occurs across firms rather than within firms and countercyclical variation is caused by a relatively small subset of filers.

To test whether polarized attention is driven by firms with few filings, we overlay a second histogram in Panel (a) that restricts to firms with at least five years of data. This adjustment greatly reduces the share of firms that never pay attention, yet over half remain polarization between always- and never-attentive firms.

We also test whether cross-industry differences in attention are responsible for polarization. Panel (b) in Figure 2.4 demeans by industry, which explains approximately one quarter of the attention variation. The distribution now contains a large mass of firms around their industry average (i.e., industries with little attention dispersion), while the remaining firms form a bimodal pattern consistent with polarization.

The presence of any inattentive firms may be surprising given that most US macroeconomic data are readily available. However, this result is consistent with a broader interpretation of attention that includes information processing, communication, and optimal response in addition to information acquisition. As highlighted in Reis (2006), firms likely require significant resources and expertise to process, summarize, and forecast macroeconomic series into sufficient statistics that inform firm decision-making. This is consistent with plant-level

[7] In this section, we illustrate patterns of countercyclicality and polarization using attention to economic conditions. Appendix Section B.1.2 reports the times series and histograms of firm attention to all 9 macro topics, which show similar patterns.

Figure 2.4: Share of filings that mention “economic conditions”

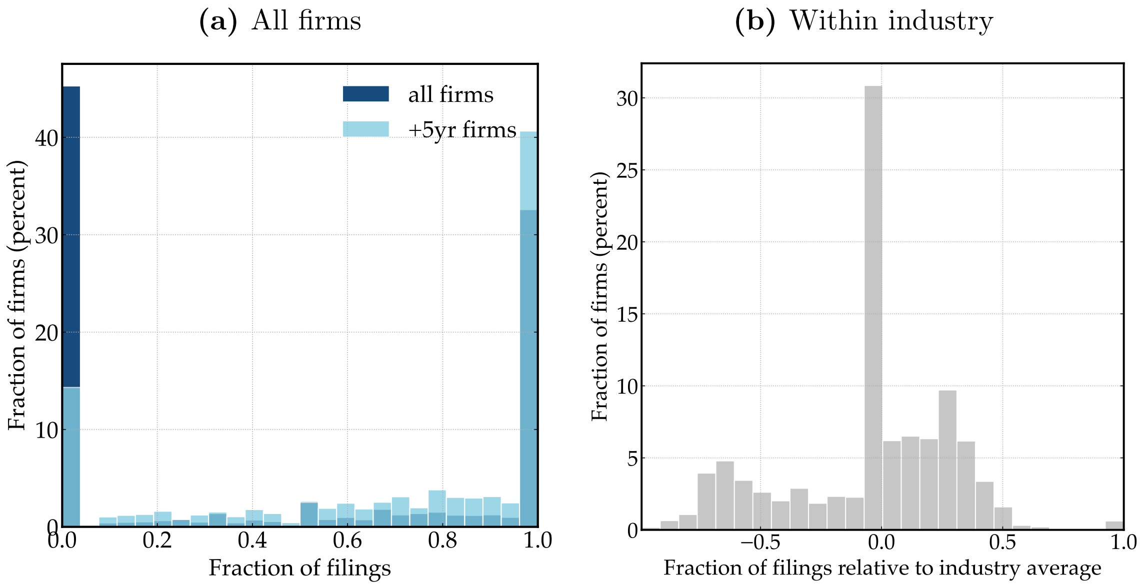

Notes: Histograms of the share of filings by a firm that mention “economic conditions.” The left panel shows the histogram of the average fraction of filings that mention the keyword “economic conditions” over the sample period of 1994–2019. Dark blue bars correspond to the distribution of all firms, and light blue bars correspond to firms appearing for at least 5 years in the sample. The right panel shows the histogram of the time series averages of the residuals of firm attention to “economic conditions” after regressing on industry fixed effects. Shares of firms on the vertical axes are reported in percentages.

evidence from Zbaracki et al. (2004). To this end, Abis and Veldkamp (2023) estimate a data production function that uses labor and capital inputs to produce knowledge from unstructured data.

Potential drivers of firm attention To understand what motivates attention and how attentive firms differ from their competitors, we turn to potentially related firm and macroeconomic factors. First, we estimate the relationship between attention and four firm variables—size, age, leverage, and productivity—both cross-sectionally and within-firm. We then examine how attention evolves alongside GDP growth and aggregate beliefs from the Survey of Professional Forecasters. Our findings suggest that size, age, productivity, economic growth, and aggregate uncertainty are important to understanding observed variation in attention.

The relationship between attention and firm characteristics is estimated with the following pair of regressions,

$$\text{Cross-firm variation:} \quad y_{i(j)t} = \alpha_j + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it} \tag{2.2}$$

$$\text{Within-firm variation:} \quad y_{i(j)t} = \nu_i + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}, \tag{2.3}$$

where $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t})\}$ represents attention by firm i in industry j and year t, and $\mathbf{X}_{it}$ is a vector of firm variables including size, age, leverage, filing length, and productivity.[8] Note that our intensity measure, s_{it}, is logged for ease of interpretation. Equation (2.2) includes time and industry (4-digit NAICS) fixed effects to highlight cross-sectional variation, while Equation (2.3) uses firm fixed effects to isolate within-firm variation. The first model clusters standard errors by industry and year, and the second model clusters by firm and year.

Data on firm characteristics are from Compustat or firms' 10-K filings. Size is measured as the log of total assets, age as the number of years that a firm has appeared in Compustat, and leverage as the debt-to-asset ratio.[9] Productivity is estimated using the control function approach from Olley and Pakes (1996) and implemented with GMM as in Wooldridge (2009).[10]

Results for this analysis are displayed in Table 2.3. The first two columns report estimates for Equation (2.2) and the second two columns do the same for Equation (2.3). Columns using our prevalence measure, d_{it}, capture changes in attention along the extensive margin, while those using our intensity measure, $\log(s_{it})$, restrict the sample to attentive firms and measure changes on the intensive margin. Both outcome variables are scaled by 100 so that units for d_{it} are percentage points and those for $\log(s_{it})$ are percents.

By and large, we find that larger, older, and more productive firms exhibit higher rates of attention. The association with size and productivity is strongest across firms, while the effect of age appears almost exclusively within-firm. Together, these results suggest that larger and more productive firms pay greater attention to aggregate conditions over time, while smaller and less productive competitors may never do so.

[8] Existing literature has found each of these characteristics to be relevant for the transmission of macroeconomic policy (Gertler and Gilchrist, 1994; Ottonello and Winberry, 2020; Cloyne et al., 2023).

[9] We exclude observations with leverage greater than 100% (about 3% of the sample) since this ratio is susceptible to extreme values. Filing length is measured as the log of total words appearing in a firm's 10-K. Unlike the other firm variables, it is intended to control for information provision, which affects the likelihood and frequency of keywords, and is therefore not included in the results below.

[10] Firm output is measured as total sales (SALE) deflated by the BEA's implicit price deflator, and labor is defined as total number of employees. Firm capital is constructed using the perpetual inventory method where capital stock for each firm is initialized as Gross Property, Plants, and Equipment (PPEGT), and annual net investment in all subsequent years is defined as the change in Net Property, Plants, and Equipment (PPENT). The capital of each period is defined as the sum of capital from the previous period and net investment. Finally, nominal capital is deflated using the BEA's investment price deflator.

Table 2.3: Firm characteristics and attention

	Cross-firm		Within-firm	
	d_{it}	$\log(s_{it})$	d_{it}	$\log(s_{it})$
Size (log total assets)	1.25***	4.25***	0.80	3.47***
	(0.22)	(0.56)	(0.48)	(0.85)
Age	0.00	0.42**	1.78***	3.33***
	(0.07)	(0.18)	(0.04)	(0.17)
Leverage	0.57	6.01*	-1.14	4.20
	(1.77)	(3.31)	(1.41)	(3.32)
Productivity (TFPR)	1.27***	2.89***	1.20***	0.78
	(0.34)	(0.73)	(0.32)	(0.67)
Observations	73101	55276	72283	54365
R^2	0.313	0.290	0.635	0.698
Industry FE	yes	yes	no	no
Firm FE	no	no	yes	yes

Notes: Columns (1) and (2) report estimates for $y_{i(j)t} = \alpha_j + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}$, while Columns (3) and (4) report estimates for $y_{i(j)t} = \nu_i + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}$. The outcome variable, $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t}))\}$, represents attention by firm i in industry j and year t, and $\mathbf{X}_{it}$ is a vector of firm variables including size, age, leverage, filing length, and productivity. The first two columns include industry fixed effects (4-digit NAICS), and the second two columns include firm fixed effects. All four columns include year fixed effects and control for filing length (log words), though not reported. Outcome variables are scaled by 100 so that units for d_{it} are percentage points and those for $\log(s_{it})$ are percents.

Next, we consider how attention varies with aggregate conditions and beliefs. We estimate the magnitude of countercyclical attention observed in Figure 2.3 by regressing our attention measures on annual real GDP growth and then see how attention comoves with three measures of aggregate uncertainty: the interquartile range of expectations, the consensus forecast error, and the absolute value of that error.[11] Each measure emphasizes a different dimension of uncertainty. The interquartile range captures disagreement among forecasters, the consensus forecast error considers how attention responds to positive or negative surprises differently, and the absolute forecast error isolates the accuracy of consensus beliefs regardless of error direction.[12]

[11]Given a series, x_t, and a sample of one-period-ahead forecasts, $\hat{x}_{it}$, the interquartile range is defined as $\mathrm{IQR}(\hat{\mathrm{x}}_{\mathrm{it}}) = \mathrm{P}_{75}(\hat{\mathrm{x}}_{\mathrm{it}}) - \mathrm{P}_{25}(\hat{\mathrm{x}}_{\mathrm{it}})$, where P_Y is the Yth percentile of the forecast sample. The consensus forecast error is defined as $\mathrm{FE}(\hat{\mathrm{x}}_{\mathrm{it}}) = \mathrm{x}_{\mathrm{t}} - \mathrm{P}_{50}(\hat{\mathrm{x}}_{\mathrm{it}})$.

[12]Data on macroeconomic expectations are from the Survey of Professional Forecasters, which has been administered on a quarterly basis since 1968. In each installment, a panel of respondents forecast several economic indicators up to one year in the future. We focus on one-quarter-ahead forecasts for real GDP growth and the unemployment rate. Uncertainty is constructed for each series at a quarterly frequency, standardized over the sample period, and then averaged into an annual composite uncertainty index. Forecast errors for unemployment are inverted so that positive values correspond to positive economic surprises.

Table 2.4: Aggregate variables and attention

	Prevalence, d_{it}				Intensity, $\log(s_{it})$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rGDP growth	-0.59***				-5.14***			
	(0.17)				(1.23)			
IQR index		3.29***				14.90**		
		(0.71)				(6.54)		
Abs(FE) index			1.33				23.05***	
			(0.83)				(6.34)	
FE index				-1.87***				-5.41
				(0.58)				(6.65)
Observations	129416	129416	129416	129416	99041	99041	99041	99041
R^2	0.651	0.651	0.651	0.651	0.749	0.747	0.750	0.744
Firm FE	yes	yes	yes	yes	yes	yes	yes	yes
Quadratic Trend	yes	yes	yes	yes	yes	yes	yes	yes

Notes: This table reports estimates of β from the model in (2.4): $y_{i(j)t} = \nu_i + \delta z_t + \beta x_{it} + \varepsilon_{it}$, where z_t represents either real GDP growth or one of three uncertainty indices: the IQR index, Abs(FE) index, or FE index. The outcome variable $y_{i(j)t}$ represents the attention of firm i in year t, x_{it} is the 10-K filing length, and ν_i represents a firm fixed effect. The first four columns use attention prevalence, d_{it}, as the outcome variable, and the last four columns use log intensity, $\log(s_{it})$. Standard errors are clustered by both firm and year.

The relationship between attention and aggregate variables is estimated with the following model,

$$y_{i(j)t} = \nu_i + \delta z_t + \beta x_{it} + \varepsilon_{it}, \tag{2.4}$$

where $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t})\}$ again represents attention, z_t is our aggregate variable of interest, x_{it} is the 10-K filing length, and ν_i represents a firm fixed effect. Standard errors are clustered by both firm and year.

The resulting estimates are reported in Table 2.4. Columns 1 and 5 show a strong, countercyclical pattern of attention, while Columns 2 and 6 show a strong positive association between forecaster disagreement and firm attention. Column 4 suggests that higher rates of attentive firms are associated with negative economic surprises, while Column 7 suggests that firms pay greater attention when consensus forecasts are less accurate.

2.2.4. Limitations and promise of textual measures

Boilerplate language is a key concern when using regulatory filings to measure firm attention. Filings are often written collaboratively between managers and legal departments, and evidence suggests that firms include certain statements within 10-K filings to appease investors or lower liability (Cao et al., 2020). Moreover, firms likely save time and resources by revising previous filings rather than starting from scratch each year.

The methods used above cannot distinguish between authentic attention to macroeconomic conditions and *cheap talk* references or recycled language that does not reflect current management practices. We address this shortcoming in Appendix B.3.1 by measuring the diversity in filing language with a Jaccard score of lexical similarity and testing whether our main findings are robust in the most linguistically diverse 10-K sections. Table B.8 confirms that our key findings are not driven by the most repetitive and standardized sections.

Even greater measurement error may arise from misidentifying firms as inattentive because they do not mention a certain keyword or discuss economic conditions when such conditions pose a financial risk.[13] *False negatives* can result from methodological limitations or variation in the amount of information that firms choose to disclose. It is worth noting that firm managers are obligated to disclose any material risk factors under SEC Regulation S-K.[14] Those who track inflation, unemployment, or any other topic because they are considered material risk factors are obligated to disclose this to their investors. For the purposes of this paper, underestimated attention would attenuate our results and imply that our current estimate for the cost of information frictions serves as a lower bound.

Text analysis methods also hold tremendous promise for uncovering a more refined depiction of firm attention and expectations formation. We illustrate these capabilities with two approaches for identifying the context in which firms discuss economic conditions. Appendix B.3.2 uses a Latent Dirichlet Allocation (LDA) unsupervised model to categorize words adjacent to each keyword and produces nine unlabeled "topics" in which keywords appear. Appendix B.3.3 uses the itemized structure of 10-K filings to identify sections that contain the most keywords. This analysis shows that keywords typically appear in sections that discuss firm risk factors (Item 1A) and business operations (Item 7A).

2.3. Illustrative Framework

This section derives a testable implication of firm attention to address a key identification challenge for our text-based measure: whether it captures *exposure* rather than attention to macroeconomic conditions. We present a stylized model in which firms are heterogeneous in both attention and exposure to an aggregate state variable and then consider how firm outcomes vary with each source of heterogeneity. The model predicts contradictory responses

[13] One reviewer compared such firms to smoke detectors, which are (ideally) always on but only beep in the presence of smoke.

[14] Konchitchki and Xie (2022) show that firms are subject to ligitation for undisclosed macroeconomic risks.

to aggregate shocks under varying attention and exposure, which guides our empirical design in Section 2.4. The model environment is kept minimal to highlight key mechanisms before Section 2.6 incorporates more realistic assumptions.

Environment The model is static. Consider a firm whose profits, $\pi(s,a)$, depend on an aggregate state variable, s, and a firm action, a. Assume that $\pi(s,a)$ is twice continuously differentiable, a single-peaked function of a and maximized at $a^* = s$. For concreteness, we think of a as the price that a monopolistically competitive firm sets and s as the exogenous optimal price determined by factors outside of that firm's control, as in Woodford (2009).

Firm profits can be approximated under a second-order log approximation around the non-stochastic steady state as[15]

$$\hat{\pi}(\hat{s},\hat{a}) = \pi_s(\bar{s},\bar{a})\bar{s}\hat{s} + \frac{1}{2}\left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\hat{s}^2 + \frac{1}{2}\pi_{aa}(\bar{s},\bar{a})\bar{a}^2(\hat{a}-\hat{s})^2, \tag{2.5}$$

where $\bar{s}$ and $\bar{a}$ denote the steady-state values; $\hat{\pi}$, $\hat{s}$, and $\hat{a}$ denote the log deviations from the steady state; and $\pi_s \equiv \frac{\partial}{\partial s}\pi(s,a)$, $\pi_{aa} \equiv \frac{\partial^2}{\partial a^2}\pi(s,a)$, and $\pi_{ss} \equiv \frac{\partial^2}{\partial s^2}\pi(s,a)$.

Lastly, assume that firm profits are increasing in s, $\pi_s > 0$, and that the profit function is concave in the own action, $\pi_{aa} < 0$.

Attention and Exposure We can now define attention and exposure in the model. A firm is more exposed to aggregate conditions if its profits are more sensitive to aggregate shocks, while a firm is more attentive if its actions are more sensitive to shocks. Definitions 1 and 2 formalize these ideas.

***Definition* 1 (attention).** *Let a firm's action be a function of the state:* $\hat{a} = f(\hat{s})$, *with* $f(0) = 0$ *and* $0 < f'(\hat{s}) \leq 1$. *Firm* i *is attentive to macroeconomic conditions if* $f_i'(\hat{s}) = 1$, *and firm* j *is inattentive to macroeconomic conditions if* $0 < f_j'(\hat{s}) < 1$.

An attentive firm reacts one-for-one with innovations to the aggregate state, whereas an inattentive firm responds less than one-for-one. The simplified definition of inattention is consistent with that in rational inattention models such as Sims (2003), which yields a steady-state Kalman gain between 0 and 1.[16]

***Definition* 2 (exposure).** *Firm* i *is more exposed to macroeconomic conditions than firm* j *if* $\pi_s^i(s,a) > \pi_s^j(s,a)$.

[15]Under this approximation, $\pi_a(s,a)$ drops out because of the first-order condition and assumption that $a^* = s$ at the optimum. Appendix B.4.1 contains detailed derivations of the approximation.

[16]In our illustrative framework, a firm's actions are a deterministic function of the aggregate state s, whereas in rational inattention models, there is noise in an agent's signals, which leads to both a Kalman gain between 0 and 1 and noise in the agent's actions conditional on the state.

Differences in attention and exposure We now derive model predictions for heterogeneity in attention and exposure that guide the empirical analysis to come.

We first construct the stock return, which is the dependent variable in our empirical analysis. As in Gorodnichenko and Weber (2016), a firm's stock price is equal to its firm value, which in the simple static setting equals its profits:

$$v = \pi(s, a).$$

Realized equity returns, measuring the log change in a firm's value around an aggregate shock, are given by

$$r = \hat{v} - \hat{v}_{-1}. \tag{2.6}$$

where $\hat{v} \equiv \log V - \log \bar{V}$ denotes the log deviation of firm value from the steady state and $\hat{v}_{-1} \equiv \log \mathbb{E}_{-1} V - \log \bar{V}$ denotes the log deviation of firm value before the shock is realized.

Proposition 1 highlights the asymmetric responses of stock returns to positive and negative aggregate shocks that result from the attention channel and the symmetric responses from the exposure channel.

***Proposition* 1.** *The return elasticity with respect to aggregate shocks for the exposure and the attention channels can be characterized as:*

(i) ***Exposure****: If firm i is more exposed to macroeconomic conditions than firm j, then, holding all else equal, the return elasticity of firm i with respect to the aggregate shock is higher than the return elasticity of firm j for all realizations of the shocks:*

$$\frac{\partial r_i}{\partial \hat{s}} > \frac{\partial r_j}{\partial \hat{s}} \quad \forall \hat{s}.$$

(ii) ***Attention****: Suppose firm i is attentive to macroeconomic conditions and firm j is inattentive. Then, holding all else equal, the return elasticity of a positive (expansionary) shock is higher for the attentive firm i than for the inattentive firm j. For negative (contractionary) shocks, the return elasticity for the attentive firm i is lower than for the inattentive firm j. For zero shocks, the return elasticities for attentive and inattentive firms equal*

$$\begin{cases} \frac{\partial r_i}{\partial \hat{s}} > \frac{\partial r_j}{\partial \hat{s}} & \textit{if } \hat{s} > 0 \\ \frac{\partial r_i}{\partial \hat{s}} = \frac{\partial r_j}{\partial \hat{s}} & \textit{if } \hat{s} = 0 \\ \frac{\partial r_i}{\partial \hat{s}} < \frac{\partial r_j}{\partial \hat{s}} & \textit{if } \hat{s} < 0 \end{cases}.$$

Figure 2.5: Model predictions for exposure vs. attention

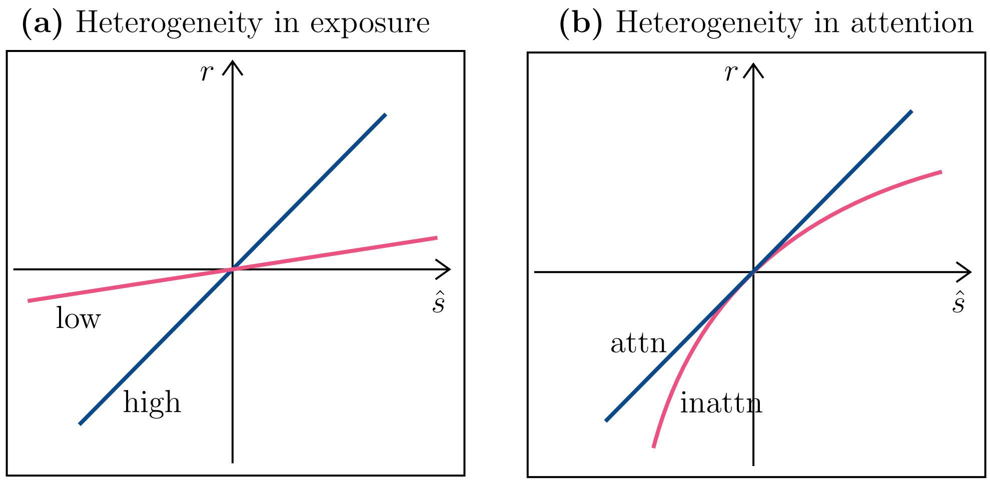

Notes: Illustration of model predictions of return elasticity with respect to aggregate shocks. Verticle axes represent conditional realized return, and horizontal axes represent the magnitude of shocks. The left panel shows return elasticity for firms that are highly exposed to macro conditions (*high*) and firms that are unexposed (*low*). The right panel shows return elasticity for attentive firms (*attn*) and inattentive firms (*inattn*). Exposure and attention are as defined in the main text.

Proof. See Appendix B.4.2. ∎

Figure 2.5 illustrates the predictions from Proposition 1. In Panel (a), firms are heterogeneous in their exposure to aggregate shocks, and those with high exposure exhibit higher return elasticities to aggregate shocks regardless of the sign of the shock. Panel (b) illustrates the mechanism of attention. Attentive firms are better at tracking the state variable, so their stock returns outperform those of inattentive firms after any aggregate disturbance. In response to a positive shock, stock returns of both attentive and inattentive firms rise, but returns of attentive firms rise more. In response to a negative shock, returns of both types of firms decrease, but returns of attentive firms drop by less.

This asymmetry in return elasticities is a unique feature of the attention channel and allows us to distinguish between the effects of firm attention and exposure to macro news. In the next section, we use this predicted asymmetry to show that our text-based measure correctly identifies firm attention and then estimate the cost of inattention based on the difference in return elasticities for positive and negative shocks.

2.4. Asymmetric Response to Monetary Shocks

We now test the hypothesis that attentive firms respond better to aggregate shocks using a high-frequency identification strategy. Shocks are constructed as plausibly exogenous monetary policy surprises following FOMC announcements, and resulting changes in firm value are measured using stock prices. We use our *prevalence* measure to estimate the relative performance of attentive firms and then test whether they fare better following both positive and negative shocks.[17] Results in this section serve the dual purpose of validating our text-based attention measure and quantifying the expected benefits of attention to economic conditions.

Stock prices are a particularly informative outcome variable because they are forward-looking and similarly high frequency as our monetary shocks. The cumulative effect of a rate surprise on expected future profits will be reflected quickly in a firm's stock price. By restricting to a narrow window around the shock, we isolate this price effect while avoiding other confounding factors. In comparison, a firm's investment and hiring decisions will be smoothed over a longer horizon and any low-frequency response is confounded by other factors that influence these choices. These limitations are exacerbated by the low statistical power of high-frequency monetary shocks, preventing precise estimates of investment and hiring responses.[18]

To best isolate the effects of attention, our baseline specification controls for firm size, age, leverage, and industry measured by 4-digit NAICS. The underlying identifying assumption is that firms have similar exposure to monetary policy shocks within a narrowly defined industry after conditioning on firm characteristics and financial structure. Residual variation in stock prices can then be attributed to firm attention rather than cross-firm variation in the exposure to monetary policy.

2.4.1. Data

Monetary policy shocks are constructed using the high-frequency identification strategy developed by Cook and Hahn (1989) and Gürkaynak et al. (2005) and used more recently in Gorodnichenko and Weber (2016), Nakamura and Steinsson (2018), and Ottonello and Winberry (2020). These shocks are measured as the change in the fed funds futures rate within a one-hour window surrounding FOMC announcements. Any changes within such a narrow window can be attributed to unanticipated changes to monetary policy as it is unlikely that other shocks occurred within the same window.

Monthly fed funds futures contracts clear at the average daily effective fed funds rate over the delivery month, so rate changes are weighted by the number of days in the month that

[17] This testable implication from Section 2.3 works for any aggregate shock with a related attention measure. We use high-frequency monetary shocks as "proof of concept" because they are familiar and well-identified. See Ramey (2016) for a comprehensive survey of alternative aggregate shocks.

[18] See Nakamura and Steinsson (2018) for further discussion of this "power problem."

are affected by the monetary policy shock. Following notation in Gorodnichenko and Weber (2016), the final shock series is defined as

$$\nu_t = \frac{D}{D-\tau}(ff^0_{t+\Delta t^+} - ff^0_{t-\Delta t^-}), \tag{2.7}$$

where t is the time of the FOMC announcement, $ff^0_{t+\Delta t^+}$ and $ff^0_{t-\Delta t^-}$ are the fed funds futures rates 15 minutes before and 45 minutes after the announcement, D is the number of days in the month of the announcement, and τ is the date of the announcement. We use the series published by Gorodnichenko and Weber (2016) and Nakamura and Steinsson (2018) for monetary shocks from 1994 to 2014. For easier interpretation of our empirical results, we normalize the sign of the monetary shock so that a positive shock is expansionary (corresponding to a decrease in interest rates).

Firm outcome and control variables are constructed using CRSP and Compustat data. Daily stock returns are measured as the open-to-close change in stock prices on the day of an FOMC announcement. Firm size, age, and industry controls are constructed as described in Section 2.2.3.

Firm attention is measured using the *prevalence* measure d_{it}, described in Section 2.2. To better suit a high-frequency methodology, firm attention at the time of an FOMC announcement is identified using the firm's most recent annual filing rather than the filing in the same year as the FOMC announcement. This modification precludes the possibility that firms are identified as attentive to the FOMC announcement that inspired their attention.

2.4.2. Methodology

We separately estimate the slope of the interaction between monetary shocks and firm attention for positive and negative shocks and then test whether these two coefficients are statistically different.

For firm i in industry j on day t, our baseline model takes the form

$$\begin{aligned} r_{it} = {} & \beta_d d_{it} + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+}\nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-}\nu_t \mathbb{1}_{\nu_t<0} \\ & + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \delta_j + \gamma_j \nu_t + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}, \end{aligned} \tag{2.8}$$

where d_{it} is the attention prevalence, ν_t is the monetary policy shock, $\mathbb{1}_{\nu_t>0}$ indicates positive monetary policy shocks, $\mathbb{1}_{\nu_t<0}$ indicates negative monetary policy shocks, and X_t is a vector of controls including the indicator variable for positive shocks and quarterly firm controls for size, age, and leverage. We also control for the interaction of monetary shocks with industry dummies and firm controls to capture the average effects of industry and firm characteristics on differential responses to monetary shocks. Standard errors are clustered by FOMC announcement to allow for correlated errors across firms at each FOMC announcement.

The coefficients of interest are $\beta_{d\nu_+}$ and $\beta_{d\nu_-}$. The theoretical framework in Section 2.3 hypothesizes $\beta_{d\nu_+}$ to be positive and $\beta_{d\nu_-}$ to be negative, implying attentive firms should outperform inattentive firms in response to both expansionary and contractionary monetary shocks. To formally test the hypothesis, we conduct a Wald test with the null hypothesis $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$.

2.4.3. Empirical results

Our baseline results are reported in Table 2.5. In the first column, we estimate the effect of high-frequency monetary shocks without our attention measures and find that a 25 basis point expansionary monetary shock is associated with about a 1% increase in stock prices. This result is consistent with existing estimates from Gorodnichenko and Weber (2016) and Nakamura and Steinsson (2018). The second column introduces the unconditional interaction between monetary shocks and firm attention. We find that attentive firms experience slightly higher stock returns than their inattentive counterparts, but our estimate is not statistically distinguishable from zero. This result is consistent with the framework outlined in Section 2.3, which remains agnostic as to the average interaction over the entire range of monetary shocks.

The main results from Equation (2.8) are presented in the third column. We test whether attention leads to differential responses to positive and negative monetary shocks. Consistent with predictions from rational inattention models, attentive firms appear to experience larger increases in stock returns following expansionary monetary shocks and smaller decreases in stock returns following contractionary monetary shocks. The coefficients are statistically different from zero, and the Wald test of whether these coefficients are equivalent is rejected at 5% significance. Column 4 shows that this result is not driven by outsized monetary surprises during the Great Recession nor unconventional monetary policy at the zero lower bound by ending the sample in 2007.

The *asymmetric* response to positive and negative shocks is inconsistent with alternative interpretations of the textual measure that predict a symmetric effect. The foremost alternative discussed in Section 2.3 is that the textual measure identifies exposure to monetary shocks rather than attention. Any such symmetric effect would also appear in the interaction coefficient $\beta_{d\nu}$ in Column 2, which is only weakly positive. Appendix B.2.2 further shows that directly estimating and controlling for exposure to monetary shocks leaves our main findings unchanged.

Suboptimal responses to monetary policy by inattentive firms reported in Table 2.5, to-

Table 2.5: Baseline results

		(1) Average	(2) Exposure	(3) Attention	(4) excl. ZLB
β_ν	Shock	5.61***	4.55*		
		(1.21)	(2.65)		
β_d	Attention		-0.01	-0.07	-0.03
			(0.05)	(0.06)	(0.06)
$\beta_{d\nu}$	Shock × Attn		1.07		
			(0.64)		
β_{ν_+}	Shock $\times \mathbb{1}_{\nu_t>0}$			4.93*	6.54**
				(2.74)	(2.75)
β_{ν_-}	Shock $\times \mathbb{1}_{\nu_t<0}$			-3.57	-0.95
				(3.72)	(3.69)
$\beta_{d\nu_+}$	Shock × Attn $\times \mathbb{1}_{\nu_t>0}$			2.02***	1.55**
				(0.72)	(0.72)
$\beta_{d\nu_-}$	Shock × Attn $\times \mathbb{1}_{\nu_t<0}$			-5.87*	-5.77*
				(3.18)	(3.30)
Observations		575667	575667	575667	432458
R^2		0.022	0.022	0.026	0.027
Clustered SE		yes	yes	yes	yes
Firm controls		yes	yes	yes	yes
4-digit NAICS FE		yes	yes	yes	yes
excl. ZLB		no	no	no	yes
Wald Test p-value				0.026	0.050

Notes: We have normalized the sign of the monetary shock ν_t so that a positive shock is expansionary (corresponding to a decrease in interest rates). Column (1) reports the average effect of monetary shocks from estimating $r_{it} = \delta_j + \beta_\nu \nu_t + \Gamma' X_t + \varepsilon_{it}$. Column (2) estimates the exposure model $r_{it} = \delta_j + \delta_j' \nu_t + \beta_\nu \nu_t + \beta_d d_{it} + \beta_{d\nu}(d_{it}\nu_t) + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$. Column (3) estimates the baseline attention model Equation (2.8): $r_{it} = \beta_d d_{it} + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \delta_j + \delta_j \nu_t + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$, where ν_t is the monetary shock, d_{it} is the prevalence attention measure, δ_j is an industry fixed effect, $\delta_j' \nu_t$ is its interaction with the shock, and X_t contains firm-level controls of size, age and leverage. The vector $X_t \nu_t$ contains the interactions between firm controls and the shock. Column (4) re-estimates Equation (2.8) on the sample ending in 2007 to exclude the zero lower bound period following the Great Recession. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

gether with the large fraction of inattentive firms documented in Figure 2.4, provide some of the first direct evidence of the empirical consequences of firm inattention in the US. We estimate that inattentive firm returns rise by 2% less following positive shocks and drop by 6% more following negative shocks compared to those of their attentive peers. These differences are substantial given the average stock return response of 5%.

Alternative sources of asymmetry We now consider alternative explanations for the asymmetric price response documented above. Each explanation is tested by augmenting our baseline model to include interaction terms for a confounding variable, c_{it}, that match those for firm attention, d_{it}. The resulting "horse-race" model takes the form

$$\begin{aligned} r_{it} = &\delta_j + \delta_j \nu_t + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} + \left[\beta_d d_{it} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0})\right] \\ &+ \left[\beta_c c_{it} + \beta_{c\nu_+}(c_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{c\nu_-}(c_{it}\nu_t \mathbb{1}_{\nu_t<0})\right] + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}, \end{aligned} \tag{2.9}$$

where, as in the baseline specification, we control for industry fixed effects, industry-specific responses to monetary shocks, a vector of firm controls and their interaction with monetary shocks. If the main result, $\beta_{d\nu_-} < 0 < \beta_{d\nu_+}$, holds true, then we rule out c_{it} as a confounding source of asymmetry.

The first factor considered is productivity. Van Nieuwerburgh and Veldkamp (2006) present a model in which higher productivity increases learning as well as production. If productivity determines both information acquisition and the response to aggregate shocks, it could explain the asymmetric result found above. Productivity is constructed as above in Section 2.2.3.

Management quality is another potential confounder that could explain both attention and firm performance. Effective managers who capitalize on opportunities during expansionary shocks and mitigate losses from contractionary shocks will generate the same asymmetric performance pattern documented in our main results. We approximate a firm's management quality using the share of board members who hold a graduate degree since existing research documents a strong relationship between education and management quality (Bloom and Van Reenen, 2010).[19] Data on the educational attainment is from BoardEx, which covers publicly traded US firms.

The third variable considered is a firm's financial performance measured using return on assets (ROA). Managers may feel compelled to cite macroeconomic conditions when explaining recent performance, and a tendency for well-performing firms to cite such conditions could generate the asymmetry observed.

Finally, we control for the length of a firm's SEC filing as a measure of its preference for information provision. Longer filings—measured using log word count—offer more opportunities for managers to mention macro keywords and signal commitment to due diligence. If thorough due diligence engenders investor confidence, then stocks should perform better following either positive or negative monetary shocks.

[19]Graduate degrees include MBA, MS, MSC, MA, JD, MD, MPA, MSE, PHD, and any degree names that include "master" or "doctor."

Table 2.6: Controlling for alternative explanations of asymmetry

Control Variable:	**Productivity (LTFP)**	**Mgmt Quality**	**Profit (ROA)**	**Filing Length**
Shock $\times \mathbb{1}_{v_t>0}$	6.88**	0.41	4.55*	-2.59
	(2.87)	(2.63)	(2.66)	(6.40)
Shock $\times \mathbb{1}_{v_t<0}$	-1.39	-10.53**	-4.04	21.86
	(4.04)	(4.37)	(3.92)	(14.24)
Attention	-0.13	-0.10*	-0.07	-0.07
	(0.10)	(0.06)	(0.06)	(0.05)
Attn $\times$ Shock $\times \mathbb{1}_{v_t>0}$	3.08**	2.59***	2.00***	1.81***
	(1.43)	(0.91)	(0.71)	(0.60)
Attn $\times$ Shock $\times \mathbb{1}_{v_t<0}$	-5.45***	-8.18**	-5.78*	-5.32*
	(1.78)	(3.34)	(3.16)	(2.95)
Control Var	0.04***	-0.07	0.04*	-0.01
	(0.01)	(0.06)	(0.02)	(0.04)
Control $\times$ Shock $\times \mathbb{1}_{v_t>0}$	0.10	1.53**	-0.97	0.80
	(0.14)	(0.73)	(1.67)	(0.49)
Control $\times$ Shock $\times \mathbb{1}_{v_t<0}$	-0.04	-2.54	-8.36**	-2.68*
	(0.22)	(2.58)	(3.93)	(1.59)
Observations	376644	324154	574804	575667
R^2	0.027	0.041	0.026	0.026
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
excl. ZLB	no	no	no	no
Wald test p-value: Attention	**0.001**	**0.003**	**0.027**	**0.031**
Wald test p-value: Control	0.387	0.143	0.126	0.041

Notes: This table augments Column (3) of Table 2.5 to control for four potential confounding sources of asymmetry. The estimated regression has the form $r_{it} = \delta_j + \delta_j \nu_t + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+}\nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-}\nu_t \mathbb{1}_{\nu_t<0} + \beta_d d_{it} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \beta_c c_{it} + \beta_{c\nu_+}(c_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{c\nu_-}(c_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$, where c_{it} represents the alternative "control" variable. As with attention, the control variable is interacted with both positive and negative monetary shocks. All other features of the model specification remain unchanged from Table 2.5. The four control variables considered are (1) firm productivity estimated as in Olley and Pakes (1996), (2) management quality approximated with board member educational attainment, (3) profit measured as earnings before extraordinary items over total assets, and (4) filing length measured as the log word count of the 10-K filing. The final two rows report p-values of Wald tests for $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$ and $H_0 : \beta_{c\nu_+} = \beta_{c\nu_-}$, respectively. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table 2.6 reports the estimates for Equation (2.9) using each factor described above: productivity, management quality, profitability, and filing length. As in our baseline results, attentive firms experience a larger increase in market value following an expansionary shock and a smaller contraction following a contractionary shock. The estimates are statistically significant, with similar magnitudes as those in Table 2.5. While some of the control variables (e.g., filing length) also display an asymmetric effect on firms' responses to monetary shocks, the explanatory power of firm attention remains under all specifications. All four Wald tests for $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$ are rejected at 5% significance. In Appendix Table B.3, we show that these results are also robust to excluding zero-lower-bound periods.

Additional robustness checks Further robustness analysis pertaining to the identification of high frequency monetary shocks can be found in Appendix B.2. Appendix B.2.3 controls for the information effect of FOMC announcements using Greenbook forecast revisions, and Appendix B.2.4 tests whether aggregate conditions confound the estimated effect of high frequency shocks. In each case, our main results remain robust.

2.5. Attention, Performance, and Aggregate Uncertainty

This section explores how attention affects firm performance under varying levels of aggregate uncertainty. One implication of our illustrative model is that the performance gap between attentive and inattentive firms widens with the magnitude of nominal demand shocks. Returns to attention should therefore increase in periods of greater uncertainty and larger shocks.[20]

We test this prediction by estimating the interaction effect between attention and uncertainty on firm performance. Aggregate uncertainty is measured using the interquartile range of quarterly forecasts for real GDP, inflation, and unemployment from the Survey of Professional Forecasters (SPF). Each series is standardized over our sample period (1994–2019) and then averaged into a composite uncertainty index.

Firm performance is measured along three dimensions: profitability, financial performance, and survival. Profitability is measured as a firm's return on assets (ROA), which we construct using earnings before extraordinary items over total assets. Financial performance is measured as return on equity (ROE) using earnings before extraordinary items over market capitalization. Finally, survival is defined as whether a firm remains in operation in the next year. Each variable is constructed using annual Compustat data, and ROA and ROE are winsorized at 1%.

[20]See Appendix B.4.3 for an extended illustrative framework that incorporates time-varying uncertainty.

Our regression model takes the form

$$y_{it} = \alpha_j + \beta d_{it} + \delta\sigma_t + \gamma d_{it} \cdot \sigma_t + \Gamma' Z_{it} + \varepsilon_{it}, \tag{2.10}$$

where y_{it} represents one of the three performance variables defined above, d_{it} is our binary attention measure, σ_t is aggregate uncertainty, α_j captures industry fixed effects with 4-digit NAICS, and Z_{it} is a vector of firm controls including size, age, and 10-K filing length (as previously defined). Standard errors are clustered by both year and industry. We extend the model to future outcomes, $y_{i,t+h}$, to capture any lagged effects of attention on performance.

Table 2.7: Effects of attention on firm performance under uncertainty

	ROE		**ROA**		**Survival**	
	Impact	Peak	Impact	Peak	Impact	Peak
Attention (general)	-0.02*	-0.02**	-0.02*	-0.03**	-0.01	-0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Uncertainty (SPF IQR)	-0.03***	-0.02*	-0.05***	-0.04*	-0.01	-0.03***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.01)
Attention × Uncertainty	0.03***	0.03***	0.05***	0.06***	0.01	0.03**
	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.01)
Observations	104507	92023	110267	97180	111637	66813
R^2	0.163	0.156	0.247	0.236	0.034	0.028
Clustered SE	yes	yes	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes	yes	yes

Notes: The table reports results from estimating (2.10), $y_{it}^h = \alpha_j + \beta d_{it} + \delta\sigma_t + \gamma d_{it} \cdot \sigma_t + \Gamma' Z_{it} + \varepsilon_{it}$, for horizons $h = 1, \cdots, 5$. The dependent variables y_t include (i) profitability measured with ROA (i.e., net income over total assets), (ii) financial performance measured with ROE (i.e., net income over equity), and (iii) an indicator variable for firm survival. Independent variables include the prevalence attention to general economic conditions, d_{it}; macroeconomic uncertainty, σ_t^2, measured as the interquartile range of quarterly growth rate forecasts for real GDP and unemployment from the SPF; the interaction between attention and uncertainty; industry fixed effects δ_j; and firm controls, Z_{it}. We standardize the interquartile range of each series over our observed sample period, take the absolute average deviation each quarter, and then average these quarterly values each year. The on-impact effect corresponds to the estimates for $h = 1$. The peak effect corresponds to the largest estimated marginal effect over the 5-year horizon. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Results from this analysis are reported in Table 2.7. On average, aggregate uncertainty reduces profitability, financial performance, and the probability of survival, which is consistent with existing models of uncertainty (e.g., Bloom et al., 2007).[21] Attention to macroeconomic conditions, however, mitigates the negative effects of uncertainty: in periods of high uncer-

[21] See Leahy and Whited (1996) for a general discussion of firm decisions under uncertainty.

tainty, attentive firms have higher profitability, better financial performance, and a higher probability of survival. Interestingly, the first row in Table 2.7 suggests that attention reduces firm performance under low uncertainty, consistent with models of imperfect information in which firms face a cost of attention and reap the benefit in states with large realized shocks (such as Reis, 2006).[22] Appendix Table B.7 further interacts attention with recession indicators and shows that attention improves firm performance mainly by reducing uncertainty.

Section 2.4 showed that attentive firms respond better to monetary shocks. This section finds that these same firms outperform less attentive competitors under elevated aggregate uncertainty. Together, they paint a picture of attentive firms as more responsive to evolving macroeconomic conditions and highlight the benefits gained for their diligence.

2.6. Quantitative Model

Our attention measure can inform model-based analysis in addition to the new empirical findings above. This section presents a quantitative model in which inattention to aggregate conditions drives monetary non-neutrality. Both the rate of attentive firms and the cost of inattention are calibrated using the prevalence measure presented in Section 2.2. This model demonstrates the importance of attention by showing that the efficacy of monetary policy depends on aggregate attention when firms face information frictions.

2.6.1. Model environment

We start with a canonical dynamic general-equilibrium model with rationally inattentive firms as in Maćkowiak and Wiederholt (2009) and Afrouzi and Yang (2021a). Time is discrete and infinite. The economy consists of a representative household, heterogeneous firms, and a central bank. Households and the central bank have full information about the economy, while firms pay a cost proportional to information obtained (measured using Shannon mutual information as in Sims, 2003). Firms differ ex-ante in their marginal costs of information, which is motivated by the heterogeneity documented in Section 2.2.3.

Household A representative household consumes a bundle of goods over the continuum of varieties $i \in [0, 1]$ and supplies labor, N_t, in a competitive labor market with wage, W_t. In

[22]Related to our findings that attentive firms appear to be "better opportunists" with state-dependent outperformance, Ahnert et al. (2021) find that banks with better information technology are more productive and spur better job creation and innovation. Furthermore, Kwon et al. (2022) find that industry concentration rises with investment intensity in information technology and research and development. Attentive firms with better information-processing technologies are better equipped to react to evolving macroeconomic conditions, which may have contributed to the rise of "superstar" firms (Autor et al., 2020).

addition to the wage income, the household has access to a one-period bond, D_t, with the interest rate ι_t and receives firms' profits, Π_t. The household maximizes its life-time utility:

$$\max_{\{C_{it},D_t,N_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\log C_t - \psi N_t), \tag{2.11}$$

$$\text{s.t.} \int_0^1 P_{it} C_{it} di + D_t \leq W_t N_t + (1+\iota_t) D_{t-1} + \Pi_t,$$

where consumption, C_t, is aggregated over each good type i with a CES aggregator, $C_t = \left(\int_0^1 C_{it}^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}}$, and ε is the elasticity of substitution. Let $Q_t \equiv P_t C_t$ denote nominal aggregate demand. The household's optimal choices are given by the following three conditions

$$C_{it} = C_t (P_{it}/P_t)^{-\varepsilon}; \qquad 1 = \beta(1+\iota_t)\,\mathbb{E}_t(Q_t/Q_{t+1}); \qquad W_t = \psi Q_t. \tag{2.12}$$

Central bank The central bank targets the aggregate money supply, $P_t C_t$, similar to Caplin and Spulber (1987) and Gertler and Leahy (2008). Nominal aggregate demand, therefore, follows

$$\Delta \log Q_t = \rho \Delta \log Q_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2), \tag{2.13}$$

Firms There is a unit measure of monopolistically competitive firms, indexed by $i \in [0,1]$. Firms operate a decreasing-returns-to-scale production technology with labor as its only input: $Y_{it} = N_{it}^{\gamma}$. They take wage and demand as given and can flexibly set prices, P_{it}, based on their information set in period t. After setting prices, they hire labor from a competitive labor market to produce the realized level of demand induced by their prices.

Firms are assumed to be *rationally inattentive*, meaning that they do not observe shocks to aggregate demand and endogenously acquire information about Q_t. In each period, firm i starts with their information set from the previous period, S_i^{t-1}, and selects the stochastic process for their new signal, s_{it}, from a set of available signals, $\mathcal{S}^t$, that vary in cost and precision. These signals satisfy the properties outlined in Definition 3.

***Definition* 3 (set of available signals).** *The set of available signals, $\mathcal{S}^t$, consists of all signal processes satisfying the following three properties:*

i. $\mathcal{S}^t$ is rich: for any posterior distribution on $\{Q_t\}_{t\geq 0}$, there is a set of signals $S^t \in \mathcal{S}^t$ that generate that posterior;

ii. Signals do not expire over time: $\mathcal{S}^t \subset \mathcal{S}^{t+h}$ for $h \geq 0$;

iii. Signals contain no information about future shocks: $S_t \perp Q_{t+h}$ *for* $S_t \in \mathcal{S}^t$ *and* $h \geq 1$.

We assume that the cost of information is linear in the Shannon mutual information:

$$2\omega_i \cdot \mathcal{I}(Q_t; s_{it}|S^{t-1}), \tag{2.14}$$

where the Shannon mutual information, $\mathcal{I}(Q_t; s_{it}|S^{t-1})$, measures the expected reduction in uncertainty about aggregate demand from observing the signal.[23] A more precise signal requires a higher flow of mutual information and is therefore more expensive.

The information represented by $\mathcal{I}(\cdot)$ can be thought of as a firm's *attention* to the economy: for each unit of mutual information (or nat), a firm pays a marginal cost, ω_i, and reduces its expected uncertainty about aggregate demand. Since firms could theoretically profit from "forgetting" information acquired in the past, we impose a no-forgetting constraint. Therefore, a firm's information set evolves according to $S_i^t = S_i^{t-1} \cup s_{it}$.

Firms are ex-ante heterogeneous in their information-processing technology and face either high or low marginal costs of attention

$$\omega_i \in \{\omega_H, \omega_L\}. \tag{2.15}$$

A fraction $\theta \in (0, 1)$ of firms are assumed to have low information-processing costs, while all remaining firms face high costs.

Firms maximize expected profits by choosing the stochastic process of the set of signals to observe over time, $\{s_{it} \in \mathcal{S}_{it}\}_{t\geq 0}$, and a pricing strategy, $P_{it}(S_i^t)$, that depends on its information set at time t containing realizations of current and past signals. The firm's problem is given by

$$\max_{\{s_{it}\in\mathcal{S}_{it}, P_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}\Big[\sum_{t=0}^{\infty} \beta^t \frac{1}{P_t C_t}\Big(\underbrace{(P_{it}Y_{it} - W_t N_{it})}_{\text{operational profits}} - \underbrace{2\omega_i \mathcal{I}(Q_t; s_{it}|S_i^{t-1})}_{\text{information costs}}\Big) \Big| S_i^{-1}\Big] \tag{2.16}$$

$$\begin{aligned} \text{s.t.}\quad & Y_{it} = Y_t (P_{it}/P_t)^{-\varepsilon} && \text{(demand for goods)} \\ & Y_{it} = N_{it}^{\gamma} && \text{(production technology)} \\ & S_i^t = S_i^{t-1} \cup s_{it}, && \text{(evolution of information)} \end{aligned}$$

where Y_t is the aggregate output, P_t is the aggregate price index, P_{it} is firm i's price, Y_{it} is the demand for the firm's goods, N_{it} is the firm's labor demand.

[23] Formally, the Shannon mutual information between random variables X and Y is defined as $\mathcal{I}(X;Y) = \int_Y \int_X p(x,y) \log \frac{p(x,y)}{p(x)p(y)} dxdy$, which measures the difference between conditional and marginal entropies. See Cover and Thomas (2006) for details.

Equilibrium Given the exogenous process for aggregate demand, $\{\Delta \log Q_t\}_{t\geq 0}$, the equilibrium consists of an allocation for the household, $\Omega^H = \{C_t, D_t, N_t, (C_{it})_{i\in[0,1]}\}_{t\geq 0}$, allocations for every firm $i \in [0,1]$ given their initial information sets S_i^{-1}, $\Omega_i^F = \{s_{it} \in \mathcal{S}_{it}, P_{it}, N_{it}, Y_{it}\}_{t\geq 0}$, a set of prices $\{\iota_t, P_t, W_t\}_{t\geq 0}$, and a stationary distribution over firms' states such that

i. Given the set of prices and firms' allocations, the household's allocation solves the problem in Equation (2.11);

ii. Given the set of prices and the household's allocation, firms' allocations solve the problem in Equation (2.16);

iii. All markets clear, that is, for $t \geq 0$ and $i \in [0,1]$, $D_t = 0$, $Y_{it} = C_{it}$, $Y_t = C_t$, and $N_t = \sum_i N_{it}$.

Solution We approximate a firm's flow profits with second-order log approximations around the full-information steady state.[24] A firm's total value under log approximation, v, is decomposed into a full-information value, v^*, representing the firm's value under optimal pricing with full information, and the imperfect information value, $\tilde{v}$, representing firm value under imperfect information.

Let lowercase letters denote log deviations from the steady state. The imperfect-information value is given by

$$\begin{aligned}
\tilde{v} = \max_{\{s_{it}\in\mathcal{S}_{it}, p_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \Big(- B(p_{it} - p_t^*)^2 - 2\omega_i \mathcal{I}(q_t; s_{it}|S_i^{t-1})\Big) \Big| S_i^{-1}\right] \qquad (2.17)\\
\text{s.t.} \quad p_t^* = \alpha p_t + (1-\alpha) q_t \\
S_i^t = S_i^{t-1} \cup s_{it},
\end{aligned}$$

where $\alpha \in (0,1)$ and $B > 0$ are constants that depend on non-information-friction parameters and relate to the degree of strategic complementarity and the curvature of the profit function, respectively.[25] p_t^* denotes the optimal price under perfect information. Since prices are fully flexible, price setting with perfect information is a static problem (see Appendix B.5.2).

[24]Appendix B.5.1 contains details of the approximation. Log-quadratic approximation is a common simplifying assumption in rational inattention models (see, e.g., Maćkowiak and Wiederholt, 2009; Afrouzi and Yang, 2021a) to address the curse of dimensionality that arises from firms having the joint distribution of prices and nominal aggregate demand as the state variable. Sims (2003) shows that the optimal distribution under Gaussian priors and quadratic payoffs is also Gaussian, so log-quadratic approximation of the profit function greatly reduces the dimensionality of the problem.

[25]See Equations (B.7) and (B.8) in the appendix.

The imperfect information problem in (2.17) is solved numerically based on the algorithm for dynamic rational inattention problems (`DRIPs`) developed in Afrouzi and Yang (2021a). Appendix B.5.3 provides detailed information on its implementation.

2.6.2. Calibration

Model parameters are divided into two sets: those that govern information frictions and all remaining parameters. In the first set, the share of attentive firms and relative cost of information between firms are calibrated to match two empirical moments using our text-based measure of attention. The cost of information among attentive firms, ω_L, is set near zero so that attentive firms have nearly full information. The second set of non-information parameters are calibrated to external sources or estimates using quarterly data on US output from the Bureau of Economic Analysis. A summary of all model parameters can be found in Appendix Table B.9.

For non-information parameters, we calibrate the model quarterly and set the discount rate to be $\beta = 0.96^{1/4}$. The stochastic process for aggregate demand, $\{\rho, \sigma_\nu\}$, is estimated using quarterly US nominal manufacturing output between 1994 and 2019. Restricting to the manufacturing sector is consistent with the within-sector results presented in our empirical analysis. The elasticity of substitution is set to $\varepsilon = 10$, implying a steady-state markup of 11%, and the disutility of labor is set to $\psi = 0.90$ to offset the steady-state distortions from monopolistic competition. Finally, we set returns to scale $\gamma = 0.93$ according to the estimate by Basu and Fernald (1997) for the US manufacturing sector.

For information parameters, the share of firms with low information costs, θ, is set to 65% to match the share of attentive firms in Figure 2.3. As in Maćkowiak et al. (2009) and Afrouzi and Yang (2021b), attention depends inversely on the ratio between attention costs and the curvature of the profit function, ω_i/B. A firm pays greater attention when the information cost is low or when its incentives to pay attention are high. We focus on calibrating information parameters and therefore fix the curvature of profit function, which depends on non-information parameters.

We set ω_L close to zero to reflect the assumption that firms with low information costs have nearly full information. The relative cost of information for high-cost firms, $\omega_H - \omega_L$, is calibrated to match the heterogeneous responses to monetary shocks estimated in Table 2.5. Stock returns in the model are defined as the log change in a firm's value, $r_{it} = \log V_{it} -$

$\log \mathbb{E}_{t-1}(V_{it})$.[26] To connect Shannon mutual information, $\mathcal{I}_{it}$, in the model with the text-based attention measure, we assume that the frequency of macro keywords in 10-K filings is strictly increasing in firm attention. This allows us to match the cross-sectional distribution of firm attention without explicitly modelling the writing process of 10-K filings. Since our main empirical analysis uses the prevalence attention measure, we define a corresponding indicator variable, $d_{it} = \mathbb{1}(\mathcal{I}_{it} > \bar{\mathcal{I}}_t)$, for firms whose attention is above the cross-sectional mean in a given period. Finally, we use ν_t as the monetary shocks.

We simulate the model for a panel of 100 firms and for 1000 quarters, discarding the first 100 quarters as burn-in. With the simulated data, we estimate

$$r_{it} = c + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+}\nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-}\nu_t \mathbb{1}_{\nu_t<0} + \beta_d d_{it} + \boldsymbol{\beta_{d\nu_+}} d_{it}\nu_{it}\mathbb{1}_{\nu_t>0} + \boldsymbol{\beta_{d\nu_-}} d_{it}\nu_{it}\mathbb{1}_{\nu_t<0} + \varepsilon_{it}.$$

We set $\omega_H - \omega_L$ to target the elasticity $\frac{1}{2}|\hat{\beta}_{d\nu_+}| + \frac{1}{2}|\hat{\beta}_{d\nu_-}|$ from Column 3 in Table 2.5, which measures the relative stock return losses of firms that do not pay attention. Appendix Figure B.7a shows how the parameter is identified. As ω_H increases and the gap between ω_L and ω_H widens, the simulated elasticity monotonically increases, implying greater heterogeneity between attentive and inattentive firms. The resulting calibration for $\omega_H - \omega_L$ is 1.11 per nat.

Discussion of the calibration strategy One primary challenge to calibrating a rational inattention model is that information costs—which determine the degree of information frictions—are unobserved in the data. Existing studies have successfully calibrated rational inattention parameters by matching moments related to aggregate consumption dynamics and monetary policy responses, while other have calibrated these parameters using survey data (Luo, 2008; Maćkowiak and Wiederholt, 2015, 2023). Our calibration strategy differs in three main ways. First, we allow heterogeneity in information costs and therefore have two parameters for information costs, ω_L and ω_H, instead of a single parameter. We focus on calibrating the heterogeneity in attention costs, $\omega_H - \omega_L$, to study its implication for monetary transmission. In doing so, we set ω_L close to zero, which implies that our calibrated model provides a lower bound on the degree of monetary non-neutrality arising from inattention.

Second, our calibration makes use of the text-based attention measure instead of macro

[26] A firm's value function in (2.16) can be expressed in recursive form as

$$V(S_i^{t-1}) = \max_{\{s_{it}\in\mathcal{S}_{it}, P_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}_t \left[\frac{1}{P_t}\Big(\big(P_{it}Y_{it} - W_t N_{it}\big) - 2\omega_i \mathcal{I}(Q_t; s_{it}|S_i^{t-1})\Big) + \beta\Lambda_{t,t+1} V(S_i^t) \,\Big|\, S_i^{t-1} \right]$$

$$\text{s.t.} \quad Y_{it} = Y_t(P_{it}/P_t)^{-\varepsilon}, \quad Y_{it} = N_{it}^{\gamma}, \quad S_i^t = S_i^{t-1} \cup s_{it}.$$

data or survey data. The measure informs attention at the granular firm level, but the tradeoff is that 10-K filings do not have a direct model counterpart. To connect the concept of Shannon mutual information with our text-based attention measure, we need to assume that the frequency of macro keywords in 10-K filings is strictly increasing in firm attention. This allows us to use the textual measure to discipline the cross-sectional distribution of firm attention.

Lastly, the relative cost of attention, $\omega_H - \omega_L$, is calibrated by targeting a micro elasticity (namely, the relative stock return losses of inattentive firms in response to monetary shocks) rather than macro moments. This is possible because our proposed attention measure is available for a large number of firms over a long sample period. It is well-known since Mehra and Prescott (1985) that standard macro models, including ours, are not designed to match the unconditional cross section of stock returns. However, our target moment in Table 2.5 is the *conditional* responses of firm values to monetary shocks, with stock returns, r_{it}, in Equation (2.8) capturing log changes in a firm's value. Appendix Table B.10 shows that our model matches heterogenous responses of firms' values to monetary shocks through relative information costs.

For further robustness, Appendix B.5.5 implements an alternative calibration strategy that targets industry-level price adjustment estimates from Figure 2.2. It finds that attention remains quantitatively important for the transmission of monetary policy.

2.6.3. Attention and the efficacy of monetary policy

Figure 2.6 plots the aggregate responses to a one standard deviation expansionary shock to nominal aggregate demand growth. Panel (a) shows that inattentive firms under-adjust prices, reflecting partial incorporation of noisy signals about demand. Attentive firms track aggregate demand better than inattentive firms and exhibit more responsive prices.[27]

Panel (b) shows that inattentive firms are responsible for increased output following an expansionary shock. These firms mischaracterize the nominal demand shock as a real shock and respond by raising output, while attentive firms correctly identify the nominal nature of the shock and respond by raising prices.

The grey solid lines represent aggregate inflation and output responses, driven by attention costs and the share of attentive firms.[28] Monetary non-neutrality increases with the degree

[27]Appendix Figure B.8 shows individual firms' impulse responses for prices, profits, attention, and stock returns (including full-information returns, imperfect-information returns, and total returns) in response to both expansionary and contractionary monetary shocks.

[28]To compare the aggregate responses with standard benchmarks, we convert the nominal aggregate demand shock to the nominal interest rate shock used in Christiano et al. (2005) by estimating the passthrough of the interest rate on the nominal aggregate demand in Appendix B.5.7. The right scale of Appendix Figure

Figure 2.6: Aggregate responses to expansionary monetary shock

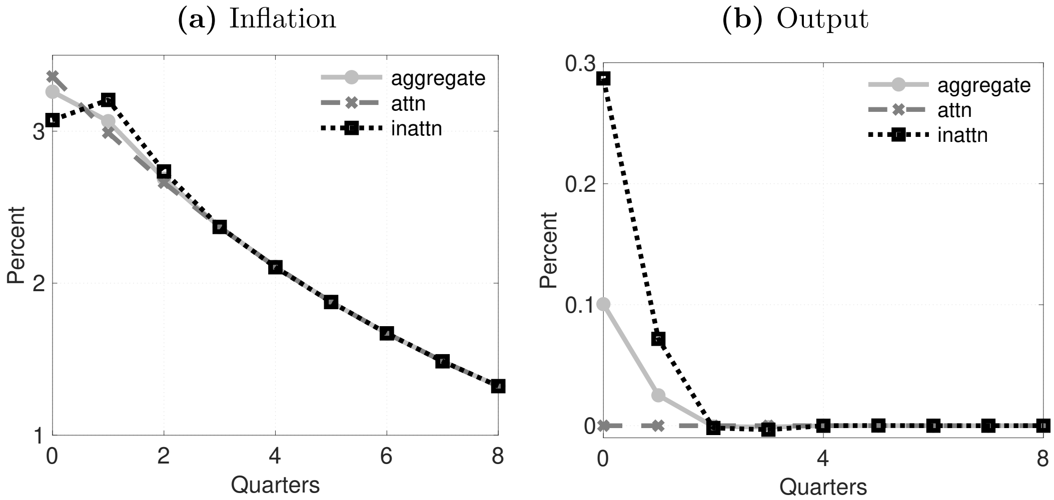

Notes: The figures report impulse responses in percent deviations from the perfect-information steady state of inflation and output for the aggregate economy, attentive firms, and inattentive firms.

of inattention in the economy. Since we assume attentive firms face near-zero attention costs ($\omega_L \approx 0$), the impulse responses in Figure 2.6 provide a lower bound on the output responses to monetary shocks and an upper bound on the inflation responses.

A key implication of Panel (b) is that the aggregate output response to monetary policy increases with the share of inattentive firms. To illustrate the quantitative scope of the effect, we exogenously vary the fraction of attentive firms and compare output responses in our baseline calibration against two alternatives, $\check{\theta} = 56\%$ and $\hat{\theta} = 73\%$, which correspond to the minimum and maximum fraction of attentive firms over the sample period.

Table 2.8 reports the aggregate responses to monetary policy change as the fraction of attentive firms in the economy changes. The response of output growth to monetary policy is 5 basis points (or 42%) weaker in the most attentive calibration compared to the least attentive calibration. This suggests that expansionary policy in the depth of a recession when more firms are paying attention will be weaker than a preemptive interest rate (i.e., leaning against the wind) when aggregate attention is lower. This pattern is consistent with existing studies on the state dependency of monetary policy (e.g., Tenreyro and Thwaites, 2016). Similarly, monetary tightening imposes a smaller contractionary effect on output when more firms are attentive to monetary news, which highlights the importance of clear central-bank communication (highlighted, e.g., by Haldane et al., 2021).

B.10 show that in response to a 25 basis point interest rate cut, output increases by 0.1% on impact, in line with the impact responses of 0.1% in Christiano et al. (2005) and smaller than the peak responses of 0.5%.

Table 2.8: Attention and monetary non-neutrality

	Least attentive	Baseline	Most attentive
Fraction of attentive firms (θ)	56%	65%	73%
Output response	0.12%	0.09%	0.07%

Notes: Dependence of output responses on the fraction of attentive firms in the economy. Output responses are calculated as percent deviations from the steady state in response to a 25 basis point rate cut. Calibration for the least and most attentive economy is described in the main text.

2.7. Conclusion

The empirical evidence of information frictions that we document in this paper, along with growing evidence in the literature (Candia et al., 2021), highlights firms' deviations from full-information rational expectations (FIRE). To discipline models without FIRE, researchers require an understanding of firms' information sets and expectation-formation processes.

In that direction, this paper presents a new text-based measure of firm attention to macroeconomic news, which will be made available publicly and updated on an ongoing basis. We validate that the measure indeed captures firm attention by testing for an asymmetric prediction of rational inattention on monetary policy transmission. We show that firms that pay attention to the FOMC have larger increases in stock returns after positive monetary shocks and smaller decreases in stock returns after negative monetary shocks, providing direct empirical evidence for the consequences of firm inattention.

The empirical measure can be used in combination with imperfect-information models to ground those theories in data. We demonstrate the value of this measure in a quantitative rational inattention model by showing that time variation in firm attention has important implications for the state dependency of monetary policy. In the model, average inattention drives the degree of monetary non-neutrality. The countercyclical nature of firm attention to macroeconomic news implies that the efficacy of monetary policy is weaker during recessions and should be considered in policy design.

CHAPTER 3

Higher Education and Labor Market Adjustment following the Great Recession

3.1. Introduction

Slow and incomplete labor market recoveries following the U.S. recessions in 1990, 2001, and 2007 acutely impacted workers without college degrees. National enrollment data suggests that a substantial portion of these workers returned to school: as the U.S. labor market lost 8.1 million jobs between fall 2007 and 2009, one million more adults ages 22 or older enrolled in post-secondary programs. Among adults ages 25-34, one new student enrolled for every 3.3 jobs lost.[1] Despite such an impressive retraining effort, we lack a complete picture of how new enrollment affected the labor market's recovery. This paper offers new detail on countercyclical enrollment over three decades and assesses how access to higher education during the Great Recession altered the trajectory of local employment rates in the years to follow.

The paper begins by documenting the historical magnitude and distribution of countercyclical enrollment in higher education. In a panel of 380 Metropolitan Statistical Areas (MSAs) between 1987 and 2019, a one percentage point decrease in the employment-population ratio was associated with a 0.12 percentage point rise in annual enrollment on average. This rate reaches 0.30 percentage points at the 90th percentile of MSAs when estimated for each MSA individually. Data on student demographics and institution characteristics reveal that countercyclical enrollment occurred almost entirely at undergraduate programs, was highest among older students, and was concentrated at public institutions, which are typically more accessible than private or for-profit alternatives. These findings offer insight into where the benefits of retraining may have resided over the past few decades.

[1] Employment estimates are from the Quarterly Workforce Indicators and enrollment estimates are from the Integrated Postsecondary Education Data System.

The next section presents an empirical strategy for estimating the causal effect of higher education access on the pace of employment recovery following the Great Recession. Since areas hit hardest by labor demand shocks experienced both higher enrollment and worse employment outcomes, this effect cannot be estimated directly. Instead, my approach isolates plausibly exogenous changes to enrollment growth using a Bartik instrument and predicts changes in the local enrollment rate based on initial enrollment levels. A remaining empirical challenge is that local enrollment rates are associated with greater educational attainment and thus affect employment outcomes. This second issue is addressed by directly conditioning on workers' educational attainment and initial enrollment rates.

Resulting estimates show that MSAs with greater access to higher education experienced persistently higher employment rates in the years following the Great Recession. On average, a one percentage point increase in the enrollment rate between 2007 and 2010 led to a 2.0 percentage point larger increase in the employment rate between 2007 and 2018. The impact of new enrollment was highest among younger workers ages 25-34 and remained substantial for workers ages 35-44. The youngest cohort, ages 19-24, experienced only weak and brief gains in the years immediately following the Great Recession, while older workers ages 55-64 showed little benefit over the entire sample period.

The main findings are largely consistent with the timing and age concentration of new enrollment during the Great Recession. Estimated employment effects appear only after an enrollment surge in 2009 and are strongest in the three years following that surge. Each age cohort's employment effect is largely consistent with its share of new enrollment, particularly among older workers who had little enrollment and the smallest employment effects. Although these findings cannot prove that my identification strategy holds, they offer some assurance that the results are not driven by selection effects that would appear from the onset of the sample period or in age groups without much new enrollment.

Related Literature These findings build upon recent evidence that retraining improves employment outcomes among displaced workers.[2] Card et al. (2018) conduct a meta analysis of 207 international active labor market programs and find that training programs improve employment of participants by 6.6 percentage points when measured at least one year after completion. Hyman (2018) uses quasi-experimental variation in Trade Adjustment Assistance (TAA) approval to estimate that participants worked an average of two more months per year and cumulatively earned $50,000 more than denied applicants. Earlier work on retraining through higher education found that displaced workers in Washington who enrolled

[2]See Heckman et al. (1999) for a comprehensive review of jobs training program evaluations in the 20th century and econometric methods using non-experimental data that emerged from this literature.

in community college saw their total hours and earnings increase by about six percent for each year completed (Jacobson et al., 2005).

This paper also extends research on how labor market conditions affect enrollment rates by revisiting previous findings using administrative data and examining the consequences of countercyclical enrollment. One area of related research documents countercyclical enrollment in higher education over the business cycle using survey data (Barr and Turner, 2013; Betts and McFarland, 1995; Dellas and Sakellaris, 2003). A second area of work uses positive labor demand shocks from natural resource extraction and the housing bubble to show that increased employment opportunity reduces educational attainment (Black et al., 2005; Cascio and Narayan, 2015; Charles et al., 2018). Complementary work on negative demand shocks finds that mass layoffs have only a minor effect on local enrollment (Hubbard, 2018; Acton, 2019; Foote and Grosz, 2019).

High enrollment rates and subsequent employment gains found in this paper are relevant to understanding how labor markets recovered following the Great Recession. One puzzle about the recovery is why so few workers sought employment opportunities in other cities. Earlier research shows that directed migration was key to the dynamism and success of the U.S. labor market in much of the 20th century (Freeman, 2007). Blanchard and Katz (1992) find that geographic mobility was the primary mechanism by which U.S. regions adjusted to differences in labor demand between the 1950 and 1990. More recent work documents a decline in overall geographic mobility since the 1980s and a diminished role of directed migration for reducing regional differences in economic growth (Molloy et al., 2011; Ganong and Shoag, 2017; Bartik, 2017). Yagan (2014) estimates that directed migration only insured workers against 7% of their local demand shock. Based on the results presented below, higher education may have provided displaced workers with a more convenient and effective means of finding new employment than migration in the wake of skill-specific shifts in labor demand.

3.2. Enrollment over the Business Cycle, 1987-2019

This section uses a panel of 380 U.S. Metropolitan Statistical Areas (MSAs) from 1987 to 2019 to estimate business cycle variation in higher education enrollment. The analysis is disaggregated by institution and student characteristics to identify schools and student populations whose enrollment is most sensitive to labor market conditions. The central finding is that most countercyclical enrollment occurs among older students, those enrolling full-time in undergraduate programs, and at public colleges and universities.

3.2.1. Data and Methods

My approach to estimating countercyclical enrollment departs from existing literature in two ways: first, enrollment and employment are measured using administrative data rather than survey data, which allows for more granular geographic units of observation that are more representative of local labor markets.[3] Second, labor market conditions are measured using the employment-population ratio rather than the unemployment rate, which avoids any confounding effects from labor force participation and allows for a more direct comparison to enrollment rates.

Enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), which cover all institutions that participate in federal student financial aid programs such as Pell Grants.[4] IPEDS collects detailed information on school characteristics, enrollment, admissions, campuses finance, and student outcomes.

Annual enrollment for each MSA is defined as the total fall headcount of students attending schools located within that MSA. Since IPEDS data is consolidated by institution, it does not distinguish between campuses that span multiple MSAs. Though uncommon, these cases are addressed using campus-specific enrollment shares from the Department of Education's College Scorecard database. Schools for which a majority of students attend remotely are excluded from the sample since we are interested in the local relationship between employment and enrollment.

The enrollment rate is constructed as the ratio of enrollment to the adult working-age population (ages 18-65) from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER). Separate rates are constructed by level of institution (four-year university, two-year college, or less-than-two-year technical college), control of institution (public, private non-profit, or private for-profit), level of student (undergraduate vs graduate), attendance status (full-time vs part-time), and student demographics (gender and age).

Employment is measured as total private employment from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW). The QCEW program covers 95% of U.S. jobs by combining state unemployment insurance administrative records and the Unemployment Compensation for Federal Employees program. Workers are geographically sorted according to their place of work. As with enrollment, the employment rate is constructed as the ratio of employment to the adult working-age population from SEER.

[3] A notable exception is Hillman and Orians (2013), which uses administrative data to estimate countercyclical enrollment at community colleges across Core Based Statistical Areas (CBSAs).

[4] IPEDS surveys were not mandatory for less-than-two-year institutions prior to 1993, which are excluded from my sample prior to that year. See https://nces.ed.gov/statprog/handbook/pdf/ipeds.pdf for more information.

The sample period begins in 1987 after IPEDS replaced its more limited predecessor and ends in 2019 before the COVID-19 pandemic.[5] The sample period covers three recessions that were followed by weak employment recoveries and is well-suited to capture business cycle variation in college enrollment. To account for enrollment trends over the sample period, each series is detrended using an HP-filter with smoothing factor of 100.[6]

Table 3.1: Employment and Enrollment Rates by Institution and Student Characteristics (%)

	N	Mean	SD	10th	25th	Median	75th	90th
Employment Rate	11,459	44.5	8.8	33.4	38.6	44.3	50.2	55.7
Total Enrollment Rate	11,459	9.2	6.7	3.7	5.3	7.2	10.2	17.3
Institution Level								
Four-year	11,459	6.0	6.9	0.0	2.0	4.1	6.8	14.3
Two-year	11,459	3.0	2.2	0.1	1.5	2.9	4.2	5.7
One-year	9,715	0.1	0.2	0.0	0.0	0.1	0.2	0.3
Institution Control								
Public	11,459	7.6	6.7	2.6	3.9	5.5	8.6	15.6
Private Non-Profit	11,459	1.3	2.1	0.0	0.0	0.7	1.9	3.3
Private For-Profit	11,459	0.2	0.3	0.0	0.0	0.1	0.3	0.6
Student Level and Attendance								
Undergraduate	11,459	8.1	5.5	3.6	4.9	6.5	9.1	15.1
Graduate	11,459	1.1	1.5	0.0	0.1	0.6	1.3	2.5
Full-time	11,459	6.0	5.7	1.8	2.8	4.1	6.5	12.5
Part-time	11,459	3.2	1.8	1.4	2.1	2.9	4.0	5.3
Student Demographics								
Men	11,459	4.0	3.4	1.5	2.1	3.0	4.4	7.9
Women	11,459	5.1	3.4	2.2	3.1	4.2	5.8	9.4
Age <25	8,984	6.2	5.3	2.2	3.2	4.5	6.7	12.4
Age 25+	8,984	3.0	1.8	1.2	1.9	2.6	3.6	5.2

Notes: This table summarizes the annual employment rate, total enrollment rate, and disaggregated enrollment rates for 380 MSAs between 1987 and 2019. All statistics are expressed in percent. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using adult working-age population (ages 18-65) from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER). Institutions that exclusively offer graduate programs are counted among four-year universities. Enrollment at one-year institutions are not available before 1993, and enrollment by student age is only available on a biannual basis between 1991 and 1998.

The relationship between local enrollment and employment rates is estimated with a univariate model,

$$s_{lt} = \alpha + \beta e_{lt} + \varepsilon_{lt}, \tag{3.1}$$

where s_{lt} is the fall enrollment rate for MSA l in year t and e_{lt} is the employment rate. Since both enrollment and employment are detrended by MSA, the specification relies ex-

[5]IPEDS began its survey in 1986 but this year is omitted due to data discrepancies with subsequent years.

[6]Detrended results are robust to using a linear time trend, quadratic time trend, or HP filter with smoothing factor of 6.25.

clusively on within-MSA variation. The model is estimated using OLS and standard errors are clustered by both MSA and year.

3.2.2. Results

Table 3.2: Enrollment Responsiveness to Local Employment Rate
(a) Enrollment by Institution Characteristics

		Institution Level			Institution Control		
	Total	Four-year	Two-year	One-year	Public	Non-profit	For-profit
Employment rate	-11.9	-3.8	-7.4	-0.7	-9.2	-0.9	-1.7
	(1.9)	(0.7)	(1.4)	(0.1)	(1.6)	(0.2)	(0.4)
Observations	11,459	11,459	11,459	9,715	11,459	11,459	11,459
R^2	0.09	0.02	0.07	0.02	0.06	0.01	0.06

(b) Enrollment by Student Characteristics

	Enrollment Level		Attendance Status		Demographics			
	Undergrad	Grad	Full-time	Part-time	Men	Women	Age <25	Age 25+
Employment rate	-11.3	-0.7	-8.8	-3.1	-5.1	-6.7	-4.7	-7.7
	(1.8)	(0.2)	(1.3)	(0.8)	(0.9)	(1.1)	(0.7)	(1.3)
Observations	11,459	11,459	11,459	11,459	11,459	11,459	8,984	8,984
R^2	0.09	0.01	0.10	0.01	0.07	0.08	0.03	0.10

Notes: This table reports the results of several univariate regressions of annual MSA enrollment rates (in basis points) on the local employment rate (in percent). Coefficients should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. Standard errors are in parenthesis. All rates are detrended using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. *Panel A* reports estimates for total enrollment, enrollment by institution program duration, and enrollment by institution control. Schools that exclusively offer graduate programs are counted among four-year programs. *Panel B* reports estimates for enrollment by student program level, attendance status, and demographics. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using adult working-age population (ages 18-65) from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER). Enrollment at one-year institutions are not available before 1993, and enrollment by student age is only available on a biannual basis between 1991 and 1998.

Table 3.2 presents the results from Equation 3.1. For ease of interpretation, enrollment responsiveness is expressed in basis points (bps) per percentage point change in employment. The first column in Panel A reports that a one percentage point decrease in an MSAs employment rate is associated with a 11.9 basis point increase in total enrollment on average. This response is even larger when restricting to MSA-year observations in which employment was below trend or had decreased since the previous year (see Appendix Figure C.1 and Table C.2).

The remaining columns in Panel A divide enrollment by institution characteristics. Columns 2-4 show that over half of countercyclical enrollment occurs at 2-year colleges despite their relatively small average share of enrollment seen in Table 3.1. Columns 5-7 show that about three quarters of cyclical enrollment is concentrated among public institutions. Community colleges alone are responsible for nearly half of cyclical enrollment (see Appendix Table C.2 for additional enrollment responsiveness results.).

Community colleges likely exhibit strong countercyclical enrollment because they accommodate students seeking career technical education. They typically offer certificate programs that require anywhere from one semester to two years of coursework, and offer awards in a variety of applied fields (Belfield and Bailey, 2017). Students who do not earn an award still exhibit labor market returns to credits in career and technical courses, and will often select high-return courses without the intention of completing an official program (Bahr, 2019).

Community colleges are also typically more accessible than four-year universities. About 95% of community colleges maintain open admissions policies, many operate geographically disperse campuses to accommodate commuting, and regularly offer night and weekend courses for students who are working while enrolled (Horn et al., 2006; Cellini, 2009; Kane and Rouse, 1999). At the start of the Great Recession, average net tuition was $2,511 at US community colleges compared to $8,533 at public flagship universities.[7] As demand for higher education rose over the Great Recession, expansions to Pell grant funding reduced average net tuition by $770 (Barr and Turner, 2013). Low costs and high accessibility are reflected in the student body: compared to students at four-year universities, community college students are more likely to be underrepresented minorities, older, parents, and full-time employees (Ma and Baum, 2016).

Panel B compares enrollment responses by student characteristics. Columns 1-4 show that undergraduates and full-time students are responsible for nearly all cyclical enrollment. Columns 5 and 6 show that women contribute more to countercyclical enrollment than men, and columns 7 and 8 show that nearly 60% of undergraduate enrollment occurs among older students who likely entered the labor force. Note that enrollment rates in Columns 5-8 are normalized by total population rather than demographic-specific population so they are consistent with other estimates in the table. They should not be interpreted as the enrollment rate change for a specific demographic population but rather the change in total enrollment that is attributable to that demographic.

Existing evidence suggests that older students reap strong returns to education in the form of higher earnings, employment, and occupational mobility. Jacobson et al. (2005) studies a sample of 97,000 displaced workers and finds that one year of community college courses raises men and women's long-run earnings by 9% and 13%, respectively. In the same sample, completing any amount of credits increased quarterly hours worked by 2.5% and 3.0% for

[7]Net tuition is calculated as tuition and fees minus student grant aid, and is measured in 2010 dollars.

men and women, respectively. Among workers who completed certificate programs in Oregon between 2007 and 2010, 22% transitioned into *healthcare services* while 11% transitioned out of *manufacturing* and 6% transitions out of *wholesale and retail trade services* (Carnevale et al., 2018). Xu and Trimble (2016) finds that certificate programs in Virginia and North Carolina improve occupational mobility and increase the likelihood of employment, though the estimated effect on earnings is negligible.

Figures 3.1 and 3.2 highlight variation in cyclical enrollment that is not captured by the average estimates in Table 3.2. To illustrate dispersion across MSAs, Figure 3.1 plots the coefficients from Equation 3.1 when estimated for each MSA separately. More than 20% of MSAs report enrollment rises by 20 basis points given a one percentage point decline in employment. Meanwhile, 11% of MSAs report *positive* average responsiveness, implying lower enrollment during worse labor market conditions on average. Most extreme values in Figure 3.1 are from MSAs dominated by major public universities.[8]

Figure 3.2 plots enrollment responsiveness for seven age cohorts spanning 18-64. Enrollment rates are constructed for each age cohort, while the employment rate is measured for the total working-age population.[9] The figure shows that enrollment responsiveness peaks among adults ages 22-24 before steadily declining. For adults ages 22-24, a one percentage point decrease in the local employment rate is associated with a 37 basis point increase in enrollment. Enrollment responsiveness for adults in their late 20s and 30s remains higher than that for college-age adults, and even adults in their 40s exhibit significant enrollment responsiveness. These findings suggest that a meaningful share of adults rely on higher education as a source of retraining and that employment gains are likely concentrated among younger workers in their 20s and early 30s.

3.2.3. Discussion

The average estimates from Table 3.1 are smaller than those in existing literature that use a similar methodology.[10] Barr and Turner (2013) find undergraduate enrollment to be nearly three times more sensitive to state-level unemployment around the Great Recession (2004-2011), though they find little evidence of countercyclical enrollment between 1978 and 2011. Dellas and Sakellaris (2003) use an earlier panel of recent high school graduates (ages 18-

[8]In the most extreme case, Bloomington, Indiana reports that a one percentage point decline in employment is associated with more than a one percentage point rise in enrollment!

[9]Figure C.3 plots enrollment responsiveness when both the enrollment and employment rates are measured within each age cohort. In this variation, enrollment responsivness is highest for adults ages 30-34 and lowest among young adults whose employment rates are less correlated with the total employment rate.

[10]The results are qualitatively consistent with work that uses other measures of countercyclical enrollment. Betts and McFarland (1995) and Hillman and Orians (2013) estimate a semi-elasticity of community college enrollment to unemployment of approximately 4 and 2.4 percent, respectively.

Figure 3.1: Average Enrollment Responsiveness by MSA

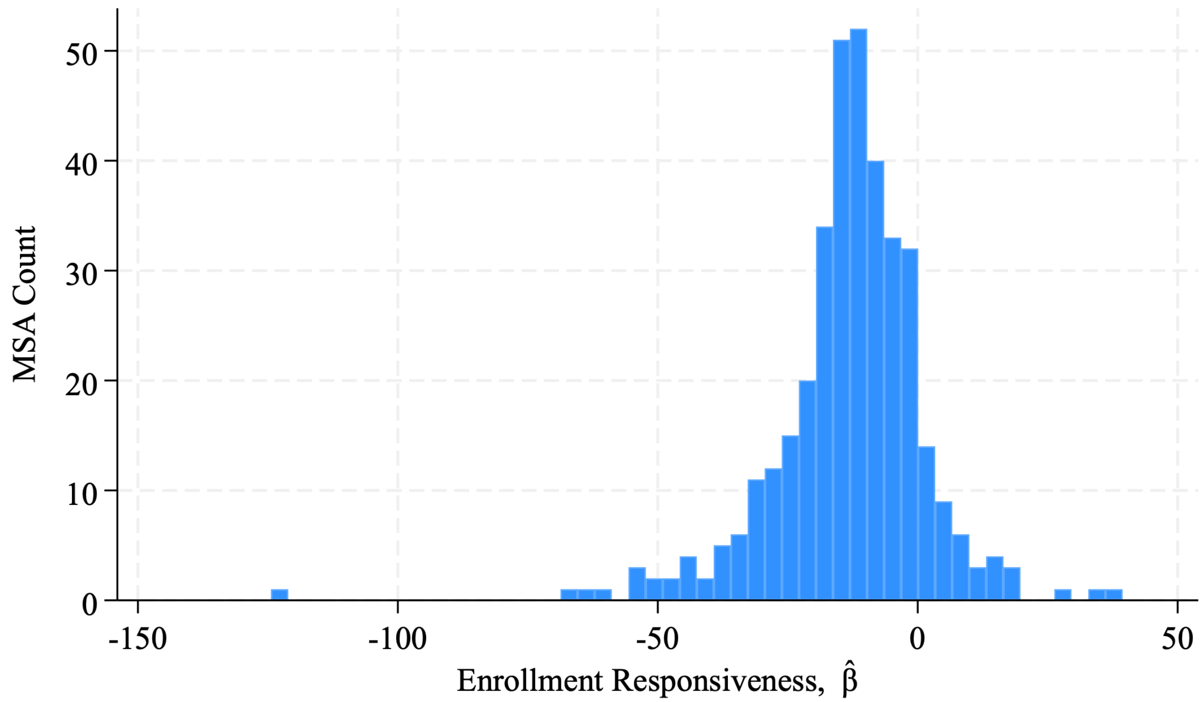

Notes: This figure plots the distribution of point estimates from Equation 3.1 when estimated on each MSA individually. Estimates should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. All rates are detrended by MSA using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. Employment data are from the Quarterly Census of Employment and Wages (QCEW), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data are from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

22) between 1968 and 1988 from the Current Population Survey and estimate that college enrollment increases by about 80 basis points for every one percentage point increase in the national unemployment rate, which is substantially larger than the 14 basis points response found in Figure 3.2. Apart from using different sample periods, the lower estimated responses above may be explained by displaced workers seeking higher education outside of their MSA, which could be picked up in state-level data or individual survey responses in the CPS.

In contrast, these estimates suggest a substantially higher enrollment rate than that implied by recent research on mass layoffs. Foote and Grosz (2019) use data from the Worker Adjustment and Retraining Notification (WARN) Act and find that two-year college enrollment increases by about three students for every one hundred workers laid off. Hubbard (2018) similarly uses mass layoffs from WARN as an instrument for the local unemployment rate and finds that a one percentage point increase in unemployment is associated with a 3.2 percentage point increase in the probability college attendance among recent high school

Figure 3.2: Enrollment Responsiveness by Age

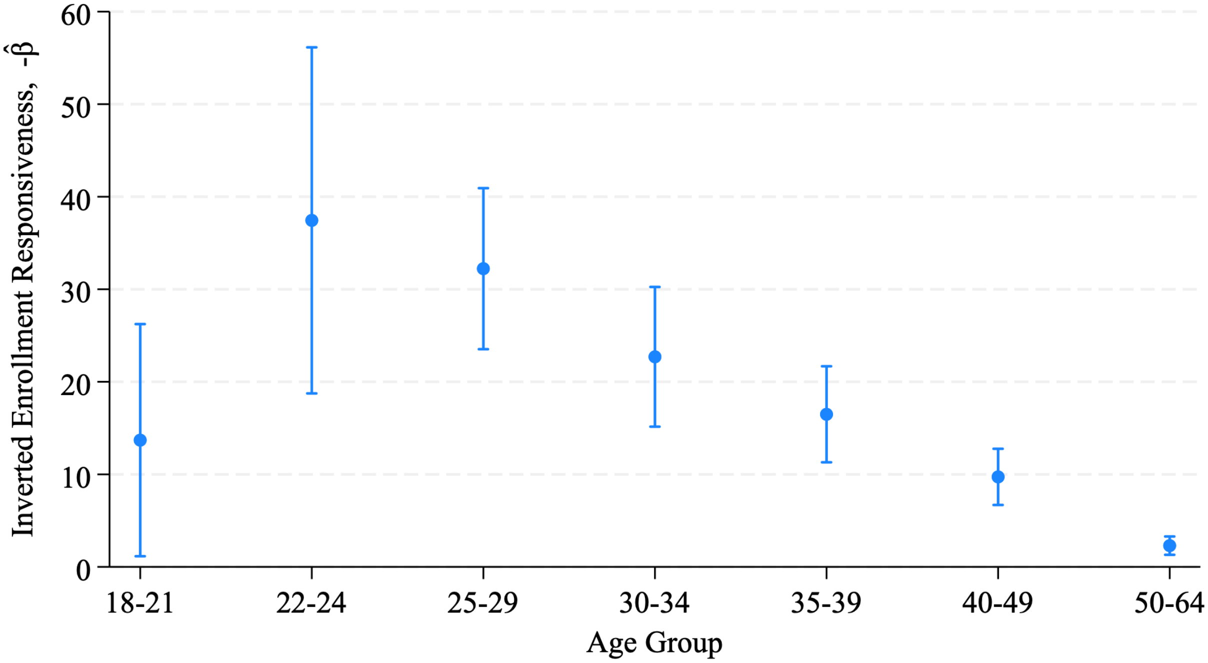

Notes: This figure plots the inverted point estimates and two standard error bars for a series of univariate regressions of annual MSA enrollment rates on the local employment rate by student age. Estimates should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. The employment rate is constructed using adult working-age population (ages 18-65), while enrollment rates are constructed using the respective population of each age group. All rates are detrended by MSA using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. Employment data are from the Quarterly Census of Employment and Wages (QCEW), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data are from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

graduates. These differences might suggest the importance of overall labor market conditions: workers who are displaced by an isolated mass layoff might find work easier than those displaced by skill-specific demand shock and are thus less likely to seek retraining.

3.3. Education and Recovery after the Great Recession

This section assesses the effect of higher education access on local employment recoveries after the Great Recession. The empirical strategy uses a scaled Bartik instrument based on local college composition to isolate new enrollment that is plausibly exogenous to local labor demand shocks. Resulting estimates show that a one percentage point rise in the enrollment rate between 2007 and 2010 is associated with a 2.0 percentage point larger increase in the employment rate by 2018. This effect steadily rises over the course of the recovery as enrollment remains above pre-recession levels and peaks 10 years after the start of the recession. Employment gains are highest among workers ages 25-34 and still apparent among those 35-44, while older age cohorts show little benefit.

3.3.1. Enrollment Surge, 2007-2010

The empirical strategy presented below is best understood in the context of how enrollment and employment evolved over the Great Recession. Figure 3.3 compares the distribution of changes in MSA employment and enrollment rates from 2007 to 2019. Enrollment patterns largely mirrored employment over the entire period, with most new enrollment occurring during the steepest employment losses between 2007 and 2010.[11] In total, 2.4 million additional students enrolled during these years including 1.0 million adults ages 22 or older. Both rates reverted towards pre-recession levels in subsequent years, though their persistence highlights the scarring effect of employment losses during the Great Recession.

Figure 3.4 shows that most new enrollment occurred at public institutions despite remarkable growth of for-profit programs.[12,13] Public universities and community colleges had some of the lowest growth rates despite attracting 70% of new enrollment, which highlights the importance of existing college systems for accommodating sudden rises in education demand. High growth rates among private and for-profit schools may be attributable to spillover demand from students who could not access public education, which would be consistent with long-run evidence that public appropriations for higher education is negatively related to the prevalence of for-profit education (Cellini, 2009; Goodman and Volz, 2020). Appendix Table C.3 tests this hypothesis with a regression of changes in enrollment on start-of-period enrollment rates and does not find evidence that for-profit programs served as complements or substitutes to public institutions in the short-run.

Table 3.3 considers additional local factors that may have affected the enrollment surge. The first three columns regress the change in total enrollment rates between 2007 and 2010 on the contemporaneous change in the employment rate, the start-of-period employment and enrollment rates, the share of adults below age 25, and the share of workers without any college experience. The results show that larger enrollment increases are associated with greater employment losses, higher initial enrollment rates, and a larger share of young adults.

The remaining columns in Table 3.3 restrict attention to community college enrollment

[11] Appendix Figure C.2 shows that the enrollment responsiveness over these years is consistent with the estimates in Table 3.2.

[12] High growth rates and allegedly predatory practices by for-profit institutions were scrutinized by journalists, regulators, and academics alike. See Cellini (2021) and Deming et al. (2012) for historical overviews.

[13] Adjusting for average enrollment duration by institution type does not significantly affect the patterns in enrollment growth found in Panel B. See Appendix C.2.1 for further discussion.

and consider three features of local public colleges: affordability, proximity, and quality. Affordability is defined as the average net price of enrollment, proximity is defined as the log average distance to college campuses within an MSA, and quality is measured as the graduation rate.[14] Surprisingly, higher net attendance cost and lower graduation rates were associated with slightly larger enrollment increases during the Great Recession. Campus proximity had no apparent relationship with enrollment increases during the Great Recession, though proximity appears important for individual students (see Kane and Rouse (1995) for an example).

The quality and availability of higher education changed during the recession as a result of state budget constraints and increased demand. Many states cut funding for public higher education, raised tuition, and increased class sizes in a manner that may have altered the quality or appeal of returning to school. Deming and Walters (2017) find that state spending freezes between 1990 and 2013 significantly impacted enrollment and completion rates, while tuition increases had no effect. Data from State Higher Education Executive Officers Association (SHEEO) shows that state higher education funding slowed in 2009 but federal stimulus through the Education Stabilization Fund and Federal Services Fund largely made up the difference until 2011. Increased funding through federal programs alongside expanded federal student aid in 2011 mitigated much of the effect of budget cuts as student enrollment surged.

Figure 3.3: Interquartile Range of Changes in MSA Employment and Enrollment, 2007-2019

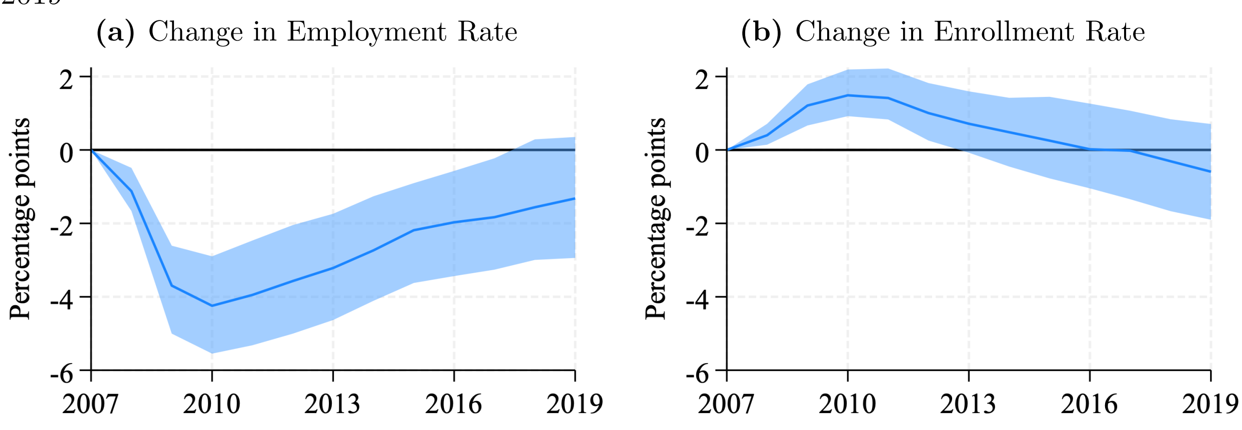

Notes: This figure plots the median and interquartile range of cumulative percentage point changes in MSA employment and enrollment relative to their 2007 values. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

[14]See Appendix C.2.2 for further detail on how these variables are defined.

Figure 3.4: Changing Enrollment by Institution Type, 2007-2010

(a) Enrollment Growth Rate **(b)** New Enrollment in Thousands

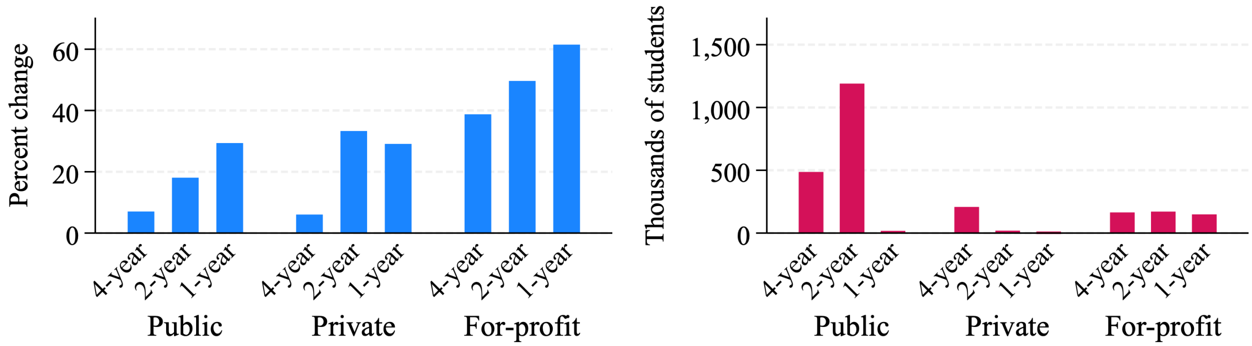

Notes: This figure plots changes in national enrollment across nine types of higher education institutions between 2007 and 2010. Schools are distinguished by their control (Public, Private, and For-Profit) and longest undergraduate program duration (4 years, 2 years, and 1 year). Schools that exclusively offer graduate programs are counted among four-year programs. Panel A presents the percent growth rate in total enrollment and Panel B presents the change in total enrollment measured in thousands of students. Enrollment data are from the Integrated Postsecondary Education Data System (IPEDS).

3.3.2. Estimating the Employment Effect of College Access

The effect of college access on employment is estimated using a first-difference model of the form,

$$\Delta e_{lt} = \alpha_t + \delta_t \Delta s_{l,2010} + \mathbf{X_l \Gamma_t} + \nu_{lt}, \tag{3.2}$$

where Δe_{lt} is the change in employment rate for MSA, l, between 2007 and year, t; $\Delta s_{l,2010}$ is the change in the enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of controls including the start-of-period employment rate, enrollment rate, and share of adults without any college attainment. The model is estimated separately for years 2008 to 2019 and produces a series of estimates, $\hat{\delta}_t$, that captures the evolving relationship between local enrollment and employment rates.

Data on employment is from the US Census's Quarterly Workforce Indicators (QWI), which covers local labor market statistics by worker and firm characteristics since 1990. Similar to QCEW data, QWI is job-level data that covers over 95% of private sector jobs. Since QWI data are constructed using Longitudinal Employer-Household Dynamics microdata, employment data is available by worker demographics not available in the QCEW.[15] Enrollment data is again from IPEDS and population data are from SEER. Unlike Section 3.2, the enrollment rate excludes adults older than 49 because they exhibited a negligible enrollment response in Figure 3.2.

[15] QWI data on employee education is constructed using the American Community Survey (ACS). Section 3.2 uses employment from the QCEW rather than the QWI because it covers a longer sample period.

Table 3.3: MSA Characteristics and the Enrollment Surge, 2007-2010

	All Institutions			Community Colleges		
	Total	Age < 25	Age 25+	Total	Age < 25	Age 25+
$\Delta e_{l,2010}$	-0.14***	-0.26***	-0.10***	-0.13***	-0.15*	-0.11***
	(0.04)	(0.08)	(0.03)	(0.04)	(0.09)	(0.03)
Employment Rate	0.01	0.02	0.01	0.01	0.02	0.00
	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)
Enrollment Rate	0.06***	0.12***	0.04***	0.06***	0.10***	0.04***
	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)
Share adults below 25	-0.12***	-0.04	-0.08***	-0.13***	-0.09	-0.09***
	(0.03)	(0.07)	(0.03)	(0.03)	(0.07)	(0.03)
Share workers without college	0.01	0.03	0.01	0.04**	0.06*	0.03**
	(0.02)	(0.03)	(0.01)	(0.02)	(0.04)	(0.01)
Net attendance cost (thousands)				0.17*	0.06	0.22***
				(0.09)	(0.19)	(0.07)
Distance to campus (log)				-0.06	0.12	-0.10
				(0.14)	(0.30)	(0.12)
Graduation rate				-0.02***	-0.02	-0.01**
				(0.01)	(0.01)	(0.00)
Constant	1.75	-1.22	1.29	1.71	-0.73	1.05
	(1.17)	(2.48)	(0.95)	(1.30)	(2.76)	(1.06)
Observations	370	370	370	308	308	308
R^2	0.14	0.13	0.10	0.17	0.09	0.14

Notes: This table reports estimates of the change in MSA enrollment rates between 2007 and 2010 on a set of local variables. The first three columns report the enrollment rate across all institutions and the last three columns focus on community college enrollment. Enrollment in each category is reported separately for younger (age < 25) and older (age 25+) students. The first five covariates include the contemporaneous change in the employment rate, $\Delta e_{l,2010}$, the local employment and enrollment rates, the share of young adults (age < 25), and the share of workers without any college experience. All rates and shares are expressed in percent. The final three factors are specific to community colleges in each MSA and include the net cost of attendance (in thousands of dollars), the log average distance to a campus, and the graduation rate in 2007. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and education data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

The precision of aggregate administrative data from IPEDS and QWI comes at the cost of potential measurement error from sample differences. Enrollment data from IPEDS is geographically sorted according to school location rather than student location, and the QWI sorts employment by employer location. Workers who attended college and found employment in different MSAs are split across separate geographic units and will add noise to estimated effects. A second limitation is that the data does not track the same cohort of students or workers over time, but rather measures enrollment and employment for fixed age ranges each year. Taking first differences will not eliminate cohort-specific fixed effects as would be the case with survey data that tracks individual workers.

Estimating the effect of college access on labor market recoveries poses two main empirical challenges. First, new enrollment is endogenous to an MSAs initial labor demand shock.

Worse employment losses during the Great Recession will induce higher enrollment rates and require greater job gains for the labor market to recover. This confounding effect will bias estimates of δ in Equation 3.2 downward if not properly addressed. Second, access to higher education alters the composition of workers' educational attainment, which affects the rate at which displaced workers find new employment. MSAs with greater college access and a higher share of workers with at least some college attainment likely experience milder initial job losses and faster recoveries. This would bias estimates of δ upward and arise before any effects from retraining might appear. Estimating the relationship between enrollment and employment without accounting for both empirical challenges would lead to biased estimates.

3.3.3. School Composition and Size as an Instrument for Enrollment

The empirical challenges described above are addressed in two steps. First, I introduce an instrument for new enrollment to reduce the confounding effects of an MSAs initial labor demand shock. The instrument predicts changes in the enrollment rate using the composition and size of local higher education. Local enrollment growth is predicted with a Bartik instrument based on the types of schools available in each MSA. Predicted enrollment growth is then scaled by an MSAs total enrollment to construct an expected change in the enrollment rate.

The Bartik instrument for enrollment growth is described following most notation from Goldsmith-Pinkham et al. (2018). For a given MSA, l, enrollment growth can be expressed as the weighted sum of growth rates across K types of higher education institutions,

$$g_l = \sum_{k=1}^{K} w_{lk} g_{lk},$$

where g_l is total enrollment growth in MSA l, w_{lk} is the share of enrollment in MSA l that occurs at type k schools, and g_{lk} is local growth rate for type k schools. Local enrollment growth by school type can be decomposed into a national and idiosyncratic term, $g_{lk} = g_k + \tilde{g}_{lk}$. The Bartik instrument for local enrollment growth is constructed by replacing observed local growth rates with the national growth rate,

$$B_l = \sum_{k=1}^{K} w_{lk} g_k.$$

This growth rate instrument alone is not strong enough to act as an instrument for changes in local enrollment rates because the scale of local enrollment is so important (see Figure 3.3 and its related discussion). Consequently, I combine B_l and start-of-period enrollment to predict the change in enrollment rates between 2007 and 2010,

$$\Delta z_{l,2010} = \frac{(1 + B_l)S_{l,2007}}{P_{l,2010}} - \frac{S_{l,2007}}{P_{l,2007}},$$

where S_{lt} and P_{lt} are total enrollment and population for MSA l in year t, respectively.

While this approach reduces bias from idiosyncratic labor demand shocks, it does not address the second empirical challenge of how the size of local higher education affects workers' educational attainment and employment outcomes. To address these concerns, Equation 3.2 includes start-of-period controls for the share of workers without college experience, enrollment rate, and employment rate.

This empirical design allows for two falsification exercises based on the timing and demographic concentration of enrollment effects. The first test uses Equation 3.2 to estimate the short-term impact of new enrollment on employment in 2008 and 2009. Since most new enrollment did not occur until 2009, we should not expect a substantial employment effect before 2010. Evidence of an earlier effect would suggest that the empirical design's identification assumptions are violated. The second falsification exercise tests whether older workers who had very low enrollment rates experienced any employment gains. Similarly, any sizeable employment effect among older adults is evidence that the identification strategy has failed. These falsification exercises cannot prove that the exclusion restriction holds but offer the opportunity to expose a clear violation.

3.3.4. Main Results

Figure 3.5 plots the OLS and IV estimates for δ from Equation 3.2. These estimates should be interpreted as the change in the employment rate between 2007 and year t that is associated with a one percentage point rise in the enrollment rate between 2007 and 2010. As expected, the OLS results in Panel A show a negative relationship between enrollment and employment in the short-run, which is consistent with a confounding effect from the initial labor demand shock. Estimates then revert toward zero and are not statistically distinguishable from zero in the long-run.

Panel B plots the IV estimates using $\Delta z_{l,2010}$ as an instrument for enrollment changes, $\Delta s_{l,2010}$. In the short-run, the IV results show no effect of enrollment on employment in 2008 and a negative effect in 2009. The estimate for 2009 is imprecise and small relative those in later years, though it may still suggest that the instrument's exclusion restriction does not hold. Estimates begin to steadily rise in 2010 and reach a peak of 2.0 in 2018, implying that a one percentage point rise in an MSAs enrollment rate between 2007 and 2010 corresponds to a 2.0 percentage point higher increase in its employment rate between 2007 and 2018.

The estimated employment response to college enrollment rose fastest between 2010 and 2012 which approximately lags the enrollment surge that peaked in 2010. According to the Department of Education's Beginning Postsecondary Students Longitudinal Study (BPS), students ages 25 and older who enrolled in public universities and community colleges in 2004 attended school for 2.2 and 1.5 years on average, respectively.[16] Since most new enrollment occurred at public schools, we should expect employment effects to appear about two year after new enrollment.

Figure 3.5: Effect of New Enrollment on Employment Recovery

(a) OLS Estimates **(b)** IV Estimates

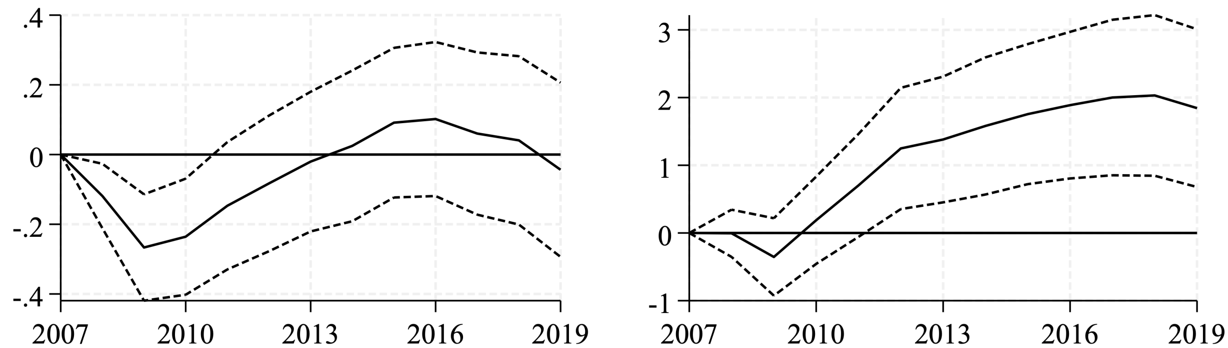

Notes: This figure plots estimates for δ_t from Equation 3.2, $\Delta e_{lt} = \alpha_t + \delta_t \Delta s_{l,2010} + \mathbf{X_l \Gamma_t} + \nu_{lt}$, for years 2008-2019. Δe_{lt} represents the change in the local employment rate for MSA, l, between 2007 and year t; $\Delta s_{l,2010}$ represents the change in the local enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of start-of-period controls including the share of employees without any college experience, the employment rate, and the enrollment rate. Coefficients should be interpreted as the percentage point change in the employment rate between 2007 and year t associated with a one percentage point increase in the enrollment rate between 2007 and 2010. Panel A plots OLS estimates of the model and Panel B plots IV estimates using $\Delta z_{l,2010}$ as an instrument for $\Delta s_{l,2010}$. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure 3.6 shows how the relationship between enrollment and employment varies with worker age by plotting the IV estimates from Equation 3.2 for five different age cohorts.[17] The effect is strongest among adults ages 25-34 who experience a 3.9 percentage point higher employment rate by 2016 given a one percentage point rise in the total enrollment rate by 2010. Adults ages 35-44 exhibit a positive and statistically significant employment gains from new enrollment, and the two oldest cohorts show only weakly positive effects over the sample period. For the oldest cohort, ages 55-64, a one percentage point increase in the total enrollment rate is associated with a 0.3 percentage point higher employment rate by the end

[16]See Appendix Table C.5 for further detail on average years of attendance by institution type and student age.

[17]The corresponding OLS estimates are plotted in Appendix Figure C.4.

of the sample period. This effect is an order of magnitude smaller than that of adults ages 25-34 and remains similarly small after standardizing changes to the employment rate across MSAs.[18] Such a muted response among the oldest cohort is consistent with low enrollment responsiveness plotted in Figure 3.2 and offers some assurance that the identification strategy is not violated by an unobserved factor that would impact young and old workers alike.

The youngest cohort of adults, ages 19-24, weakly benefited from college access in the years immediately following the Great Recession between 2010 and 2013. Since employment is measured for a different cohort of workers each year, it is not surprising that the positive effect among these workers dissipates in later years as adults who were not yet working-age during the Great Recession transition into the youngest age group. Workers who *were* ages 19-24 during the recession eventually age into the next oldest age cohort, 25-34, and are likely responsible for the strong employment effects observed in that group. A continuation of these compositional dynamics may explain why estimates for adults ages 25-34 peaked in 2016 while those for adults ages 35-44 continued to rise.

[18] After standardizing, the point estimate is 0.15 for adults ages 55-64 compared to 0.95 for adults ages 25-34. This result is available upon request.

Figure 3.6: Effect of New Enrollment on Employment by Age (IV Estimates)

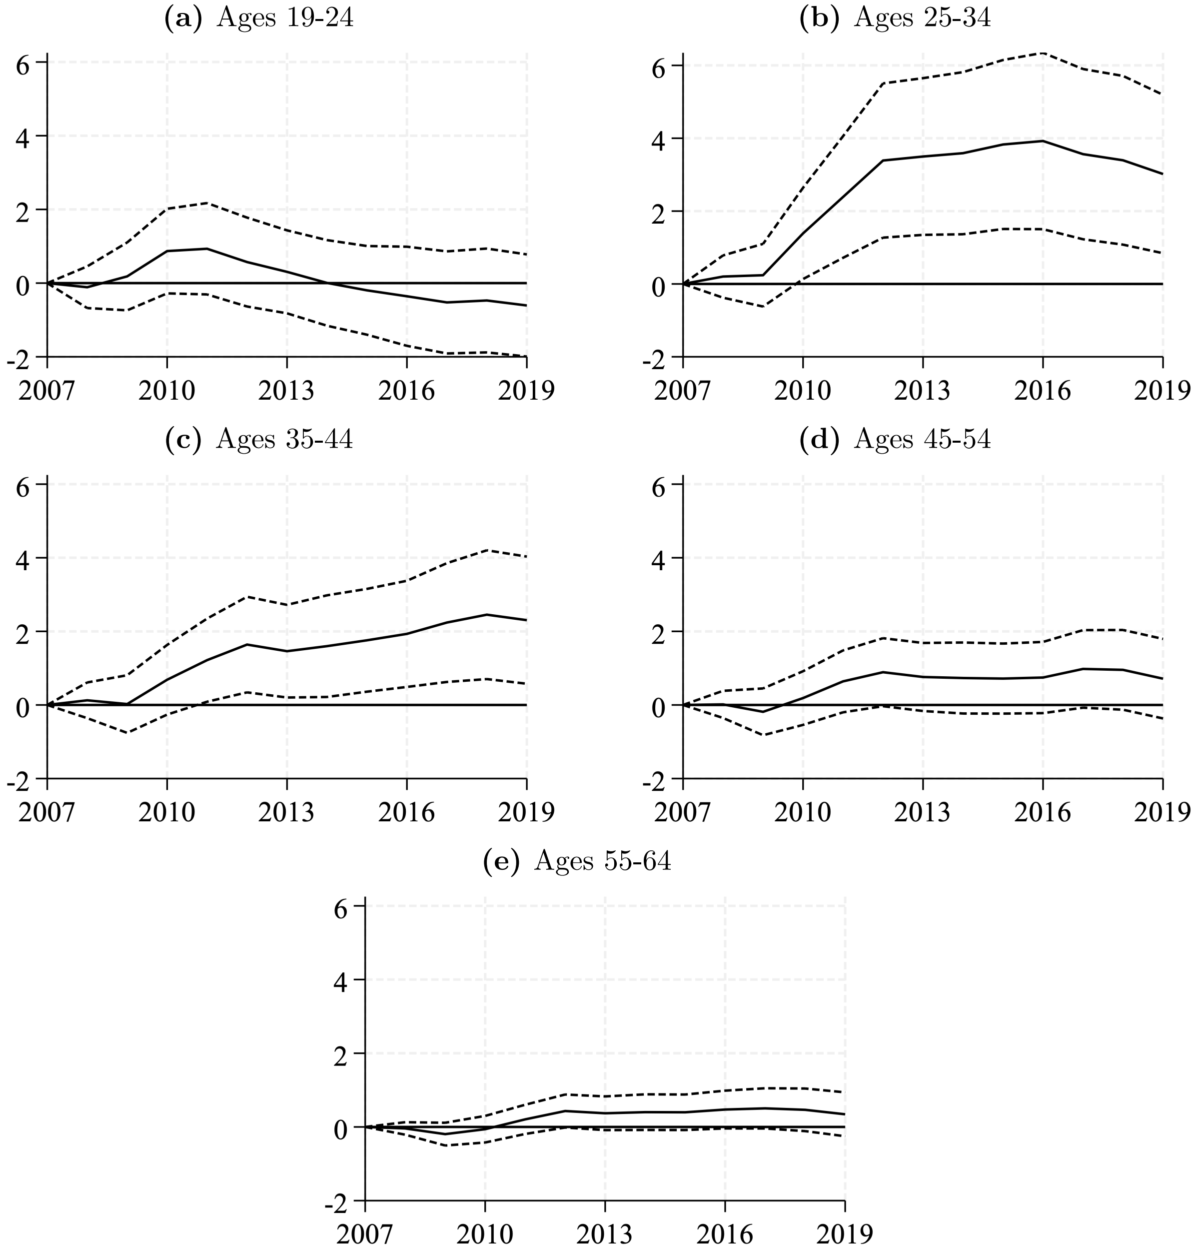

Notes: This figure plots IV estimates for δ_t from Equation 3.2, $\Delta e_{lt} = \alpha_t + \delta_t \Delta s_{l,2010} + \mathbf{X_l \Gamma_t} + \nu_{lt}$, for years 2008-2019 and five separate ages cohorts. Δe_{lt} represents the change in the local employment rate for each age cohort in MSA, l, between 2007 and year t; $\Delta s_{l,2010}$ represents the change in the local total enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of start-of-period controls including the share of employees without any college experience, the age-specific employment rate, and the age-specific enrollment rate. Coefficients should be interpreted as the percentage point change in the employment rate between 2007 and year t associated with a one percentage point increase in the enrollment rate between 2007 and 2010. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

3.3.5. Gender Differences in Enrollment and Employment Effects

Table 3.2 shows that men and women contribute similarly to cyclical enrollment, with women accounting for slightly more than half of the average enrollment response to a change in the local employment rate. However, this analysis does not account for differences in employment rate volatility between men and women. When Equation 3.1 is estimated using the employment rate for men and women separately, women appear substantially more responsive to local employment conditions. Table 3.4 shows that a one percentage point decrease in the employment rate for women is associated with a 13.5 basis point higher average enrollment rate, while the average enrollment rate increase among men is only 7.8 basis points.[19]

This gender gap is also shown in Figure 3.7, which compares enrollment responsiveness by gender across several age cohorts. Men and women in the youngest cohort respond similarly to changes in their respective employment rates, while enrollment for all other age cohorts is higher among women than men. Results for the youngest cohort suggest that gender differences are less relevant for those enrolling for the first time or deciding whether to remain enrolled. In contrast, women in older age cohorts appear more likely to pursue higher education in response to worsening labor market conditions.

Despite a stronger enrollment response, women appear to have benefited less from access to higher education during the Great Recession. Figure 3.8 reproduces the main results from Figure 3.5 using separate employment and enrollment rates for men and women. The OLS estimates in Panel A show a similar negative association between new enrollment and employment in both groups, which is likely caused by the confounding effect of the initial labor demand shock. Panel B plots the IV estimates for Equation 3.2 and shows that higher education access improved employment rates for both men and women following the Great Recession. Areas with larger increases in enrollment between 2007 and 2010 began to see higher employment rates among both men and women around 2010. The average employment effect for men steadily grows over the remaining event horizon, while the effect for women plateaus around 2015. By the end of the sample horizon, the average effect of a one percentage increase in enrollment for men was about twice as large as that for women.

[19] This analysis relies on QWI data across 380 MSAs beginning in 1990 since the QCEW data used in Figure 3.2 does not separate employment by gender. By and large, both data sources produce similar estimates of cyclical enrollment: total enrollment rate responsiveness estimated using QWI data is -12.4(2.2) in column 1 of Table 3.4 compared -11.9(1.9) for QCEW data reported in Table 3.2.

Table 3.4: Enrollment Responsiveness to Local Employment Rate by Gender

	(1) Total	(2) Men	(3) Women
Employment Rate	-12.4	-7.8	-13.5
	(2.2)	(1.4)	(2.6)
Observations	8,264	8,264	8,264
R^2	0.09	0.10	0.09

Notes: This table reports estimates of a univariate regression of annual MSA enrollment rates (in basis points) on the local employment rate (in percent). Column 1 uses total MSA enrollment, employment, and adult working-age population, while Columns 2 and 3 separate enrollment, employment, and population by men and women, respectively. Coefficients should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. Standard errors are in parenthesis. All rates are detrended using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1990 and 2019. Employment data are from the Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure 3.7: Enrollment Responsiveness by Age and Gender

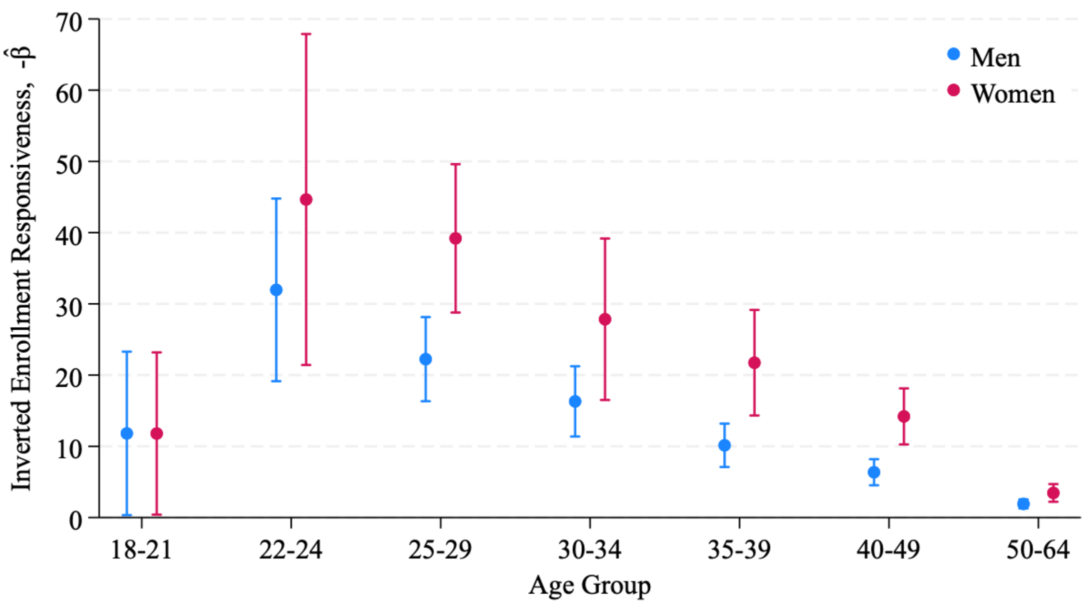

Notes: This figure plots the inverted point estimates and two standard error bars for a series of univariate regressions of annual MSA enrollment rates on the local employment rate by student age and gender. Estimates should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. The employment rates for men and women are constructed using their respective adult working-age population (ages 18-65), while enrollment rates are constructed using the respective population of each age group by gender. All rates are detrended by MSA using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1990 and 2019. Employment data are from the Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data are from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure 3.8: Effect of New Enrollment on Employment Recovery by Gender

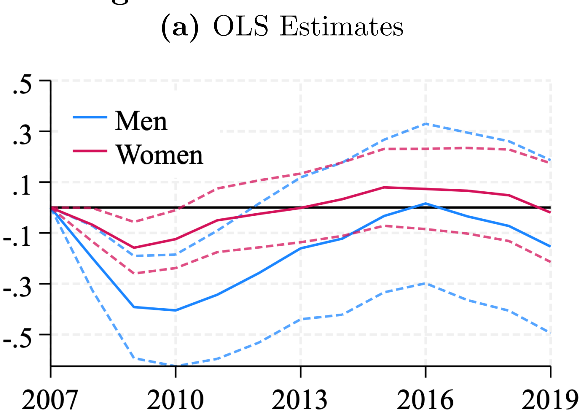

Notes: This figure plots estimates for δ_t by gender from Equation 3.2, $\Delta e_{lt} = \alpha_t + \delta_t \Delta s_{l,2010} + \mathbf{X_l \Gamma_t} + \nu_{lt}$, for years 2008-2019. Δe_{lt} represents the change in the local employment rate for MSA, l, between 2007 and year t; $\Delta s_{l,2010}$ represents the change in the local enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of start-of-period controls including the share of employees without any college experience, the employment rate, and the enrollment rate. Coefficients should be interpreted as the percentage point change in the employment rate between 2007 and year t associated with a one percentage point increase in the enrollment rate between 2007 and 2010. Panel A plots OLS estimates of the model and Panel B plots IV estimates using $\Delta z_{l,2010}$ as an instrument for $\Delta s_{l,2010}$. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

3.3.6. Sensitivity to Net Migration

Enrollment and employment rates in the above analysis are constructed using contemporaneous population to avoid spurious correlation from population trends. This approach may bias the estimated effect of enrollment on employment by removing population changes that are endogenous to labor market conditions or access to higher education. If out-migration was higher in MSAs with larger labor demand shocks then the effect of enrollment would be biased upward; if college access deterred out-migration or attracted students from other MSAs then estimates would be biased downward. This section investigates potential bias from constructing rates with contemporaneous population and re-estimates the main analysis from Figure 3.5 using start-of-period population and controlling for population growth. It shows that the results are qualitatively similar but lose statistical significance likely due to a weakened instrument.

Panel A in Figure 3.9 illustrates the effect of population changes on the reduced form relationship between enrollment and employment. It compares estimates of θ_1 from the univariate model, $\Delta e_{lt} = \theta_0 + \theta_1 \Delta s_{l,2010} + \varepsilon_l$, when rates are constructed with start-of-period and contemporaneous population. Rate changes based on start-of-period population are defined as, $\Delta x_t = \frac{X_t - X_{2007}}{P_{2007}}$, and those constructed with contemporaneous population are defined as $\Delta \tilde{x}_t = \frac{X_t}{P_t} - \frac{X_{2007}}{P_{2007}}$. Data and variable definitions are otherwise identical to those described in Section 3.3.2. Resulting estimates based on the start-of-period population are higher than those based on contemporaneous population over the entire event horizon and show weak countercyclical enrollment during the Great Recession. This difference does not suggest that using contemporaneous population biases the estimated effect of enrollment upward due to out migration. Instead, it is consistent with a spurious correlation between employment and enrollment rates when not accounting for population trends.

The potential effect of a population-driven spurious correlation is shown in Panel B of Figure 3.9, which plots the share of variation in MSA employment rates since 2007 that is explained by population growth. The share of variation is measured as the R^2 of the model, $\Delta e_{lt} = \kappa_0 + \kappa_1 p_{lt} + \eta_l$, where p_{lt} is the population growth rate for MSA l between 2007 and year t. Resulting estimates show that a substantial and increasing share of changes in local employment rates is explained by population growth when using start-of-period population. By the end of the event horizon, more than two-thirds of variation in employment rate changes can be explained by population changes. For comparison – and largely by construction – local population growth explains at most 7% of changes in employment rates when constructed with contemporaneous population.

Figure 3.10 tests whether the results in Figure 3.5 are sensitive to adjusting for population changes. Unlike Figure 3.5, employment and enrollment rates are constructed using start-of-period population rather than contemporaneous population, and MSA population growth is included as a covariate to control for its common effect on local employment and enrollment rates. The resulting estimates are qualitatively similar to those in Figure 3.5 but suggest that the analysis is sensitive to accounting for population growth. The OLS estimates in Panel A exhibit an initial negative relationship between enrollment and employment rate changes between 2007 and 2010 – likely due to the confounding effect of the initial labor demand shock – before reverting in later years. The IV estimates in Panel B follow a similar pattern to those in Figure 3.5, though estimates have larger magnitudes and are less precise. Point estimates during the Great Recession are larger, which suggests that the instrument fails to identify new enrollment that is unrelated to local labor demand shocks. First-stage estimates in Column 6 of Appendix Table C.4 show that using start-of-period population greatly reduces the instrument's power.

Figure 3.9: Potential Bias from Population Adjustment

(a) Employment Rate Changes Associated with New Enrollment, $\hat{\theta}_1$

(b) Variation in Employment Rates Explained by Population Growth, R^2

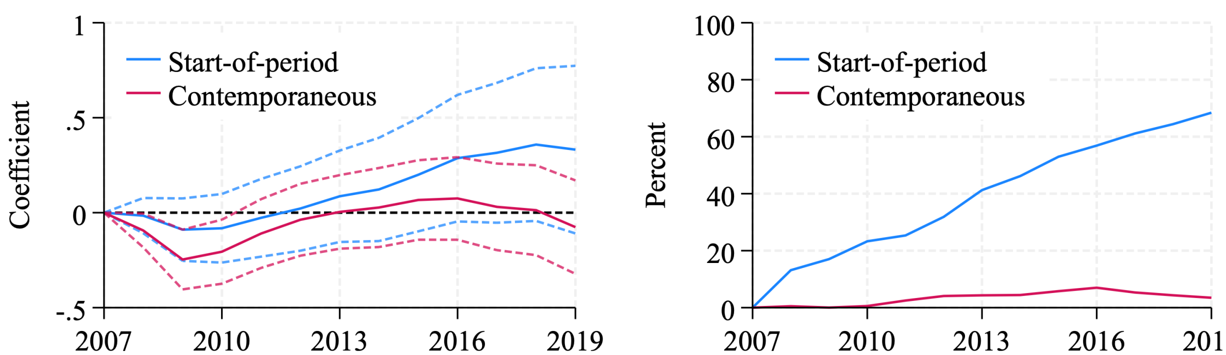

Notes: This figure presents two panels that illustrate the effect of adjusting population changes in first-differences analysis. Each panel compares estimates based on employment and enrollment rates that are constructed using start-of-period population (blue) and contemporaneous population (red). Panel A plots estimates of θ_1 from the univariate model, $\Delta e_{lt} = \theta_0 + \theta_1 \Delta s_{l,2010} + \varepsilon_l$, where Δe_{lt} represents the change in the employment rate between 2007 and year t for MSA l and $\Delta s_{l,2010}$ represents the change in enrollment rate between 2007 and 2010. Estimates for θ_1 should be interpreted as the percentage point change in the employment rate associated with a one percentage point change in the enrollment rate. Panel B plots the R^2 of the model, $\Delta e_{lt} = \kappa_0 + \kappa_1 p_{lt} + \eta_l$, where p_{lt} is the population growth rate for MSA l between 2007 and year t. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure 3.10: Effect of New Enrollment on Employment using Start-of-Period Population

(a) OLS Estimates **(b)** IV Estimates

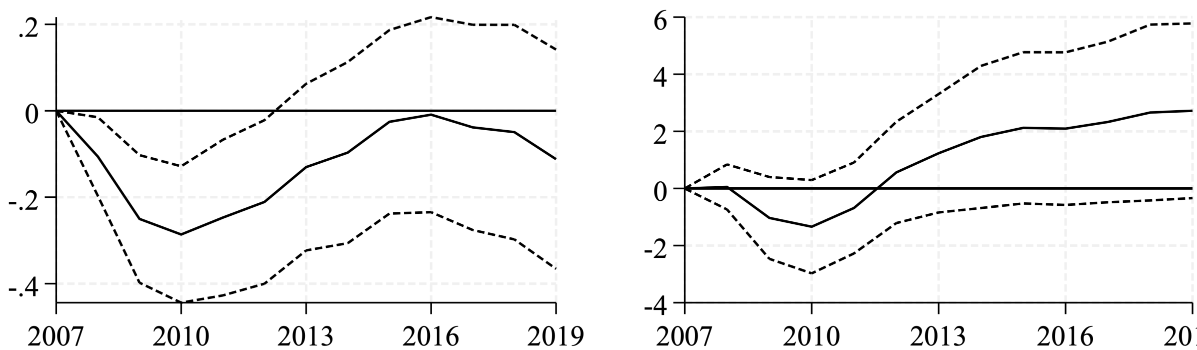

Notes: This figure plots estimates for δ_t from Equation 3.2, $\Delta e_{lt} = \alpha_t + \delta_t \Delta s_{l,2010} + \mathbf{X_l \Gamma_t} + \nu_{lt}$, for years 2008-2019. Δe_{lt} represents the change in the local employment rate for MSA, l, between 2007 and year t; $\Delta s_{l,2010}$ represents the change in the local enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of start-of-period controls including the share of employees without any college experience, the employment rate, and the enrollment rate. Coefficients should be interpreted as the percentage point change in the employment rate between 2007 and year t associated with a one percentage point increase in the enrollment rate between 2007 and 2010. Panel A plots OLS estimates of the model and Panel B plots IV estimates using $\Delta z_{l,2010}$ as an instrument for $\Delta s_{l,2010}$. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using MSA population in 2007 from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

3.3.7. Discussion

The above findings suggest that higher education served as an important means of local labor market adjustment following the Great Recession. Much of its importance may have stemmed from declining demand for routine occupations and the subsequent need for new training among affected workers. Jaimovich and Siu (2012) calculate that 88% of net job losses in these occupations coincided with recessions since the mid-1980s and that routine employment losses accounted for nearly all job loss in the Great Recession. Multiple sectors that employed workers of similar, imperfectly transferable skills were shocked during the Great Recession and left workers with fewer alternative employment opportunities than in previous recessions (Grigsby, 2019). Higher education may have offered displaced workers the opportunity to find new employment by training for new occupations and pivoting careers.

The popularity and success of retraining through higher education may partially explain why so few workers "moved to opportunity" during the Great Recession (Yagan, 2014). Enrolling locally allowed workers to find new employment without moving to a different area and offered the prospect of better job opportunities. The apparent decline in geographic mobility could suggest that displaced workers in the Great Recession were more reluctant to remain in their occupation, or that returning to school became a more appealing path to new employment.

3.4. Conclusion

This paper offers new detail on countercyclical enrollment in higher education and shows that access to higher education during the Great Recession altered the trajectory of local employment rates. It presents an empirical strategy that isolates plausibly exogenous increases in enrollment during the Great Recession based on the local composition of schools and estimates that a one percentage point increase in total enrollment between 2007 and 2010 is associated with a 2.0 percentage point larger increase in the employment rate between 2007 and 2018. The estimated employment effect is largely consistent with the timing and age composition of new enrollment during the Great Recession.

APPENDIX A

Appendix to Chapter 1

A.1. Additional Tables and Figures

Table A.1: Number of OSHA Inspections by Reason for Inspection

	All Inspections		Manufacturing	
	N	%	N	%
Programmed	2,916,122	58.6	707,781	53.9
Complaint/Referral	1,522,561	30.6	413,370	31.5
Monitoring/Follow-up	329,564	6.6	139,236	10.6
Accident	211,549	4.2	52,867	4.0
Total	4,979,796	100.0	1,313,254	100.0

Notes: This table reports the number of OSHA inspections by reason for inspection between 1972 and 2023. *Programmed inspections* are scheduled based on "objective or neutral criteria"; *Complaint/Referral inspections* follow notice of an alleged safety hazard by a current employee (complaint) or external source (referral); *Monitoring/Follow-up inspections* revisit worksites with previously inspected hazards; and *Accident inspections* follow any worker fatalities or hospitalizations. The second column of inspection counts restrict to manufacturing worksites. Data on OSHA inspections are available through the Department of Labor's Enforcement Data Catalog.

Table A.2: Number of OSHA Inspections by Industry

	All Inspections		Programmed	
	N	%	N	%
Construction	2,249,005	45.2	1,592,289	54.6
Manufacturing	1,313,254	26.4	707,781	24.3
Services	522,527	10.5	196,361	6.7
Transp, Comm, Utilities	228,609	4.6	98,659	3.4
Retail Trade	184,756	3.7	74,525	2.6
Public Admin	164,910	3.3	80,499	2.8
Wholesale Trade	135,043	2.7	64,661	2.2
Agriculture	111,308	2.2	69,102	2.4
Mining	37,344	0.7	23,401	0.8
Finance, Real Estate	33,040	0.7	8,844	0.3
Total	4,979,796	100.0	2,916,122	100.0

Notes: This table reports the number of OSHA inspections by industry between 1972 and 2023. Industry categories are based on SIC divisions and NAICS sectors. The second column of inspection counts restrict to programmed inspections, which are scheduled in advance according to “neutral or objective” selection criteria. Data on OSHA inspections are available through the Department of Labor’s Enforcement Data Catalog.

Table A.3: OSHA-Census Match Quality

Sample	Variable	N	Mean	S.D.	25th	Median	75th
Full	Employment	240,465	100.4	414.6	14	32	85
	Violations	240,465	5.8	7.0	1	4	8
	Total Penalty	240,465	1,809.4	21,462.5	0	178	1,168
Matched	Employment	122,000	96.8	287.3	19	39	92
	Violations	122,000	6.5	7.2	2	5	9
	Total Penalty	122,000	2,104.0	18,860.0	0	316	1,505

Notes: This table compares the full sample of OSHA programmed inspections at manufacturing establishments between 1979 and 1997 to a subsample of inspections that is matched to establishment records from the US Census Bureau. *Employment* reports the number of employees at the inspected establishment; *Violations* reports the number of violations discovered during the inspection; and *Total Penalty* reports the total fine assessed against the inspected establishment. Penalties are deflated to 2012 dollars using the BEA investment price deflator. Employment data for the matched sample is from the US Census Bureau’s Longitudinal Business Database, and all other data are from the Department of Labor’s Enforcement Data Catalog.

Figure A.1: Robustness Exercises for Employment Response

(a) ASM Sample

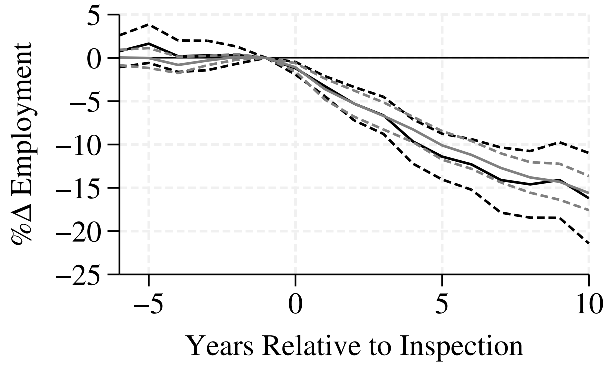

(b) Targeted Inspections (1998-2013)

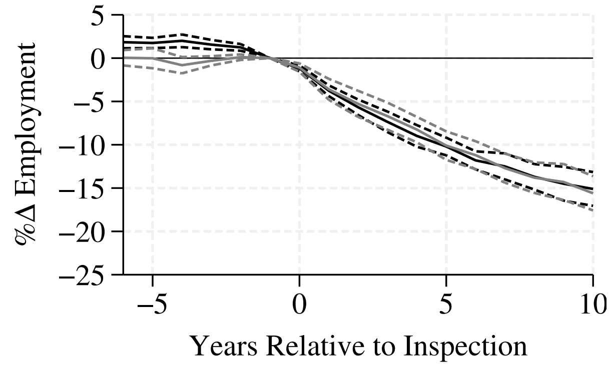

(c) Soon-to-be-Inspected

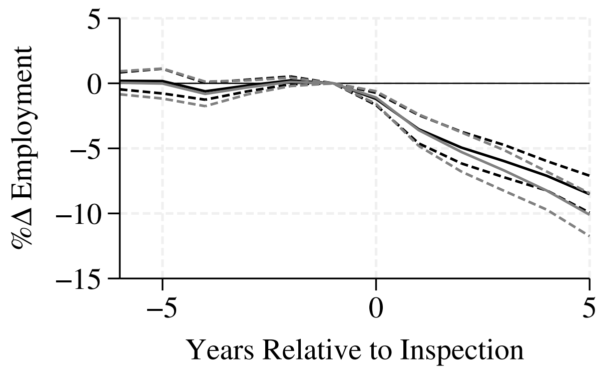

(d) Balanced Panel

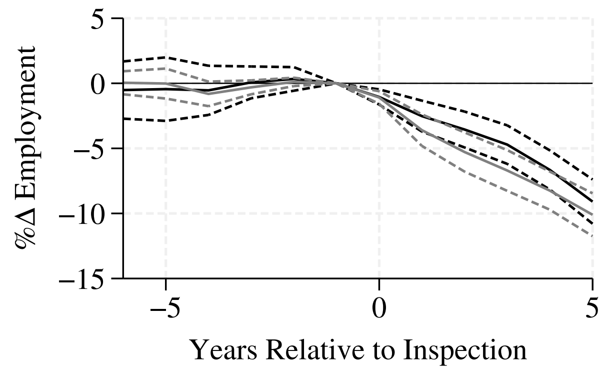

Notes: This figure plots four robustness checks (black) to the baseline employment response (gray) in Panel (a) of Figure 1.1. Each panel plots the average employment response to randomly scheduled OSHA inspections at manufacturing establishments between 1979 and 1997. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in employment due to inspection since the year preceding inspection. Panel (a) restricts the sample to establishments in the ASM; Panel (b) uses programmed inspections after OSHA began non-randomly targeting unsafe worksites in 1998; Panel (c) restricts control units to establishments that were inspected within five years of the treated units; and Panel (d) restricts to a balanced panel of establishments that appear in every year of the event study. Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

Figure A.2: Effect of OSHA Inspections on Real Output

(a) Random Inspections

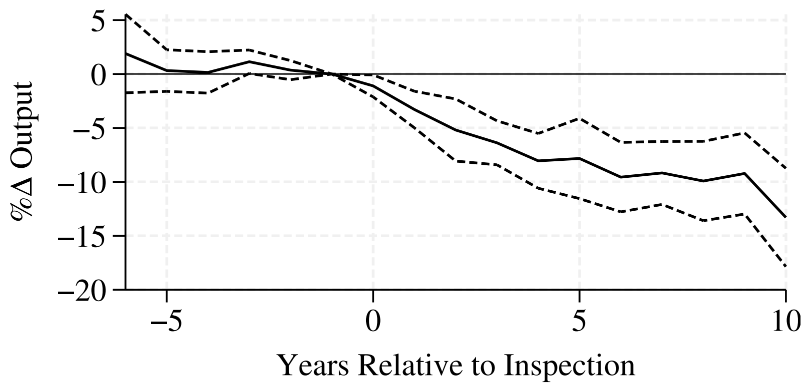

(b) Complaint/Referral Inspections

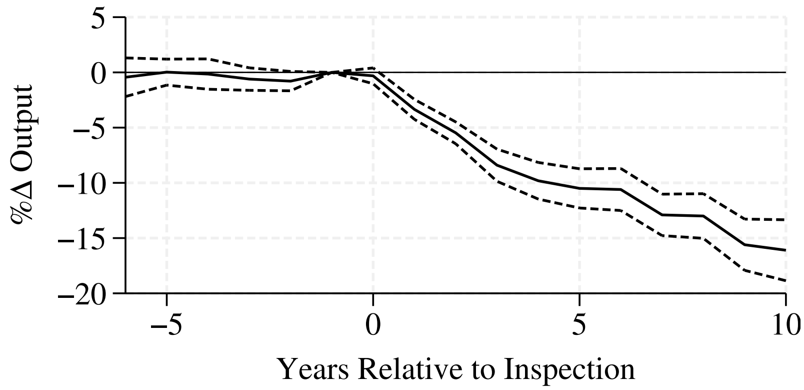

(c) Accident Inspections

Notes: This figure plots the average output response to three samples of OSHA inspections at manufacturing establishments. Responses are estimated using a local projection diff-in-diff model and can be interpreted as the percent change in real output due to inspection since the year preceding inspection. Nominal output is defined as total value of shipments adjusted for resales and changes in inventories. Real output is deflated using industry-specific shipments deflator, *piship*, published by the NBER-CES Manufacturing Industry Database. *Random Inspections* in Panel (a) covers randomly scheduled programmed inspections from 1979-1997; *Complaint/Referral Inspections* in Panel (b) covers inspections following an employee complaint or external referral from 1977-2013; and *Accident Inspections* in Panel (c) covers inspections following a workplace accident from 1977-2013. Dashed lines denote the 95% confidence interval. Data on OSHA inspections are from the Department of Labor, and data on manufacturing operations are from the US Census Bureau.

APPENDIX B

Appendix to Chapter 2

B.1. Additional Tables and Figures

B.1.1. Construction of the textual measure

Table B.1 contains the list of keywords used in frequency search under each topic. The keywords are based on Econoday, which provides notifications for major economic news and is the service behind Bloomberg economic calendar.

Table B.1: Macroeconomic topics and keywords

Topic	**Keywords**
General	economic conditions
Output	GDP, economic growth, macroeconomic condition, construction spending, national activity, recession
Employment	unemployment, JOLTS, labor market, jobless claims, jobs report, non-farm payroll, ADP employment report, empoyment cost index
Consumption	consumer confidence, consumer credit, consumer sentiment, durable goods, personal income, retail sales
Investment	business inventories, manufacturing survey, factory orders, business outlook survey, manufacturing index, industrial production, business optimism, wholesale trade
Monetary	FOMC, monetary policy, quantitative easing
Housing	home sales, home prices, housing starts, housing market
Inflation	price index, price level, consumer price index, CPI, PMI, PPI, inflation, inflationary, disinflation, disinflationary, hyperinflation, hyperinflationary
Oil	oil prices, oil supply, oil demand

Notes: Dictionary of keywords used in constructed text-based attention measures. Keywords are based on names of macroeconomic releases from `Econoday`, complemented with macroeconomic words and phrases from popular press.

Table B.2: Summary statistics of firm characteristics by attention

	Mean	**Median**	**SD**	**N**
Attentive				
Size (million)	7,312	538	65,275	102,493
Age	11.57	10.00	7.37	103,312
Leverage	0.30	0.20	0.46	101,981
Inattentive				
Size	2,873	104	35,004	33,277
Age	7.78	7.00	4.98	33,796
Leverage	0.35	0.17	0.69	32,955
All firms				
Size	6,224	371	59,333	135,770
Age	10.64	9.00	7.05	137,108
Leverage	0.31	0.19	0.53	134,936

Notes: This table reports summary statistics for firm-year observations that are identified as attentive or inattentive according our our "general" attention topic. Firm size is measured by total assets, age is measured as the number of years since the firm first appeared in our sample, leverage is defined as the ratio of total debt to market equity. Values for leverage are winsorized at 1%.

B.1.2. Firm attention to macroeconomic topics

Figure B.1: Time series of firm attention to macro topics

(a) General

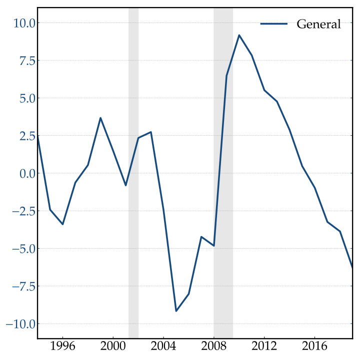

(b) Output

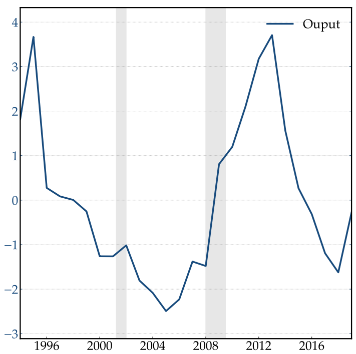

(c) Employment

(d) Consumption

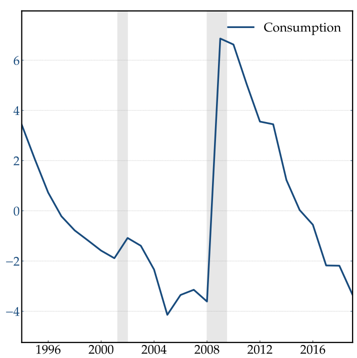

(e) Investment

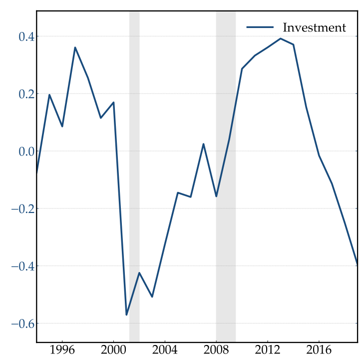

(f) Monetary

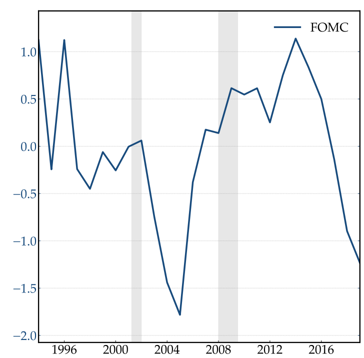

(g) Housing

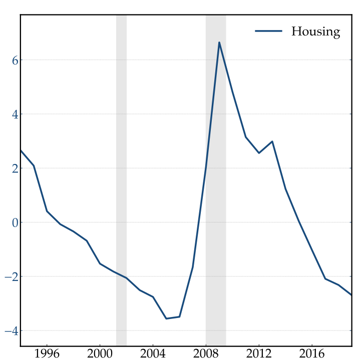

(h) Inflation

(i) Oil

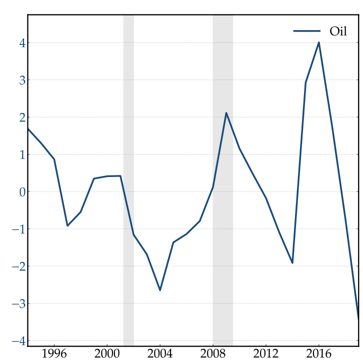

Figure B.2: Cross-sectional distribution of firm attention to macro topics

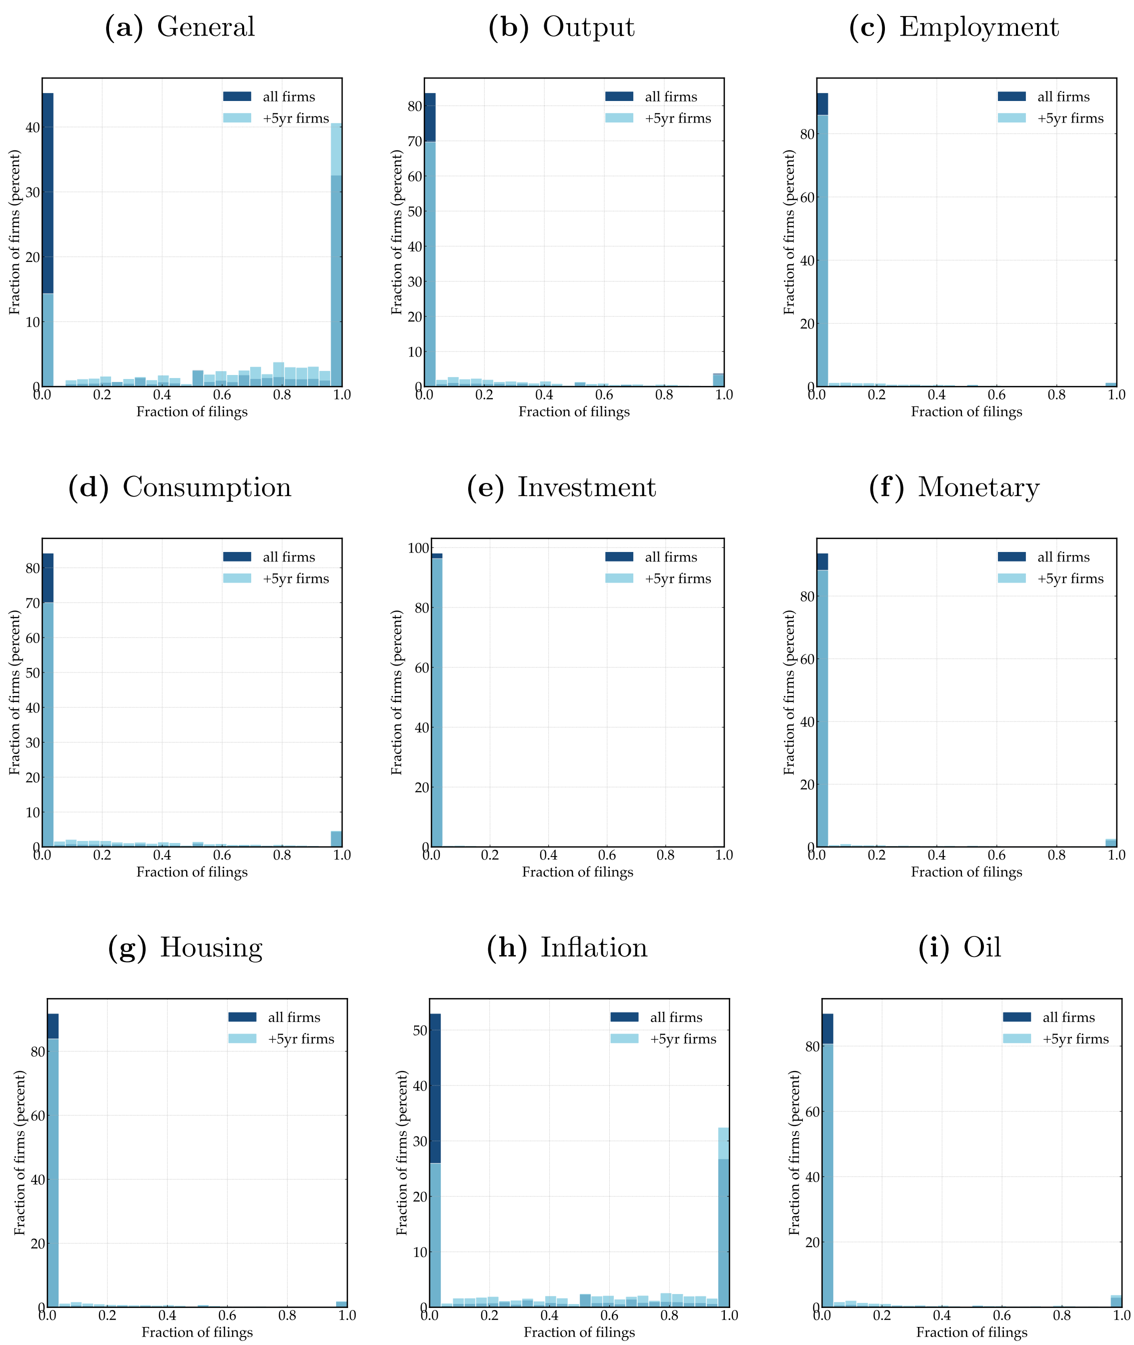

B.2. Additional Robustness and Results

This appendix includes additional robustness tests of our main results in Table 2.5 and additional empirical results.

B.2.1. Controlling for alternative explanations of asymmetry (excluding ZLB)

Table B.3: Controlling for alternative explanations of asymmetry (excl. ZLB)

Control Variable:	**Productivity (LTFP)**	**Mgmt Quality**	**Profit (ROA)**	**Filing Length**
Shock $\times \mathbb{1}_{v_t>0}$	8.54***	2.13	6.01**	3.01
	(2.79)	(2.82)	(2.65)	(5.33)
Shock $\times \mathbb{1}_{v_t<0}$	1.33	-8.05*	-1.52	27.23*
	(3.89)	(4.52)	(3.60)	(13.93)
Attention	-0.09	-0.05	-0.03	-0.04
	(0.10)	(0.07)	(0.06)	(0.06)
Attn $\times$ Shock $\times \mathbb{1}_{v_t>0}$	2.38	2.35***	1.55**	1.51**
	(1.52)	(0.85)	(0.72)	(0.64)
Attn $\times$ Shock $\times \mathbb{1}_{v_t<0}$	-5.85***	-8.21**	-5.75*	-5.09*
	(1.69)	(3.42)	(3.26)	(3.06)
Control Var	0.04***	-0.05	0.89***	-0.00
	(0.01)	(0.07)	(0.16)	(0.03)
Control $\times$ Shock $\times \mathbb{1}_{v_t>0}$	0.02	1.06	-3.70*	0.38
	(0.13)	(0.67)	(1.91)	(0.34)
Control $\times$ Shock $\times \mathbb{1}_{v_t<0}$	-0.14	-3.01	0.54	-2.97*
	(0.20)	(2.57)	(4.59)	(1.54)
Observations	280381	192900	431724	432458
R^2	0.038	0.032	0.027	0.027
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
excl. ZLB	yes	yes	yes	yes
Wald Test p-value: Attention	**0.001**	**0.005**	**0.048**	**0.060**
Wald Test p-value: Control	0.294	0.173	0.460	0.035

Notes: This table augments Column 3 of Table 2.5 to control for four potential confounding sources of asymmetry. The estimated regression is specified in (2.9) as in Table 2.6. Estimates in this table are for the sample up to 2007 to exclude zero-lower-bound periods. (1) firm productivity estimated as in Olley and Pakes (1996), (2) management quality approximated with board member educational attainment, (3) profit measured as earnings before extraordinary items over total assets, and (4) filing length measured as the log word count of the 10-K filing. The final two rows report p-values of Wald tests for $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$ and $H_0 : \beta_{c\nu_+} = \beta_{c\nu_-}$, respectively.

B.2.2. Results robust to controlling for monetary exposure

The theoretical prediction of asymmetry from Section 2.3 confirms the baseline effects in Table 2.5 to be driven by firm attention rather than firm exposure to monetary policy. Nevertheless, we conduct an additional robustness test by directly controlling for firms' exposure to monetary policy.

To measure a firm's exposure to the monetary policy at date τ, we estimate the sensitivity of its stock prices to prior FOMC announcements over a 5-year rolling window using $t \in [\tau - 1826, \tau)$:

$$\begin{aligned}
&\text{Baseline model:} \quad & r_{it} &= \alpha_{i\tau} + \beta_{i\tau}^{\texttt{baseline}} \nu_t + \varepsilon_{it} \\
&\text{CAPM model:} \quad & r_{it} - r_t^f &= \alpha_{i\tau} + \beta_{i\tau}^{\texttt{capm}} \nu_t + \beta_{i\tau}^{M} (r_t^M - r_t^f) + \varepsilon_{it} \\
&\text{FF3 model:} \quad & r_{it} - r_t^f &= \alpha_{i\tau} + \beta_{i\tau}^{\texttt{ff3}} \nu_t + \beta_{i\tau}^{1} (r_t^M - r_t^f) + \beta_{i\tau}^{2} SMB_t + \beta_{i\tau}^{3} HML_t + \varepsilon_{it}
\end{aligned}$$

where ν_t is the high-frequency monetary shock, and r_{it} is the close-to-close returns of firm i at date t. We also estimate sensitivity while controlling for the market factor (r^M) and Fama-French 3 factors (r^M, SML, and HML) using daily data on factors from Kenneth French's website. Exposure is defined as the absolute value of estimated sensitivity,

$$\theta_{i\tau}^{\lambda} = |\hat{\beta}_{i\tau}^{\lambda}| \quad \text{for } \lambda \in \{\text{baseline, CAPM, FF3}\}$$

Table B.4 presents our two interaction coefficients of interest after controlling for exposure, θ_{it}^{λ}. The Wald tests for our null hypothesis, $\beta_{d\nu_+} = \beta_{d\nu_-}$, remains rejected at 5% for all three exposure measures. This confirms that our results are not driven by firms' exposure to monetary policy.

Table B.4: Controlling for exposure to monetary policy

	(1)	(2)	(3)
Shock $\times$ Attn $\times \mathbb{1}_{v_t>0}$	2.03***	2.03***	2.03***
	(0.73)	(0.72)	(0.72)
Shock $\times$ Attn $\times \mathbb{1}_{v_t<0}$	-5.99*	-5.99*	-5.94*
	(3.25)	(3.25)	(3.24)
Observations	572884	571708	568169
R^2	0.026	0.026	0.026
Clustered SE	yes	yes	yes
Firm controls	yes	yes	yes
4-digit NAICS FE	yes	yes	yes
Monetary sensitivity control	baseline model	CAPM model	FF3 model
Wald Test p-value	0.027	0.027	0.027

Notes: Results from estimating the baseline specification (2.8) with additional controls for monetary exposure, θ_{it}^{λ}, $\lambda \in \{\text{baseline, CAPM, FF3}\}$. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

B.2.3. Results not driven by information effect of monetary policy

Nakamura and Steinsson (2018) documents that FOMC announcements release information about the economic fundamentals, in addition to monetary policy. Following Miranda-Agrippino and Ricco (2021), we control for the information effects of monetary policy by including as controls the Greenbook forecast revisions between FOMC meetings. We obtain data on Greenbook forecasts from the Federal Reserve Bank of Philadelphia. Table B.5 show that our main results are robust to controlling for Greenbook forecast revisions.

Table B.5: Controlling for Greenbook forecast revisions

	(1)	(2)	(3)	(4)
Shock $\times$ Attn $\times \mathbb{1}_{v_t>0}$	2.02***	1.88**	1.94***	1.94***
	(0.72)	(0.75)	(0.72)	(0.72)
Shock $\times$ Attn $\times \mathbb{1}_{v_t<0}$	-5.87*	-5.47	-5.71	-5.71
	(3.18)	(3.58)	(3.68)	(3.68)
Observations	575667	575667	575667	575667
R^2	0.026	0.026	0.026	0.026
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
Greenbook rev controls		rgdp	rgdp infl	rgdp infl unemp
Wald Test p-value	0.026	0.070	0.063	0.063

Notes: Results from estimating the baseline specification (2.8) with additional controls for Greenbook forecast revisions. Column (1) displays the baseline results from Table 2.5. Columns (2) - (4) adds Greenbook forecast revisions for real GDP, inflation, and unemployment iteratively. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

B.2.4. Results robust to controlling for macro fluctuations

While high-frequency monetary shocks ν_t are considered exogenous, we conduct additional robustness controlling for business-cycle fluctuations. Macro controls include: lagged real GDP growth, unemployment rate, and inflation, obtained from `FRED`. Column (1) of Table B.6 displays our baseline results without macro controls. Column (2) includes macro controls, controlling for aggregate fluctuations. Column (3) includes macro controls and their interactions with the monetary shock, controlling for differential firm sensitivity to aggregate fluctuations. Column (4) includes macro controls and their separate interactions with expansionary and contractionary monetary shocks, controlling for asymmetric firm sensitivity to aggregate fluctuations. Our main results are robust under all specifications.

Table B.6: Controlling for macroeconomic variables

	(1)	(2)	(3)	(4)
Shock × Attn $\times \mathbb{1}_{v_t>0}$	2.02***	2.06***	1.74**	1.74**
	(0.72)	(0.73)	(0.78)	(0.71)
Shock × Attn $\times \mathbb{1}_{v_t<0}$	-5.87*	-6.27*	-5.38	-7.31**
	(3.18)	(3.21)	(3.34)	(3.31)
Observations	575667	575667	575667	575667
R^2	0.026	0.028	0.028	0.028
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
Macro controls	no	yes	yes	yes
+ interactions	no	no	yes	no
+ asym interactions	no	no	no	yes
Wald Test p-value	0.026	0.021	0.060	0.014

Notes: Results from estimating the baseline specification (2.8) with an additional vector of macro control Z_{t-1}, where Z_{t-1} include lagged real GDP growth, unemployment rate, and inflation. Column (1) displays the baseline results from Table 2.5. Column (2) includes macro controls Z_{t-1},. Column (3) includes Z_{t-1} and $Z_{t-1}\nu_t$. Column (4) includes Z_{t-1} and $Z_{t-1}\nu_t\mathbb{1}_{\nu_t>0}$, and $Z_{t-1}\nu_t\mathbb{1}_{\nu_t<0}$. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

B.2.5. Attention, uncertainty, and business cycles

This section revisits the results presented in Table 2.7 and tests whether attentive firms' better performance under greater uncertainty is driven by business cycle conditions rather than uncertainty itself. To test this hypothesis, we simultaneously interact attention with uncertainty and real GDP growth. For each horizon, $h = 1, \cdots, 5$, the estimating equation takes the form

$$z_{it}^{h} = \alpha_j + \beta_d d_{it} + \beta_\sigma \sigma_t + \beta_y y_t + \beta_{d\sigma} d_{it} \sigma_t + \beta_{dy} d_{it} y_t + \Gamma' Z_{it} + \varepsilon_{it}, \tag{B.1}$$

where the dependent variables considered are profitability (ROA), financial performance (ROE), and an indicator variable for firm survival (as in Section 2.5). Attention, d_{it}, is defined as the prevalence measure for general economic conditions; macroeconomic uncertainty, σ_t, is defined as the interquartile range of quarterly growth rate forecasts for real GDP and unemployment from the Survey of Professional Forecasters; and business cycle conditions, y_t, are measured using annual real GDP growth. The model also controls for industry fixed effects, δ_j, and a vector of firm controls, Z_{it}.

Table B.7 shows that the findings in Table 2.7 are robust to conditioning on business cycle conditions. Columns labeled *Impact* report estimates for outcomes in the following year, $h = 1$, and those labeled *Peak* report estimates for the same horizons reported under *Peak* in Table 2.7. The first four columns show that the estimated effect of attention under uncertainty on financial performance slightly strengthens after controlling for real GDP growth. The estimated interaction effect between attention and real GDP growth in the same columns is quite imprecisely but suggests that attentive firms may fare slightly better under good business cycle conditions. Overall, these findings support the hypothesis that attentive firms perform better relative to their inattentive peers because of their ability to navigate uncertainty.

Table B.7: Attention, uncertainty, and business cycles

	ROE		**ROA**		**Survival**	
	Impact	Peak	Impact	Peak	Impact	Peak
Attention (general)	-0.02	-0.03	-0.02	-0.04	-0.03	-0.00
	(0.02)	(0.02)	(0.04)	(0.03)	(0.02)	(0.01)
Uncertainty (SPF IQR)	-0.02	-0.02	-0.04	-0.04	-0.03	-0.02**
	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.01)
rGDP growth	2.00	0.14	1.77	-0.36	-2.46*	0.82
	(1.59)	(1.36)	(2.35)	(2.09)	(1.28)	(1.14)
Attention × Uncertainty	0.03*	0.04**	0.06*	0.06**	0.02	0.02
	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.01)
Attention × rGDP growth	0.51	0.57	0.43	0.85	1.52	-1.17
	(1.44)	(1.25)	(2.20)	(1.91)	(1.22)	(0.76)
Observations	104507	92023	110267	97180	111637	66813
R^2	0.165	0.156	0.248	0.236	0.034	0.028
Clustered SE	yes	yes	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes	yes	yes

Notes: The table reports results from estimating

$$z_{it}^h = \alpha_j + \beta_d d_{it} + \beta_\sigma \sigma_t + \beta_y y_t + \beta_{d\sigma} d_{it}\sigma_t + \beta_{dy} d_{it} y_t + \Gamma' Z_{it} + \varepsilon_{it}, \tag{B.2}$$

for horizons $h = 1, \cdots, 5$. The dependent variables z_t include (i) profitability measured with ROA (i.e., net income over total assets), (ii) financial performance measured with ROE (i.e., net income over equity), and (iii) an indicator variable for firm survival. Independent variables include the prevalence attention to general economic conditions, d_{it}; macroeconomic uncertainty, σ_t, measured as the interquartile range of quarterly growth rate forecasts for real GDP and unemployment from the Survey of Professional Forecasters; real GDP growth, y_t; interaction terms; industry fixed effects δ_j; and firm controls, Z_{it}. We standardize the interquartile range of each series over our observed sample period, take the absolute average deviation each quarter, and then average these quarterly values each year. The on-impact effect corresponds to the estimates for $h = 1$. The peak effect horizons are the same as those in Table 2.7. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

B.3. Additional Results from Textual Analysis

This appendix contains a set of additional results using natural language processing to investigate the context in which firms discuss macro keywords in 10-K filings and to provide further validation of the text-based measures.

B.3.1. Lexical similarity

Our measure of lexical similarity is a Jaccard score, $J(y_{it}, y_{it-1})$, which measures the share of unique non-stop words that appear between the current year's 10-K (y_i) compared to the previous year's 10-K (y_{it-1}).

$$J(y_i, y_{it-1}) = \frac{|y_i \cap y_{it-1}|}{|y_i \cup y_{it-1}|}$$

The Jaccard score is bounded by the unit interval, and is decreasing with the "uniqueness" of the text. Figure B.3 reports the average Jaccard score for each section of 10-K filings.

Figure B.3: Lexical similarity by section of 10-K filings

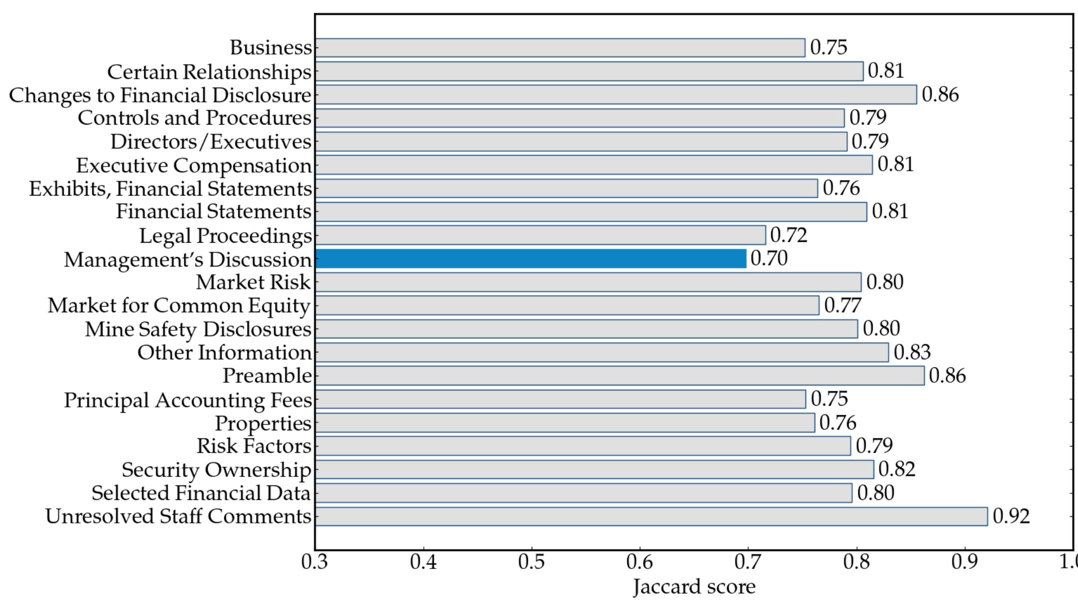

Notes: Average Jaccard scores for sections in 10-K filings. The Jaccard score is bounded by the unit interval. A high Jaccard score represents high lexical similarity between filings. The Management's Discussion section has the lowest level of lexical similarity in all 10-K sections.

We then restrict the attention measures to keywords mentioned in low Jaccard score sections: Business (Item 1) and Managment's Discussion (Item 7). We exclude Legal Proceedings

Table B.8: Restricting attention to low lexical similarity 10-K sections

	(1) Average	(2) Exposure	(3) Attention	(4) excl ZLB
Shock	5.62***	4.13*		
	(1.22)	(2.42)		
Attention		-0.03	-0.08	-0.05
		(0.04)	(0.05)	(0.05)
Shock × Attn		0.02		
		(0.45)		
Shock × $\mathbb{1}_{v_t>0}$			4.55*	6.21**
			(2.65)	(2.66)
Shock × $\mathbb{1}_{v_t<0}$			-4.16	-1.45
			(3.72)	(3.69)
Shock × Attn ×$\mathbb{1}_{v_t>0}$			0.79	0.50
			(0.56)	(0.54)
Shock × Attn ×$\mathbb{1}_{v_t<0}$			-5.24**	-4.95*
			(2.48)	(2.53)
Observations	546596	546596	546596	409889
R^2	0.018	0.023	0.026	0.027
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
excl. ZLB	no	no	no	yes
Wald Test p-value			0.030	0.058

Notes: Results from variants of estimating the baseline specification in (2.8), restricting to 10-K items that discuss firm operations (Items 1 and 7). Standard errors are in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

(Item 3) that has a low Jaccard score to avoid false positives from legal languages. Regression results with attention restricted to low lexical similarity 10-K sections are reported in Table B.8.

B.3.2. LDA: context of macro discussions

To enable automated context detection, we use the Latent Dirichlet Allocation (LDA) model to uncover topics firms tend to discuss in conjunction with macro news. LDA (Blei et al., 2003) is an unsupervised learning algorithm aimed at grouping words in documents into meaningful topics. We apply LDA to texts in earning filings within 20 words surrounding a macroeconomic keyword and set the number of topics to be 10.

Following Hansen et al. (2018), we pre-process texts of 10-K filings for LDA as follows:

we remove numbers and words that are only one character. Then we lemmatize to combine different word forms (for example, "operated" and "operates" are lemmatized to "operate"). The advantage of lemmatizing over stemming is that the resulting LDA outputs are more friendly to interpret. Our corpus include words and bigrams which appear for at least 20 times. We filter out words that occur in less than 20 documents or more than 50% of the documents. Then we transform the texts through bag-of-words representation.

We model topics surrounding each of the nine macro categories for the attention measure, as well as an aggregate category containing keywords from all categories. Figures B.4 and B.5 visualize the LDA output surrounding keywords in all categories.Figure B.4 shows the heat map of LDA outputs. Each row represent a topic clustered by LDA, and the darkness of the cell within a topic represent the likelihood of a word to appear in the topic. Figure B.5 highlights the word cloud of selected topics in B.4.

Although LDA output does not label topics, it is natural to characterize some of the topics. Topic 1 relates to business operations, as firms discuss how macro conditions feed into into their daily operations; Topic 2 relates to demand, as firms track and gauge the aggregate demand; Topic 6 relate to financing costs, as firms pay attention to how monetary policy affect their financial costs, investment decisions, and portfolio holdings; Topic 10 relates to labor costs, as firms assess the tightness of the labor market. Rest of the topics relate to housing, currency, and risk factors.

Figure B.4: LDA output for texts surrounding all macro keywords

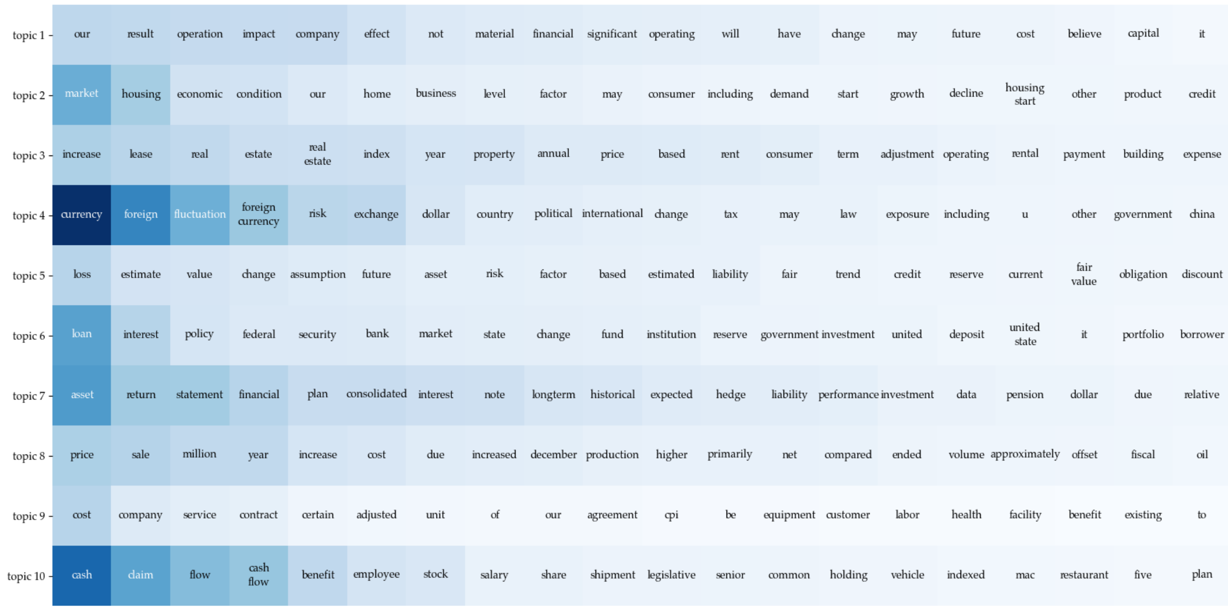

Figure B.5: LDA output for texts surrounding all macro keywords: Selected topics

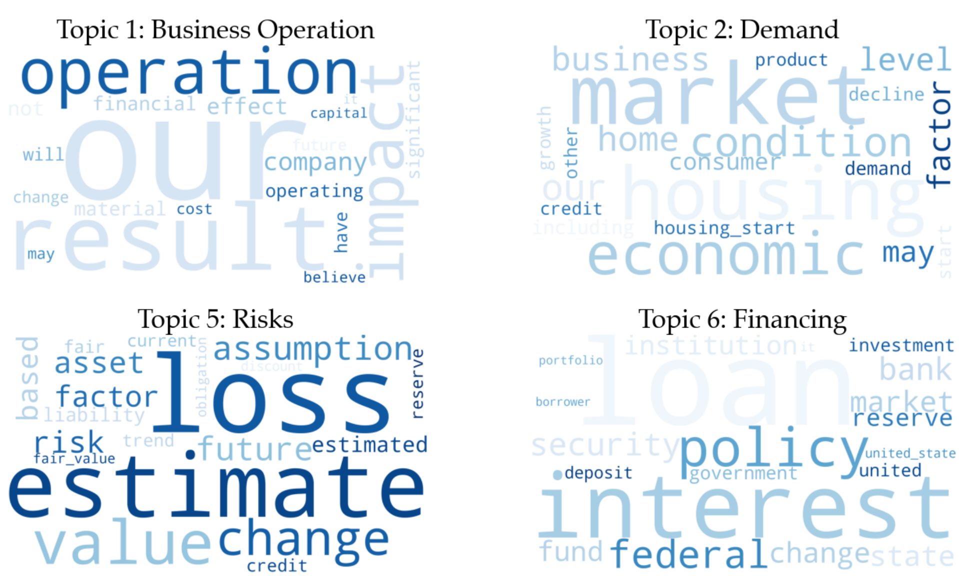

B.3.3. Itemized frequency search

10-K filings have standard formats and are organized in sections. We perform refined frequency counts for each of the section, or "items", to see where attention is concentrated in. Results of frequency counts of macroeconomic keywords by filing item are shown in Figure B.6 in the Appendix. Discussions of the macroeconomy are concentrated in Description of Business (Item 1), Risk Factors (Item 1A) and Management Discussion and Analysis of Financial Condition and Results of Operations (Item 7A).

Figure B.6: Firm attention by filing items

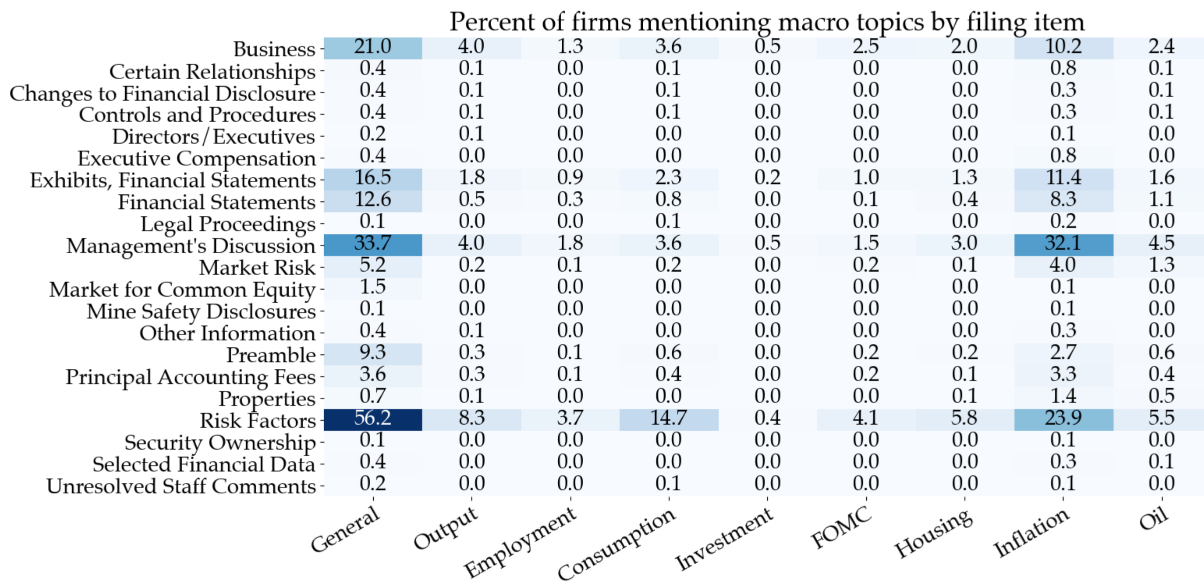

Notes: Heat map of firm attention by filing items. Each row represents a section ("item") of 10-K, and each column represents a macroeconomic topic. Darkness represents a higher fraction of firms that pay attention to a macroeconomic topic in an item.

Results in Figure B.6 show that firms pay attention to macro news to assess the impact on their business operations and risks, consistent with assumptions that firms mentioning a macroeconomic topic do so in order to incorporate the news into their decision making.

B.4. Additional Details for the Stylized Model

B.4.1. Approximation of firm profits in the stylized model

Under second-order approximation around the non-stochastic steady state, the log approximation of a firm's profits, denoted by $\hat{\pi}(s_t, a_t)$, is given by:

$$\begin{aligned}\hat{\pi}(s_t, a_t) &= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \pi_a(\bar{s}, \bar{a})\bar{a}\hat{a}_t + \frac{1}{2}\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{a}_t^2 + \pi_{sa}(\bar{s}, \bar{a})\bar{s}\bar{a}\hat{s}_t\hat{a}_t \\ &= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \frac{1}{2}\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{a}_t^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}\bar{s}\hat{a}_t\hat{s}_t \\ &= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \frac{1}{2}\left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2(\hat{a}_t - \hat{s}_t)^2\end{aligned}$$

In the second line, $\pi_a(\bar{s}, \bar{a}) = 0$ because of optimal choice. In addition, the assumption that $a = s$ under full information yields $\pi_a(a, a) = 0\ \forall a$, which implies $\pi_{sa}(\bar{s}, \bar{a}) = -\pi_{aa}(\bar{s}, \bar{a})$. The third line added and subtracted $\frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{s}_t^2$ to complete squares and used the fact that $\bar{a} = \bar{s}$ in the steady state. The resulting expression is equation (2.5).

B.4.2. Proof of Proposition 1

Proof. We consider the responses of returns to an aggregate shock ε. Holding all else equal, that is, $\pi_{ss}^k(s, a) = \pi_{ss}(s, a)$ and $\pi_{aa}^k(s, a) = \pi_{aa}(s, a)$ for all firms k, we can show the following for heterogeneity in exposure and in attention.

(i) **Exposure**: Let firms be heterogeneous in exposure and homogeneous in attention. Specifically, suppose firm i is more exposed to macro conditions than firm j, that is, $\pi_s^i > \pi_s^j > 0$. We consider how heterogeneity in exposure affects return elasticity for cases in which both firms are attentive and both are inattentive.

  (a) Case 1 (both firms attentive): When firms are both attentive, $\hat{a}_t = \hat{s}_t$. Then by equation (2.5) we can derive the return elasticity with respect to the aggregate shock to be:

$$\frac{\partial r_k}{\partial \varepsilon} = \frac{\partial \hat{\pi}_k}{\partial \varepsilon} = \pi_s^k(\bar{s}, \bar{a})\bar{s} + \left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\varepsilon \quad \text{for firm } k = i, j.$$

  Therefore, the return elasticity for firms i is larger for the return elasticity for firm j for all magnitudes of shocks

$$\frac{\partial r_i}{\partial \varepsilon} - \frac{\partial r_j}{\partial \varepsilon} = \pi_s^i(\bar{s}, \bar{a})\bar{s} - \pi_s^j(\bar{s}, \bar{a})\bar{s} > 0$$

because $\pi_s^i > \pi_s^j > 0$.

(b) Case 2 (both firms inattentive): When both firms are inattentive, the return elasticity with respect to the shock can be expressed as:

$$\frac{\partial r_k}{\partial \varepsilon} = \frac{\partial \hat{\pi}_k}{\partial \varepsilon} = \pi_s^k(\bar{s},\bar{a})\bar{s} + \left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\varepsilon + \pi_{aa}(\bar{s},\bar{a})\bar{a}^2(f_k(\varepsilon)-\varepsilon)(f_k'(\varepsilon)-1) \quad \text{for firm } k = i, j.$$

Since firms are only heterogeneous in exposure, the second and third term in the above expression for return elasticity is the same for both firms. Therefore:

$$\frac{\partial r_i}{\partial \varepsilon} - \frac{\partial r_j}{\partial \varepsilon} = \pi_s^i(\bar{s},\bar{a})\bar{s} - \pi_s^j(\bar{s},\bar{a})\bar{s} > 0$$

which is also independent of the magnitude of ε.

(ii) **Attention**: Now instead let firms be heterogeneous in attention and homogeneous in exposure, so the attentive firm i has $f_i'(\varepsilon) = 1$, the inattentive firm j has $f_j'(\varepsilon) < 1$, and both firms have $\pi_s^i = \pi_s^j$. The return elasticity for attentive and inattentive firms can be expressed as:

$$\frac{\partial r_i}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} + \left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\varepsilon \tag{B.3}$$

$$\frac{\partial r_j}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} + \left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\varepsilon + \pi_{aa}(\bar{s},\bar{a})\bar{a}^2(f_j(\varepsilon)-\varepsilon)(f_j'(\varepsilon)-1) \tag{B.4}$$

since firms are homogeneous in exposure: $\pi_s^i = \pi_s^j = \pi_s$. The relative magnitude of return elasticities between attentive and inattentive firms depends on the sign of the shock ε. Specifically, we consider three cases.

(a) Zero shock ($\varepsilon = 0$): Since $f(0) = 0$, (B.3) and (B.4) lead to:

$$\frac{\partial r_i}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} = \frac{\partial r_j}{\partial \varepsilon}$$

(b) Positive shock ($\varepsilon > 0$): Since $\varepsilon_t > f_j(\varepsilon_t) > 0$,

$$\frac{\partial r_j}{\partial \varepsilon} - \frac{\partial r_i}{\partial \varepsilon} = \underbrace{\pi_{aa}(\bar{s},\bar{a})\bar{a}^2}_{<0}\underbrace{(f_j(\varepsilon)-\varepsilon)}_{<0}\underbrace{(f_j'(\varepsilon)-1)}_{<0} < 0$$

(c) Negative shock ($\varepsilon < 0$) Since $\varepsilon_t < f_j(\varepsilon_t) < 0$,

$$\frac{\partial r_j}{\partial \varepsilon} - \frac{\partial r_i}{\partial \varepsilon} = \underbrace{\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2}_{<0} \underbrace{(f_j(\varepsilon) - \varepsilon)}_{>0} \underbrace{(f_j'(\varepsilon) - 1)}_{<0} > 0$$

■

B.4.3. Model with time-varying uncertainty

We provide a framework to illustrate the effects of attention on firm profits when macroeconomic uncertainty is time-varying. This gives rise to the test we perform in Table 2.7.

Environment The aggregate state variable, y_t, follows an autoregressive process:

$$y_t = \rho y_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \nu_t^2),$$

where ν_t^2 denotes the volatility of the aggregate state and is time-varying.

The firm's objective is to track the state variable as closely as possible and set its prices, x_t, accordingly. The loss function is given by

$$\mathcal{L} = (x_t - y_t)^2.$$

To track the macroeconomy, the firm chooses a noisy signal centered around the true state:

$$s_t = y_t + u_t, \quad u_t \sim N(0, \tau_t^2).$$

Following Sims (2003), the level of noise contained in the signal implies the flow of information of

$$\kappa_t = \frac{1}{2} \log_2 \frac{\tau_t^2 + \nu_t^2}{\tau_t^2}.$$

Firms are constrained by its cognitive bandwidth to process information

$$\kappa_t \leq \kappa.$$

Optimization Given this set up, the price a firm sets given the signal is

$$\begin{aligned} x_t = \mathbb{E}[y_t|y_{t-1}, s_t] &= \frac{\tau_t^2}{\tau_t^2 + \nu_t^2}\rho y_{t-1} + \frac{\nu_t^2}{\tau_t^2 + \nu_t^2} \\ &= 2^{-2\kappa_t}\rho y_{t-1} + (1 - 2^{-2\kappa_t})(y_t + u_t). \end{aligned}$$

A firm chooses the level of information to obtain in order to minimize the expected loss,

$$\begin{aligned} \min_{\kappa_t} \mathbb{E}\,\mathcal{L} &= \mathbb{E}[2^{-2\kappa_t}\varepsilon_t - (1 - 2^{-2\kappa_t})u_t]^2 \\ &= \mathrm{Var}\,[2^{-2\kappa_t}\varepsilon_t - (1 - 2^{-2\kappa_t})u_t] \\ &= 2^{-2\kappa_t}\nu_t^2, \end{aligned}$$

subject to its bandwidth constraint

$$\kappa_t \leq \kappa.$$

Therefore, the firm's realized loss is

$$\mathcal{L} = 2^{-2\kappa}\nu_t^2.$$

As the economy becomes more uncertain, a firm's loss increases:

$$\frac{\partial \mathcal{L}}{\partial \nu_t^2} = 2^{-2\kappa} > 0.$$

However, attentive firms suffer a smaller loss compared to inattentive firms as uncertainty rises:

$$\frac{\partial \frac{\partial \mathcal{L}}{\partial \nu_t^2}}{\partial \kappa} = -(2\log 2)2^{-2\kappa} < 0.$$

B.5. Additional Details for the Quantitative Model

B.5.1. Approximation of firms' value function

A firms' value function for its operating profits can be expressed as

$$
\begin{aligned}
V^{op} &= \max \sum_{t=0}^{\infty} \beta^t \, \mathbb{E}\left[\Pi(P_{it}, P_t, Q_t) | S_i^{-1}\right] \\
&= \max \sum_{t=0}^{\infty} \beta^t \, \mathbb{E}\left[\frac{\Pi(P_{it}, P_t, Q_t)}{\Pi^*(P_{it}^*, P_t, Q_t)} \Pi^*(P_{it}^*, P_t, Q_t) | S_i^{-1}\right] \\
&= \max \sum_{t=0}^{\infty} \beta^t \Pi^*(P_{it}^*, P_t, Q_t) \, \mathbb{E}\left[L(P_{it}, P_{it}^*, P_t, Q_t) | S_i^{-1}\right]
\end{aligned}
\tag{B.5}
$$

where $\Pi(P_{it}, P_t, Q_t)$ denotes the firm's operating profits, and $L(P_{it}, P_{it}^*, P_t, Q_t) \equiv \frac{\Pi(P_{it},P_t,Q_t)}{\Pi^*(P_{it}^*,P_t,Q_t)}$ denotes the loss from imperfect information relative to full-information profits $\Pi^*(P_{it}^*, P_t, Q_t)$. The last equality follows the fact that L is homogeneous of degree 1.

Under the second-order log approximation around the non-stochastic steady state, we can express the loss as:

$$
\begin{aligned}
&L(P_{it}, P_{it}^*, P_t, Q_t) \equiv \frac{\Pi(P_{it}, P_t, Q_t)}{\Pi^*(P_{it}^*, P_t, Q_t)} \\
&\approx \bar{L} + p_{it}\bar{P}\bar{L}_1 + p_{it}^*\bar{P}\bar{L}_2 + p_t\bar{P}\bar{L}_3 + q_t\bar{Q}\bar{L}_4 + \frac{1}{2}p_{it}^2\bar{P}^2\bar{L}_{11} + \frac{1}{2}p_{it}^{*2}\bar{P}^2\bar{L}_{22} + \frac{1}{2}p_t^2\bar{P}^2\bar{L}_{33} + \frac{1}{2}q_t^2\bar{Q}^2\bar{L}_{44} \\
&\quad + p_{it}p_{it}^*\bar{P}^2\bar{L}_{12} + p_{it}p_t\bar{P}^2\bar{L}_{13} + p_{it}q_t\bar{P}\bar{Q}\bar{L}_{14} + p_{it}^*p_t\bar{P}^2\bar{L}_{23} + p_{it}^*q_t\bar{P}\bar{Q}\bar{L}_{24} + p_tq_t\bar{P}\bar{Q}\bar{L}_{34} \\
&= \frac{\bar{\Pi}}{\bar{\Pi}} + p_{it}\bar{P}\frac{\bar{\Pi}_1}{\bar{\Pi}} - p_{it}^*\bar{P}\frac{\bar{\Pi}_1}{\bar{\Pi}} + p_t\bar{P}\cdot 0 + q_t\bar{Q}\cdot 0 + \frac{1}{2}p_{it}^2\bar{P}^2\frac{\bar{\Pi}_{11}}{\bar{\Pi}} - \frac{1}{2}p_{it}^{*2}\bar{P}^2(\frac{\bar{\Pi}_{11}}{\bar{\Pi}} - \frac{2\bar{\Pi}_1^{\,2}}{\bar{\Pi}^2}) \\
&\quad + \frac{1}{2}p_t^2\bar{P}^2\frac{2\bar{\Pi}_2 - 2\bar{\Pi}_2^2}{\bar{\Pi}^2} + \frac{1}{2}q_t^2\bar{Q}^2\frac{2\bar{\Pi}_3 - 2\bar{\Pi}_3^2}{\bar{\Pi}^2} - p_{it}p_{it}^*\bar{P}^2\frac{\bar{\Pi}_1^2}{\bar{\Pi}^2} + p_{it}p_t\bar{P}^2(\frac{\bar{\Pi}_{12}}{\bar{\Pi}} - \frac{\bar{\Pi}_1\bar{\Pi}_2}{\bar{\Pi}^2}) \\
&\quad + p_{it}q_t\bar{P}\bar{Q}(\frac{\bar{\Pi}_{13}}{\bar{\Pi}} - \frac{\bar{\Pi}_1\bar{\Pi}_3}{\bar{\Pi}^2}) + p_{it}^*p_t\bar{P}^2(\frac{\bar{\Pi}_1\bar{\Pi}_2}{\bar{\Pi}^2} - \frac{\bar{\Pi}_{12}}{\bar{\Pi}}) + p_{it}^*q_t\bar{P}\bar{Q}(\frac{\bar{\Pi}_1\bar{\Pi}_3}{\bar{\Pi}^2} - \frac{\bar{\Pi}_{13}}{\bar{\Pi}}) + p_tq_t\bar{P}\bar{Q}\cdot 0 \\
&= \frac{1}{2}(p_{it}^2 - p_{it}^{*2})\bar{P}^2\frac{\bar{\Pi}_{11}}{\bar{\Pi}} + (p_{it} - p_{it}^*)p_t\bar{P}^2\frac{\bar{\Pi}_{12}}{\bar{\Pi}} + (p_{it} - p_{it}^*)q_t\bar{P}\bar{Q}\frac{\bar{\Pi}_{13}}{\bar{\Pi}} + \text{ terms independent of } p_{it} \\
&= \frac{1}{2}\bar{P}^2\frac{\bar{\Pi}_{11}}{\bar{\Pi}}(p_{it} - p_{it}^*)^2 + \text{ terms independent of } p_{it},
\end{aligned}
\tag{B.6}
$$

where lowercase letters denote log deviations from the steady state, $\bar{L}$ is the short hand for $L(\bar{P}, \bar{P}, \bar{P}, \bar{Q})$, and $\bar{\Pi}$ is the short hand for $\Pi(\bar{P}, \bar{P}, \bar{Q})$. The first two (approximate) equalities are second-order log approximations. The third equality uses the fact that $\Pi_1 = 0$ from optimal choices. In addition, $\Pi_1(P_{it}^*, P_t, Q_t) = 0$ implies $p_{it}^*\bar{P}\Pi_{11}(\bar{P}, \bar{P}, \bar{Q}) + p_t\bar{P}\Pi_{12}(\bar{P}, \bar{P}, \bar{Q}) +$

$q_{it}\bar{Q}\Pi_{13}(\bar{P},\bar{P},\bar{Q}) = 0$, which leads to the last equality.

Therefore, a firm's problem under second-order log approximation is given by

$$\begin{aligned} \max_{\{s_{it}\in\mathcal{S}_{it}, p_{it}(S_i^t)\}_{t\geq 0}} \ & \mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\Big(-B(p_{it}-p_{it}^*)^2 - 2\omega_i\mathcal{I}(p_t^*; s_{it}|S_i^{t-1})\Big)\Big|S_i^{-1}\right] \\ \text{s.t.} \quad & p_t^* = \alpha p_t + (1-\alpha)q_t \\ & S_i^t = S_i^{t-1}\cup s_{it}, \end{aligned}$$

where

$$B \equiv -\frac{1}{2}\bar{P}^2\frac{\bar{\Pi}_{11}}{\bar{\Pi}} = -\frac{1}{2}\big(\varepsilon(\varepsilon-1) - \frac{\psi\varepsilon(\varepsilon+\gamma)}{\gamma^2}\eta^{-\frac{1}{(1-\alpha)\gamma}}\big)(1-\psi\eta^{-\frac{1}{(1-\alpha)\gamma}})^{-1} > 0. \tag{B.7}$$

B.5.2. Optimal price under full information

Under full information, a firm's operating profit in period t is given by

$$\Pi(P_{it},P_t,Q_t) = \frac{1}{Q_t}(P_{it}Y_{it} - W_tN_{it}) = P_{it}^{1-\varepsilon}P_t^{\varepsilon-1} - \psi Q_t^{\frac{1}{\gamma}}P_{it}^{-\frac{\varepsilon}{\gamma}}P_t^{\frac{\varepsilon-1}{\gamma}}.$$

The first-order condition with respect to P_{it}:

$$(\varepsilon-1)P_{it}^{-\varepsilon}P_t^{\varepsilon-1} + \frac{\varepsilon\psi}{\gamma}Q_t^{\frac{1}{\gamma}}P_{it}^{-\frac{\varepsilon}{\gamma}-1}P_t^{\frac{\varepsilon-1}{\gamma}} = 0$$

implies that

$$P_{it}^* = \left(\frac{\varepsilon\gamma}{(\varepsilon-1)\gamma}\right)^{\frac{\gamma}{-\varepsilon\gamma+\varepsilon+\gamma}} P_t^{\frac{-\varepsilon\gamma+\varepsilon+\gamma-1}{-\varepsilon\gamma+\varepsilon+\gamma}} Q_t^{\frac{1}{-\varepsilon\gamma+\varepsilon+\gamma}}.$$

Log linearization around the full-information nonstochastic steady state yields

$$p_{it}^* = \alpha p_t + (1-\alpha)q_t,$$

where

$$\alpha = \frac{-\varepsilon\gamma+\varepsilon+\gamma-1}{-\varepsilon\gamma+\varepsilon+\gamma}. \tag{B.8}$$

B.5.3. Numerical solution

We solve the rational inattention problem based on the `DRIPs` algorithm developed by Afrouzi and Yang (2021a). Under log-quadratic approximation, the firm's problem is given by

$$\begin{aligned} p_{it}^* &= \alpha p_t + (1-\alpha) q_t, \\ \Delta q_t &= \rho \Delta q_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2) \\ p_t &= \int_0^1 p_{it} di, \end{aligned}$$

where $\alpha = \frac{-\varepsilon\gamma+\varepsilon+\gamma-1}{-\varepsilon\gamma+\varepsilon+\gamma}$. Differencing out the unit root allows us to obtain the Wold representation of p_{it}^* as

$$p_{it}^* = (1-L)\Phi(L)\tilde{\nu}_t, \quad \tilde{\nu}_t = (1-L)^{-1}\nu_t = \sum_{j=0}^{\infty} \nu_{t-j},$$

where $\Phi(L)$ is the lag operator.

We specify the length of truncation to be $L = 40$ and define $x_t = (\tilde{\nu}_t, \tilde{\nu}_{t-1}, \cdots, \tilde{\nu}_{t-(L+1)})$. Then, the state-space representation of the system is given by

$$\begin{aligned} x_t &= \mathbf{A} x_{t-1} + \mathbf{Q} \nu_t \\ q_t &= \mathbf{H}_q' x_t \\ p_{it}^* &= \mathbf{H}' x_t, \end{aligned}$$

where

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_\nu \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \text{ and } \quad \mathbf{H}_q = \begin{bmatrix} 1 \\ \rho \\ \rho^2 \\ \vdots \\ \rho^{L-1} \end{bmatrix}.$$

To solve for $\mathbf{H}$, we proceed as follows. In the n-th iteration, we start with the guess from the previous iteration, $\mathbf{H}_{k,(n-1)}$, where the subscript $k \in \{l, h\}$ indexes firms with low and high marginal costs of information. Then, we solve the rational inattention problem and obtain an updated guess. The optimal price is given by

$$p_t^* = \alpha p_t + (1-\alpha) q_t = (1-\alpha) \sum_{j=0}^{\infty} \alpha^j q_t^{(j)} = (1-\alpha) \sum_{j=0}^{\infty} \alpha^j (\theta q_{lt}^{(j)} + (1-\theta) q_{ht}^{(j)}),$$

where $q_{kt}^{(j)}$ is the j-th order belief of type-k firms on average, and $k \in \{l, h\}$. Now we guess and verify the expression for $q_{kt}^{(j)}$. Suppose there exists a matrix $\mathbf{X_{kj}}$ such that $q_{kt}^{(j)} = \mathbf{H}_q'\mathbf{X}_{kj}x_t$. Then, we can solve $q_t^{(j+1)}$ forward as

$$\begin{aligned} q_t^{(j+1)} &= \int_\theta \mathbb{E}_{it,l}\, q_{lt}^{(j)} di + \int_{(1-\theta)} \mathbb{E}_{it,h}\, q_{ht}^{(j)} di \\ &= \mathbf{H}_q'\big(\theta \mathbf{X}_{lj}\mathbf{X}_{l,(n)}x_t + (1-\theta)\mathbf{X}_{hj}\mathbf{X}_{h,(n)}\big)x_t, \end{aligned}$$

where $\mathbf{X}_{k,(n)} = \sum_{j=0}^{\infty}((\mathbf{I} - \mathbf{K}_{k,(n)}\mathbf{Y}_{k,(n)}')\mathbf{A})^j\mathbf{K}_{k,(n)}\mathbf{Y}_{k,(n)}'\mathbf{M}'^j$. Matrices $\mathbf{K}$ and $\mathbf{Y}$ are Kalman gains and loadings of optimal signals solved from the rational inattention problem, as specified in Afrouzi and Yang (2021a); $\mathbf{M}$ is a shift matrix.

Setting $\mathbf{X}_{kj} = \mathbf{X}_{k,(n)}^j$ for all j implies that

$$q_{kt}^{(j)} = \mathbf{H}_q'\mathbf{X}_{k,(n)}^j x_t,$$

which verifies the guess for $q_{kt}^{(j)}$.

We, therefore, obtain the updated guess $\mathbf{H}_{k,(n)} = (1-\alpha)\mathbf{X}_{k,p,(n)}'\mathbf{H}_q$, where $\mathbf{X}_{k,p,(n)} = \sum_{j=0}^{\infty} \alpha^j \mathbf{X}_{(n)}^j$, and iterate until convergence.

B.5.4. Details for model calibration

This section provides additional details on model calibration.

B.5.4.1. Parameter values

Table B.9 summarizes parameter values for the quantitative model as described in Section 2.6.2.

Table B.9: Calibration

Parameter	Value	Source or Targeted Moment
Standard parameters		
Discount rate (β)	$0.96^{1/4}$	Quarterly frequency
Shock persistence (ρ)	0.89	Estimates from US nominal output 1994–2019
Shock standard deviation (σ_ν)	0.034	Estimates from US nominal output 1994–2019
Elasticity of substitution (ε)	10	Steady-state markup of 11%
Disutility of labor (ψ)	0.90	Offset steady-state monopolistifc inefficiency
Returns to scale (γ)	0.93	Basu and Fernald (1997)
Information parameters		
Low cost of information (ω_L)	1×10^{-6}	Assigned to near zero
Fraction of attentive firms (θ)	65%	Average share of attentive firms (Figure 2.3)
Relative costs of information ($\omega_H - \omega_L$)	1.11	Heterogenous responses to monetary shocks (Table 2.5)

B.5.4.2. Parameter identification plots

Figure B.7 provides details on how calibrated parameters are identified. In Panel (a), we simulate the model for a range of ω_H. As ω_H increases and the gap between ω_L and ω_H widens, the simulated elasticity monotonically increases, implying greater heterogeneity between attentive and inattentive firms. In Panel (b) we report the sensitivity of ω_L and ω_H separately. As ω_H increases, the magnitude of both semi-elasticities increases monotonically.

Figure B.7: Sensitivity of simulated moments to calibrated parameters

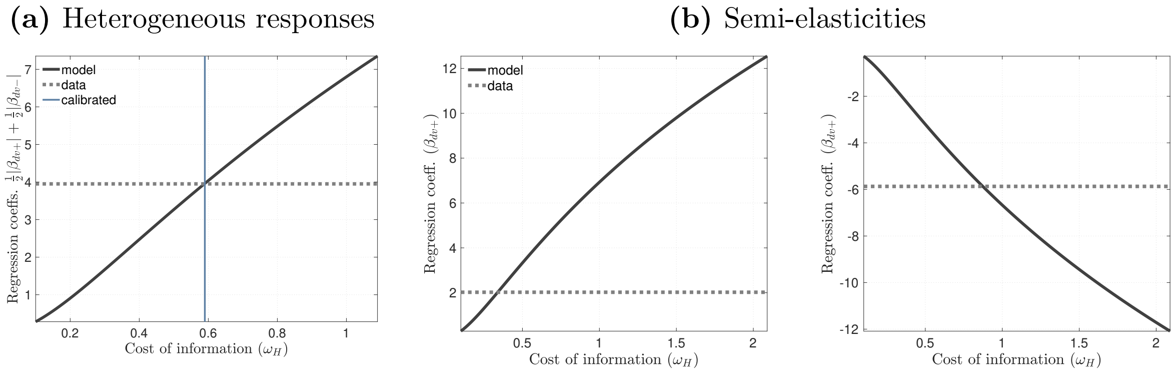

Notes: Simulated moments for a range of parametrization for labor share γ and cost of information ω_H. We simulate models for a panel of 100 firms and for 1000 periods with 100 period burn-ins. Simulated moments are generated with regressions discussed in the text. Panel (a) shows the sensitivity of average absolute values of $\beta_{d\nu_+}$ and $\beta_{d\nu_-}$ to changes in parameter values of ω_H; Panel (b) shows the sensitivity of $\beta_{d\nu_+}$ and $\beta_{d\nu_-}$ separately.

B.5.4.3. Model fit

Table B.10 presents data and model moments. In Column 1, we present average and marginal semi-elasticities to monetary shocks observed in the data. In Column 2, we present targeted semi-elasticities generated by our model compared with their empirical counterparts. We calibrate $\omega_H - \omega_L$ to match $\frac{1}{2}|\beta_{d\nu_+}| + \frac{1}{2}|\beta_{d\nu_-}|$; the other elasticities are untargeted.

Table B.10: Empirical and model moments

	Data	**Model**	
Average moment			
β_v	5.61	4.73	untargeted
	(1.21)		
Marginal moments			
$\frac{1}{2}\lvert\beta_{dv+}\rvert + \frac{1}{2}\lvert\beta_{dv-}\rvert$	3.95	3.95	targeted
β_{dv+}	2.02	4.04	untargeted
	(0.72)		
β_{dv-}	-5.87	-3.86	untargeted
	(3.18)		

Notes: Data moments correspond to estimates in Columns (1) and (3) from Table 2.5, with standard errors in parentheses. Model moments are generated with corresponding regressions with simulated data for a panel of 100 firms and for 1000 periods with 100 periods burn-ins.

B.5.5. Alternative calibration with price adjustment

In this section, we provide an alternative calibration strategy for the quantitative model. We calibrate the model to match the speed of price adjustment rather than return elasticities from Section 2.4, as in the baseline case.

Table B.11: Calibration with price adjustment

Parameter	Description	Value
Assigned parameters		
β	discount rate	$0.96^{1/4}$
ρ	shock persistence	0.89
σ_ν	shock std. dev.	4.23×10^{-2}
ε	elasticity of substitution	7
ψ	disutility of labor	0.86
γ	returns to scale	0.83
Information parameters		
θ	fraction of attentive firms	65%
ω_L	cost of information	58
ω_H	cost of information	53×10^3

We calibrate the model quarterly. The unit of analysis, i, represents a 4-digit NAICS sub-industry within manufacturing.[1] Table B.11 shows the parameter values. Assigned parameters that are unrelated to information frictions follow the baseline calibration in Section 2.6.2 with two exceptions. We set the elasticity of substitution $\varepsilon = 7$ to capture a lower elasticity of substitution across industries than across firms, following Gorodnichenko and Weber (2016). We set returns to scale $\gamma = 0.83$ according to the estimate by Basu and Fernald (1997) for the US private economy.

We then calibrate costs of information, ω_L and ω_H, to target industry inflation responses to monetary shocks. To obtain empirical targets, we re-estimate Equation (2.1) for manufacturing sectors at quarterly frequency using

$$\Delta \log P_{s,t} = \alpha_s + \alpha_t + \beta_\nu \nu_t^{\mathrm{M}} + \beta_d d_{st} + \beta_{d\nu} d_{st} \nu_t^{\mathrm{M}} + \Gamma' Z_t + \varepsilon_{st},$$

where $P_{s,t}$ is the PPI of industry s (4-digit NAICS) in quarter t; ν_t is monetary shocks; d_{st} is sector average prevalence attention; Z_t and $\{\alpha_s, \alpha_t\}$ are our standard controls and fixed effects, respectively.

[1] For example, NAICS 3331 represents “agricultural, construction and mining machinery manufacturing,” and NAICS 3332 represents “industrial machinery manufacturing.”

Table B.12: Targeted moments

Targeted Moment	Data	Model
attentive price semi-elasticity to monetary shocks ($\hat{\beta}_\nu + \hat{\beta}_{d\nu}$)	8.79	8.79
inattentive price semi-elasticity to monetary shocks ($\hat{\beta}_\nu$)	3.41	3.41

Table B.12 shows that in response to a 100-basis point expansionary monetary shock, attentive sectors raise prices by 8.8% in the first quarter, while inattentive sectors raise prices by 3.4%.

To match these empirical targets, we simulate the model for a range of ω values and obtain the impulse responses to a monetary shock equivalent to a 100-basis point rate cut. We set ω_L so that the simulated inflation of attentive industries in response to a monetary shock matches $\hat{\beta}_\nu + \hat{\beta}_{d\nu}$, the observed semi-elasticity of attentive industries. Similarly, we set ω_H so that the simulated inflation responses of inattentive industries match $\hat{\beta}_\nu$, the observed semi-elasticity of inattentive industries. Table B.12 shows that the calibrated model matches the semi-elasticities of industry inflation in response to monetary shocks.

As in the baseline model, we quantify the importance of attention on the efficacy of monetary policy by varying the share of attentive firms. Table B.13 shows the response of output growth is 23 basis points (or 14%) weaker in the most attentive calibration compared to the least attentive calibration, which shows a quantitative importance consistent with the baseline calibration.

Table B.13: Attention and monetary non-neutrality

	Least attentive	Baseline	Most attentive
Fraction of attentive firms (θ)	56%	65%	73%
Output response	1.65%	1.52%	1.42%

Notes: Dependence of output responses on the share of attentive firms in the economy. Output responses are calculated as a percent deviation from the steady state in response to a monetary shock equivalent to a 25-basis point rate cut. Calibration for the least and most attentive economies is described in Section 2.6.3 in the main text.

B.5.6. Firm impulse responses

Figure B.8 shows the impulse responses of individual firms to monetary shocks of one standard deviation. Panel (a) shows that as nominal aggregate demand rises, inattentive firms under-adjust prices, reflecting partial incorporation of noisy signals about demand. Attentive firms track aggregate demand better than inattentive firms and exhibit more responsive prices. Because of the imprecise information about the aggregate demand, inattentive firms experience greater losses in flow profits from the full-information benchmark in response to both expansionary and contractionary shocks. With a constant marginal cost of information, firms' equilibrium choice of attention is not time-varying. Even though inattentive firms pay less attention, the higher marginal costs they face result in higher total information costs.

Panel (b) shows the responses of stock returns. Following an expansionary monetary shock, full-information equity returns of both attentive and inattentive firms increase since firms are monopolistically competitive and have decreasing returns to scale. Returns of attentive firms increase by more than those of inattentive firms because attentive firms track the optimal price more closely. Returns of an imperfect-information firm are lower than those of a full-information firm that sets the optimal price. Following a contractionary shock, returns of attentive firms drop by less than those of inattentive firms.

Figure B.8: Firm impulse responses to monetary shocks

(a) Firm prices and operating profits

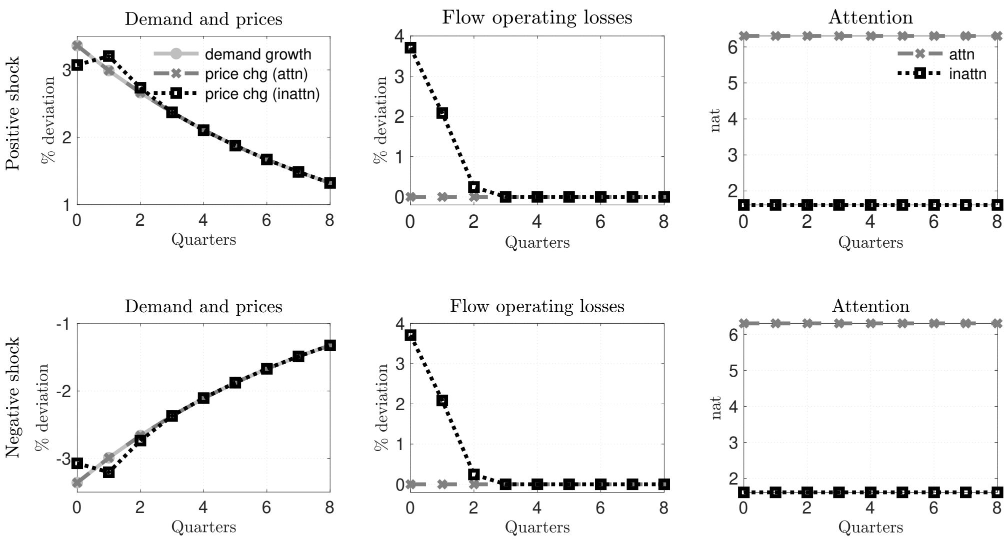

(b) Conditional realized returns

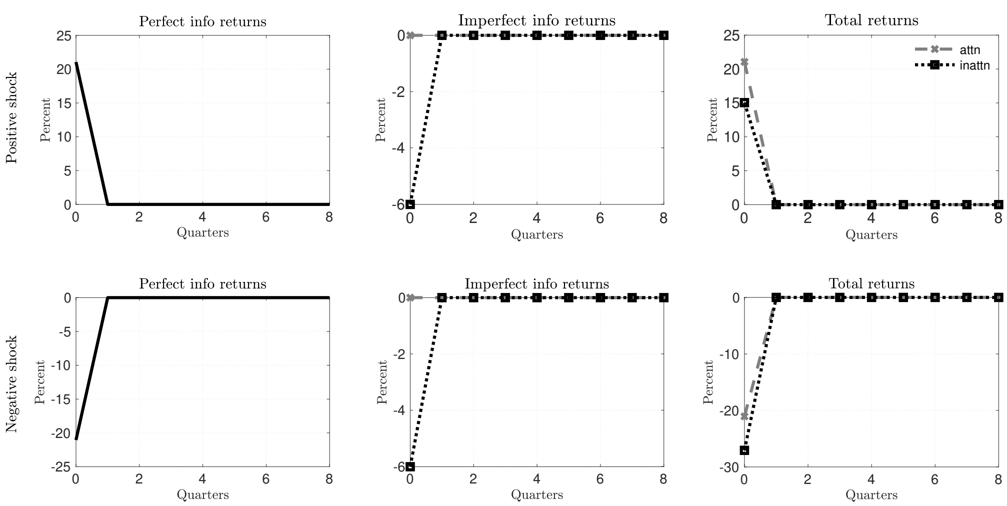

Notes: Firm impulse responses to a one standard deviation positive (expansionary) monetary shock and negative (contractionary) shock. Impulse responses are in percent deviations from the perfect-information steady state.

B.5.7. Passthrough regressions

The passthrough of nominal interest rate change to nominal demand change is estimated with local projections (Jordà, 2005). We estimate the following model for horizons $h = 1, 2, \cdots, 20$:

$$\Delta_h y_{t-1,t+h} = \alpha_h + \beta_h \varepsilon_t^{\mathtt{i}} + u_{th}$$

where $\Delta_h y_{t-1,t+h}$ is average percent changes in the variable of interest, and $\varepsilon_t^{\mathtt{i}}$ is the high-frequency monetary policy shock. The dependent variables are U.S. manufacturing output over the sample period of 1994 to 2019. We estimate the responses of manufacturing prices, real output and nominal output. Path of β_h informs the average cumulative changes in the dependent variable in response to the interest rate shock.

Figure B.9: Passthrough of rates to nominal demand

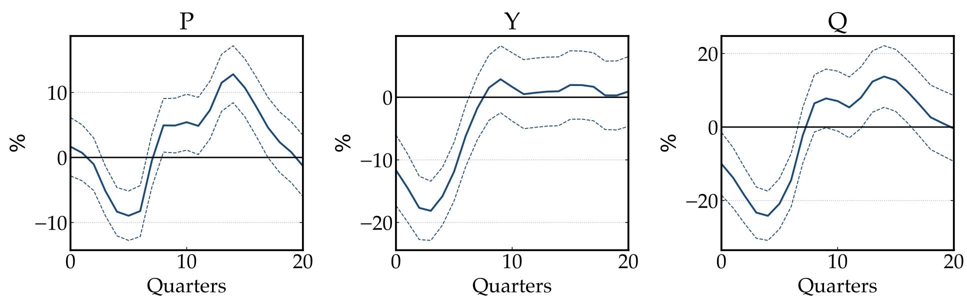

Figure B.9 shows the results of the local projection. A 25 basis point point expansionary shock to the interest rate leads to about 3.3 percent peak increase in nominal demand.

B.5.8. Aggregate responses to rate cuts

Figure B.10: Aggregate responses to expansionary monetary shock

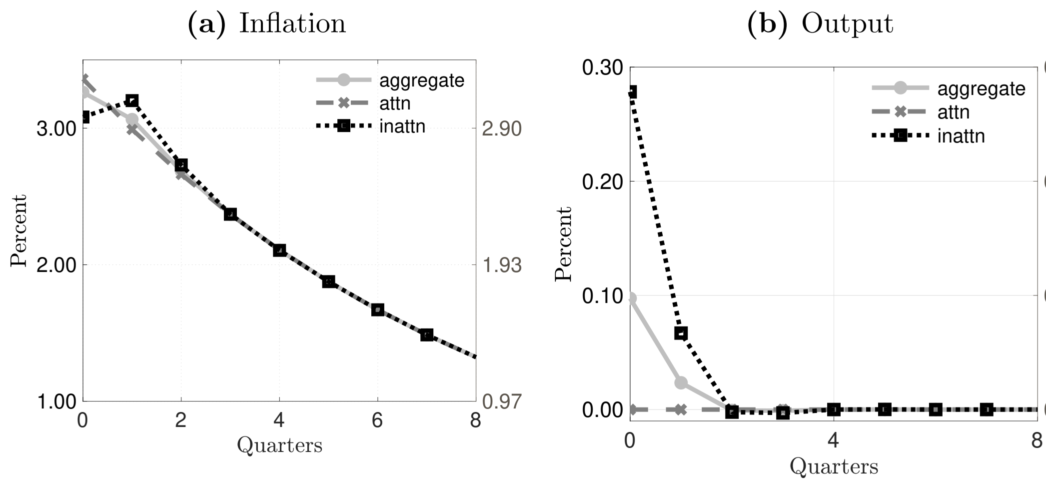

Notes: Impulse responses of inflation and output. The left scales show the impulse responses to a one standard deviation expansionary monetary shock, and the right scales show the impulse responses to an equivalent of 25 basis point expansionary monetary policy shock. Impulse responses are in percent deviations from the perfect-information steady state. "attn" refers to the impulse responses of attentive firms, "inattn" refers to the impulse responses of inattentive firms, and "aggregate" refers to the aggregate impulse responses.

APPENDIX C

Appendix to Chapter 3

C.1. Additional Tables and Figures

Table C.1: College Enrollment and Local Labor Market Conditions

(a) Total Enrollment Response by Institution Type

	Public			Private non-profit			For-profit		
	4-year	2-year	1-year	4-year	2-year	1-year	4-year	2-year	1-year
Employment rate	-2.5***	-6.6***	-0.1**	-0.8***	-0.1***	-0.0	-0.5***	-0.7***	-0.6***
	(0.5)	(1.3)	(0.1)	(0.2)	(0.0)	(0.0)	(0.2)	(0.1)	(0.1)
Observations	11,459	11,459	9,715	11,459	11,459	9,715	11,459	11,459	9,715
R^2	0.01	0.06	0.00	0.00	0.00	0.00	0.01	0.03	0.03

(b) Age 25+ Enrollment Response by Institution Type

	Public			Private non-profit			For-profit		
	4-year	2-year	1-year	4-year	2-year	1-year	4-year	2-year	1-year
Employment rate	-1.4***	-4.5***	-0.1	-0.6***	-0.0	-0.0**	-0.4***	-0.4***	-0.3***
	(0.3)	(0.9)	(0.0)	(0.2)	(0.0)	(0.0)	(0.1)	(0.1)	(0.0)
Observations	8,984	8,984	7,910	8,984	8,984	8,271	8,984	8,984	8,630
R^2	0.01	0.07	0.00	0.01	0.00	0.00	0.01	0.04	0.04

Notes: This table reports the results of several univariate regressions of annual MSA enrollment rates (in basis points) on the local employment rate (in percent). Standard errors are in parenthesis and coefficients should be interpreted as the basis point change enrollment associated with a one percentage point increase in the local employment rate. All rates are detrended using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. *Panel A* reports estimates for total enrollment by institution program duration and control. Schools that exclusively offer graduate programs are counted among four-year programs. *Panel B* reports estimates for enrollment among students ages 25 or older by institution program duration and control. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using adult working-age population (ages 18-65) from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER). Enrollment at one-year institutions are not available before 1993, and enrollment by student age is only available on a biannual basis between 1991 and 1998.

Figure C.1: Binscatter of Detrended Enrollment and Employment (%)

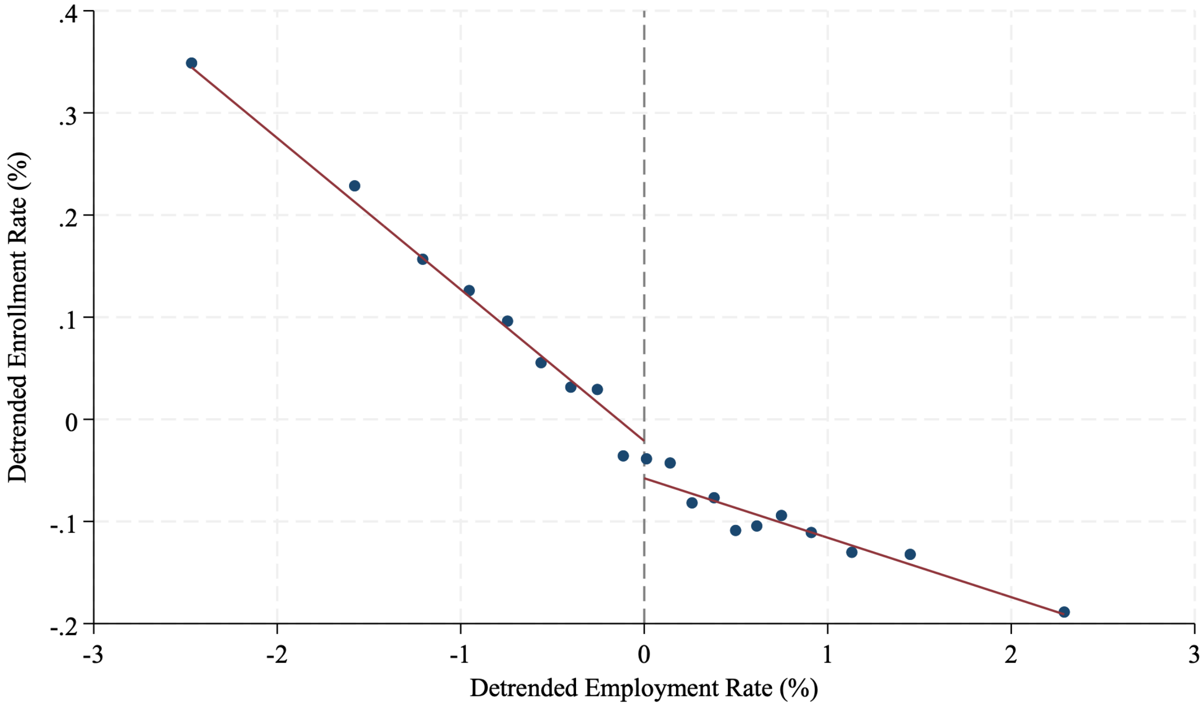

Notes: This figure plots the binscatter of detrended annual enrollment and employment rates and the best fit line with a discontinuity at zero. The employment rate is constructed using adult working-age population (ages 18-65), while enrollment rates are constructed using the respective population of each age group. All rates are detrended using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. Employment data are from the Quarterly Census of Employment and Wages (QCEW), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data are from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Table C.2: Total Enrollment Responsiveness by Employment Conditions

	Change in Employment		Relative to Trend	
	Positive	Negative	Above	Below
Employment rate	-10.5***	-13.9***	-5.8***	-14.8***
	(2.1)	(2.4)	(1.2)	(3.3)
Observations	7,174	4,285	6,072	5,387
R^2	0.07	0.11	0.01	0.06

Notes: This table reports the results of univariate regressions of the total enrollment rate (in basis points) on the local employment rate (in percent). Standard errors are in parenthesis and coefficients should be interpreted as the basis point change enrollment associated with a one percentage point increase in the local employment rate. All rates are detrended using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1987 and 2019. The first two columns report enrollment responsiveness when splitting the sample by the change in employment since the previous year. The last two columns do the same after splitting the sample by whether employment was above or below trend. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using adult working-age population (ages 18-65) from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure C.2: Changes in Employment and Enrollment, 2007-2010

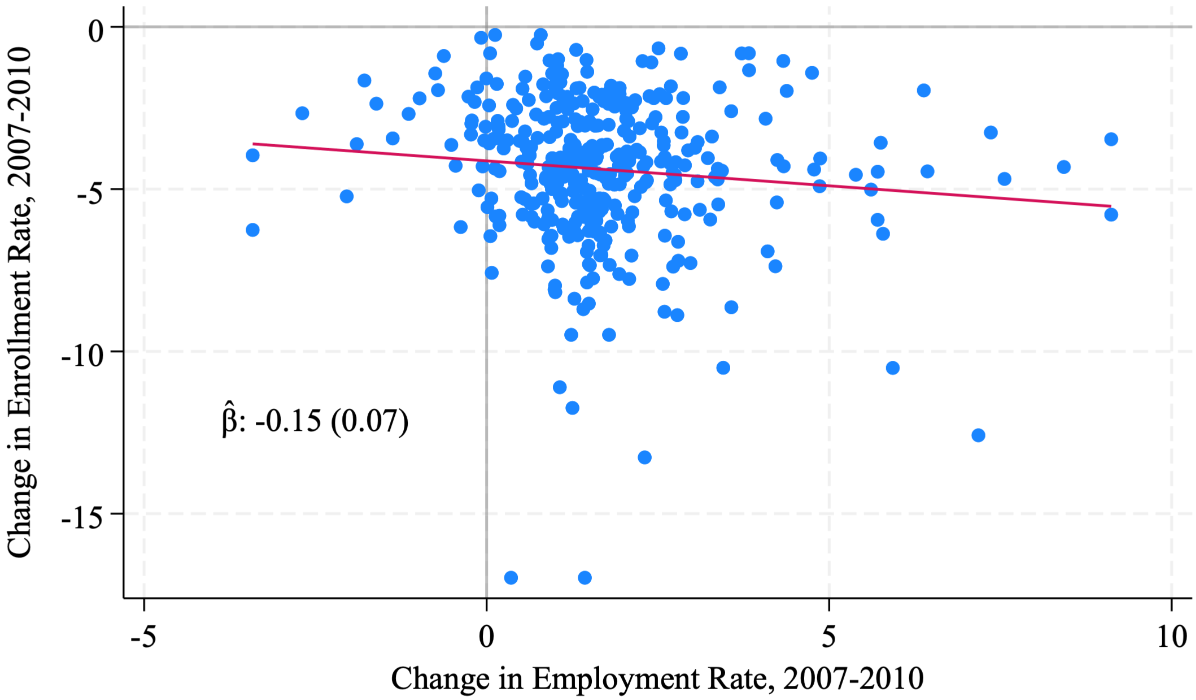

Notes: This figure plots the percentage point change in MSA employment and enrollment rates between 2007 and 2010. The figure includes a linear best-fit line and reports the estimated coefficient and standard error for that line. Employment data are from the Quarterly Census of Employment and Wages (QCEW) and enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Table C.3: Enrollment Rates and Changes in Enrollment by Institution Control, 2007-2010

	(1) Δ Public	(2) Δ Private	(3) Δ For-Profit
Public Enrollment	0.04***	-0.00	-0.00
	(0.0)	(0.0)	(0.0)
Private Enrollment	0.01	0.05***	0.00
	(0.0)	(0.0)	(0.0)
For-Profit Enrollment	-0.11	-0.04	0.25***
	(0.1)	(0.0)	(0.0)
Constant	0.01***	0.00	0.00***
	(0.0)	(0.0)	(0.0)
Observations	370	370	370
R^2	0.08	0.25	0.23

Notes: This table reports estimates from the model, $\Delta s^c_{l,2010} = \delta_0 + \sum_{c=1}^{C} \delta_c s^c_l + \varepsilon_l$, where $\Delta s^c_{l,2010}$ is the percentage point change in the enrollment rate at institution type, c, for MSA, l, between 2007 and 2010; and s^c_l is the enrollment rate at institution type, c, in MSA, l, in 2007. Enrollment type is defined according to institution control: public, private non-profit, and private for-profit. Enrollment data are from the Integrated Postsecondary Education Data System (IPEDS) and all rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure C.3: Enrollment Responsiveness by Age-Specific Employment Rate

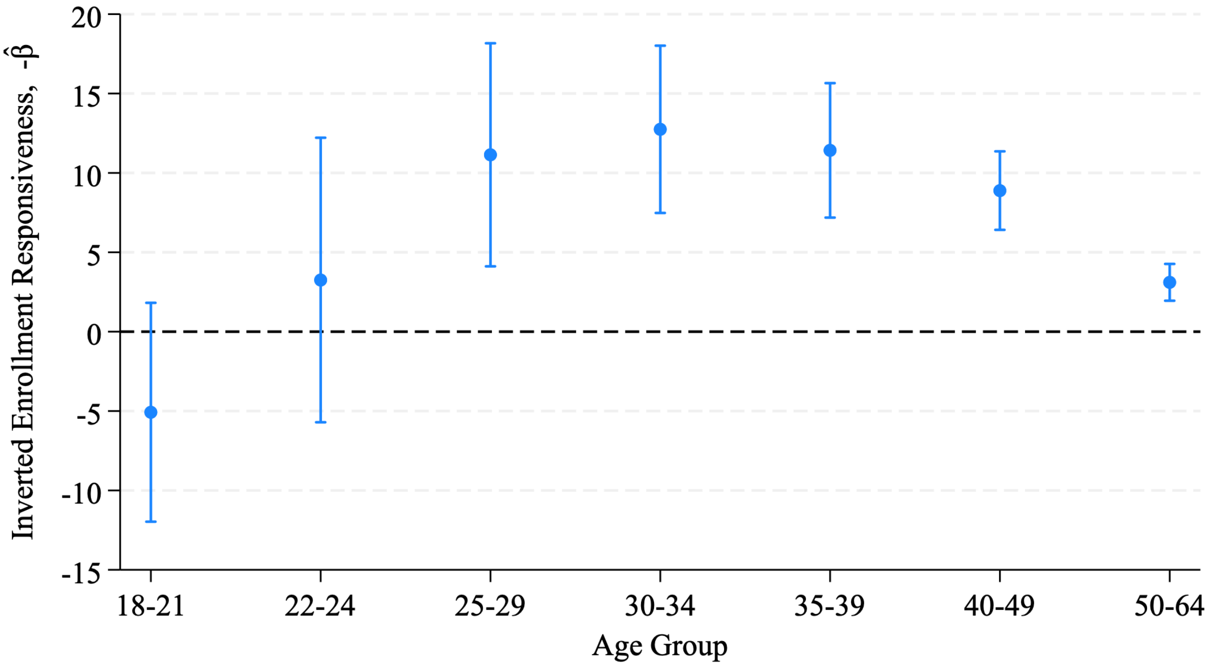

Notes: This figure plots the inverted point estimates and two standard error bars for a series of univariate regressions of annual MSA enrollment rates on the local employment rate by student age. Estimates should be interpreted as the basis point change in the enrollment rate associated with a one percentage point increase in the local employment rate. The employment and enrollment rate are both constructed within the corresponding age group, and all rates are detrended by MSA using an HP-filter with smoothing factor $\lambda = 100$. The sample covers 380 MSAs between 1990 and 2019. Employment data are from the Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS), and population data are from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Table C.4: First Stage Estimates

	(1) Total	(2) 25-34	(3) 35-44	(4) 45-54	(5) 55-64	(6) Alt Rates
$\Delta z_{l,2010}$	0.57***	0.44***	0.43***	0.46***	0.66***	0.44**
	(0.11)	(0.11)	(0.10)	(0.10)	(0.10)	(0.20)
Share no college	0.03*	0.03**	0.03*	0.03*	0.02	0.03**
	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Employment Rate	0.02**	0.02***	0.02***	0.02***	0.01	0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Enrollment Rate	0.02	0.04***	0.13***	0.15***	-0.03	0.03**
	(0.01)	(0.01)	(0.04)	(0.04)	(0.09)	(0.02)
Population Growth						0.08**
						(0.03)
Constant	-0.02*	-0.02**	-0.02**	-0.02**	-0.01	-0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	370	370	370	370	370	370
R^2	0.15	0.17	0.18	0.17	0.14	0.17
Instrument t-stat	5.26	3.98	4.15	4.47	6.88	2.26

Notes: This table reports first-stage estimates for the IV analysis presented in Section 3.3. The first stage model takes the form, $\Delta s_{l,2010} = \pi_0 + \pi_1 \Delta z_{l,2010} + X_l \Pi + \nu_l$, where $\Delta s_{l,2010}$ represents the change in the local enrollment rate between 2007 and 2010; $\Delta z_{l,2010}$ is the instrument for changes in local enrollment rates; and X_l is a vector of start-of-period controls including the share of employees without any college experience, the employment rate, and the enrollment rate. Column 1 reports first-stage estimates for the main results in Figure 3.5, columns 2-5 present estimates by age cohort for Figure 3.6, and column 6 reports first-stage estimates for Figure 3.10, which uses employment and enrollment rates based on start-of-period population and controls for population growth separately. The t-statistic for each instrument coefficient is reported in the last row the table. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

Figure C.4: Effect of New Enrollment on Employment by Age (OLS Estimates)

(a) Ages 19-24

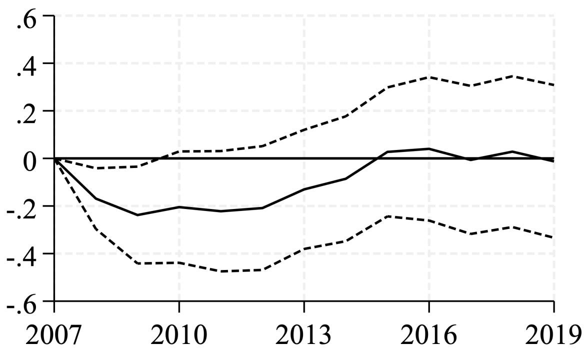

(b) Ages 25-34

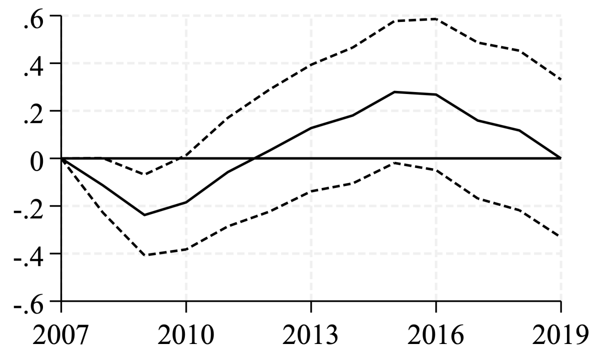

(c) Ages 35-44

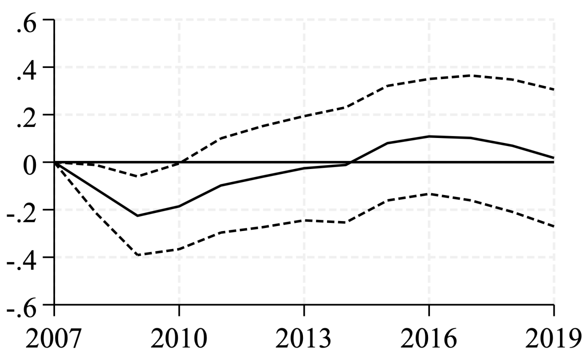

(d) Ages 45-54

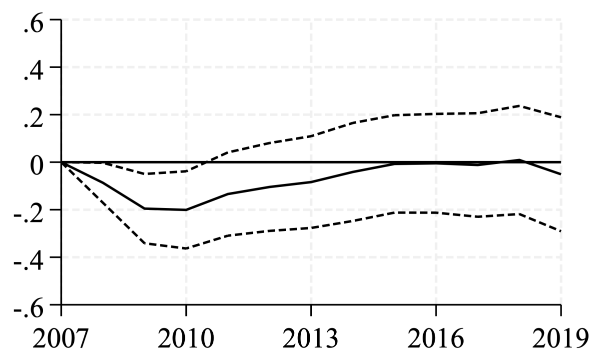

(e) Ages 55-64

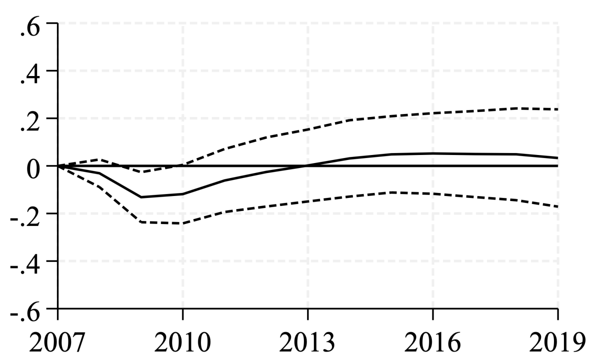

Notes: This figure plots OLS estimates for δ from Equation 3.2, $\Delta e_{l,t} = \alpha + \delta \Delta s_{l,2010} + \mathbf{X_l \Gamma} + \nu_l$, for years 2008-2019 and five separate ages cohorts. $\Delta e_{l,t}$ represents the change in the local employment rate for each age cohort in MSA, l, between 2007 and year t; $\Delta s_{l,2010}$ represents the change in the local total enrollment rate between 2007 and 2010; and $\mathbf{X_l}$ is a vector of start-of-period controls including the share of employees without any college experience, the age-specific employment rate, and the age-specific enrollment rate. Coefficients should be interpreted as the percentage point change in the employment rate between 2007 and year t associated with a one percentage point increase in the enrollment rate between 2007 and 2010. Employment data are from the US Census' Quarterly Workforce Indicators (QWI), enrollment data are from the Integrated Postsecondary Education Data System (IPEDS). All rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

C.2. Additional Methods and Details

C.2.1. Approximating the share of adults who enrolled in college

Enrollment rates offer an imprecise approximation for the share of new adults who sought higher education during the Great Recession. Annual enrollment only captures the flow of students in a given year and adding enrollment over multiple years leads to double-counting.[1] This section constructs an alternative approximation for the share of new adults who enrolled in higher education to test whether unadjusted changes in enrollment rates offer an acceptable approximation. My analysis focuses on the period from 2007 to 2010, which covers a surge in new enrollment during the Great Recession.

This share is constructed by differencing annual enrollment rates with their pre-recession levels, adjusting for the average years that students attended, and summing over the sample period. Average attendance is measured separately for different schools types using the Department of Education's Beginning Postsecondary Students Longitudinal Study (BPS). See Table C.5 for average years of attendance by school type and student age.

Using 2007 as the base year, new enrollment in MSA l and year t is approximated as,

$$\tilde{S}_{lt} = \sum_{k=1}^{K} \frac{S_{lt} - S_{l,2007}}{\gamma_k},$$

where S_{lt} represents total annual enrollment and γ_k is the average years of attendance for school type k. Total new students is approximated by summing $\tilde{S}_{lt}$ for all years between 2007 and final year, T. The share of adults enrolled is then constructed by dividing by population. Figure C.5 plots the raw change in enrollment between 2007 and 2010 used in the analysis above with the alternative approximation for total new enrollment constructed in this section. It shows that these two measures are strongly correlated and move approximately one-for-one with each other.

C.2.2. Further Details on the Enrollment Surge, 2007-2010

Average proximity to community college campus is measured as the population-weighted geodesic distance between each 2000 Census block centroid and the list of community college campuses from the U.S. Department of Education Campus Safety and Security database. This database lists the locations of separate campuses and branches that are otherwise

[1] IPEDS does report first-time, first-year students for the entire student body. However, this measure does not account for workers who are returning to school (ie not first-time students) and does not distinguish student age, which is of interest here.

Table C.5: Average Years of Attendance

Institution Type	Ages 18-24	Ages 25+	Total
Public, 4-year	3.1	2.2	3.0
Private, 4-year	2.9	2.1	2.8
For-profit, 4-year	2.1	2.0	2.0
Public, 2-year	2.2	1.5	1.9
Private, 2-year	2.0	2.0	2.0
For-profit, 2-year	1.5	1.4	1.5
Public, 1-year	1.1	1.1	1.1
Private, 1-year	1.6	1.2	1.4
For-profit, 1-year	1.2	1.1	1.1

Notes: This table reports the average number of years of higher education attendance by institution type and student age. Data are from the Department of Education's Beginning Postsecondary Students Longitudinal Study (BPS) survey for the cohort of students who first enrolled in 2004.

consolidated under the main campus in IPEDS. College affordability is measured using the average net price students paid to attend community college within each MSA in 2007. Net price is reported in IPEDS and measured as tuition plus fees minus average grant funding that students received. IPEDS defines graduation rates as the share of students who completed their college program within one and a half times the normal period of time for that program.

Figure C.5: Comparison of Raw and Adjusted Enrollment Rates

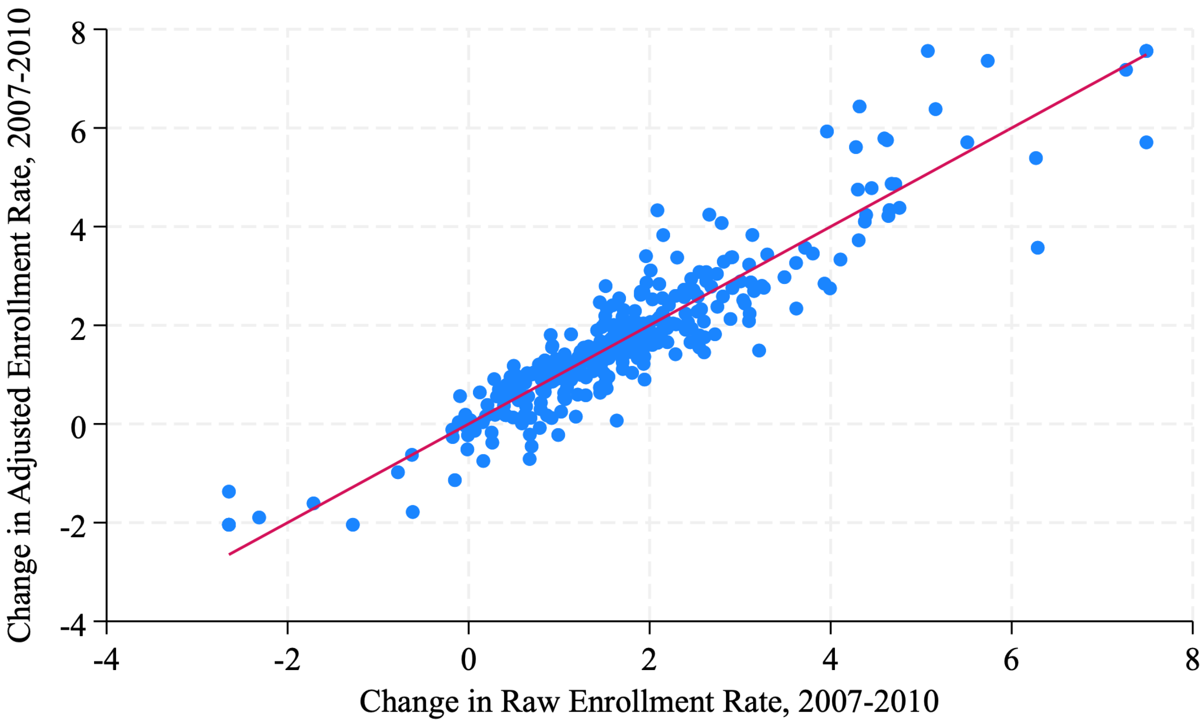

Notes: This figure compares the change in raw and adjusted enrollment rates between 2007 and 2010. The change in adjusted enrollment is constructed by Enrollment data are from the Integrated Postsecondary Education Data System (IPEDS) and rates are constructed using population data from the National Institutes of Health Surveillance, Epidemiology, and End Results Program (SEER).

BIBLIOGRAPHY

Abis, Simona and Laura Veldkamp, "The Changing Economics of Knowledge Production," *The Review of Financial Studies*, 08 2023, p. hhad059.

Acharya, Viral V, Matteo Crosignani, Tim Eisert, and Christian Eufinger, "Zombie Credit and (Dis-)Inflation: Evidence from Europe," 2020.

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer, "Identification properties of recent production function estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Acton, Riley K, "Community College Program Choices in the Wake of Local Job Losses," 2019.

Afrouzi, Hassan, "Strategic Inattention, Inflation Dynamics, and the Non-neutrality of Money," *Manuscript*, 2020.

_ **and Choongryul Yang**, "Dynamic Rational Inattention and the Phillips Curve," 2021.

_ **and** _ , "Selection in Information Acquisition and Monetary Non-neutrality," 2021.

Ahnert, Toni, Sebastian Doerr, Nicola Pierri, and Yannick Timmer, "Does IT Help? Information Technology in Banking and Entrepreneurship," Technical Report, International Monetary Fund 2021.

Albrizio, Silvia, Tomasz Kozluk, and Vera Zipperer, "Environmental policies and productivity growth: Evidence across industries and firms," *Journal of Environmental Economics and Management*, 2017, *81*, 209–226.

Alpay, Ebru, Joe Kerkvliet, and Steven Buccola, "Productivity growth and environmental regulation in Mexican and US food manufacturing," *American journal of agricultural economics*, 2002, *84* (4), 887–901.

Amirault, David, Naveen Rai, and Laurent Martin, "A Reference Guide for the Business Outlook Survey," Technical Report, Bank of Canada 2020.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Technical Report, National Bureau of Economic Research 2020.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Bahr, Peter Riley, "The Labor Market Returns to a Community College Education for Noncompleting Students," *The Journal of Higher Education*, 2019, *90* (2), 210–243.

Barr, Andrew and Sarah E Turner, "Expanding enrollments and contracting state budgets: The effect of the Great Recession on higher education," *The ANNALS of the American Academy of Political and Social Science*, 2013, *650* (1), 168–193.

Bartel, Ann P and Lacy Glenn Thomas, "Direct and indirect effects of regulation: A new look at OSHA's impact," *The Journal of Law and Economics*, 1985, *28* (1), 1–25.

_ **and** _ , "Predation through regulation: the wage and profit effects of the occupational safety and health administration and the environmental protection agency," *The Journal of Law and Economics*, 1987, *30* (2), 239–264.

Bartik, Alexander W, "Worker adjustment to changes in labor demand: Evidence from longitudinal census data," *Job Market Paper*, 2017.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka, "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition," *American Economic Review*, 2016, *106* (6), 1437–75.

Basu, Susanto and John G Fernald, "Returns to Scale in US Production: Estimates and Implications," *Journal of Political Economy*, 1997, *105* (2), 249–283.

Becker, Randy and Vernon Henderson, "Effects of air quality regulations on polluting industries," *Journal of political Economy*, 2000, *108* (2), 379–421.

Belfield, Clive and Thomas Bailey, "The Labor Market Returns to Sub-Baccalaureate College: A Review. A CAPSEE Working Paper.," *Center for Analysis of Postsecondary Education and Employment*, 2017.

Berman, Eli and Linda TM Bui, "Environmental regulation and productivity: evidence from oil refineries," *Review of Economics and Statistics*, 2001, *83* (3), 498–510.

Betts, Julian R and Laurel L McFarland, "Safe port in a storm: The impact of labor market conditions on community college enrollments," *Journal of Human resources*, 1995, pp. 741–765.

Black, Dan A, Terra G McKinnish, and Seth G Sanders, "Tight labor markets and the demand for education: Evidence from the coal boom and bust," *ILR Review*, 2005, *59* (1), 3–16.

Blanchard, Olivier and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, (1), 1–75.

Blei, David, Andrew Ng, and Michael Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, 2003, *3* (Jan), 993–1022.

Bloom, Nicholas and John Van Reenen, "Why do Management Practices Differ Across Firms and Countries?," *Journal of Economic Perspectives*, 2010, *24* (1), 203–24.

_ , Stephen Bond, and John Van Reenen, "Uncertainty and Investment Dynamics," *Review of Economic Studies*, 2007, *74* (2), 391–415.

Buehlmaier, Matthias and Toni Whited, "Are Financial Constraints Priced? Evidence from Textual Analysis," *Review of Financial Studies*, 2018, *31* (7), 2693–2728.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of econometrics*, 2021, *225* (2), 200–230.

_ , Andrew Goodman-Bacon, and Pedro HC Sant'Anna, "Difference-in-differences with a continuous treatment," *arXiv preprint arXiv:2107.02637*, 2021.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko, "The Inflation Expectations of US Firms: Evidence from a New Survey," Technical Report, National Bureau of Economic Research 2021.

Cao, Sean, Wei Jiang, Baozhong Yang, and Alan L Zhang, "How to Talk When a Machine is Listening: Corporate Disclosure in the Age of AI," Technical Report, National Bureau of Economic Research 2020.

Caplin, Andrew and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 1987, *102* (4), 703–725.

_ , Mark Dean, and John Leahy, "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice," *Review of Economic Studies*, 2019, *86* (3), 1061–1094.

Card, David, Jochen Kluve, and Andrea Weber, "What works? A meta analysis of recent active labor market program evaluations," *Journal of the European Economic Association*, 2018, *16* (3), 894–931.

Carnevale, Anthony P, Neil Ridley, and L Fasules Megan, "Certificates in Oregon: A model for workers to jump-start or reboot careers," 2018.

Cascio, Elizabeth U and Ayushi Narayan, "Who needs a fracking education? The educational response to low-skill biased technological change," Technical Report, National Bureau of Economic Research 2015.

Cellini, Stephanie Riegg, "Crowded colleges and college crowd-out: The impact of public subsidies on the two-year college market," *American Economic Journal: Economic Policy*, 2009, *1* (2), 1–30.

_ , "For-profit colleges in the United States: Insights from two decades of research," *The Routledge Handbook of the Economics of Education. Routledge*, 2021, *2021*, 512–523.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo, "Housing booms and busts, labor market opportunities, and college attendance," *American Economic Review*, 2018, *108* (10), 2947–94.

Chiang, Yu-Ting, "Strategic Uncertainty over Business Cycles," *Manuscript*, 2021.

Chow, Melissa C, Teresa C Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T Kirk White, "Redesigning the longitudinal business database," Technical Report, National Bureau of Economic Research 2021.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico, "Monetary Policy, Corporate Finance, and Investment," *Journal of the European Economic Association*, 2023, p. jvad009.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 2015, *105* (8), 2644–78.

_ , _ , and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 2018, *108* (9), 2671–2713.

_ , _ , and Tiziano Ropele, "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics*, 2020, *135* (1), 165–219.

Cook, Timothy and Thomas Hahn, "The Effect of Changes in the Federal Funds Rate Target on Market Interest Rates in the 1970s," *Journal of Monetary Economics*, 1989, *24* (3), 331–351.

Cover, Thomas and Joy Thomas, *Elements of Information Theory*, Wiley-Interscience, 2006.

Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf, "Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity," *Review of Income and Wealth*, 2022.

Dellas, Harris and Plutarchos Sakellaris, "On the cyclicality of schooling: theory and evidence," *oxford Economic papers*, 2003, *55* (1), 148–172.

Deming, David and Chris Walters, "The impacts of price and spending subsidies on US postsecondary attainment," *NBER Working Paper*, 2017, *23736.*

Deming, David J, Claudia Goldin, and Lawrence F Katz, "The for-profit postsecondary school sector: Nimble critters or agile predators?," *Journal of Economic Perspectives*, 2012, *26* (1), 139–164.

Dube, Arindrajit, Daniele Girardi, Oscar Jorda, and A Taylor, "A local projections approach to difference-in-differences event studies," *NBER working paper*, 2022.

Dufour, Charles, Paul Lanoie, and Michel Patry, "Regulation and productivity," *Journal of Productivity Analysis*, 1998, *9* (3), 233–247.

Flynn, Joel P and Karthik Sastry, "Attention Cycles," *Manuscript*, 2022.

Foote, Andrew and Michel Grosz, "The effect of local labor market downturns on post-secondary enrollment and program choice," *Education Finance and Policy*, 2019, pp. 1–50.

Fort, Teresa C, Shawn D Klimek et al., "The effects of industry classification changes on US employment composition," *Tuck School at Dartmouth*, 2016.

Freeman, Richard B, *America Works: Thoughts on an exceptional US labor market*, Russell Sage Foundation, 2007.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar, "Expectations with Endogenous Information Acquisition: An Experimental Investigation," *Review of Economics and Statistics*, 2018, pp. 1–54.

Gabaix, Xavier, "Behavioral Inattention," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 2, Elsevier, 2019, pp. 261–343.

Ganong, Peter and Daniel Shoag, "Why has regional income convergence in the US declined?," *Journal of Urban Economics*, 2017, *102*, 76–90.

Gertler, Mark and John Leahy, "A Phillips Curve with an *Ss* Foundation," *Journal of Political Economy*, 2008, *116* (3), 533–572.

_ **and Simon Gilchrist**, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics*, 1994, *109* (2), 309–340.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," 2018.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Goodman, Sarena and Alice Henriques Volz, "Attendance spillovers between public and for-profit colleges: Evidence from statewide variation in appropriations for higher education," *Education Finance and Policy*, 2020, *15* (3), 428–456.

Gorodnichenko, Yuriy and Michael Weber, "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review*, 2016, *106* (1), 165–99.

Gray, Wayne B, "The cost of regulation: OSHA, EPA and the productivity slowdown," *The American Economic Review*, 1987, *77* (5), 998–1006.

_ **and John M Mendeloff**, “The declining effects of OSHA inspections on manufacturing injuries, 1979–1998,” *ILR Review*, 2005, *58* (4), 571–587.

_ **and John T Scholz**, “Does regulatory enforcement work-a panel analysis of OSHA enforcement,” *Law & Soc’y Rev.*, 1993, *27*, 177.

_ **and Ronald J Shadbegian**, “Plant vintage, technology, and environmental regulation,” *Journal of Environmental Economics and Management*, 2003, *46* (3), 384–402.

Greenstone, Michael, “The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures,” *Journal of political economy*, 2002, *110* (6), 1175–1219.

_ **, John A List, and Chad Syverson**, “The effects of environmental regulation on the competitiveness of US manufacturing,” Technical Report, National Bureau of Economic Research 2012.

Grigsby, John, “Skill Heterogeneity and Aggregate Labor Market Dynamics,” Technical Report, Working Paper, University of Chicago 2019., Erik Hurst, and Ahu Yildirmaz ... 2019.

Guner, Nezih, Gustavo Ventura, and Yi Xu, “Macroeconomic implications of size-dependent policies,” *Review of economic Dynamics*, 2008, *11* (4), 721–744.

Gürkaynak, Refet S, Brian Sack, and Eric T Swanson, “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 2005.

Haldane, Andrew, Alistair Macaulay, and Michael McMahon, “The Three E’s of Central-Bank Communication with the Public,” in Ernesto Pastén and Ricardo Reis, eds., *Independence, Credibility, and Communication of Central Banking*, 2021.

Handley, Kyle and J Frank Li, “Measuring the Effects of Firm Uncertainty on Economic Activity: New Evidence from One Million Documents,” Technical Report, National Bureau of Economic Research 2020.

Hansen, Stephen, Michael McMahon, and Andrea Prat, “Transparency and Deliberation within the FOMC: A Computational Linguistics Approach,” *Quarterly Journal of Economics*, 2018, *133* (2), 801–870.

Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, “Firm-level Political Risk: Measurement and Effects,” *Quarterly Journal of Economics*, 2019, *134* (4), 2135–2202.

Haviland, Amelia M, Rachel M Burns, Wayne B Gray, Teague Ruder, and John Mendeloff, “A new estimate of the impact of OSHA inspections on manufacturing injury rates, 1998–2005,” *American journal of industrial medicine*, 2012, *55* (11), 964–975.

He, Guojun, Shaoda Wang, and Bing Zhang, "Watering down environmental regulation in China," *The Quarterly Journal of Economics*, 2020, *135* (4), 2135–2185.

Heckman, James J, Robert J LaLonde, and Jeffrey A Smith, "The economics and econometrics of active labor market programs," in "Handbook of labor economics," Vol. 3, Elsevier, 1999, pp. 1865–2097.

Hillman, Nicholas W and Erica Lee Orians, "Community colleges and labor market conditions: How does enrollment demand change relative to local unemployment rates?," *Research in Higher Education*, 2013, *54*, 765–780.

Horn, Laura, Stephanie Nevill, and James Griffith, "Profile of Undergraduates in US Postsecondary Education Institutions, 2003-04: With a Special Analysis of Community College Students. Statistical Analysis Report. NCES 2006-184.," *National Center for Education Statistics*, 2006.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

Hubbard, Daniel, "The impact of local labor market shocks on college choice: Evidence from plant closings in Michigan," in "Unpublished manuscript. Retrieved from https://www. aeaweb. org/conference/2019/preliminary/paper/Q67dtN77" 2018.

Hyman, Benjamin G, "Can displaced labor be retrained? Evidence from quasi-random assignment to trade adjustment assistance," in "Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association," Vol. 111 JSTOR 2018, pp. 1–70.

Jacobson, Louis, Robert LaLonde, and Daniel G Sullivan, "Estimating the returns to community college schooling for displaced workers," *Journal of Econometrics*, 2005, *125* (1-2), 271–304.

Jaimovich, Nir and Henry E Siu, "The trend is the cycle: Job polarization and jobless recoveries," Technical Report, National Bureau of Economic Research 2012.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

Kane, Thomas J and Cecilia Elena Rouse, "Labor-market returns to two-and four-year college," *The American Economic Review*, 1995, *85* (3), 600–614.

_ and _ , "The community college: Educating students at the margin between college and work," *Journal of economic Perspectives*, 1999, *13* (1), 63–84.

Kehrig, Matthias, "The cyclical nature of the productivity distribution," *Earlier version: US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*, 2015.

Ko, Kilkon, John Mendeloff, and Wayne Gray, "The role of inspection sequence in compliance with the US Occupational Safety and Health Administration's (OSHA) standards: Interpretations and implications," *Regulation & Governance*, 2010, *4* (1), 48–70.

Konchitchki, Yaniv and Jin Xie, "Undisclosed Material Inflation Risk," *Manuscript*, 2022.

Kwon, Spencer Yongwook, Yueran Ma, and Kaspar Zimmermann, "100 Years of Rising Corporate Concentration," *Manuscript*, 2022.

Larsen, Vegard H, Leif Anders Thorsrud, and Julia Zhulanova, "News-Driven Inflation Expectations and Information Rigidities," *Journal of Monetary Economics*, 2021, *117*, 507–520.

Leahy, John V. and Toni M. Whited, "The Effect of Uncertainty on Investment: Some Stylized Facts," *Journal of Money, Credit and Banking*, 1996, *28* (1), 64–83.

Levine, David I, Michael W Toffel, and Matthew S Johnson, "Randomized government safety inspections reduce worker injuries with no detectable job loss," *Science*, 2012, *336* (6083), 907–911.

Loughran, Tim and Bill McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 2011, *66* (1), 35–65.

Luo, Yulei, "Consumption Dynamics under Information Processing Constraints," *Review of Economic Dynamics*, 2008, *11* (2), 366–385.

Ma, Jennifer and Sandy Baum, "Trends in community colleges: Enrollment, prices, student debt, and completion," *College Board Research Brief*, 2016, *4*, 1–23.

Macaulay, Alistair, "Cyclical Attention to Saving," in "UniCredit Foundation Working Papers Series 140" 2020.

Maćkowiak, Bartosz and Mirko Wiederholt, "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 2009, *99* (3), 769–803.

_ and _ , "Business Cycle Dynamics under Rational Inattention," *Review of Economic Studies*, 2015, *82* (4), 1502–1532.

_ and _ , "Rational Inattention and the Business Cycle Effects of Productivity and News Shocks," 2023.

_ , Emanuel Moench, and Mirko Wiederholt, "Sectoral price data and models of price setting," *Journal of Monetary Economics*, 2009, *56*, S78–S99.

_ , Filip Matějka, and Mirko Wiederholt, "Forthcoming, Rational Inattention: A Review," *Journal of Economic Literature*, Forthcoming.

MacLaury, Judson Edward, "Special Anniversary Feature: The Division of Labor Standards: Laying the Groundwork for OSHA," *Applied Industrial Hygiene*, 1988, *3* (12), F–8.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Mankiw, Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Marcus, Michelle and Pedro HC Sant'Anna, "The role of parallel trends in event study settings: An application to environmental economics," *Journal of the Association of Environmental and Resource Economists*, 2021, *8* (2), 235–275.

Matějka, Filip and Alisdair McKay, "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 2015, *105* (1), 272–98.

McCaffrey, David P, "An assessment of OSHA's recent effects on injury rates," *The Journal of Human Resources*, 1983, *18* (1), 131–146.

McKay, Alisdair and Johannes F Wieland, "Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy," *Econometrica*, 2021, *89* (6), 2717–2749.

Mehra, Rajnish and Edward C Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 1985, *15* (2), 145–161.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 74–107.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak, "Internal migration in the United States," *Journal of Economic perspectives*, 2011, *25* (3), 173–96.

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Nonneutrality: The Information Effect," *Quarterly Journal of Economics*, 2018.

Olley, G Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 1996, *64* (6), 1263–1297.

OSHA, "Scheduling System for Programmed Inspections," Technical Report 1995.

_ , "Field Operations Manual (FOM)," Technical Report 2023.

Ottonello, Pablo and Thomas Winberry, "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Porter, Michael E and Claas van der Linde, "Toward a new conception of the environment-competitiveness relationship," *Journal of economic perspectives*, 1995, *9* (4), 97–118.

Ramey, Valerie, "Macroeconomic Shocks and Their Propagation," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 71–162.

Reis, Ricardo, "Inattentive Producers," *Review of Economic Studies*, 2006, *73* (3), 793–821.

Reutskaja, Elena, Rosemarie Nagel, Colin Camerer, and Antonio Rangel, "Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study," *American Economic Review*, 2011, *101* (2), 900–926.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023.

Ruser, John W and Robert S Smith, "Reestimating Osha's effects: have the data changed?," *Journal of Human Resources*, 1991, pp. 212–235.

Sims, Christopher, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Siskind, Frederic B, "Twenty Years of OSHA Federal Enforcement Data: A Review and Explanation of the Major Trends," 1993.

Smith, Robert Stewart, "The impact of OSHA inspections on manufacturing injury rates," *Journal of Human Resources*, 1979, pp. 145–170.

Tanaka, Mari, Nicholas Bloom, Joel David, and Maiko Koga, "Firm Performance and Macro Forecast Accuracy," *Journal of Monetary Economics*, 2019.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Van Nieuwerburgh, Stijn and Laura Veldkamp, "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 2006, *53* (4), 753–772.

_ and _ , "Information Immobility and the Home Bias Puzzle," *Journal of Finance*, 2009, *64* (3), 1187–1215.

Vavra, Joseph, "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *Quarterly Journal of Economics*, 2014, *129* (1), 215–258.

Vike, James, "The Bureaucracy as a Battleground: Contentious Politics Surrounding OSHA 1980-2004," *Politics & Policy*, 2007, *35* (3), 570–607.

Viscusi, W Kip, "The impact of occupational safety and health regulation," *The Bell Journal of Economics*, 1979, pp. 117–140.

Weil, David, "If OSHA is so bad, why is compliance so good?," *The RAND Journal of Economics*, 1996, pp. 618–640.

Weitzner, Gregory and Cooper Howes, "Bank Information Production Over the Business Cycle," *Manuscript*, 2021.

Woodford, Michael, "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics*, 2009, *56*, S100–S124.

Wooldridge, Jeffrey M, "On estimating firm-level production functions using proxy variables to control for unobservables," *Economics letters*, 2009, *104* (3), 112–114.

Xu, Di and Madeline Trimble, "What about certificates? Evidence on the labor market returns to nondegree community college awards in two states," *Educational Evaluation and Policy Analysis*, 2016, *38* (2), 272–292.

Yagan, Danny, "Moving to opportunity? Migratory insurance over the Great Recession," *Job Market Paper*, 2014.

Yang, Choongryul, "Rational Inattention, Menu Costs, and Multi-product Firms: Micro Evidence and Aggregate Implications," *Journal of Monetary Economics*, 2022, *128*, 105–123.

Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 2004, *86* (2), 514–533.