

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays on the Economics of Recent Public Policies

Permalink

<https://escholarship.org/uc/item/8062m7h9>

Author

Kennedy, Patrick J

Publication Date

2023

Peer reviewed|Thesis/dissertation

Essays on the Economics of Recent Public Policies

by

Patrick J. Kennedy

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jesse Rothstein, Chair

Professor Alan Auerbach

Professor Danny Yagan

Spring 2023

Essays on the Economics of Recent Public Policies

Copyright 2023
by
Patrick J. Kennedy

Abstract

Essays on the Economics of Recent Public Policies

by

Patrick J. Kennedy

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Jesse Rothstein, Chair

This dissertation empirically studies the economics of several recent important public policies in the United States, including the corporate tax cuts enacted by the Tax Cuts and Jobs Act of 2017; the Opportunity Zone program enacted by the same legislation; housing policy changes in the city of San Francisco; and tariffs enacted during the ongoing US-China Trade war. The chapters offer both causal and descriptive evidence and build on previous economic research in the fields of public finance, labor economics, and international trade. Overall, the dissertation highlights the importance of economic research in evaluating the tradeoffs inherent in setting policy and in empirically evaluating policy impacts.

Contents

Contents	i
List of Figures	iii
List of Tables	v
1 The Efficiency-Equity Tradeoff of the Corporate Income Tax: Evidence from the Tax Cuts and Jobs Act	3
1.1 Introduction	4
1.2 Institutional Setting: The Tax Cuts and Jobs Act	9
1.3 Data	12
1.4 Empirical Analysis	20
1.5 Revenue Impacts, Excess Burden, and Incidence	51
1.6 Conclusion	62
2 The TCJA and Domestic Corporate Tax Rates	64
2.1 Introduction	65
2.2 Overview of TCJA Corporate Tax Provisions	67
2.3 Description of the Corporate Tax Data	68
2.4 Results	70
2.5 Conclusions	77
3 Neighborhood-Level Investment from the U.S. Opportunity Zone Program	79
3.1 Introduction	80
3.2 Opportunity Zones: Brief Background	82
3.3 Descriptive Evidence on OZ Investment	86
3.4 New Panels from IRS Microdata	100
3.5 Conclusion	105
4 High-End Housing and Gentrification: Evidence from a San Francisco Lottery	107
4.1 Introduction	108
4.2 Background and Setting	111

4.3	Data	115
4.4	Design	119
4.5	Results	124
4.6	Neighborhood Effects	133
4.7	Conclusion	140
5	The Return to Protectionism	141
5.1	Introduction	142
5.2	Data and Timeline	145
5.3	Trade Framework and Identification	156
5.4	Estimation	161
5.5	Aggregate and Regional Impacts	173
5.6	Conclusion	186
	Bibliography	188

List of Figures

1	Marginal Income Tax Rates and Taxes Per Worker	11
2	Firm Size Distributions and Industry Composition	18
3	Event Studies: Marginal Tax Rates and Taxes Paid	22
4	Event Study: Firm Sales, Costs, and Pre-Tax Profits	24
5	Event Study: After-Tax Profits and Shareholder Payouts	29
6	Event Studies: Labor Market Outcomes	32
7	Earnings Quantile Regressions	33
8	Event Studies: Executive Pay	34
9	Event Studies: Net Investment	38
10	Firm Heterogeneity	42
11	Corporate Entity-Type Switching, 2013-2019	49
12	Geographic Incidence	60
13	The Efficiency-Equity Tradeoff in Context	62
1	Effective Tax Rates for C-Corporations: 2010 to 2019	71
2	Distribution of ETRs Before and After TCJA	72
3	Coefficient Estimates: TCJA provisions and ETR Changes	77
1	Tract-Level Distribution of OZ Investment	88
2	Demographic Correlates of OZ Investment	91
3	OZ Investment in 25 Top Commuting Zones	93
4	Mapping OZ Investment in Six Illustrative Cities	95
5	OZ Investment, Population Density, and Distance	96
6	Income Distribution of QOF Investors and the US Population	98
7	Geography of QOF Investors	99
8	IRS Measures vs. ACS Measures	102
9	Employment in IRS and LODES Data	103
10	IRS Correlates of OZ Investment	105
1	Gentrification in San Francisco and Other U.S. Cities	113
2	Lottery Applicants and Winners	116
3	Lottery Probability of Winning	117

4	Map of Lottery Applicants and Winners	118
5	First-Stage: Effect of Winning Lottery on Condominium Conversions	125
6	Effect of Winning the Lottery on Property Values and Permits	127
8	Effect of Winning the Lottery on Homeownership, Sales, and Evictions	130
10	Spillovers on Nearby Property Values	132
12	Neighborhood Heterogeneity	139
1	Variety Event Study: Imports	154
2	Variety Event Study: Exports	155
3	Dynamic Specification	169
4	Regional Variation in U.S. and Retaliatory Tariffs	180
5	Model Simulation of Real Wage Impacts from U.S. and Retaliatory Tariffs	182
6	Tariff Changes vs. 2016 Republican Vote Share	184
7	Model-Simulated Tradeable Real Wage Impact vs. 2016 Republican Vote Share	185

List of Tables

1	Summary Statistics	19
2	Marginal Tax Rates and Taxes Paid	25
3	Sales, Costs, and Pre-Tax Profits	27
4	After-Tax Profits and Shareholder Payouts	30
5	Labor Market Outcomes	35
6	Net Investment	39
7	Robustness to Alternate Specifications	44
8	Robustness to Alternate Samples	47
9	Quantification Moments and Parameters	57
10	Revenue and Welfare Estimates	58
11	Incidence Estimates	61
1	Exposure to TCJA Provisions	75
1	Tract Summary Statistics	85
2	Investment in Opportunity Zones Over Time	86
3	Investment in Opportunity Zones by Type	87
4	Industry Composition of Funds and Recipient Firms	89
1	Balance Table	119
2	Effect of Winning the Lottery on Property Values	128
3	Effects of Lottery Winners on Neighborhood Demographics	136
4	Neighborhood Effects (Business Counts)	137
1	The 2018 Trade War	148
2	Sector Variation in Tariff Rate Changes	151
3	Tests for Pre-Existing Trends	162
4	Variety Import Demand (σ) and Foreign Export Supply (ω^*)	163
5	Product Elasticity η	165
7	Foreign Import Demand σ^*	167
6	Sector Elasticity κ	172
8	Aggregate Impacts	178

Acknowledgments

I am deeply indebted to my committee chair, Jesse Rothstein. Early in graduate school, someone suggested to me that I should find an advisor who exemplified the type of person and economist that I'd aspire to be like. I'm very glad and fortunate to have found one.

Guidance from my committee advisors, Alan Auerbach and Danny Yagan, as well as from Emmanuel Saez and Gabriel Zucman, only confirmed what I already knew: that Berkeley is the best place in the world to study public economics.

I benefitted immensely from nearly a decade of mentorship from Amit Khandelwal, who showed me how to actually do economics. Who knows where I'd be otherwise.

Research is both more productive and more fun with a team. I'm grateful to my brilliant coauthors, who have taught me so much: Pablo Fajgelbaum, Penelopi Goldberg, and Amit Khandelwal; Christine Dobridge, Paul Landefeld, and Jake Mortenson; and Harry Wheeler.

I am thankful to many advisors, teachers, peers, and colleagues: Felipe Arteaga, Nano Barahona, Tom Barthold, Michael Bruce, David Card, Benjamin Faber, Cecile Gaubert, Leticia Juárez, Patrick Kline, Michael Love, Emi Nakamura, Cristóbal Otero, Sebastián Otero, Andrés Rodríguez-Clare, Jesse Rothstein, Benjamin Schoefer, David Sraer, Elif Tasar, Damián Vergara, seminar participants at Berkeley, the California Policy Lab, and staff at the Joint Committee on Taxation.

I acknowledge generous financial support from the National Science Foundation Graduate Research Fellowship Program, the Burch Center for Public Finance, the Berkeley Opportunity Lab, the Clausen Center for International Business and Policy, and the Center for Equitable Growth.

Most of all, I am grateful to my family and friends for their love and support no matter how my dissertation turned out. I am so lucky.

Introduction

This dissertation empirically studies the economics of several recent important public policies in the United States, including the corporate tax cuts enacted by the Tax Cuts and Jobs Act of 2017; the Opportunity Zone program enacted by the same legislation; housing policy changes in the city of San Francisco; and tariffs enacted during the ongoing US-China Trade war. The chapters offer both causal and descriptive evidence and build on previous economic research in the fields of public finance, labor economics, and international trade.

In Chapter 1, Christine Dobridge, Paul Landefeld, Jake Mortenson and I study the effects of an historically large federal corporate income tax cut on U.S. firms and workers, leveraging quasi-experimental policy variation from the 2017 law known as the Tax Cuts and Jobs Act. To identify causal effects, we use employer-employee matched federal tax records and an event study design comparing similarly-sized firms in the same industry that faced divergent tax changes due to their pre-existing legal status. We find that reductions in marginal income tax rates cause increases in sales, profits, investment, and employment, with responses driven by firms in capital-intensive industries. Workers' earnings gains are concentrated in executive pay and in the top 10% of the within-firm income distribution, while workers in the bottom 90% of the distribution see no change in earnings. Interpreted through the lens of a stylized model, our estimates imply that a \$1 marginal reduction in corporate tax revenue generates an additional \$0.10 in output, with 78% of gains flowing to the top 10% of the income distribution.

In Chapter 2, again with Dobridge, Landefeld, and Mortenson, we document descriptive trends in the tax positions for U.S. C-corporations following passage of the TCJA. We document several novel facts, including that effective tax rates (ETRs) *increased* on average for privately held, domestic firms and for firms in the bottom 90% of the firm size distribution after the TCJA. By contrast, public, multinational, and large firms saw substantial ETR cuts. We show that changes to the corporate tax rate and treatment of net operating losses have the strongest correlation with ETR changes, and discuss avenues for future research.

In Chapter 3, Harrison Wheeler and I use de-identified federal tax records from tax years 2019 and 2020 to document the first available evidence on the short-run response of financial capital to the Opportunity Zone (OZ) program, a federal place-based policy that provides tax incentives for capital investments in more than 8,000 low-income neighborhoods across the United States. We observe \$41.5 billion of aggregate cumulative OZ investments by tax year 2020. Using a sub-sample of electronically filed returns covering 78% of total

observed investment, we document three emerging patterns in the data. First, OZ capital is highly spatially concentrated. Second, among OZ-designated neighborhoods, investors report greater equity and property investments in neighborhoods with relatively higher incomes, home values, educational attainment, and pre-existing income and population growth. Third, OZ investors have extremely high incomes relative to the US population, implying that the direct distributional incidence of the tax subsidy benefits households in the 99th percentile of the national income distribution.

In Chapter 4, Wheeler and I study the effects of housing regulations on gentrification. In major cities throughout the United States and the world, contentious debates over housing policy frequently focus on the extent to which high-end developments are a cause of gentrification. We shed new light on this question empirically by studying a unique lottery in the city of San Francisco that allowed a limited number of property owners to convert their buildings into high-end condominiums. Relying solely on exogenous variation from the lottery, we study the long-run effects of these developments on local home prices, demographics, and new business entry. Compared to losing lottery applicants, winners are vastly more likely to invest in alterations and renovations in their properties, to see their property values increase, to rent their properties to new tenants, and to sell them to new owners. The home values of adjacent buildings also increase. At the neighborhood level, conversions lead to higher resident incomes, home values, rental prices, and shares of the population that are White and college-educated.

In Chapter 5, Pablo Fajgelbaum, Penelopi Goldberg, Amit Khandelwal and I study the aggregate and distributional effects of the import and export tariffs enacted during the ongoing US-China trade war. After decades of supporting free trade, in 2018 the United States raised import tariffs and major trade partners retaliated. We analyze the short-run impact of this return to protectionism on the U.S. economy. We find that import and retaliatory tariffs caused large declines in imports and exports. Prices of imports targeted by tariffs did not fall, implying complete pass-through of tariffs to duty-inclusive prices. The resulting losses to U.S. consumers and firms that buy imports was \$51 billion, or 0.27% of GDP. We embed the estimated trade elasticities in a general-equilibrium model of the U.S. economy. After accounting for tariff revenue and gains to domestic producers, the aggregate real income loss was \$7.2 billion, or 0.04% of GDP. Import tariffs favored sectors concentrated in politically competitive counties, and the model implies that tradeable-sector workers in heavily Republican counties were the most negatively affected due to the retaliatory tariffs.

Overall, the dissertation highlights the importance of economic research in evaluating the tradeoffs inherent in setting policy and in empirically evaluating policy impacts.

Chapter 1

The Efficiency-Equity Tradeoff of the Corporate Income Tax: Evidence from the Tax Cuts and Jobs Act

In Chapter 1, Christine Dobridge, Paul Landefeld, Jake Mortenson and I study the effects of an historically large federal corporate income tax cut on U.S. firms and workers, leveraging quasi-experimental policy variation from the 2017 law known as the Tax Cuts and Jobs Act. To identify causal effects, we use employer-employee matched federal tax records and an event study design comparing similarly-sized firms in the same industry that faced divergent tax changes due to their pre-existing legal status. We find that reductions in marginal income tax rates cause increases in sales, profits, investment, and employment, with responses driven by firms in capital-intensive industries. Workers' earnings gains are concentrated in executive pay and in the top 10% of the within-firm income distribution, while workers in the bottom 90% of the distribution see no change in earnings. Interpreted through the lens of a stylized model, our estimates imply that a \$1 marginal reduction in corporate tax revenue generates an additional \$0.10 in output, with 78% of gains flowing to the top 10% of the income distribution.

This chapter is coauthored with Christine Dobridge, Paul Landefeld, and Jake Mortenson. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee. The views and opinions expressed here are the authors' own. They are not necessarily those of the Board of Governors of the Federal Reserve System, its members, or its staff.

1.1 Introduction

How large are the output distortions from the corporate tax, and what are its distributional implications? We study the effects of corporate income tax cuts on firms and workers in the United States, where in 2017 Congress enacted the most sweeping and significant legislation on American federal business taxation in a generation. Commonly known as the Tax Cuts and Jobs Act (TCJA), the legislation introduced reforms to corporate marginal income tax rates, investment incentives, and taxation of foreign income, among several other provisions of the tax code. Collectively, the breadth of these provisions and the magnitude of the tax rate changes constitute the largest overhaul of American business taxation since the Tax Reform Act of 1986, providing a rare and sharp natural experiment to shed light on contemporary research and policy debates.

Even as governments around the globe have dramatically reduced corporate income tax rates over the past half-century — from an unweighted average country worldwide statutory tax rate of 40% in 1980 to just 23% in 2021 ([Tax Foundation 2021](#)) — policymakers and researchers today fiercely debate the costs and benefits of declining corporate tax burdens. Advocates for tax cuts argue that lower rates increase investment, growth, and workers' living standards, while opponents argue they do little to boost growth and primarily benefit the wealthy.

In this paper we bring new evidence to these debates. Our empirical analysis specifically studies the core provisions of TCJA affecting firms' statutory marginal income tax rates, using a rich employer-employee matched panel dataset constructed from large random samples of firm- and worker-level federal tax records. The data allow us to observe a holistic set of firm outcomes — such as sales, profits, shareholder payouts, and investment — and to merge them with worker-level data on employment and annual labor earnings.

Our main empirical strategy leverages an event study design to compare the outcomes of similarly-sized firms in the same industry that faced divergent changes in their tax treatment. In particular, TCJA cut the top marginal tax rate for a legal entity type of firms known as C-corporations from 35% to 21% (a 40% reduction). At the same time, TCJA cut the implied top marginal tax rate for a separate legal entity type of firms known as S-corporations from 39.6% to 37% (a 6% reduction), and also introduced a new tax deduction that, for many of these firms, further reduced the top marginal rate to 29.6% (for a cumulative 25% reduction).¹ C-corporations and S-corporations operate in the same industries, overlap in their firm size distributions, and faced broadly similar tax burdens prior to TCJA, inviting a natural comparison between the two.

¹The top marginal tax rate for S-Corporations is *implied* because, unlike for C-Corporations, it must be computed as a weighted average of the marginal tax rates faced by each firms' individual owners and cannot be directly inferred from firms' tax records. The newly introduced tax deduction for S-corporations is known as the Qualified Business Income (QBI) deduction. We discuss both points in greater detail in Sections 1.2 and 1.3. As we will describe, the main differences between the firm types are that C-corporations may have unlimited shareholders and pay taxes directly to the federal government, whereas S-corporations face greater shareholder restrictions and pay taxes indirectly to the government via the individual tax filings of their shareholders.

We exploit the fact that the average C-corp received a significantly larger tax cut than the average S-corp to provide the first exhaustive evidence of these corporate tax changes on firms' sales, profits, shareholder payouts, investment, and employment, as well as on workers' annual earnings. As in [Yagan \(2015\)](#), the identifying assumption of our research design is not random assignment of C or S status; rather, it is that outcomes for C- and S-corps would have trended similarly in the absence of the tax cuts. Event studies indicate that outcomes of comparable C- and S-corps were on similar trends prior to TCJA, and we further implement a series of robustness checks to validate that our causal estimates are driven by changes in top marginal tax rates rather than other features of the law, superficial tax shifting behaviors, or unrelated economic forces differentially affecting C- and S-corps at the same time as TCJA.

Our benchmark regression specifications, which compare trends in outcomes of C- versus S-corps controlling for firm and industry-size-year fixed effects, indicate that corporate income tax cuts cause economically and statistically significant increases in firms' sales, profits, payouts to shareholders, employment, and real investment in capital goods. Responses are concentrated in capital intensive industries, and are not larger for smaller or cash-constrained firms, suggesting that effects are driven by a reduction in the cost of capital rather than by liquidity effects.

Our benchmark estimate of the federal corporate elasticity of taxable income, a key parameter for measuring the magnitude of tax distortions, is 0.38 (s.e.=0.13). This elasticity is smaller than most comparable estimates generated from variation in state and local corporate taxes, but larger than most estimates on personal income taxes. Since businesses are less mobile at the federal level than at the state or local level, and since a large literature documents that personal labor supply elasticities are small, we interpret this evidence as consistent with the common economic intuition that tax distortions vary proportionally with factor mobility.

Moving to the worker-level evidence, we study impacts on firm wage quantiles and show that annual earnings do not change for workers in the bottom 90% of the within-firm distribution, but do increase for workers in the top 10%, and increase particularly sharply for firm managers and executives. Unlike other outcomes such as employment and investment, we find that executive earnings increase in both capital and non-capital intensive industries. Moreover, executive pay increases are only weakly correlated with changes in firm sales, profits, or sales growth relative to other firms in the same industry. Synthesizing this evidence, we estimate that approximately 10% of the executive pay bump is driven by improved firm performance, while the remaining 90% is plausibly attributable to rent-sharing or executive capture.

Descriptively, relative to the population of workers in our sample, the executives and workers in the top 10% who benefit from higher earnings are typically older, have longer employment tenures at the firm, and are more likely to be men. However, we find little evidence that earnings effects vary heterogeneously by gender, age, or tenure after controlling for workers' place in the within-firm earnings distribution.

To evaluate the effects of corporate tax cuts on tax revenue and output, we combine the

reduced-form elasticities from the empirical analysis with a stylized model of firm owners and workers. Using the model, we estimate that a \$1 marginal reduction in corporate tax revenue generates an additional \$0.10 increase in output. Corporate tax revenues decline by \$0.87, with behavioral responses of firms and workers modestly blunting mechanical revenue losses, and consistent with the notion that contemporary top corporate marginal income tax rates in the US are below the revenue-maximizing rate.

To assess distributional impacts, we estimate the short-run incidence of corporate tax cuts on several factor groups — firm owners, executives, and high- and low-paid workers — as the share of total output gains accruing to each factor. Combining our reduced form elasticities with moments from the tax data, we find that approximately 56% of gains flow to firm owners, 12% flow to executives, 32% flow to high-paid workers, and 0% flow to low-paid workers. We then go beyond factor incidence to estimate effects across the income distribution, accounting for the empirical fact that many workers are also firm owners (that is, they hold equity portfolios) and many firm owners also work. Using data on the distribution of capital ownership, we find that approximately 78% of the gains from tax cuts accrue to the top 10% of earners, and 22% of gains flow to the bottom 90%. Leveraging the empirically observable geographic distribution of workers and income, we further find that these benefits are disproportionately concentrated in the Northeastern and Western regions of the United States, particularly among workers in large and high-income cities.

This paper builds on a large body of research that studies the effects of corporate taxes on profits, investment, shareholder payouts, employment, wages, and executive compensation.² Early seminal studies use aggregate or firm-level panel data and estimate two-way fixed effect models to study policy variation across countries or industries (Hall and Jorgenson 1967; Cummins, Hassett, and Hubbard 1994; Cummins, Hassett, and Hubbard 1996; Goolsbee 1998; Hassett and Hubbard 2002). More recent contributions use detailed administrative microdata and modern econometric methods to exploit geographic policy variation (Link, Menkhoff, Peichl, and Schüle 2022; Duan and Moon 2022; Garrett, Ohrn, and Suárez Serrato 2020; Giroud and Rauh 2019; Fuest, Peichl, and Sieglöcher 2018; Suárez Serrato and Zidar 2016; Becker, Jacob, and Jacob 2013), industry-level variation in exposure to tax deductions or credits (Curtis, Garrett, Ohrn, Roberts, and Serrato 2021; Ohrn 2022; Dobridge, Landefeld, and Mortenson 2021; Ohrn 2018b; Zwick and Mahon 2017; House and Shapiro 2008), and firm-level policy variation induced by plausibly arbitrary legal or circumstantial distinctions (e.g., Boissel and Matray 2022; Moon 2022; Carbonnier, Malgouyres, Py, and Urvoy 2022; Bachas and Soto 2021; Risch 2021; Alstadsæter, Jacob, and Michaely 2017;

²Other outcomes studied in the literature include: establishment counts (e.g., Suárez Serrato and Zidar 2016; Giroud and Rauh 2019); consumer prices (Baker, Sun, and Yannelis 2020); innovation and the mobility of inventors (Akcigit, Grigsby, Nicholas, and Stantcheva 2021); international tax competition (Devereux, Lockwood, and Redoano 2008); the location and investment decisions of multinational firms (Becker, Fuest, and Riedel 2012; Devereux and Griffith 2003); tax avoidance and profit shifting (Garcia-Bernardo, Janský, and Zucman 2022; Desai and Dharmapala 2009; Auerbach and Slemrod 1997; Slemrod 1995; Hines and Rice 1994); and macroeconomic performance (Cloyne, Martinez, Mumtaz, and Surico 2022; Zidar 2019; Romer and Romer 2010, Lee and Gordon 2005). These outcomes are beyond the scope of this paper.

Patel, Seegert, and Smith 2017; Yagan 2015; Devereux, Liu, and Loretz 2014).

Despite major advances in recent research, there are natural reasons to question whether existing evidence is generalizable to understanding the effects of corporate tax cuts in the context of TCJA. Evidence from subnational governments, small developing countries, or small firms may have limited applicability to major reforms in a large advanced economy such as the United States (Auerbach 2018). This concern is especially salient with respect to the U.S. federal corporate income tax, where the tax base is broader, top tax rates are higher, revenues are orders of magnitude larger, and factors of production are considerably less mobile. Moreover, economic theory predicts that alternate tax instruments — such as dividend taxes, capital gains taxes, or narrowly targeted corporate tax deductions and credits — have very different effects than the corporate income tax (Auerbach 2002; Hassett and Hubbard 2002). In this light, it is not surprising that, due to differences in both empirical and normative worldviews, debates over the effects of TCJA remain hotly contested by researchers and policymakers (Barro and Furman 2018).

Empirical evidence on the effects of the federal corporate income tax has remained scarce for three reasons. First, federal tax reforms are rare historical events, leaving limited policy variation for researchers to study. Second, digitized administrative microdata was previously unavailable to researchers, constraining the scope and precision of empirical analyses. Third, even when countries do change their tax rates, it is difficult for researchers to establish credible counterfactuals for causal inference, particularly as the parallel trends assumption underlying cross-country difference-in-difference analyses are challenging to defend in disparate socioeconomic and institutional settings.

This paper overcomes these limitations to provide transparent evidence on the effects of corporate tax cuts on firms and workers. In doing so, we make four main contributions to the literature.

First, we study a rare policy change that generated historically large within-country variation in federal corporate income tax rates, and moreover generated variation even across similarly-sized firms in the same industry. As a share of GDP, the TCJA tax cut is orders of magnitude larger than previous studies that focus, for example, on changes in state or local corporate taxes, which tend to have lower rates and a smaller tax base (e.g., Giroud and Rauh 2019; Fuest, Peichl, and Siegloch 2018; Suárez Serrato and Zidar 2016). The large magnitude of the tax cut is relevant on both theoretical grounds (according to the conventional view that tax distortions are proportional to the square of the tax rate, as in Harberger 1964) and on purely empirical grounds (since ex-ante it is unclear whether existing evidence can be extrapolated to the case of an outlier).

Second, we complement the large shock with detailed employer-employee matched tax records that allow us to observe an unusually holistic set of firm- and worker-level outcomes. We build on frontier research that uses employee-level data to provide a nuanced account of corporate tax incidence on different types of workers (Carbonnier, Malgouyres, Py, and Urvoy 2022; Risch 2021; Dobridge, Landefeld, and Mortenson 2021; Fuest, Peichl, and Siegloch 2018), and extend existing work by empirically estimating geospatial incidence and incidence on firm owners. In contrast to studies that do not directly observe profits (e.g.,

Suárez Serrato and Zidar 2016), the richness of our data allows us to estimate incidence using fewer assumptions than are typically required when data availability are more limited.

Third, we contribute to a growing literature that seeks to understand the effects of TCJA on the U.S. economy. Researchers have studied impacts on macroeconomic performance (Gale and Haldeman 2021; Gale, Gelfond, Krupkin, Mazur, and Toder 2019; Kumar 2019; Barro and Furman 2018; Mertens 2018), international and intertemporal profit shifting (Garcia-Bernardo, Janský, and Zucman 2022; Dowd, Giosa, and Willingham 2020; Clausing 2020), pass-through businesses (Goodman, Lim, Sacerdote, and Whitten 2021), executive compensation (De Simone, McClure, and Stomberg 2022), capital structures (Carrizosa, Gaertner, and Lynch 2020), and regional or local economic outcomes (Kennedy and Wheeler 2022; Altig, Auerbach, Higgins, Koehler, Kotlikoff, Terry, and Ye 2020). Our study differs from existing research in that we specifically study the effects of TCJA’s marginal corporate income tax cuts on firm- and worker-level outcomes using rich administrative microdata and a quasi-experimental research design leveraging cross-firm policy variation.

Finally, we contextualize our findings from this historical episode in broader debates about efficiency and equity in national tax and transfer systems (Carbonnier, Malgouyres, Py, and Urvoy 2022; Bachas and Soto 2021; Risch 2021; Hendren and Sprung-Keyser 2020; Fuest, Peichl, and Siegloch 2018; Suárez Serrato and Zidar 2016; Devereux, Liu, and Loretz 2014; Arulampalam, Devereux, and Maffini 2012; Gruber and Rauh 2007). With respect to efficiency, our model-based estimates of the marginal output gains from cutting the federal corporate income tax are approximately 1.5 to 2 times as large as the literature-implied marginal gains from cutting personal income or payroll taxes. With respect to equity, our results contrast with much existing research in that we find the incidence of the corporate tax falls heavily on capital and highly-paid workers. Assessing incidence across the income distribution, we estimate that corporate income tax cuts are similarly regressive relative to personal income tax cuts, but markedly less progressive than payroll tax cuts. We note that our results capture short-run responses and do not account for potential changes in government spending or after-tax redistribution, which are important considerations for policymakers but beyond the scope of this research.

The rest of the paper proceeds as follows. Section 1.2 summarizes key features of the Tax Cuts and Jobs Act, including its legislative history, institutional context, and major policy changes. Section 1.3 describes data sources and variable definitions. Section 1.4 details our empirical strategy and presents results. Section 1.5 presents a stylized model that we use to estimate the revenue impacts, excess burden, and incidence of TCJA’s corporate tax cuts. Section 3.5 concludes with a discussion of the results.

1.2 Institutional Setting: The Tax Cuts and Jobs Act

Legislative History

In 2017 Congress took on the task of reforming federal business tax policy, with the stated aims of increasing capital investment, economic growth, and international competitiveness.³ Following several months of political negotiations and policy proposals, in December 2017 Congress and the President enacted Public Law 115-97, more commonly known as the Tax Cuts and Jobs Act, or TCJA. The law included provisions affecting many aspects of the federal business tax code, including corporate income tax rates, investment incentives, and taxation of foreign income. Most policy changes were implemented beginning in tax year 2018, although some provisions, such as the investment incentives that we will later discuss, were applied to tax year 2017. Our aim below is not to exhaustively detail TCJA's numerous reforms — for reviews of significant provisions see [Auerbach \(2018\)](#) and [Joint Committee on Taxation \(2018b\)](#) — but rather to illuminate the key institutional details and policy variation that we leverage in our empirical analysis.

C-Corporations vs. S-Corporations

At the heart of TCJA was an overhaul of the income tax schedules facing two legally distinctive types of businesses, known as C-Corporations and S-Corporations. Combined, C- and S-corps account for approximately 70% of total U.S. employment and 74% of total payrolls, with government, non-profits, and non-corporate private businesses comprising the remainder ([Census Bureau 2019](#)). Our analysis focuses exclusively on the corporate sector, as other entity types face different tax and regulatory regimes, and are beyond the scope of this paper. Below we describe salient legal differences between C- and S-corps.

C-Corporations

C-corps are required to pay income taxes directly to the federal government, may be private or public, and are subject to both corporate income taxes (paid on corporate profits) and dividend taxes (paid by shareholders on profits distributed as dividends). Prior to TCJA, C-corps faced a progressive tax schedule with eight income brackets and a top marginal rate of 35%. After TCJA, these brackets collapsed to a single uniform 21% tax rate. The Online Appendix documents the evolution of top marginal income tax rates for C-corps in the United States since 1909, illustrating the historic nature of this large and rare tax cut, and details the collapse of the progressive corporate income brackets following TCJA. The Online Appendix also places the U.S. corporate tax in a global perspective, and benchmarks the magnitude of the TCJA corporate tax cut against other recent studies in the literature.

³The policy reforms were first proposed in a blueprint document released by Republicans in the House of Representatives in June 2016, available [here](#).

S-Corporations

S-corps do not pay taxes directly to the federal government. Rather, the firms' profits are distributed to the individual owners of the firm, who pay taxes on profits as ordinary income and can deduct any losses. S-corps may have up to 100 shareholders, all of whom must be U.S. citizens and not businesses or institutional investors, and are not permitted to sell shares on publicly traded stock exchanges. Unlike C-corps, S-corps do not face corporate income taxes, nor are their distributed profits subject to the dividend tax.

Prior to 2018, owners of S-corps faced a top marginal income tax rate of 39.6%. TCJA then provided two distinct types of tax relief to owners of S-corps. First, it reduced the top personal income tax rate from 39.6% to 37%. Second, it introduced a 20% tax deduction on qualified business income that further reduced the effective marginal tax rate on S-Corp income for most high-income tax-payers from 37% to 29.6%. This tax deduction — known as the Qualified Business Income (“QBI”) deduction, or as “Section 199A” after the applicable section of the internal revenue code — is claimable by most but not all owners of S-corps. Since the QBI limitations are complex and not crucial for our empirical analysis, we abstract from details here and provide more details in the Online Appendix.

Entity Type Choice and Switching

Firms must elect either C or S status upon incorporating. The decision to choose one corporate form over the other may reflect a variety of considerations, including access to capital (recall that S-corps may not be publicly traded) and tax planning (recall that C-corps must pay entity-level taxes and are subject to dividend taxes on distributed profits). After electing C or S status, switching entity types is costly, rare, and subject to regulatory restrictions. Thus, a firm's entity type prior to TCJA is strongly related to the tax rate change it faced after TCJA, and endogenous switching is not a concern for our analysis.


Policy Variation in Marginal Income Tax Rates

Figure 1 shows the evolution of top marginal income tax rates and tax burdens for C- and S-corps in the years before and after TCJA. Panel A shows the sharp reduction in top statutory marginal income tax rates for C-corps, as well as the change in implied top statutory marginal income tax rates for S-corps depending on whether or not they are eligible for the QBI deduction.


Panel B shows the change in observed marginal tax rates from our analysis sample of large firms with at least 100 employees. Entity-level tax rates and taxes paid are imputed for S-corps by linking to returns of S-corp owners, as we will describe in detail in the following section. Observed average marginal tax rates are lower than top statutory rates for several reasons. First, in any given year some firms will have non-positive taxable income (for example, if they earn zero or negative profits) and thus face a marginal tax rate of zero. Second, C-corps prior to TCJA faced a graduated tax rate schedule. Third, our measure of the marginal tax rate for S-corps is computed as a weighted-average of the tax rates faced by their owners, some of whom may not be in the top tax bracket.

FIGURE 1: MARGINAL INCOME TAX RATES AND TAXES PER WORKER

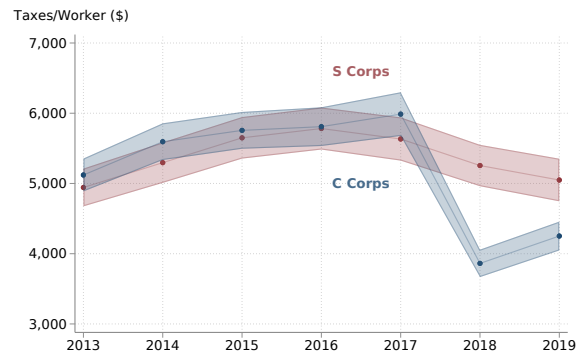
Panel A: Top MTR (Statutory)



Panel B: Average MTR (Observed)



Panel C: Taxes Per Worker (Observed)



Notes: Panel A shows top statutory marginal income tax rates for C and S Corporations before and after enactment of TCJA. Panel B shows the average MTRs observed in our data analysis sample of large firms with at least 100 employees; we discuss the data construction and variable definitions in Section 1.3. Panel C shows the change in taxes per worker paid by C- and S- corps observed in the data over the sample period.

Panel C underscores the economic significance of the tax cuts in dollar terms, documenting the observed average change in corporate taxes per worker for C- and S-corps. The panel shows that average taxes per worker declined from 2016 to 2019 by approximately \$1,600 per worker ($\approx 28\%$) for C-corps, and by approximately \$800 per worker ($\approx 13\%$) for S-corps.

Most importantly, all three panels show that, on average, C-corps received a significantly larger tax cut than S-corps due to TCJA, illustrating the key policy variation that we use in our empirical analysis to identify causal effects.

1.3 Data

We use a panel of employer-employee matched annual federal tax records from tax years 2013 to 2019. We begin the sample period in 2013, allowing us to compare trends in the outcomes of C- and S-corps several years before TCJA, and end the sample in 2019, prior to the onset of the COVID-19 pandemic in 2020. Below we describe the data sources and sample construction, provide variable definitions, and present descriptive statistics. We provide additional details about the data cleaning procedures in the Online Appendix.

Corporate Tax Returns

We study firms in the corporate Statistics of Income (SOI) files produced by the U.S. Internal Revenue Service (IRS). The corporate SOI files include stratified random samples of corporate tax returns from both C-corps (from IRS Form 1120) and S-corps (from IRS Form 1120S). IRS produces and cleans these random samples to estimate aggregate statistics and to provide government agencies with essential data for development of legislation and policy analysis. The corporate tax returns allow us to observe firms' domestic sales, costs, profits, investment, and taxes paid, as well as their year of incorporation and industry. The IRS over-samples large firms with known probability weights, and the samples are designed as rolling panels so as to allow for longitudinal analyses.⁴

We impose the following two sample restrictions on the SOI panel, yielding an analysis sample of approximately 11,600 unique firms and 81,000 distinct firm-year observations.

First, we restrict the sample to large firms, defined as those with at least 100 employees and \$1 million in sales in every year of our pre-treatment period from 2013 to 2016. There are two reasons for restricting the sample to large firms. Large firms account for the lion's share of corporate economic activity, comprising approximately 90% of corporate sales, 70% of corporate taxes, and 67% of corporate employment.⁵ Moreover, many small C-corps faced tax increases (rather than tax cuts) after TCJA due to the flattening of corporate income tax brackets to a uniform 21% rate. Including smaller firms would thus require significantly complicating our empirical design, which simply compares outcomes of similar C- and S-corps

⁴For additional details on construction of the SOI samples, see documentation provided by the IRS [here](#).

⁵Authors' calculations using IRS SOI data.

over time. The large firm restriction both allows us to study the most economically significant firms and to employ a more transparent and credible research design.

Second, we balance the panel and drop firms that ever switch entity types from C to S or from S to C over the course of our sample period. Balancing the panel ensures that our results are not driven by the changing composition of firms in the SOI samples. Because entity-switching is rare, dropping switchers from our sample excludes only approximately 4% of firms, collectively comprising less than 0.5% of corporate sales or profits.

Individual Tax Returns

We complement the sample of corporate tax records with several sources of individual-level administrative records.

First, we merge the sample of corporate tax returns with the universe of worker-level filings of IRS Form W-2, which provides information on workers' annual wage earnings from each of their employers. Employers are required each year to share copies of form W-2 both with their workers and with the IRS, allowing us to observe the earnings of all workers even if they had no federal tax liability or did not file a personal income tax return.

Second, we collect information about the owners of S-corps in our sample from the universe of filings of IRS Form 1099-K1, which provides data on the income received by owners of S-corps from each of their pass-through businesses each year, including pass-through income from non-corporate partnerships. As we will describe below, we complement this information with data from IRS Forms 1040 to compute implied marginal income tax rates and federal taxes paid for S corporations.

Finally, we observe individuals' age and gender from the Master Database maintained by the Social Security Administration (SSA). We also observe their residential location using data from [Kennedy and Wheeler \(2022\)](#).

Variable Definitions and Measurement

Our empirical analysis uses information on firm-level tax rates, taxes paid, sales, profits, investment, employment, and shareholder payouts. We also use data on workers' employment and annual wage and salary earnings. We take care to measure these variables consistently over time, such that our outcomes are not affected, for example, by changes in the tax base or in reporting requirements on tax forms. We provide additional details on variable definitions, including specific forms and line item numbers, in the Online Appendix.

Marginal Tax Rates

Our primary explanatory variable of interest is the marginal income tax rate paid by firms. For C-corps, we observe taxable income and directly infer each firm's marginal income tax rate using the federal corporate income tax schedules reproduced in the Online Appendix. For S-corps, we observe each owner's taxable income from their personal tax returns, and

directly infer their personal marginal income tax rate using the federal personal income tax schedules as reproduced in the Online Appendix. We then compute the implied corporate marginal tax rate for the firm as a weighted average of the marginal personal income tax rates faced by the firms' owners, where the weights are given by the share of ordinary business income distributed to each owner from that firm. For example, if an S-corp has two owners who receive an equal share of that firm's business income, facing marginal tax rates on their individual income of 25% and 35%, respectively, then we compute the implied corporate marginal tax rate as $(.5 * .25) + (.5 * .35) = 30\%$.

Taxes Paid

For C-corps, we directly observe total tax payments to the federal government on Form 1120. For S-corps, which do not pay entity-level taxes, we must estimate tax payments using information from the individual-level tax records of the firms' owners. To do so, we first compute each owner's average tax rate from Form 1040 as total federal tax divided by taxable income. We also record each owner's total net ordinary business income from Form 1040 Schedule E, and estimate total business taxes paid on this income by multiplying it by the owner's average tax rate. We bottom code total business taxes at zero, ensuring in our calculations that owners do not pay tax on business losses. For each owner, we allocate her total business tax payments to each firm that she owns in proportion to the share of ordinary business income received from that business. Finally, we sum up the total tax payments of each firm's owners to record an estimate of total firm-level tax payments. We provide additional details about these computations in the Online Appendix.

Sales, Costs, and Profits

We measure firm sales as gross receipts. Pre-tax profits are defined as sales minus cost of goods sold, which includes both material and labor inputs. An advantage of this profit measure is that it is simple, transparent, consistent over time, and invariant to tax law and corporate form. As a robustness check, we also construct a harmonized measure of earnings before interest, taxes, depreciation, and amortization (EBITDA), described in the Online Appendix. After-tax profits are equal to pre-tax profits minus taxes paid.

Dividends, Share Buybacks, and Total Payouts

Dividends are defined as total cash and property payments to shareholders. Share buybacks are defined as non-negative changes in treasury stock, and total payouts are measured as the sum of dividends and share buybacks.

Investment

Net investment is defined as the change in the dollar value of capital assets, where capital assets are equal to the book value of tangible investment minus capital asset retirements and

accumulated book depreciation. We also report results on new investment, defined as the sum of capital expenditures reported on IRS Form 4562. These tax forms include information on firms' purchases of new capital assets such as machinery, computers, vehicles, office furniture, and structures. Firms report these investments according to the lifespan of the investment, which affects the horizon of capital tax deductions available to the firm. We decompose new investment into "short-life" equipment with depreciation schedules of less than or equal to 10 years (such as light machinery, computers, and vehicles), "long-life" equipment with longer depreciation schedules (such as heavy machinery), and structures (such as new factories or office buildings).

Employment, Earnings, and Executive Compensation

We measure firm employment as the total number of unique individuals with a W-2 issued by the firm. Firms with complex ownership structures often use multiple employer identification numbers, and we use crosswalks to improve the linkage between W-2s and their ultimate parent companies (see [Joint Committee on Taxation 2022](#)). Workers' annual earnings are defined as Medicare wages from the W-2, which capture wage, tip, and salary income even if it is not taxable. Total firm payrolls are the sum of workers' annual earnings. Because employment, earnings, and payrolls are always strictly positive in our sample, we take logs of these outcomes in the empirical analyses.

Firms report compensation of officers on Forms 1120 and 1120s, which we use to measure executive pay. Officer designations are determined by state tax law, and reported compensation captures several but not all components of executive pay, including: wage, salary, and bonus income; stock options and grants, when exercised; and non-qualified deferred compensation. However, this measure does not capture stock options or grants before they are exercised, and does not include qualified incentive-based compensation plans.⁶ The measure thus represents a lower bound on executive compensation.

We also construct an alternate proxy measure of executive compensation as the combined annual W-2 earnings of the top five highest paid workers at the firm. This measure captures compensation of high-ranking employees who may not qualify as officers for tax purposes.

Additional Firm Characteristics

We group firms into four time-invariant size bins with approximately comparable numbers of observations based on their average employment in the pre-period years prior to 2017, where the bins are: 100-199 employees; 200-499 employees; 500-999 employees; and 1000+

⁶Designation of officers is determined by the laws of the state or country where the firm is incorporated. Qualified deferred compensation plans include 401(k) and similar investment vehicles; these plans are subject to contribution limits and regulatory restrictions, and investments are generally risk-free to workers. Non-qualified deferred compensation plans are not subject to contribution limits and in principle are at risk if the firm declares bankruptcy, although such losses are empirically rare. Qualified incentive-based compensation plans have a maximum deferral of \$100,000 per year and are taxed as long-term capital gains, and thus are not reported on the W-2.

employees. We also classify firms into time-variant industries using the NAICS-3 codes they report on Forms 1120 and 1120s. In the resulting data we observe 86 distinct industries and 280 distinct industry-size bins.

Firm age is inferred from the firm's year of incorporation, reported on the 1120. Firms are classified as multinationals if their foreign sales share is greater than 1%, where foreign sales are defined as the sum of gross receipts from all Controlled Foreign Corporations (that is, foreign subsidiaries) reported on Form 5471. We measure capital intensity at the industry level as capital assets divided by sales. Firms are classified as capital intensive if the mean of this ratio in the pre-period is greater than the sample median, and others are classified as non-capital intensive. The Online Appendix shows that capital intensive firms are approximately equally distributed across the firm size bins, and are not exclusively manufacturing firms.

Data Processing

We scale several outcomes — taxes, sales, costs, and profits — by firm sales in 2016, our baseline year prior to the passage of TCJA. While it is common in economic research to estimate elasticities by transforming regression outcomes using natural logs, doing so in our case is problematic because taxes and profits are often zero or negative in a given year. Scaling firm variables by baseline sales permits a natural economic interpretation of the regression coefficients in our empirical analyses, allows us to study a range of outcomes such that they can be consistently and easily compared, and is standard in the literature. In accordance with economic theory and prior research, investment results are scaled by lagged capital, although for consistency across results we also report results scaled by baseline sales in the Online Appendix. We also follow the literature in winsorizing the top and bottom 0.1% of the scaled outcomes separately for C- and S-corps in each year. Winsorizing ensures that our results are not driven by outliers or by measurement error, and improves statistical precision. We also show that the empirical results are robust to alternate winsorizing thresholds.

Descriptive Statistics

Panels A and B of Figure 2 show the distributions of log firm sales and log firm employment in our sample, and illustrate broad overlap in the size distributions of C- and S-corps. The panels make clear that the firm size distributions are strongly right-skewed, and that this skewness is more pronounced for C-corps than S-corps. In robustness checks, we show that our empirical results are insensitive to the inclusion or exclusion of very large C-corps; since they are qualitatively irrelevant to the results, we include them in the main analysis sample.

Panel C of Figure 2 shows the NAICS-2 industry composition of the sample, and again reveals broad overlap of C- and S-corps. Most industries have comparable shares of C- and S-corps. Some sectors, such as management and professional services, have a relatively higher proportion of C- than S-corps, while the reverse is true for others, such as construction


and retail trade. Because our event study analysis will use industry-size-year fixed effects to compare C- and S-corps in the same industry and employment size bin, the observed sectoral overlap in the sample is more than sufficient for our empirical design. In robustness checks, we show that results are insensitive to the exclusion of industries in which the firm share of C-corps or S-corps exceeds 80%.

Table 1 presents descriptive statistics for our analysis sample from 2016. The mean firm in the sample has annual sales of \$1,046.9 million, earns pre-tax profits of \$385.0 million, pays \$18.1 million in federal taxes, and makes real investments of \$49.6 million per year. Mean firm employment is 2,968 workers, and the average worker earns approximately \$63,700 per year. Consistent with Figure 2, columns 3-10 again underscore the right-skewness of firm size, especially of C-corps, such that mean outcomes are significantly higher than medians and outcomes for C-corps are higher and more variable than for S-corps.⁷


⁷All medians and other quantile statistics reported in this paper are fuzzed to protect taxpayer privacy; see Online Appendix ?? for details.

FIGURE 2: FIRM SIZE DISTRIBUTIONS AND INDUSTRY COMPOSITION


Panel A: 2016 Sales



Panel B: 2016 Employment



Panel C: Industry Composition



Notes: Panels A and B show the distribution of 2016 log firm sales and employment, respectively, for C- and S-corps in the analysis sample. Panel C shows the NAICS-2 industry composition of firms in the sample.

TABLE 1: SUMMARY STATISTICS

	All Firms		C Corporations				S Corporations			
	(1) Mean	(2) SD	(3) Mean	(4) SD	(5) p50	(6) p90	(7) Mean	(8) SD	(9) p50	(10) p90
Taxes										
Marginal Tax Rate	0.248	0.161	0.215	0.168	0.340	0.350	0.310	0.124	0.380	0.396
Federal Tax (mil)	18.1	164.7	26.1	202.7	0.6	24.7	2.9	13.1	0.6	6.8
Federal Tax Per Worker	6,050	11,503	6,382	12,480	1,229	17,567	5,415	9,326	1,669	15,187
Sales and Profits										
Sales (mil)	1,046.9	5,976.2	1,467.0	7,327.1	152.7	2,276.4	244.4	636.3	117.6	463.5
Costs (mil)	654.5	4,667.9	905.8	5,733.4	69.5	1,130.1	174.5	521.9	72.5	345.3
Pre-Tax Profit (mil)	385.0	1,902.1	549.9	2,325.6	55.2	937.6	69.9	216.8	31.9	134.0
After-Tax Profit (mil)	366.9	1,799.7	523.8	2,199.9	52.2	906.3	67.0	211.7	30.2	127.1
EBITDA (mil)	155.9	1,240.2	227.0	1,525.2	11.6	273.0	20.1	64.4	8.1	38.6
Shareholder Payouts										
Dividends (mil)	34.6	444.8	47.6	548.1	0.0	14.6	9.8	33.8	2.9	20.5
Share Buybacks (mil)	19.4	351.2	29.5	433.2	0.0	0.8	0.3	4.6	0.0	0.0
Total Payouts (mil)	56.5	708.4	80.8	873.0	0.0	24.9	10.1	34.1	3.0	20.9
Real Investment										
Net Investment (mil)	17.9	404.4	26.6	498.8	0.0	23.3	1.3	15.0	0.0	5.3
Net Investment / Lagged Capital	0.16	1.36	0.15	1.26	0.01	0.39	0.18	1.53	0.01	0.42
New Investment (mil)	66.7	1,650.6	98.2	2,036.5	3.2	73.0	6.5	36.5	1.2	12.4
Employment and Earnings										
Employment	2,968	21,533	4,041	26,241	510	5,986	918	5,262	340	1,454
Payroll (mil)	173	975	242	1,192	31	402	39	146	18	70
Mean Annual Earnings (thous)	63.7	59.2	68.2	62.8	56.5	111.3	55.1	50.3	49.0	80.9
Median Annual Earnings (thous)	46.2	25.4	49.6	28.0	42.9	84.8	39.7	17.7	37.6	60.2
Executive Pay										
Executives' Earnings (thous)	5,606	26,572	7,319	32,077	1,651	13,654	2,334	8,555	988	4,806
Mean Top 5 Earnings (thous)	1,187	3,329	1,496	3,950	458	3,203	596	1,388	341	1,059
Firm Characteristics										
Firm Age	35	23	33	24	28	64	40	21	37	67
Multinational	0.24		0.32				0.10			
Private	0.84		0.76				1.00			
Capital Intensive	0.50		0.55				0.41			
N Firms	11,647		7,645				4,002			

Notes: Table shows summary statistics from 2016 for firms in the analysis sample. Medians and centile statistics are fuzzed to protect taxpayer privacy. For data sources and variable definitions see Section 1.3.

1.4 Empirical Analysis

Empirical Strategy

We implement a transparent research design comparing trends in outcomes of C- and S-corps in the same industry-size bin before and after TCJA. Our event study specification is given by:

$$y_{ft} = \sum_{t \neq 2016} \beta_t C_f * \mathbf{1}(\text{year} = t) + \gamma_f + \alpha_{is(f),t} + \epsilon_{ft} \quad (1.1)$$

where y_{ft} is an outcome for firm f in year t ; C_f is a binary variable equal to 1 if firm f is a C-corp or 0 if it is an S-corp; γ_f is a firm fixed effect; and $\alpha_{is(f),t}$ is an industry-size-year fixed effect, where the industry-size bins are constructed as described in Section 1.3. The coefficients of interest, β_t , capture the average differences in outcomes between C- and S-corps in the same industry-size bin in year t . We use 2016 as the reference year, allowing us to compare C- and S-corp trends for several years prior to TCJA and also to observe any potential anticipatory tax-shifting behaviors beginning in 2017. Standard errors are clustered by firm.

The key identifying assumption permitting a causal interpretation of the β_t coefficients is that the outcomes of C- and S-corps would have trended similarly in the absence of TCJA's changes to firms' marginal income tax rates. While this assumption is not directly empirically testable, there are several reasons that parallel trends is likely to hold in our setting. First, Congressional passage of TCJA was widely unexpected prior to the 2016 federal elections, and so firms had limited scope to anticipate the reform and to adjust their behavior endogenously to the policy changes. Second, our narrowly defined industry-size-year fixed effects imply that we make comparisons among C- and S-corps that compete in similar product markets and are subject to the same industry-by-size specific supply and demand shocks. Third, [Yagan \(2015\)](#) finds that C- and S-corp trends in real outcomes were statistically indistinguishable for all years in his sample period from 1996-2008, implying that C- and S-corps have historically responded similarly to macroeconomic shocks and trends. Fourth, as we will show, our event studies show parallel trends in the outcomes of C- and S-corps in the years directly prior to the policy reform. Lastly, in Section 1.4, we carefully consider additional identification threats, and present a series of robustness checks to ensure that our causal estimates are not driven by non-MTR features of the law, anticipation effects, superficial tax-shifting behaviors, or unrelated economic shocks differentially affecting C- and S-corps at the same time as TCJA.

Our goal of assessing the efficiency impacts and distributional effects of TCJA's corporate tax cuts will require that we obtain elasticities of profits, investment, and earnings with respect to the net-of-tax rate. To estimate these key elasticities, we pool outcomes in the post-period and use two-stage least squares. The reduced form, first-stage, and structural equations are given, respectively, by:

$$y_{ft} = \lambda C_f * Post_t + \gamma_f + \alpha_{is(f),t} + \epsilon_{ft} \quad (1.2)$$

$$\Delta \ln(1 - \tau_f) * Post_t = \delta C_f * Post_t + \gamma_f + \alpha_{is(f),t} + \epsilon_{ft} \quad (1.3)$$

$$y_{ft} = \varepsilon \Delta \ln(1 - \tau_f) * Post_t + \gamma_f + \alpha_{is(f),t} + \epsilon_{ft} \quad (1.4)$$

where $\Delta \tau_f$ is the 2016 to 2019 change in the marginal income tax rate for firm f , $Post_t$ is an indicator equal to 1 for years after 2018, and the fixed effects are the same as in equation 1.1. Intuitively, we instrument for firms' net-of-tax change using their pre-existing entity type status as a C-or S-corps. The identifying assumptions underlying this empirical strategy are well known: exogeneity, relevance, monotonicity, and exclusion. We do not claim strict exogeneity in our setting – that is, we do not claim there is random assignment of C or S status – but rather rely on the weaker claim of parallel trends in the outcome absent the changes in the tax rate (see [Conley, Hansen, and Rossi 2012](#)). We examine the relevance and monotonicity conditions below, and return to a discussion of the exclusion restriction when we evaluate mechanisms.

We begin the empirical analysis with a presentation of average responses, and then turn to heterogeneity tests and robustness checks. We conclude the empirical analysis with a discussion of mechanisms, where the focus is naturally related to the task of disentangling the impacts of TCJA's marginal tax rate cuts from other concurrent policy changes. First, however, our goal is more modest: to provide clear evidence on how TCJA differentially affected C- and S-corps.


Marginal Tax Rates and Taxes Paid

Figure 3 plots the β_t coefficients and 95% confidence intervals from estimating equation 1.1, using the firms' marginal tax rates and taxes paid as outcomes. We scale taxes paid (and other outcomes that we will report below) by the firm's baseline 2016 sales for the reasons discussed in Section 1.3. In the bottom of each left panel we also report the 2016 sample outcome mean to contextualize the economic scale and significance of the estimated coefficients.


Panel A of Figure 3 shows that the observed marginal tax rates of C- and S-corps trended similarly prior to TCJA, but diverged sharply thereafter. On average, the marginal tax rate of C-corps fell by approximately -5.2 percentage points (s.e.=0.2) compared to S-corps in the sample; relative to the 2016 outcome mean in levels of 0.25, this represents a $-5.2/0.25 \approx 20.8\%$ decline in the marginal tax rate facing C-corps relative to S-corps. The panel also makes clear that firms' tax burdens began to decline in 2017, even though the bulk of TCJA's provisions did not take effect until 2018. This pattern provides suggestive evidence that firms engaged in intertemporal shifting behaviors to minimize tax liability, such as reporting costs in 2017 rather than 2018 so that those costs could be deducted at a higher tax rate. We discuss shifting behaviors in greater detail later in Section 1.4.

FIGURE 3: EVENT STUDIES: MARGINAL TAX RATES AND TAXES PAID


Panel A: Marginal Tax Rate




Panel B: Log Net-of-Tax Rate



Panel C: Tax / 2016 Sales



Panel D: Tax Per Worker



Notes: The unit of analysis is a firm-year. The panels plot the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm and error bands show 95% confidence intervals. The outcome in Panel A is the firm's marginal tax rate, τ_f^{MTR} , and the outcome in Panel B is the log net-of-tax rate, $\ln(1 - \tau_f^{MTR})$. The outcome in Panel C is tax per worker, reported in dollars, and the outcome in Panel D is tax scaled by the firms' baseline 2016 sales. Marginal tax rates for S-corps are defined as the weighted average of the shareholders' individual marginal tax rates, where the weights are given by the ownership shares. See Section 1.3 for details on the measurement of tax payments for S-corps. For data sources and variable definitions see Section 1.3.

Panel B of Figure 3 shows an analogous version of Panel A, where the outcome is transformed as the log net-of-tax rate, $(1-\tau_j^{MTR})$. We show this transformation because economic theory predicts that firms respond to the net-of-tax rate when optimizing profits. The figure shows that, on average, C-corps saw their net-of-tax rate increase by approximately 6.8% (s.e.=0.2) relative to S-corps following TCJA. Below, we use this result to scale other reduced form effects, allowing us to estimate elasticities of key outcomes with respect to changes in the log net-of-tax rate.

Panel C of Figure 3 shows that the differences in tax cuts also translated into differences in taxes paid, with C-corps paying approximately -1.0 percentage points ($\approx 15.0\%$; s.e.=0.3) less in federal tax in 2019 relative to their baseline sales when compared to S-corps. Panel D illustrates that the magnitude of this effect is economically large: on average, C-corps paid approximately \$2,200 (s.e.=\$436) less in tax per worker than comparable S-corps following TCJA.

Columns 1 to 4 of Table 2 report the $C \times Post$ estimates produced from estimating equation 1.2. Similar to the event studies, these coefficients capture the average difference between C- and S-corps in the pre- and post-periods for each outcome after controlling for firm and industry-size-year fixed effects.

The results in Figure 3 and Table 2 provide evidence of a strong first stage, demonstrating an economically meaningful and statistically powerful differential effect of TCJA on the tax rates and tax payments of C-corps versus S-corps. These results also show that the relevance and monotonicity assumptions underlying equation 1.3 are satisfied in this setting.

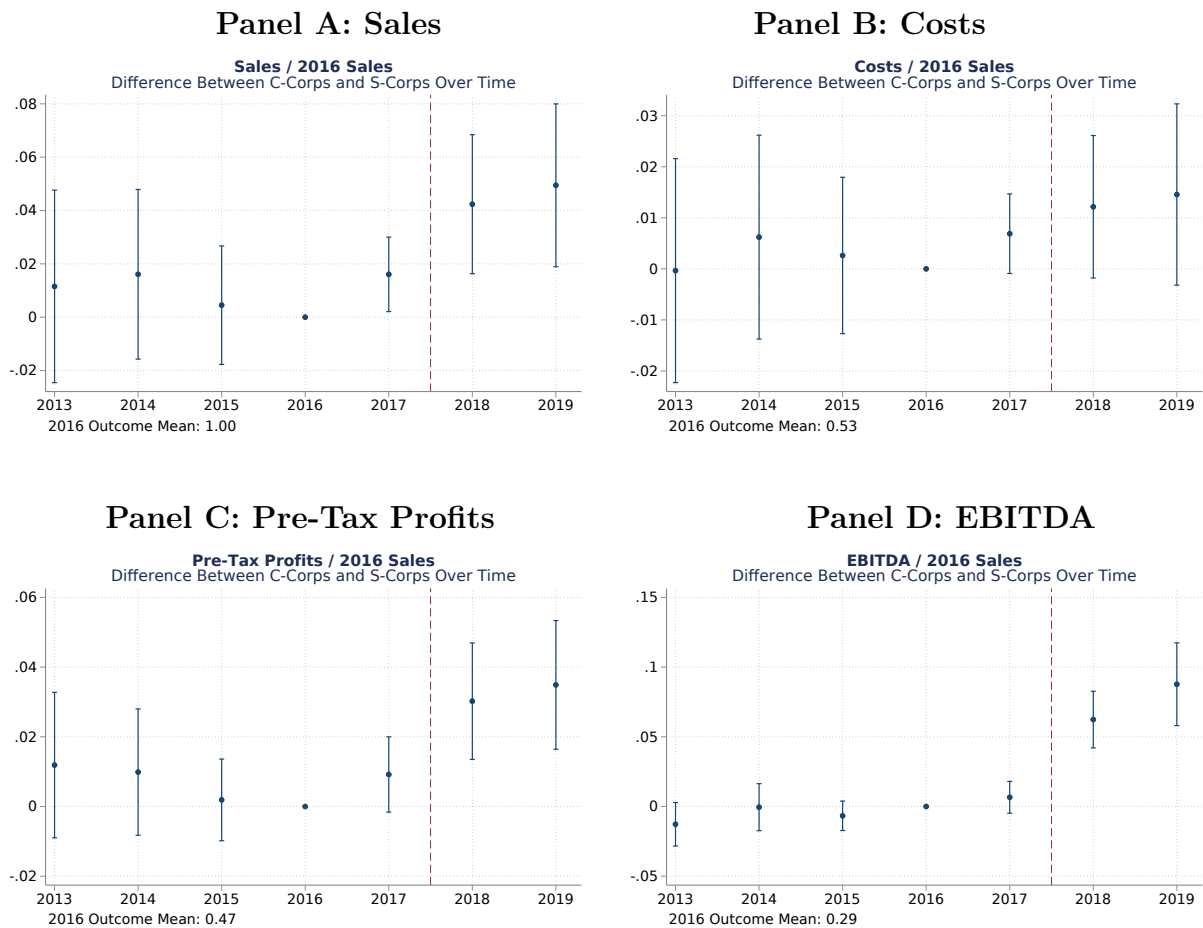
Sales, Costs, Pre-Tax Profits, and EBITDA

Figure 4 plots the results from estimating equation 1.1 to assess trends in the sales, costs, and pre-tax profits, and EBITDA of C- and S-corps over time. The figure shows trends in these outcomes were statistically similar before TCJA, again lending support to the parallel trends assumption underlying the identification strategy. After TCJA, however, C-corps' sales increased markedly relative to S-corps, by approximately 3.6 percentage points (s.e.=1.5) by 2019. The effect is precisely estimated and economically significant: using values from Table 1, the coefficient implies that the average C-corp increased its sales by approximately \$40 million relative to comparable S-corps.

C-corps also faced higher costs, as shown in Panel B, although the magnitude of the cost increase is smaller than for sales and, on average, is not statistically significant. Later we show that this average effect masks important heterogeneity, with both sales and costs increasing predominantly in capital intensive industries.

Given sharply increasing sales and only modestly increasing costs, Panel C shows that the average pre-tax profits of C-corps also increased relative to S-corps, by 2.6 percentage points (s.e.=0.9). Panel D shows an alternate measure of pre-tax profits, using the harmonized EBITDA measure, and again reveals a clear increase in the profits of C-corps relative to S-corps. These results provide initial evidence that firms expanded in response to tax cuts,

FIGURE 4: EVENT STUDY: FIRM SALES, COSTS, AND PRE-TAX PROFITS



Notes: Unit of analysis is firm-year. The panels plot the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm, and error bands show 95% confidence intervals. Sales are gross receipts. Costs are equal to cost of goods sold, including both material and labor costs. Pre-tax profits are sales minus costs. EBITDA is a harmonized measure of earnings before interest, taxes, depreciation, and amortization; see Section 1.3 and the Online Appendix for details.

TABLE 2: MARGINAL TAX RATES AND TAXES PAID

	(1)	(2)	(3)	(4)
	τ_f^{MTR}	$\ln(1 - \tau_f^{MTR})$	Tax Per Worker	Tax/Sales ₂₀₁₆
C × Post	-0.052*** (0.002)	0.068*** (0.002)	-2205.482*** (436.454)	-0.010*** (0.003)
2016 Outcome Mean	0.25	-0.31	6,046	0.06
Firm FE	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes
R2	0.72	0.73	0.59	0.83
N	81,515	81,515	81,515	81,515
N Firms	11,645	11,645	11,645	11,645

Notes: The unit of analysis is a firm-year. The table shows the $C \times Post$ coefficients from equation 1.2. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. The outcome in column 1 is the firm’s marginal tax rate, τ_f^{MTR} , and the outcome in column 2 is the log net-of-tax rate, $\ln(1 - \tau_f^{MTR})$. The outcome in column 3 is tax per worker, reported in nominal dollars, and the outcome in column 4 is tax scaled by the firms’ baseline 2016 sales. Marginal tax rates for S-corps are defined as the weighted average of the shareholders’ individual marginal tax rates, where the weights are given by the ownership shares. See Section 1.3 for details on the measurement of tax payments for S-corps. Standard errors are clustered by firm.

consistent with the standard notion that taxes induce economic distortions and may generate deadweight loss.

Columns 1 to 4 of Table 3 show the $C \times Post$ coefficients associated with the event studies in Figure 4. In column 5, we estimate the elasticity of pre-tax profits with respect to the net-of-tax rate using equation 1.4. Scaling the reduced form estimate in column 3 by the first-stage estimate from column 2 of Table 2 yields an elasticity of 0.38 (s.e.= 0.13).

This elasticity – known as the elasticity of the tax base, or alternately as the elasticity of taxable income (ETI) – is a key parameter in the analysis. As shown by Feldstein (1999) and reviewed by Saez, Slemrod, and Giertz (2012), under plausible assumptions it is a sufficient statistic that can be used to estimate the welfare impacts and efficiency costs of tax changes. In general, a larger taxable income elasticity implies greater deadweight loss, since it implies a larger distortion of economic activity resulting from the tax.

Our estimate of the federal corporate ETI, 0.38, is on the lower end of corporate elasticities identified from policy variation in small open economies. For example, Giroud

and Rauh (2019) estimate an elasticity of establishment growth (a proxy for the tax base) of approximately 0.50 with respect to state corporate taxes in the United States; Suárez Serrato and Zidar (2016) estimate an elasticity of establishment growth of approximately 0.9 for U.S. state corporate taxes over an analogous time horizon to ours; and Bachas and Soto (2021) estimate large taxable income elasticities of 3.0-5.0 for small firms in Costa Rica. On the other hand, our estimate of the corporate ETI is on the higher end of most existing estimates of the ETI for personal incomes, which Saez, Slemrod, and Giertz (2012) find in a literature review ranges from approximately 0.14 to 0.40, with a central estimate of 0.25.

Viewed in the context of other research, we view our corporate ETI estimate as consistent with the common economic intuition that tax distortions vary with factor mobility. Firms and workers are less mobile at the federal level than at the state and local level, mitigating distortions from the federal corporate tax relative to the state and local corporate tax. However, many forms of capital are internationally mobile relative to workers (Kotlikoff and Summers 1987), suggesting that federal taxes on labor income, the primary source of personal income tax revenue, may be less distortive than the federal corporate tax.

In Section 1.4 we perform extensive robustness checks on our ETI estimate, and in Section 3.5 we discuss its significance in the context of the broader national tax and transfer system.

TABLE 3: SALES, COSTS, AND PRE-TAX PROFITS

	(1)	(2)	(3)	(4)	(5)
	Sales	Costs	Pre-tax π	EBITDA	Pre-tax π
$C \times Post$	0.036** (0.015)	0.010 (0.009)	0.026*** (0.009)	0.078*** (0.011)	
$\Delta \ln(1 - \tau_f) \times Post$					0.379*** (0.127)
2016 Outcome Mean	1.00	0.53	0.47	0.29	0.47
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes	Yes
R2	0.40	0.65	0.62	0.84	n.a.
N	81,515	81,515	81,515	81,515	81,515
N Firms	11,645	11,645	11,645	11,645	11,645
First-Stage F					409.3

Notes: The unit of analysis is a firm-year. Columns 1-4 show the $C \times Post$ coefficients from equation 1.2. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. All outcomes are scaled by 2016 baseline sales. Sales are gross receipts. Costs are equal to cost of goods sold, including both material and labor costs. Pre-tax profits are sales minus costs. EBITDA is a harmonized measure of earnings before interest, taxes, depreciation, and amortization. Column 5 reports the elasticity of pre-tax profits with respect to the net-of-tax rate, computed by scaling the reduced form outcome in column 3 by the first stage coefficient from column 2 of Table 2. Standard errors are clustered by firm. For additional information on data sources and variable definitions see Section 1.3 and the Online Appendix.

After-Tax Profits and Shareholder Payouts

We use the same empirical strategy to evaluate trends in firms' after-tax profits and payouts to shareholders. Consistent with the increases in pre-tax profits and decline in tax liability, Panel A of Figure 5, and Column 1 of Table 4, shows that the after-tax profits of C-corps increased relative to S-corps following TCJA, by 3.6 percentage points ($\approx 10.7\%$; s.e.=0.9). The magnitude of this effect is economically and statistically significant, and underscores that tax cuts are highly lucrative to firm owners. The elasticity of after-tax profits with respect to the net-of-tax rate, estimated in column 4 of Table 4 using equation 1.4, is 0.52

(s.e.= 0.13). Later, we use this elasticity to assess the incidence of TCJA's tax cuts on firm owners.

We also find that firms returned some of these excess profits to their shareholders via dividends and share buybacks, the sum of which we refer to as total shareholder payouts. Because shareholder payouts are infrequent events (approximately half of the payout observations in our sample are zero), we study both the extensive and intensive margins. The outcome in Panel B of Figure 5 is equal to one if total payouts are greater than zero (that is, the extensive margin), and shows an increase of 4.0 percentage points (s.e.=0.8) in the payout probability of C-corps relative to S-corps following TCJA. In Panel C we show the intensive margin, where the outcome is log total payouts, and find that payouts of C-corps relative to S-corps increase by 21.9% (s.e.=5.0). Consistent with this increase in shareholder payouts, in the Online Appendix we find that C-corps do not increase their issuance of equity or debt relative to S-corps after TCJA. The results are consistent because, if firms need external financing to fund operations, they generally do not simultaneously distribute cash to shareholders.⁸

Overall, the results from Figure 5 and Table 4 provide evidence that firm owners bear a substantial portion of the short-run economic incidence of the corporate income tax.

Labor Market Outcomes and Executive Pay

We again use equation 1.1 to examine the labor market outcomes of workers at C- and S-corps before and after TCJA. Figure 6 shows the results from estimating equation 1.1 to assess trends in log employment, total payroll, and annual earnings for selected groups of workers.


Figure 6 shows that the labor market outcomes of C- and S-corps followed similar trends prior to TCJA. After TCJA, employment in C-corps increased modestly relative to S-corps, by 0.4% (s.e.=0.8) on average, but the difference is not statistically distinguishable from zero. Total payrolls, shown in Panel B, also increased modestly in C-corps relative to S-corps, by 1.2% (s.e.=0.8), and again the difference is not statistically significant. However, later we show that these average effects mask important heterogeneity across firms.

Panels C and D move beyond total payrolls to shed light on the distributional impacts of TCJA on workers' earnings. Panel C shows that the earnings of the median worker at the firm evolved similarly for both C- and S-corps over the entire sample period, and implies that corporate tax cuts did not have a statistically significant effect on earnings for the typical


⁸In assessing the effects of TCJA on shareholder payouts, it is also relevant that enactment of the TCJA levied a one-time, mandatory tax on the previously untaxed foreign earnings of C-corporations through section 965 of the Internal Revenue Code. The American Jobs Creation Act in 2005 included a similar provision in which c-corporations could voluntarily repatriate foreign earnings at a reduced rate. Research on the 2005 repatriation holiday has suggested that the primary effect of this provision was to increase shareholder payouts (Dharmapala, Foley, and Forbes 2011) and potentially increased investment among some credit constrained firms (Faulkender and Petersen 2012). This one time repatriation tax is a potential confounder in our setting, but our results are robust to inclusion of controls for foreign earnings as well as the exclusion of large multinational firms, suggesting that is not the primary driver of our findings.

FIGURE 5: EVENT STUDY: AFTER-TAX PROFITS AND SHAREHOLDER PAYOUTS


Panel A: After-Tax Profits



Panel B: Shareholder Payouts (Extensive Margin)



Panel C: Shareholder Payouts (Intensive Margin)



Notes: The unit of analysis is a firm-year. The panels plot the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm and error bands show 95% confidence intervals. In Panel A, after-tax profits are defined as pre-tax profits minus tax, and are scaled by 2016 baseline sales. In Panel B, the outcome is an indicator equal to 1 if shareholder payouts are positive (i.e., the extensive margin), where payouts are defined as the sum of cash and property distributions to shareholders. In Panel C, the outcome is log shareholder payouts (i.e., the intensive margin). For additional information on data sources and variable definitions see Section 1.3 and the Online Appendix.

TABLE 4: AFTER-TAX PROFITS AND SHAREHOLDER PAYOUTS

	(1)	(2)	(3)	(4)
	Post-Tax π	Payouts (0/1)	Log Payouts	Post-Tax π
$C \times \text{Post}$	0.036*** (0.009)	0.034*** (0.006)	0.246*** (0.034)	
$\Delta \ln(1 - \tau_f) \times \text{Post}$				0.521*** (0.129)
2016 Outcome Mean	0.41	0.54	1.45	0.41
Firm FE	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes
R2	0.69	0.76	0.86	n.a.
N	81,515	81,515	81,515	81,515
N Firms	11,645	11,645	11,645	11,645
First-Stage F				409.3

Notes: The unit of analysis is firm-year. Columns 1-3 show the $C \times \text{Post}$ coefficients from equation 1.2. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. After-tax profits are defined as pre-tax profits minus tax, and are scaled by 2016 baseline sales. In column 2, the outcome is an indicator equal to 1 if shareholder payouts are positive (i.e., the extensive margin), where payouts are defined as the sum of cash and property distributions to shareholders. In column 3 the outcome is log shareholder payouts (i.e., the intensive margin). Column 4 reports the elasticity of after-tax profits with respect to the net-of-tax rate, computed as in equation 1.4. Standard errors are clustered by firm. For additional information on data sources and variable definitions see Section 1.3 and the Online Appendix.

worker.⁹ By contrast, Panel D shows that the earnings of higher-income C-corp workers increased sharply relative to their counterparts in S-corps.

To more comprehensively evaluate the effects of TCJA on the distribution of workers' earnings, we estimate quantile regression specifications of equation 1.1, where the outcome $y_{ft}(p)$ is log annual earnings of workers in firm f and year t at centile p . For example, $y_{ft}(p = 50)$ uses log median earnings as the outcome, as shown in Panel C of Figure 6, and $y_{ft}(p = 99)$ uses the 99th percentile of log worker earnings as the outcome, as in Panel D.

Figure 7 plots the β_{2019} coefficients from these quantile regressions along with their corresponding 95% confidence intervals. The figure shows that the relative earnings of workers in C- and S-corps below the 90th percentile are statistically identical following TCJA; we cannot reject that the coefficients are statistically distinguishable from zero.

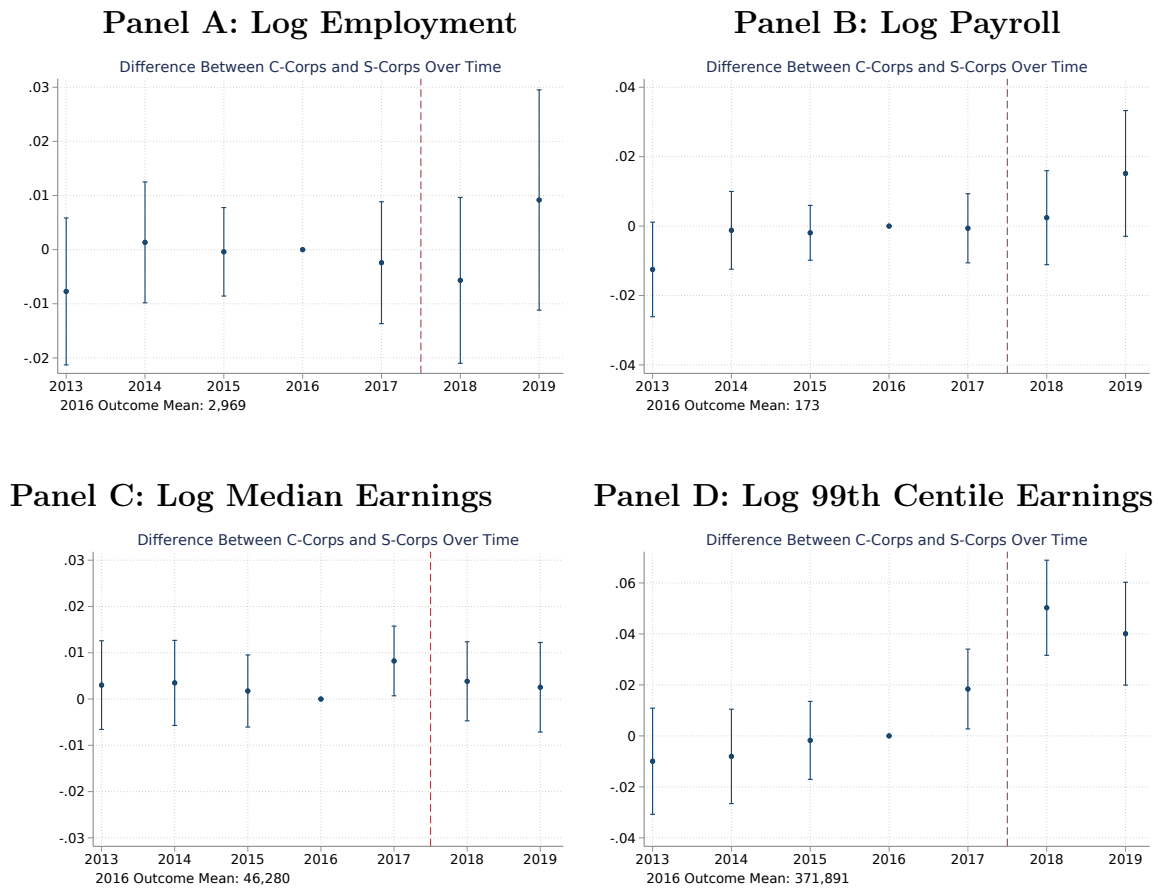
However, Figure 7 reveals a very different pattern for workers in the top 10% of the earnings distribution. Workers in C-corps at the 90th percentile of the within-firm distribution see their relative earnings increase by 1.0% (s.e.=0.3), and these impacts grow steadily larger and statistically sharper further up the distribution. At the 95th percentile, we estimate a relative earnings increase of 1.2% (s.e.=0.4) for C-corp workers, and this magnitude climbs to 4.5% (s.e.=0.8) at the 99th percentile.

We further assess the impacts of MTR cuts on executive pay. Figure 8 estimates equation 1.1 using as outcomes log officer compensation (observed on IRS Forms 1120 and 1120s) and a proxy variable constructed as the log mean earnings of the top five highest-paid workers at the firm (observed from IRS Form W2). In Panel A, we estimate that the relative earnings of executives increased by 4.6% (s.e.=1.2) at C-corps relative to S-corps, and in Panel B we estimate a quantitatively comparable effect for the earnings of the top 5 highest paid workers at the firm of 4.1% (s.e.=0.8). Because the tax data do not allow us to observe all components of executive compensation, such as awarded but unvested stock grants, these estimates likely represent a lower bound on the effects of TCJA on executive pay. The fact that executive earnings increase in 2017, before TCJA fully took effect, is consistent with firms intertemporally shifting forward executive compensation, perhaps in the form of bonuses, so that these costs could be deducted at a higher tax rate prior to the corporate rate cut beginning in 2018.

Panel A of Table 5 reports the $C \times Post$ coefficients from equation 1.2, as well as the dependent variable means in the baseline year and implied elasticities with respect to the net-of-tax rate. For workers at the 95th percentile, we estimate an earnings elasticity of 0.18% (s.e.=0.05), and for executives we estimate a larger earnings elasticity of 0.65% (s.e.=0.17). The mean baseline earnings of these workers and executives are high: the average worker in the sample at the 95th percentile of the within-firm distribution earns \$157,652 per year, the average worker in the top five earns \$1,186,896 per year, and the average annual combined


⁹Models of perfect labor competition predict that, in response to an increase in labor demand, wages should increase equally for workers of both C- and S-corps. It is possible that these general equilibrium effects are absorbed in our industry-size-year fixed effects. Our main results abstract from this potential response. We discuss theoretical and empirical strategies for considering potential general equilibrium effects of corporate tax cuts on outcomes such as earnings, profits, and investment in the Online Appendix.

FIGURE 6: EVENT STUDIES: LABOR MARKET OUTCOMES



Notes: Unit of analysis is firm-year. The panels plot the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm and error bands show 95% confidence intervals. Employment, payrolls, and annual earnings are computed by matching worker-level W-2's with firm-level tax returns. For additional details on data sources and variable definitions see Section 1.3.

FIGURE 7: EARNINGS QUANTILE REGRESSIONS




Notes: Unit of analysis is firm-year. Figure plots the β_{2019} coefficients from equation 1.1, where the outcomes are centiles of the distribution of workers' wages within the firm. For example, centile 50 measures the annual earnings of the median worker within the firm, and centile 90 captures the annual earnings of the worker at the 90th percentile of the within-firm earnings distribution. For additional details on data sources and variable definitions, see Section 1.3. Standard errors are clustered by firm and error bands show 95% confidence intervals.

earnings of firm executives is \$6,290,490. Applying the baseline sample levels, the average firm net-of-tax-rate change in the sample, and the estimated net-of-tax elasticity, the results imply that average combined executive earnings increased by approximately \$270,000 per year. Similar computations yield that average earnings for workers at the 95th percentile increased by approximately \$1,900 per year, and that earnings changes for workers below the 90th percentile are statistically indistinguishable from zero.

What drives the sharp increases in executive pay? The coinciding increases in the sales, pre-tax and after-tax profits, and investment of C-corps compared to S-corps suggest there is plausible scope for managerial decisionmaking and effort to drive increased firm productivity. Moreover, in some cases firm owners may incentivize managerial effort by explicitly compensating executives on the basis of firm performance metrics (e.g., [Bebchuk and Fried 2003](#); [Murphy 1999](#); [Jensen and Murphy 1990](#)). On the other hand, to the extent

FIGURE 8: EVENT STUDIES: EXECUTIVE PAY



Notes: Unit of analysis is firm-year. The panels plot the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm and error bands show 95% confidence intervals. The right panels are constructed such that the distance between the C-corp and S-corp lines in each year is equal to the corresponding β_t coefficient in the left panel, and such that the observation-weighted average of the two lines is equal to the unweighted sample average of the outcome in each year. For data sources and variable definitions see Section 1.3.

TABLE 5: LABOR MARKET OUTCOMES

Panel A: Labor Market Outcomes								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Emp	Payroll	p50	p90	p95	p99	Top5	Executives
$C \times Post$	0.004 (0.008)	0.012 (0.008)	-0.000 (0.004)	0.010*** (0.003)	0.012*** (0.004)	0.045*** (0.008)	0.041*** (0.008)	0.046*** (0.012)
2016 Outcome Mean	2,969	173	46,280	113,856	157,652	371,891	1,186,896	6,290,490
ϵ^{NTR}	0.05	0.18	-0.00	0.14	0.18	0.64	0.59	0.65
s.e.	0.12	0.11	0.05	0.05	0.05	0.12	0.12	0.17
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R2	0.97	0.97	0.95	0.96	0.95	0.88	0.92	0.92
N	81,515	81,515	81,515	81,515	81,515	81,515	81,515	72,385
N Firms	11,645	11,645	11,645	11,645	11,645	11,645	11,645	10,678

Panel B: Executive Compensation

Outcome is log executive pay				
	Controls for:			
	Benchmark	Sales	Profits	Relative Sales
$C \times Post$	0.046*** (0.012)	0.041*** (0.012)	0.041*** (0.012)	0.042*** (0.012)
Firm FE	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes
R2	0.92	0.92	0.92	0.92
N	72,400	72,400	72,400	72,198
NF	10,680	10,680	10,680	10,678

Panel C: Worker Characteristics

	Bottom 90%	Top 10%	Executives
Mean Wage (2016)	39,688	175,957	1,063,511
Female (Share)	0.45	0.30	0.12
Age (Years)	39.5	46.7	53.5
Firm Tenure (Years)	4.2	5.7	6.0
N Workers	32,386,875	3,924,113	72,117

Notes: Unit of analysis is firm-year. Panel A reports results the $C \times Post$ coefficients obtained from estimating equation 1.2, where the outcomes are log employment, log payroll, log annual earnings of workers at various centiles of the within-firm income distribution, log mean earnings of the top 5 highest paid workers at the firm, and log executive compensation. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. Panel B estimates variations of equation 1.2 where the outcome is log executive compensation, and adds time-varying controls for several measures of firm performance. Standard errors in Panels A and B are clustered by firm. Panel C presents descriptive statistics for the individual characteristics of workers in the bottom 90% of the distribution, in the top 10%, and of the top five highest paid workers, where we use the latter as a proxy for executives.

that executives have significant bargaining power vis-a-vis shareholders, they may be in a position to extract a portion of after-tax profits even in the absence of improvements in managerial productivity.

In Panel B of Table 5 we perform a series of empirical tests developed by [Ohrn \(2022\)](#) to evaluate the relevance of these competing mechanisms, which are not necessarily mutually exclusive. The outcome in all columns is log officer compensation as reported on Forms 1120 and 1120s. The first column shows the benchmark specification given by equation 1.2. In the remaining columns, we respectively add controls for three measures of firm performance: sales growth, profit growth, and sales growth relative to other firms in the same industry.¹⁰ To the extent that executive pay is correlated with these performance metrics, we may expect the $C \times Post$ coefficient to shrink as we add the controls.

The results in columns 2-4 show that the $C \times Post$ coefficient on executive pay shrinks only modestly as we add controls for the firm performance metrics. The benchmark estimate of 4.6% declines to a minimum of 4.1% when we add controls for sales growth, and shrinks by a similar amount when controlling for profit growth or sales growth relative to other firms in the same industry. Taken at face value, a plausible estimate is that $1 - (4.1/4.6) \approx 10\%$ of the increase in executive pay is plausibly attributable to improved firm performance, while the remaining 90% may reflect rent-sharing mechanisms. The results are similar when we use our proxy measure of executive pay, reported in the Online Appendix.

These tests are not dispositive — the econometric problems with conditioning on post-treatment outcomes are well-known (e.g., [Imbens 2020](#)), and increasing managerial productivity may not be fully reflected in the firm performance metrics over our limited time horizon — but they are suggestive, and give an approximate sense of plausible orders of magnitude. The results are consistent with empirical evidence from [Ohrn \(2022\)](#), who finds that executive pay in publicly traded firms is highly responsive to narrowly targeted corporate tax breaks, and that pay increases are driven by rent-sharing rather than a higher marginal product of labor. The results are also consistent with [Bertrand and Mullainathan \(2001\)](#), who find that executives are often rewarded for positive shocks to the firm even if those shocks are clearly beyond the manager’s control.

To provide additional insight into the distributional impacts of corporate tax cuts on workers’ earnings, Panel C of Table 5 presents descriptive statistics on the individual characteristics of workers in the bottom 90% of the firm wage distribution, in the top 10% of the firm wage distribution, and in the group of top five highest paid workers at the firm, which we use a proxy for identifying executives (we do not directly observe which individuals are executives from Form W-2). In our sample, 88% of executives are men, and on average these workers are 53 years old, have worked for their employer for 6 years, and earn over \$1 million in annual labor income. By contrast, just 55% of workers in the bottom 90% of the distribution are men, and these workers on average are 39 years old, have worked for their employer for 4 years, and earn less than \$40,000 in annual labor income.

¹⁰As in [Ohrn \(2022\)](#), another natural performance metric would be earnings per share; however, we do not observe this information for S-corporations or for private C-corporations.

Collectively, the findings from Table 5 imply that the short-run effects of corporate tax cuts are regressive, increasing earnings only for workers at the top of the within-firm distribution. The results also demonstrate that the distributional impacts of corporate tax cuts do not affect all demographic groups equally; rather, the beneficiaries of the tax cuts are disproportionately likely to be men, to be older, and to have longer tenures at their current employer.

Our results are consistent with studies finding evidence of rent-sharing with high-income workers in response to tax or productivity shocks (Ohrn 2022; Carbonnier, Malgouyres, Py, and Urvoy 2022; Gale and Thorpe 2022; Dobridge, Landefeld, and Mortenson 2021; Kline, Petkova, Williams, and Zidar 2019), but are in tension with other studies finding that the incidence of business income taxes falls substantially on low-income and marginally attached workers (Risch 2021; Fuest, Peichl, and Siegloch 2018). Notably, the former studies finding gains for high-income workers focus on tax cuts, whereas the latter studies finding that costs borne by low-income workers focus on tax increases. A plausible reconciliation of the literature is that tax cuts and tax hikes may have asymmetric labor market effects (see also discussion in Fuest, Peichl, and Siegloch 2018); we view this hypothesis as a fruitful area for future research.

In the Online Appendix, we study whether causal effects vary by demographic characteristics, and find no evidence of heterogeneous effects after controlling for workers' initial place in the income distribution. We explore additional aspects of heterogeneity in Section 1.4, and return to broader issues of assessing the incidence of the corporate income tax in Section 1.5.


Investment

Figure 9 shows the results from estimating equation 1.1 to assess relative trends in real net investment of C- and S-corps. As with shareholder payouts, investment is a statistically volatile outcome, and so we investigate both the extensive and intensive margin responses. In Panel A, the outcome is an indicator equal to one if net investment is positive (that is, the extensive margin), where net investment is defined as the change in book value of depreciable capital assets less accumulated book depreciation. In any given year, net investment is negative for approximately half the firms in our sample. The figure shows that C- and S-corps have similar trends in this outcome over the pre-period, but after TCJA positive net investment among C-corps increases by approximately 2.2 percentage points (s.e.=0.9) relative to S-corps.

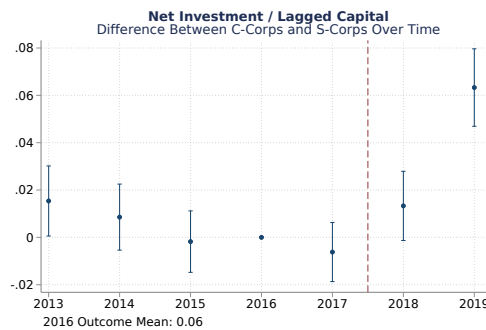
On the intensive margin, Panel B of Figure 9 shows trends in the investment rate, defined as net investment scaled by the lagged capital stock. We find that C-corps increase net investment by 3.5% (s.e.=0.9) relative to S-corps after TCJA. For consistency with other previously reported outcomes, in Panel C we also show net investment scaled by baseline 2016 sales. The figure again shows that investment of C-corps increases relative to S-corps, by approximately 0.8 percentage points (s.e.=0.5). In the Online Appendix, we provide additional results on investment using alternate measures.

FIGURE 9: EVENT STUDIES: NET INVESTMENT


Panel A: Positive Net Investment (0/1)



Panel B: Net Investment / Lagged Capital



Panel C: Net Investment / 2016 Sales



Notes: Unit of analysis is firm-year. The figure plots the β_t coefficients estimated from equation 1.1. These coefficients capture average differences in outcomes between C- and S-corps over time after controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm and error bands show 95% confidence intervals. Net investment is defined as the change in book value of depreciable capital assets minus accumulated book depreciation. The outcome in Panel A is an indicator equal to 1 if net investment is positive. The outcomes in Panels B and C, are net investment scaled by lagged capital and by baseline 2016 sales, respectively. For data sources and variable definitions see Section 1.3.

TABLE 6: NET INVESTMENT

	(1)	(2)	(3)	(4)
	$I_t > 0$	I_t/K_{t-1}	$I_t//textSales_{2016}$	I_t/K_{t-1}
$C \times Post$	0.022**	0.035***	0.008***	
	(0.009)	(0.005)	(0.003)	
$\Delta \ln(1 - \tau_f) \times Post$				0.516***
				(0.082)
2016 Outcome Mean	0.52	0.06	0.01	0.06
Firm FE	Yes	Yes	Yes	Yes
Industry-Size-Year FE	Yes	Yes	Yes	Yes
R2	0.29	0.22	0.26	n.a.
N	81,515	81,515	81,515	81,515
N Firms	11,645	11,645	11,645	11,645
First-Stage F				403.3

Notes: The unit of analysis is a firm-year. Columns 1-3 report the $C \times Post$ coefficients from equation 1.2. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm. Net investment is defined as the change in book value of depreciable capital assets minus accumulated book depreciation. The outcome in column 1 is an indicator equal to 1 if net investment is positive. The outcomes in columns 2 and 3 are net investment scaled by lagged capital and by baseline 2016 sales, respectively. Column 4 reports the elasticity of net investment with respect to the net-of-tax rate, computed from equation 1.4.

The elasticity of investment may have implications for economic growth (e.g., [Romer and Romer 2010](#)) and, in our setting, has direct implications for assessing the incidence of the corporate tax on capital suppliers (e.g., as in [Goolsbee 1998](#)). In column 4 of Table 6, we estimate an investment elasticity of 0.52 (s.e.=0.08), implying that a 1% increase in the net-of-tax rate causes an approximately 0.52% increase in investment. Later, we also estimate the elasticity of investment with respect to the cost of capital. Before doing so, however we must first investigate whether changes in investment are driven by changes in the cost of capital, as in a standard model, or by others channels such as liquidity effects. We thus turn now to a battery of heterogeneity and robustness tests, and then turn to an explicit discussion of mechanisms.

Firm Heterogeneity

Figure 10 presents our benchmark difference-in-difference estimates (i.e., the $C \times Post$ coefficients and 95% confidence intervals from equation 1.2) across several dimensions of firm heterogeneity, focusing on the following outcomes: pre-tax profits, costs, employment, payroll, and net investment. We focus on these outcomes due to their natural relevance in assessing whether our estimates are driven by changes real economic activity or by tax and profit shifting behaviors by firms; we also present results on other outcomes in the Online Appendix.

Existing research has emphasized that the effects of tax changes may vary by firm size (where smaller firms may be better able to engage in tax shifting, as in [Giroud and Rauh 2019](#)); by liquidity (where low-cash firms may face borrowing constraints and thus respond more elastically to tax changes, as in [Zwick and Mahon 2017](#) and [Saez, Schoefer, and Seim 2019](#)); by factor intensity (where capital-intensive firms may be most responsive to a shock, as in [Acemoglu and Guerrieri 2008](#)); by firm profitability (where highly profitable firms may be managed more effectively, as suggested by [Bloom and Van Reenen 2007](#)); by unionization rates (where highly unionized firms may reduce firms' profits and investment, as studied in [Card, Devicienti, and Maida 2014](#)); and by industry concentration (where highly concentrated firms may be better able to pass the costs of tax increases to their input suppliers, as in [Fuest, Peichl, and Sieglöcher 2018](#) and [Juarez 2022](#)).

Here we focus on heterogeneity across the first three of these characteristics — firm size, liquidity, and capital intensity — and report additional heterogeneity tests in the Online Appendix. Firm size is defined using the the pre-TCJA employment bins used in our main analysis, although here we exclude very large firms (defined as those with greater than \$1 billion in sales or greater 10,000 employees in 2016) to ensure that results in the largest firm size bin are not driven by C-corps with no comparably sized S-corps. We measure cash as the sum of the firm's liquid assets, and classify firms as high-cash if their average cash-to-assets ratio in the pre-period is greater than the median value for the sample. Capital intensity is defined as in Section 1.3.

When we test for heterogeneity across firm size we include only industry-year fixed effects in our regression specifications; for all other specifications we include industry-size-year fixed

effects. To obtain the point estimates in Figure 10 we run the model separately for each subsample of firms.

Looking across the outcomes in Figure 10, we find no clear pattern with respect to firm size. With respect to liquidity, we similarly do not find statistically significant differences between high- and low-cash firms; although, taking the point estimates at face value, the results suggest that high-cash firms are if anything weakly *more* responsive than low-cash firms in our sample. These results contrast with [Zwick and Mahon \(2017\)](#), who find that small and financially-constrained firms are most likely to increase investment and payrolls in response to bonus depreciation incentives. While we find that smaller firms have larger point estimates than larger firms, these differences are not statistically significant.

Several factors may explain why our findings differ from [Zwick and Mahon \(2017\)](#). First, Zwick and Mahon study countercyclical policies enacted during U.S. recessions, when financial constraints are most likely to be binding. By contrast, TCJA was enacted during a long macroeconomic expansion with low interest rates and favorable financial conditions. Second, Zwick and Mahon find that responses are largest for the smallest firms in their sample. By contrast, our sample includes only medium-to-large sized firms. Finally, Zwick and Mahon use an identification strategy that exploits cross-industry exposure to bonus depreciation incentives. By contrast, our identification strategy exploits within-industry variation in tax policy, and as such our industry-size-year fixed effects may absorb any time-varying policy impacts that affect both C- and S-corps similarly.

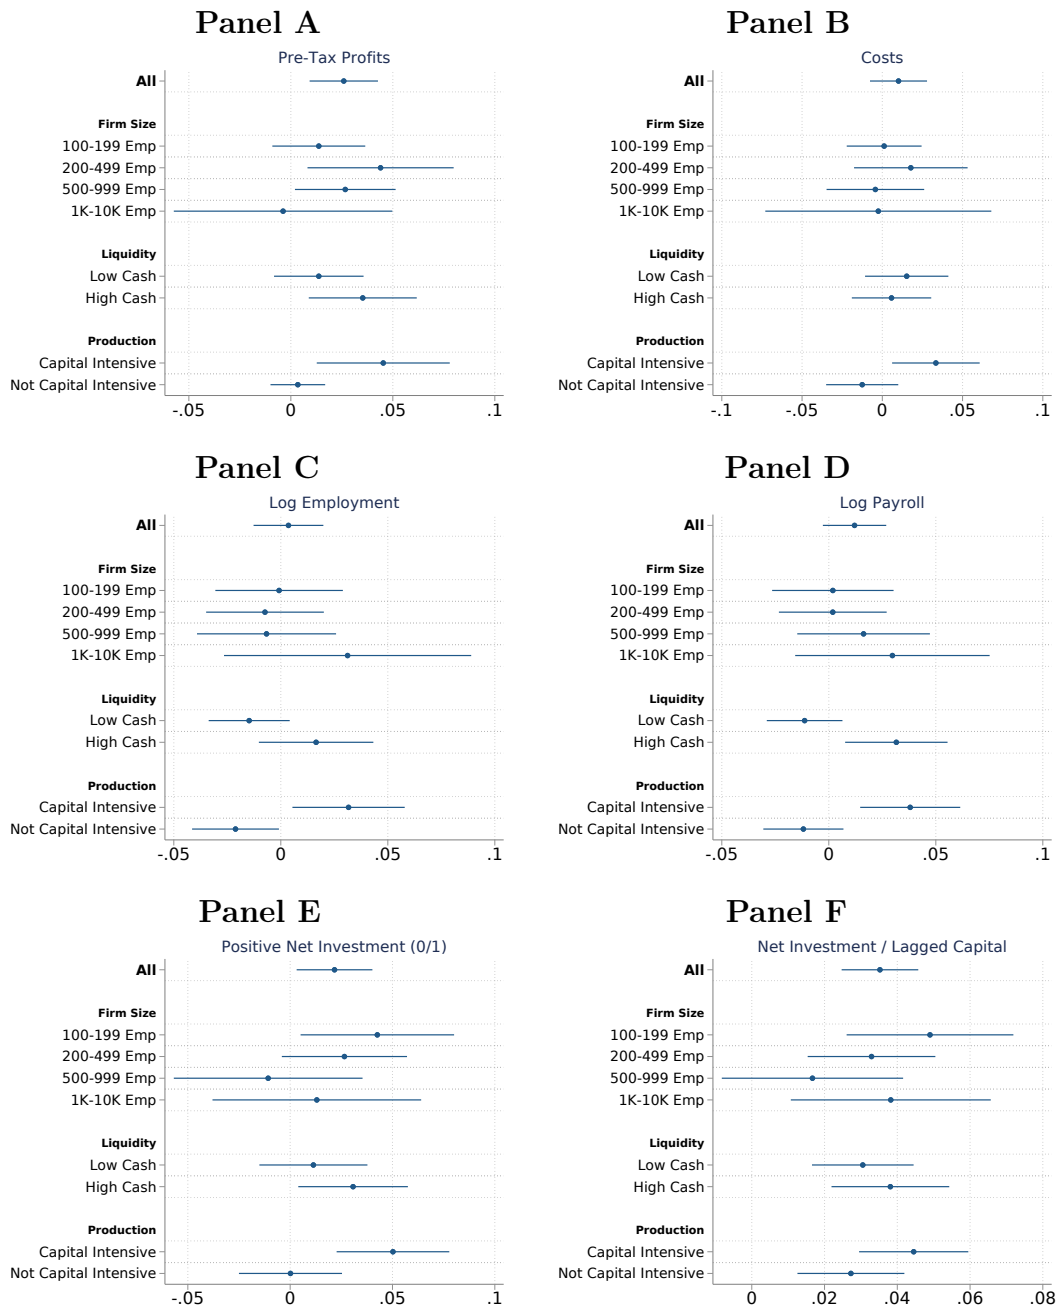
Although policy impacts do not appear to vary with firm size or liquidity, Figure 10 shows evidence that they do appear to vary with capital intensity. C-corps in capital intensive industries are much more likely than comparable S-corps to increase their pre-tax profits, costs, employment, and total payrolls in the years following TCJA, and the differences are both economically and statistically significant. The evidence also suggests that capital-intensive C-corps are somewhat more likely than comparable S-corps to increase investment, although the differences are not statistically significant. The fact that the investment responses for these firms are similar, even while the profit and cost elasticities are larger, suggests that the elasticity of output with respect to capital inputs is likely larger in capital intensive industries, as in [Acemoglu and Guerrieri \(2008\)](#).

In the Online Appendix we show the full event studies separately for capital-intensive and non-capital-intensive firms, and present additional heterogeneity tests. In general, we do not find clear evidence that policy effects vary with firm profitability, unionization rates, or market concentration.

Robustness

Having presented the main results, we turn now to assessing their robustness and to addressing potential threats to our identification strategy. The identifying assumption underlying our elasticity estimation strategy is that the outcomes of C- and S-corps would have trended similarly in the absence of TCJA's marginal income tax rate cuts. While this assumption is not directly testable, the event studies above show that C- and S-corps followed

FIGURE 10: FIRM HETEROGENEITY



Notes: The unit of analysis is a firm-year. The table shows the $C \times Post$ coefficients from equation 1.2. These coefficients estimate average differential changes in outcomes between C- and S-corps before and after TCJA, controlling for firm and industry-size-year fixed effects. Standard errors are clustered by firm.

broadly similar trends prior to TCJA, lending support to its plausibility. However, if the differential tax cuts affecting C- and S-corps following TCJA are correlated with simultaneous supply or demand shocks — for example, due to other provisions of the legislation or external events differentially affecting C- and S-corps — then our causal estimates would be biased. To address concerns about robustness and identification, we present specifications with alternate controls and samples, and explicitly consider how other provisions of TCJA, unrelated external events, anticipation effects, and tax-shifting behaviors may affect our analysis.

Alternate Specifications

Row 1 of Table 7 reports the key net-of-tax elasticities obtained from estimating equation 1.1 using our benchmark specification. We focus on the outcomes that will serve as key inputs in our analyses of welfare and incidence: pre- and after-tax profits; the earnings of executive and high- and low-paid workers; and net investment. In the remaining rows, we examine the sensitivity of these estimates to alternate modeling choices. Unless otherwise noted, all specifications include firm and industry-size-year fixed effects.

In row 2 we control for cohort-by-year fixed effects, where cohorts are defined using the firms' year of incorporation. This specification implies that the elasticities are identified from comparisons of C- and S-corps that are the same age.

In row 3 we control for state-by-year fixed effects, where a firm's state is defined using the address reported on Form 1120 or 1120s. In practice, the operations of large or exporting firms may span many states. Therefore, this specification controls for location-specific trends to the extent that firm performance is influenced by time-varying shocks associated with the firm's reported tax address.

Although we do not observe any systematic evidence that C- and S-corps were on different trends prior to TCJA, in row 4 we nevertheless probe the sensitivity of the estimates to adding pre-trend controls. Specifically, we control for lagged outcomes in 2013, 2014, 2015, and 2016, and interact each of the lagged outcomes with year indicators.

In row 5 we show the results using a log transformation of the outcome. These specifications implicitly exclude observations in which the outcome is zero or negative. For outcomes where the benchmark specification was already logged, the results are the same as in row 1.

In column 6, we winsorize the outcomes at the 5th and 95th percentiles. These specifications provide information about the extent to which the estimates are driven by changes in outcomes at the tails of the outcome distribution.

In general, the elasticity estimates across all specifications and outcomes are stable and within the 95% confidence interval of the benchmark specification in row 1.

TABLE 7: ROBUSTNESS TO ALTERNATE SPECIFICATIONS

Specification	(1)	(2)	(3)	(4)	(5)	(6)
	ε^B Pre-Tax π	ε^π After-Tax π	$\varepsilon^{w_{p50}}$ p50 w	$\varepsilon^{w_{p95}}$ p95 w	$\varepsilon^{w_{exec}}$ Exec w	ε^I I_t/K_{t-1}
1. Benchmark	0.379 (0.127)	0.521 (0.129)	-0.001 (0.052)	0.176 (0.053)	0.639 (0.169)	0.515 (0.082)
2. Age Controls	0.247 (0.119)	0.397 (0.122)	-0.010 (0.050)	0.150 (0.050)	0.468 (0.161)	0.519 (0.078)
3. Location Controls	0.361 (0.130)	0.502 (0.132)	-0.016 (0.052)	0.165 (0.053)	0.685 (0.168)	0.486 (0.081)
4. Pre-Trend Controls	0.434 (0.117)	0.471 (0.121)	0.091 (0.048)	0.262 (0.050)	1.103 (0.172)	0.535 (0.075)
5. Log Outcome	0.234 (0.136)	0.372 (0.137)	-0.001 (0.052)	0.176 (0.053)	0.639 (0.169)	0.790 (0.452)
6. Winsorize 5-95	0.257 (0.051)	0.279 (0.050)	-0.009 (0.041)	0.148 (0.043)	0.761 (0.150)	0.302 (0.066)

The table shows net-of-tax elasticities for key outcomes estimated from variations on equation 1.4. Unless otherwise indicated, all specifications include firm and industry-size-year fixed effects. The outcomes in columns 1 and 2 are scaled by baseline 2016 sales, the outcomes in columns 3-5 are logged, and the outcome in column 6 is scaled by lagged capital. Row 1 shows the benchmark specification. Row 2 controls for firm age by including cohort-by-year effects, where cohorts are defined as the firms' year of incorporation. Row 3 controls for state-by-year fixed effects, using the firm's reported address on its tax return. Row 4 controls for pre-trends, interacting the lagged pre-TJCA outcomes with year indicators. Row 5 shows a log transformation of the outcome rather than scaling by baseline sales; four columns 3-5, the results are the same as in row 1. Row 6 winsorizes the outcomes at the 5th and 95th percentiles. Standard errors clustered by firm.

Alternate Samples

Table 7 shows robustness results for the same key elasticities using alternate samples. Row 2 excludes firms with 2016 sales greater than \$1 billion or 2016 employment greater than 10,000. This sample restriction effectively excludes C-corps that are larger than the largest S-corps. Row 3 excludes "mismatched" industries, defined as those in which C-corps account for greater than 80% or less than 20% of the firms in the sample. Row 4 excludes manufacturing firms, which may have been more affected, for example, by the US-China trade war that occurred during our sample period (Fajgelbaum, Goldberg, Kennedy, and Khandelwal 2020). In all the samples, the magnitudes are broadly similar to the benchmark specification.

Other Provisions of TCJA

Beyond reducing marginal corporate income tax rates, TCJA also introduced several new policies affecting various business tax deductions, the taxation of foreign business income, and capital investment incentives. To assess the extent to which our estimates may be driven by these policy changes rather than by the rate cuts, below we briefly summarize the major provisions of TCJA affecting corporations, and then present several additional robustness checks. For more details on these reforms, see [Auerbach \(2018\)](#) and [Joint Committee on Taxation \(2018b\)](#).

- **Net Operating Loss (NOL) Deductions:** TCJA limited NOL deductions to 80% of a corporation's taxable income, eliminated NOL carrybacks, and allowed indefinite NOL carryforwards.
- **Domestic Production Activities Deduction (DPAD):** TCJA repealed DPAD, which provided a tax deduction to corporations that produce manufactured goods within the United States.
- **Alternative Minimum Tax (AMT):** TCJA repealed the corporate AMT, which imposed a minimum tax of 20% on corporations' relevant taxable income in excess of a \$40,000 exemption threshold, excluding the firm's AMT foreign tax credit.
- **Interest Deductions:** TCJA limited the interest payment deductions to 30% of adjusted taxable income.
- **Bonus Depreciation:** TCJA temporarily allowed corporations to immediately deduct 100% of the cost of newly purchased eligible capital investments (known as "full expensing"), an increase from 50% prior to TCJA, but scheduled to phase out beginning in 2023.
- **Taxation of Foreign Income:** TCJA introduced several changes to taxation of corporations' income earned abroad. The most significant changes include: (a) a one-time tax on previously accumulated foreign income and an elimination of tax on repatriated income; (b) a minimum tax on foreign income above a threshold return on tangible assets (known as Global Intangible Low-Taxed Income, or GILTI); (c) a minimum tax on deductible related-party transactions to U.S. subsidiaries — known as the Base Erosion and Anti-Abuse Tax, or BEAT); and (d) a lower tax rate on income earned from foreign sales, known as Foreign Derived Intangible Income, or FDII).

Ex-ante, it may seem unlikely that these other provisions would drive our results, for two reasons. First, because TCJA's non-rate policy changes broadly applied to both C- and S-Corps, our difference-in-difference design implicitly controls for them to the extent that C- and S-corps were similarly affected. Second, the legislative budget scoring report by the Congressional Joint Committee on Taxation ([2017](#)) projected that the rate cuts would do

more than any other business tax provision of TCJA to reduce tax revenues, making those rate cuts natural suspects.

Nevertheless, these considerations do not rule out that other provisions of TCJA may affect our estimates. For example, if C- and S-corps were differentially exposed to these policy changes — for example, perhaps because C-corps are more likely to earn foreign income than S-corps and thus more likely to be affected by the international provisions — then our net-of-tax elasticities may be biased. In the case of bonus depreciation, theory suggests that the effect of the tax rate may interact with the expensing rate; we discuss this possibility in greater detail in Section 1.4.

To assess the sensitivity of our estimates to alternate policy changes, in rows 5 to 9 of Table 8 we implement a series of additional robustness tests in which we exclude the firms most likely to be affected by each respective provision of TCJA. In row 5, we exclude industries where net operating losses are most common, defined as those where the share of firms reporting a loss in the pre-period is greater than the sample median. We similarly define and exclude industries where firms were most likely to claim the DPAD deduction (row 6) or exceed the interest limitation threshold (row 7). In row 8 we exclude industries where firms were most exposed to changes in bonus depreciation, defined as those where the ratio of bonus-eligible investment to sales is greater than the sample median in the pre-period, and in row 9 we exclude the multinational firms in our sample. Across the different samples, the results remain within the confidence interval of the benchmark estimates in row 1.

Anticipation Effects

If businesses expected the federal government to cut corporate taxes long before TCJA was formally enacted, they may have adjusted their behavior in anticipation of actual policy changes. In that case, a naive empirical strategy that compares outcomes of firms before and after TCJA risks underestimating the absolute magnitude of treatment effects, since a portion of the treatment effects would be captured in the pre-period data.

However, a careful consideration of the legislative history of TCJA suggests that anticipation effects are unlikely to bias our elasticity estimates. Pre-election polling and betting market spreads, as well as post-election stock market responses and media coverage, indicated that the November 2016 federal election outcome was difficult to predict and largely unexpected by the public.¹¹ Because members of the two major U.S. political parties generally have different preferences over business tax policy, the fact that the election was widely unexpected implies that firms and workers could not have significantly adjusted their behavior long in advance of TCJA. While it is possible that our empirical results may capture some anticipations effects during 2017 while policy negotiations were ongoing, in the Online Appendix we report elasticity estimates based on changes in outcomes between 2016 and 2019 — where the former is long before the legislative details of TCJA were promulgated —

¹¹For pre-election polling, see a composite of surveys compiled by Real Clear Politics [here](#). For betting spreads, see time series data from PredictIt [here](#). For examples of media coverage, see [here](#). For stock market responses, see [here](#).

TABLE 8: ROBUSTNESS TO ALTERNATE SAMPLES

Sample	(1) ε^B Pre-Tax π	(2) ε^π After-Tax π	(3) $\varepsilon^{w_{p50}}$ p50 w	(4) $\varepsilon^{w_{p95}}$ p95 w	(5) $\varepsilon^{w_{exec}}$ Exec w	(6) ε^I I_t/K_{t-1}
1. All Firms	0.379 (0.127)	0.521 (0.129)	-0.001 (0.052)	0.177 (0.053)	0.645 (0.169)	0.516 (0.082)
2. Exclude Large C-Corps	0.376 (0.142)	0.540 (0.144)	0.021 (0.056)	0.207 (0.058)	0.778 (0.178)	0.547 (0.089)
3. Exclude Mismatch Industries	0.277 (0.132)	0.435 (0.134)	0.002 (0.054)	0.149 (0.053)	0.529 (0.173)	0.523 (0.085)
4. Exclude Mfg Industries	0.344 (0.149)	0.545 (0.152)	0.093 (0.069)	0.245 (0.068)	0.814 (0.222)	0.618 (0.109)
5. Exclude NOL Industries	0.213 (0.138)	0.390 (0.143)	-0.023 (0.051)	0.097 (0.056)	0.353 (0.176)	0.319 (0.083)
6. Exclude DPAD Industries	0.244 (0.168)	0.495 (0.173)	0.123 (0.076)	0.170 (0.074)	0.640 (0.238)	0.608 (0.119)
7. Exclude High-Debt Industries	0.789 (0.243)	1.120 (0.252)	0.077 (0.090)	0.474 (0.106)	0.918 (0.309)	0.813 (0.166)
8. Exclude Bonus Industries	0.320 (0.135)	0.364 (0.134)	0.078 (0.071)	0.208 (0.074)	0.697 (0.241)	0.559 (0.119)
9. Exclude Multinationals	0.221 (0.116)	0.406 (0.118)	0.012 (0.056)	0.217 (0.059)	0.773 (0.187)	0.575 (0.093)
10. Exclude Small Firms	0.442 (0.163)	0.601 (0.165)	0.022 (0.062)	0.186 (0.060)	0.644 (0.203)	0.452 (0.093)
11. Exclude Single-Owner S-Corps	0.378 (0.122)	0.478 (0.124)	0.008 (0.050)	0.162 (0.052)	0.640 (0.162)	0.477 (0.079)

The table shows net-of-tax elasticities for key outcomes estimated from equation 1.4 using alternate samples. All specifications include firm and industry-size-year fixed effects. The outcomes in columns 1 and 2 are scaled by baseline 2016 sales, the outcomes in columns 3-5 are logged, and the outcome in column 6 is scaled by lagged capital. Row 1 shows the benchmark specification. Row 2 excludes firms with 2016 employment of greater than 10,000 or 2016 sales greater than \$1 billion. Row 3 excludes industries where the firm share of C-corps is less than 20% or greater than 80%. Row 4 excludes manufacturing industries. Row 5 excludes industries where net operating losses are most common in the pre-period. Row 6 excludes industries most likely to claim the Domestic Production Activities Deduction in the pre-period. Row 7 excludes industries where firms are highly leveraged. Row 8 excludes industries most likely to claim bonus depreciation. Row 9 excludes multinational firms. Row 10 excludes firms with fewer than 199 employees, in our smallest firm-size bin. Row 11 excludes S-corps with only a single owner.

and the results are similar. Moreover, because our difference-in-difference design compares relative changes in the trends of C- and S-corps, any anticipatory responses prior to 2017 affecting all firms are absorbed in our industry-size-year fixed effects.

Tax Shifting and Evasion

Tax-shifting behaviors are strategies employed by firms to minimize their tax burdens without significantly altering their broader economic behavior. Recent research emphasizes that taxable income elasticities must be interpreted with caution to the extent that firms engage in tax-shifting behaviors (e.g., Gorry, Hubbard, and Mathur 2021; Saez, Slemrod, and Giertz 2012). These behaviors may include intertemporal shifting (e.g., firms accelerate deductions or delay income in the years directly before and after a tax cut to minimize their tax burdens, for example as in Dowd, Giosa, and Willingham 2020) or shifting across tax bases (e.g., when owners of pass-through firms shift income between the corporate and individual sectors, for example as in Auerbach and Slemrod 1997 and Slemrod 1995).

Estimation strategies that do not account for shifting may yield misleading elasticities. In the case of intertemporal shifting, if revenue leakage in one year is offset by revenue gains in subsequent years (or vice versa), the choice of measurement years may materially impact the elasticity magnitudes. In the case of shifting across tax bases — in essence, a form of fiscal externality — elasticities that measure only a single tax base may exaggerate the decline in taxable income resulting from a tax increase (or vice versa for a tax cut).


The elasticities that we report in Tables 3, 6, and 5 compare the relative changes in outcomes of C- and S-corps before and after TCJA. To assess the sensitivity of these estimates to shifting behaviors, in the Online Appendix we show results where the elasticities are estimated from the differential change in C- and S-corps from 2016 to 2019. These specifications avoid capturing intertemporal tax-shifting behaviors in 2017 and 2018, the years where intertemporal shifting is likely to be most significant. The results are statistically similar to the benchmark results from equation 1.2.

In principle, shifting across tax bases is possible across several different margins. For example, firms may change their entity-type, switching from C to S status or vice versa. S-corps owners may choose to reclassify their wages as profits to maximize the value of the QBI deduction. The incentive to reclassify wages in this way is strongest for S-corps with just one owner: because wage costs are deductible, in firms with more than one owner, this form of shifting comes at the expense of the other owners, who are unlikely to approve of it. Multinational firms may have incentives to book their profits domestically rather than abroad in response to the new international tax provisions. Finally, to the extent that tax cuts reduce incentives for tax evasion (that is, illegal misreporting of income), this is most likely to be true among small firms, which are least likely to be audited.

To evaluate the extent of entity-type switching, Figure 11 shows the profit-weighted share of entity-type switchers in each year of our sample both before and after TCJA. The combined switching rate of both C and S corps prior to TCJA was less than 0.3%, and this share increased only trivially after TCJA to less than 0.4%. Thus, although we document a

clear increase in entity-type switching after TCJA, this form of tax shifting is negligible and does not bias our elasticities.

FIGURE 11: CORPORATE ENTITY-TYPE SWITCHING, 2013-2019



Notes: Figure shows the profit-weighted share of firms that switch their legal entity type from C-to-S or from S-to-C over our sample period. Entity switching is very rare, and increased only modestly after TCJA.

To further assess the sensitivity of our estimates to firms with potentially high shifting or evasion propensities, in Table 8 we run our benchmark specification from equation 1.4 on samples that: exclude multinational firms (row 6); exclude small firms (row 10); and exclude S-corps with only one owner (row 11). The point estimates are statistically indistinguishable from the benchmark specification and do not suggest that income shifting across tax bases is a significant concern in our setting. These findings are also consistent with contemporaneous research. With respect to S-corps, [Goodman, Lim, Sacerdote, and Whitten \(2021\)](#) study a large sample of de-identified tax returns of pass-through businesses and find that S-corps mostly did not engage in wage-to-profit shifting in response to the QBI deduction. With respect to multinationals, [Garcia-Bernardo, Jansky, and Zucman \(2022\)](#) find only small changes in the share of US multinational profits booked abroad following TCJA.

Mechanisms and the Cost of Capital

In Section 1.4 we provided evidence that our empirical results are unlikely to be driven by liquidity effects. In this section we argue that our findings are consistent with dynamic theories in which firms are responsive to both current and future changes in the cost of

capital. To illuminate these mechanisms, we first begin with a static model of the corporate income tax. Suppose firms optimize after-tax profits π :

$$\pi = F(K, L)(1 - \tau) - wL(1 - \tau) - rK(1 - \theta\tau) \quad (1.5)$$

where τ is the corporate income tax rate, and $\theta \in [0, 1]$ is an expensing rate parameter capturing the share of production costs that are tax deductible. These costs may include fully deductible capital purchases (such as durable equipment eligible for bonus depreciation), partially deductible capital purchases (such as structures), and non-deductible costs (such as managerial effort, to the extent that it is not reflected in cost-deductible compensation). We assume that $F(K, L)$ varies across firms due to heterogeneous productivities, such that some firms are able to produce greater output than others with a fixed set of inputs. The first-order condition for profit maximization with respect to capital yields:

$$\underbrace{\frac{\partial F}{\partial K}}_{\text{MRPK}} = \underbrace{\frac{(1 - \theta\tau)}{1 - \tau} r}_{\text{cost of capital}} \equiv \phi \quad (1.6)$$

where the left side of the equation is the marginal revenue product of capital, and the right side expression is defined as the user cost of capital, ϕ . The expression shows that, in general, either decreasing the tax rate τ or increasing the expensing parameter θ lowers the cost of capital ϕ . However, if all production costs are tax deductible (corresponding to the “full expensing” case where $\theta = 1$), then the tax rate does not affect the cost of capital, and thus does not affect capital demand.

In our setting, TCJA permanently cut the corporate tax rate and, due to the bonus incentives discussed in Section 1.4, also temporarily increased the expensing rate (for some assets, up to 100%). In response to the lower cost of capital, standard models predict that firms will demand more capital (that is, invest more), demand more labor (to the extent that capital and labor are complements), and increase their scale (in part because a higher capital stock may make the firm more productive).

How should we make sense of the simultaneous changes in the tax rate and depreciation allowances in our setting? Which policy instrument is driving the results? In a two-period model where the changes in expensing are assumed to be permanent, such that $\theta_{pre} < 1$ and $\theta_{post} = 1$, equation 1.6 implies that the relative change in the cost of capital wedge for C-versus S-corps is given by the following expression:

$$\frac{\phi_{post}^C}{\phi_{pre}^C} - \frac{\phi_{post}^S}{\phi_{pre}^S} = \frac{(1 - \theta_{pre}^C)}{1 - \tau_{pre}^C} - \frac{(1 - \theta_{pre}^S)}{1 - \tau_{pre}^S} \quad (1.7)$$

In this expression, the change in the cost of capital wedge is only a function of the initial levels of the tax rates, τ_{pre}^C and τ_{pre}^S , and is not affected by the changes in tax rates or by the levels of the rates in the post-period. Because initial marginal income tax rates in our sample

are higher for S-corps than for C-corps — that is, $\tau_{pre}^S > \tau_{pre}^C$ — equation 1.7 implies that, at least for capital assets eligible for bonus depreciation, S-corps faced a larger reduction in the cost of capital than C-corps following TCJA. As a result, the model would predict that S-corps should increase investment, employment, and sales relative to C-corps — the exact opposite of our empirical findings.

That the qualitative predictions from this model differ from our empirical results suggests that the simple model does not fully capture economic reality. In practice, some capital expenditures, such as structures, do not qualify for bonus depreciation. Moreover, existing studies show that a large fraction of firms do not claim bonus depreciation even when eligible to do so (Kitchen and Knittel 2016; see also Joint Committee on Taxation 2021), suggesting the corporate rate change may remain relevant for firm behavior.

A complementary plausible reconciliation of evidence and theory in our setting appeals to richer models that incorporate dynamics and adjustment costs, as in Auerbach and Hassett (1992). In these models, firms respond not only to current tax rates but also to the future path of policy. In the presence of high adjustment costs, firms will be highly sensitive to future policy, since investment will depend not only on the user cost of capital today but on the user cost of capital in the future. Because TCJA’s provisions primarily favoring C-corps were made permanent (the large corporate rate cut), while the provisions favoring S-corps were made temporary (the QBI deduction), and the provisions benefitting both were also temporary (the increase in expensing), our empirical results are consistent with the view that firms considered the future path of tax policy in response to TCJA.

In the Online Appendix, we adopt the Auerbach and Hassett (1992) model and use cost-of-capital estimates from Foertsch (2018) to estimate structural elasticities with respect to (net of) effective marginal tax rates. The model can flexibly incorporate salient features of firm behavior and of the U.S. business tax provisions before and after TCJA, including: forward-looking expectations; adjustment costs in investment; different tax rates on the income of shareholders of C- versus S-corps, including individual-level taxes on dividends, capital gains, interest income, and distributions of nonqualified annuities; the expansion and phase-out of bonus depreciation, as well as incomplete take-up of bonus; and the sunset of the QBI deduction. Our net-of-tax elasticities using this method are very similar to our benchmark results reported in row 1 of Table 7.

1.5 Revenue Impacts, Excess Burden, and Incidence

In this section we leverage the reduced form elasticities estimated in Section 1.4 to evaluate the short-run revenue impacts, excess burden, and incidence of TCJA’s corporate tax cuts. We adopt a transparent framework in the style of Feldstein (1999), such that elasticities of key outcomes with respect to the net-of-tax rate are sufficient to estimate the aggregate welfare consequences of changes in tax policy. As discussed in Saez, Slemrod, and Giertz (2012), the empirical validity of this approach rests on two key assumptions: (a) negligible tax shifting, and (b) negligible income effects. In Section 1.4 we presented several empirical tests

suggesting that shifting behaviors are unlikely to drive our estimate of the corporate taxable income elasticity. Moreover, our heterogeneity tests in Section 1.4 showed that high- and low-liquidity firms responded similarly to corporate tax cuts, suggesting income effects are indeed negligible in our setting.

For clarity and to facilitate comparison with existing literature, when interpreting the results we focus only on the core provisions of TCJA relating to firms' marginal income tax rates, and abstract from issues relating to changes in personal income taxes, deficit financing, public goods provision, consumer prices, and dynamics.

Revenue Impacts

Starting from the firm problem in equation 1.5, let the corporate tax base B be defined as firm revenues less deductible costs:

$$B = F(K, L) - wL - \theta rK \quad (1.8)$$

Total corporate tax revenues T are the product of the tax base B and the corporate tax rate τ :

$$T = \tau B \quad (1.9)$$

In the absence of behavioral responses, mechanical changes in tax revenue from a change in the corporate net-of-tax rate $1 - \tau$ are given by holding the tax base constant:

$$dM = -Bd(1 - \tau) \quad (1.10)$$

The additional change in tax revenue generated by behavioral responses is given by:

$$dB = \frac{\varepsilon^B B \tau}{1 - \tau} d\tau \quad (1.11)$$

where $\varepsilon^B = \frac{\partial B/B}{\partial(1-\tau)/(1-\tau)}$ is the elasticity of taxable income with respect to the net-of-tax rate, equivalently called the elasticity of pre-tax profits or the elasticity of the corporate income tax base. Intuitively, the extent to which revenue losses from tax cuts are offset by an expanding tax base is directly proportional to the taxable income elasticity ε^B . The total change in tax revenue is given by:

$$dT \equiv dM + dB = dM \left[1 - \frac{\tau \varepsilon^B}{1 - \tau} \right] \quad (1.12)$$

where dT is the sum of the mechanical and behavioral responses.¹²

¹²For simplicity, here we abstract from the effects of corporate tax changes on personal income tax revenues.

Welfare and Excess Burden

Define aggregate welfare W as the sum of after-tax private income Y and public tax revenues T :

$$W = Y + T \quad (1.13)$$

where taxes are defined as in equation 1.9, and Y is the sum of private income received by firm owners (π^K) and different groups of workers (π^{L^j}), indexed by j :

$$Y = \pi^K + \sum_j \pi^{L^j} \quad (1.14)$$

We use this definition of welfare, which corresponds approximately to GDP or total output, because it is transparent, can be objectively measured in the data, and can be easily compared with existing estimates in the literature. In general, however, changes in output and welfare will not be equivalent if, for example, there is curvature in individuals' utility functions. Rather than welfare, an alternate interpretation of W is that it quantifies the market value of the output distortion from the corporate tax.

Guided by our empirical results, we classify three groups of workers: low-paid workers, high-paid workers, and executives. Low-paid workers are defined as those in the bottom 90% of the within-firm wage distribution, and high-paid workers as those in top 10%. Workers and executives optimize consumption $C^L = w^j L^j$, where w^j is the wage for workers of type j and L^j is labor supply. The indirect utility function for workers is given by $U^j(w^j) = w^j L^j$, and the change in utility from a change in wages is given by:

$$dU^j(w^j) = L^j dw^j = w^j L^j \varepsilon^{L^j} d(1 - \tau) \quad (1.15)$$

where ε^{L^j} is the elasticity of earnings for workers of type j with respect to the net-of-tax rate. Because firm owners are assumed to optimize their demands for factor inputs, by application of the envelope theorem the change in firm owners' profits is given by:

$$d\pi^K = -d\tau B - \sum_j (1 - \tau)(dw^j L^j) \quad (1.16)$$

The first term $-d\tau B$ implies that a reduction in the tax rate increases profits. The second term $\sum_j (1 - \tau)(dw^j L^j)$ captures the effects of factor price adjustments: while higher wages improve welfare for workers, they reduce welfare for firm owners, whose production costs increase. In practice, these offsetting adjustments are implicitly embedded in the elasticity of after-tax profits to the net-of-tax rate, which we have estimated in the empirical analysis. We can thus compute the change in welfare for firm owners as:

$$d\pi^K = \pi^K \varepsilon^\pi d(1 - \tau) \quad (1.17)$$

where ε^π is the elasticity of after-tax profits with respect to the net-of-tax rate, and π represents after-tax profits in the baseline year. The total change in welfare is given by:

$$\begin{aligned}
 dW &= dY + dT & (1.18) \\
 &= d\pi^K + \sum_j dU^j + dT
 \end{aligned}$$

We combine the elasticities estimated in Section 1.4 with moments from the tax data to compute the total change in welfare as expressed in equation 1.18. Finally, the marginal excess burden from the corporate tax cut is given by:

$$\frac{dW}{dT} = \frac{dT + d\pi^K + \sum_j dU^j}{dT} \quad (1.19)$$

which expresses the marginal welfare loss from raising an additional dollar of corporate tax revenue — or, in our setting, the marginal welfare gain from an additional dollar of foregone tax revenue.

Incidence

To assess distributional impacts of TCJA’s corporate tax changes, we adapt the framework developed in [Suárez Serrato and Zidar \(2016\)](#) and [Fuest, Peichl, and Sieglöcher \(2018\)](#) to estimate the share of productive factors in the total corporate tax burden. Our analysis differs from these studies in two respects. First, the detailed microdata used in this study allows us to observe returns to firm owners, and thus allow us to empirically estimate how these returns are affected by changes in the corporate tax rate. Second, we are able to estimate the effects of corporate tax changes on the full distribution of workers’ earnings. Using these two sets of estimates, we are able to evaluate the incidence of corporate taxes using weaker assumptions than are required when these outcomes are not empirically observable.

We are also able to extend our analysis to assess corporate tax incidence not only on factors of production — that is, on firm owners and workers, as is standard in the literature — but also to approximate incidence over the income distribution, taking account of the empirical fact that many low-income workers own capital and most capital owners also work. Doing so allows us to speak directly to research and policy debates about the progressivity of the corporate income tax. Because we observe workers’ locations, we are also able to evaluate the geographic incidence of corporate income tax cuts.¹³

In evaluating incidence, we make the standard assumptions of a representative consumer and equal redistribution of tax revenues to all citizens. The former assumption rules

¹³In the incidence analysis of [Fuest, Peichl, and Sieglöcher \(2018\)](#) studying workers’ wages, the effects of corporate tax changes on returns to firm owners are unobservable, and changes in rental rates are assumed to be negligible. This assumption is likely to be appropriate in their analysis of tax changes in German municipalities, which they characterize as small open economies. In the incidence analysis of [Suárez Serrato and Zidar \(2016\)](#), returns to firm owners are unobservable but inferred via structural estimation. These studies evaluate impacts of corporate tax changes on median and mean worker wages, respectively, but do not directly assess impacts over the earnings distribution.

out distributional impacts through changes in consumer prices, which are unobservable in our data. The latter assumption, while strong, allows us to avoid making even stronger alternative assumptions about the future path of fiscal policy. TCJA's corporate tax cuts were deficit financed, and the future trajectory of federal tax policy is always uncertain in a democracy.

Accounting for changes in factor prices, the change in welfare for firm owners is given by:

$$d\pi^K = -d\tau B - \sum_j (1 - \tau)(dw^j L^j) \quad (1.20)$$

We can thus compute the change in welfare for firm owners as:

$$I^F = \frac{d\pi^K}{d\pi^K + \sum_j dU^j} \quad (1.21)$$

Similarly, the share of workers in the total tax burden, I^{L_j} , is given by replacing the numerator in equation 1.21 with dU^j .

We expand on the traditional analysis of factor incidence in two extensions. First, we evaluate incidence with respect to the population distribution of income. Estimating distributional incidence allows us to account for the empirical fact that many workers are also firm owners (because they may hold equity portfolios, as emphasized in [Auerbach 2006](#)) and that many firm owners also earn labor income (as documented in [Smith, Yagan, Zidar, and Zwick 2019](#)). We assume that everyone works, and ascribe firm and capital ownership to workers using data on the distribution of equity and wealth ownership from the Distributional Financial Accounts (DFA) produced by the Federal Reserve Board ([2018](#)). We assume that executives are in the top 1% of the distribution, that high-paid workers are in the 90-99th percentiles, and that low-paid workers comprise the bottom 90%. Letting ω^j represent the capital ownership share of workers of type j , we have:

$$I^{L_j} = \frac{dU^j + \omega^j d\pi^K}{d\pi^K + \sum_j dU^j} \quad (1.22)$$

which measures incidence across the income distribution for all workers, inclusive of both labor and capital income.

Second, we combine the distributional estimates from equation 1.22 with the observed locations of workers, inferred from zip codes reported on IRS Form W-2, to estimate the geographic incidence of corporate income tax cuts across Census regions and commuting zones. Letting $\rho^{j(r)}$ represent the share of workers of type j living in region r , and N^r represent the region's population, we compute:

$$\frac{dY^r}{N^r} = \frac{\sum_j \rho^{j(r)}(dU^j + \omega^j d\pi^K)}{N^r} \quad (1.23)$$

which provides an estimate of the effect of federal corporate tax changes on per capita income in region r . To the extent that firm ownership and employment are unequally spatially distributed across the country, the gains from corporate tax cuts are likely to be unequally geographically distributed as well.

Quantification Moments and Parameters

Table 9 summarizes the key inputs that we use to quantify the welfare and incidence implications of TCJA's corporate income tax cuts. Panel A includes information on the empirically observed average tax rates and changes faced by C- and S-corps in our sample, and Panel B shows their aggregate 2016 levels of tax liabilities, after-tax profits, and payrolls for different groups of workers. Panel C reports the distribution of capital ownership as observed in the Distributional Financial Accounts data, and Panel D reviews the key net-of-tax elasticities we have estimated in the empirical analysis.

Revenue, Income, and Welfare Impacts

Panel A of Table 10 shows our estimates of the impacts of corporate income tax cuts on government tax revenues. To generate these estimates, we use the empirically estimated elasticities and key moments from our sample of tax returns. We show estimates of the mechanical effects on tax revenue (that is, holding the tax base constant), as well as estimates of the total effects (taking account of behavioral responses). We present the estimates as dollar values, percentage changes, and as a share of 2016 GDP.

Panel A of Table 9 shows that, on average, TCJA reduced the marginal tax rate by 10 percentage points for C-corps and by 4 percentage points for S-corps. In the absence of behavioral responses, Panel A of Table 10 shows that this would lead to a \$101 billion (34%) reduction in corporate tax revenues, corresponding to approximately 0.47% of 2016 GDP. However, because firms respond to the tax cut by expanding their operations and increasing pre-tax profits, the total reduction in tax revenue is modestly attenuated, instead \$88 billion (30%), or 0.41% of GDP.

Panel B of Table 10 shows our estimates of the change in private income from TCJA's corporate tax changes. We estimate that private income increases by \$97 billion, or 0.46% of GDP. Approximately \$54 billion of these gains accrue to firm owners and \$43 billion accrue to workers.

Panel C shows our estimates of our welfare and the marginal excess burden of the corporate tax. In our stylized framework, welfare increases linearly in private income and public revenues. Private income gains of \$97 billion combined with revenue losses of \$88 billion imply a net increase in total welfare of \$9 billion, or 0.04% of 2016 GDP. Our estimate is of a similar order of magnitude to Barro and Furman 2018, who structurally simulate the effects of TCJA on GDP using a fully parameterized Ramsey model.

Panel C provides our estimate of the marginal excess burden of the corporate income tax, $\frac{dW}{dT}$. We find that a marginal dollar of foregone revenue from corporate income tax cuts generates an additional \$0.10 in output. Viewed through the lens of the model, the results thus imply substantial efficiency gains from corporate tax cuts. However, as we show below, these aggregate gains mask significant distributional effects.

TABLE 9: QUANTIFICATION MOMENTS AND PARAMETERS

	All Corps	C-Corps	S-Corps
Panel A: MTRs (Sales-Weighted)			
Mean 2016 τ	0.25	0.24	0.31
Mean $\Delta\tau$	-0.09	-0.10	-0.04
Mean $\Delta\tau/\tau_{t-1}$	-0.33	-0.47	0.34
Mean $\ln(1 - \tau)$	0.13	0.14	0.06
Panel B: Firm Aggregates (bil)			
Tax	299	255	44
Taxable Income	1,163	914	250
After-Tax Profit	882	658	224
Executive Payroll	151	105	47
Top 10 Payroll	767	673	94
Bottom 90 Payroll	1,403	1,211	192
Panel C: Distribution of Capital			
Top 1%	0.27		
91-99%	0.34		
Bottom 90%	0.39		
Panel D: Net-of-Tax Elasticities			
Pre-Tax Profit	0.38		
After-Tax Profit	0.52		
Executive Pay	0.65		
Top 10 Earnings	0.32		
Bottom 90 Earnings	0.00		

Table shows the inputs that we use to quantify the revenue, income, and welfare impacts of TCJA's corporate tax cuts. Data in Panels A and B are directly observed in our sample of tax records, and data in Panel C are from the 2018 Federal Reserve Board Distributional Financial Accounts. The parameters in Panel E are estimated in the empirical analysis.

TABLE 10: REVENUE AND WELFARE ESTIMATES

	bil	% GDP	%
Panel A: Tax Revenues			
Mechanical, dM	-101.2	-0.47	
Total, dT	-88.4	-0.41	
Panel B: After-Tax Private Income			
Total Income, dY	97.3	0.45	
Capital Income, $d\pi^K$	54.5	0.25	
Labor Income, $d\pi^L$	42.9	0.20	
Panel C: Welfare and Excess Burden			
Welfare, dW	8.9	0.04	
Marginal Excess Burden, dW/dT			10.1

This table shows estimated revenue, income, and welfare impacts from TCJA's changes in marginal corporate income tax rates. Outcomes are scaled in billions of dollars in column 1. The denominators for percentage changes in tax revenues are: 2016 federal corporate tax revenues in Panel A; 2016 private income of corporate firm owners and employees in Panel B; and 2016 GDP in Panel C. Outcomes in column 3 are scaled as a percent of GDP. The marginal excess burden is defined as the ratio of the change in welfare to the change in tax revenues. See Section 1.5 for details.

Incidence

Panel A of Table 11 shows our estimates of changes in private income for firm owners, executives, and high- and low-paid workers. Combining our reduced form elasticities from Section 1.4 with the moments from the tax data, we find that approximately 56% of the gains from TCJA's corporate tax cuts flow to firm owners; 12% flow to executives; 32% flow to high-paid workers; and 0% of the gains flow to low-paid workers.

Panel B reports our estimates of incidence over the income distribution. When we allocate the gains of firm owners to workers using data from the Distributional Financial Accounts, we find that approximately 27% of the gains from corporate tax cuts accrue to the top 1% of the earnings distribution; 51% accrue to the 90-99th percentiles; and 22% accrue to the bottom 90%. These results highlight the importance of considering the joint impacts of changes on both capital and labor income when assessing the distributional effects of corporate tax

changes.

Panel C of Table 11 reports our estimates of geographic incidence across Census regions, produced from equation 1.23.¹⁴ Because firm owners and highly-paid workers are relatively highly concentrated in the Northeast and West Coast regions of the United States, we find that gains from the corporate tax cuts disproportionately accrue to those regions. For example, our estimate of the per capita income gain for residents of the Northeast (\$33) is approximately 77% larger than for residents of the South (\$19).

Panel A of Figure 12 maps the variation in our estimates of geographic incidence across commuting zones. Beyond the regional patterns highlighted in Table 11, the map highlights substantial within-region variation, with larger and higher-income commuting zones generally seeing larger gains from the corporate tax cuts. The patterns are most clearly illustrated in Panel B, which plots the estimated change in income against the 2016 average earnings of corporate-sector employees, and where the bubbles are proportional in size to each commuting zone's population. Relative to the median commuting zone gain of approximately \$120 per capita, we estimate that gains are approximately 4 times larger in Houston or New York City; 6-7 times larger in Seattle; and roughly 10 times larger in the San Francisco Bay Area. The results imply that corporate income tax cuts not only increase income inequality across workers, but also contribute to growing inequality across regions and commuting zones (Gaubert, Kline, Vergara, and Yagan 2021).

The Efficiency-Equity Tradeoff of the Corporate Income Tax

Our estimate of corporate taxable income elasticity — a key parameter in the literature for measuring the distortion of a tax — implies substantial efficiency gains from cutting corporate taxes (or, equivalently, implies substantial losses from increasing the corporate tax rate). However, we also find that corporate tax cuts disproportionately benefit those with high incomes, with 78% of the gains flowing to just 10% of the population.

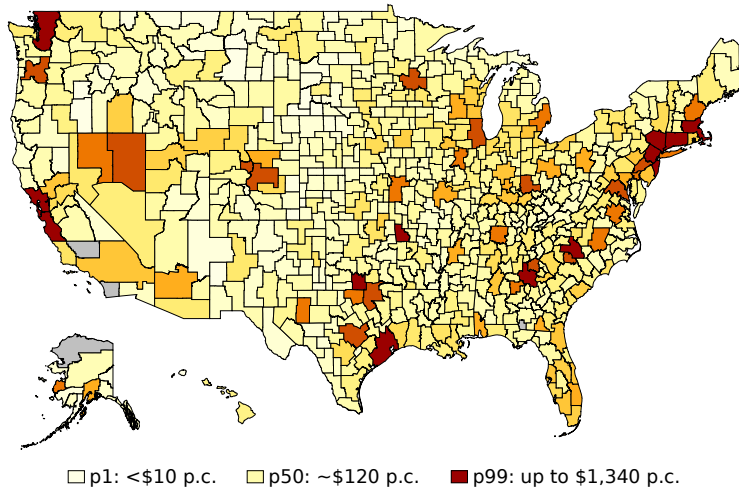
Given that the federal government must raise some level of revenue to finance its operations, how should we interpret our results on the corporate income tax in the context of the broader national tax and transfer system? In Figure 13, we benchmark our findings against estimates from the literature on personal income and payroll taxes — the two other largest sources of federal tax revenues in the United States. The X-axis shows $-\varepsilon^B$, where ε^B is the elasticity of the tax base to an increase in the net-of-tax rate for each policy instrument. A larger magnitude of ε^B implies a larger distortion from the tax. The Y-axis shows the share of tax burden borne by the top 10% of the income distribution, where higher shares imply that the tax is more progressive.

The estimates of ε^B for the personal income tax (0.25) and payroll tax (0.21) are from Saez, Slemrod, and Giertz (2012) and Saez, Schoefer, and Seim (2019), respectively. The former is based on a comprehensive literature review of the voluminous empirical evidence on


¹⁴We classify states into regions using the definitions from the U.S. Census, provided [here](#), with the minor modification of classifying Delaware, DC, and Maryland as belonging to the Northeast rather than to the South.

FIGURE 12: GEOGRAPHIC INCIDENCE

Panel A: Change in Per Capita Income across Commuting Zones



Panel B: Change in Per Capita Income vs. Initial Worker Earnings



The unit of analysis is a commuting zone. Panel A illustrates geographic variation in our estimates of changes in per capita private income due to the corporate tax cuts, generated from equation 1.23. Income gains are proportional to color intensity in the map, with darker colors representing larger gains. Panel B plots the estimated change in income for each commuting zone against the 2016 average earnings of corporate-sector workers. The size of the bubbles is proportional to each commuting zone’s 2016 population.

TABLE 11: INCIDENCE ESTIMATES


	\$	% Change	% Incidence
Panel A: Factors (\$ bil)			
Firm Owners	54.5	6.2	56
Executives	11.3	7.5	12
High-Paid Workers	31.6	4.1	32
Low-Paid Workers	0.0	0.0	0
Panel B: Distributional (\$ bil)			
Top 1%	26.0	6.7	27
91-99th%	50.1	4.7	51
Bottom 90%	21.2	1.2	22
Panel C: Geographic (\$ per capita)			
Northeast	423	0.7	33
Midwest	263	0.5	21
South	239	0.5	19
West	346	0.6	27

Table shows the estimated incidence of corporate tax cuts on firm owners, executives, and high- and low-paid workers. To compute distributional incidence, we allocate gains of firm owners to workers using data on capital ownership from the Federal Reserve Distributional Financial Accounts. The denominators for percent change shown in column 2 are: 2016 private income of corporate firm owners and employees in Panel A; private labor and corporate capital income in Panel B; and real mean personal income by Census region according to the 2016 Census American Community Survey in Panel C.

personal income taxes, while the latter is, less satisfyingly, based on evidence of employment effects from a payroll tax reform in Sweden. However, it is, to our knowledge, the best available estimate for payroll taxes in the literature. We compute estimates for the share of personal and payroll tax burdens borne by the top 10% and top 1% of the income distribution using data from the Congressional Budget Office.

Viewed in the context of the literature, the results in Figure 13 suggest that the corporate income tax is approximately 1.5 times less efficient than the personal income tax. In Panel A, where our measure of equity is the share of the tax burden borne by the top 10%, the corporate income tax is similarly progressive to the personal income tax. In the Online Appendix, we use a more extreme measure of equity, the share of the tax burden borne by the top 1%, and find that the corporate income tax is both less efficient and less progressive than the personal income tax. In either case, the corporate tax appears 3-4 times more progressive than the payroll tax, although it is twice as inefficient.

FIGURE 13: THE EFFICIENCY-EQUITY TRADEOFF IN CONTEXT



The figure contextualizes our results on the corporate income tax, compared here with the personal income and payroll taxes, the two other largest sources of federal tax revenue in the United States. The elasticity of taxable income, shown on the X-axis, is a key parameter for measuring tax distortions. The share of the tax burden borne by the top 10% of the income distribution, shown on the Y-axis, is a measure of progressivity.

1.6 Conclusion

This paper analyzes the short-run impacts of historically large federal corporate income tax cuts on large U.S. firms and their workers. Exploiting tax policy variation that allows us to compare trends in outcomes of similarly sized firms operating in the same industry, we find that tax cuts cause firms to increase their sales, profits, and investment. These responses are predominantly concentrated in capital-intensive industries. Labor earnings increase for workers in the top 10% of the within-firm earnings distribution, and rise particularly sharply for executives, but do not change for workers in the bottom 90%. We do not find evidence that firms' responses are driven by liquidity effects, and interpret the results as consistent with dynamic models in which firms are responsive to both current and future changes in the cost of capital.

We empirically estimate key elasticities of firm- and worker-level outcomes, and combine these elasticities with a stylized model to estimate the revenue impacts, welfare gains, and incidence of TCJA's corporate tax changes. We find that private incomes increase by \$97 billion and tax revenues decline by \$88 billion, implying a net aggregate output gain of \$9

billion, equivalent to approximately 0.04% of GDP. In the the model, reducing corporate tax revenues by \$1 generates an additional \$0.10 in output, implying substantial efficiency gains from corporate tax cuts.

We also find that the gains from corporate tax cuts disproportionately flow to those with high incomes. We estimate that approximately 56% of the gains accrue to firm owners, 12% accrue to executives, 32% accrue to high-paid workers, and 0% flow to low-paid workers. When we adjust these calculations to allow for the empirical fact that many workers hold equity portfolios, we estimate that 78% of the gains flow to the top 10% of the earnings distribution, and 22% flow to the bottom 90%.

In a benchmarking exercise, we find that the efficiency gains from corporate tax cuts are 1.5 to 2 times as large as personal income or payroll tax cuts, even while the distributional effects are similarly regressive. Holding all else equal, the results imply that, on the margin, adjusting the composition of federal revenues toward a larger share of personal income taxes and a lower share of corporate income taxes may yield significant efficiency gains without sacrificing progressivity.

We conclude with important caveats. Our results do not capture a range of potentially important channels through which corporate tax cuts may affect welfare. For example, in the long-run, higher investment may increase productivity and broadly increase workers' wages. While we do not find clear effects of tax cuts on productivity in our data (see the Online Appendix), and estimate zero effect on low-income workers' earnings, it is possible that such gains may materialize over a longer time horizon. On the other hand, reductions in tax revenues may lead to a deterioration in the provision of public services (such as education, health, or infrastructure spending), or reduce redistributive transfers, with potentially adverse implications for equity and efficiency. We believe these are important topics for future research.

Chapter 2

The TCJA and Domestic Corporate Tax Rates

In Chapter 2, Christine Dobridge, Paul Landefeld, Jake Mortenson and I document descriptive trends in the tax positions for U.S. C-corporations following passage of the TCJA. We document several novel facts, including that effective tax rates (ETRs) *increased* on average for privately held, domestic firms and for firms in the bottom 90% of the firm size distribution after the TCJA. By contrast, public, multinational, and large firms saw substantial ETR cuts. We show that changes to the corporate tax rate and treatment of net operating losses have the strongest correlation with ETR changes, and discuss avenues for future research.

This chapter is coauthored with Christine Dobridge, Paul Landefeld, and Jake Mortenson. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee. The views and opinions expressed here are the authors' own. They are not necessarily those of the Board of Governors of the Federal Reserve System, its members, or its staff.

2.1 Introduction

In this paper we aim to answer several fundamental descriptive questions about firms' tax positions before and after the law commonly known as the Tax Cuts and Jobs Act (TCJA): How did the effective tax rates (ETRs) paid by C corporations change after TCJA? How do changes in ETRs vary with firm characteristics, such as firm size and whether a firm is public, private, domestic, or multinational? And which of TCJA's many provisions appear most significant in explaining changes in firms' ETRs?

Despite the importance of these questions for understanding and analyzing policy, existing studies have not been able to comprehensively answer them, primarily due to data limitations. While previous empirical studies have focused on publicly traded firms, which account for the majority of employment and taxes paid, research suggests that smaller private domestic firms and startups are engines of innovation and growth, especially in economically important sectors such as technology and health (Decker et al., 2014). Moreover, large private multinational C corporations outnumber public multinationals five-to-one, and comprise approximately 20 percent of total C corporation employment. Understanding trends in both public and private firms is thus likely to be important for analyzing TCJA's effects. Yet due to limitations in commonly used databases of public firms, basic information such as, for example, the share of firms subject to the corporate alternative minimum tax prior to its post-TCJA repeal, is not well-documented in existing research.

We fill these gaps in existing research by using a representative sample of de-identified corporate tax returns from 2010-2019. The rolling panel from the Internal Revenue Service's Statistics on Income (SOI) files includes both public and private firms, and both domestic and multinational firms, allowing us to provide a holistic picture of the changing landscape of firms' tax positions before and after TCJA.

The analysis proceeds in three parts. First, we document times series trends in average domestic cash ETRs for different groups of firms, highlighting especially the divergence in trends for public/private, domestic/multinational, and small/large firms. Second, we show how the full distribution of firms' ETRs shifted before and after TCJA. Third, we document basic evidence on the share of firms likely to have been affected by TCJA's various provisions (which we call "exposure measures"), and show how these exposure measures correlate with firms' ETR changes after TCJA. We note that the descriptive analysis is not intended to provide evidence on the causal effects of TCJA's individual provisions; rather, our aim is to highlight previously undocumented patterns, to provide clear evidence on the relative scope and economic significance of TCJA's various provisions, and to discuss implications for policy analysis and future research.

From the analysis, we document three descriptive facts. First, we show that ETRs declined sharply for publicly traded firms and for privately owned multinationals. These firms pay the lion's share of federal corporate income taxes and account for approximately 69% of C corporation employment. This result is broadly consistent with Henry and Sansing (2018); Dyreng et al. (2020); Wagner et al. (2020); however, our analysis paints a more nuanced picture relative to existing research because we present effects on domestic cash

ETRs for multinational firms, as opposed to global ETRs.

The second fact — which, from our perspective, appears to be less commonly understood in public and academic discourse — is that, in contrast to publicly traded and multinational firms, average ETRs increased for private domestic firms. While private domestic firms do not pay a large share of federal corporate income taxes, they do account for the vast majority of corporate firms and 31% of corporate employment. Our analysis shows, therefore, that a large majority (approximately 81%) of U.S. C corporations and workers at those firms experienced a corporate income tax increase following TCJA.

Third, when assessing the relative importance of the TCJA’s various corporate provisions, we find the exposure to the marginal corporate income tax rate changes (i.e., increase at the bottom, decrease at the top) and the net operating loss (NOL) restriction are most strongly correlated with changes in firms’ ETRs after TCJA. By contrast, the alternative minimum tax, interest limitation, and multinational provisions were moderately correlated with changes in firms’ ETRs while bonus depreciation-related changes and repeal of the domestic production activities deduction are only very weakly associated with changes in firms’ ETRs.

These findings represent contributions to collective understanding of this historically large corporate tax reform along two primary dimensions. First, we use a set of data that is representative of public and private firms in the United States. Virtually all prior related work has used data on publicly traded companies or non-representative samples of private companies, and private companies comprise the majority of U.S. corporations (Lisowsky and Minnis, 2020) and make up an economically important share of U.S. employment and innovative activity. Second, our analysis of the components of the corporate tax reform provides information on the scope and relative importance of different provisions. This is a valuable addition to the literature given the breadth of reforms that comprise the TCJA, beyond the corporate rate changes.

This paper also contributes to the broad set of research in economics, finance, and accounting that studies trends in and determinants of corporate tax rates and tax-related financial positions both domestically and globally (Slemrod, 2004; Clausing, 2007; Dyreng et al., 2017; Wier and Zucman, 2022). In related studies examining corporate tax trends around the TCJA, Dowd et al. (2020) use IRS data and find that large firms had considerable tax-related behavioral responses to the legislation. Dyreng et al. (2020) and Wagner et al. (2020) use data from Securities and Exchange Commission filings to study trends in global ETRs around the TCJA and examine the association between ETR changes and firm characteristics. Henry and Sansing (2018) study scaled cash tax differences for public firms in addition to ETRs find that the TCJA had no effect in tax-favored status for public firms overall but reduced the share of tax-favored profitable firms.¹ Several concurrent papers also examine certain provisions of the TCJA in isolation, including the interest limitation

¹The scaled cash tax difference is defined as year-to-year changes in cash taxes paid divided by the market value of assets, which can be calculated for firms with losses. The authors define tax-favored status as whether a firm’s cash ETR continued to be less than the statutory rate or whether scaled cash tax differences continued to be negative after the TCJA.

(Carrizosa et al., 2020) and the international provisions (Clausing, 2020; Garcia-Bernardo et al., 2022).

Our work is also related to the growing literature studying effects of the TCJA on firm behavior.² The majority of the current research study effects of the TCJA on public firm behavior. However, research on private firms suggests considerable differences in tax-related incentives and behavior of private firms compared to public firms (Mills and Newberry, 2001; Graham et al., 2014; Hoopes et al., 2020; Dobridge et al., 2021). Given the different post-TCJA trends that we document in tax positions for private domestic firms, our work raises the question of how other outcomes may have differed for private firms compared to public firms after the TCJA as well. Overall, our results underscore that TCJA’s policy effects for private domestic firms are likely to be very different from public and multinationals.

2.2 Overview of TCJA Corporate Tax Provisions

In this section, we provide a short summary of the numerous domestic and international corporate tax changes in the TCJA.³ Beginning with the domestic business-related changes, the Joint Committee on Taxation estimated these provisions would reduce federal revenue by \$650 billion over the 10-year budget window following the TCJA’s fiscal years 2018 to 2027. The most substantial change, in terms of revenue consequences, was the change in the corporate tax rate schedule from an approximately graduated statutory tax schedule, with a top marginal tax rate of 35 percent on income in excess of \$10 million, to a flat tax schedule with a rate of 21 percent. The TCJA also repealed the corporate alternative minimum tax (AMT) while allowing firms to offset tax liability or receive a credit for unused AMT credits.

Three aspects of the TCJA directly affected deductions for capital expenditures, interest expenses, and business activity associated with domestic production. First, under the TCJA, the depreciation schedule for most capital investment was changed to full expensing through 2022. Second, the TCJA imposed a new limit on the deductibility of interest expenses to 30 percent of adjusted taxable income.⁴ Third, the TCJA also repealed the domestic production

²See, for example, work on corporate actions and statements following TCJA enactment (Hanlon et al., 2019), contributions to defined benefit pension plans (Gaertner et al., 2020), changes in executive compensation (De Simone et al., 2021), reclassification of business costs (?), organizational structures (Henry et al., 2018), earnings management (Kubick et al., 2021), and employment, worker earnings, and capital investment (Dobridge et al., 2022). Gale et al. (2019) examined the TCJA’s effects on aggregate domestic activity, including nonresidential fixed investment, non-farm employment, and real earnings.

³A detailed description of all the statutory changes enacted under the TCJA and an explanation of the prior law’s provisions are presented in Joint Committee on Taxation (2018c), estimated budget effects are included in Joint Committee on Taxation (2018a) and additional discussion of the legislative motive for the changes is contained in Dowd et al. (2020).

⁴Through 2021, adjustable taxable income was defined as taxable income excluding business interest income and expense, depreciation, amortization, depletion, and NOLs. After 2021, adjustable taxable income was defined as taxable income excluding business interest income and expense and NOLs. Other TCJA provisions related to deductions included the limitation of deductions for fringe benefits and expanding the definition of executive compensation for the purposes of Section 162(m).

activity deduction (DPAD), which had provided taxpayers deduction of nine percent for income from qualifying activities related to producing goods and services domestically.⁵

There were also key provisions related to the tax treatment of losses. Prior to the TCJA, firms were permitted to apply NOLs against taxable income for 2 years prior (“carryback”) or apply NOLs against taxable income as many as 20 years into the future (“carryforward”). The TCJA eliminated NOL carrybacks and changed the limitation on carryforwards to 80 percent of available NOLs for 20 years.

The TCJA made extensive changes to tax provisions related to foreign income and operations as well. As a sense of scale, the JCT estimated that the total international tax reform changes would raise about \$325 billion from 2018 to 2027.

First, the legislation changed the treatment of foreign income from a worldwide system of taxation, whereby foreign earnings were generally taxed only when they were repatriated to the United States, to a territorial system that eliminates the tax on repatriated or unrepatriated foreign earnings (*i.e.*, gives a 100 percent deduction for dividends received from foreign subsidiaries). The legislation also included a one-time transition tax on the previously untaxed earnings of U.S. multinationals of eight percent on non-cash assets and 15.5 percent on cash assets held overseas.

Several new international tax provisions were implemented related to firm income-shifting incentives, particularly: the base erosion and anti-abuse tax (BEAT), the global intangible low-taxed income (GILTI) provision, and the foreign derived intangible income (FDII) provision. The switch to a territorial international tax system increased incentives, to some extent, for firms to move operations and income overseas because any tax savings from doing so would be permanent instead of deferred. The BEAT, GILTI, and FDII were intended to counteract these effects and decrease incentives to move business activity overseas. The BEAT imposed a minimum 10 percent U.S. tax on deductible transactions made between related parties (*i.e.*, payments made from a U.S. parent to a controlled foreign corporation). GILTI was intended to reduce incentives for firms to relocate operations to lower-tax countries by levying a minimum 10.5 percent tax on a firm’s foreign earnings greater than 10 percent of total foreign tangible assets. FDII reduced firm taxes on U.S. earnings derived from foreign sales and was intended to incentivize firms to locate intangible capital domestically.

2.3 Description of the Corporate Tax Data

The corporate tax return data we use for this analysis are sourced from Form 1120 corporate tax filings and related schedules, as provided by the Statistics of Income (SOI) Division of the Internal Revenue Service (IRS). In particular, we use data from a stratified random sample of corporate tax returns that is created, cleaned, and edited each year by the SOI (U.S. Internal Revenue Service, 2013). We refer to these data as the “SOI sample.” For the

⁵See [Dobridge et al. \(2021\)](#) and [Ohrn \(2018a\)](#) for a detailed discussion of the DPAD.

main analysis, we use data on tax payments from the front page of Form 1120 and data on book income from Schedules M-1 and M-3.⁶

Our primary variable of interest is the domestic U.S. corporate effective cash tax rate (ETR), which we define as total taxes paid over domestic pre-tax book income. We calculate this for all firms with positive book income, mimicking the approaches of Hoopes et al. (2020) and Dobridge et al. (2021).⁷ Detailed definitions for all variables used in this study, including IRS form and line numbers, are including in Appendix A.

We split the SOI sample – which contains about 27,000 firms – along several firm characteristics to study the heterogeneous effects of the TCJA. The characteristics include public ownership, private ownership, multinational operations, domestic-only operations, and size bins of domestic sales. We define firms as public if they ever report having publicly traded voting common stock on Schedule M-3, line 3a, during the sample period from 2010 to 2019, and categorize all other firms as private firms.⁸ We define a firm as multinational if it ever reports having a controlled foreign corporation (*i.e.*, a foreign subsidiary) by filing Form 5471 during the sample period. Finally, to study firms of different sizes, we separate firms across categories of the firm domestic sales distribution, using total domestic gross receipts reported on line 1c of Form 1120.

We construct two different analysis samples of firms in this paper. For our analysis of the time trends in U.S. domestic cash ETRs, we require a firm to have non-negative and non-zero book income in a given year to be included in the sample for that year, such that we are able to calculate a cash ETR for each firm (designated “ETR sample”). For our analysis of changes in the domestic cash ETR from 2016 to 2019, we require a firm to have a non-missing cash ETR in both of those years (designated “ETR change sample”). Table ?? presents summary statistics for the ETR sample of firm characteristics in 2016, prior to TCJA enactment, for all firms and separately for multinational, domestic, public and private firms. All samples are weighted using SOI weights to produce firm-weighted population averages and standard deviations.

⁶Note that total taxes paid will include some U.S. taxes on foreign source income which is subject to tax. Prior to the TCJA this included dividends from CFCs, foreign branch income, and subpart F income. After TCJA, this could also include tax owed on GILTI and tax on the one-time transition tax for unrepatriated foreign earnings.

⁷Specifically, when a firm reports attaching a Schedule M-3 to the Form 1120 (Box A4), we calculate the ETR as taxes paid (Form 1120, line 31) divided by pre-tax financial statement income, where pre-tax financial statement income is the sum of net income (Schedule M-3, Part I, line 11), U.S. current income tax expense (Schedule M-3, Part III, line 1), and U.S. deferred income tax expense (Schedule M-3, Part III, line 2). However, only firms with assets above \$10 million are required to file the Schedule M-3. When a firm does not report attaching a Schedule M-3, we define pre-tax financial statement income as the sum of net income per book (Schedule M-1, line 1) and federal income tax per books (Schedule M-1, line 2).

⁸This method of identifying publicly traded firms mis-classifies as private any publicly traded firm that was not required to file a schedule M-3.

2.4 Results

Trends in Domestic Cash ETRs

We begin our analysis by presenting trends in average ETRs from 2010 to 2020 in Figure 1. Figure 1(a) separately displays ETRs for different types of firms: public multinational, public domestic, private multinational, and private domestic. Figure 1(b) presents ETRs at different points in the size distribution by total total sales: firms above the 95th percentile, firms in the 90th to 95th percentile, and firms below the 90th percentile.⁹


In Figure 1(a), we observe a similar trend in average firm domestic cash ETRs for public C corporations as documented by Dyreng et al. (2020) for the average firm global cash ETR of public U.S. firms. After fluctuating near 18 percent between 2012 and 2017, cash ETRs for public multinational and domestic firms declined sharply in 2018 after passage of the TCJA, to about 11 percent. In 2019, ETRs declined further for public multinationals (to 10 percent) and increased somewhat for public domestic C corporations (to 13 percent). Private multinational firms experienced a similar change as public multinational firms, though their average ETRs were one or two percentage points higher than public firms in the five years leading up to TCJA passage, and they experienced a somewhat smaller drop in ETRs after, to 13.5 percent in 2018 and 12 percent in 2019.

Changes in ETRs for private domestic firms were much different than public firms. After a gradual increase in ETRs over most of the pre-TCJA sample period, ETRs increased discontinuously in 2018 and ticked up again in 2019. By 2019, the average ETR for domestic private firms was *higher* than for public firms or for private multinational corporations—about 14 percent for private domestics compared to about 10 percent for public multinationals, for example—a striking reversal of the pre-TCJA trend when average ETRs were considerably lower for private domestic firms than for public firms and private multinationals.


In 2020, we observe trends that were likely materially affected by the Covid-19 pandemic; it is difficult, therefore, to interpret these moves as directly related to the TCJA in light of the pandemic's effect on economic activity and corporate profitability. Amongst public firms, the ETR for multinationals ticked back up in 2020 and the ETR for domestic firms rose further as well. In contrast, for both private multinationals and private domestic firms, ETRs declined notably in 2020 compared to 2019.

⁹In Appendix Table 3, we also present summary statistics on firm characteristics and ETRs in 2016 (pre-TCJA), on ETRs in 2019 (post-TCJA) and on ETR changes between those two years. As in Figure 1, the statistics are weighted to reflect the C corporation population.

FIGURE 1: EFFECTIVE TAX RATES FOR C-CORPORATIONS: 2010 TO 2019




(B) BY THE DISTRIBUTION OF SALES




This figure presents trends in average domestic cash effective tax rates from 2010 to 2020 for U.S. C corporations. Panel A presents trends for firms overall, for public multinationals (MNEs), public domestic firms, private MNEs and private domestic firms. Panel B presents trends for various bins of the firm size distribution, by total sales. All panels were created by the authors using IRS administrative tax data and variables are defined in Appendix A.

FIGURE 2: DISTRIBUTION OF ETRs BEFORE AND AFTER TCJA

(A) DISTRIBUTION OF ETRs, 2016 AND 2019



(B) ETRs BY FIRM TYPE, 2016 AND 2019



This figure presents the distribution of ETRs prior to TCJA passage, in 2016, and following passage, in 2019. Panel A shows the distribution for all firms, and Panel B shows the distributions by firm type. All panels were created by the authors using IRS administrative tax data and variables are defined in Appendix A.

Figure 1(b) decomposes trends by firm size, and illustrates that the largest tax cuts were concentrated amongst the biggest firms (i.e., those that most benefited from the rate structure flattening). For firms in the top five percent of the sales distribution, average ETRs declined from 23% in 2016 and to 16% in 2019. Firms in the 90 to 95th percentile in terms of sales had relatively small average ETR declines, from 17% in 2016 to 16% in 2019, while small firms experienced an average ETR increase. We observe that in 2020, first across the size distribution experienced a decline in ETR from 2019.

To further investigate the dispersion in ETR changes across firms, we examine the distribution of ETRs before and after the TCJA. Figure 2(a) presents a histogram of the distribution of ETRs in 2016 and 2019. We choose 2016 as our primary pre-TCJA year of comparison due to evidence that firms undertook behavioral responses in 2017 in anticipation of the TCJA enactment, as studied by Dowd et al. (2020). We find that annual ETR averages mask considerable variation in ETRs across firms. In both 2016 and 2019, the distribution of ETRs is bimodal, with the most frequent observations around zero ETR in both years, reflecting the large fraction of firms with positive book income but zero U.S. federal cash taxes paid. In 2016, the distribution shows a second mass around 15 percent: the marginal tax rate for firms with taxable income less than \$50,000.¹⁰ In 2019 the distribution shows a second mass around 21 percent: the new, flat corporate tax rate for all C corporations.

Figure 2(b) shows the distribution of ETRs in 2016 and 2019 by firm type. Consistent with the results in figure 1, the panel shows that the distribution of ETRs shifted to the left for public firms and for private multinationals, but shifted to the right for private domestic firms. Given that private domestic firms make up the majority of U.S. C corporations, the figure reveals a surprising fact about the distribution of ETRs before and after the TCJA: the vast majority of U.S. C corporations with positive book income experienced an ETR increase between 2016 and 2019 (approximately 81 percent of all C corporations), despite the reduction in marginal tax rates for most of the 2016 tax brackets.

The picture is considerably different for public firm and private multinationals. The majority of these firms experienced ETR declines from 2016 to 2019, similar to findings for public multinationals documented by Wagner et al. (2020) for the distribution of changes in firms' global cash ETRs after the TCJA.

The Relative Importance of Different TCJA Corporate Provisions

In this section, we move beyond studying the ETR – an “overall” measure of tax liability – and examine the relative importance of specific TCJA provisions. Given the lack of transparency and granular information about many firms' financial and tax positions, it can be challenging for researchers and policymakers to understand how exposed different types of firms were to the various categories of legislative changes in the bill. Therefore, we evaluate firms' tax-related characteristics in 2016, one year prior to the passage and two years prior to the enactment of the TCJA, as a proxy for exposure to different provisions prior to the TCJA.

¹⁰Average and marginal tax rates are the same in the first bracket of a graduated rate structure.

Panel A of Table 1 presents a list of major TCJA provisions, with proxies for firm exposure to each provision listed beneath. The fraction of firms for which each measure applied is presented for public multinationals in column 1, public domestic firms in column 2, private multinationals in column 3, and private domestic firms in column 4. To provide a sense of economic scale, Panel B reports the share of total C corporate firms, sales, tax, and employment corresponding to each firm type.

The centerpiece of the TCJA corporate provisions was the change from an approximately graduated corporate tax schedule, with a maximum rate of 35 percent, to a flat corporate income tax rate of 21 percent. To better understand the relative importance of the corporate tax rate change, we examine the share of firms in each tax bracket in 2016. From that baseline, a substantial fraction of private domestic U.S. firms would have had higher ETRs under the TCJA's new tax rate schedule (*ceteris paribus*): 88 percent of private domestic firms paid no tax and 8 percent were in the first tax bracket (with a 15% tax rate) in 2016. In contrast, 58 percent of private multinationals and only a quarter of public firms paid no tax. Approximately zero percent of public firms were in the first tax bracket. For firms with positive taxable income, the largest fraction of public domestic and private multinationals were in the 34 percent bracket (with between \$335 thousand and \$10 million in taxable income) and for public firms, 58 percent were in the 35 percent tax bracket, with greater than \$18 million in taxable income.

The corporate AMT repeal and DPAD repeal were more likely to affect multinationals and public firms than domestic and private firms. In 2016, the AMT affected few private domestic firms (less than 1 percent), but around 11 and 6 percent of public and private multinationals, respectively, and about 16 percent of public domestic firms. The DPAD was most utilized by multinationals, claimed by 48 and 15 percent of public and private multinationals, respectively, in 2016, compared to just 11 percent of public domestic firms and 1 percent of private domestic firms.

The other two provisions with wide applicability were the NOL limitations and the capital expensing provisions. 59 percent of public multinationals and 47 percent of public domestic firms were able to claim an NOL carryforward, as well as 28 percent and 9 percent of multinational and domestic private firms, respectively. A large majority of public firms and private multinational firms also had depreciable capital expenditures in 2016, which would have likely been eligible for immediate, full bonus depreciation (and lower current-year taxable income) under the TCJA. This was more common for public multinationals and domestic firms (95 percent and 87 percent, respectively) than private multinationals (73 percent). Of private domestic firms, 38 percent had depreciable capital expenditures in 2016.¹¹

The interest limitation – which was based on the amount of interest expense relative to EBITDA – would have affected a small share of U.S. firms. While a large fraction of firms had positive interest expenses, a small fraction had positive taxable income and interest

¹¹Note that prior to the TCJA, many smaller firms were already able to expense a portion of their capital expenditures through Section 179.

TABLE 1: EXPOSURE TO TCJA PROVISIONS

Panel A: 2016 Exposure Measures

TCJA Provision	Public Multinational	Public Domestic	Private Multinational	Private Domestic
Paid No Tax	0.263	0.286	0.586	0.880
1st MTR Bracket	0.000	0.001	0.028	0.081
AMT	0.111	0.159	0.062	0.003
NOL	0.594	0.469	0.278	0.087
Bonus	0.953	0.873	0.729	0.381
DPAD	0.479	0.110	0.153	0.012
Interest Limited	0.073	0.018	0.032	0.007
Foreign Earnings	0.699	0.000	0.388	0.000
Repatriated Earnings	0.318	0.029	0.061	0.000
Foreign Earnings > Tangible Assets	0.576	0.000	0.334	0.000

Panel B: Summary Statistics by Firm Type

Share of C Corp	Public Multinational	Public Domestic	Private Multinational	Private Domestic
Firms	0.002	0.001	0.012	0.985
Sales	0.517	0.039	0.222	0.221
Tax	0.665	0.051	0.147	0.136
Employment	0.454	0.057	0.176	0.314

Notes: Panel A presents descriptive statistics of the share of firms exposed to various TCJA corporate tax provisions, in 2016 prior to enactment. Statistics are provided for four types of firms in columns (1) to (4): Public multinationals, public domestic firms, private multinationals, and private domestic firms. Panel B provides context about the share of C corporate firms, sales, tax, and employment corresponding to each firm type. Statistics are firm-weighted using SOI weights to be representative of the U.S. population of C corporations. The figure was created by the authors using IRS administrative tax data and variables are defined in Appendix A.

expense above the limitation in 2016: Approximately 7 percent and 2 percent of public and private multinationals, for example.

Lastly, we examine the relative importance of the corporate international provisions.¹² A cornerstone of the multinational policy changes was the switch from a worldwide tax system, whereby foreign earnings were taxed when repatriated to the United States, to a territorial tax system, whereby foreign earnings would be taxed at foreign, local rates. Approximately 70 percent of public multinationals and 39 percent of private multinationals had positive foreign earnings overall in 2016. About 32 percent of public multinationals and 6 percent of private ones repatriated earnings to the United States in that year, which manifests as

¹²The international provisions generally may lead to a change in the amount of U.S. corporate tax owed by a corporation without any change in their domestic book income.

a U.S. parent receiving a dividend from a controlled foreign corporation (CFC). Finally, we also evaluate the global intangible low-taxed income (GILTI) provision. GILTI exacts at least a 10.5 percent tax on a firm's foreign earnings that are greater than 10 percent of qualified business asset investments (QBAI). In 2016, we observe that 58 percent of public multinationals and 33 percent of private multinationals had foreign earnings in excess of tangible assets, a proxy for QBAI.

Correlates of ETR Changes

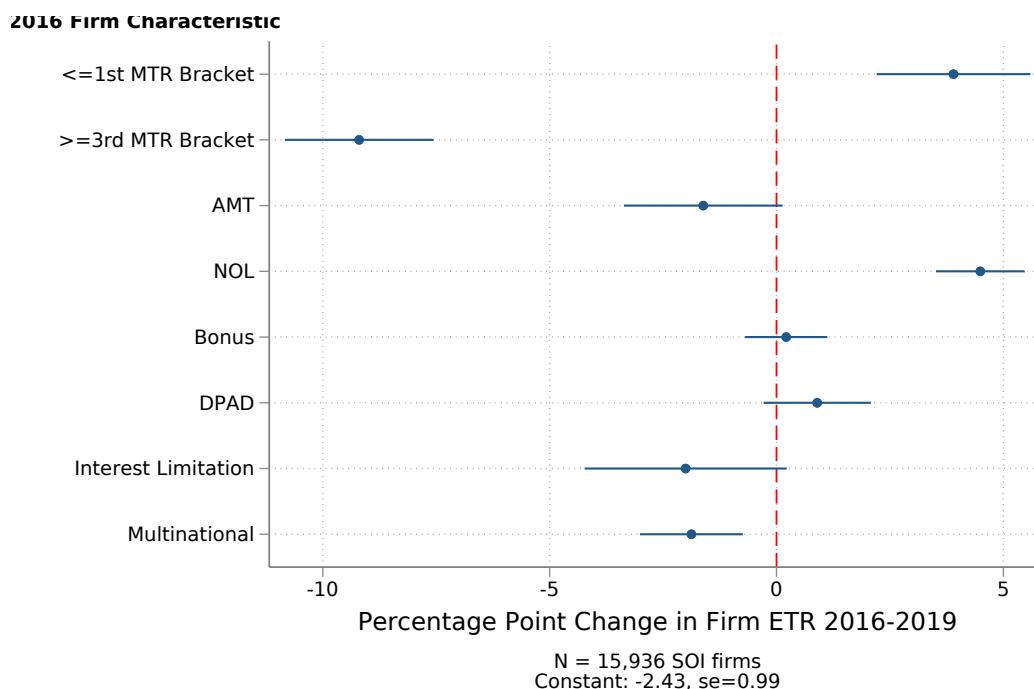
How was exposure to the various TCJA provisions correlated with changes in firms' effective tax rates? This section roughly approximates the relative salience of different provisions to observed ETR changes documented in Section 2.4. We regress firm-level ETR changes from 2016 to 2019 on indicators for firm exposure to certain aspects of the TCJA, using the following specification:

$$\Delta ETR_f = \alpha + \beta_1 \cdot \leq 1stMTRBracket_f + \beta_2 \cdot \geq 3rdMTRBracket_f + \beta_3 AMT_f + \beta_4 NOL_f + \beta_5 Bonus_f + \beta_6 DPAD_f + \beta_7 InterestLimitation_f + \beta_8 MNC_f + \epsilon_f \quad (2.1)$$

In the regression, the outcome ΔETR_f is the 2016 to 2019 change in the cash ETR, measured at the firm level (f), and all of the regressors are defined using firm characteristics in 2016. $\leq 1stMTRBracket$ is an indicator variable equal to 1 if a firm was in the first (15 percent) tax bracket or paid zero tax; $\geq 3rdMTRBracket$ is an indicator variable equal to 1 if a firm was in the third tax tax bracket or higher; AMT is an indicator variable equal to 1 if a firm paid the corporate AMT, NOL is an indicator variable equal to 1 if a firm used an NOL carryforward; $Bonus$ is an indicator variable equal to 1 if a firm claimed accelerated depreciation from Form 4562; $DPAD$ is an indicator equal to 1 if a firm recorded a DPAD deduction; $InterestLimitation$ is a dummy variable equal to 1 if interest payments exceed 30 percent of the sum of taxable income, net depreciation, and interest paid; and MNC is a dummy variable equal to 1 if a firm has a Controlled Foreign Corporation. Standard errors are clustered by firm.

Figure 3 displays the estimated coefficients and 95 percent confidence intervals for the policy-specific indicator variables in equation 2.1. These coefficient estimates suggest exposure to the statutory marginal tax rate changes and the NOL changes have the largest conditional correlations with ETR changes.

FIGURE 3: COEFFICIENT ESTIMATES: TCJA PROVISIONS AND ETR CHANGES



The figure presents coefficient estimates and 95 percent confidence intervals of a regression of the change in ETRs from 2016 to 2019 on indicators measuring a firm's exposure to various TCJA provisions. The figure was created by the authors using IRS administrative tax data and variables are defined in Appendix A.

Being in the the first MTR bracket or paying zero taxes in 2016 is associated with having a four percentage point larger increase in ETR from 2016 to 2019, compared to firms in the second MTR bracket (the 25% percent bracket—the omitted group). In contrast, being in the third MTR bracket or above in 2016 is associated with about an 8 percentage point lower ETR change as compared to firms in the second bracket. Using an NOL carryforward in 2016 is associated with about a 4.5 percentage point higher ETR change. Pre-TCJA AMT exposure, interest limitation exposure, and multinational firm status all have the expected negative association with the ETR change, but much smaller magnitudes than the rate changes and NOL changes. Exposure to bonus depreciation had almost no conditional correlation with ETR changes, while DPAD exposure was weakly and positively correlated.

2.5 Conclusions

In this paper, we use a representative sample of U.S. C-corporation tax returns to summarize the direct effects of TCJA corporate tax provisions on the tax positions of corporations. While the decline in public firm ETRs after the TCJA has been well-documented by other

researchers and the TCJA's corporate tax cuts are commonly discussed in media and other public venues (*e.g.*, see [Rubin and Francis \(2020\)](#) and [Rubin and Francis \(2021\)](#)), we present several novel facts about changes in ETRs and the relative importance of several aspects of this complex reform.

We document that a majority of U.S. firms saw an increase in their ETR after passage of the TCJA and that the average domestic ETR increased for U.S. firms overall as well. Both trends were driven by an increase in ETRs for private domestic firms. We show that similar to public MNCs, private MNCs saw a decline in ETRs after the TCJA, but a smaller decline than that experienced by public multinational and domestic firms. Finally, we find that the change in statutory tax rates and the limitation on net operating losses had the strongest correlation with changes in firm ETRs from 2016 to 2019. Exposure to the corporate AMT repeal, the interest limitation, and the multinational had smaller, negative estimated correlations, while exposure to the bonus depreciation and DPAD changes had little effect.

Our work contributes to the body of existing research by presenting a more complete picture of the effects of the TCJA on tax positions of U.S. corporations, the majority of which are privately held, domestically based, and small in terms of total assets and employment. Given the opaqueness of tax positions of privately-held firms that do not file public financial statements, as well as limited information even about tax positions included in public firm financial statements, it is challenging to evaluate changes in the U.S. tax code. This study provides a resource for researchers, policymakers, and the general public to better understand the direct effects of this substantial reform, including its component parts.

Chapter 3

Neighborhood-Level Investment from the U.S. Opportunity Zone Program

In Chapter 3, Harrison Wheeler and I use de-identified federal tax records from tax years 2019 and 2020 to document the first available evidence on the short-run response of financial capital to the Opportunity Zone (OZ) program, a federal place-based policy that provides tax incentives for capital investments in more than 8,000 low-income neighborhoods across the United States. We observe \$41.5 billion of aggregate cumulative OZ investments by tax year 2020. Using a sub-sample of electronically filed returns covering 78% of total observed investment, we document three emerging patterns in the data. First, OZ capital is highly spatially concentrated. Second, among OZ-designated neighborhoods, investors report greater equity and property investments in neighborhoods with relatively higher incomes, home values, educational attainment, and pre-existing income and population growth. Third, OZ investors have extremely high incomes relative to the US population, implying that the direct distributional incidence of the tax subsidy benefits households in the 99th percentile of the national income distribution.

This chapter is coauthored with Harrison Wheeler.

3.1 Introduction

Socioeconomic disparities across regions and neighborhoods are pervasive in the United States (Gaubert, Kline, Vergara, and Yagan 2021; Reardon and Bischoff 2011), and recent research documents that these disparities are likely to have causal effects on individuals' productivity (Moretti 2012), health (Chandra and Skinner 2003), intergenerational economic mobility (Chetty, Hendren, and Katz 2016; Chetty, Hendren, Kline, and Saez 2014), and propensity for innovation (Bell, Chetty, Jaravel, Petkova, and Van Reenen 2019). However, researchers and the public disagree about which policies, if any, are effective means to improving these outcomes (Glaeser and Gottlieb 2009; Busso, Gregory, and Kline 2013; Kline and Moretti 2014; Neumark and Simpson 2015; Gaubert, Kline, and Yagan 2021).

In this paper, we use de-identified federal business tax records to study the short-run response of financial capital to the U.S. Opportunity Zone (OZ) program, a federal place-based policy enacted in 2017 as part of the Tax Cuts and Jobs Act. As a result of this recent legislation, equity and property investments in more than 8,000 designated census tracts across the United States are eligible for highly favorable tax treatment of income accrued from capital gains.

The scale of the Opportunity Zone (OZ) program is unique in the modern landscape of place-based federal policies, both in terms of its expansive geographic scope and significant federal cost. OZs are located in urban, suburban, and rural areas across all 50 states, covering approximately 12% of all U.S. census tracts. The breadth of the program offers a natural setting to consider how place-based policies impact heterogeneous neighborhoods. The Congressional Joint Committee on Taxation has estimated that the OZ program will cost the government \$1.6 billion annually in foregone tax revenue, more than any other existing federal place-based policy.¹

A nascent literature on Opportunity Zones studies short-run impacts of the program on real estate prices (Chen, Glaeser, and Wessel 2019), job postings (Atkins, Hernandez-Lagos, Jara-Figueroa, and Seamans 2020), and employment (Freedman, Khanna, and Neumark 2021), finding null or modest effects. An exception is Arefeva, Davis, Ghent, and Park (2020), who estimate substantial increases in employment from establishment-level data. However, a key missing link in the early evidence is data on the response of financial investors to the tax subsidy. Existing studies estimate intent-to-treat (ITT) effects based on tract-level binary indicators for OZ status, but without information on program take-up are unable to estimate average treatment effects (ATE). While both of these parameters are of clear and natural interest to policymakers and researchers, a richer and more complete understanding of the evidence requires data on how investors have responded to the capital tax subsidy.

In particular, if investor behavior is only weakly responsive to the OZ tax subsidies, then small or null ITT effects are perhaps unsurprising, and policymakers may wish to consider if or how alternative policy mechanisms might attract investment to low-income neighborhoods. On the other hand, if investors are highly responsive to the subsidy and yet over time we

¹See The Joint Committee on Taxation (2020) estimates of federal tax expenditures from 2020-2024.

do not observe desirable downstream effects on labor market outcomes, then policymakers may wish to shift budget priorities away from capital tax subsidies and consider alternative policy levers that may be more effective.

This paper fills a gap in the existing research by documenting the first available evidence on tract-level financial investment associated with the OZ program. Our data is based on de-identified electronically-filed federal business tax records from tax years 2019 and 2020, the first two years in which OZ investors were required to report detailed information on the location and recipients of their investments to the IRS. We emphasize that these data are preliminary, and do not yet incorporate data from an estimated \$9.0 billion (approximately 22%) of cumulative OZ investments filed via paper tax returns. Throughout the paper we explicitly discuss limitations of these early data, and we will continue to update this working paper as more up-to-date information becomes available.

We highlight three main findings from the early evidence.

First, OZ investment is highly spatially concentrated. The vast majority of designated Opportunity Zone tracts in our sample, 63%, receive zero OZ capital. However, among tracts where investing firms report positive investment, the average value is substantial, at approximately \$3,313 per resident. The distribution is strongly skewed even among these tracts with positive investment, such that the median value is \$386, approximately one-ninth of the average.

Second, we correlate reported OZ investment with demographic and firm characteristics, and show that OZ capital gravitates toward eligible neighborhoods with relatively higher educational attainment, incomes, home values, population density, and concentrations of professional and amenity services. These patterns are strongest for neighborhoods with pre-existing upward trends in population, income, and home values, and declining shares of elderly and non-white residents. On the firm side, we show that reported OZ investment is overwhelmingly concentrated in equity investments in businesses that specialize in real estate, construction, and finance.

Third, OZ investors have extremely high incomes relative to the US population. We identify a large sample of OZ investors and estimate their median and average 2019 household income to be greater than \$741,000 and \$4.9 million, respectively.² These estimates imply that the direct distributional incidence of the tax subsidy is likely to benefit households in the 99th percentile of the US income distribution.

In the final section, we geocode the universe of individual and business tax records to construct novel measures of tract-level household and family income, employment, commuting, firm growth, and real investment. We demonstrate that these estimates closely match corresponding measures from publicly available data and describe advantages of our new measures relative to existing data. As more comprehensive data on OZ investment become available, we plan to use these data to evaluate the causal effect of the OZ tax subsidies on local labor market and real investment outcomes.

²Throughout the paper, all centile statistics are computed as centile averages to protect taxpayer privacy. For example, medians are computed as the average of all taxpayers in the 49th to 51st percentiles.

The rest of the paper proceeds as follows. Section 3.2 highlights the OZ program's goals and objectives as described by its authors in Congress, describes the process by which neighborhoods were nominated and selected, and provides details on the program's capital tax subsidies. Section 3.3 presents the first available descriptive evidence on the spatial distribution of OZ investment across the United States, based on electronic business tax filings in tax years 2019 and 2020. Section 3.4 presents new tract-level estimates of wages, family income, firm growth, and real investment based on IRS microdata, and relates these measures with the available data on OZ investment. Section 3.5 concludes with a discussion of this new evidence in relation to other recent studies, and provides roadmap for future research.

3.2 Opportunity Zones: Brief Background

Historical Overview

In February of 2017, a bipartisan group of U.S. Senators and Representatives introduced the Investing in Opportunity Act, which was later incorporated into the Tax Cuts and Job Act enacted by Congress in December of that year. The Congressional authors of the legislation described the goals of the OZ program in a joint public statement:

*"Too many American communities have been left behind by widening geographic disparities and increasingly uneven economic growth. [...] Americans should have access to economic opportunity regardless of their zip code. The Investing in Opportunity Act will unlock new private investment for communities where millions of Americans face the crisis of closing business, lack of access to capital, and declining entrepreneurship. [...] With this bill, we will dramatically expand the resources to restore economic opportunity, job growth, and prosperity for those who need it most."*³

The legislative focus on capital subsidies, rather than wage or employment subsidies, distinguishes the OZ program from the federal Empowerment Zone initiative launched in 1995, and is more ambitious in scope but similar in spirit to the federal New Markets Tax Credit Program enacted in 2000. In accordance with Congressional goals, an important aim of this research is to estimate how responsive investment has been to the OZ tax subsidy, how investment has affected local workers and businesses, and how these impacts may be heterogeneous across individuals who live, work, and invest in Opportunity Zones.

³Statement by Senators Cory Booker and Tim Scott and Representatives Ron Kind and Pat Tiberi, February 2, 2017.

Tract Eligibility and Nomination Process

The primary geographic units of the OZ program are *census tracts*, which we interchangeably refer to as *neighborhoods* or just *tracts*. Census tracts are small spatial units of approximately 4,000 residents, with coverage spanning the entirety of the United States.

Congress determined that tracts would be eligible for OZ designation if they could be classified as a *low income community* (LIC), defined as a tract with a poverty rate above 20% or median family income (MFI) less than 80% of the area median.⁴ In practice, policymakers used estimates of tract poverty rates and median family income from the 2015 5-Year American Community Survey (ACS) to assess eligibility.

Congress also allowed for a small number of tracts to be eligible for OZ designation even if they did not meet the poverty or income thresholds. Tracts classified as high-migration rural communities or low-population communities were deemed eligible, as were tracts with median family income of less than 125% of an adjacent eligible low income community.⁵ However, the vast majority of designated OZ tracts (97%) were deemed eligible on the basis of their poverty rate or median family income in the 2015 American Community Survey rather than these alternate criteria.

After Treasury and IRS determined which tracts were policy-eligible, state governors were given three months to nominate tracts for OZ designation. States could nominate up to 25% of their eligible tracts, and less populated states were granted a minimum of 25 OZs. Treasury accepted all state nominations from April to June of 2018, and ultimately designated 8,764 tracts ($\approx 12\%$ of all tracts) as Opportunity Zones. We explore the characteristics of eligible and chosen OZ tracts in greater detail in Section 3.2.

OZ Tax Subsidies

The central policy instruments of the OZ program are capital subsidies — specifically, highly favorable tax treatment of income accrued from capital gains. Investors intending to claim the tax benefit must (a) register their business as a Qualifying Opportunity Fund (QOF) with the IRS, (b) liquidate an existing asset, and (c) re-invest the capital gains into qualifying OZ assets. There are three main tax advantages conferred on these investments, which we summarize below.

First, taxes owed on capital gains from liquidating the initial asset are deferred until the fund sells its subsequent OZ investments or until the end of 2026, whichever is sooner. Since investors may redeploy this taxable income into income-bearing assets until the tax is due,

⁴For rural tracts, the area MFI is taken to be the statewide median family income. For urban tracts, the area MFI is the larger of the statewide MFI and the metropolitan area MFI.

⁵A high-migration rural community is defined as a census tract located within a high-migration rural county whose median family income was 85% of the statewide median family income. High migration rural counties are those that have had net outmigration of greater than 10% over the period 1990-2010. A low-population tract is a tract within an empowerment zone, contiguous to at least one LIC, with a population of less than 2,000. No more than 5% of a state's tracts could be nominated on the basis of meeting the adjacency criteria.

the deferral is potentially lucrative. Second, investors who hold qualifying OZ assets are eligible for a step-up in basis on their initial capital gains after 5 years (10%) and 7 years (15%), directly reducing tax liability. Finally, investors who hold qualifying OZ assets for at least 10 years may claim a 100% reduction in capital gains tax on appreciation of those OZ assets. The capital gains tax rate typically ranges from 15-20%, and so full elimination of the tax represents a large and significant subsidy.⁶

Broadly, QOF funds may invest in two categories of assets: (1) stock and partnership interests in qualifying operating businesses (QOB), and (2) qualifying property (QOP), which can be leased or owned. Qualifying OZ businesses (that is, firms receiving investment from QOFs) must meet regulatory criteria requiring that their core economic activities occur within the boundaries of a designated OZ tract, and property investors are generally legally required to demonstrate “substantial” capital improvements in real estate assets.⁷ These regulations were introduced by the Treasury Department to curb tax evasion, and to increase the likelihood that OZ investments spur real economic activity and opportunity for OZ workers and residents.

Private-sector investors estimate that, under a range of plausible assumptions about discount rates and rates of return on OZ capital, investors who maximally leverage the OZ policy incentives may ultimately increase their after-tax return by approximately 40%.⁸ The OZ program thus introduces a large spatial capital tax wedge that varies sharply even across neighborhoods within the same city.

Summary Statistics

Table 1 shows tract-level demographic summary statistics to illustrate differences between OZ-eligible tracts (Column 1), OZ-designated tracts (Column 2), and the country as a whole (Column 3). The data are from the 5-Year 2015 American Community Survey, and corresponds to the data used by the IRS and Treasury to determine which tracts were eligible to be nominated by states as Opportunity Zones. Column 4 shows differences between OZ-designated tracts and OZ-eligible tracts, and Column 5 calculates the relevant p-values. The table shows that designated OZ tracts (Column 1) tend to have lower incomes, home values, and education attainment – and higher poverty rates and non-white population shares – relative to tracts that were eligible for the OZ tax subsidy but were not nominated by the states (Column 2). Eligible tracts are, in turn, socioeconomically disadvantaged relative to the country as a whole (Column 3).

This evidence, corroborated by other researchers, is consistent with the view that both federal and state lawmakers generally intended to target OZ investment toward populations most in need. In the following section, we present new evidence on take-up and the spatial distribution of OZ investment across tracts.

⁶IRS provides further details on capital gains tax rates [here](#).

⁷The IRS provides further details on these regulatory requirements [here](#).

⁸See e.g. [Weinstein and Glickman \(2020\)](#).

TABLE 1: TRACT SUMMARY STATISTICS

	(1) OZ Tracts	(2) Eligible, Not Chosen	(3) All Tracts	(4) Diff (1-2)	(5) p-val
Population	3,999 (1,908)	4,041 (1,860)	4,326 (2,129)	-42	0.07
Rural	0.21 (0.41)	0.18 (0.39)	0.16 (0.37)	0.03	0.00
Median Age	35.6 (7.3)	35.9 (7.5)	38.9 (7.7)	-0.3	0.00
% White	0.58 (0.29)	0.63 (0.28)	0.73 (0.25)	-0.05	0.00
% Black	0.26 (0.30)	0.21 (0.27)	0.14 (0.22)	0.05	0.00
% Foreign Born	0.15 (0.16)	0.17 (0.17)	0.14 (0.14)	-0.02	0.00
% High School	0.49 (0.12)	0.51 (0.12)	0.58 (0.12)	-0.02	0.00
% College	0.11 (0.08)	0.12 (0.09)	0.20 (0.14)	-0.01	0.00
Median Family Income	38,978 (15,401)	46,000 (16,317)	68,357 (33,997)	-7022	0.00
% Poverty Rate	0.29 (0.13)	0.25 (0.11)	0.16 (0.12)	0.04	0.00
Median Home Values (1000s)	696 (495)	748 (465)	1,021 (632)	-52	0.00
Household Gini	0.46 (0.06)	0.44 (0.06)	0.42 (0.06)	0.02	0.00
N	8,638	23,699	74,001		

Notes: Unit of analysis is 74,001 census tracts. Demographic data are from the 2015 5-Year American Community Survey (ACS). Table excludes tracts with missing ACS data. The table shows means and standard deviations in parentheses. OZ tracts (Column 1) are socioeconomically disadvantaged relative to eligible-but-not-chosen tracts (Column 2), which are in turn disadvantaged relative to the country as a whole (Column 3). Column 4 computes the difference in means between Columns 1 and 2, and Column 5 presents p-values testing the null hypothesis that these means are equal. These data are consistent with the view that policymakers intended to target populations most in need.

3.3 Descriptive Evidence on OZ Investment

OZ Data in Federal Tax Records

We measure OZ program investment for all businesses that filed an electronic copy of IRS Form 8996 in tax years 2019 and 2020. This form requires QOF funds to identify the firms and census tracts in which they are investing, as well the corresponding dollar values. These data do not yet cover OZ investments from businesses that submitted paper copies of their tax returns, nor do they cover data from subsequent tax years. We provide details about how line items in Form 8996 correspond to definitions in this paper in the Online Appendix.

The first two columns of Table 2 show that QOF businesses reported approximately \$26.7 billion in OZ-subsidized capital investment flows in 2019 and an additional \$14.8 billion in flows in 2020, for a cumulative total of \$41.5 billion by the end of 2020.

TABLE 2: INVESTMENT IN OPPORTUNITY ZONES OVER TIME

Tax Year	Total Annual Flows (mil)	Cumulative Flows (mil)	Cumulative E-Filed (mil)	E-Filed Share	Tracts (#)	QOF (#)	QOB (#)
2019	26,670	26,670	18,779	0.70	1,347	2,526	2,224
2020	14,789	41,459	32,504	0.78	3,242	3,514	3,281

Notes: Data in the first two columns are from the universe of 8996 QOF returns. Data in all remaining columns cover only electronically-filed 8996 tax returns. The final four columns show cumulative values. We provide additional details about these data in the Online Appendix.

We calculate that the electronic Form 8996 returns in our analysis sample cover approximately 78% of the cumulative value of QOF investments in tax year 2020. Among the sub-sample of e-filers for which we have detailed data, the number of OZ tracts receiving any investment more than doubled from 2019 and 2020 and the cumulative number of QOFs (investors) and QOBs (investees) increased substantially as well.

That our sample is limited to electronic filers naturally invites the question: how representative are electronic filers of all 8996 filers? Since the OZ program and its associated tax forms are new, historical patterns provide limited guidance in assessing possible differences between electronic and paper filers. Even if electronic filers on average make similar investment decisions to paper filers, the descriptive estimates presented below should nevertheless be interpreted as providing a lower bound on aggregate OZ investment by tax year 2020.

Caveats aside, the existing data from electronic filers provide an emerging picture of OZ investment to date. In what follows, we describe the data sources, present aggregate summary statistics, break out investment by industry and geography, and correlate investment flows with demographic, industry, and firm characteristics. Overall, the data show that OZ investment is highly spatially concentrated, is directed toward the real estate

and construction sectors, and gravitates toward tracts with relatively higher educational attainment, income, density, and pre-existing upward income and population growth trends.

OZ Investment is Spatially Concentrated

We begin with a broad overview of the Form 8996 data. Table 3 shows that businesses filing electronic 8996 returns reported approximately \$32.5 billion in cumulative OZ-subsidized capital investments by 2020. In total, we observe 3,953 QOF funds investing in 3,677 QOB businesses across 3,242 OZ census tracts. Panel A reveals that this investment is highly concentrated in a small share of tracts: in fact, 5,522 of 8,764 OZ tracts in our sample (63%) appear to receive zero investment. We also find that approximately \$3.2 billion of this investment (10%) is not associated with a designated Opportunity Zone tract; this may reflect regulatory guidance allowing QOF funds to invest a fraction (10%) of their assets in non-OZ tracts, as well as taxpayer or administrative error.

TABLE 3: INVESTMENT IN OPPORTUNITY ZONES BY TYPE


	(mil)	Share	# OZ Tracts	# QOF	# QOB
Panel A: By Tract Type					
OZ tract, >0 investment	29,267	0.90	3,242	3,514	3,281
OZ tract, no investment	0	0	5,522	0	0
Unmatched tract	3,236	0.10	n.a.	511	501
Panel B: By Investment Type					
Stock or Partnership Interest	32,478	1.00	3,573	3,953	3,677
Owned or Leased Property	25	0.00	6	0	0
Panel C: By QOB Entity Type					
Partnership	23,217	0.71	3,102	2,623	2,149
Other	9,286	0.29	1,121	1,558	1,528
Total	32,504	1.00	3,242	3,953	3,677

Notes: Data based on IRS records of [Form 8996](#) from electronic filers in tax year 2019. Columns need not always sum to totals. The table shows that OZ investment is highly concentrated in a small number of Opportunity Zones and in partnership interests.

Panel B of Table 3 shows that OZ investment is virtually entirely concentrated in equity and partnership interests (100%) rather than property (0%). In Panel C we identify the legal entity type of the QOB businesses receiving investment from OZ funds, and confirm that this investment is overwhelmingly concentrated in partnerships (71%). Structuring a business as a partnership offers owners several legal and economic advantages over alternative entity types, but in our setting perhaps the most important is that partnerships allow taxable depreciation deductions (such as those resulting from real estate depreciation) to flow through to the investors.

Although Table 3 shows that the vast majority of OZ tracts do not attract any investment, the tracts that do receive investment report large and economically significant amounts. Panel A of Figure 1 shows the distribution of investment for tracts that received at least \$5,000 by 2020, and Panel B shows these values normalized on a per-resident basis. Among these tracts, median OZ investment is \$1.6 million, or \$386 per resident. Overall, the distribution of investment across all OZ tracts is highly skewed, such that the top 5% of tracts receive 78% of total investment, and the top 1% of tracts receive 42% of total investment.

FIGURE 1: TRACT-LEVEL DISTRIBUTION OF OZ INVESTMENT



Notes: N=3,230 census tracts with at least \$5,000 of OZ investment. Data based on electronic filers of IRS form 8996 in tax year 2020. Panel A shows the distribution of total investment, and Panel B shows the distribution of investment per capita. We use log scales on the x-axes and exclude tracts with less than \$5,000 of investment to improve the data visualization. The figures underscore that OZ investment is highly spatially concentrated: the top 5% of OZ tracts receive 78% of total investment, and the top 1% of tracts receive 42% of total investment. As shown in Table 3, the bottom 63% of tracts receive zero investment. Among tracts that receive >0 investment, the median investment of \$386 per resident is economically large relative to existing federal place-based programs.

In summary, many neighborhoods have received no OZ investment, but for those that do, the amount of investment can be quite large. Low tract-level take-up rates may help to explain estimates of modest or null intent-to-treat effects in existing research (e.g., Chen, Glaeser, and Wessel 2019; Atkins, Hernandez-Lagos, Jara-Figueroa, and Seamans 2020; Freedman, Khanna, and Neumark 2021). Among tracts where QOFs do report investment, the extent to which these financial investments translate into physical capital expenditures that would not have occurred in the absence of the OZ tax subsidy is a question we are investigating in ongoing research.

Industry Composition of OZ Investment

We next examine how OZ investment varies across industries. Panels A and B of Table 4 show the NAICS-2 composition of QOF funds and QOB businesses, respectively. Both

QOF investor funds and recipient QOB firms are mainly in the business of real estate, with smaller but significant shares in related industries such as construction, finance, and management. Panel B shows that approximately 52% of OZ dollars are invested in real estate firms, while 11% is invested in construction firms, and 9% in finance. In the Online Appendix, we further decompose industry composition of funds and recipient firms using finer 6-digit industry codes, and show that both residential and non-residential real estate businesses attract considerable OZ investment.

TABLE 4: INDUSTRY COMPOSITION OF FUNDS AND RECIPIENT FIRMS

Panel A: QOF Investor Funds

NAICS	Industry	# Funds	(mil)	Share
53	Real Estate, Renting, and Leasing	1,943	13,738	0.42
52	Finance and Insurance	1,238	9,376	0.29
23	Construction	384	1,805	0.06
55	Management of Companies	200	1,724	0.05
–	Other	188	5,860	0.18
	Total	3,953	32,504	1.00

Panel B: QOB Firms Receiving Investment

NAICS	Industry	# Targets	(mil)	Share
53	Real Estate, Renting, and Leasing	2,066	16,778	0.52
23	Construction	448	3,733	0.11
52	Finance and Insurance	326	2,985	0.09
55	Management of Companies	81	888	0.03
72	Lodging and Restaurants	111	842	0.03
54	Professional Services	73	823	0.03
31	Manufacturing	54	325	0.01
–	Other	214	4,220	0.13
–	Unknown	304	1,909	0.06
	Total	3,677	32,504	1.00

Notes: Data based on electronic filers of IRS form 8996 in tax years 2019 and 2020. Panel A shows the number of QOF investor funds, dollar values, and dollar share of OZ investment, and Panel B shows analogous measures for the QOB firms receiving investment. QOF investment is highly concentrated in real estate, with smaller but significant shares in related industries such as construction, finance, and management..

Several factors help to explain why OZ funds exhibit a preference for real estate investments. First, real estate is a highly capital-intensive sector. The Bureau of Economic Analysis estimates that residential and non-residential structures account for approximately

39% of annual private fixed asset investment.⁹ Second, investment in real estate is geographically versatile and thus well suited to benefit from a tax subsidy that applies broadly to heterogeneous neighborhoods. Virtually any area of the country with population growth is likely to need new housing and commercial structures.¹⁰ By contrast, other capital-intensive sectors such as oil refineries or manufacturing plants are unlikely to sprout up, for example, in dense urban areas. Third, specialists in the real estate sector may be uniquely situated to facilitate financing and reduce transaction costs associated with investment. This specialization is reflected in part by the large number of real estate funds in Panel A of Table 4. Similarly, local real estate developers may have portfolios of potential projects that can be prioritized or de-prioritized depending on the price and availability of capital financing. Fourth, the widespread availability of data on real estate price trends may help investors to identify investments likely to have higher returns and lower risk. Finally, legal and regulatory considerations also favor investments in real estate over other sectors; see [Hadjiligiou, Lutz, and Bruno \(2021\)](#) for a review.

Demographic Correlates of OZ Investment

In this section we explore how OZ investment is correlated with tract demographics. Figure 2 compares demographic characteristics for three groups of census tracts: (1) OZ tracts receiving positive investment from QOFs; (2) OZ tracts receiving zero investment from QOFs; and (3) all tracts nationally. In Panels A and B, these demographic characteristics are computed from the 2017 American Community Survey, while in Panel C we use data from the 2016 Census Longitudinal Employer-Household Dynamics LODES data. We standardize the variables to have mean zero and standard deviation one, and report how OZ tracts that receive QOF investment differ in standardized units from all tracts and from OZ tracts that do not receive investment. The confidence bars report 95% confidence intervals computed using robust standard errors.

Panel A of Figure 2 shows that, relative to the general population, OZ tracts that receive investment have on average fewer residents with a college degree, lower incomes, and higher poverty rates. Conversely, when compared to other OZ tracts with zero investment, tracts that receive investment have relatively high educational attainment, home values, and incomes, as well as lower unemployment and higher shares of prime-age workers. The interpretation of the coefficients is, for example, that the share of college graduates in tracts that received OZ investment is on average 0.55 standard units lower than the national average, and 0.22 standard units higher than in OZ tracts that did not receive any investment. For reference, we report the raw mean and standard deviation of these variables for all tracts on the right-hand side of the figure, and also report the raw means for each of these outcomes and groups of tracts in tables in the Online Appendix.

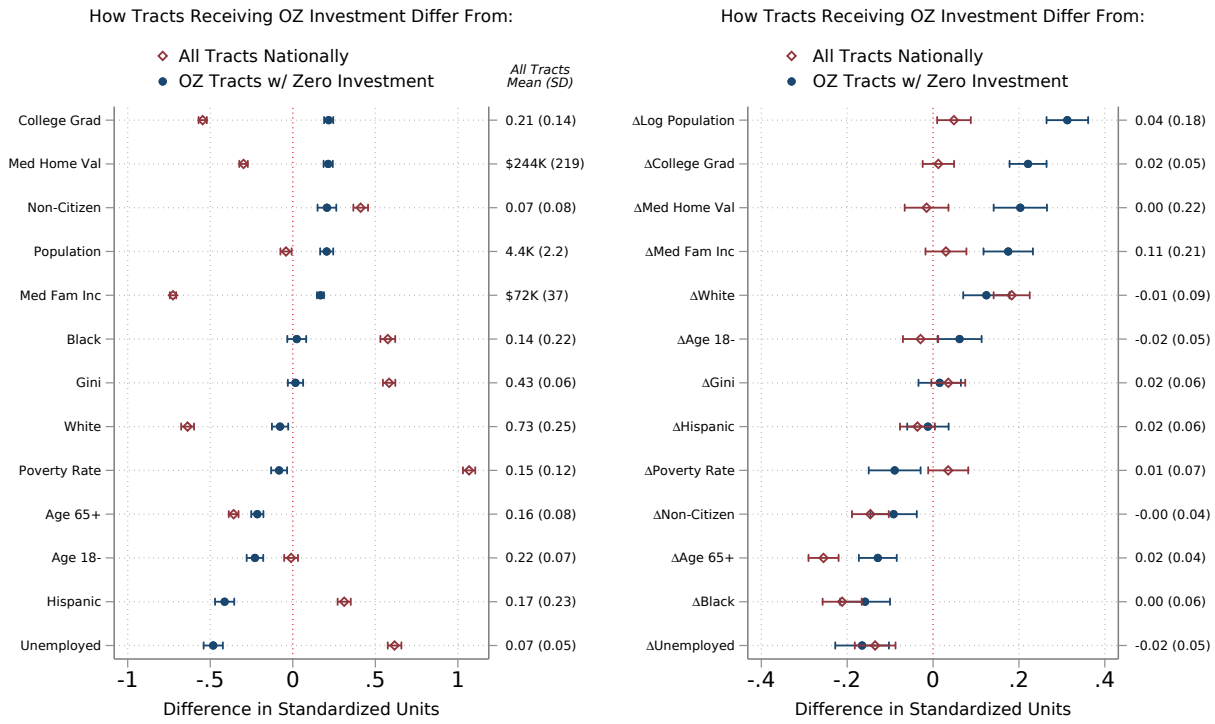
⁹See BEA [Table 5.10](#): Changes in Net Stock of Produced Assets (Fixed Assets and Inventories).

¹⁰In Section 3.3 we show that population growth is indeed a strong predictor of OZ investment.


FIGURE 2: DEMOGRAPHIC CORRELATES OF OZ INVESTMENT

Panel A: 2017 Demographics

Panel B: 2010-2017 Demographic Trends



Panel C: 2016 Tract Workforce Industry Composition



Panel B of Figure 2 correlates OZ investment with 2010-2017 demographic trends. Among OZ tracts, QOF funds invested in neighborhoods where incomes, population, and the share of college educated residents have increased sharply over the past decade, and where the non-white and elderly share of the population have declined. However, the figure also shows that trends in these neighborhoods are similar to trends in the rest of the US. As in Panel A, these results overall point towards investment in tracts with relatively greater pre-existing economic opportunity.

Lastly, Panel C of Table 2 compares the 2016 industry composition of the workforce across these three groups of tracts. On average, tracts with higher 2016 shares of workers in professional and amenity services – such as finance, management, restaurants, and the arts – attracted more OZ capital by 2020 relative to other OZ tracts. By contrast, QOF funds were less likely to invest in OZ tracts with higher workforce shares in healthcare, manufacturing, education, or retail. Relative to all tracts, tracts receiving OZ investment have a significantly larger share of government workers and a smaller share of construction workers.

Taken together, the three panels in Figure 2 paint a consistent picture. Although all OZ tracts are relatively disadvantaged in comparison to the rest of country, the tracts that received investment were the least disadvantaged of those granted OZ status. Moreover, the preliminary descriptive evidence suggests that OZ capital may disproportionately benefit a narrow subset of tracts in which economic conditions were already improving prior to implementation of the tax subsidy.

In the Online Appendix, we show variations on Panels A and B to illustrate how the characteristics of tracts receiving OZ investment have changed from 2019 to 2020, relative to OZ tracts that did not receive investment in either year. While QOF investment in 2020 continued to favor relatively well-off OZ neighborhoods, the figure shows that this pattern was attenuated relative to 2019.

Geographic Patterns in OZ Investment


We now explore geographic patterns in OZ Investment. Panels A and B of Figure 3 show total investment and investment per OZ resident, respectively, for the top 25 commuting zones.¹¹ The diverse list of commuting zones in Panel A reflects that QOF funds reported investment in virtually every region of the country and, not surprisingly, that the most populous commuting zones such as New York and Los Angeles generally received the most investment. Panel B shows that, on a per capita basis, mid-size commuting zones like Salt Lake City, Nashville, and Tampa received the most investment, although QOF investors did not neglect larger commuting zones like Denver, San Francisco, and Phoenix. OZ investment in Hunstville, Alabama is especially large and appears to be an outlier relative to other eligible OZ labor markets.

In Figure 4 we zoom in at a finer level of detail and map the spatial distribution of OZ investment in six illustrative cities: Brooklyn, Philadelphia, Detroit, Nashville, Hunstville,


¹¹The Online Appendix shows these statistics for the top 50 commuting zones.

FIGURE 3: OZ INVESTMENT IN 25 TOP COMMUTING ZONES

Panel A: Total Investment



Panel B: Investment Per OZ Resident



Notes: Panel A shows total OZ investment by commuting zone, and Panel B shows investment per OZ-resident, normalizing by the population of tracts with >0 investment. We compute investment from electronically-filed business tax records of Form 8996 in tax years 2019 and 2020. The panels present data for the top 25 commuting zones, excluding those with few QOF funds and/or QOB businesses to protect taxpayer privacy. The figure shows that QOF’s invested in diverse labor markets in nearly every region of the country. The Online Appendix shows these statistics for the top 50 commuting zones.

and Los Angeles. Dark red areas on the maps indicate OZ tracts with >0 investment, and pink areas indicate OZ tracts that receive zero investment. Grey areas indicate tracts that are not Opportunity Zones. These illustrative examples suggest that OZ investment gravitated toward dense city centers and central business districts (or, in Brooklyn, the neighborhoods most proximate to Manhattan).

We confirm generalizable relationships between investment, population density, and distance from the city or commuting zone center in Panel A of Figure 5, which shows how tracts receiving positive OZ investment differ in economic geography from all tracts

and from OZ tracts that did not receive investment. Tracts receiving OZ investment are on average more densely populated and urban relative to other OZ tracts and relative to non-OZ tracts. These tracts are also closer to the centers of commuting zones relative to other tracts.¹² OZ investment is also decreasing in the distance between investor funds and OZ tracts. Panel B of Figure 5 shows the distance distribution between OZ funds and the census tracts in which they invest, and Panel C plots fund-by-tract-level investment against the log distance between funds and OZ tract. Consistent with empirically and theoretically documented linkages between spatial proximity and economic activity, investment between QOFs and QOBs is declining in distance. In the next section we further explore how the locations of not only QOB businesses, but also QOF investors, may have implications for understanding the geographic incidence of the OZ program.

¹²We define the commuting zone center as the census tract with the largest number of jobs in the municipality (commuting zone) in which the tract located.

FIGURE 4: MAPPING OZ INVESTMENT IN SIX ILLUSTRATIVE CITIES

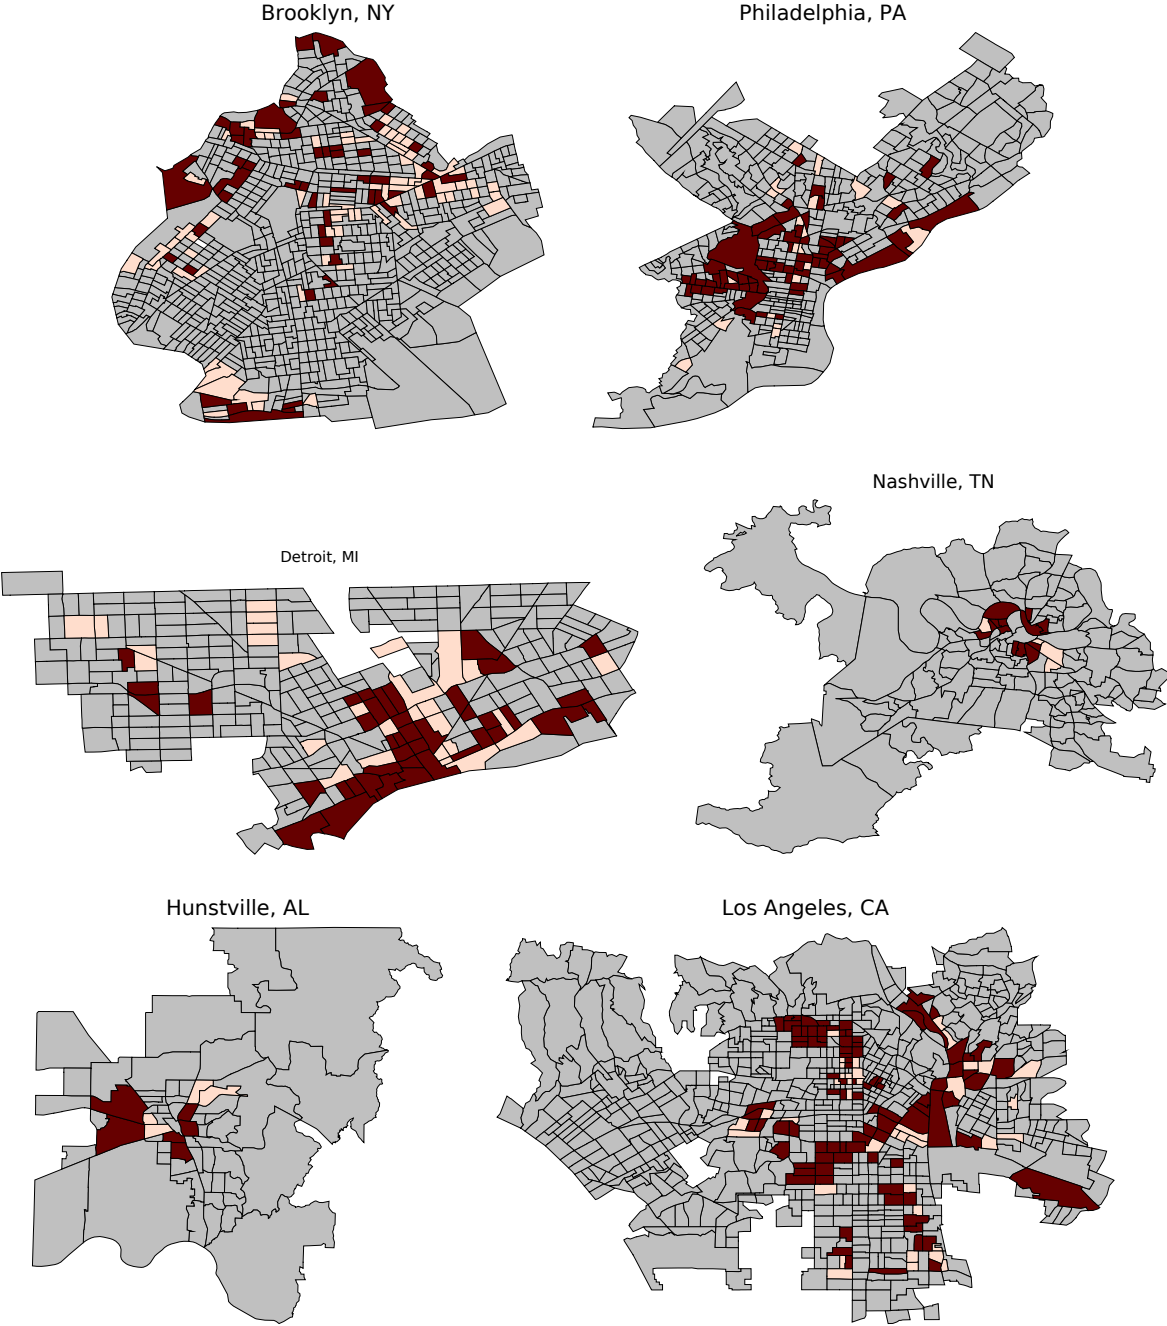




FIGURE 5: OZ INVESTMENT, POPULATION DENSITY, AND DISTANCE


Panel A: Population Density, Distance from City and CZ Center



Panel B: Distance Between Fund-Tract Pairs



Panel C: Investment vs. Fund-Tract Distance



Notes: The sample is N=74,288 census tracts. Panel A shows differences in economic geography for three groups of tracts: (1) OZ tracts receiving positive investment; (2) OZ tracts receiving zero investment; and (3) all tracts nationally. On average, QOF funds invest more heavily in densely populated, urban neighborhoods closer to city and commuting zone centers. Panel B reports the distribution of distances between fund-tract pairs, and Panel C plots fund-level investment against distance from OZ tracts using a smooth polynomial fit. The plots highlight that OZ investor funds tend to be located (or, set up ex-post) in locations very close to OZ tracts.


Income and Geography of QOF Investors

We have focused so far on describing the economic and sectoral characteristics of QOF investments, as well as the demographic and geographic characteristics of neighborhoods that receive those investments. Apart from residents of OZ neighborhoods, the incidence of the OZ program will naturally also fall in part on QOF investors, who are likely to most directly benefit from the tax incentives described in Section 3.2. In this section we briefly describe the income and geographic profiles of QOF investors in the available data.

To estimate the household income of QOF investors, we link QOF partnerships to their partners using the universe of 1065-K1 information return filings, which must be reported to IRS annually for all partners. Partnership ownership structures can be complex — for example, higher-tier partnerships may include both individuals and/or lower-tier partnerships as partners — rendering a complete match of these data to be difficult. Nevertheless, we are able to match approximately 89% of the partners of higher-tier QOF partnerships to individuals, who we then link to our household income database. In Figure 6, we show the distribution of household income for these QOF investors relative to the general US population.

The plot shows that, on average, QOF investors have substantially higher household income relative to the general US population. We estimate 2019 median and average household income for QOF investors to be \$741,000 and \$4,852,000, respectively — an order of magnitude higher than the national median and average household incomes of \$69,000 and \$117,000, respectively. While tax benefits to QOF investors will ultimately depend on the extent to which their investments appreciate in value over time, these results suggest that the direct tax incidence of the OZ program is likely to benefit households in the 99th percentile of the national household income distribution.

FIGURE 6: INCOME DISTRIBUTION OF QOF INVESTORS AND THE US POPULATION




Notes: The plot shows the distribution of 2019 household income for QOF investors relative to the general US population. We identify QOF investors by linking QOZ partnerships to their partners using IRS Form 1065-K1, an information return that must be filed annually for all partners. Household income computations are described in the Online Appendix. We winsorize the top 1% of the QOF income distribution to improve the data visualization, and exclude households with negative income. Median and average household income for QOF investors is approximately \$741,000 and \$4,852,000, respectively, relative to the national median and average of \$69,000 and \$117,000, respectively.

Finally, in Figure 7, we link QOF investors to their state of residence, and estimate total the value of QOF investments coming from each state. To perform this computation, we again focus on QOF partnerships, and further make the simplifying assumption that all partners of a fund are equally invested in it. Panel A shows the resulting aggregate QOF investment that we assign to each state, scaled in million of dollars, and shows that the bulk of QOF dollars flow from populous and relatively wealthy states such as California, Texas, Florida, New York, and New Jersey.


Panel B shows these aggregate totals scaled by state population, and shows that investors disproportionately reside in the Northeast and Pacific Coast, as well as a few states in the Mountain West such as Nevada, Utah, and Colorado. The maps highlight that the geographic incidence of the OZ program depends not only on which OZ tracts receive QOF investments, but also on the residential locations of QOF investors.

FIGURE 7: GEOGRAPHY OF QOF INVESTORS

Panel A: Total QOF Investment (mil \$), by Investors' State of Residence



Panel B: QOF Investment Per Capita, by Investors' State of Residence



Notes: We link QOF investors to their state of residence by linking QOZ partnerships to their partners using IRS Form 1065-K1, an information return that must be filed annually for all partners. Panel A shows the resulting aggregate QOF investment that we assign to each state, scaled in million of dollars, and shows that the bulk of QOF dollars flow from populous and relatively wealthy states such as California, Texas, Florida, New York, and Illinois. Panel B shows these aggregate totals scaled by state population, and shows that investors disproportionately reside in the Northeast and Pacific Coast, as well as a few states in the Mountain West and Great Plains such as Nevada, Utah, and Colorado. The maps highlight that the geographic incidence of the OZ program depends not only on which OZ tracts receive QOF investments, but also on the residential locations of QOF investors.

3.4 New Panels from IRS Microdata

We construct new annual panels of individual tract-level outcomes using rich data from federal tax records. We provide an overview of our data sources below, and provide more detailed discussion of our data processing in the Online Appendix. We then show how these measures correlate with the available data on OZ investment.

Individual- and Business-Level Federal Tax Records

We leverage the universe of de-identified federal individual- and business-level tax records from the Internal Revenue Service (IRS) to construct novel tract- and block-level measures of economic activity. On the individual side, our work builds on [Larrimore, Mortenson, and Splinter \(2019\)](#), who map virtually all individuals residing in the United States to household identifiers using address data from 1040's and information returns. While taking stringent precautions to protect taxpayer privacy, we use open-source and commercial geocoding services to match household addresses with latitude and longitude coordinates, and to locate households within 2010 census tract boundaries.

We use the individual-level tax data to construct measures of household and family income, poverty, employment, wages, migration, and commuting that closely correspond to analogous measures from publicly available data. The new measures incorporate data from both income tax returns and information returns (such as W2s and 1099s), and thus allow us to observe income even for individuals and households that do not file income tax returns.

On the business side, our sample of firms includes the universe of corporations and partnerships, and excludes self-proprietorships. We link all businesses to their parent companies using the crosswalks constructed by [Dobridge, Landefeld, and Mortenson \(2019\)](#), and geocode them based on the address information provided on the cover form of their annual tax returns. We further link firms to their employees using W2s, and construct firm-level measures of real investment from Form 4562 following [Yagan \(2015\)](#). These measures of real investment capture firm spending on tax-deductible depreciable assets such as buildings, machinery, computer, vehicles, and office furniture.

A limitation of the business tax data is that we are unable to observe the establishment locations of multi-establishment firms. This means, for example, that if a large national retail chain were to purchase new buildings in multiple states, we would be unable to observe the location of such investments. Thus, the firm-level measures must be interpreted with caution. When aggregating the firm-level data, we differentiate firms based on firm size: since smaller firms are less likely to have multiple establishments, they may provide a more geographically accurate picture of local economic conditions even if they are not representative of all firms.

In total, we geocode more than one billion individual- and business-level tax returns from 2010 to 2019. We aggregate our resulting measures to the census-tract level and, in the following section, evaluate their validity in relation to publicly available datasets. We then correlate these measures with the available evidence on OZ investment.


Evaluating the New Tract-Level Measures Against Public Data

We probe the validity of our new measures by comparing them with analogous measures from publicly available data. Figure 8 compares our tract-level 2017 estimates of income, poverty, and population based on IRS data with survey-based estimates of these measures from the 2017 5-year Census American Community Survey (ACS). The 2017 5-year ACS pools together and averages survey responses from five consecutive years of 1% national population surveys from 2013-2017, which allows the Census to estimate population demographics at the tract-level using larger sample sizes. By contrast, our IRS-based measures are based on the universe of federal tax returns from a single tax-filing year.


In Panels A and B, we use binscatter plots to compare our tract-level IRS measures of median household income (MHI) and median family income (MFI), respectively, with the ACS data. Each point in these plots represents a simple average of an approximately equal number of census tracts. The plots also report the regression coefficient, standard error, and R-squared obtained from regressing the IRS measure on the ACS measure using OLS. The slopes of the lines are close to one, implying that a 1% increase in the ACS income is on average associated with an approximately 1% increase in the IRS income. Our IRS-based estimates are systematically higher than the ACS estimates, due primarily to the fact that our IRS measures are based only on data from tax-year 2017, whereas the ACS is based on five-year averages from 2013-2017. The upward level-shift thus represents income growth and inflation relative to the ACS measure.

FIGURE 8: IRS MEASURES VS. ACS MEASURES


Panel A: Median Family Income




Panel B: Median Household Income



Panel C: Poverty Rate



Panel D: Population



Notes: $N=72,349$ census tracts with non-missing IRS and ACS measures. The figures compare 2017 5-year ACS tract-level outcomes (horizontal axis) with their corresponding 2017 IRS measure (vertical axis). The blue line shows the line of best fit, and the red line shows the 45-degree line. In Panels A and B, our IRS measures of median household and family income are systematically higher than the ACS measures, since the latter represent five-year 2013-2017 pooled averages whereas the former are based on data only from 2018. We find higher income at the bottom of the tract-level income distribution, due to underreporting of wage and private-retirement income among low-income households in ACS relative to what we observe from information returns in IRS data. This pattern is also reflected in systematically lower poverty rates in our IRS-based measures relative to ACS in Panel C. In Panel D, we find modestly lower population counts in our sample relative to the ACS, driven by households that we are unable to locate from our geocoding procedure.

Panels A and B also reveal that the IRS data yield higher estimates of income at the lower end of the tract-level distribution relative to ACS. This difference reflects underreporting of wage and private-retirement income among low-income households in ACS relative to what we observe from information returns in IRS data (Bee and Rothbaum, 2017; Larrimore, Mortenson, and Splinter, 2020). Consistent with this result, in Panel C we estimate systematically lower poverty rates in the IRS data relative to the survey-based ACS measures.

Finally, Panel D compares our IRS population sample with the ACS estimate of tract

population. The gap between the IRS and ACS population estimates is driven by individuals for whom we are unable to assign a census tract using the geocoding procedure discussed in the Online Appendix. Overall, our geocoding procedure captures approximately 81% of the total US population, and approximately 85% of the population that does not report a PO Box address on their tax returns. The close alignment of the IRS- and ACS-based measures of tract-level income in Panels A and B suggest that any biases resulting from non-random biases in the geocoding procedure are likely to be small.

FIGURE 9: EMPLOYMENT IN IRS AND LODES DATA

Panel A: Employment by Tract of Residence

Panel B: Employment by Tract of Workplace



Notes: $N=71,809$ census tracts with non-missing IRS and LODES data. Panels A and B compare tract employment counts by residential and workplace tract locations, respectively, using our IRS measures and measures from the Census LODES data. The red line is a 45-degree line, and the blue line is the line of best fit. Panel A plots employment counts based on employees' tract of residence. Panel B plots employment counts based on employees' tract of workplace. Our IRS-based measure of workplace employment only covers small businesses with 1-49 employees, since the workplace location data for these businesses is likely to be more reliable; this definitional difference leads to a consistent gap between the IRS and LODES-based measures of workplace employment. Nevertheless, the high correlations between the IRS and LODES measures in both panels lend credence to the validity of the geocoding procedure.

As a final validity exercise, we calculate employment totals by census tracts of residence and workplace location. We compare these counts from our IRS measures with those available in the census LODES data, shown as binscatter plots in Figure 9. In Panel A, employment counts by tract of residence align well with the corresponding counts from the LODES data. We underpredict employment by residence for areas with little employment, reflecting poorer geocoding coverage among individuals in sparsely populated areas. In Panel B, employment counts by tract of workplace are highly correlated with those seen in the LODES data, although a sizeable gap exists between the two estimates. This gap reflects that we only tabulate workplace employment counts for small businesses with 1-49 employees – that is, businesses that are more likely to have only a single establishment and whose employees are

thus more likely to work at the same physical location as the address reported on the firm's tax filings. While differences between administrative and survey sources are natural, we find the high R-squared in both plots to be a reassuring signal of the quality of geocoding.

IRS Correlates of OZ Investment


Panels A and B of Figure 10 perform the same analysis as in Section 3.3 using our IRS measures. Panel A uses IRS-based measures from 2017, while Panel B uses changes in those measures from 2010 to 2017.¹³ As before, we standardize the variables to have mean zero and standard deviation one, and report how OZ tracts that receive QOF investment differ in standardized units from all tracts (in red) and from OZ tracts that do not receive investment (in blue). The confidence bars report 95% confidence intervals computed using robust standard errors.

The evidence presented in Figure 10 is broadly consistent with the evidence from Section 3.3. Among OZ-designated tracts, QOF funds invested in neighborhoods with higher wages, lower poverty rates, more employment, more firms, and higher levels of real investment. Still, these tracts receiving investment are economically disadvantaged relative to tracts nationally. Panel B uses 2010 to 2017 changes in the IRS measures to assess the extent to which OZ investment is correlated with recent neighborhood-level trends. While the magnitudes are smaller than those seen in Section 3.3, we find that QOF investment favored neighborhoods with higher income and firm growth. These patterns are most pronounced when the comparison group is OZs with no investment, but OZs also had higher rates of employment and median family income growth relative to all tracts nationally. The raw means for these figures can be found in the Online Appendix.

This evidence suggests that OZ tracts receiving investment from QOF funds were experiencing substantially different trends in economic activity relative to all tracts nationally and relative to OZ tracts that did not receive investment. A natural implication is that research designs that compare trend growth in OZ and non-OZ tracts to assess the causal impacts of the policy must be interpreted with care and caution. Comparable tracts should be balanced on a broad set of demographic and economic characteristics and trends to avoid spuriously conflating pre-existing trends with the causal effects of the OZ tax subsidy.

¹³For median family income and poverty rates, we use a shorter difference of 2015-2017, since we have not yet extended our IRS sample back to 2010.

FIGURE 10: IRS CORRELATES OF OZ INVESTMENT



Notes: N=74,001 census tracts. The figure shows differences in IRS measures for three mutually exclusive groups of census tracts: (1) OZ tracts receiving positive investment; (2) OZ tracts receiving zero investment; and (3) all other tracts. The data in Panels A and B are constructed from IRS microdata as described in Section 3.4 and the Online Appendix. All variables are standardized to have mean zero and standard deviation one. Error bands show 95% confidence intervals with robust standard errors. Among OZ tracts eligible for the tax subsidy, QOFs typically invested in neighborhoods with more firms, more employment, higher wages and income, and lower poverty rates.

3.5 Conclusion

We provide the first available evidence on the response of QOF investors to the OZ tax subsidy. We emphasize that this evidence is preliminary and does not yet incorporate data from paper tax filers, who we estimate account for approximately 78% of QOF investment dollars. The OZ investment data are based on business tax returns from tax years 2019 and 2020, the first two years that detailed OZ reporting requirements have made this analysis possible. We also emphasize that the patterns of investment described in this paper may evolve over time, perhaps particularly in response to the coronavirus pandemic beginning in 2020.

Caveats aside, the early evidence shows several striking patterns. We find that OZ investments are highly spatially concentrated in a relatively small number of census tracts, and are heavily concentrated in the real estate sector. Among tracts designated as OZs, investors favored neighborhoods with higher income, educational attainment, home values, and pre-existing population and income growth. These neighborhoods have also experienced significant changes in their demographic composition over the past decade, with increasing shares of college educated adults and declining shares of non-white residents. However, tracts that receive OZ investment are nevertheless considerably economically disadvantaged relative to all tracts nationally. We presented evidence consistent with these findings using a broad range of demographic measures from publicly available ACS data, and corroborated the results using a new panel of IRS-based tract-level measures. Finally, we find that the direct incidence of the OZ tax subsidy is likely to benefit taxpayers in the 99th percentile of national income distribution.

Our results help to contextualize findings from other recent studies on Opportunity Zones. As we have noted, a nascent research literature generally finds modest or null intent-to-treat (ITT) effects of the OZ program on neighborhood-level economic outcomes such as real estate prices, employment, job growth (Chen, Glaeser, and Wessel 2019; Atkins, Hernandez-Lagos, Jara-Figueroa, and Seamans 2020; Freedman, Khanna, and Neumark 2021). Our research raises the possibility that these null intent-to-treat effects may be explained by the fact that a majority of OZ tracts have not received any investment from QOF investors. However, existing research does not yet answer the question whether the OZ program induced positive economic changes in the set of neighborhoods that did receive investment from QOF investors.

An important goal for future research, then, is to estimate not only intent-to-treat effects, but also average treatment effects (ATE). Conditional on receiving OZ investment, what are the causal effects of the OZ program on real investment and local labor markets? In other words, to what extent has financial investment from QOF investors translated into business growth, employment, wage growth, and physical capital expenditures that would not have otherwise occurred in the absence of the OZ tax subsidy? The answers to these questions will be of central importance public and scholarly understanding of the Opportunity Zone program and of place-based policies more broadly.

Chapter 4

High-End Housing and Gentrification: Evidence from a San Francisco Lottery

In Chapter 4, Harrison Wheeler and I study the effects of housing regulations on gentrification. In major cities throughout the United States and the world, contentious debates over housing policy frequently focus on the extent to which high-end developments are a cause of gentrification. We shed new light on this question empirically by studying a unique lottery in the city of San Francisco that allowed a limited number of property owners to convert their buildings into high-end condominiums. Relying solely on exogenous variation from the lottery, we study the long-run effects of these developments on local home prices, demographics, and new business entry. Compared to losing lottery applicants, winners are vastly more likely to invest in alterations and renovations in their properties, to see their property values increase, to rent their properties to new tenants, and to sell them to new owners. The home values of adjacent buildings also increase. At the neighborhood level, conversions lead to higher resident incomes, home values, rental prices, and shares of the population that are White and college-educated.

This chapter is coauthored with Harrison Wheeler.

4.1 Introduction

Policymakers in cities across the United States and around the world are grappling with how best to address surging rents and home prices. Since the 1990s, a steady influx of high-income workers to major cities has led to rapid increases in home prices, changes in demographic characteristics, and shifts in the composition of local businesses --- a process commonly referred to as gentrification.

In response to these trends, many economists and policymakers have advocated relaxing regulatory barriers that restrict market-rate housing developments in expensive cities (e.g., [Furman 2015](#); [Glaeser 2017](#); [Hsieh and Moretti 2017, 2019](#)). Yet despite its popularity among economists, this policy prescription has proven highly controversial among the broader public. From the perspective of many observers, market-rate housing development seems to make problems worse: luxury condominiums sprout up in low-income neighborhoods, high-income residents continue to stream in, and local price growth continues unabated. Rather than taming the excesses of gentrification, opponents argue that market-rate housing developments cause and exacerbate it.

In this paper, we provide new evidence about the extent to which high-end, market-rate housing developments are a cause of gentrification. Empirical strategies to answer this question must address a fundamental econometric concern: Precisely because high-end developments are most likely to occur in neighborhoods with strong housing demand, comparisons between locations with and without such developments will overstate their role in driving neighborhood change. The econometric challenge is thus to isolate the causal effect of high-end developments independently from local shocks that may simultaneously induce neighborhood change.

To overcome this challenge, we study a unique administrative lottery in the city of San Francisco that permitted a limited number of property owners to legally convert buildings into high-end condominiums. In 1979, San Francisco banned condominium conversions due to widespread public concern about their effects on local home prices and the supply of rentable units. However, in response to growing cross-pressure from opponents of the ban, city officials in 1981 struck a compromise: Each year the city would run a lottery allowing a maximum of 200 winning units to convert their properties into condominiums.

Beyond the econometric appeal of the lottery, our focus on San Francisco and on condominiums is motivated by their central place in national debates about gentrification and housing policy. As we will discuss in greater detail, San Francisco has been a poster city of skyrocketing home prices and demographic change in recent decades, providing an ideal empirical setting for this research. At the same time, condominiums are front-and-center in controversies about housing policy not only in San Francisco, but in cities across the United States and around the globe. As a legal structure designed to facilitate ownership of units within multi-family buildings, condominiums are often attractive to high-income workers with preferences for living in dense urban areas. Condominiums are also generally exempt from rent control and tenant eviction protections that apply to other multi-family buildings, further fueling concerns about displacement of low-income residents.

To shed new light on these controversies, we study the long-run effects of lottery-induced condominium conversions in San Francisco on local home prices, demographics, and new business entry. To do so, we use annual lottery data from 2001 to 2013, including applicant information from both lottery winners and losers. We supplement the lottery panel with a rich and detailed suite of data sources and outcomes: address-level data on evictions, building permits, building characteristics, assessed home values, property sales, and homeowner-occupied status from the San Francisco Assessor’s Office; block-level data on new business and establishment entry from SF Open Data; and tract- and block-level demographic data from the U.S. Census Bureau.

Our empirical strategy combines exogenous variation from the lottery with a stacked difference-in-differences design to estimate the causal effects of condominium conversions on key outcomes. This framework allows us to assess the validity of the research design using balance tests; to combine information across lotteries; to consider treatment dynamics over a ten-year period; and to move seamlessly from estimating the intention-to-treat (ITT) effect of winning the lottery to the local average treatment effect (LATE) of converting to a condominium.

Lottery winners and losers are statistically indistinguishable on key outcomes and characteristics as far back as 12 years before the lottery, implying the lottery was successfully randomized. After the lottery, we find that winners invest in costly new alterations and renovations to their properties, and on average see their home values increase by 45% 15 years later. Condominium converters see their home values increase by 53% over the same horizon. Lottery winners shift towards renting their units, and hold their properties in the near-term before selling at higher rates in the long-run.

A central controversy in policy debates focuses on the extent to which new condominium supply affects nearby home and rental prices. To address this question, we turn towards estimating price spillovers on nearby properties. Our empirical design offers a convenient setting for estimating these effects, by comparing nearby properties of lottery winners to nearby properties of lottery losers. We augment our main empirical specification with controls for the expected number of nearby lottery winners at various distances. This procedure addresses endogeneity concerns stemming from the fact that being close to a lottery winner is, in part, a function of location; the location is, in turn, potentially correlated with unobservable characteristics or shocks determining home values (Borusyak and Hull, 2023). Following a condominium conversion, we find that home values for parcels nearby lottery winners increase by 11%.

The finding that nearby home values increase is suggestive that condominium conversions may play a larger role in neighborhood change. To further explore this possibility, we adapt our lottery design to exploit exogenous variation in neighborhood-level exposure to condominium conversions. Over a horizon spanning approximately 15 to 20 years, we find that an additional lottery winner increases home values, rents, the population of high-income residents, and shares of the population that are White and college-educated. Using data on new business formation, we also find that lotteries lead to increases in establishments specializing in education, real estate, and professional services.

Heterogeneity analyses suggest that the effects of condominium conversions on neighborhood-level home price appreciation are smaller in neighborhoods with initially higher poverty rates and Hispanic population shares, and larger in neighborhoods with initially higher shares of college graduates. These results are consistent with a wealth of qualitative research in sociology arguing that demographic characteristics such as income, race, and education are key mediators of gentrification in American cities (Zukin 1987, Lees et al. 2013, Freeman 2005). Overall, the results imply that supply-side housing policies play a significantly larger role in gentrification than has been previously documented in the existing literature.

Our study contributes to a growing body of research on the determinants of housing prices and gentrification in cities. Existing studies have emphasized that demand for low-income neighborhoods reflects changes in employment and amenity opportunities in the city core (Diamond, 2016; Almagro and Dominguez-Iino, 2021; Couture et al., 2019). At the same time, a budding and complementary literature has examined the role of supply-side drivers of local home price appreciation, like new construction (Asquith et al., 2019; Pennington, 2021). Policy debate further focuses on how the quality and price of housing responds to urban development policies such as zoning, rent control, and eviction protections (Autor et al., 2015; Diamond et al., 2019). Our paper contributes to frontier research that uses highly credible empirical variation, detailed data, and detailed analysis of counterfactuals to study the impacts of key urban policies on local home prices and demographics.

To our knowledge, Boustan et al. (2019) is the only other paper to consider the role of condominiums in urban change. Boustan et al. (2019) instrument for city-level variation in condominium density with regulatory changes governing conversions. They find no relationship between condominiums and resident income, education, or race. Their work leaves open the possibility that condominiums could affect the distribution of individuals and incomes within the city. In our study, we focus on the experience of one city — San Francisco — but provide credible estimates of the within-city housing and neighborhood effects of condominium conversions. Our work is closely related to Diamond et al. (2019) and Pennington (2021), both of which consider the setting of San Francisco, and study the local effects of rent-control and new construction, respectively. Our paper is also similar in spirit to Greenstone et al. (2010), who study the spillover effects of plausibly exogenous commercial plant openings on local labor markets. By contrast, we study the spillover effects of high-end residential buildings on local home prices and demographics.

The rest of the paper is organized as follows. Section 4.2 discusses the history of condominium conversions and our institutional setting. Section 4.3 discusses the data, and section 4.4 discusses how we implement the lottery design in a regression framework. Section 4.5 presents our estimates of the effect of winning the lottery on winning and nearby parcels. Section 4.6 studies the impacts of condominium conversions on neighborhood outcomes, finding a significant effect on demographic outcomes normally associated with gentrification. Section 4.7 concludes.

4.2 Background and Setting

Changes in Demographics and Housing Markets in San Francisco and Other Major U.S. Cities

San Francisco provides an ideal setting for an empirical study of gentrification and housing supply, for two reasons.

First, like many major American cities, San Francisco has undergone dramatic demographic changes since the 1990s (Couture and Handbury 2023). Panel A of Figure 1 plots time series for San Francisco, selected major American cities, and the U.S. national average in rental prices, median family incomes, the share of the population that is college-educated, and the share of the population that is Black or Hispanic. All values are indexed to 100 in 1990. Relative to the national average, San Francisco, New York, Washington DC, Los Angeles, and Boston have seen meteoric growth in rents, median family income, and the share of the population that is college-educated over the last three decades. Moreover, while the country as a whole became increasingly diverse over this period, the share of the population that is Black or Hispanic declined in San Francisco and other major U.S. cities.

Second, Panel B shows that San Francisco, like other major cities, has failed to increase the per capita supply of housing at a pace consistent with the national average. In a world of scarce housing, changes to the existing stock of units, like condominium conversions, could play an outsized role in shaping broader demographic trends.

The Rise of Condominiums

A condominium is a legal form of ownership for housing units, often thought of as apartments, within a multi-unit building. An individual owns the unit itself, while common spaces (such as elevators, hallways, stairwells, and yards), building infrastructure (such as heating and water pipes), and the land under the building are jointly owned by residents. In the U.S., laws governing this ownership structure were first passed in Puerto Rico in 1958. The Federal Housing Administration began to insure mortgages on condominiums as part of the National Housing Act of 1961 (Kerr, 1963). Over the next few years, states passed their own laws authorizing condominiums en masse. By 1969, all U.S. states had passed such statutes (Boustan et al., 2019).

Condominiums can be added to the local housing supply either through new construction or by converting existing units. In the 1970s, city lawmakers became concerned that conversions were drastically reducing supply in the rental housing market, leading to higher prices and displacing renters. In response, several cities passed ordinances to limit or prohibit this behavior (Boustan et al., 2019).¹ Public debates over condominium developments and

¹The timing of these laws serve as an instrumental variable for city condominium density in Boustan et al. 2019.

conversions have recently reemerged. For example, in response to growing public concern, New York state legislators passed a law in 2019 requiring approval from current tenants to convert a building into condominiums.² Following passage of the law, conversions in New York City declined sharply by 80%.³

San Francisco Lottery

By 1979, condominium conversions were commonplace in the Bay Area, doubling each year since 1975. For approximately 15% of Bay Area municipalities, conversions comprised 10% or more of the existing rental stock (Ichino, 1979). Following several high-profile conversions that lead to the eviction of long-term tenants, San Francisco city officials moved to regulate the practice.⁴ New regulations were designed to “prevent the displacement of existing tenants,” “reduce the impact of conversions on nonpurchasing tenants who may be required to relocate,” and “prevent the effective loss of the City’s low or moderate income housing stock” (San Francisco Board of Supervisors, 2004). Starting in 1981, the City prohibited conversions for buildings with more than six units. Buildings with fewer than six units could still convert, but owners were required to apply and win the right to do so through a lottery process⁵.

City officials in San Francisco limited the total number of units eligible to convert by lottery each year to 200 units. A lottery applicant is an entire building. The lottery consists of two separate applicant pools vying for 100 units of eligible conversions: Pool A and Pool B. Pool A contains applicants who applied to (and lost) three or more prior lotteries, and imposes some restrictions on tenant eviction history, ownership, and occupancy. Prior to 2006, a simple lottery was run among Pool A applicants if the the number of units in pool A exceeded 100 units. Any unallocated units were added to Pool B. From 2006 until 2013, applicants were grouped and ranked by the number of times they had previously lost the lottery. If the first group included fewer than 100 units, all lottery applicants were allowed to convert. Any remaining units were allocated to the second group, and so on, until the final group’s total number of units was larger than the remaining number of units available in Pool A. At that point, those units were randomly allocated among that group’s applicants. In Pool B, each applicant receives tickets equal in number to the times they have previously lost the lottery. Tickets were then drawn randomly until 100 units were deemed eligible for conversion.⁶ Prior to 2006, the number of tickets was limited to be at most five (San Francisco Board of Supervisors, 2005).

²Wall Street Journal. “New York Condo Conversions Near the End, a Casualty of Rent Reform.” 2019.

³The Real Deal. “Rental-to-condo Conversions Drop 80% After 2019 Rent Law: Report.” 2021.

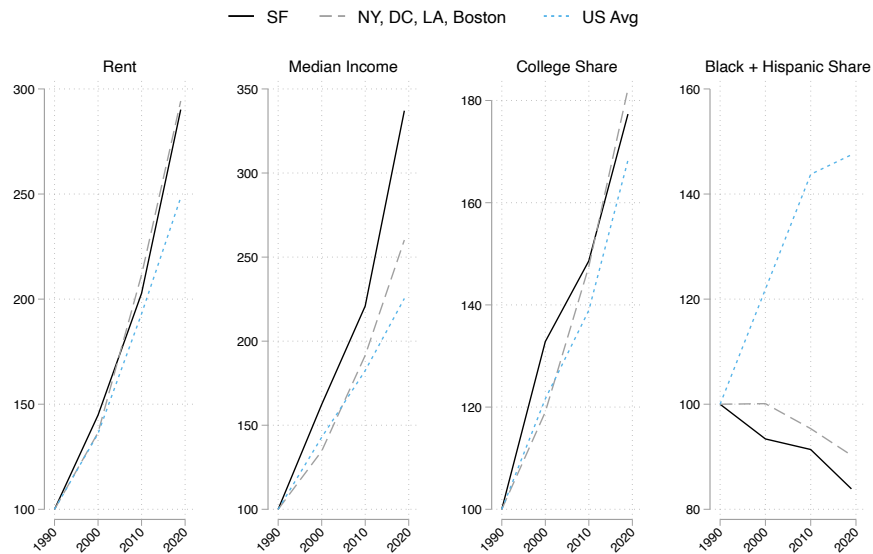
⁴Discussions over condominium conversions were regularly in the newspaper: “Condos Put Squeeze on Rentals” (San Francisco Chronicle, December 1977), “S.F. Problem That’s Hard to Live With” (San Francisco Chronicle, March 1979).

⁵Sirkin-Law. “Summary of San Francisco Condominium Conversion Rules.” 2022.


⁶Sirkin-Law. “San Francisco’s Condo Conversion Lottery System.” 2022.

FIGURE 1: GENTRIFICATION IN SAN FRANCISCO AND OTHER U.S. CITIES

PANEL A: RENTS AND DEMOGRAPHICS



PANEL B: HOUSING UNITS PER CAPITA



Notes: The unit of analysis is a census tract, and data are from the Neighborhood Change Database. In Panel A, data values are indexed to 100 in 1990, the earliest year consistently available for these outcomes. In Panel B, data values are indexed to 100 in 1970.

Lottery tickets were priced at \$250 each, but upon winning, the conversion application required additional fees. Inspection plus application fees, on average, totaled approximately \$13,000. For some 5-6 unit buildings, an additional charge of \$1,700 was levied by the state of California. A mandatory engineering survey of the building costed at least \$8,000.⁷

After a growing backlog of lottery applicants, San Francisco halted the lottery program in 2013. City officials replaced it with the Expedited Conversion Program (ECP) beginning in 2015. Under this new program, TICs satisfying certain ownership and occupancy requirements would be eligible to convert. Buildings that had been owned continuously for longer would be eligible first. A new Expedited Conversion Fee of \$22,500 per unit would also be charged.⁸ Buildings with renters are required to offer a lifetime lease upon conversion; due to legal challenges, the city stopped accepting conversion applications from buildings with renters in 2017.

Application Behavior

We now summarize descriptive patterns in lottery applications. Figure ?? describes how the probability of winning the lottery and reapplying to the lottery varies with the number of tickets received for the lottery. By design, the probability of winning the lottery increases with the number of tickets. Given high demand for the lottery, the probability of winning is still low (<30%) even after applying five times before. If a building had applied seven times previously, and consequently was awarded eight lottery tickets, the probability of winning was close to 90%. Reapplication rates increase in the number of tickets a building receives.

Figure ?? plots how the probability of applying and winning varies with whether a building had applied a certain number of years ago. The light blue coefficient above an x-axis value of one captures how much more likely an individual is to apply if they had applied one year ago. The dark blue coefficients show the same coefficients for whether an individual won the lottery. The plot shows that most individuals that apply also reapply. About 12% of applicants win in a given lottery year, but 82% of lottery applicants reapply the following year. This suggests that over 90% of lottery losers reapply. Applicants dynamically selecting into lotteries is thus a minor concern in our setting, given that reapplication rates are so high. Second, over 60% of applicants had won a lottery within seven years.

This latter point is confirmed in the Appendix, which reports the probability that a lottery loser in year $t=0$ wins in a lottery any number of years later. Many lottery losers reapply and become lottery winners soon after. Simple comparisons in outcomes between lottery winners and losers are likely to bias down the effect of a condominium conversion, since many losers ultimately convert. This fact motivates the regression design we discuss in Section 4.4.

⁷Cost estimates from: [Sirkin-Law](#). “San Francisco’s Condo Conversion Lottery System.” 2022.

⁸[GMH - Real Estate Law](#). “Condominium Conversion in San Francisco.” 2019.

4.3 Data

Our primary outcome of interest is home values, for which we use the assessed value of land and structures for a parcel as given in the San Francisco Assessor's Office annual files. We further merge information about parcel-level building permits and evictions, as well as information about lottery applicants and winners. For neighborhood outcomes, we rely on data from the ACS and business registrations.

Data Sources

Our main data come from four main sources, most available through data.sfgov.org.

Property Tax Rolls (1999 - 2019) : The San Francisco Office of the Assessor-Recorder makes the years 2007 through 2019 publicly available through their website. To this, we merge in years 1999 through 2006 provided to us directly by the assessor's office. This dataset contains information about the property location, type and construction type, number of bathrooms, bedrooms, rooms, stories, and units, local zoning, property area, whether the property is homeowner-occupied, most recent sale, and assessed value of land and improvements (structures).

Building Permits (1983 - 2019): This dataset contains information on the parcel number, date, estimated cost, and type of building permits.

Evictions (1997 - 2019): This dataset contains each eviction in the city of San Francisco, with its location, file date, and the reason for the eviction.

Lottery Information (2001 - 2013): Lottery information on the applicants, the number of tickets they were assigned, and the winners was provided as part of a public records request (#17-1329 accessible through <https://sanfrancisco.nextrequest.com/>). Importantly, the number of tickets allows us to infer whether an applicant was in Pool A or Pool B of the lottery.


Census & ACS Data (2000, 2013-2019): We collect ACS census block group-level data for years 2013 to 2019 for the city of San Francisco. This dataset contains information about demographics, income, home values, rents, and education. We also use tract-level outcomes from the 2000 census, concorded to 2010 census tracts in the Neighborhood Change Database.

Business Registrations (2000 - 2019): The Office of the Treasurer and Tax Collector contains information on business registrations, including their location, sector, date of registration, and whether the business is still active or not. We aggregate this data to the census block group level, and tabulate counts of establishments in different sectors. We tally new establishments as well as the stock of active ones.

Summary Statistics

Our lottery data begins in 2001 and continues until the lottery ended in 2013. Figure 2 presents time series for the number of applicants (right axis) and the number of winners (left axis) for each lottery in our sample. The number of winners remains flat at 60 to 65 per year. Each winner on average is a 3-4 unit building, leading to cumulative totals of approximately 200 units per year. Over the study period, the number of applicants nearly doubled. This dramatic increase ensured that winners were randomized even amongst pool A applicants for most years.

FIGURE 2: LOTTERY APPLICANTS AND WINNERS




Notes: The figure plots the number of lottery applicants (right axis) and winners (left axis) for each year of the sample. Data are from the City of San Francisco Assessor’s Office.

Figure 3 plots our estimate of an applicant’s probability of winning, based on rules stipulated in city ordinances, and the empirical probability that the applicants actually won. The plot shows that we are able to replicate the lottery’s randomization procedure, with the line of fit precisely on the 45 degree line. The years 2001, 2006, 2010, and 2012 saw a sizeable fraction of applicants guaranteed winning in pool A — that is, these applicants had

probability “1” of winning. These buildings accounted for 93 of the total 812 winners we observe. Our results on lottery winners, which fully match on the lottery propensity score, will effectively ignore variation from these applicants.

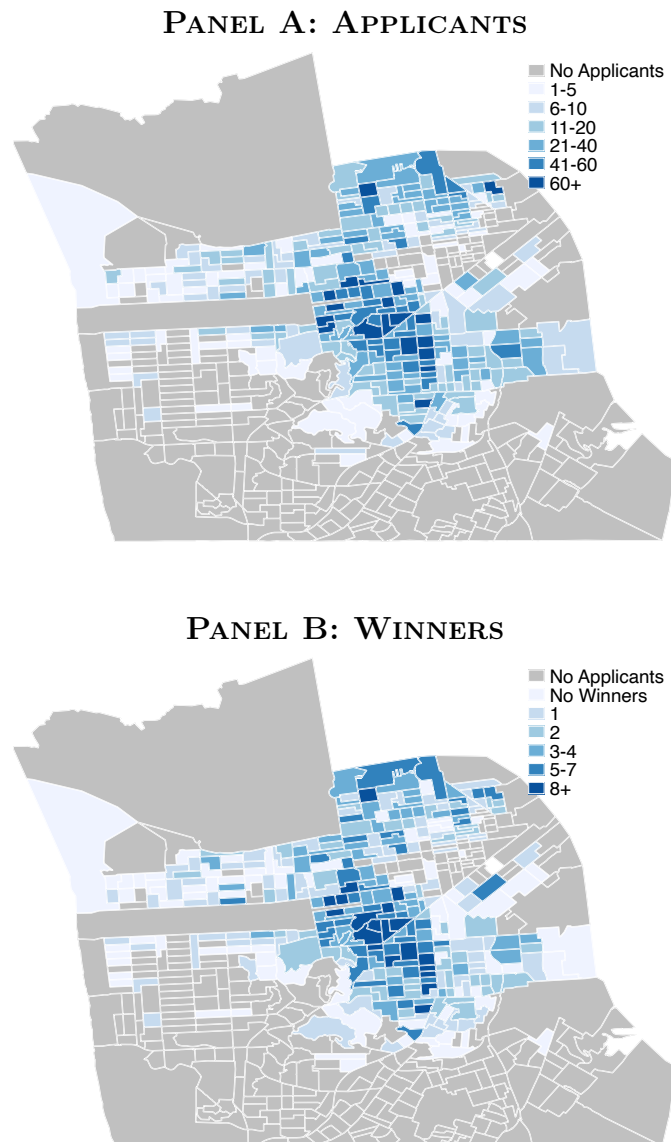
FIGURE 3: LOTTERY PROBABILITY OF WINNING



Notes: The unit of analysis is a lottery applicant, and marker sizes are proportional to the number of applicants. Data are from the City of San Francisco Assessor’s Office. The x-axis reports each applicant’s predicted probability of winning the lottery, based on the lottery rules and regulations described in section 4.2, and the y-axis reports the corresponding share of actual lottery winners. The dashed line shows the linear best-fit, which lies precisely on the 45-degree line.

Figure 4 maps the geographic distribution of lottery applicants and winners. Neighborhoods like North Panhandle, Haight Ashbury, Duboce Triangle, and the Mission saw high demand for conversions as well as many lottery winners. Russian Hill had many applicants but few winners. Inner Sunset had few applicants but a surprising number of winners. This random geographic variation will be central to estimating the neighborhood effects of condominium conversion in Section 4.6. Our main set of findings will rely on locations with or near to lottery applicants. Consequently, neighborhoods like Outer Parkside and South of Market will be largely excluded from the analysis.

FIGURE 4: MAP OF LOTTERY APPLICANTS AND WINNERS



Notes: The figure illustrates geographic variation in the cumulative number of lottery applicants (Panel A) and lottery winners (Panel B) across Census block groups over the sample period.

Summary statistics for winning and losing applications are presented in Table 1. The calculations are for two years prior to an application. Applicant buildings have around 15 rooms and 1300 sq. ft. per unit. The average assessed home value was nearly 1 million dollars for winners. Column (5) calculates the difference between building characteristics

for winners and losers after controlling for lottery p-score and year fixed effects, and column (7) has the corresponding p-value. Reassuringly, no building characteristics are significantly different between lottery winners and losers prior to their application.

TABLE 1: BALANCE TABLE

Parcel Characteristic	Summary Statistics						
	Winners		Losers		Difference		
	mean	s.d.	mean	s.d.	diff	s.e.	pval
Value (1000s USD)	999	696	1,231	1,191	21	35	0.55
Year Built	1915	20	1915	20	-0	1	0.60
Homeowner	0.67	0.47	0.58	0.49	0.00	0.02	0.95
Evictions	0.03	0.26	0.02	0.24	0.01	0.01	0.25
Units	3.23	1.09	3.24	1.21	0.02	0.06	0.68
Rooms	14.99	4.52	14.82	4.60	-0.00	0.23	0.98
Sq. Ft. Per Unit	1,321	473	1,284	475	17	25	0.48
Beds	1.21	2.69	1.16	2.72	0.22	0.14	0.13
Baths	3.60	1.35	3.58	1.39	0.07	0.07	0.35
Permits	0.47	1.16	0.63	1.58	-0.06	0.08	0.43
Permit Costs (1000s USD)	5.27	26.53	8.25	45.30	-1.43	2.06	0.49
N	812		6023		6835		

Notes: The table reports the means and standard deviations of building characteristics for winning and losing applicants two years prior to the lottery. The difference in means reported in Column (5) controls for the probability of winning the lottery and year fixed effects. Sample includes lotteries from 2001 to 2013.

4.4 Design

Our setting provides a unique randomized experiment for studying the effects of condominium conversions on building home values, investment, sales, and renting behavior. Framing this design within an econometric framework provides some complications however, driven largely by repeated applications of losing properties and multiple event times. We now discuss our data structure and framework for estimating the relevant treatment effects.

Simple Lottery Design

We begin by discussing a regression implementation of the simple lottery design before generalizing to our setting. Let y_{it} denote assessed home value, one of our key outcomes, for parcel i in year t . Let $\tau_t(k)$ indicate whether year t is k years from the lottery and ν_i indicate whether parcel i wins the lottery.

Applicants have unequal probabilities of winning the lottery according to how many tickets they purchase and the ticket composition of other applicants. The number of tickets that can be purchased depends on how many previous times an applicant has lost. Comparisons between lottery winners and losers may give misleading estimates of the effect of condominium conversions if, for example, applicants with the greatest expected housing price appreciation apply more frequently. Controlling for the probability of winning the lottery ensures that we rely on random variation generated by the lottery, rather than endogenous selection into the lottery. As such, we include $\chi_i(p)$ fixed effects, which indicate whether parcel i had probability p (from support set \mathcal{P}) of winning the lottery.

We model home values using the following equation:

$$y_{it} = \sum_{k \neq -1} \beta_k^{ITT} \tau_t(k) \nu_i + \sum_k \sum_{p \in \mathcal{P}} \gamma_{pk} \tau_t(k) \chi_i(p) + \varepsilon_{it} \quad (4.1)$$

The second summation term ensures that comparisons are made between parcels with the same probability of winning the lottery.⁹ Following [Abdulkadiroğlu et al. \(2017\)](#), this type of full propensity-score matching ensures that the coefficients β_k^{ITT} capture a convex-weighted average of the causal effects of winning in one's specific lottery-strata. The coefficients β_k^{ITT} map out the full set of intention-to-treat (ITT) effects — that is, the effect of winning on home values k years before or after the lottery. Years prior to lottery implementation serve as balance tests, allowing to evaluate whether, within a lottery-strata, outcomes are similar prior to the lottery. Throughout, we cluster standard errors at the level of lottery applicant.

The ITT effects provide a transparent assessment of the lottery design. However, we are primarily interested in the effect on home values from parcel i converting to a condominium at time t . We therefore augment the above design to an instrumental-variables setting, where we instrument for whether a parcel ever converts to a condominium after the lottery (given by κ_i) with whether a building won the conversion lottery.¹⁰ The second-stage equation that relates home values to condominium status is given by the following equation:

$$y_{it} = \sum_{k \geq 0} \beta_k^{LATE} \tau_t(k) \kappa_i + \sum_k \sum_{p \in \mathcal{P}} \gamma_{pk} \tau_t(k) \chi_i(p) + \varepsilon_{it} \quad (4.2)$$

⁹The probability of winning a lottery is the same for almost all applicants who purchase the same number of tickets in a given lottery; consequently, specifications with lottery by ticket fixed effects produce nearly identical regression results.

¹⁰We instrument for this variable, rather than if a parcel is currently a condominium, so that we do not need keep track of two sets of event times: one for lottery application and one for conversion.

We instrument for condominium conversion with winning the lottery in the following first-stage regressions:

$$\tau_t(k)\kappa_i = \sum_{k \geq 0} \beta_k^{FS} \tau_t(k) \nu_i + \sum_k \sum_{p \in \mathcal{P}} \gamma'_{pk} \tau_t(k) \chi_i(p) + \varepsilon'_{it} \quad (4.3)$$

We use this instrumental variables (IV) approach to assess how the magnitude and significance of the effects change when treatment is defined as a condominium conversion, rather than winning the lottery. Thus, we estimate and report β_k^{LATE} for years after the lottery is implemented.

Intuitively, this empirical strategy is an instrumented difference-in-differences design comparing the outcomes of applicant lottery winners and losers with the same ex-ante probability of winning. The lottery design ensures independence of the instrument, but in this setting, only mean-independence of the potential outcomes and treatment assignment with respect to the lottery is required (Hudson et al., 2017). Testing for parallel-trends in the fully-interacted ITT specification offers a diagnostic to assess this assumption.

The exclusion restriction requires that winning or losing the lottery does not directly affect home values — that is, it requires that the lottery can only affect home values through the condominium conversion itself. This is a natural assumption since the lottery’s sole purpose is to allow or prohibit condominium conversions. Monotonicity is guaranteed in this setting, as is one-sided non-compliance — regulation for these buildings prohibited condominium conversions except through the lottery while it was active. In Section 4.5, we show that the first-stage is strong, with more than 80% of winning properties converting. Consequently, the coefficient β_k^{LATE} can be interpreted as the causal effect of converting to a condominium on home values at time k after the lottery (Imbens and Angrist, 1994).

Dynamic Lottery Design

Parcels that lost the lottery but had the same probability of winning are a natural control group for lottery winners. We now extend the above framework to a setting with repeated lotteries and where losing parcels continue to apply. We also rely on recent research on difference-in-differences designs with heterogeneity in treatment timing in order to implement our econometric analysis.

We first create thirteen (one for each lottery) simple lottery designs, composed of each applicant for each lottery year from 2001 to 2013. We then stack these observations according to each lottery m . The $\tau_{mt}(k)$ denotes whether year t for lottery m is k years since the lottery was run. The ν_{im} is an indicator for whether the parcel won that lottery. The $\chi_{im}(p)$ are indicators that the parcel in a given lottery had probability p of winning. Our new outcome y_{imt} denotes assessed home value for parcel i in lottery m in year t .

Parcels have histories duplicated according to how many times they have applied. Consistent with Cengiz et al. (2019) and Baker et al. (2021), we construct a clean set of controls by only including observations for control parcels that are yet to convert to a

condominium. This ensures that we do not compare outcomes of previous winners and later winners, which would naturally bias down our results. If we were interested in the effect of winning the lottery, it would be reasonable to include observations from later winners in the control group, as those are downstream effects from losing the lottery. A dynamic approach like [Cellini et al. \(2010\)](#) could then be used to estimate the desired treatment effects. Given that our object of interest is the effect of condominium conversion rather than winning the lottery, our approach is more natural.

The ITT version of our main specification is as follows.

$$y_{imt} = \sum_{k \neq -1} \beta_k^{ITT} \tau_{mt}(k) \nu_{im} + \sum_k \sum_m \sum_{p \in \mathcal{P}} \gamma_{pmk} \tau_{mt}(k) \chi_{im}(p) + x'_{imt} \zeta + \alpha_{im} + \eta_t + \varepsilon_{imt} \quad (4.4)$$

While not necessary for a causal interpretation of the coefficients β_k^{ITT} , we include parcel by lottery fixed effects α_{im} in our main specification, for two reasons. First, the fixed effects reduce residual variation in the errors, which increases precision. Second, we drop observations for lottery losers that ultimately convert to a condominium. This occurs if lottery losers convert through the Expedited Conversion Program (implemented in 2013 and discussed in [Section 4.2](#)) or win in a subsequent lottery. The parcel fixed effects help address (i) possible selection bias among losing parcels that later convert, and (ii) an unbalanced sample stemming from different event-time coverage for each lottery. The fact that most losing applicants reapply, as shown in [Section 4.2](#), lessens the first concern. Nevertheless, we consider several robustness exercises to explore these issues in [Section 4.5](#). With parcel-lottery fixed effects, the coefficients β_k^{ITT} can be interpreted as a convex-weighted average of the underlying treatment effects for each lottery ([Sun and Abraham, 2020](#)).

While we do not include them in our main specification, we allow for additional controls x_{it} , like neighborhood trends. These controls may adjust for random imbalances between lottery winners and losers, and allow us to assess the importance of sample attrition in the control group later in event time. We consider how robust our results are to their inclusion in [Section 1.4](#). Additionally, an attractive feature of the dynamic setting is that it allows us to disentangle event-time effects from calendar-time effects, given by η_t , and to control for them separately.

This design corresponds to the “stacked” difference-in-differences design of [Cengiz et al. \(2019\)](#) and [Baker et al. \(2021\)](#). While other approaches to multiple event timings have been suggested ([Borusyak et al., 2021](#); [Callaway and SantAnna, 2020](#)), the stacked difference-in-differences design offers greater transparency and most naturally accomodates our IV and spillovers analyses.

The LATE implementation estimates the following second-stage regression.

$$y_{imt} = \sum_{k \geq 0} \beta_k^{LATE} \tau_{mt}(k) \kappa_{im} + \sum_k \sum_m \sum_{p \in \mathcal{P}} \gamma_{pmk} \tau_{mt}(k) \chi_{im}(p) + x'_{imt} \zeta + \alpha_{im} + \eta_t + \varepsilon_{it} \quad (4.5)$$

We instrument for winning the lottery and eventually converting to a condominium κ_{im} with winning the lottery through the following first-stage.

$$\tau_{mt}(k)\kappa_{im} = \sum_{k \geq 0} \beta_k^{FS} \tau_{mt}(k) v_{im} + \sum_k \sum_m \sum_{p \in \mathcal{P}} \gamma'_{pmk} \tau_{mt}(k) \chi_{im}(p) + x'_{imt} \zeta' + \alpha'_{im} + \eta'_t + \varepsilon'_{it} \quad (4.6)$$

As in the ITT specification, attrition in our control group due to ECP and later lottery winners might raise concerns over the independence of the instrument. However, the IV difference-in-differences relaxes the necessary assumptions to maintain a causal interpretation. We report coefficients β_k^{LATE} for event times after the lottery was conducted. We cluster errors at the applicant level. This is particularly important in this context since applicant histories appear multiple times according to how many times they have entered lotteries.

Spillovers Design

Evaluating the effects of condominium conversions on nearby home values is complicated by two features of our setting. First, treatment will depend on the number of winners at various distances. Second, while winning the lottery may be random, being located near a winner is likely not. For example, parcels in the city center are more likely to be close to winners than parcels on the periphery, and being closer to the city center is likely correlated with unobservable shocks that determine home values. We extend our approach in the previous section to account for these facts.

We consider all residential parcels j within a fixed distance of any lottery applicant. For lottery m , we calculate the number of winning applicants within distance band d given by $\tilde{\nu}_{jm}(d)$. We also sum the probabilities for each distance band that nearby lottery applicants win and stack them into the vector $\tilde{\chi}_{jm}$. This variable approximates the expected number of lottery winners at different distances from a given parcel. The ITT version of our main spillovers specification is as follows.

$$y_{jmt} = \sum_d \sum_{k \neq -1} \beta_{dk}^{ITT} \tau_{mt}(k) 1(\tilde{\nu}_{jm}(d) > 0) + \sum_d \sum_k \sum_m \tau_{mt}(k) f(\tilde{\chi}_{jm}, \gamma_{dmk}) + x'_{jmt} \zeta + \alpha_{jm} + \eta_t + \varepsilon_{jmt} \quad (4.7)$$

In our main specification, f captures all linear terms and first order interactions of the coordinates of $\tilde{\chi}_{jm}$. The vector γ_{dmk} contains the coefficients on the terms in the function f . These parametric controls adjust for the fact that tracts near lottery applicants and winners are unlikely to be comparable with tracts that were not near lottery applicants and winners. This is the same insight as in [Borusyak and Hull \(2023\)](#), ensuring that we still rely on lottery variation to estimate the spillover effects while controlling for endogeneity due to a parcel's location. We focus on having any lottery winner a certain distance away as the treatment.¹¹ The parameters of interest β_{dk}^{ITT} map the full set of spillover dynamics for each

¹¹The vast majority of parcels are at most near one winner.

distance band d . Parcels are included as controls until their nearby applicants convert to a condominium, if ever. We map parcels to their closest applicant, and cluster errors at that location.

The IV model is estimated in the same way as before. We instrument whether the nearby winner ever converts to a condominium after the lottery with whether the nearby parcel wins the lottery. We estimate these coefficients for all event times greater than zero.

4.5 Results

In this section, we leverage the lottery design to estimate the causal impact of condominium conversions on winning and nearby properties. We find that winning property owners see large increases in assessed home values. Owners renovate their properties, rent them out, and eventually sell them 7 years on from the lottery. We also find large and highly localized price spillovers on nearby properties. Parcels within 25 meters of a winner see their home values increase 11% after 15 years, with this effect becoming insignificant farther away.

First-Stage: Effects of Winning the Lottery on Condominium Conversions


We first document that lottery winners overwhelmingly convert their properties into condominiums. Figure 5 plots the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5, respectively, along with their associated 95% confidence intervals. In both specifications, the outcome is an indicator equal to one if the property is legally registered as a condominium, and zero otherwise.

The β_k^{ITT} coefficients in Figure 5 trace the dynamic treatment effects of winning the lottery over time. In years prior to the lottery, winning and losing applicants are equally (un)likely to convert to condominiums. This is unsurprising, since conversions are legally prohibited unless property owners win the lottery. After the lottery, winning property owners are generally unable to immediately convert their properties into condominiums, since the process for preparing and approving applications is costly and takes time. However, over time, the share of conversions steadily increases. Approximately 20% of winners convert their properties within the first full year, and 60% convert within 3 years. At longer time horizons, the conversion rate surpasses 80% within 10 years, and stabilizes at around 83% within 15 years, which is the end of our sample horizon. That the vast majority of lottery winners eventually convert is again unsurprising, since the sole purpose of the applying to the lottery is to obtain legal permission to do so.¹²In the Online Appendix, we show that the

¹²For a negligible share of observations (less than 2%), we observe properties classified as condominiums prior to implementation of the lottery. One possible explanation for this fact is that, several years before the lottery, property owners may have converted in the opposite direction (that is, they changed from condominiums to another legal form), and then applied to the lottery in order to change back. Another perhaps more likely possibility is measurement error in the administrative data.

propensity for lottery winners to convert their properties is statistically similar irrespective of initial neighborhood characteristics.

FIGURE 5: FIRST-STAGE: EFFECT OF WINNING LOTTERY ON CONDOMINIUM CONVERSIONS



Notes: The unit of analysis is a property-year, and the sample includes properties whose owners apply to the lottery. Data are from the City of San Francisco Assessor’s Office. The outcome is an indicator for converting to a condominium. The figure reports the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5, respectively. These specifications compare trends in conversions of lottery winners versus losers. In the IV model, the endogenous variable is a time-invariant indicator for properties that ever convert to a condominium, interacted with years to/from the lottery. Standard errors are clustered by lottery applicant, and error bands show 95% confidence intervals.

In Figure 5, the endogenous variable in the IV model is a time-invariant indicator equal to one if the property ever converts to a condominium in our sample period. In the figure, the β_k^{IV} coefficients are informative of the share of lottery-induced compliers who have converted within k years from the lottery. For example, three years after the lottery, the $\beta_{k=3}^{IV}$ coefficient of $0.74 = (0.61/0.83)$ implies that 74% of lottery winners that will ever convert have already done so. By construction, this share converges to 1 by the end of our sample period.


Overall, the high condominium conversion rates in Figure 5 provide compelling evidence of an economically and statistically strong first-stage in our instrumental variables design.

Effects on Winning Properties


We now study the effects of the lottery on winning properties. The panels in Figure 6 plot the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5 for a suite of key outcomes: property values and building and renovation permits. Across all the outcomes, the panels show that winning and losing properties were on common trends prior to lottery, consistent with effective random assignment from the lottery and with the balance tests presented in Table 1.

FIGURE 6: EFFECT OF WINNING THE LOTTERY ON PROPERTY VALUES AND PERMITS


(A) PANEL A: LOG PROPERTY VALUE




(B) PANEL B: TOTAL BUILDING PERMITS



(C) PANEL C: IHS VALUE OF RENOVATIONS



(D) PANEL D: IHS CUMULATIVE VALUE OF RENOVATIONS



Notes: The unit of analysis is a property-year, and the sample includes properties whose owners apply to the lottery. Data are from the City of San Francisco Assessor’s Office. The figure reports the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5, respectively. These specifications compare trends in outcomes of lottery winners versus losers. In the IV model, the endogenous variable is a time-invariant indicator for properties that ever convert to a condominium, interacted with years to/from the lottery. Standard errors are clustered by lottery applicant, and error bands show 95% confidence intervals.

Panel A of Figure 6 shows that, over time, winning properties increase dramatically in value relative to losing properties. Within 15 years, the IV estimates indicate that condominium conversions on average cause property values to appreciate by 53%. Since the average property in the sample is worth approximately \$1 million, this implies that condominium conversion was on average worth more than \$500,000 during our sample period.

Table 2 shows that this result is robust to alternate specifications that control for local time trends by neighborhood, census tract, or block-group.

TABLE 2: EFFECT OF WINNING THE LOTTERY ON PROPERTY VALUES

	(1)	(2)	(3)	(4)
	$\Delta_{t=-1}^{15}$ Log Value	$\Delta_{t=-1}^{15}$ Log Value	$\Delta_{t=-1}^{15}$ Log Value	$\Delta_{t=-1}^{15}$ Log Value
Condo Conversion	0.534*** (0.0684)	0.536*** (0.0692)	0.503*** (0.0757)	0.517*** (0.0813)
Observations	116,086	116,016	115,922	115,547
Model	IV	IV	IV	IV
Geography x Year FE	None	Nbhd	Tract	Block Group

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: The unit of analysis is a property-year, and the sample includes properties whose owners apply to the lottery. Data are from the City of San Francisco Assessor’s Office, and the outcome is log property value. The table reports the $\beta_{k=15}$ coefficient from equation 4.5, comparing the values of properties that win the lottery versus those that lose. Column 1 reports the benchmark specification. Columns 2-4 include controls for local trends by neighborhood (Column 2), Census tract (Column 3), and Census block group (Column 4). Standard errors are clustered by lottery applicant.

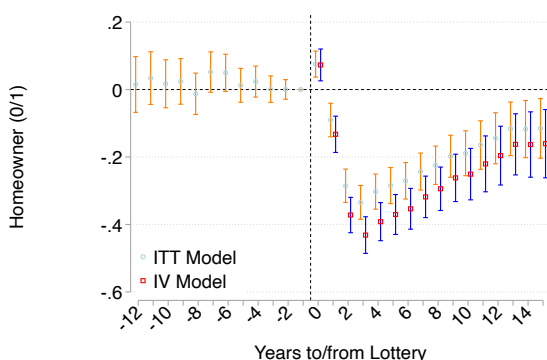
Panels B, C, and D of Figure 6 document the effects of winning the lottery on permits permits for alterations and renovations (Panel B), the inverse hyperbolic sine transformation of the estimated cost of these renovations (Panel C), and their cumulative value, summed from 1999 forwards (Panel D). The number of permits and their value increase sharply in the two years immediately after the lottery, as winning owners make new investments and improvements in their properties. These improvements likely explain part of the increase in property values documented in Panel A. Permit effects are entirely concentrated within the first two years, after which they decline to levels below the level for lottery losers. Panel D demonstrates that the overall value of alterations and renovations for winning owners remain 29% above losing owners through the study period.

The panels in Figure 8 plot the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5 for a second set of outcomes: tenant evictions, homeownership, and property sales. Panel A shows that winning properties are more likely to be occupied by a homeowner in the year following the lottery. This implies that the winning property owners are more likely to live in their units when they are making renovations and alterations. However, these owners then quickly move out of the units after the first full year, and instead rent the properties to new tenants. The owners are likely to benefit from higher rental prices, both because


condominiums are not subject to rent control and because renters are willing to pay more for the recently renovated (and presumably higher quality) units. Three years after the lottery, the IV estimates imply that lottery-induced converted units are 43percentage points more likely to be occupied by renters compared to losing lottery units. The magnitude of this effect steadily attenuates as units are sold to new homeowners over time, such that winning units are 16percentage points more likely to be occupied by renters 15 years after the lottery.

FIGURE 8: EFFECT OF WINNING THE LOTTERY ON HOMEOWNERSHIP, SALES, AND EVICTIONS


(A) PANEL A: HOMEOWNER-OCCUPIED (0/1)



(B) PANEL B: PROPERTY SALES (0/1)



(C) PANEL C: EVICTIONS



Notes: The unit of analysis is a property-year, and the sample includes properties whose owners apply to the lottery. Data are from the City of San Francisco Assessor’s Office. The figure reports the β_k^{ITT} and β_k^{IV} coefficients from equations 4.4 and 4.5, respectively. These specifications compare trends in outcomes of lottery winners versus losers. In the IV model, the endogenous variable is a time-invariant indicator for properties that ever convert to a condominium, interacted with years to/from the lottery. Standard errors are clustered by lottery applicant, and error bands show 95% confidence intervals.

Panel B of Figure 6 traces the effects of winning and converting on property sales, defined as an indicator equal to one if any unit in the parcel is sold.¹³ Winning property owners are modestly less likely than losing applicants to sell their units in the years immediately

¹³The comparison of lottery winners vs. losers is apples-to-apples because units of condominiums can be sold (winners) as can units of tenancy-in-commons (losers; see section 4.2 for more details).

following the lottery — this result is consistent with the finding that owners are more likely to be renovating and renting their properties during these years. However, 7 years after the lottery, winning applicants that convert are on average 4percentage points more likely to sell their properties than losing applicants. Within 15 years, this effect increases modestly to 7percentage points.

Lastly, contrary to the concerns of many policymakers and voters, Panel C reveals that condominium conversions do not cause a statistically discernable change in eviction rates. However, we caution that this result does not necessarily imply there is no turnover or displacement of tenants, since in some cases owners may induce tenants to move out without resorting to the legal eviction process.Heterogeneity


In the Appendix, we explore whether the effects of condominium conversions on property values vary with neighborhood or property characteristics. The Online Appendix shows that average property appreciation of winning properties does not vary systematically with initial neighborhood characteristics. The Online Appendix also shows that appreciation is modestly larger for older buildings with intially lower values, consistent with significant value-added from alterations and renovations.

Spillover Effects on Nearby Properties


We now use the econometric design described in Section 4.4 to evaluate the spillover effects of condominium conversions on nearby properties. Panel A of Figure 10 plots the β_k^{ITT} and β_k^{IV} coefficients from estimating equation 4.7, where the outcome is the assessed value of homes within a 25-meter radius from the winning property. The 25-meter bandwidth typically includes 3 to 6 properties that are either immediately adjacent to the winning property, across or behind the street, or two to three doors down. Before the lottery, price trends are similar for homes nearby winning and losing applicants, as expected from random treatment assignment. After the lottery, homes located nearby the winners do not immediately increase in value, but do increase beginning approximately 8 years later. This timing is consistent with the treatment dynamics we observed for property sales in Figure 6. Within 15 years of the lottery, these nearby home prices appreciate by 11% (s.e.=4.8%) relative to homes nearby losing lottery applicants. The Online Appendix shows that this result is qualitatively similar when including controls for local time trends by neighborhood, census tract, or census block group.

FIGURE 10: SPILLOVERS ON NEARBY PROPERTY VALUES

(A) PANEL A: LOG VALUE OF PROPERTIES WITHIN 25 METERS



(B) PANEL B: LOG VALUE OF PROPERTIES WITHIN 25-50 METERS



Notes: The unit of analysis is a property-year, and the sample includes "nearby" properties within 25 meters (Panel A) or 25-50 meters (Panel B) of a lottery winner. Data are from the City of San Francisco Assessor's Office. The figure reports the β_k^{ITT} and β_k^{IV} coefficients from equation 4.7. The specifications compare trends in log home prices of properties located near lottery winners versus properties located near lottery losers. The regressions control for predicted probabilities of treatment, as in [Borusyak and Hull \(2023\)](#); see section 4.4 for details. In the IV model, the endogenous variable is a time-invariant indicator for properties that ever convert to a condominium, interacted with years to/from the lottery. Standard errors are clustered by lottery applicant, and error bands show 95% confidence intervals.

Panel B of Figure 10 shows results from estimating the same equation but using a different bandwidth: the outcome in Panel B is the assessed value of homes within a 25- to 50-meter bandwidth from the winning property. These parcels are typically on the same block as the winning property, but further down the street. In this bandwidth, we do not find compelling evidence of spillovers after the lottery: the estimated treatment effects are small and positive but imprecisely estimated ($\beta_{k=15}^{IV}=5\%$, s.e. = 5%).

The evidence from Figure 10 thus suggests that the spillover effects of condominium conversions are highly localized. Properties within 25 meters of a condominium conversion appreciate in value, while those farther away see weaker or no appreciation. This empirical pattern is informative of the mechanisms that may be driving the spillovers. First, nearby properties might appreciate due to an aesthetic externality from the improved and renovated condominiums, which fades at further distances. Second, nearby properties might appreciate due to a form of behavioral benchmarking, whereby prospective home buyers face imperfect information in the real estate market and so use information about highly local properties a signal of underlying value. Third, to the extent that local demographic characteristics affect home values, changes in the composition of residents living in the converted condominiums may also play a role. We explore these issues in greater detail in the following section.

4.6 Neighborhood Effects

We now turn to addressing whether conversions lead to changes in demographic or economic outcomes in affected neighborhoods. In the previous section we documented that conversions have a large effect on the prices of winning properties, and induce substantial turnover in resident composition as owners move in, move out, rent, and eventually sell their properties. We also documented that conversions increase the prices of adjacent buildings. However, directly affected properties comprise only a small share of the neighborhood housing stock – approximately 2.4% of local residential units, on average. Thus, if conversions have broad-based effect on neighborhood-wide home prices and demographics, they must have spillover effects on the surrounding area.

Condominium conversions could play a role in broader neighborhood-level gentrification through a series of self-reinforcing mechanisms. First, the new, higher-income condominium residents may increase demand for local goods and services, putting upward pressure on local prices and changing the composition of local businesses. Second, the changing quality and prices of local goods and services — as well as the changing composition of the residents themselves — may in turn affect which residents find the neighborhood attractive. For example, low-income residents may not value the new neighborhood amenities given the prices, whereas high-income residents may find the neighborhood increasingly desirable. Demographic homophily, racism, and network effects may also play an important role in residential sorting patterns, as has been documented extensively in existing research.¹⁴ As

¹⁴For example, immigrants are more likely to move to neighborhoods with other immigrants that share their ethnicity (Abramitzky and Boustan 2017). In general, to the extent that social and information

more high-income residents enter the neighborhood, the pattern reinforces itself, fueling broader neighborhood change.

Below we explore to what extent condominium conversions in San Francisco caused neighborhood-wide gentrification, and evaluate whether the evidence is consistent with these mechanisms.

Neighborhood Design

To study neighborhood outcomes, we first adapt our lottery design to a setting in which a cross-section of census block groups are impacted by the cumulative number of lottery winners from 2001 to 2013. We focus on the long-run, aggregate effect of condominium conversions on census block group outcomes. In addition to being a natural time horizon for studying a slow-moving process such as gentrification, census block group data is largely only available after 2013. Consequently, we focus on neighborhood outcomes from 2013 to 2019 and combine information across all lotteries. Census block groups are the smallest geographic unit for which the Census and the American Community Survey (ACS) release information.¹⁵ Being a small, contiguous set of city blocks, they are an intuitive definition of a neighborhood.

To estimate the effect of a continuous treatment (conversions) on neighborhood outcomes, we rely on generalized propensity score methods (Imbens, 2000; Hirano and Imbens, 2004). In particular, let y_{gt} be an outcome for census block group g in year t . The variable ν_g denotes the number of lottery winners from all of the thirteen lotteries in our sample and $p_g(\nu_g)$ denotes the probability that the census block group had ν_g winners, given the lottery design.¹⁶ The $p_g(\nu_g)$ captures a neighborhood's demand for the lottery, and by extension, demand for condominium conversions. Once adequately controlled for, we can leverage random variation in condominium conversions through the number of actual winners in the lottery.

We follow the Hirano and Imbens (2004) multi-step procedure in two parts. First, we model neighborhood outcomes as arising from a quadratic in the actual number of conversion lottery winners and the probability of having that number of winners, as follows:

$$y_{gt} = \gamma_0\nu_g + \gamma_1\nu_g^2 + \gamma_2p_g(\nu_g) + \gamma_3p_g(\nu_g)^2 + \gamma_4p_g(\nu_g)\nu_g + x'_g\zeta + \varepsilon_{gt} \quad (4.8)$$

networks are segregated by education, class, and/or race, this is also likely to affect sorting (Jackson 2021; DiMaggio and Garip 2012). For a treatment of the history of racism and segregation in housing markets, see Rothstein 2017.

¹⁵While their size varies, they correspond to a median of nine city blocks (and most between six and ten blocks) in San Francisco. In our sample, census block groups have a median of 1400 individuals and 600 residential units in 2019.

¹⁶These probabilities are calculated as follows. For each lottery, we permute winners according to the lottery probabilities. We then repeat this process across a large number of simulations. We then calculate $p_g(\nu)$ as the fraction of the simulations in which census block group g received ν winners.

The x_g denote additional controls. In our main specification, we include all relevant outcomes in the year 2000 as controls i.e. $x_g = \mathbf{y}_{g,2000}$. For demographic outcomes available from the ACS, this information is only available at the tract level.¹⁷ For outcomes of establishment counts le from business permits and registrations data, this information is available at the census block group level. These controls serve dual purposes. First, our estimate of the generalized propensity score does not fully capture dynamic behavior among lottery applicants. We have argued elsewhere that this is not a substantial concern in our context, and including the baseline controls further mitigates any potential for imbalance. Second, the controls improve the precision of our estimates.

In the second step, we use equation 4.8 to map out the entire dose-response function $\hat{\mathbb{E}}[y_{gt}|\nu, p_g(\nu)]$. The dose-response function can be used to estimate the marginal change in neighborhood outcomes from one additional lottery winner at every level of the treatment and for every block group. These differences can be averaged to get an overall effect as follows:

$$\hat{\beta} = \frac{1}{G} \sum_g \sum_{\nu \neq 0} p_g(\nu) \cdot (\hat{\mathbb{E}}[y_{gt}|\nu, p_g(\nu)] - \hat{\mathbb{E}}[y_{gt}|\nu - 1, p_g(\nu - 1)]) \quad (4.9)$$

Standard errors are calculated by bootstrapping the entire procedure, clustering on neighborhoods. We run this regression only for census block groups that had at least one lottery application over the study period. There are six census block groups (out of 326 that had at least one lottery applicant) with 10 or more lottery winners. The linear specification in lottery winners is sensitive to their inclusion; the quadratic specification in the main specification above is far more stable.

Results

Table 3 presents our main results for how the number of lottery winners impacts demographic change in neighborhoods. All outcomes are scaled so that the coefficients can be interpreted as a percentage point (pp) change.

¹⁷Census tract outcomes in the year 2000 using 2010 boundaries come from the Neighborhood Change Database. We have all outcomes for census tracts in 2000 with the exception of the lower and upper quartiles of rent and home values. We simply control for the median home value and the median rent in 2000. The percentage of the population that is hispanic also does not appear in our census tract data, so we calculate it as 100 minus the percent of the population that is White, Asian, or Black. For consistency, we also use this definition in our census block group data.

TABLE 3: EFFECTS OF LOTTERY WINNERS ON NEIGHBORHOOD DEMOGRAPHICS

<i>Panel A: Population and Income</i>				
	100 × log Pop.	Poverty (pp)	College Deg. (pp)	100 × log MHI
Lottery Winners	0.28 (2.68)	-0.30 (0.31)	0.87* (0.49)	2.47* (1.46)
<i>Panel B: Demographics</i>				
	White (pp)	Asian (pp)	Black (pp)	Hispanic (pp)
Lottery Winners	1.57*** (0.51)	-1.46*** (0.44)	-0.09 (0.25)	-0.03 (0.31)
<i>Panel C: Home Values</i>				
	100 × log(Home Val. Q25)	100 × log(Home Val. Q50)	100 × log(Home Val. Q75)	
Lottery Winners	2.34* (1.33)	2.22** (1.01)	2.19*** (0.81)	
<i>Panel D: Rent</i>				
	100 × log(Rent Q25)	100 × log(Rent Q50)	100 × log(Rent Q75)	
Lottery Winners	0.70 (1.60)	2.57* (1.44)	2.18** (0.92)	
N	2200	2200	2200	2200
Year FE	✓	✓	✓	✓
Block Group Pairs	✓	✓	✓	✓
Tract Value in 2000	✓	✓	✓	✓

Notes: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Notes: The unit of analysis is a census block group, and the sample includes all block groups with lottery applicants over the sample period. Data are from the City of San Francisco Assessor’s Office and from the 2000 and 2019 American Community Surveys. The table reports the β coefficients from equation 4.9. The specifications estimate the effect of an additional lottery winner on the long-run change in neighborhood demographics. See Section 4.6 for details.

We find that a conversion lottery winner increases median household income by 2.47pp and the share of the population with a college degree by 0.87pp both (significant at the 10% level). Lower quartile rents are unchanged, but median and upper quartile rents increase by 2.57pp and 2.18pp respectively. Aligning with our results on assessed property value effects and spillovers, home values increase by more than 2pp across quartiles. The population undergoes a significant demographic change as well. The White population increases by 1.57pp, while the Asian population declines by 1.46pp.

We now consider whether the count and sectoral composition of businesses in these neighborhoods changed as a result of condominium conversions. Our main outcomes will be the total number of new establishments and the total stock of active establishments, broken down by sector. We consider eight large sector groups: Food, Retail, Education, Arts and Entertainment, Professional Services, Manufacturing, Real Estate, and Construction. Many of the neighborhoods in our sample are residential, so our variables contain a large number of zeros. For this reason, we take the inverse hyperbolic sine transformation, and multiply

the outcome by 100 so that coefficients can be interpreted as percentage point changes. However, we stress that the extensive margin response is important for interpreting the magnitude of the effects.

TABLE 4: NEIGHBORHOOD EFFECTS (BUSINESS COUNTS)

<i>Panel A: Local Sectors (New)</i>				
	$100 \times IHS(\text{New Food})$	$100 \times IHS(\text{New Retail})$	$100 \times IHS(\text{New Educ.})$	$100 \times IHS(\text{New Arts \& Entertainment})$
Lottery Winners	-1.75 (2.65)	-1.31 (1.84)	2.96 (1.94)	-0.82 (1.76)
<i>Panel B: Local Sectors (Total)</i>				
	$100 \times IHS(\text{Tot. Food})$	$100 \times IHS(\text{Tot. Retail})$	$100 \times IHS(\text{Tot. Educ.})$	$100 \times IHS(\text{Tot. Arts \& Entertainment})$
Lottery Winners	-6.41* (3.59)	-3.24 (2.74)	7.70*** (2.98)	3.40 (2.98)
<i>Panel C: Other Sectors (New)</i>				
	$100 \times IHS(\text{New Prof. Services})$	$100 \times IHS(\text{New Manufacturing})$	$100 \times IHS(\text{New Real Estate})$	$100 \times IHS(\text{New Construction})$
Lottery Winners	5.07** (2.31)	-1.13 (1.09)	2.34 (1.75)	-4.10*** (1.54)
<i>Panel D: Other Sectors (Total)</i>				
	$100 \times IHS(\text{Tot. Prof. Services})$	$100 \times IHS(\text{Tot. Manufacturing})$	$100 \times IHS(\text{Tot. Real Estate})$	$100 \times IHS(\text{Tot. Construction})$
Lottery Winners	5.10* (2.65)	-1.34 (3.02)	2.20* (1.25)	-9.93*** (3.30)
N	2200	2200	2200	2200
Year FE	✓	✓	✓	✓
Gen. P-score	✓	✓	✓	✓
Block Value in 2000	✓	✓	✓	✓

Notes: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Notes: The unit of analysis is a census block group, and the sample includes all block groups with lottery applicants over the sample period. Data are from the City of San Francisco Assessor’s Office and SF Open Data. The table reports the β coefficients from equation 4.9. The specifications estimate the effect of an additional lottery winner on new business entry. See section 4.6 for details.

These results are contained in Table 4. Panel A uses counts of new establishments for the first four sector groups as outcomes. Panel B uses total counts of active establishments for those same sector groups as outcomes. Panel C and Panel D are structured similarly for the other four sector groups. We find that an additional lottery winner induces a 6.41pp decline in total food establishments (significant at the 10% level) and a 9.93pp decline in total construction establishments. Education, real estate, and professional services establishments increase by 7.70pp, 2.20pp, and 5.10pp respectively. We find no effect for the retail, arts, and manufacturing sectors.

The story that emerges is consistent with condominium conversions inducing gentrification in neighborhoods. White, college-educated, high-income individuals move in. Asian individuals move out. Consistent with our findings in Section 4.5, home values increase.

The fact that rents increase is consistent with condominiums being rent de-controlled, and also consistent with pass-through from increased home values to rents. New establishments in the education sector enter to meet increasing local demand. We also find that some professional service and real estate businesses move in as the neighborhood gentrifies. All of these neighborhood changes likely magnify and reinforce one another.


Heterogeneity

The effects of condominium conversions on demographic outcomes may vary with pre-lottery neighborhood characteristics. To assess this possibility, we estimate a fully-interacted version of equation 4.8 with census tract covariates in 2000. All covariates are normalized to have mean zero and standard deviation one. We then augment the estimand in equation 4.9 to estimate the additional effect of a lottery winner on a neighborhood one standard deviation from the mean of covariate x . The new estimand is given below. As before, standard errors are calculated by bootstrapping the entire procedure.

$$\hat{\beta}_x = \frac{1}{G} \sum_g \sum_{\nu \neq 0} p_g(\nu) \cdot \left(\frac{\hat{\mathbb{E}}[y_{gt} | \nu, p_g(\nu), x = 1] - \hat{\mathbb{E}}[y_{gt} | \nu - 1, p_g(\nu - 1), x = 0]}{\hat{\mathbb{E}}[y_{gt} | \nu, p_g(\nu), x = 1] - \hat{\mathbb{E}}[y_{gt} | \nu - 1, p_g(\nu - 1), x = 0]} \right) \quad (4.10)$$

These estimates are plotted in Figure 12 for each of ten baseline covariates: percentages of the population that are White, Asian, Hispanic, Black, college-educated, living in poverty, and the logs of the population, median household income, median home values, and rents. We focus on the critical outcomes of home value appreciation, and perform the same analysis separately for home values at the 50th and 75th percentile of the neighborhood distribution. Given the modest number of census tracts in our sample (n=123), confidence intervals are plotted at the 90% level.

FIGURE 12: NEIGHBORHOOD HETEROGENEITY



Notes: The unit of analysis is a block-group-year, and the sample includes block groups with lottery applicants over the sample period. Data are from the City of San Francisco Assessor’s Office and the 2000 and the American Community Survey, and the outcome is log property value. The figure reports the β_x coefficients from equation 4.10. These specifications estimate how the effects of lottery winners on property appreciation vary with initial neighborhood characteristics. Standard errors are clustered by census block group, and error bands show 90% confidence intervals.

Panels A and B of Figure 12 provide suggestive evidence that the neighborhood-level effects of condominium conversions on home price appreciation are smaller in areas with

higher poverty rates and Hispanic population shares, and larger in areas with a higher share of college graduates. The results are consistent with the hypothesis that demographic characteristics such as income, race, and education are key mediators of gentrification in US cities.

4.7 Conclusion

Condominium conversions, at the core, entail a change in the legal form of how building units are owned. But conversions bundle several other changes that are important for understanding their effects on neighborhoods: the units are attractive to homeowners in part because they are no longer subject to rent control, face less stringent evictions protections, and are more liquid on the real estate market. We find that conversions also induce property owners to renovate and upgrade units, further causing home values to appreciate. Consequently, nearby parcels become more attractive. In the long-run, the neighborhood gentrifies: rents and home values increase; more educated, higher-income, and white individuals move in; the sectoral composition of local businesses changes; and demographic minorities move out.

Housing supply in many major U.S. cities, and particularly so for San Francisco, has failed to keep pace with housing demand (Figure 1). In housing markets without new construction, how the existing housing stock is renovated, used, sold, and rented is especially important. In this paper, we study the effect of one such supply behavior: condominium conversions. By analyzing the direct and indirect effects of conversions, our findings contribute to the debate across the U.S. on how best to regulate them. While our results are difficult to extrapolate to housing markets where supply is more elastic, even here, condominiums are a legal form commonly chosen by developers. Our paper sheds light on the role that condominiums play in cities and neighborhoods more generally.

Chapter 5

The Return to Protectionism

In Chapter 5, Pablo Fajgelbaum, Penelopi Goldberg, Amit Khandelwal and I study the aggregate and distributional effects of the import and export tariffs enacted during the ongoing US-China trade war. After decades of supporting free trade, in 2018 the United States raised import tariffs and major trade partners retaliated. We analyze the short-run impact of this return to protectionism on the U.S. economy. We find that import and retaliatory tariffs caused large declines in imports and exports. Prices of imports targeted by tariffs did not fall, implying complete pass-through of tariffs to duty-inclusive prices. The resulting losses to U.S. consumers and firms that buy imports was \$51 billion, or 0.27% of GDP. We embed the estimated trade elasticities in a general-equilibrium model of the U.S. economy. After accounting for tariff revenue and gains to domestic producers, the aggregate real income loss was \$7.2 billion, or 0.04% of GDP. Import tariffs favored sectors concentrated in politically competitive counties, and the model implies that tradeable-sector workers in heavily Republican counties were the most negatively affected due to the retaliatory tariffs.

This chapter is coauthored with Pablo Fajgelbaum, Penelopi Goldberg, and Amit Khandelwal. An edited version of this paper was published in 2020 in the *Quarterly Journal of Economics*.

5.1 Introduction

After more than a half-century of leading efforts to lower international trade barriers, in 2018 the U.S. enacted several waves of tariff increases on specific products and countries. Import tariffs increased from 2.6% to 16.6% on 12,043 products covering \$303 billion (12.7%) of annual U.S. imports. In response, trade partners imposed retaliatory tariffs on U.S. exports. These counter-measures increased tariffs from 7.3% to 20.4% on 8,073 export products covering \$127 billion (8.2%) of annual U.S. exports.

This return to protection is unprecedented in the post-war era due to the sizes of the countries involved, the magnitudes of the tariff increases, and the breadth of tariffs across sectors. What were the short-run impacts on the U.S. economy? Classical trade theory dictates that the effects depend on the incidence of tariffs. Consumers and firms who buy foreign products lose from higher tariffs. Reallocations of expenditures into or away from domestic products induced by the U.S. and retaliatory tariffs may lead to changes in U.S. export prices relative to import prices—that is, terms-of-trade effects—and generate tariff revenue. The trade war may also have distributional consequences across sectors, and therefore across regions with different patterns of specialization.

Very little is known about tariff incidence, despite its central role in policy analysis. In this paper, we first estimate the impacts of tariffs on U.S. trade quantities and prices. We estimate a U.S. demand system that accommodates reallocations across imported *varieties* (defined as country-product pairs), across imported *products* (defined as 10-digit Harmonized System product codes), and between imported and domestic products within a *sector* (defined as a 4-digit NAICS industry code). We combine this system with foreign export supply curves for each variety. The estimation leverages the property that if changes in tariffs are uncorrelated with demand and supply shocks, then a tariff can be used to simultaneously instrument *both* the import demand and foreign export supply curves.¹ We exploit panel variation at the variety level, and aggregate tariffs to construct instruments that identify elasticities of substitution at the product and sector levels. Tests for pre-existing trends, tariff anticipation, and an event-study framework validate using tariffs as a source of identification.

We find large declines in imports when the tariffs were implemented. Imports of varieties targeted by U.S. tariffs fell on average 31.7%; imports of targeted products fell 2.5%; and imports in targeted sectors fell 0.2%. The event study reveals no differential change in before-duty import prices between targeted and untargeted source countries exporting the same product. These results imply that we cannot reject horizontal foreign export supply curves. We estimate elasticities of substitution across origins (i.e., varieties) within a product, across imported products, and between domestic goods and imports within a sector of 2.53, 1.53, and 1.19, respectively.

On the export side, we find that retaliatory tariffs resulted in a 9.9% decline in U.S. exports within products. We estimate a roughly unitary elastic foreign demand for

¹This estimation approach was first applied by Romalis (2007) to study the effects of NAFTA and recently formalized by Zoutman et al. (2018).

U.S. varieties (1.04), and also find complete pass-through of retaliatory tariffs to foreign consumers. As with the import side, we demonstrate that these elasticities are not confounded by pre-existing trends or anticipation of the retaliations.

The findings imply complete pass-through of tariffs to duty-inclusive import prices, a finding that is systematic across products with heterogeneous characteristics. The resulting real income loss to U.S. consumers and firms who buy imports can be computed as the product of the import share of value added (15%), the fraction of U.S. imports targeted by tariff increases (13%), and the average increase in tariffs among targeted varieties (14%). This decline is \$51 billion, or 0.27% of GDP.

The previous results have two important caveats. First, our analysis considers short-run effects, but relative prices could change over longer horizons. Second, our estimation controls for country-time and product-time effects, and therefore is unable to capture import price declines due to relative wage changes across countries or sectors.² In other words, the results do not imply that the U.S. is a small open economy unable to affect world prices, as terms-of-trade effects could have occurred through wage adjustments at the country-sector level.

We combine the previously estimated parameters with a supply side model of the U.S. economy to gauge some of these effects. The model imposes upward sloping industry supply curves in the U.S. and predicts changes in sector-level prices in the U.S. due to demand reallocation induced by tariffs. We impose perfect competition, flexible prices, and flexible adjustment of intermediate inputs. To assess regional effects, we assume immobile labor and calibrate the model to match specialization patterns across U.S. counties.³ In the model, U.S. tariffs reallocate domestic demand into U.S. goods, raising total demand and therefore U.S. export prices, while retaliatory tariffs have the opposite effect. These price changes are qualitatively consistent with suggestive evidence that U.S. tariffs led to increases in the PPI and that sector-level export prices fell with retaliatory tariffs.

We obtain a ballpark estimate of the aggregate and regional effects of the 2018 tariff waves. We estimate producer gains of \$9.4 billion, or 0.05% of GDP. Adding up these gains, tariffs revenue, and the losses from higher import costs yields a short-run loss of the 2018 tariffs on aggregate real income of \$7.2 billion, or 0.04% of GDP. Hence, we find substantial redistribution from buyers of foreign goods to U.S. producers and the government, but a small net loss for the U.S. economy as a whole (which is not statistically significant at conventional levels after accounting for the parameters' standard errors). While we cannot reject the null

²Influential work by [Bagwell and Staiger \(1999\)](#) demonstrates that trade agreements serve to deal with terms-of-trade externalities.

³Our model-based calculations abstract from imperfect competition in international transactions, although incorporating variable markups would imply incomplete pass-through, which we do not observe. We measure input-output linkages at the 4-digit industry level observed in BEA IO tables and impose unitary elasticities as in [Caliendo and Parro \(2015\)](#). The aggregate impacts could be larger under tariff uncertainty ([Handley and Limão, 2017](#)) or different assumptions on the input-output structure ([Antràs and De Gortari, 2017](#); [Baqaee and Farhi, 2019](#)). See [Freund et al. \(2018\)](#), [Altig et al. \(2018\)](#) and [Bellora and Fontagné \(2019\)](#) for analyses that incorporate some of these forces in the context of the 2018 trade war.

hypothesis that the aggregate losses are zero, the results strongly indicate large consumer losses from the trade war. If trade partners had not retaliated, the economy would have experienced a modest (and also not statistically significant) gain of \$0.5 billion.

The small net effect also masks heterogeneous impacts across regions driven by patterns of specialization across sectors. If capital and labor are regionally immobile—a reasonable assumption over this short time horizon—sectoral heterogeneity in U.S. and foreign tariffs generates unequal regional impacts. Our counterfactuals imply that all counties experienced reductions in tradeable real wages. Using the model, we find a standard deviation of real wages in the tradeable sectors across counties of 0.5%, relative to an average decline of 1.0%.

We show that U.S. import protection was biased toward products made in electorally competitive counties, as measured by their 2016 Presidential vote share, suggesting a potential *ex ante* electoral rationale for the pattern of tariffs increases. This structure of U.S. protection is consistent with the view that trade policies determined by electoral competition tend to favor voters who are likely to be closer to an indifference point between candidates (Mayer, 1984; Dixit and Londregan, 1996; Grossman and Helpman, 2005). In contrast, retaliations disproportionately targeted agricultural sectors, which tend to be concentrated in Republican-leaning counties. The model-based results suggest that tradeable-sector workers in heavily Republican counties were the most negatively because of this pattern of tariff retaliations.

A large literature studies the impacts of changes in trade costs or foreign shocks through empirical and quantitative methods (e.g., Eaton and Kortum (2002), Arkolakis et al. (2012) and Autor et al. (2013)). We focus instead on trade policy, and on tariffs in particular, since they are the primary policy instrument of the 2018 trade war.

One approach to studying the impacts of trade policy uses *ex post* variation in tariffs across sectors to assess impacts on sectors (e.g., Attanasio et al. 2004), regions (e.g., Topalova (2010), Kovak (2013), and Dix-Carneiro and Kovak (2017)), firms (e.g., Amiti and Konings (2007), Goldberg et al. (2010), and Bustos (2011)), or workers (e.g., Autor et al. (2014) and McCaig and Pavcnik (2018)). A complementary approach uses quantitative models to simulate aggregate impacts of tariffs, such as the Nash equilibrium of a global trade war (Ossa, 2014) or regional trade liberalizations (e.g., Caliendo and Parro (2015) and Caliendo et al. (2015)).⁴

A key challenge in the empirical literature is to address the potential endogeneity of tariff changes, and we devote significant attention to these concerns in our analysis. In quantitative models, the parametrization of how trade volumes change with trade policy plays a key role, and we use the observed changes in tariffs to estimate these trade elasticities.⁵

⁴Goldberg and Pavcnik (2016) and Ossa (2016) survey the recent literature studying the impacts of trade policy.

⁵Some papers use time-series variation in tariffs to estimate trade elasticities; e.g., see Romalis (2007), Spearot (2013) and Spearot (2016). Hillberry and Hummels (2013) and Head and Mayer (2014) review alternative approaches typically used to estimate demand elasticities, including gravity estimates of the relationship between trade and prices or proxies of marginal costs (e.g., Eaton and Kortum (2002), Atkin and Donaldson (2015), Simonovska and Waugh (2014) and Donaldson (2018)) or GMM identification via

Finally, our finding of complete pass-through deserves some discussion.⁶ [Amiti et al. \(2019\)](#) and [Cavallo et al. \(2019\)](#) also find complete tariff pass-through to border prices in this trade war, and [Flaaen et al. \(2019\)](#) estimate high tariff pass-through to retail prices for washing machines. Yet, a large literature has estimated incomplete pass-through, in particular for exchange rates (e.g., [Goldberg and Knetter, 1997](#)). An exception is [Feenstra \(1989\)](#), who finds symmetry in the pass-through between tariffs and exchange-rate movements. Several hypotheses could reconcile our findings with the exchange rate pass-through literature. The persistence of the tariff shocks may cause before-duty import prices to eventually decline as time elapses. Our results are also consistent with incomplete exchange rate pass-through if import prices are sticky and denominated in dollars ([Gopinath et al., 2010](#)). Inspecting the precise mechanism underlying the complete tariff pass-through finding deserves further exploration in future research.

The remainder of the paper is structured as follows. Section 5.2 summarizes the data used for the analysis. Section 5.3 outlines the demand-side framework that guides the estimation of the elasticities and discusses the identification strategy. Section 5.4 presents the empirical results. Section 5.5 presents the model-based aggregate and distributional effects. Section 5.6 concludes.

5.2 Data and Timeline

This section describes the data, provides a timeline of key events, and presents an event study of the impact of tariffs. The details about the dataset construction are available in the Online Appendix.

Data

We build a monthly panel dataset of U.S. statutory import tariffs using public schedules from the U.S. International Trade Commission (USITC). Prior to 2018, USITC released annual “baseline” tariff schedules in January and a revised schedule in July. In 2018, by contrast, USITC issued 14 schedule revisions, reflecting a rapid series of tariff increases. These ad-valorem tariff increases were predominantly set at the 8-digit Harmonized System (HS) level, and swiftly implemented within 1-3 weeks following a press release by the Office of the U.S. Trade Representative.⁷ As we work with monthly data and the tariffs were

heteroskedasticity of supply and demand shocks (e.g., [Feenstra \(1994\)](#) and [Broda et al., 2008](#)). Our elasticities are lower than those obtained from cross-sectional variation but in the range of estimates from time series estimation (see [Hillberry and Hummels 2013](#)).

⁶The few papers studying the impact of tariffs on import prices include [Irwin \(2014\)](#) for U.S. sugar duties in late 19th and early 20th centuries, [Winkelmann and Winkelmann \(1998\)](#) for 1980’s tariff reductions in New Zealand, and [Feenstra \(1989\)](#) for U.S. duties on Japanese compact trucks in the 1980s.

⁷We ignore a small number of changes in import tariffs in 2018:1, 2018:7, and 2019:1 that are the result of pre-existing treaty commitments. Thus, we use only the tariff changes due to the trade war as identifying variation.

implemented in the middle of months, we scale the tariff increases by the number of days of the month they were in effect.

We compile retaliatory tariffs on U.S. exports from official documents released by the Ministry of Finance of China, the Department of Finance of Canada, the Office of the President of Mexico, and the World Trade Organization (covering the EU, Russia, and Turkey). These tariffs were also entirely *ad valorem* and went into effect shortly after the announcement dates. To construct the retaliatory tariffs, we use the annual WTO database of Most Favored Nation (MFN) tariff rates, and compute the retaliatory tariff rate for each country-product as the sum of the MFN rate and the announced tariff rate change. We measure export tariffs at the HS-6 level, since HS-8 codes are not directly comparable across countries. As with the import tariffs, we scale the retaliations based on the day of the month they go into affect.

We use publicly-available monthly administrative U.S. import and export data from the U.S. Census Bureau that record values and quantities of trade flows at HS-10 level, which we refer to as *products*.⁸ Country-product pairs are referred to as *varieties*. Our sample period covers 2017:1 to 2019:4, and covers the universe of HS-10 codes and countries. For imports, we directly observe the value of duties collected. Unit values are constructed as the ratio of values to quantities, and duty-inclusive unit values are constructed as $(\text{value} + \text{duties})/\text{quantity}$. We do not observe the duties collected by foreign governments on U.S. exports, so we construct duty-inclusive unit values for exports as the unit value multiplied by $(\text{one} + \text{the } ad\ valorem\ retaliatory\ statutory\ rate)$.

We define *sectors* as NAICS-4 codes. We use the Federal Reserve G17 Industrial Production Index as a measure of domestic sector output, and the BLS PPI, MPI and XPI indices of producer prices, import prices and export prices, respectively. These sector-level panels are available at monthly frequency. We use the 2016 Bureau of Economic Analysis (BEA) annual “use” tables from the national input-output (I-O) accounts to construct I-O linkages between sectors.

To analyze regional exposure we use the Census County Business Patterns (CBP) database, which provides annual industry employment and wage data at the county-by-sector level for all non-farm sectors. For county-level data covering the farm sector, we use the BEA Local Area Personal Income and Employment database. From both data sources, we use 2016 data to compute the industry employment share of each county. Finally, we obtain county-level demographic statistics from the 2016 five-year American Community Survey and county-level voting data from the U.S. Federal Election Commission.

Timeline

Table 1 provides a timeline of events, and Figure ?? plots the tariff increases. Panel A of Table 1 reports the total scope of affected imports, and shows that U.S. import tariffs targeted 12,043 distinct HS-10 products. In 2017, these imports were valued at \$303 billion,

⁸These data are available at <https://usatrade.census.gov/>.

or 12.7% of imports. The average statutory tariff rate increased from 2.6% to 16.6%. An important feature of these tariffs is that they were discriminatory across countries, which allows us to exploit variation in tariff changes across varieties within products.⁹

⁹The U.S. authorized the tariffs through Section 201 of the Trade Act of 1974, Section 301 of the Trade Act of 1974, and Section 232 of the Trade Expansion Act of 1962. These laws permit the president to apply protectionist measures under different justifications, including “serious injury” to domestic industries, threats to national security, or retaliations for allegations of unfair trade practices.


TABLE 1: THE 2018 TRADE WAR

Panel A: Tariffs on U.S. Imports Enacted by U.S.						
Tariff Wave	Date Enacted	Products	2017 Imports		Tariff (%)	
		(# HS-10)	(mil USD)	(%)*	2017	Post-War
Solar Panels	Feb 7, 2018	8	5,782	0.2	0.0	30.0
Washing Machines	Feb 7, 2018	8	2,105	0.1	1.3	32.2
Aluminum	Mar-Jun, 2018	93	17,685	0.7	2.0	12.8
Iron and Steel	Mar-Jun, 2018	757	30,655	1.3	0.0	24.8
E.U.	October 18, 2019	313	11,819	0.5	4.8	28.7
China	Jul '18 - Sep '19	16,403	352,563	14.7	3.2	16.8
Total		17,295	420,608	17.6	2.9	17.8


Panel B: Retaliatory Tariffs on U.S. Exports Enacted by Trading Partners						
Retaliating Country	Date Enacted	Products	2017 Exports		Tariff (%)	
		(# HS-10)	(mil USD)	(%)*	2017	Post-war
Mexico	Jun 5, 2018	232	6,746	0.4	9.4	27.9
Turkey	Jun 21, 2018	248	1,554	0.1	8.8	31.6
European Union	Jun 22, 2018	303	8,244	0.5	4.4	28.9
Canada	Jul 1, 2018	325	17,818	1.2	2.1	20.2
Russia	Aug 6, 2018	165	268	0.0	5.2	37.2
India	Jun 16, 2019	65	1,280	0.1	13.2	27.5
China	Apr '18 - Sep '19	7,757	98,016	6.3	8.7	19.5
Total		8,400	133,926	8.7	7.7	20.8

Notes: “*” Values indicate percentage point tariff increases. Panels display unweighted monthly HS10-country average statutory tariff rates. 2017 tariff rates computed as the annual average; 2018 tariff rates computed using data from December 2018. Total tariff rates are computed as the trade-weighted average of table row values. The denominator for import (export) share is the total 2017 annual USD value of all U.S. imports (exports). The U.S. government announced import tariffs on aluminum and steel products on March 23 but granted exemptions for Canada, Mexico, and the European Union; those exemptions were lifted on June 1. The dates of Chinese retaliations are: April 6, July 2, August 23, and September 24. See text for data sources.

Panel A: Tariffs on U.S. Imports



Panel B: Retaliatory Tariffs on U.S. Exports



Notes: Figure shows the unweighted average tariff rate of targeted import and export varieties for each tariff wave before and after they are targeted. Import tariffs constructed from U.S. International Trade Commission (USITC) documents, and retaliatory tariffs constructed using official documents from foreign finance and trade ministries.

The first wave of tariff increases began in February 2018, when the U.S. increased tariffs on \$8 billion of solar panel and washing machine imports. A second wave of tariffs, implemented in March 2018, targeted iron, aluminum, and steel products. The largest tranches of import tariffs targeted approximately \$247 billion worth of imports from China. In March 2018 the U.S. implemented tariffs on approximately \$50 billion of Chinese imports, and the scope and value of targeted products on China expanded with subsequent tariffs waves implemented

in July and September. Rows 5-7 indicate that tariffs on China targeted 11,207 imported products worth \$247 billion, and increased tariffs, on average, from 3.0% to 15.5%. A total of 48.8% of imports from China were targeted with tariff increases.

Panel B of Table 1 reports the retaliatory tariffs imposed on U.S. exports by trade partners. Canada, China, Mexico, Russia, Turkey, and the E.U. enacted retaliatory tariffs against the U.S., and collectively these retaliations covered \$127 billion (8.2%) of annual U.S. exports across 7,763 products. The average statutory tariff rate on these exports increased from 7.3% to 20.4%.

Structure of Protection across Sectors

Table 2 reports summary statistics for targeted import and export varieties across NAICS-3 codes. For imports, we report the number of targeted HS-10 products and varieties, and the means and standard deviations of tariff increases across targeted varieties within NAICS-3 codes. In sectors where only China was targeted, the number of targeted products equals the number of targeted varieties. The table also reports the corresponding statistics for the retaliatory tariffs on U.S. exports.

TABLE 2: SECTOR VARIATION IN TARIFF RATE CHANGES

Tariff Changes for Targeted Import Varieties and Export Products, by NAICS3 Code										
Sector (1)	NAICS-3 (2)	Imports (U.S. Tariffs)				Exports (Retaliatory Tariffs)				
		# Varieties (3)	# Products (4)	Δ Tariffs		# Products (7)	# Varieties (8)	Δ Tariffs		
				Mean (5)	Std. (6)			Mean (9)	Std. (10)	
Crop and Animal Production	111-2	456	456	0.10	0.00	303	380	0.24	0.11	
Forestry and Logging	113	71	71	0.10	0.00	79	79	0.12	0.07	
Fishing, Hunting and Trapping	114	486	486	0.10	0.00	247	247	0.24	0.03	
Oil and Gas Extraction	211	17	17	0.10	0.00	8	8	0.22	0.07	
Mining (except Oil and Gas)	212	103	103	0.10	0.00	89	92	0.10	0.05	
Food	311	732	732	0.10	0.00	622	1,014	0.17	0.09	
Beverage and Tobacco Product	312	64	64	0.10	0.00	55	379	0.23	0.06	
Textile Mills	313	1,502	1,502	0.10	0.00	468	494	0.12	0.06	
Textile Product Mills	314	176	176	0.10	0.00	122	235	0.16	0.08	
Apparel	315	92	92	0.10	0.00	325	1,082	0.20	0.07	
Leather and Allied Product	316	237	237	0.10	0.00	196	357	0.16	0.08	
Wood Product	321	424	424	0.10	0.00	194	194	0.10	0.03	
Paper	322	335	335	0.12	0.05	239	388	0.12	0.07	
Printing and Related Activities	323	14	14	0.10	0.00	46	74	0.13	0.09	
Petroleum and Coal Products	324	74	74	0.13	0.06	64	64	0.23	0.05	
Chemical	325	1,730	1,730	0.12	0.05	1,159	1,411	0.12	0.08	
Plastics and Rubber Products	326	251	251	0.15	0.07	171	196	0.10	0.07	
Nonmetallic Mineral Product	327	354	354	0.11	0.03	225	632	0.18	0.08	
Primary Metal	331	1,147	14,093	0.19	0.07	495	1,738	0.20	0.07	
Fabricated Metal Product	332	583	852	0.14	0.06	404	1,236	0.18	0.09	
Machinery	333	1,344	1,344	0.20	0.07	1,075	1,218	0.11	0.06	
Computer and Electronic Product	334	617	878	0.21	0.07	458	506	0.11	0.07	
Electrical Equipment and Appliances	335	414	594	0.18	0.08	326	656	0.16	0.08	
Transportation Equipment	336	429	429	0.15	0.07	273	680	0.21	0.08	
Furniture and Related Product	337	160	160	0.10	0.01	37	244	0.21	0.07	
Miscellaneous	339	231	231	0.13	0.06	393	608	0.16	0.09	
Total		12,043	25,699	0.12	0.03	8,073	14,212	0.17	0.07	

Notes: Table shows the mean and standard deviation of tariff increases across 3-digit NAICS sectors. A tariff change of 0.10 indicates a 10 percentage point increase. Sectors with the same number of targeted varieties and products in Columns 3 and 4 reflect import tariffs exclusively targeting Chinese products. Means and standard deviations in the final row computed as the simple average of table row values. Import tariffs constructed from U.S. International Trade Commission (USITC) documents, and retaliatory tariffs constructed using official documents from foreign finance and trade ministries.

The table conveys three facts. First, U.S. sectors that receive the most protection are primary metals, machinery, computer products, and electrical equipment and appliances. These sectors contain a large share of intermediate inputs, comprise a large share of targeted varieties and products, and saw steep tariff increases relative to most other sectors.¹⁰ Second, U.S. trade partners concentrated retaliatory tariffs on different sets of products and sectors; the sector-level correlation between import and retaliatory tariffs is 0.46. For example, retaliatory tariff increases on U.S. agriculture exports are on average more than double the U.S. tariff increases in the crop, fishing, and beverage and tobacco sectors. Third, column 5 shows that the mean tariff increases on targeted import varieties are similar across sectors, and column 6 shows that the standard deviation of U.S. tariff changes within sectors is low (and most often, zero).

¹⁰The Online Appendix provides a breakdown of the targeted products by final versus intermediate goods. For this table, we manually construct a match HS-10 products to BLS Consumer Price Index product codes. This match suggests that 87% of targeted products within these sectors are intermediate goods (in value), compared to 72% of targeted products in all other sectors.

Since [Johnson \(1953\)](#), an extensive literature on optimal tariffs has argued that governments can maximize national income by setting higher tariffs on sectors with more inelastic foreign export supply, and [Broda et al. \(2008\)](#) offer empirical support. However, the tariff changes observed in the 2018 trade war are highly similar across sectors. The Online Appendix illustrates this point by plotting the distribution of tariff changes for targeted varieties. The left panel shows that, during the trade war, the U.S. applied only five tariff rate changes to targeted varieties: 10%, 20%, 25%, 30%, and 50%. Virtually all varieties (99.8%) were hit with either 10% or 25% tariff changes. The right panel shows that most of the retaliatory rate increases were concentrated at 10% or 25% as well. These patterns suggest that neither the U.S. nor retaliating countries were likely driven by a terms-of-trade rationale, since in that case we would expect tariff changes to vary across sectors. The Online Appendix plots average 2018 sector-level tariff rates against the foreign export supply elasticities estimated by [Broda et al. \(2008\)](#) and reveals a negative (and statistically insignificant) relationship (the correlation is -0.10).

This lack of variation across sectors also suggests that the tariff changes are unlikely to have been driven by sector-specific interest groups. Explanations in this tradition argue that sectors make political campaign contributions and engage in costly lobbying activities to secure import protection from policymakers ([Grossman and Helpman, 1994](#); [Goldberg and Maggi, 1999](#)). However, these explanations also rely on variation in protection across sectors. The Online Appendix plots financial campaign contributions made to candidates for the U.S. House of Representatives in the 2016 election against tariff changes at the sector-level, and reveals a negative, rather than a positive, correlation. While this evidence is only suggestive, it appears unlikely that campaign contributions were the main determinant of the U.S. tariff structure in the trade war.

Event Study

We visualize the impacts of the tariff war on trade using an event-study framework. To assess impacts, we compare the trends of targeted varieties (those directly affected by a tariff increase) to varieties not targeted in the following specification:

$$\ln y_{igt} = \alpha_{ig} + \alpha_{gt} + \alpha_{it} + \sum_{j=-6}^6 \beta_{0j} I(\text{event}_{igt} = j) + \sum_{j=-6}^6 \beta_{1j} I(\text{event}_{igt} = j) \times \text{target}_{ig} + \epsilon_{igt}. \quad (5.1)$$

This specification includes variety (α_{ig}), country-time (α_{it}) and product-month (α_{gt}) fixed effects. Varieties targeted by tariffs are captured by the target_{ig} dummy. The inclusion of α_{gt} fixed effects implies that the β_{1j} coefficients are identified using variation between targeted and non-targeted varieties within product-time. The event time coefficients are captured by the indicator variables. In these specifications, we assign the event date of targeted varieties to be the nearest full month to the actual event date, using the 15th of the month as the

cutoff date.¹¹ Non-targeted varieties within the same HS-10 product as a targeted variety are assigned the earliest event date within that product code. For all other non-targeted varieties, we assign the event date to be the earliest month of a targeted variety within the same NAICS-4 sector. If a non-targeted variety does not share the same NAICS-4 as any targeted varieties, we sequentially use NAICS-3 and NAICS-2 codes, and otherwise assign the event month to be the earliest month of the trade war (February 2018 for imports, and April 2018 for exports). We bin event times ≥ 6 together and exclude event time ≤ -7 . For import outcomes, standard errors are clustered by country and HS-8, since these are generally the levels at which the tariffs are set.¹² For export outcomes, standard errors are clustered by HS-6 and country; here, we use HS-6 because that is the finest level at which product codes are comparable across countries and the level at which we code the retaliatory tariffs. We plot the β_{1j} dummies that capture the relative trends of targeted varieties.

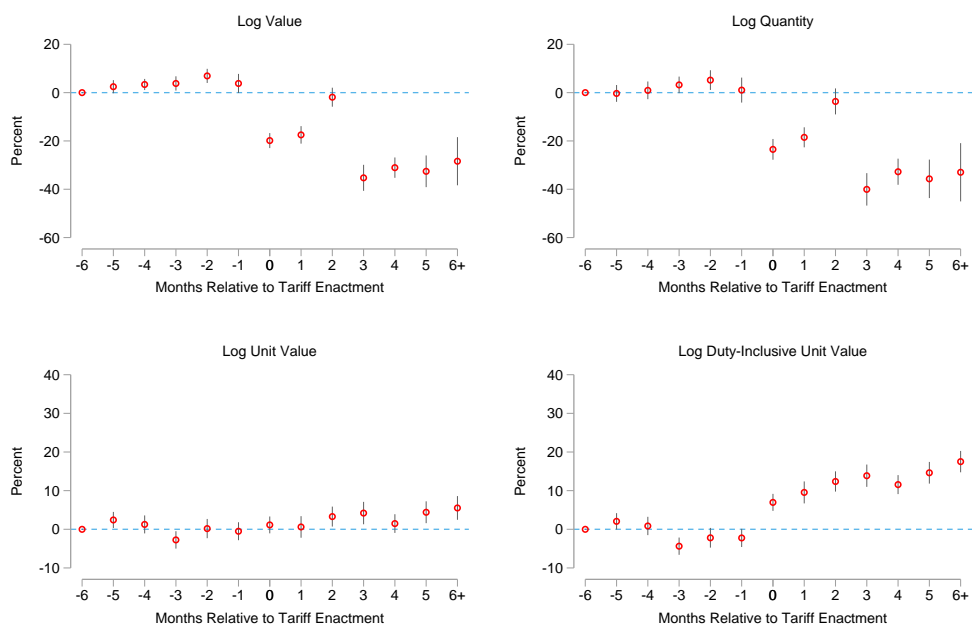
Figure 1 reports the impacts on imported varieties. The top two panels trace the impact of tariffs on import values and quantities, and the bottom panels show the effects on unit values, both exclusive and inclusive of duties. Upon impact, we detect large declines in imports. Import values decline on average by 20% and quantities decline by 23%.¹³ In the bottom-left panel, before-duty unit values do not change. However, duty-inclusive unit values increase sharply for targeted varieties. These two panels provide initial evidence of complete pass-through of the tariffs to import prices at the variety level.

¹¹The event date varies by both product and country since some varieties within the same product code are targeted before others. For example, the U.S. imposed steel tariffs on Canada, Mexico, and the EU three months after imposing steel tariffs on other countries.

¹²In a small number of cases, tariffs vary within HS-8 codes at the HS-10 level. See the Online Appendix.

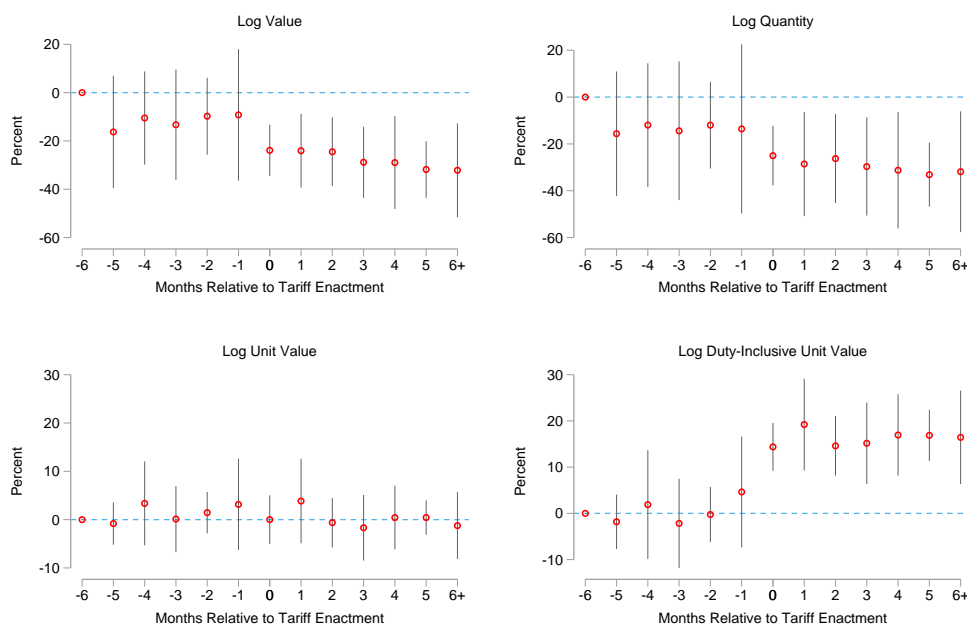
¹³The figure reveals a temporary increase in import values and quantities in event period +2 for targeted varieties. In the Online Appendix we show that this increase is driven by imports in December 2018 as a result of a September 2018 announcement that the U.S. would increase tariffs on \$200 billion of already-targeted Chinese varieties from 10% to 25% on January 1, 2019. A plausible reason why we observe large anticipation effects only in this instance is that, unlike in previous U.S. tariff waves, the January 2019 tariff escalation was announced long in advance and was perceived to be credible given the previous tariff waves. The U.S. ultimately implemented this threat in May 2019, which is beyond our sample period.

FIGURE 1: VARIETY EVENT STUDY: IMPORTS



Notes: Figure plots event time dummies for targeted varieties relative to untargeted varieties. Regressions include country-product, product-time, and country-time fixed effects. Standard errors clustered by country and HS-8. Event periods before -6 are dropped, and event periods ≥ 6 are binned. Error bars show 95% confidence intervals. In the Online Appendix we provide evidence that the temporary surge in imports during event period 2 reflects an anticipation response to additional tariff threats on a subset of Chinese varieties. Sample: Monthly variety-level import data from U.S. Census. Sample period is 2017:1 to 2019:4.

FIGURE 2: VARIETY EVENT STUDY: EXPORTS



Notes: Figure plots event time dummies for targeted varieties relative to untargeted varieties. Regressions include country-product, product-time, and country-time fixed effects. Standard errors clustered by country and HS-6. Event periods before -6 are dropped, and event periods ≥ 6 are binned. Error bars show 95% confidence intervals. Sample: Monthly variety-level export data from U.S. Census. Sample period is 2017:1 to 2019:4.

The event study also addresses concerns of tariff anticipation that would complicate the elasticity estimates. The figure reveals anticipatory effects occurring before the tariff changes, but they are quantitatively small. Hence, the concern that importers shifted forward their purchases in order to avoid paying tariffs is mild. Below, we further assess tariff anticipation through dynamic specifications.

Figure 2 reports the impacts of the retaliatory tariffs on U.S. exports. The patterns are similar to what we observe for imports. We find that, at the month of implementation, export values decline on average by 24% and quantities fall by 25%. Again, we observe no change in the before-duty unit values, suggesting complete pass-through of the retaliatory tariffs to foreigners' imports of U.S. varieties. We also observe no clear pattern of anticipation for U.S. exports.

5.3 Trade Framework and Identification

We now describe the trade framework that guides the estimation. We defer supply-side and general-equilibrium assumptions to Section 5.5.

U.S. Import Demand

There are S traded sectors corresponding to 4-digit NAICS sectors (collected in the set \mathcal{S} and indexed by s). Within each traded sector, aggregate demand (from producers and consumers) is structured according to a 3-tier CES demand system. In the upper nest there is differentiation between domestic and imported goods. Within each of these two nests of sector s there are G_s products (collected in the set \mathcal{G}_s and indexed by product g) corresponding to an HS-10 level of aggregation. Within the nest of imported products, varieties are differentiated by country of origin. The U.S. trades with I countries (collected in the set \mathcal{I} and indexed by country i).

The CES utility functions and price indexes are presented in the Online Appendix. This structure gives U.S. import demand in each tier as a function of prices. The value of imports in sector s is:

$$P_{M_s}M_s = E_s A_{M_s} \left(\frac{P_{M_s}}{P_s} \right)^{1-\kappa}, \quad (5.2)$$

where E_s are aggregate U.S. expenditures in sector s from both final consumers and firms, A_{M_s} is an import demand shock, P_{M_s} is the import price index defined in equation (??) in the Online Appendix, and P_s is the sector price index defined in equation (??).

The value of imports for product g in sector s is

$$p_{M_g}m_g = P_{M_s}M_s a_{M_g} \left(\frac{p_{M_g}}{P_{M_s}} \right)^{1-\eta}, \quad (5.3)$$

where a_{M_g} is an import demand shock and p_{M_g} is the import price index of product g defined in equation (??).

Finally, the quantity imported of product g 's variety from country i is:

$$m_{ig} = m_g a_{ig} \left(\frac{p_{ig}}{p_{M_g}} \right)^{-\sigma}, \quad (5.4)$$

where a_{ig} is a demand shock and p_{ig} is the domestic price of the variety ig . The U.S. imposes ad-valorem tariffs τ_{ig} on the CIF price p_{ig}^* , so the domestic price is:

$$p_{ig} = (1 + \tau_{ig}) p_{ig}^*. \quad (5.5)$$

The previous demand equations depend on three elasticities: across imported varieties within product (σ), across products (η), and between imports and domestic products within a sector (κ).¹⁴

Foreign Export Supply and Import Demand

Trade partners are represented with export-supply and import-demand curves at the variety level. We allow for import price effects of U.S. trade policy through potentially upward sloping foreign export supply. The inverse foreign export supply curve is:

$$p_{ig}^* = z_{ig}^* m_{ig}^{\omega^*}, \quad (5.6)$$

where z_{ig}^* is a foreign marginal cost shifter that could also include a bilateral iceberg trade cost. The parameter ω^* is the *inverse* foreign export supply elasticity. It is a key determinant of the effects of U.S. trade policy, as it drives the magnitude of the reduction in foreign prices when U.S. tariffs are imposed. Before-duty import prices p_{ig}^* fall more sharply the larger is ω^* .

Each foreign country demands a quantity x_{ig} of U.S. exports of good g . Foreign import demand for U.S. varieties is similar to (5.4) on the U.S. side, but with a potentially different demand shifter and demand elasticity:

$$x_{ig} = a_{ig}^* \left((1 + \tau_{ig}^*) p_{ig}^X \right)^{-\sigma^*}, \quad (5.7)$$

where x_{ig} is the U.S. exports of product g to country i , p_{ig}^X is export price received by exporters, τ_{ig}^* is the *ad valorem* tariff set by country i on U.S. exports of good g , and a_{ig}^* is a foreign demand shock.

Identification

This section discusses the identification strategy for the elasticities and its potential threats.

U.S. Import and Foreign Export Variety Elasticities (σ, ω^*)

We use variation in U.S. import tariffs to estimate the variety import demand and export supply elasticities simultaneously. The strategy of identifying two elasticities with one

¹⁴This demand system is also used by Broda et al. (2008). In our setting, it is motivated by the available monthly public data: variety- and product-level imports and exports, and sector-level domestic production data. With this nesting structure, it is sufficient to observe the import shares of expenditures within each sector s to estimate the elasticities and implement counterfactuals. It does not require information on import shares within each product g , which are not observed in publicly available data but would be required under alternative nesting assumptions. A potential shortcoming is that the imports m_g of any particular product g in sector s impact the domestic expenditures of that same product only through sector-level shifters. Inverting the order of the top two nests does not matter for the estimation of the lowest tier elasticities (σ , σ^* and ω^*), and it would not matter for the implementation of counterfactuals if κ and η were equal.

instrument was applied by Romalis (2007) in a trade context and studied by Zoutman et al. (2018) in the context of applications to public finance. Intuitively, tariffs create a wedge between what the importer pays and what the exporter receives. A tariff shifts down the demand curve for any given price received by the exporter, tracing the supply curve. Similarly, a tariff shifts up the supply curve for any given price paid by the consumer, tracing the demand curve. Hence, data on changes in prices, tariffs, and quantities is sufficient to trace both the demand and supply curves simultaneously.

Adding a time subscript and log-differencing over time, equations (5.4) and (5.6) can be written as:

$$\Delta \ln m_{igt} = \eta_{gt}^m + \eta_{it}^m + \eta_{is}^m - \sigma \Delta \ln p_{igt} + \varepsilon_{igt}^m, \quad (5.8)$$

$$\Delta \ln p_{igt}^* = \eta_{gt}^{p^*} + \eta_{it}^{p^*} + \eta_{is}^{p^*} + \omega^* \Delta \ln m_{igt} + \varepsilon_{igt}^{p^*}, \quad (5.9)$$

where, $y = \{p^*, m\}$, the η_{gt}^y are product-time fixed effects, the η_{it}^y are country-time fixed effects, and the η_{is}^y are country-sector fixed effects (s is the sector of product g). For now, suppose that tariffs are uncorrelated with unobserved import demand and export supply shocks entering in the residuals, an issue we return to at the end of this sub-section. Then, the import demand elasticity σ is identified by instrumenting the duty-inclusive price Δp_{igt} with the tariff $\Delta \tau_{igt}$ in (5.8). The export supply ω^* is identified by instrumenting imports with $\Delta \tau_{igt}$ in (5.9).¹⁵

Product Elasticity (η)

The elasticity η across products is identified by aggregating variety-specific tariffs to the product level. From (5.3), adding a time subscript and log-differencing over time, we have

$$\Delta \ln s_{Mgt} = \psi_{st} + (1 - \eta) \Delta \ln p_{Mgt} + \varepsilon_{Mgt}, \quad (5.10)$$

where $s_{Mgt} \equiv \frac{p_{Mgt} m_{gt}}{P_{Mst} M_{st}}$ is the import share of product g in sector s . The parameter $\psi_{st} \equiv -(1 - \eta) \Delta \ln (P_{Mst})$ is a sector-time fixed effect that controls for the overall sector import price index, and ε_{Mgt} is a residual that captures the imported product demand shock. The elasticity η can be estimated from a regression of changes in import expenditure shares of product g on sector-time fixed effects and changes in the import price index p_{Mgt} .

We build the import price index from the variety-level data accounting for the entry and exit of varieties by applying the variety correction from Feenstra (1994). Combining (??) and (5.4) we obtain the following exact expression for the change in the product price index:

$$\Delta \ln p_{Mgt} = \frac{1}{1 - \sigma} \ln \left(\sum_{i \in \mathcal{C}_{gt}} s_{igt} e^{(1-\sigma) \Delta \ln (p_{igt}^* (1 + \tau_{igt})) + \Delta \ln a_{igt}} \right) - \frac{1}{1 - \sigma} \ln \left(\frac{S_{g,t+1}(\mathcal{C}_{gt})}{S_{g,t}(\mathcal{C}_{gt})} \right), \quad (5.11)$$

¹⁵Our model assumes flexible prices and abstracts from sticky prices, so we interpret ω^* as the slope of the supply curve.

where s_{igt} is the share of continuing variety i in all continuing varieties, \mathcal{C}_{gt} is the set of continuing imported varieties in product g between t and $t + 1$, and $S_{g,t}(\mathcal{C})$ is the share of the varieties in the set \mathcal{C} in the total imports of product g at time t .¹⁶ The price index includes two pieces from the estimation in the previous step: the estimated σ and the residuals, which reflect mean-zero demand shocks $\Delta \ln(a_{igt})$.

According to our model, the change in the product price index p_{Mgt} is correlated with the unobserved demand shock ε_{Mgt} . Using the same logic applied at the previous stage that tariffs are uncorrelated with demand shocks, we can instrument $\Delta \ln p_{Mgt}$ using the tariffs. Since using value weights may induce mechanical correlations with the left-hand side of equation (5.10), we construct an instrument that is a simple average of changes in tariffs across the continuing varieties:

$$\Delta \ln Z_{Mgt} = \ln \left(\frac{1}{N_{gt}^{\mathcal{C}}} \sum_{i \in \mathcal{C}_{gt}} e^{\Delta \ln(1+\tau_{igt})} \right), \quad (5.12)$$

where $N_{gt}^{\mathcal{C}}$ is the number of continuing varieties in product g between t and $t + 1$.

Import Elasticity (κ)

We further aggregate to the top tier within a sector to estimate the elasticity κ between domestic and imported products within sectors. The import expenditures $P_{Mst}M_{st}$ defined in (5.2), relative to the expenditures in domestically produced goods $P_{Dst}D_{st}$, are a function of the import price index P_{Mst} relative to the price index of domestically produced goods P_{Dst} , defined in equations (??) and (??):

$$\Delta \ln \left(\frac{P_{Mst}M_{st}}{P_{Dst}D_{st}} \right) = \psi_s + \psi_t + (1 - \kappa) \Delta \ln \left(\frac{P_{Mst}}{P_{Dst}} \right) + \varepsilon_{st}. \quad (5.13)$$

The fixed effects and residual components capture demand shocks. We proceed analogously to the previous step to construct the sector import price index, P_{Mst} , and to instrument by aggregating product-level tariff instruments. The import price index of sector s changes according to:

$$\Delta \ln P_{Mst} = \frac{1}{1 - \eta} \ln \left(\sum_{g \in \mathcal{C}_t^s} s_{gt} e^{(1-\eta)\Delta \ln p_{gMt} + \Delta \ln(a_{gMt})} \right) - \frac{1}{1 - \eta} \ln \left(\frac{S_{t+1}^s(\mathcal{C}_t^s)}{S_t^s(\mathcal{C}_t^s)} \right), \quad (5.14)$$

where s_{gt} is the import share of continuing product g in continuing products imported in sector s , $S_t^s(\mathcal{C})$ is the share of the products in the set \mathcal{C} in imports of sector s at time t , and \mathcal{C}_t^s is the set of continuing imported products in sector s between t and $t + 1$.

¹⁶I.e., $s_{igt} \equiv \frac{p_{igt}m_{igt}}{\sum_{i' \in \mathcal{C}_{gt}} p_{i'gt}m_{i'gt}}$ and $S_{g,t}(\mathcal{C}) \equiv \frac{\sum_{i' \in \mathcal{C}} p_{i'gt}m_{i'gt}}{\sum_{i' \in \mathcal{I}} p_{i'gt}m_{i'gt}}$.

We construct $\Delta \ln P_{Mst}$ using the residuals $\varepsilon_{Mgt} = \Delta \ln(a_{gMt})$ estimated from (5.10). We instrument for the relative price of imports, $\Delta \ln(P_{Mst}/P_{Dst})$, using simple averages:

$$\Delta \ln Z_{Mst} \equiv \ln \left(\frac{1}{N_{st}^C} \sum_{g \in \mathcal{C}_t^s} e^{\Delta \ln Z_{gMt}} \right), \quad (5.15)$$

where Z_{Mst} is the instrument defined in (5.12) at the product level and N_{st}^C is the number of continuing products in sector s between t and $t + 1$.

Foreign Import and U.S. Export Variety Elasticities (σ^*, ω)

The foreign import demand is estimated using an analogous equation to (5.8). We consider how U.S. exports respond to retaliatory tariffs. From (5.7), decomposing the log-change of the foreign demand shifter into a product-time effect η_{gt}^x , a country-time effect η_{it}^x , a country-sector effect η_{is}^x , and a residual ε_{igt}^x , we obtain:

$$\Delta \ln x_{igt} = \eta_{gt}^x + \eta_{it}^x + \eta_{is}^x - \sigma^* \Delta \ln((1 + \tau_{igt}^*) p_{igt}^X) + \varepsilon_{igt}^x, \quad (5.16)$$

where p_{igt}^X is the before-duty price observed in the U.S. If the retaliatory tariffs τ_{igt}^* are uncorrelated with foreign import demand shocks ε_{igt}^x , we can identify σ^* by instrumenting the change in the duty-inclusive price, $p_{igt}^{X*} \equiv p_{igt}^X(1 + \tau_{igt}^*)$, with the change in retaliatory tariffs.

We estimate the U.S. variety inverse export supply curve using a specification analogous to (5.9):

$$\Delta \ln p_{igt}^X = \eta_{gt}^p + \eta_{it}^p + \eta_{is}^p + \omega \Delta \ln x_{igt} + \varepsilon_{igt}^p, \quad (5.17)$$

where ω is the inverse export supply elasticity to each destination from the U.S., after controlling for the fixed effects. We instrument for changes in exports with the changes in retaliatory tariffs.

Threats to Identification

There are three main identification threats when using tariffs to estimate the elasticities.

First, the simultaneous identification of demand and supply requires that the tariff affects importers' willingness to pay. If importers can evade the tariff or do not base their demand on duty-inclusive prices, the tariffs will not cause inward shifts of the import demand curve. In our setting, we do not believe either concern is of first order. While sales taxes may not be salient to consumers because retail prices are quoted in before-tax prices (e.g., [Chetty et al. 2009](#)), tariffs are paid at the border and importers observe the after-tariff prices. Tariff evasion is a larger concern in developing countries (e.g., [Sequeira 2016](#)).

Second, as previously mentioned, we require that tariff changes are uncorrelated with unobserved import demand and export supply shocks. The system of equations is estimated in first differences and controls flexibly for unobserved demand and supply shocks at each

step, which mitigates this concern. The event study figures suggest that targeted import and export varieties were not on statistically different trends prior to the war. In the next section, we implement additional checks for pre-trends that support this key identification assumption.

Third, importers may have anticipated looming tariffs in the months before implementation. If they shifted forward their imports, this could bias the elasticities because of a mismatch in the timing of imports and tariff changes.¹⁷ The event study suggests that tariff anticipation is not a concern, and in the next section we implement dynamic specifications that allow lags and leads of tariffs to test formally for anticipation effects.

The identification strategy is not threatened if the tariff changes reflect differences in preferences for redistribution towards specific sectors between the policymakers elected in 2016 and the previous policymakers. Rather, the identification only requires those changes in preferences to be uncorrelated with unobserved shocks to demand and supply over the time period in which the tariff changes take place.

5.4 Estimation

This section addresses threats to identification, presents the elasticity estimates, and examines the robustness of the results.

Pre-existing Trends

To identify the elasticities, tariff changes must be uncorrelated with import demand and export supply shocks. The event studies suggest that targeted varieties were not on different trends prior to the trade war. We further assess concerns about pre-trends by correlating import and export outcomes before the 2018 trade war—values, quantities, unit values, and duty-inclusive unit values—with the subsequent tariff changes.

We compute these outcomes as the average monthly change in 2017 and regress them against the changes in the import tariff rates between 2017 and 2018¹⁸

$$\overline{\Delta \ln y_{ig,2017}} = \alpha_g + \alpha_{is} + \beta \Delta \ln(1 + \tau_{ig}) + \epsilon_{ig}. \quad (5.18)$$

These regressions control for HS-10 product (α_g) and country-sector (α_{is}) fixed effects, since the estimating equations derived in Section 5.3 exploit tariff variation controlling for these fixed effects. Standard errors are clustered by country and HS-8 (for imports) or HS-6 (for exports).

¹⁷Coglianesse et al. (2017) emphasize this point in the context of estimating the demand for gasoline.

¹⁸We examine pre-trends between the start of the Trump Administration in 2017:1 and 2017:12, which pre-dates the first round of the trade war by two months. The Online Appendix reports tests for pre-existing trends over a longer time horizon by correlating average monthly outcomes between 2013:1 and 2017:12 with the tariff changes during the war. There is no evidence that the import tariff and retaliatory changes were biased towards import or export trends over this longer horizon.

The top panel of Table 3 reports the pre-trend tests for imports. We do not observe any statistically significant relationship across import outcomes, suggesting that targeted import varieties were not on differential trends prior to the war. The bottom panel reports the analogous results for U.S. export outcomes and show a similar pattern: pre-war export trends are uncorrelated with retaliatory tariffs.

TABLE 3: TESTS FOR PRE-EXISTING TRENDS

Panel A: U.S. Import Trends				
	(1)	(2)	(3)	(4)
	$\frac{\Delta \ln p_{ig}^* m_{ig}}{\Delta \ln p_{ig}^*}$	$\frac{\Delta \ln m_{ig}}{\Delta \ln p_{ig}^*}$	$\frac{\Delta \ln p_{ig}^*}{\Delta \ln p_{ig}^*}$	$\frac{\Delta \ln p_{ig}}{\Delta \ln p_{ig}^*}$
$\Delta_{17-18} \ln(1 + \tau_{ig})$	0.12 (0.11)	-0.04 (0.19)	0.18 (0.15)	0.18 (0.15)
Country \times Sector FE	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes
R2	0.14	0.14	0.14	0.14
N	180,744	149,173	149,173	149,173
Panel B: U.S. Export Trends				
	(1)	(2)	(3)	(4)
	$\frac{\Delta \ln p_{ig}^X x_{ig}}{\Delta \ln p_{ig}^X}$	$\frac{\Delta \ln x_{ig}}{\Delta \ln p_{ig}^X}$	$\frac{\Delta \ln p_{ig}^X}{\Delta \ln p_{ig}^X}$	$\frac{\Delta \ln p_{ig}^X (1 + \tau_{ig}^*)}{\Delta \ln p_{ig}^X}$
$\Delta_{17-18} \ln(1 + \tau_{ig}^*)$	0.07 (0.06)	0.11 (0.09)	-0.03 (0.07)	-0.03 (0.07)
Country \times Sector FE	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes
R2	0.11	0.12	0.12	0.12
N	207,840	163,181	163,181	163,181

Notes: Table reports pre-trend tests for import (Panel A) and export (Panel B) variety- level trade outcomes. Table reports regressions of the 2017:1-2017:12 average monthly changes in values, quantities, unit values, and tariff-inclusive unit values against the 2018 tariff changes. Standard errors clustered by country and HS-8 (imports) or HS-6 (exports). Significance: *** .01; ** 0.05; * 0.10.

U.S. Imports and Foreign Exports at the Variety Level

This subsection estimates the elasticity of variety import demand and foreign export supply following the approach described in Section 5.3.

Table 4 reports the responses of U.S. imports to the tariff changes. Columns 1-4 of Table 4 report the results of regressing the four outcomes –values (p^*m), quantities (m), unit values (p^*), and duty-inclusive unit values (p)– on the tariffs. Each specification is run in first-differences and includes fixed effects for product-time, country-time and country-sector.

The specification exploits variation in variety-level tariffs over time to identify the elasticities while controlling for seasonality, time-varying country factors (such as exchange rates), and country-sector time trends. Standard errors are two-way clustered by country and HS-8.

TABLE 4: VARIETY IMPORT DEMAND (σ) AND FOREIGN EXPORT SUPPLY (ω^*)

	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln p_{igt}^* m_{igt}$	$\Delta \ln m_{igt}$	$\Delta \ln p_{igt}^*$	$\Delta \ln p_{igt}$	$\Delta \ln p_{igt}^*$	$\Delta \ln m_{igt}$
$\Delta \ln(1 + \tau_{igt})$	-1.52*** (0.18)	-1.47*** (0.24)	0.00 (0.08)	0.58*** (0.13)		
$\Delta \ln m_{igt}$					-0.00 (0.05)	
$\Delta \ln p_{igt}$						-2.53*** (0.26)
Product \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country \times Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
1st-Stage F					36.5	21.2
Bootstrap CI					[,]	[,]
R2	0.13	0.13	0.11	0.11	0.00	.
N	2,993,288	2,454,023	2,454,023	2,454,023	2,454,023	2,454,023

Notes: Table reports the variety-level import responses to import tariffs. Columns 1-4 report import values, quantities, before-duty unit values, and duty-inclusive unit values regressed on the statutory tariff rate. Column 5 reports the foreign export supply curve IV regression, $\hat{\omega}^*$, from equation (5.9); the first stage is column 2. Column 6 reports the import demand curve IV regression, $\hat{\sigma}$, from equation (5.8); the first stage is column 4. All regressions include product-time, country-time and country-sector fixed effects. The coefficient in column (4) is not 1 plus the coefficient in column (3) because the duty-inclusive unit value is constructed using actual duties collected by U.S. customs data. Standard errors clustered by country and HS-8. 90% Bootstrap confidence intervals constructed from 1000 samples. Significance: * 0.10, ** 0.05, *** 0.01. Sample: monthly variety-level import data from 2017:1 to 2019:4.

Column 1 shows that import values drop sharply with tariff increases. Column 2 shows that the decline in import values is closely matched by a commensurate decline in quantities. Column 3 indicates no impact of tariff increases on before-duty unit values. This is the key result providing evidence that the incidence of import tariffs is borne by the U.S. economy, which is consistent with the event study in Figure 1.¹⁹ The reduced-form regressions suggest

¹⁹The elasticity of the duty-inclusive unit value in column 4 is not one plus the coefficient in column 3 because the duty-inclusive value p_{igt} is computed using actual duties collected by U.S. customs rather than imputing from the statutory rate. Columns 1 to 4 of Table ?? report regressions of the variables in columns 2 to 5 on the applied tariff instrumented by the statutory rate. It also reveals complete tariff pass-through. In these regressions, the coefficient on duty-inclusive prices (column 5) is one plus the coefficient on the before-duty price (column 4) and the coefficient on import quantities (column 3) is very close to the estimated σ in column 4 of Table 4.

a complete pass-through of tariffs to duty-inclusive import prices.

We report the variety import demand and foreign export supply elasticities $\{\sigma, \omega^*\}$ using the IV equations in (5.8) and (5.9) in columns 5-6. Column 5 reports the supply curve elasticity $\hat{\omega}^*$; the first stage is column 2. The coefficient is small and imprecisely estimated, $\hat{\omega}^* = -0.002$ (*se* 0.05). This estimate implies that we cannot reject a horizontal supply curve and supports the reduced-form evidence of complete pass-through. Column 6 reports the estimated import demand elasticity σ ; the first stage is column 4. The estimate implies $\hat{\sigma} = 2.53$ (*se* 0.26). The bootstrapped 90% confidence interval, formed from 1000 samples, is [1.75, 3.02]. With these elasticities, using the solution to the system of supply and demand equations (5.8) and (5.9) in columns 5-6, the average change in import values of targeted varieties is:

$$\overline{\Delta \ln(p_{igt}^* m_{igt})} = - \underbrace{\hat{\sigma} \frac{1 + \hat{\omega}^*}{1 + \hat{\omega}^* \hat{\sigma}}}_{2.54} \underbrace{\overline{\Delta \ln(1 + \tau_{igt})}}_{12.5\%} = 31.7\%.$$

Product Level Imports

Table 5 presents estimates of the product elasticity (η) from equation (5.10), following the steps described in Section 5.3. The procedure aggregates the import data to the product-time level, and the regressions are run in first differences controlling for sector-time pair fixed effects, as dictated by the model. Standard errors are clustered at the HS-8 level. We construct the price index from equation (5.11) using $\hat{\sigma} = 2.53$ and the demand shocks from the import variety demand equation from column 6 of Table 4. We build the instrument $\Delta \ln Z_{gMt}$ using (5.12).

Column 1 regresses the change in product shares, $\Delta \ln(s_{Mgt})$, on the instrument $\Delta \ln Z_{Mgt}$ (i.e., the reduced form). Higher product-level tariffs lower the product import share. Column 2 reports the first-stage, a regression of the duty-inclusive product-level price index $\Delta \ln(p_{Mgt})$ on the instrument. The sign is consistent with higher tariffs raising the product price index. Column 3 reports the IV estimate which regresses the change in product shares on the change in the instrumented price index. The estimate implies $\hat{\eta} = 1.53$ (*se* 0.27). The bootstrapped confidence interval for η , which accounts for the variance of $\hat{\sigma}$ and the demand shocks from the lowest tier, is [1.15, 1.89].

The reduction in imports of targeted products implies that imports from untargeted origins did not fully offset import declines from targeted origins. Hence, re-routing of imports or reallocation of producers to untargeted countries does not seem to be a first-order force over the time horizon that we consider. We also find that import tariffs did not lower before-tariff product-level prices. We construct the before-duty product-level price index using (5.11) but excluding duties. This before-duty product-level price index includes a Feenstra variety correction, so it accounts for reallocations towards new source countries. Regressing that index against the tariff instrument $\Delta \ln Z_{gMt}$ with sector-time fixed effects yields a positive coefficient of 0.91 (*se* 0.40).

TABLE 5: PRODUCT ELASTICITY η

	(1)	(2)	(3)
	$\Delta \ln s_{Mgt}$	$\Delta \ln p_{Mgt}$	$\Delta \ln s_{Mgt}$
$\Delta \ln Z_{Mgt}$	-0.81**	1.52***	
	(0.39)	(0.40)	
$\Delta \ln p_{Mgt}$			-0.53*
			(0.27)
Sector-Time FE	Yes	Yes	Yes
1st-Stage F			14.6
$\hat{\eta}$ (se[$\hat{\eta}$])			1.53 (0.27)
Bootstrap CI			[,]
R2	0.01	0.10	.
N	371,916	371,916	371,916

Notes: Table reports the product-level import responses to import tariffs. Column 1 reports the reduced form regression of the imported product's share within sectoral imports, s_{gt} , on the product-level instrument, Z_{gt} . Column 2 reports the first stage: the regression of the product-level import price index P_{gt} on Z_{gt} . Column 3 reports the IV regression with the implied $\hat{\eta}$ and its standard error noted at the bottom of the table in column 3. The product-level import price index is constructed using $\hat{\sigma}$ from column 6 of Table 4 according to (5.11), and the instrument is constructed using the statutory tariffs using equation 5.12. All regressions include sector-time fixed effects. 90% Bootstrap confidence intervals constructed from 1000 samples. Regressions clustered by HS-8. Significance: * 0.10, ** 0.05, *** 0.01. Sample: monthly product-level import data from 2017:1 to 2019:4.

We complement this product-level analysis with a reduced-form approach that does not rely on the CES nesting structure. The Online Appendix regresses the product-time fixed effects from the variety-level regressions on the product-time component of variety-level tariffs. Consistent with the previous results, we find a decline in product-level imports of targeted products, suggesting that re-routing is not an important concern, and no statistically significant decline in the before-duty import price, suggesting complete pass-through at the product level.²⁰

²⁰Flaen et al. (2019) argue that, in response to discriminatory anti-dumping duties of 2012 and 2016, producers of washing machines reallocated products to untargeted countries with lower marginal costs. Our

Using this elasticity and the average change in product-level statutory import tariffs, these estimates imply that import values for targeted products within imported sectors fell, on average, 2.5% across targeted products. This number is the average change in import values for targeted products obtained from $\Delta \ln p_{Mgt} m_{gt} = -(\hat{\eta} - 1) \Delta \ln Z_{gMt}$, where $\hat{\eta} = 1.53$ and $\overline{\Delta \ln Z_{gMt}} = 4.7\%$.

Sector Level Imports

Table 6 reports estimates of the sector elasticity (κ) following the steps described in Section 5.3. The regressions control for sector and time fixed effects, and cluster standard errors at the sector level. As shown in (5.13), estimating this elasticity requires data on changes in imports and domestic expenditures at the sectoral level.

The monthly change in U.S. expenditures in domestically produced goods, $\Delta \ln (P_{Dst} D_{st})$, is not directly observed. We measure it as the difference between the changes in sectoral production and exports. We also need data on the price index of domestically produced goods, $\Delta \ln (P_{Dst})$. We assume that the change in the price index of domestically produced goods equals the change in PPI, $\Delta \ln p_{st}$, plus a mean-zero shock: $\Delta \ln P_{Dst} = \Delta \ln p_{st} + \Delta \ln \varepsilon_{st}^P$.²¹ Then, we can implement equation (5.13) using the PPI instead of the consumer price index of domestically produced goods, which we do not observe. Hence, our specification uses $\Delta \ln (P_{Mst}/p_{st})$ instead of $\Delta \ln (P_{Mst}/P_{Dst})$ in (5.13). The change in the price index, $\Delta \ln P_{Mst}$, is constructed from (5.14) using the estimated $\hat{\sigma}$ and $\hat{\eta}$ from the previous two steps, and the demand shocks constructed from these regressions.

Column 1 is the reduced form specification that projects relative imports on the instrument, column 2 is the first stage and column 3 is the IV estimate. The coefficient is negative, suggesting that price propagation of the tariff through input-output linkages is strong and causes the domestic PPI to increase, but is noisy. The estimate implies a statistically significant $\hat{\kappa} = 1.19$ (*se* 0.49). The bootstrapped confidence interval for $\hat{\kappa}$, which takes into account the estimated $\{\hat{\sigma}, \hat{\eta}\}$ and demand shocks from the previous stages, is [0.89, 1.71].

Using this elasticity and the average change in sector-level statutory import tariffs, these estimates imply that import values for targeted across targeted sectors fell, on average, 0.2%. This number is the average change in import values for targeted sectors obtained from $\Delta \ln \left(\frac{P_{Mst} M_{st}}{P_{Dst} D_{st}} \right) = (1 - \hat{\kappa}) \Delta \ln Z_{Mst}^{stat}$, where $\hat{\kappa} = 1.19$ and $\overline{\Delta \ln Z_{gMt}} = 1.0\%$.

U.S. Exports at the Variety Level

This subsection implements the analysis in Section 5.3. These regressions examine the change in U.S. export outcomes, at the variety level, in response to changes in retaliatory tariffs.

estimate of the elasticity of substitution within products ($\hat{\sigma} = 2.53$) is far from perfect substitution. As we have argued, we also find a decline in the import share of targeted products, and no decline in the before-duty product-level import prices.

²¹This assumption is consistent with the production structure we assume in the general equilibrium model.

The regressions include product-time, destination-time and destination-sector fixed effects, and cluster standard errors by destination country and HS-6.

We first report regressions of the four export outcomes on the retaliatory tariffs in columns 1-4 of Table 7. We observe a statistically significant decline in both export values and quantities. In column 3 we find no evidence that the retaliatory tariffs, on average, caused U.S. exporters to lower (before-duty) product level unit values. Rather, column 4 implies that the duty-inclusive export price rises approximately one-for-one with the tariff.

TABLE 7: FOREIGN IMPORT DEMAND σ^*

	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln p_{igt}^X x_{igt}$	$\Delta \ln x_{igt}$	$\Delta \ln p_{igt}^X$	$\Delta \ln p_{igt}^X (1 + \tau_{igt}^*)$	$\Delta \ln p_{igt}^X$	$\Delta \ln x_{igt}$
$\Delta \ln(1 + \tau_{igt}^*)$	-0.99*** (0.28)	-1.00*** (0.36)	-0.04 (0.16)	0.96*** (0.16)		
$\Delta \ln x_{igt}$					0.04 (0.16)	
$\Delta \ln p_{igt}^X (1 + \tau_{igt}^*)$						-1.04*** (0.32)
Product \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Country \times Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
1st-Stage F					7.8	38.2
Bootstrap CI					[.]	[.]
R2	0.07	0.07	0.06	0.06	.	0.51
N	3,306,766	2,564,731	2,564,731	2,564,731	2,564,731	2,564,731

Notes: Table reports the variety-level export responses to retaliatory tariffs. Columns 1-4 report reduced form regressions of export values, quantities, before-duty unit values, and duty-inclusive unit values on $\Delta \ln(1 + \tau_{igt}^*)$, the change in retaliatory export tariffs. Column 5 reports the IV regression that estimates the U.S. export supply elasticity $\hat{\omega}$; the first stage is column 2. Column 6 reports the IV regression that estimates the foreign import demand elasticity σ^* ; the first stage is column 4. All regressions include product-time, country-time and country-sector fixed effects. Standard errors clustered by country and HS-6. 90% Bootstrap confidence intervals constructed from 1000 samples. Significance: * 0.10, ** 0.05, *** 0.01. Sample: monthly variety-level export data from 2017:1 to 2019:4.

Column 5 reports the IV regression that estimates the U.S. export supply curve at the variety level. This is the analog to the variety level supply curve (5.6) on the export side. The first stage is column 2. The estimate is imprecisely measured and we cannot reject the null that foreigners face a horizontal U.S. export supply curve. Column 6 reports the IV estimate of equation (5.16). The first-stage is column 4. We estimate $\hat{\sigma}^* = 1.04$ (*se* 0.32). The bootstrapped confidence interval is [0.73, 1.39].

Using the estimated elasticity and the average change in retaliatory tariffs, these estimates imply that U.S. export values for varieties targeted by trade partners fell, on average, 9.9%. This number is the average change in export values for targeted varieties obtained from $\Delta \ln(p_{igt}^X x_{igt}) = -\hat{\sigma}^* \Delta \ln(1 + \tau_{igt}^*)$ where $\Delta \ln(1 + \tau_{igt}^*) = 9.5\%$.

Robustness Checks

This section explores the robustness of the results. We first assess concerns that underlying trends or tariff anticipation bias the estimates. We also explore heterogeneity across sectors, compare the pass-through of tariffs to unit values at different time horizons, and examine how the results change with alternative sets of fixed effects.

Trends and Dynamic Specifications

Section 5.4 documents that pre-existing trends and anticipation effects are unlikely to threaten our identification. In this section, we provide further evidence that the elasticities are not sensitive to concerns over pre-existing trends or anticipation effects.

The first robustness check controls for trends through panel fixed effects. We re-estimate the variety-level specifications to include variety fixed effects and report the analog to Table 4 in Table ?? and the analog to Table 7 in Table ?. We assess long-run trends by re-estimating the specifications with variety fixed effects using data from 2013:1 to 2019:4 in Tables ?? and ?. The results are essentially unchanged and remain consistent with the prior evidence that pre-existing trends are unlikely to be confounding factors.

The second concern is that importers may have anticipated the changes in tariffs and shifted their purchasing decisions forward to avoid the duties. This would imply that, even though tariffs have real impacts on trade, regressions identified from contemporaneous changes in tariffs may produce biased elasticities. We check for anticipatory and delayed effects by allowing for leads and lags in variety-level reduced-form regressions:

$$\Delta \ln y_{igt} = \alpha_{gt} + \alpha_{it} + \alpha_{is} + \sum_{m=-6}^{m=6} \beta_m^y [\ln(1 + \tau_{ig,t-m}) - \ln(1 + \tau_{ig,t-1-m})] + \epsilon_{igt}, \quad (5.19)$$

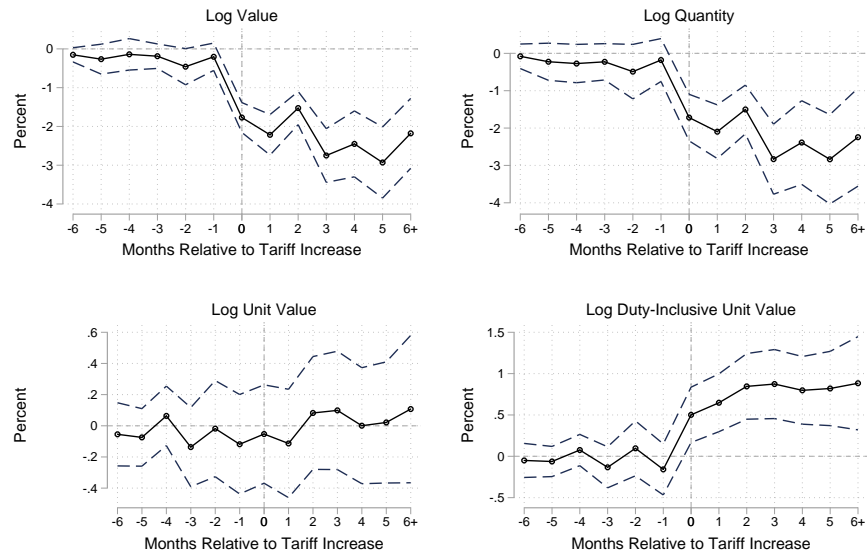
where we allow for leads and lags up to six months before and after the tariff changes.²²

Figure 3 reports the cumulative estimated coefficients for import values, quantities, unit values, and duty-inclusive unit values. The results reveal no quantitatively large patterns of tariff anticipation. Additionally, the results show no evidence of before-duty price declines occurring after the tariffs are implemented. Finally, the cumulative magnitudes displayed in the figure are quantitatively similar to the reduced form estimates from the static regressions. These results reassure us that the elasticities are not biased due to anticipation effects, and that the variety-level pass-through findings.

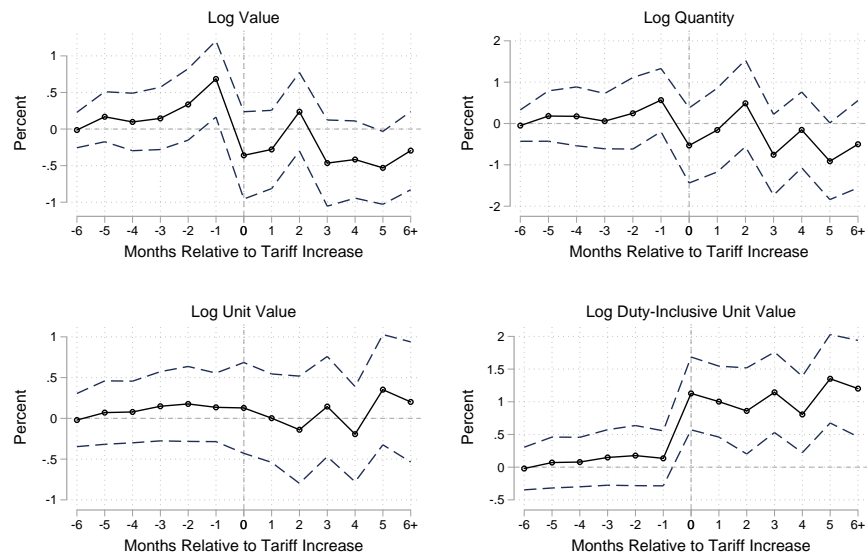
²²Since the dynamic specification requires a balanced panel in event time, we replace missing leading and lagged tariff changes with zeros, and include indicators for those missing values. This is equivalent to assuming that the price of a variety does not change when we do not observe it in the data.

FIGURE 3: DYNAMIC SPECIFICATION

Panel A: Tariffs on U.S. Imports



Panel B: Retaliatory Tariffs on U.S. Exports



Notes: Figures plot cumulative sum of β coefficients from the regression:

$$\Delta \ln y_{igt} = \alpha_{gt} + \alpha_{it} + \alpha_{is} + \sum_{m=-6}^{m=6} \beta_m^y [\ln(1 + \tau_{ig,t-m}) - \ln(1 + \tau_{ig,t-1-m})] + \epsilon_{igt}$$

where i denotes countries, g denotes products, and t denotes time. Standard errors clustered by county and HS-8 for imports, and by country and HS-6 for exports. Error bands show 95% confidence intervals. Monthly variety-level import and export data from U.S. Census. Sample period is 2017:1 to 2019:4.

The bottom panel of Figure 3 reports the results for exported varieties. In this specification, we do find some evidence of tariff anticipation as U.S. export values increase in the month before tariffs change; however, we do not observe this pattern for export quantities. We also do not observe cumulative declines in before-duty export prices occurring as time elapses after the retaliations are implemented.

Heterogeneity

The baseline specifications impose common elasticities across sectors. However, these specifications potentially mask variation in the tariff pass-through result. For example, we may expect more differentiated products or less competitive sectors to exhibit less-than-complete pass-through. We may also expect demand elasticities to depend on inventories and whether goods are durable, which may allow buyers to postpone sales more easily. Other product characteristics, such as variation in price stickiness or stocks of inventories, could also induce heterogeneous pass-through.

In this subsection, we explore this potential heterogeneity by focusing on the complete pass-through finding at the variety level. To do so, we implement reduced-form specifications that regress changes in the import (export) unit values on changes in import (retaliatory) tariffs, controlling for product-time, country-time and country-sector fixed effects. Each specification interacts the tariff change with different measures of product or sector characteristics that have been used widely in the literature.

The Online Appendix reports results from interacting the tariff changes with three different classifications of final versus intermediate goods. The top panel examines the pass-through of the import tariffs to import unit values, and the bottom panel examines pass-through of the retaliatory tariffs to U.S. export unit values. Column 1 uses the Broad Economic Categories (BEC) classification that categorizes sectors according to their end use. Column 2 uses an indicator for whether or not the HS product matches an entry line item in the BLS Consumer Price Index (CPI). Column 3 uses an indicator for whether each HS-10 product description contains the word “part” or “component”. While each classification is imperfect, the results do not show statistical differences between final and intermediate goods.

The Online Appendix examines interactions across 11 different measures of product or sector characteristics that have been used in the literature: 1) quality ladders from [Khandelwal \(2010\)](#); 2) markups estimated by [De Loecker et al. \(2018\)](#); 3) the coefficient of price variation; 4) elasticities of substitution estimated by [Broda and Weinstein \(2006\)](#); 5) trade elasticities estimated by [Caliendo and Parro \(2015\)](#); 6) contract intensity measured by [Nunn \(2007\)](#); 7) inverse frequency of price adjustments from micro-data tabulated by [Nakamura and Steinsson \(2008\)](#) (a higher value indicates less frequent price adjustments); 8) measures of industry upstreamness by [Antràs et al. \(2012\)](#); 9) inventory to sales ratios constructed from Census data; 10) differentiation developed by [Rauch \(1999\)](#); and, 11) an indicator for durable goods using the BEC classifications. The top panel reports the results for imports. The pattern across the 11 different metrics, along with the previous table,

suggests that there is no meaningful heterogeneity in the complete pass-through result, at least with respect to the characteristics we have examined. The Online Appendix reports the results for export unit values. We again find no systematic evidence of heterogeneity with respect to observable characteristics.

Horizons

The variety-level complete pass-through results may be a short-run phenomenon. We assess the incidence of the tariffs at different time horizons by aggregating the data to the 2-month, 3-month, and 4-month level, taking first differences, and re-implementing the reduced-form regression specifications. The results for before-duty import and export unit values are reported in the Online Appendix, with the baseline estimates from the monthly data replicated in column 1 of each panel. Even at these medium-term frequencies, we do not observe downward pressure on before-duty unit values in response to the tariff changes.

Alternative Fixed Effects

Our baseline set of fixed effects — that is, product-time, country-time and country-sector fixed effects — control for potentially confounding import demand and export supply shocks. However, if the trade war induces global or country-specific general equilibrium responses to wages, these fixed effects may mask tariff pass-through effects in our regressions. The Online Appendix reports the elasticity of before-duty unit values against the tariff changes controlling for different sets of fixed effects to explore this possibility. The top panel reports the import results and the bottom panel reports the export results, and column 1 of each panel reports the baseline estimates to facilitate comparisons across the alternate specifications. We do not observe any impact of the tariff on before-duty unit values across 8 different sets of fixed effects, some of which exclude country-time or product-time fixed effects, in both the import and export data. We also extract the country-time α_{it} fixed effect in the baseline specification and regress them against monthly exchange rates, and do not find a statistically significant relationship (estimate is 0.11 (*se* 0.19)). We also do not find a relationship between the country-time fixed effects and exchange rate changes for China (estimate is 0.04 (*se* 0.04)).

TABLE 6: SECTOR ELASTICITY κ

	(1)	(2)	(3)
	$\Delta \ln\left(\frac{P_{Mst}M_{st}}{P_{Dst}D_{st}}\right)$	$\Delta \ln\left(\frac{P_{Mst}}{p_{st}}\right)$	$\Delta \ln\left(\frac{P_{Mst}M_{st}}{P_{Dst}D_{st}}\right)$
$\Delta \ln Z_{Mst}$	0.30 (0.36)	-1.59 (3.49)	
$\Delta \ln\left(\frac{P_{Mst}}{p_{st}}\right)$			-0.19 (0.49)
Sector FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
1st-Stage F			0.2
$\hat{\kappa}$ (se $[\hat{\kappa}]$)			1.19 (0.49)
Bootstrap CI			[.]
R2	0.24	0.67	.
N	2,041	2,041	2,041

Notes: Table reports the sector-level import responses to import tariffs. The sample is at the sector-time level from 2017:1 to 2019:4. Column 1 reports the reduced form regression of the imported sector's share within total sectoral expenditures imports, s_{gt} , on the product-level instrument, Z_{gt} . Column 2 reports the first stage: the regression of the product-level import price index P_{gt} on Z_{gt} . Column 3 reports the IV regression with the implied $\hat{\eta}$ and its standard error noted at the bottom of the table in column 3. The sector import price index is constructed using $\hat{\sigma}$ from column 6 of Table 4 and $\hat{\eta}$ from column 3 of Table 5 according to (5.14), and the instrument is constructed using the statutory tariffs using (5.15). All regressions include sector fixed effects. Regressions clustered by sector. 90% Bootstrap confidence intervals constructed from 1000 samples. Significance: * 0.10, ** 0.05, *** 0.01. Sample: monthly sector-level import data from from 2017:1 to 2019:4.

5.5 Aggregate and Regional Impacts

Before turning to the full model in the next section, it is instructive to perform a few back-of-the-envelope calculations using the estimated parameters to gauge the magnitude of the aggregate impacts of the trade war.

First, given complete tariff pass-through, the first-order approximation to the impact on U.S. consumer surplus is the product of three terms: the import share of value added (15%), the fraction of U.S. imports targeted by tariff increases (13%), and the average import price increase among targeted varieties (14%). This calculation implies buyers of imports lost in aggregate 0.27% of GDP, or \$50.8 billion at an 2016 annual basis.

Second, under some assumptions the elasticities can be used to compute the impact on aggregate real income. Specifically, in the absence of changes in both U.S. import and export prices, starting from free trade, and provided the environment satisfies neoclassical assumptions, the (second-order) approximation to the aggregate equivalent variation is $12(\Delta \mathbf{m})' \Delta \tau$, where $\Delta \mathbf{m}$ is the change in the vector of imports and $\Delta \tau$ is the change in per-unit import tariffs. Using the estimates for the changes in variety-level imports estimated in Section 5.4, and assuming that the fixed-effects in those regressions are unresponsive to the tariffs, this calculation yields a real GDP loss of \$11 billion, or 0.059% of GDP.²³

These approximations are computed assuming complete tariff pass-through. However, our empirical analysis at the variety level does not rule out terms-of-trade effects through changes in prices at the country or sector levels. Also, these calculations do not consider the impacts of retaliatory tariffs. We now combine the previously estimated parameters with a supply side of the U.S. economy. The model imposes upward sloping industry curves and predicts changes in sector level prices in the U.S. due to the demand reallocation induced by import and retaliatory tariffs. Through this channel, it generates additional aggregate and regional effects not captured by the previous measurements.

General Equilibrium Structure

We use a static general equilibrium model of the U.S., imposing strong assumptions such as perfect competition, flexible prices, and an input-output structure with unitary elasticities. We implement counterfactuals that keep constant the wages in foreign countries.

The U.S. is divided into R counties (collected in the set \mathcal{R} and indexed by r). In each region r there are L_r workers. In addition to the traded sectors there is one non-traded sector. Traded sectors are freely traded within the U.S. but face trade costs internationally.²⁴

²³Baqae and Farhi (2019) show that, under these assumptions, this back of the envelope is also the effect on real GDP. In terms of our previous notation, we compute $\frac{1}{2} \sum_s \sum_{g \in \mathcal{G}_s} \sum_i p_{gi}^* m_{gi} \Delta \ln m_{gi} \Delta \tau_{gi}$, where the change in imports of product g from country i is $\Delta \ln m_{gi} = -\hat{\sigma} \Delta \ln(1 + \tau_{gi})$ with $\sigma = 2.53$. Using the error in the estimation of $\hat{\sigma}$, the 90% bootstrap confidence interval around this aggregate loss is [-\$13.1b, -\$7.6b].

²⁴The assumption of free internal trade sidesteps the need to pin down the location of production of HS-10 products within the U.S., for which we do not have data. It also ensures that the aggregate import demand system that we have previously estimated is consistent with the model we use for simulations. Caliendo

Consumption in county r results from maximizing aggregate utility,

$$\beta_{NT} \ln C_{NT,r} + \sum_{s \in \mathcal{S}} \beta_s \ln C_{sr}, \quad (5.20)$$

where $C_{NT,r}$ is consumption of a homogeneous non-traded good, C_{sr} is consumption of tradeable sector s , and the β 's add up to 1. The price of the non-traded good is $P_{NT,r}$ and the price index of sector s is P_s .

Production of tradeable goods in each sector-region uses workers, intermediate inputs, and a fixed factor (capital and structures). Since we are looking at short-run outcomes we assume that the primary factors of production, capital and labor, are immobile across regions and sectors, but intermediate inputs can be freely adjusted. We also consider the implications of perfect labor mobility across sectors.²⁵ The domestic production of tradeable sector s in region r is

$$Q_{sr} = Z_{sr} \left(\frac{I_{sr}}{\alpha_{I,s}} \right)^{\alpha_{I,s}} \left(\frac{L_{sr}}{\alpha_{L,s}} \right)^{\alpha_{L,s}}, \quad (5.21)$$

where Z_{sr} is local productivity, I_{sr} is a bundle of intermediate inputs, and L_{sr} is the number of workers. The production share of the fixed factor is $\alpha_{K,s} \equiv 1 - \alpha_{I,s} - \alpha_{L,s}$. Intermediate inputs in sector s are also aggregated using a Cobb-Douglas technology. We let $\alpha_s^{s'}$ be the share of input s' in total sales of sector s . The cost of the intermediates bundle used by sector s is:

$$\phi_s \propto \prod_{s' \in \mathcal{S}} P_{s'}^{\alpha_s^{s'} / \alpha_{I,s}}. \quad (5.22)$$

The owners of fixed factors choose the quantities I_{sr} and L_{sr} to maximize profits Π_{sr} . Letting p_s be the producer price in tradeable sector s and w_{sr} be the wage per worker in sector s and region r , the returns to the fixed factors are:

$$\Pi_{sr} = \max_{Q_{sr}} p_s Q_{sr} - (1 - \alpha_{K,s}) \left(\frac{\phi_s^{\alpha_{I,s}} w_{sr}^{\alpha_{L,s}}}{Z_{sr}} Q_{sr} \right)^{\frac{1}{1 - \alpha_{K,s}}}, \quad (5.23)$$

giving the supply curve and the national supply in sector s , $Q_s = \sum_{r \in \mathcal{R}} Q_{sr}$. Non-traded output in region r uses labor: $Q_{NT,r} = Z_{NT,r} L_{NT,r}$, where $L_{NT,r}$ is the employment in the non-traded sector in region r .

Production by sector and region, defined above in (5.21), is allocated across products at a constant marginal rate of transformation. Letting q_g be the national output of good g in sector s , the feasibility constraint for products in sector s is:

$$\sum_{g \in \mathcal{G}_s} \frac{q_g}{z_g} = Q_s, \quad (5.24)$$

et al. (2017) combine input-output linkages with internal trade costs in a quantitative analysis of the U.S. economy.

²⁵The system of equilibrium conditions in changes in Online Appendix ?? is defined for both immobile and mobile labor.

where z_g is a product-level productivity shock. We assume this production structure because we observe employment by region at the sector level (NAICS-4 in our data) but not at the product level (HS-10 in our data). The model equilibrium does not pin down where each good g is produced, and this information is not needed to implement counterfactuals.²⁶

Assuming perfect competition, the price of the domestically produced variety of good g is $p_{Dg} = \frac{p_s}{z_g}$. Given iceberg costs δ_{ig} , the price faced by importer country i of product g is $p_{ig}^X = \delta_{ig} p_{Dg}$. Hence, market clearing in the U.S. variety of product g implies

$$q_g = \underbrace{(a_{Dg} D_s) \left(\frac{p_{Dg}}{P_{Ds}} \right)^{-\eta}}_{d_g} + \sum_{i \in \mathcal{I}} \delta_{ig} \underbrace{a_{ig}^* \left((1 + \tau_{ig}^*) p_{ig}^X \right)^{-\sigma^*}}_{x_{ig}}, \quad (5.25)$$

where d_g is the U.S. demand of product g resulting from the CES structure in the Online Appendix, where a_{Dg} is a demand shock, D_s is the aggregate U.S. consumption of domestic goods in sector s defined in (??), and P_{Ds} is the price index of domestically produced goods defined in (??). x_{ig} is the foreign import demand defined in (5.7).

To close the model, we assume that labor income and profits are spent where they are generated. Total tariff revenue R is distributed to each region in proportion b_r equal to its national population share. We allow for aggregate income D derived from ownership of foreign factors, owned by region r also in proportion to its population. By aggregate accounting, D equals the trade deficit. Final consumer expenditures in region r therefore are²⁷

$$X_r = w_{NT,r} L_{NT,r} + \sum_{s \in \mathcal{S}} w_{sr} L_{sr} + \sum_{s \in \mathcal{S}} \Pi_{sr} + b_r (D + R). \quad (5.26)$$

A general equilibrium given tariffs consists of import prices p_{ig}^* , U.S. prices p_{Dg} , traded wages w_{sr} , non-traded wages $w_{NT,r}$, and price indexes $(P_s, P_{Ds}, P_{Ms}, p_{Mg}, \phi_s)$ such that: i) given these prices, final consumers, producers, and workers optimize; ii) local labor markets clear for every sector and region, international markets clear for imports and exports of every variety, and domestic markets for final goods and intermediates clear; and iii) the government budget constraint is satisfied. The foreign demand and supply shifters z_{ig}^* and a_{ig}^* in (5.6) and (5.7) are taken as given.

Implementation

To compute the impacts of the tariffs we derive a system of first-order approximations to the impact of tariff shocks around the pre-war equilibrium. Since the U.S. predominantly

²⁶This product-level supply structure is consistent with the export variety elasticity $\omega = 0$ estimated in Section 5.4.

²⁷We now have an explicit expression for the aggregate demand shifters E_s entering previously in the import demand defined in (5.2): $E_s = \sum_{r \in \mathcal{R}} \beta_s X_r + \sum_{r \in \mathcal{R}} \sum_{s' \in \mathcal{S}} \alpha_{s'}^s p_{s'} Q_{s'r}$. The first term adds up the regional expenditures of final consumers, and the second term adds up the regional expenditures of producers in each sector.

increased tariffs on varieties with initially zero tariffs, we use a higher-order approximation to the change in tariff revenue. The system is fully characterized by equations (??)-(??) in the Online Appendix. In response to a simulated shock to U.S. and foreign tariffs, the system gives the change in every outcome as a function of the elasticities $\{\sigma, \sigma^*, \omega^*, \eta, \kappa\}$ estimated from tariff variation in Section 5.4, the preference and technology parameters $\{\beta_{NT}, \beta_s, \alpha_{L,s}, \alpha_{I,s}, \alpha_s^{s'}\}$, distributions of sales and employment across sectors and counties, and imports and exports across varieties. We obtain the non-estimated parameters and variables from input-output (IO) tables from 2016 (the most recent year before the tariff war for which this information is available), the 2016 County Business Patterns database, and the customs data we used in the estimation. The Online Appendix describes the implementation and parameterization in more detail.²⁸

Impact of Tariffs on U.S. Prices

We now explain the mechanisms through which U.S. and retaliatory tariffs induce price effects in the general-equilibrium model. Since we consider the short run impact of tariffs, we assume no primary factor mobility across sectors and regions. Sector-level quantities only change with intermediate inputs. As a result, the sector-level supply of U.S. goods is upward-sloping with the price. At the sector level, the price of U.S. goods is determined by the intersection between the U.S. supply resulting from (5.23) and its world demand (from both the U.S. and foreign countries) resulting from adding up the right-hand side of (5.25) over all varieties within a sector.

The U.S. experiences a terms-of-trade gain in a sector if the price of products in that sector (some of which are exported) increases compared to the price of its imports. U.S. and foreign tariffs affect these prices by shifting world demand. When the U.S. imposes a tariff on the imports of a particular product from some origin (e.g., wooden kitchen tables from China), U.S. consumers reallocate to the U.S. variety of that product. This reallocation increases the world demand for U.S. production in this sector, and reduces world demand for foreign production. Hence, there is a terms-of-trade gain in the furniture sector. Similarly, when a foreign country imposes tariffs on U.S. varieties, foreign consumers reallocate away from U.S. production, lowering the price in the sector where foreign tariffs are imposed.

The extent of these price changes due to tariffs depends on the elasticities of both U.S. and foreign demands, which we have estimated, and on the the sector-level elasticities of U.S. supply, which we have imposed through the model assumptions and the calibration. The Online Appendix discusses in more detail the determinants of sector-level prices in the general equilibrium model.

The terms-of-trade effects implied by the model operate at the sector level, and are therefore not captured by our previous empirical analysis. Qualitatively, these terms-of-trade effects are corroborated by an analysis of sector-level producer, export, and import price

²⁸Under the “hat algebra” of Dekle et al. (2008), the outcomes depend on endogenous variables in exact relative changes. Our solutions are a special case of Baqaee and Farhi (2019).

indexes published by the Bureau of Labor Statistics. The Online Appendix reports regressions of each price index on a simple average of import and retaliatory tariffs within sector. The table shows that: i) the PPI increases with sector-level import tariffs; ii) U.S. export prices fall with retaliatory tariffs; and iii) there are no impacts of the tariffs on sector-level import prices, which is consistent with the evidence in the previous empirical sections and with our model assumptions.

Aggregate Impacts

We use the model to quantify the impacts of the tariff war. For each primary factor (capital and labor), the equivalent variation is the change in income at initial prices (before the tariff war) that would have left that factor indifferent with the changes in tariffs that took place. Adding up the equivalent variations across all primary factors (capital and labor in each region), we obtain the aggregate equivalent variation EV , or change in aggregate real income. This term can be written as a function of initial trade flows and price and revenue changes (Dixit and Norman, 1980):

$$EV = \underbrace{-\mathbf{m}'\Delta\mathbf{p}^M}_{EV^M} + \underbrace{\mathbf{x}'\Delta\mathbf{p}^X}_{EV^X} + \Delta R \quad (5.27)$$

where \mathbf{m} is a column vector with the imported quantities of each variety before the war, \mathbf{x} collects the quantities exported of each product to each destination, $\Delta\mathbf{p}^M$ are changes in duty-inclusive import prices, and $\Delta\mathbf{p}^X$ are changes in export prices.²⁹ EV^M is the increase in the duty-inclusive cost of the pre-war import basket, EV^X is the increase in the value of the pre-war export basket, and ΔR is the change in tariff revenue. The pre-tariff war levels of imports and exports in (5.27) are directly observed, while the estimated model gives the responses of import and export prices to the simultaneous change in U.S. and retaliatory tariffs.

The top panel of Table 8 shows each of the components of EV in response to the 2018 U.S. and retaliatory tariff waves of the trade war. The first row of each panel reports the monetary equivalent on an annual basis at 2016 prices and the second row reports numbers relative to GDP. The point estimates are calculated using the model elasticities estimated in Section 5.4, $\{\hat{\sigma} = 2.53, \hat{\eta} = 1.53, \hat{\kappa} = 1.19, \hat{\omega}^* = -0.002, \hat{\delta}^* = 1.04\}$, and bootstrapped confidence intervals are computed for each component using the 1,000 bootstrapped parameter estimates.

²⁹In our previous notation, $\mathbf{m}'\Delta\mathbf{p}^M \equiv \sum_{s \in \mathcal{S}} \sum_{g \in \mathcal{G}_s} \sum_{i \in \mathcal{I}} m_{ig} \Delta p_{ig}$ and $\mathbf{x}'\Delta\mathbf{p}^X \equiv \sum_{s \in \mathcal{S}} \sum_{g \in \mathcal{G}_s} \sum_{i \in \mathcal{I}} x_{ig} \Delta p_{ig}^X$.

TABLE 8: AGGREGATE IMPACTS

	EV^M	EV^X	ΔR	EV
	(1)	(2)	(3)	(4)
2018 Trade War				
Change (\$ b)	-51.0	9.4	34.3	-7.2
	[-54.8,-47.2]	[4.1,15.6]	[32.3,36.1]	[-14.4,0.8]
Change (% GDP)	-0.27	0.05	0.18	-0.04
	[-0.29,-0.25]	[0.02,0.08]	[0.17,0.19]	[-0.08,0.00]
2018 U.S. Tariffs and No Retaliation				
Change (\$ b)	-50.9	16.6	34.8	0.5
	[-52.9,-49.0]	[13.2,20.3]	[32.8,36.5]	[-4.0,5.7]
Change (% GDP)	-0.27	0.09	0.19	0.00
	[-0.28,-0.26]	[0.07,0.11]	[0.18,0.20]	[-0.02,0.03]

Notes: Table reports the aggregate impacts in column 4, and the decomposition into EV^M , EV^X , and tariff revenue (ΔR) in columns 1-3. The top panel reports the impacts from the 2018 trade war. The bottom panel simulates a hypothetical scenario where trade partners do not retaliate against U.S. tariffs. The first row in each panel reports the overall impacts of each term in billions of USD. The third row scales by 2016 GDP. These numbers are computed using the model described in Section 5.5 with $\{\hat{\sigma} = 2.53, \hat{\eta} = 1.53, \hat{\kappa} = 1.19, \hat{\omega}^* = -0.00, \hat{\sigma}^* = 1.04\}$. Bootstrapped 90% confidence intervals based on 1000 simulations of the estimated parameters reported in brackets.

The first column, which reports EV^M , shows that U.S. buyers of imports lost in aggregate \$51 billion (0.27% of GDP). Since our estimation finds a foreign supply elasticity ω^* very close to zero, this number remains very close (but not identical since ω^* is not exactly zero) to the number we reported at the beginning of this section. Using the error around our parameter estimates, we can reject the null hypothesis that EV^M is zero.

The second column shows the EV^X component. This second term depends on the export price changes implied by the general equilibrium model. Export prices increase if the reallocation of domestic and foreign demand into U.S. goods induced by tariffs is stronger than the reallocations away from these goods. As discussed in the last subsection, the intensity of these reallocations depend on the combination of the estimated elasticities and the supply-side model assumptions.

We estimate a (statistically significant) increase of EV^X of \$9.4 billion (0.05% of GDP). This aggregate number equals a model-implied 0.7% increase in the export price index times a 7.4% observed share of exports of manufacturing and agricultural sectors in GDP. Since import prices essentially do not change, these export price changes mean terms-of-trade improvements at the country level. The model predicts a 0.1% average (nominal) wage

increase for tradeable-sector workers in the U.S. relative to its trade partners.

The final component of the decomposition is the increase in tariff revenue. The model matches a tariff revenue share of 0.2% of GDP and yields an increase in tariff revenue of \$34.3 billion, or 0.18% of GDP. This increase is larger than the \$29.1 billion increase in actual tariff revenue between 2017 and 2018. It is not exactly the same because the model isolates the revenue increases solely from tariffs (as opposed to other shocks).

These numbers imply large and divergent consequences of the trade war on consumers and producers. However, the effects approximately balance out, leading to a small aggregate loss for the U.S. as a whole. Column 4 sums the three components of EV to obtain the aggregate impacts of the war on the U.S. economy. We estimate an aggregate loss of \$7.2 billion, or 0.04% of GDP. While we cannot reject the null that the aggregate losses are zero, we can conclude that the consumer losses from the trade war were large.³⁰

The second panel reports the aggregate outcomes of a hypothetical scenario where foreign trade partners did not retaliate against the U.S. In this scenario, the export price index would have increased by 1.2% and the aggregate impact would have resulted in a modest gain to the U.S. economy of \$0.5 billion (also not statistically significant). The difference operates through export prices: by lowering demand for U.S. exports, our computations imply a 75.9% larger producer gains without retaliation.

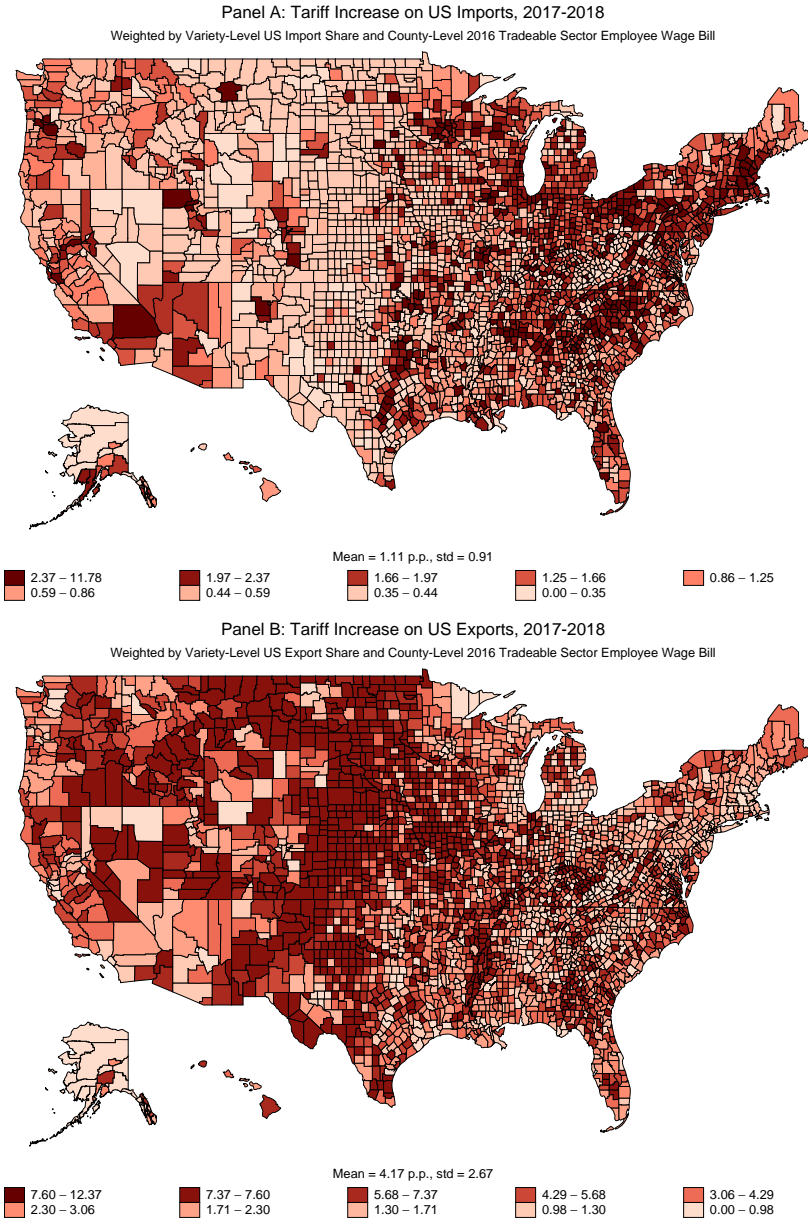
Regional Impacts

We now examine the distributional impacts of the trade war across regions. Tariffs raise the price of consumption for everyone, but also benefit workers in protected sectors through the producer and export price increases we previously discussed. At the same time, tariffs increase the costs of intermediate inputs, which were heavily targeted (see the Online Appendix) and are used more intensively by some regions than others. The ultimate regional impact also depends on the structure of the retaliatory tariffs.

We examine real wages implied by the model. There are three reasons why we do not examine county-level wages directly. First, monthly earnings data are available only at the sector level and for a subset of sectors. Second, even if such data were available, the model would still be necessary to construct the impact of the tariffs on the level of wages. The Online Appendix illustrates that the wage effects are a complex function of shocks in general equilibrium. Third, the model allows us to compare wages under different counterfactual scenarios, such as shutting down foreign retaliations.

³⁰We find similar results assuming mobile labor across sectors. In that case, the overall loss is \$4 billion, with the breakdown for $\{EV^M, EV^X, \Delta R\}$ as $\{\$ - 51b, \$12.7b, \$34.3b\}$.

FIGURE 4: REGIONAL VARIATION IN U.S. AND RETALIATORY TARIFFS



Notes: Figure shows county-level exposure to U.S. import tariff changes (Panel A) and retaliatory tariff changes (Panel B) due to the trade war, weighted by variety-level 2013-2017 U.S. trade shares (constructed from Census data) and by 2016 county-level tradeable sector employee wage bill (constructed from County Business Patterns). Darker shades indicate higher tariff exposure. Values indicate percentage point tariff increases.

Figure 4 illustrates large variation in exposure to the trade war across counties in the U.S. The top panel shows county-level exposure to U.S. tariffs, and the bottom panel

shows county-level exposure to retaliatory tariffs. We construct the county-level exposure of tradeable sectors by first computing the trade-weighted import and retaliatory tariff changes by NAICS sector and then mapping them to counties based on counties' employment structure.³¹ The maps show a clear contrast between the regional structure of U.S. protection and retaliation. The Great Lakes region of the Midwest and the industrial areas of the Northeast received higher tariff protection, while rural regions of the Midwestern plains and Mountain West received higher tariff retaliation.

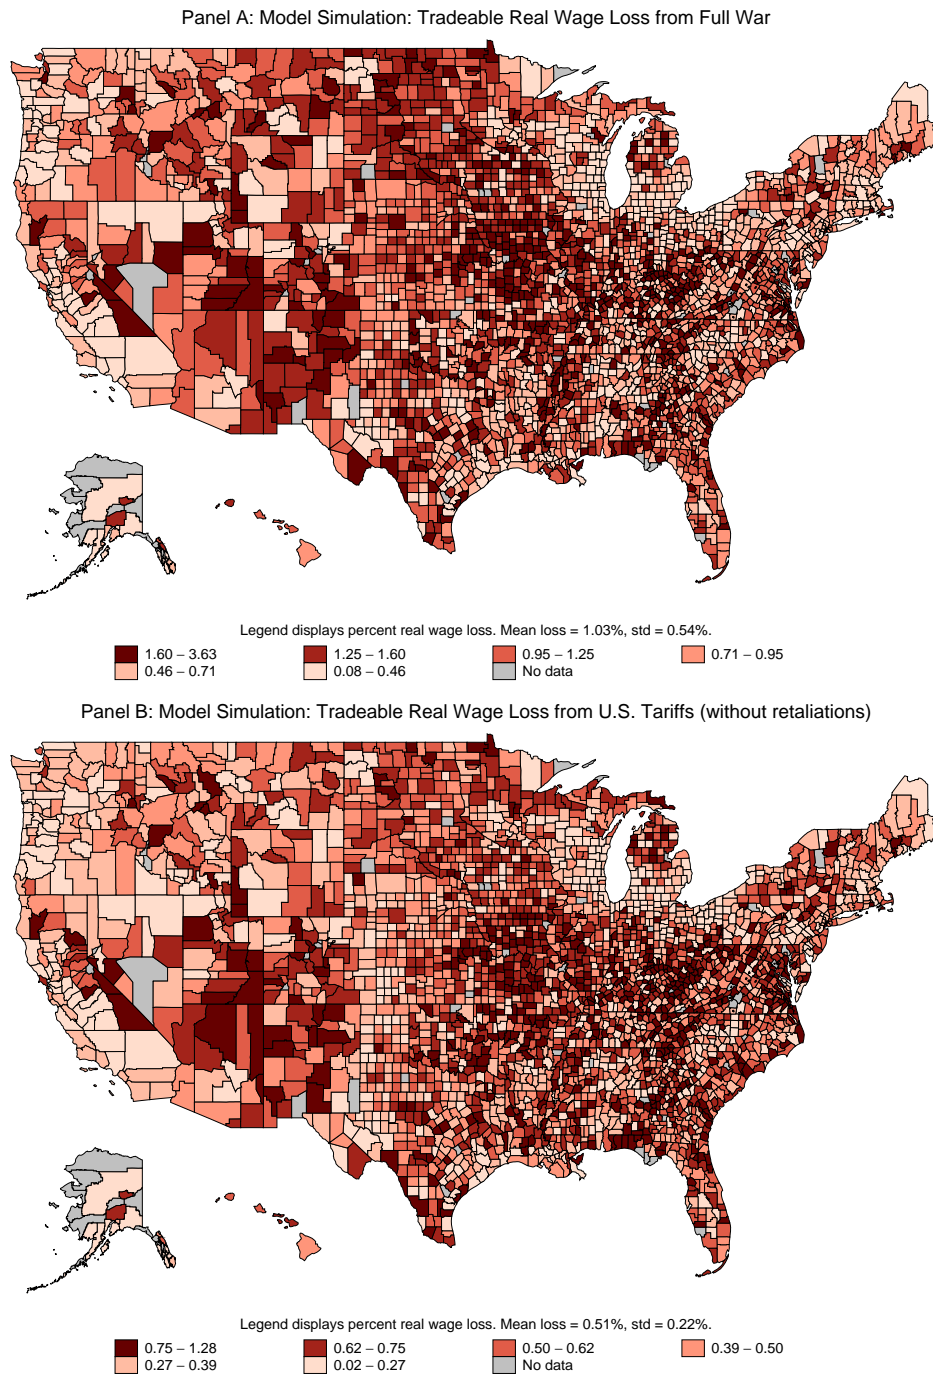
We construct the model-implied impacts across counties in response to the tariff war.³² On average across counties, the nominal wages for workers in tradeable sectors increase by 0.1% (*sd* 0.4%). However, these income gains at initial prices are more than offset by a higher cost of living, as the CPI of tradeable goods increases by 1.1% on average across sectors, partly due to an average 2.0% increase in import prices. As a result, real wages in the tradeable sector fall by 1.0% (*sd* 0.5%), on average. We do not observe a meaningful change in the gini coefficient across counties.

Figure 5 shows the impacts of the trade war across counties. The first map shows the county-level reduction in real wages in tradeable sectors in a hypothetical scenario where U.S. trade partners did not retaliate, and the second map shows real wage losses from the full war. Every county experiences a reduction in the tradeable real wage. Counties with smaller relative losses are concentrated in the Rust Belt region as well as the Southeast. These patterns map imperfectly with the direct protection received through import tariffs shown in Figure 4 because of input-output linkages across sectors. The counties hit hardest by the war are those concentrated in the Midwestern Plains, largely due to the structure of the retaliatory tariffs.

³¹We compute the NAICS-level import and export tariff shock as the import and export-weighted averages of the variety level U.S. and retaliatory tariff changes using average 2013-2016 trade shares. We then construct the county-level import and export tariff shocks as the labor-compensation weighted average of the NAICS-level tariff shocks. In the notation of the model, the import tariff shock (due to U.S. tariffs) is $\Delta\tau_r^j = \sum_{s \in \mathcal{S}} \left(\frac{w_{sr} L_{sr}}{w_r^T L_r^T} \right) \frac{\sum_{g \in \mathcal{G}_s} \sum_{i \in \mathcal{I}} p_{ig}^* m_{ig} \Delta\tau_{ig}}{\sum_{g' \in \mathcal{G}_s} \sum_{i' \in \mathcal{I}} p_{i'g'}^* m_{i'g'}}$ and the export tariff shock (due to retaliatory tariffs) is $\Delta\tau_s^* = \sum_{s \in \mathcal{S}} \left(\frac{w_{sr} L_{sr}}{w_r^T L_r^T} \right) \frac{\sum_{g \in \mathcal{G}_s} \sum_{i \in \mathcal{I}} p_{ig}^x x_{ig} \Delta\tau_{ig}^*}{\sum_{g \in \mathcal{G}_s} \sum_{i \in \mathcal{I}} p_{ig}^x x_{ig}}$, where $w_r^T L_r^T$ are total tradeable sector wages in county r .

³²The real tradeable wage change in region r is defined as $w_{\hat{T},r} - \hat{P}_r$, where $w_{\hat{T},r}$ is the nominal wage increase in the tradeable sector, and where $\hat{P}_r = \beta_{NT} \hat{P}_{NT,r} + \sum_{s \in \mathcal{S}} \beta_s \hat{P}_s$ is the change in the local price index. Equation (??) gives the solution for the wage change as function of price changes. Equations (??) to (??) characterize the block of the model with the solution to the price changes as function tariffs and expenditure shifters.

FIGURE 5: MODEL SIMULATION OF REAL WAGE IMPACTS FROM U.S. AND RETALIATORY TARIFFS



Notes: Figure shows county-level mean tradeable wage losses as simulated from the model. Panel A shows losses accounting for both import and retaliatory tariffs. Panel B shows losses in the counterfactual scenario that U.S. trade partners did not retaliate. Darker shades indicate greater losses. Values indicate percent wage declines.

Tariff Protection, Wages, and Voting Patterns


As discussed in Section 5.2, the pattern of tariff changes across sectors does not *a priori* support the view that protection was driven by incentives to maximize national income or by contributions of special interests. We probe a third hypothesis from the political economy of trade protection literature, namely that policy-makers pursued an electoral strategy when setting tariffs by targeting regions according to their voting patterns. We examine the relationship between the county-level tariff exposure shown in Figure 4 and voting patterns in the 2016 presidential election. The logic of majority voting suggests that tariffs set by an electorally motivated incumbent government should be higher in sectors that are disproportionately located in regions where voters are likely to be pivotal in elections.³³ We then contrast the *ex ante* incentives of policymakers suggested by the relationship between tariffs and voting with the *ex post* distributional consequences of their policies.

Figure 6 presents a non-parametric plot of county-level import and retaliatory tariff changes against the Republican (GOP) vote share, weighted by county population. The county-level tariffs are constructed within tradeables, and therefore do not reflect differences in shares of tradeable activity across counties. The figure reveals two different patterns of protection for U.S. and retaliatory tariffs. For U.S. tariffs, we observe an inverted-U shape, implying that counties with a 40-60% Republican vote share received more protection than heavily Republican or Democratic counties. Hence, U.S. tariffs appear targeted toward sectors concentrated in politically competitive counties. By contrast, trading partners retaliated by targeting exports in sectors concentrated in heavily Republican counties.³⁴ We explore how these targeting patterns vary with other demographic and economic variables in the Online Appendix.

³³Helpman (1995) characterizes optimal tariffs under majority voting in a specific factors model, showing that tariffs are higher in sectors where the median voter has larger factor ownership. Grossman and Helpman (2018) emphasize that psychological benefits to voters from tariff protection (e.g., increased self-esteem from mutually recognized group membership) may underlie a shift to protectionism.

³⁴This finding is also shown by Fetzer and Schwarz (2019).


FIGURE 6: TARIFF CHANGES VS. 2016 REPUBLICAN VOTE SHARE



Notes: Figure plots county-level import and retaliatory tariff changes against the 2016 Republican presidential two-party vote share, using a non-parametric fit weighted by county population. County-level tariff changes weighted by variety-level 2013-2017 U.S. trade shares and by 2016 county-level tradeable sector employee wage bill. Vote shares constructed from Federal Election Commission data. Unit of analysis is 3,111 U.S. counties.

We use the general-equilibrium model to assess if the tradeable real wages of electorally competitive counties indeed experience the largest (relative) gains. Figure 7 plots tradeable real wages against the county Republican vote share for two different scenarios. The black solid curve shows the actual impacts of the war. The dashed curve reflects the impact under a hypothetical scenario where U.S. trade partners did not retaliate. The figure reveals that in the (hypothetical) scenario where trade partners did not retaliate, the impacts would have been fairly even across electorally competitive counties. There is no sharp peak, and the relationship plateaus between a 35% and 50% vote share. Relative to a heavily Democratic county (a 5-15% vote share), the losses in a heavily Republican county (85-95% vote share) are 6% larger.

FIGURE 7: MODEL-SIMULATED TRADEABLE REAL WAGE IMPACT VS. 2016 REPUBLICAN VOTE SHARE



Notes: Figure plots model-simulated county-level tradeable real wage changes due to the trade war against the 2016 Republican presidential two-party vote share, using a non-parametric fit weighted by county population. Vote shares constructed from Federal Election Commission data. Unit of analysis is 3,049 U.S. counties.

The black curve reveals the impacts from the full war. The peak shifts leftward and is more pronounced. The war relatively favored tradeable workers in Democratic-leaning counties with a 2016 Presidential vote share of roughly 35%. Moreover, workers in Republican counties (85-95% vote share) bore the largest cost of the full war.³⁵ The losses in these counties are 32% larger than in a heavily Democratic county (a 5-15% vote share). This asymmetry between Republican and Democratic counties is further illustrated in the Online Appendix, which plots across counties the simulated tradable wage change from the full trade war against the hypothetical scenario where U.S. trade partners did not retaliate. Retaliatory tariffs had a disproportionately negative impact on Republican counties, as is

³⁵ Auer et al. (2018) suggest heavy Republican districts would lose more from revoking NAFTA. Ma and McLaren (2018) provide evidence that tariff changes in the years leading up to NAFTA were biased towards industries located in swing states.

illustrated by the mass of red counties that fall far below the 45-degree line. In contrast, the model implies that Democratic-leaning counties were not as harshly affected by retaliations.

5.6 Conclusion

This paper analyzes the impacts of the 2018 trade war on the U.S. economy. We estimate key elasticities of trade outcomes using import and retaliatory tariff variation. We find large impacts of the war on imports and exports. Before-duties import prices faced by the U.S. did not fall in response to tariffs over the time horizon that we consider, implying complete pass-through of tariffs to duty-inclusive import prices.

These estimates imply an annual loss for the U.S. of \$51 billion due to higher import prices. However, a general equilibrium model imposing neoclassical assumptions implies a small aggregate real income loss of \$7.2 billion. Hence, we find substantial redistribution from buyers of foreign goods to U.S. producers and the government, but a small net effect for the U.S. economy as a whole. We also document that U.S. tariffs protected sectors concentrated in electorally competitive counties, while foreign retaliations affected sectors concentrated in Republican counties. These spatial patterns generate heterogeneous impacts of the trade war, and through model simulations we find that tradeable sectors in heavily GOP counties experienced the largest losses. Therefore, even though the aggregate impacts are small, the distributional effects are substantial.

We close with four important caveats. First, our analysis does not include an analysis of U.S. retail prices paid by final consumers. Second, we do not consider the impacts of trade policy uncertainty on the business climate. Third, our framework does not allow for country-level wage effects in foreign countries that would further impact the terms-of-trade. Finally, our analysis does not examine long-run impacts of the trade war. We believe these are important topics for future research.

Conclusion

This dissertation has studied the economics of several recent important public policies, including business tax policy, place-based policy, housing policy, and trade policy. Using tools from the fields of public finance, labor economics, and trade, this research aims to shed light on the aggregate and distributional effects of taxes and regulations, and to inform researchers, policymakers, and the public about the tradeoffs inherent in raising and redistributing revenue, and in setting policy more broadly.

Bibliography

- Abdulkadiroğlu, A., J. Angrist, Y. Narita, and P. Pathak (2017). Research design meets market design: Using centralized assignment for impact evaluation. *Econometrica* 85(5), 1373–1432.
- Abramitzky, R. and L. Boustan (2017). Immigration in american economic history. *Journal of Economic Literature* 55(4), 1311–1345.
- Acemoglu, D. and V. Guerrieri (2008). Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy* 116(3), 467–498.
- Akcigit, U., J. Grigsby, T. Nicholas, and S. Stantcheva (2021, 06). Taxation and Innovation in the Twentieth Century. *The Quarterly Journal of Economics* 137(1), 329–385.
- Almagro, M. and T. Dominguez-Iino (2021, Mar). Location Sorting and Endogenous Amenities:Evidence from Amsterdam. Technical report.
- Alstadsæter, A., M. Jacob, and R. Michaely (2017). Do Dividend Taxes Affect Corporate Investment? *Journal of Public Economics* 151, 74–83.
- Altig, D., A. Auerbach, P. Higgins, D. Koehler, L. Kotlikoff, E. Terry, and V. Ye (2020). Did the 2017 Tax Reform Discriminate Against Blue-State Voters? *National Tax Journal* 73(4), 1087–1108.
- Altig, D., N. Bloom, S. Davis, B. Meyer, and N. Parker (2018). Are tariff worries cutting into business investment? Federal Reserve of Atlanta Blog.
- Amiti, M. and J. Konings (2007, December). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review* 97(5), 1611–1638.
- Amiti, M., S. Redding, and D. Weinstein (2019). The impact of the 2018 trade war on u.s. prices and welfare. CEPR Discussion Paper 13564.
- Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012). Measuring the upstreamness of production and trade flows. *American Economic Review Papers and Proceedings* 102(3).
- Antràs, P. and A. De Gortari (2017). On the geography of global value chains. Technical Report 23456, National Bureau of Economic Research.

- Arefeva, A., M. A. Davis, A. C. Ghent, and M. Park (2020). Who benefits from place-based policies? job growth from opportunity zones. *Job Growth from Opportunity Zones (July 7, 2020)*.
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *The American Economic Review* 102(1), 94–130.
- Arulampalam, W., M. P. Devereux, and G. Maffini (2012). The Direct Incidence of Corporate Income Tax on Wages. *European Economic Review* 56(6), 1038–1054.
- Asquith, B., E. Mast, and D. Reed (2019, Dec). Supply Shock Versus Demand Shock: The Local Effects of New Housing in Low-Income Areas. [Online; accessed 10. May 2021].
- Atkin, D. and D. Donaldson (2015, July). Who’s getting globalized? the size and implications of intra-national trade costs. Working Paper 21439, National Bureau of Economic Research.
- Atkins, R., P. Hernandez-Lagos, C. Jara-Figueroa, and R. Seamans (2020). What is the impact of opportunity zones on employment outcomes? *Available at SSRN*.
- Attanasio, O., P. K. Goldberg, and N. Pavcnik (2004). Trade reforms and wage inequality in colombia. *Journal of development Economics* 74(2), 331–366.
- Auer, R., B. Bonadio, and A. A. Levchenko (2018). The economics and politics of revoking nafta. Technical Report 25379, National Bureau of Economic Research.
- Auerbach, A. J. (2002). Taxation and Corporate Financial Policy. *Handbook of Public Economics* 3, 1251–1292.
- Auerbach, A. J. (2006). Who Bears the Corporate Tax? A Review of What We Know. *Tax Policy and the Economy* 20, 1–40.
- Auerbach, A. J. (2018). Measuring the Effects of Corporate Tax Cuts. *Journal of Economic Perspectives* 32(4), 97–120.
- Auerbach, A. J. and K. Hassett (1992). Tax Policy and Business Fixed Investment in the United States. *Journal of Public Economics* 47(2), 141–170.
- Auerbach, A. J. and J. Slemrod (1997). The Economic Effects of the Tax Reform Act of 1986. *Journal of Economic Literature* 35(2), 589–632.
- Autor, D., C. Palmer, and P. Pathak (2015, Jul). Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts. *Journal of Political Economy*.
- Autor, D. H., D. Dorn, and G. H. Hanson (2013, October). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review* 103(6), 2121–68.

- Autor, D. H., D. Dorn, G. H. Hanson, and J. Song (2014). Trade adjustment: Worker-level evidence. *The Quarterly Journal of Economics* 129(4), 1799–1860.
- Bachas, P. and M. Soto (2021). Corporate taxation Under Weak Enforcement. *American Economic Journal: Economic Policy* 13(4), 36–71.
- Bagwell, K. and R. W. Staiger (1999). An economic theory of gatt. *American Economic Review* 89(1), 215–248.
- Baker, A., D. Larcker, and C. Wang (2021, Mar). How Much Should We Trust Staggered Difference-In-Differences Estimates? [Online; accessed 11. May 2021].
- Baker, S. R., S. T. Sun, and C. Yannelis (2020). Corporate Taxes and Retail Prices. Technical report, National Bureau of Economic Research.
- Baqaae, D. and E. Farhi (2019). Networks, barriers, and trade. Technical Report 26108, National Bureau of Economic Research.
- Barro, R. J. and J. Furman (2018). Macroeconomic Effects of the 2017 Tax Reform. *Brookings Papers on Economic Activity* 2018(1), 257–345.
- Bebchuk, L. A. and J. M. Fried (2003). Executive Compensation as an Agency Problem. *Journal of Economic Perspectives* 17(3), 71–92.
- Becker, B., M. Jacob, and M. Jacob (2013). Payout Taxes and the Allocation of Investment. *Journal of Financial Economics* 107(1), 1–24.
- Becker, J., C. Fuest, and N. Riedel (2012). Corporate Tax Effects on the Quality and Quantity of FDI. *European Economic Review* 56(8), 1495–1511.
- Bee, A. and J. Rothbaum (2017). Do older americans have more income than we think? Technical report, SESHD Working Paper.
- Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics* 134(2), 647–713.
- Bellora, C. and L. Fontagné (2019). Shooting oneself in the foot? trade war and global value chains.
- Bertrand, M. and S. Mullainathan (2001). Are CEOs Rewarded for Luck? The Ones Without Principals Are. *The Quarterly Journal of Economics* 116(3), 901–932.
- Bloom, N. and J. Van Reenen (2007). Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics* 122(4), 1351–1408.

- Boissel, C. and A. Matray (2022, September). Dividend Taxes and the Allocation of Capital. *American Economic Review* 112(9), 2884–2920.
- Borusyak, K. and P. Hull (2023). Efficient estimation with non-random exposure to exogenous shocks. *Working paper*.
- Borusyak, K., X. Jaravel, and J. Spiess (2021). Revisiting event study designs: Robust and efficient estimation. *Working paper*.
- Boustan, L. P., R. Margo, M. Miller, J. Reeves, and J. Steil (2019, Aug). Does Condominium Development Lead to Gentrification? *NBER*.
- Broda, C., N. Limao, and D. E. Weinstein (2008, December). Optimal tariffs and market power: The evidence. *American Economic Review* 98(5), 2032–65.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *The Quarterly Journal of Economics* 121(2), 541–585.
- Busso, M., J. Gregory, and P. Kline (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review* 103(2), 897–947.
- Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American economic review* 101(1), 304–40.
- Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2015). Tariff reductions, entry, and welfare: Theory and evidence for the last two decades. Technical Report 21768, National Bureau of Economic Research.
- Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies* 82(1), 1–44.
- Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2017). The impact of regional and sectoral productivity changes on the US economy. *The Review of economic studies* 85(4), 2042–2096.
- Callaway, B. and P. H. SantAnna (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- Carbonnier, C., C. Malgouyres, L. Py, and C. Urvoy (2022). Who benefits from tax incentives? The heterogeneous wage incidence of a tax credit. *Journal of Public Economics* 206, 104577.
- Card, D., F. Devicienti, and A. Maida (2014). Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data. *Review of Economic Studies* 81(1), 84–111.
- Carrizosa, R., F. Gaertner, and D. P. Lynch (2020). Debt and Taxes? The Effect of TCJA Interest Limitations on Capital Structure. *Journal of the American Taxation Association*.

- Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2019). Tariff pass-through at the border and at the store: Evidence from us trade policy.
- Cellini, S. R., F. Ferreira, and J. Rothstein (2010). The value of school facility investments: Evidence from a dynamic regression discontinuity design. *The Quarterly Journal of Economics* 125(1), 215–261.
- Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019, Aug). The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics* 134(3), 1405–1454.
- Census Bureau (2019). 2019 SUSB Annual Data Tables by Establishment Industry. Technical report.
- Chandra, A. and J. Skinner (2003). Geography and racial health disparities. Technical report, National bureau of economic research.
- Chen, J., E. L. Glaeser, and D. Wessel (2019). The (non-) effect of opportunity zones on housing prices. Technical report, National Bureau of Economic Research.
- Chetty, R., N. Hendren, and L. F. Katz (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review* 106(4), 855–902.
- Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics* 129(4), 1553–1623.
- Chetty, R., A. Looney, and K. Kroft (2009, September). Salience and taxation: Theory and evidence. *American Economic Review* 99(4), 1145–77.
- Clausing, K. A. (2007). Corporate tax revenues in oecd countries. *International tax and public finance* 14(2), 115–133.
- Clausing, K. A. (2020). Profit Shifting Before and After the Tax Cuts and Jobs Act. *National Tax Journal* 73(4), 1233–1266.
- Cloyne, J., J. Martinez, H. Mumtaz, and P. Surico (2022). Short-Term Tax Cuts, Long-Term Stimulus. Technical report, National Bureau of Economic Research.
- Coglianese, J., L. W. Davis, L. Kilian, and J. H. Stock (2017). Anticipation, tax avoidance, and the price elasticity of gasoline demand. *Journal of Applied Econometrics* 32(1), 1–15.
- Conley, T. G., C. B. Hansen, and P. E. Rossi (2012). Plausibly Exogenous. *Review of Economics and Statistics* 94(1), 260–272.
- Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2019, Aug). Income Growth and the Distributional Effects of Urban Spatial Sorting. *NBER*.

- Couture, V. and J. Handbury (2023, May). Neighborhood change, gentrification, and the urbanization of college graduates. *Journal of Economic Perspectives* 37(2), 29–52.
- Cummins, J. G., K. A. Hassett, and R. G. Hubbard (1994). A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments. *Brookings Papers on Economic Activity* 1994(2), 1–74.
- Cummins, J. G., K. A. Hassett, and R. G. Hubbard (1996). Tax Reforms and Investment: A Cross-Country Comparison. *Journal of Public Economics* 62(1-2), 237–273.
- Curtis, E. M., D. G. Garrett, E. C. Ohrn, K. A. Roberts, and J. C. S. Serrato (2021). Capital Investment and Labor Demand. Technical report, National Bureau of Economic Research.
- De Loecker, J., J. Eeckhout, and G. Unger (2018, August). The rise of market power and the macroeconomic implications. Working Paper 23687, National Bureau of Economic Research.
- De Simone, L., C. McClure, and B. Stomberg (2021). Examining the effects of the tcja on executive compensation. *Kelley School of Business Research Paper* (19-28).
- De Simone, L., C. McClure, and B. Stomberg (2022). Examining the Effects of the Tax Cuts and Jobs Act on Executive Compensation. *Contemporary Accounting Research*.
- Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014). The role of entrepreneurship in us job creation and economic dynamism. *Journal of Economic Perspectives* 28(3), 3–24.
- Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: measuring the burden of adjustment. *IMF Staff Papers* 55(3), 511–540.
- Desai, M. A. and D. Dharmapala (2009). Corporate Tax Avoidance and Firm Value. *The Review of Economics and Statistics* 91(3), 537–546.
- Devereux, M. P. and R. Griffith (2003). Evaluating Tax Policy for Location Decisions. *International Tax and Public Finance* 10(2), 107–126.
- Devereux, M. P., L. Liu, and S. Loretz (2014). The Elasticity of Corporate Taxable Income: New Evidence from UK Tax Records. *American Economic Journal: Economic Policy* 6(2), 19–53.
- Devereux, M. P., B. Lockwood, and M. Redoano (2008). Do Countries Compete Over Corporate Tax Rates? *Journal of Public Economics* 92(5-6), 1210–1235.
- Dharmapala, D., C. F. Foley, and K. J. Forbes (2011). Watch what i do, not what i say: The unintended consequences of the homeland investment act. *The Journal of Finance* 66(3), 753–787.

- Diamond, R. (2016, Mar). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review* 106(3), 479–524.
- Diamond, R., T. McQuade, and F. Qian (2019, Sep). The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco. *American Economic Review* 109(9), 3365–94.
- DiMaggio, P. and F. Garip (2012). Network effects and social inequality. *Annual Review of Sociology* 38, 93–118.
- Dix-Carneiro, R. and B. K. Kovak (2017). Trade liberalization and regional dynamics. *American Economic Review* 107(10), 2908–46.
- Dixit, A. and J. Londregan (1996). The determinants of success of special interests in redistributive politics. *the Journal of Politics* 58(4), 1132–1155.
- Dixit, A. and V. Norman (1980). *Theory of International Trade: A Dual, General Equilibrium Approach*. Cambridge University Press.
- Dobridge, C., P. Landefeld, and J. Mortenson (2019). Corporate taxes and the wage distribution: Effects of the domestic production activities deduction. Technical report, Finance and Economics Discussion Series 2021-081. Board of Governors of the Federal Reserve System.
- Dobridge, C., P. Landefeld, and J. Mortenson (2021). Corporate Taxes and the Earnings Distribution: Effects of the Domestic Production Activities Deduction.
- Dobridge, C., R. Lester, and A. Whitten (2021). Ipos and corporate taxes.
- Dobridge, C. L., P. Kennedy, P. Landefeld, and J. Mortenson (2022). The efficiency-equity tradeoff of the corporate income tax: Evidence from the tax cuts and jobs act. Technical report.
- Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review* 108(4-5), 899–934.
- Dowd, T., C. Giosa, and T. Willingham (2020). Corporate Behavioral Responses to the TCJA for Tax Years 2017-2018. *National Tax Journal* 73(4), 1109–1134.
- Duan, Y. and T. S. Moon (2022). Tax Cuts, Firm Growth, and Worker Earnings: Evidence from Small Businesses in Canada. *Working Paper*.
- Dyreng, S., F. B. Gaertner, J. L. Hoopes, and M. Vernon (2020). The effect of us tax reform on the tax burdens of us domestic and multinational corporations. *Available at SSRN 3620102*.

- Dyreng, S. D., M. Hanlon, E. L. Maydew, and J. R. Thornock (2017). Changes in corporate effective tax rates over the past 25 years. *Journal of Financial Economics* 124(3), 441–463.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The Return to Protectionism. *The Quarterly Journal of Economics* 135(1), 1–55.
- Faulkender, M. and M. Petersen (2012). Investment and capital constraints: Repatriations under the american jobs creation act. *The Review of Financial Studies* 25(11), 3351–3388.
- Feenstra, R. C. (1989). Symmetric pass-through of tariffs and exchange rates under imperfect competition: An empirical test. *Journal of International Economics* 27(1), 25 – 45.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.
- Feldstein, M. (1999). Tax Avoidance and the Deadweight Loss of the Income Tax. *Review of Economics and Statistics* 81(4), 674–680.
- Fetzer, T. and C. Schwarz (2019). Tariffs and politics: Evidence from trump’s trade wars. *CEPR Discussion Paper No. DP13579*.
- Flaen, A. B., A. Hortaçsu, and F. Tintelnot (2019). The production relocation and price effects of us trade policy: The case of washing machines. Technical Report 25767, National Bureau of Economic Research.
- Foertsch, T. (2018). U.s. effective marginal tax rates on new investment under prior law and the tax cuts and jobs act.
- Freedman, M., S. Khanna, and D. Neumark (2021). The impacts of opportunity zones on zone residents. Technical report, National Bureau of Economic Research.
- Freeman, L. (2005). Displacement or succession? residential mobility in gentrifying neighborhoods. *Urban Affairs Review* 40(4), 463–491.
- Freund, C., M. Ferrantino, M. Maliszewska, and M. Ruta (2018). China us trade war scenarios: Impacts on global trade and income. *World Bank Working Paper*.
- Fuest, C., A. Peichl, and S. Siegloch (2018). Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany. *American Economic Review* 108(2), 393–418.
- Furman, J. (2015). Barriers to shared growth: The case of land use regulation and economic rents.

- Gaertner, F. B., D. P. Lynch, and M. E. Vernon (2020). The effects of the tax cuts and jobs act of 2017 on defined benefit pension contributions. *Contemporary Accounting Research* 37(4), 1990–2019.
- Gale, W. G., H. Gelfond, A. Krupkin, M. J. Mazur, and E. J. Toder (2019). Effects of the Tax Cuts and Jobs Act: A Preliminary Analysis. *National Tax Journal* 71(4), 589–612.
- Gale, W. G. and C. Haldeman (2021). The Tax Cuts and Jobs Act: Searching for Supply-Side Effects. *National Tax Journal* 74(4), 895–914.
- Gale, W. G. and S. Thorpe (2022). Rethinking the corporate income tax: The role of rent sharing. Available at SSRN 4093219.
- Garcia-Bernardo, J., P. Janský, and G. Zucman (2022). Did the Tax Cuts and Jobs Act Reduce Profit Shifting by US Multinational Companies? Technical report, National Bureau of Economic Research.
- Garrett, D. G., E. Ohrn, and J. C. Suárez Serrato (2020). Tax policy and Local Labor Market Behavior. *American Economic Review: Insights* 2(1), 83–100.
- Gaubert, C., P. Kline, D. Vergara, and D. Yagan (2021). Trends in US Spatial Inequality: Concentrating Affluence and a Democratization of Poverty. In *AEA Papers and Proceedings*, Volume 111, pp. 520–25.
- Gaubert, C., P. M. Kline, and D. Yagan (2021). Place-based redistribution. Technical report, National Bureau of Economic Research.
- Giroud, X. and J. Rauh (2019). State taxation and the Reallocation of Business Activity: Evidence from Establishment-Level Data. *Journal of Political Economy* 127(3), 1262–1316.
- Glaeser, E. (2017). Reforming land use regulations. *Washington, DC: Brookings Institution*.
- Glaeser, E. L. and J. D. Gottlieb (2009). The wealth of cities: Agglomeration economies and spatial equilibrium in the united states. *Journal of economic literature* 47(4), 983–1028.
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported intermediate inputs and domestic product growth: Evidence from india. *The Quarterly journal of economics* 125(4), 1727–1767.
- Goldberg, P. K. and M. M. Knetter (1997, September). Goods Prices and Exchange Rates: What Have We Learned? *Journal of Economic Literature* 35(3), 1243–1272.
- Goldberg, P. K. and G. Maggi (1999). Protection for sale: An empirical investigation. *American Economic Review* 89(5), 1135–1155.

- Goldberg, P. K. and N. Pavcnik (2016). The effects of trade policy. In *Handbook of commercial policy*, Volume 1, pp. 161–206. Elsevier.
- Goodman, L., K. Lim, B. Sacerdote, and A. Whitten (2021). How Do Business Owners Respond to a Tax Cut? Examining the 199A Deduction for Pass-through Firms. Technical report, National Bureau of Economic Research.
- Goolsbee, A. (1998). Investment Tax Incentives, Prices, and the Supply of Capital Goods. *The Quarterly Journal of Economics* 113(1), 121–148.
- Gopinath, G., O. Itskhoki, and R. Rigobon (2010, March). Currency choice and exchange rate pass-through. *American Economic Review* 100(1), 304–36.
- Gorry, A., G. Hubbard, and A. Mathur (2021). The Elasticity of Taxable Income in the Presence of Intertemporal Income Shifting. *National Tax Journal* 74(1), 45–73.
- Graham, J. R., M. Hanlon, T. Shevlin, and N. Shroff (2014). Incentives for tax planning and avoidance: Evidence from the field. *The Accounting Review* 89(3), 991–1023.
- Greenstone, M., R. Hornbeck, and E. Moretti (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy* 118(3), 536–598.
- Grossman, G. and E. Helpman (2018, December). Identity politics and trade policy. Working Paper 25348, National Bureau of Economic Research.
- Grossman, G. M. and E. Helpman (1994). Protection for sale. *The American Economic Review* 84(4), 833–850.
- Grossman, G. M. and E. Helpman (2005). A protectionist bias in majoritarian politics. *The Quarterly Journal of Economics* 120(4), 1239–1282.
- Gruber, J. and J. Rauh (2007). How Elastic is the Corporate Income Tax Base. *Taxing Corporate Income in the 21st Century*, 140–163.
- Hadjilogiou, S., J. Lutz, and C. Bruno (2021). Highlights from the final opportunity zone regulations. *National Law Review* XI(98).
- Hall, R. E. and D. W. Jorgenson (1967). Tax Policy and Investment Behavior. *The American Economic Review* 57(3), 391–414.
- Handley, K. and N. Limão (2017). Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states. *American Economic Review* 107(9), 2731–83.
- Hanlon, M., J. L. Hoopes, and J. Slemrod (2019). Tax reform made me do it! *Tax Policy and the Economy* 33(1), 33–80.

- Harberger, A. C. (1964). The Measurement of Waste. *The American Economic Review* 54(3), 58–76.
- Hassett, K. A. and R. G. Hubbard (2002). Tax Policy and Business Investment. In *Handbook of Public Economics*, Volume 3, pp. 1293–1343. Elsevier.
- Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, Volume 4, pp. 131–195. Elsevier.
- Helpman, E. (1995). Politics and trade policy. Technical report, National Bureau of Economic Research.
- Hendren, N. and B. Sprung-Keyser (2020). A Unified Welfare Analysis of Government Policies. *The Quarterly Journal of Economics* 135(3), 1209–1318.
- Henry, E., G. A. Plesko, and S. Utke (2018). Tax policy and organizational form: Assessing the effects of the tax cuts and jobs act of 2017. *National Tax Journal* 71(4), 635–660.
- Henry, E. and R. Sansing (2018). Corporate tax avoidance: Data truncation and loss firms. *Review of Accounting Studies* 23(3), 1042–1070.
- Hillberry, R. and D. Hummels (2013). Trade elasticity parameters for a computable general equilibrium model. In *Handbook of computable general equilibrium modeling*, Volume 1, pp. 1213–1269. Elsevier.
- Hines, J. R. and E. M. Rice (1994). Fiscal Paradise: Foreign Tax Havens and American Business. *The Quarterly Journal of Economics* 109(1), 149–182.
- Hirano, K. and G. Imbens (2004). The propensity score with continuous treatments. *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, 73–84.
- Hoopes, J. L., P. Langetieg, E. L. Maydew, and M. Mullaney (2020). Is tax planning best done in private? Available at SSRN 3420362.
- House, C. L. and M. D. Shapiro (2008). Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation. *American Economic Review* 98(3), 737–68.
- Hsieh, C.-T. and E. Moretti (2017). How local housing regulations smother the us economy. *New York Times*.
- Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11(2), 1–39.
- Hudson, S., P. Hull, and J. Liebersohn (2017, Sep). Interpreting Instrumented Difference-in-Differences. Technical report.

- Ichino, S. (1979). Condominium conversions in the bay area. *California Agencies* (Paper 299).
- Imbens, G. (2000). The role of the propensity score in estimating dose-response function. *Biometrika* 87(3), 706–710.
- Imbens, G. and J. Angrist (1994, Mar). Identification and Estimation of Local Average Treatment Effects. *Econometrica* 62(2), 467–475.
- Imbens, G. W. (2020). Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics. *Journal of Economic Literature* 58(4), 1129–79.
- Irwin, D. A. (2014). Tariff incidence: Evidence from us sugar duties, 1890-1930. Technical Report 20635, National Bureau of Economic Research.
- Jackson, M. (2021). Inequality’s economic and social roots: The role of social networks and homophily. Available at SSRN 3795626.
- Jensen, M. C. and K. J. Murphy (1990). Performance Pay and Top-Management Incentives. *Journal of Political Economy* 98(2), 225–264.
- Johnson, H. G. (1953). Optimum tariffs and retaliation. *The Review of Economic Studies* 21(2), 142–153.
- Joint Committee on Taxation (2017). Estimated Budget Effects of the Conference Agreement for H.R.1, the Tax Cuts and Jobs Act. *JCX-67-17*.
- Joint Committee on Taxation (2018a). *Estimated budget effects of the conference agreement for H.R. 1, The "Tax Cuts and Jobs Act."*. Government Printing Office.
- Joint Committee on Taxation (2018b). General Explanation of Public Law 115-97. *JCS-1-18*.
- Joint Committee on Taxation (2018c). *General Explanation of Public Law 115-97*. Government Printing Office.
- Joint Committee on Taxation (2021). Tax Incentives for Domestic Manufacturing. *JCX-15-21*.
- Joint Committee on Taxation (2022). Linking Entity Tax Returns and Wage Filings. *JCX-5-22*.
- Juarez, L. (2022). Buyer Market Power and Exchange Rate Pass-Through. *Working Paper*.
- Kennedy, P. and H. Wheeler (2022). Neighborhood-Level Investment from the U.S. Opportunity Zone Program: Early Evidence. *Working Paper*.

- Kerr, W. (1963). Condominium-statutory implementation. *St. John's Law Review* 38(1).
- Khandelwal, A. K. (2010). The long and short (of) quality ladders. *Review of Economic Studies* 77(4), 1450–1476.
- Kitchen, J. and M. Knittel (2016). Business Use of Section 179 Expensing and Bonus Depreciation, 2002–2014. *Office of Tax Analysis Working Paper 110*, 46.
- Kline, P. and E. Moretti (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly journal of economics* 129(1), 275–331.
- Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who Profits from Patents? Rent-Sharing at Innovative Firms. *The Quarterly Journal of Economics* 134(3), 1343–1404.
- Kotlikoff, L. J. and L. H. Summers (1987). Tax incidence. In *Handbook of Public Economics*, Volume 2, pp. 1043–1092. Elsevier.
- Kovak, B. K. (2013, August). Regional effects of trade reform: What is the correct measure of liberalization? *American Economic Review* 103(5), 1960–76.
- Kubick, T. R., G. B. Lockhart, and J. R. Robinson (2021). Taxes and earnings management: Evidence from the tax cuts and jobs act of 2017. *Available at SSRN 3891629*.
- Kumar, A. (2019). Did Tax Cuts and Jobs Act Create Jobs and Stimulate Growth? Early Evidence Using State-Level Variation in Tax Changes. *Early Evidence Using State-Level Variation in Tax Changes (November 15, 2019)*.
- Larrimore, J., J. Mortenson, and D. Splinter (2019). Household Incomes in Tax Data: Using Addresses to Move from Tax Unit to Household Income Distributions. *Journal of Human Resources*, 0718–9647R1.
- Larrimore, J., J. Mortenson, and D. Splinter (2020). Presence and persistence of poverty in us tax data. Technical report, National Bureau of Economic Research.
- Lee, Y. and R. H. Gordon (2005). Tax Structure and Economic Growth. *Journal of Public Economics* 89(5-6), 1027–1043.
- Lees, L., T. Slater, and E. Wyly (2013). *Gentrification*. Routledge.
- Link, S., M. Menkhoff, A. Peichl, and P. Schüle (2022). Downward Revision of Investment Decisions after Corporate Tax Hikes.
- Lisowsky, P. and M. Minnis (2020). The silent majority: Private us firms and financial reporting choices. *Journal of Accounting Research* 58(3), 547–588.

- Ma, X. and J. McLaren (2018). A swing-state theorem, with evidence. Technical report, National Bureau of Economic Research.
- Mayer, W. (1984). Endogenous tariff formation. *American Economic Review* 74(5), 970–85.
- McCaig, B. and N. Pavcnik (2018, July). Export markets and labor allocation in a low-income country. *American Economic Review* 108(7), 1899–1941.
- Mertens, K. (2018). The Near Term Growth Impact of the Tax Cuts and Jobs Act.
- Mills, L. F. and K. J. Newberry (2001). The influence of tax and nontax costs on book-tax reporting differences: Public and private firms. *Journal of the American Taxation Association* 23(1), 1–19.
- Moon, T. S. (2022). Capital Gains Taxes and Real Corporate Investment: Evidence from Korea. *American Economic Review* 112(8), 2669–2700.
- Moretti, E. (2012). *The new geography of jobs*. Houghton Mifflin Harcourt.
- Murphy, K. J. (1999). Executive Compensation. *Handbook of Labor Economics* 3, 2485–2563.
- Nakamura, E. and J. Steinsson (2008, 11). Five Facts about Prices: A Reevaluation of Menu Cost Models*. *The Quarterly Journal of Economics* 123(4), 1415–1464.
- Neumark, D. and H. Simpson (2015). Place-based policies. In *Handbook of regional and urban economics*, Volume 5, pp. 1197–1287. Elsevier.
- Nunn, N. (2007, 05). Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade*. *The Quarterly Journal of Economics* 122(2), 569–600.
- Ohrn, E. (2018a). The effect of corporate taxation on investment and financial policy: Evidence from the dpad. *American Economic Journal: Economic Policy* 10(2), 272–301.
- Ohrn, E. (2018b). The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD. *American Economic Journal: Economic Policy* 10(2), 272–301.
- Ohrn, E. (2022). Corporate Tax Breaks and Executive Compensation. *American Economic Journal: Economic Policy*.
- Ossa, R. (2014). Trade wars and trade talks with data. *American Economic Review* 104(12), 4104–46.
- Ossa, R. (2016). Quantitative models of commercial policy. In *Handbook of commercial policy*, Volume 1, pp. 207–259. Elsevier.
- Patel, E. S., N. Seegert, and M. G. Smith (2017). At a Loss: The Real and Reporting Elasticity of Corporate Taxable Income. *Comparative Political Economy: Fiscal Policy eJournal*.

- Pennington, K. (2021, Apr). Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco. Technical report.
- Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics* 48(1), 7 – 35.
- Reardon, S. F. and K. Bischoff (2011). Income inequality and income segregation. *American journal of sociology* 116(4), 1092–1153.
- Risch, M. (2021). Does Taxing Business Owners Affect Employees? Evidence from a Change in the Top Marginal Tax Rate. *Working Paper*.
- Romalis, J. (2007). Nafta’s and cusfta’s impact on international trade. *The Review of Economics and Statistics* 89(3), 416–435.
- Romer, C. D. and D. H. Romer (2010, June). The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review* 100(3), 763–801.
- Rothstein, R. (2017). *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing.
- Rubin, R. and T. Francis (2020, Jan). Did the u.s. tax overhaul do what it promised? *Wall Street Journal*.
- Rubin, R. and T. Francis (2021, Sep). In house tax bill, companies get return of higher rates but not the breaks. *Wall Street Journal*.
- Saez, E., B. Schoefer, and D. Seim (2019). Payroll Taxes, Firm behavior, and Rent Sharing: Evidence from a Young Workers’ Tax Cut in Sweden. *American Economic Review* 109(5), 1717–63.
- Saez, E., J. Slemrod, and S. H. Giertz (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature* 50(1), 3–50.
- San Francisco Board of Supervisors (2004). Ordinance no. 281-04.
- San Francisco Board of Supervisors (2005). Ordinance no. 281-05.
- Sequeira, S. (2016, October). Corruption, trade costs, and gains from tariff liberalization: Evidence from southern africa. *American Economic Review* 106(10), 3029–63.
- Simonovska, I. and M. E. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics* 92(1), 34–50.
- Slemrod, J. (1995). Income Creation or Income Shifting? Behavioral Responses to the Tax Reform Act of 1986. *The American Economic Review* 85(2), 175–180.

- Slemrod, J. (2004). Are corporate tax rates, or countries, converging? *Journal of Public Economics* 88(6), 1169–1186.
- Smith, M., D. Yagan, O. Zidar, and E. Zwick (2019). Capitalists in the Twenty-first Century. *The Quarterly Journal of Economics* 134(4), 1675–1745.
- Spearot, A. (2016). Unpacking the long-run effects of tariff shocks: New structural implications from firm heterogeneity models. *American Economic Journal: Microeconomics* 8(2), 128–67.
- Spearot, A. C. (2013). Variable Demand Elasticities and Tariff Liberalization. *Journal of International Economics* 89(1), 26–41.
- Suárez Serrato, J. C. and O. Zidar (2016). Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms. *American Economic Review* 106(9), 2582–2624.
- Sun, L. and S. Abraham (2020, Dec). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.
- Tax Foundation (2021). Corporate Tax Rates Around the World, 2021. Technical report.
- The Board of Governors of the Federal Reserve System (2018). Distributional Financial Accounts. Technical report.
- The Joint Committee on Taxation (2020, November). Estimates of federal tax expenditures for fiscal years 2020-2024. Technical report.
- Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from india. *American Economic Journal: Applied Economics* 2(4), 1–41.
- Wagner, A. F., R. J. Zeckhauser, and A. Ziegler (2020). The tax cuts and jobs act: Which firms won? which lost? Technical report, National Bureau of Economic Research.
- Weinstein, I. and S. Glickman (2020). The guide to making opportunity zones work. Technical report, CohnReznick.
- Wier, L. S. and G. Zucman (2022). Global profit shifting, 1975-2019. Technical report, National Bureau of Economic Research.
- Winkelmann, L. and R. Winkelmann (1998). Tariffs, quotas and terms-of-trade: The case of new zealand. *Journal of International Economics* 46(2), 313–332.
- Yagan, D. (2015). Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut. *American Economic Review* 105(12), 3531–63.

- Zidar, O. (2019). Tax Cuts For Whom? Heterogeneous Effects of Income Tax Changes on Growth and Employment. *Journal of Political Economy* 127(3), 1437–1472.
- Zoutman, F. T., E. Gavrilova, and A. O. Hopland (2018). Estimating both supply and demand elasticities using variation in a single tax rate. *Econometrica* 86(2), 763–771.
- Zukin, S. (1987). Gentrification: culture and capital in the urban core. *Annual review of sociology* 13(1), 129–147.
- Zwick, E. and J. Mahon (2017). Tax Policy and Heterogeneous Investment Behavior. *American Economic Review* 107(1), 217–48.