Essays in the Economics and Law of Taxation

by

Eleanor R. Wilking

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2018

Doctoral Committee:

Professor Joel B. Slemrod, Chair
Professor Charles C. Brown
Professor James R. Hines, Jr.
Professor J.J. Prescott

Eleanor R. Wilking

ewilking@umich.edu

ORCID iD: 0000-0001-7234-2681

DEDICATION

This dissertation is dedicated to three individuals, each of whom have had an enormous influence over my career and life, albeit in vastly different ways. My late grandfather, Lawrence A. Leclerc, did not live to see me complete this work, but his unshakable—at times, blatantly partisan—faith in my abilities helped ensure that I would finish. In addition to shaping how I think about economics, my advisor Joel B. Slemrod has modeled a rare combination of integrity and humor to which I continually aspire. And Max Kapustin, my partner, has been an unwavering pillar of support. Our relationship was forged, in part, in the crucible of this degree. He has shared in the victories and the occasional, sometimes devastating, defeats. His seemingly infinite patience, advice and encouragement were, and continue to be, indispensable.

ACKNOWLEDGMENTS

I regret that so few of those who aided in the completion of this project may be named below.

COMMITTEE

I am grateful to my committee, Prof. Joel B. Slemrod, Prof. James R. Hines, Jr., Prof. J.J. Prescott, and Prof. Charlie Brown, for providing valuable insights and comments on my research. This project could not have succeeded without their vision, guidance, and insistence.

I owe particular thanks to Jim for both asking me the big questions and always believing that I could (eventually) answer them. J.J. showed me how I might straddle the divide between law and economics and contribute meaningfully to both disciplines. Charlie's comments, full of intricate and complex meaning, are always delivered with cheery aplomb.

A very generous thank you must go to Joel for tolerating my bizarre labor supply patterns (at some point, he generously, if not subtly, suggested that we begin meeting in the afternoon). Notably, in his time at UM, Joel has chaired no fewer than 20 committees for female students—I am the 21st—and mentored countless others. In my view, few have done more to advance women in public finance.

COLLABORATORS

I have been immensely fortunate to work with several astute collaborators. Teju Velayudhan conceived the research question in Chapter 3 and has been my partner in its execution. We are joined by Yeliz Kacamak. In addition to sharing her mastery of microeconomic theory, she is a generous friend and even endured rooming with me for this final year. I also thank Max Risch and Katie Lim for joining me, not only on Chapter 2, but in the broader investigation of the tax implications of contract work. Special thanks go to Alicia Miller, an invaluable liaison with the IRS, for her collaboration on Chapter 2. Her persistence allowed us to acquire essential data and resources.

COLLEAGUES

I have also benefitted from the feedback of several colleagues. Will Boning, Brett Collins, Ed Fox, Jacob Goldin, John Guyton, “Steve” Hou Fei, Sarah Johnston, Johannes Norling, Kyle Rozema, Danny Schaffa, Alex Turk, Xiaoqing Zhou and attendees of the public finance seminars at UM for their comments. In addition to his feedback, Reuven Avi-Yonah at UM Law, also provided keen professional advice.

I am especially beholden to Charisse Willis, my first friend at Michigan. Much of the text in this dissertation has benefitted from her beautiful—and lightning-fast—editing.

OTHER THANKS

I thank the exemplary and caring UM staff: Laura Flak and Lauren Pulay in the Dept. of Economics and Mary Ceccanese at Ross/OTPR. As director of graduate studies, Prof. Linda Tesar provided both advice and logistical aide in navigating the requirements of both programs.

I gratefully acknowledge financial support from the Rackham Merit Fellowship, the Office of Tax Policy Research, the Michigan Institute for Teaching and Research in Economics, and the National Science Foundation.

Finally, I thank my family, writ large, for their love and encouragement: Eliot and Virginia Wilking, Steven Leclerc, Sam and Sofia Kapustin, and the late Kathryn C. and Lawrence A. Leclerc. Special thanks are owed to my parents, Martha K. Leclerc and Leo F.J. Wilking, and my partner, Max Kapustin.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

FIGURE

ABSTRACT

Firms are critical components in the success of virtually all modern tax systems. This dissertation uses theoretical and empirical methods to investigate firm behavioral responses to changes in tax law that affect firms' substantive tax burden, their role as the government's intermediary, or both. Such responses can change key tax system parameters, including the tax burden and government revenues through a change in compliance with the law or, in some cases, a change in the economic fundamentals like consumption or labor demand.

In Chapters 1 and 3, I explore this premise in the context of consumption taxes. Chapter 2 considers the income tax context.

Two themes emerge from these projects. First, the nature of firms as entities in the tax system is changing, and these changes may affect the traditional role of firms as the locus of information for economic activity. Chapters 1 and 3 consider the phenomenon of "platform" firms, which shift income generation or consumption from traditional employers or retailers to individuals. Additionally, information on these taxable transactions is concentrated in the firm facilitating the transaction. Chapters 1 and 3 explore how this combination affords the platform firm significant leverage in negotiating the terms of its cooperation with local tax authorities in remitting taxes. In Chapter 2, I document the expansion of labor income generated outside of the traditional employer relationship, which also has a potentially disruptive effect as the government does not have the resources to intensively audit individual business deductions.

Second, collectively, these papers highlight the relevance of statutory features—reporting requirements, remittance and withholding obligations, size-based exemptions—which, historically, have not featured prominently in the optimal taxation or incidence literatures in public finance. In both Chapters 1 and 3, I find that shifting the remittance duty affects the absolute tax burden of consumers. In Chapter 2, I find evidence that a firm size-based exemption distorts the decisions of firms to locate, or at least to report their location, close to the threshold number of employees. These findings join those in an increasingly populated field of papers that recognize the importance of legal and regulatory dimensions for the design of modern tax systems.

Chapter 1. Tax Incidence with Heterogeneous Firm Evasion: Evidence from Airbnb Remittance Agreements

How does assignment of the remittance obligation affect consumption tax incidence? In classical tax theory, the responsibility of transferring tax revenue has no effect on which party bears the economic burden of a consumption tax. I explore this prediction in the context of agreements between city governments and a large digital platform firm that shifted the obligation to remit hotel taxes from independent renters to the platform firm itself. Using variation in the location and timing of such agreements, I identify a substantial increase in advertised tax-inclusive rental prices—a violation of remittance invariance—but comparatively modest declines in completed reservations. A contemporaneous increase in hotel tax revenue collections suggests that the policy was an effective tax increase assessed on previously non-compliant renters. I explore heterogeneity in pass-through using several proxies for host price-setting sophistication. Pass-through of the effective tax increase was lowest among full-space, frequent renters who likely faced smaller optimization frictions relative to more amateur renters. My results indicate that shifting the remittance obligation to the platform increases after-tax prices and raises revenue, suggesting that consumers bear a greater share of the tax burden when the remittance obligation is shifted to a party with fewer evasion opportunities.

Chapter 2. Independent Contractor or Employee? The Changing Relationship Between Firms and Their Workforce and Potential Consequences for the U.S. Income Tax

The number of U.S. workers classified as independent contractors has risen dramatically over the last two decades. While this trend, in part, reflects technological changes in how work is carried out, some of the increase may also reflect firms and workers taking advantage of the legal ambiguity between classifications to obtain preferential tax treatment or to avoid complying with regulations. To study this phenomenon, we exploit a sharp discontinuity in the marginal cost of hiring an employee, created by the 2010 Affordable Care Act. We use U.S. tax returns from 1997-2015 to link firm filings to associated employees and independent contractors. Here, we find preliminary evidence that firms substitute independent contractors for employees in order to qualify for size-based regulatory exemptions. We discuss empirical strategies to distinguish whether this substitution reflects re-organization in the production process (a real response) or misclassification (an evasion response).

Chapter 3. Does the Elasticity of the Sales Tax Base Depend on Enforcement? Evidence from U.S. states' Voluntary Collection Agreements

In addition to taxpayer preferences, elasticity of taxable income has been shown to depend on parameters of the tax system—including the costs and expected penalty of tax evasion and the costs of tax avoidance. However, less is known about how consumption elasticities change in response to enforcement. The theory of statutory neutrality predicts that structuring as a 'use tax'—where the consumer remits—or, as a 'sales tax'—under which the retailer remits—should have no effect on the fundamental parameters of the tax system. We test this in the context of U.S. state restructuring of the remittance regime governing online sales shipped to state residents. Using detailed purchase data from the Nielsen Consumer Panel and monthly, zip-code level information on local sales tax rates, we find that consumers reduce their online expenditure in response to Voluntary Collection Agreements (VCA). However, we do not find evidence of a large change in elasticity of the tax base with respect to tax changes. We conclude that shifting the remittance duty to the party with fewer evasion opportunities, akin to an enforcement increase, could affect the responsiveness of the tax base to future tax rate changes, but that the effect of the enforcement on online retailers is too small to measure.

CHAPTER I

Tax Incidence with Heterogeneous Firm Evasion: Evidence from Airbnb Remittance Agreements

How does assignment of the remittance obligation affect consumption tax incidence? In classical tax theory, the responsibility of transferring tax revenue has no effect on which party bears the economic burden of a consumption tax. I explore this prediction in the context of agreements between city governments and a large digital platform firm that shifted the obligation to remit hotel taxes from independent renters to the platform firm itself. Using variation in the location and timing of such agreements, I identify a substantial increase in advertised tax-inclusive rental prices-a violation of remittance invariance-but comparatively modest declines in completed reservations. A contemporaneous increase in hotel tax revenue collections suggests that the policy was an effective tax increase, assessed on previously non-compliant renters. I explore heterogeneity in pass-through using several proxies for host price-setting sophistication. Pass-through of the effective tax increase was lowest among full-space, frequent renters who likely

---

Thanks to Joel Slemrod, Jim Hines, Charlie Brown, and J.J. Prescott for their consistent support and feedback. I also thank Max Kapustin, Sarah Johnston, Yeliz Kacamak, Xiaoqing Zhou, Daniel Reck and Anne Brockmeyer. In addition to participants at the Michigan public finance lunch seminar, I thank participants at the following conferences for their feedback: Conference of Empirical Legal Studies at Duke University (Nov. 2016, Durham, NC); Midwestern Law and Economics Association Conference at Marquette (Nov. 2017, Milwaukee, WI); National Tax Association, Annual Conference on Taxation (Nov. 2016, Baltimore, MD); Conference on Business Taxation at Oxford University (Sep. 2016, Oxford U.K.); IIPF Annual Congress (Aug. 2016, Reno, NV); Mannheim University/IZA Doctoral Conference (Jun. 2016, Mannheim, Germany)

faced smaller optimization frictions relative to more amateur renters. My results indicate that shifting the remittance obligation to the platform increases after-tax prices and raises revenue, suggesting that consumers bear a greater share of the tax burden when the remittance obligation is shifted to a party with fewer evasion opportunities.

1.1 Introduction

A fundamental tenet of classical tax theory is that the economic burden of a tax is independent of which side of the market *remits,* or transfers tax monies to the government (Myles 1989; Weyl and Farbinger 2013). Tax incidence, so the theory goes, depends only on the relative demand and supply elasticities: the less elastic party bears more of the tax.

However, recent studies have identified at least two circumstances in which who remits affects incidence in practice. First, who remits matters when one side of the market has access to differential evasion opportunities (Kopczuk 2009; Goolsbee 2012). Second, who remits matters when agents face optimization frictions (Chetty 2009; Finkelstein 2009). For example, a consumer who wholly ignores the presence of sales tax or incorrectly calculates the tax-inclusive price of an item will bear a larger share of a tax increase, on average, than a consumer who perceives the tax and performs this calculation correctly. In situations where, on average, consumers and suppliers face different optimization frictions, shifting the remittance duty from consumers to suppliers, or vice versa, will affect tax incidence.

I extend this literature by studying a context in which both optimization frictions and differential evasion opportunities are likely present. In a simple model, I show how these characteristics separately and jointly affect equilibrium prices. To study this phenomenon empirically, I exploit plausibly exogenous variation in the timing of bilateral remittance agreements, called Voluntary Collection Agreements (VCAs), between Airbnb and city governments in the United States. These VCAs shifted the responsibility to remit hotel taxes from individual suppliers to the Airbnb platform itself. I conclude that which side of the market remits can have an economically meaningful effect on equilibrium prices, tax collections, and the characteristics of market entrants.

The paper makes two contributions. First, I show that shifting the remittance duty substantially increased after-tax prices and that this effect likely stemmed from the elimination of a differential evasion opportunity available to suppliers. Intuitively, suppliers that previously evaded the tax adjust their pre-tax price downward by less than the amount of the tax in response to the policy, passing some or all of the tax on to consumers. In contrast, suppliers that previously complied with the tax will respond to it by lowering their pre-tax price by the amount of the tax. This practice, as classical tax theory predicts will happen when switching the remittance obligation, leaves consumer prices unchanged.

To identify the effect VCA adoption has on consumer prices, I employ two complementary estimation techniques that rely on separate identifying assumptions. First, I exploit variation in the timing and location—both across and within metropolitan areas—of VCA adoption to estimate a triple difference specification. The identifying assumption is that—prior to the policy—consumer prices in treated cities were moving in parallel with respect to two sets of controls: metropolitan areas that did not adopt VCAs and neighboring jurisdictions within metropolitan areas that did not adopt VCAs

Second, I take advantage of detailed data on the locations of listings to estimate a geographic regression discontinuity (RD) design, comparing those listings just within the municipal border of a VCA adopting city to listings just outside that border. Reassuringly, I find similar estimates using both methods. On average, for each one percentage point of the local hotel tax rate, the price paid by consumers rises by approximately 0.9 percent. Using the same sources of variation in timing and location, I find that hotel tax collections increase in proportion to the size of the Airbnb market before the policy, which I interpret as circumstantial evidence that failure to remit was widespread.

This result complements the main finding in Kopczuk, Marion, Muehlegger, and Slemrod (2016) which states that the economic incidence of a quantity tax on diesel fuel depends on the point of collection within the supply chain. As the remittance obligation moves "up" the chain from retailers to distributors and prime suppliers, the pass-through rate of diesel taxes to the retail price increases, as do tax revenues. This suggests that differential evasion opportunities afforded to these agents explain the relevance of a tax's collection point.

My second contribution is to provide evidence suggesting that the effect of VCA adoption may be heterogeneous with respect to suppliers' attentiveness to the policy and the existence of hotel taxes. Although some suppliers may have purposely chosen not to comply with the tax prior to the adoption of the VCAs, other suppliers may have been unaware of the hotel taxes' existence or their obligation to remit them. I therefore model supplier behavior as being characterized by their "attentiveness" and also allow for the possibility that inattentive hosts are not only less informed about the policy environment but may systematically err in their demand forecasts as well, a hypothesis for which I find empirical support.

I document heterogeneity in the effect of VCA adoption on consumer prices by several supplier characteristics, including responsiveness to local demand shocks, experience, and concentration of competitors. I do this by re-estimating event study and difference-in-difference models while interacting the policy variable with characteristics of suppliers and their surroundings. For example, I find that a one percentage point increase in the correlation between a host's prices and those of local hotels—a proxy for price-setting sophistication—results in a 0.2 percent reduction on the overall increase in consumer prices following adoption of the VCA.

One interpretation of this finding is that attention to local demand conditions and attention to the tax regime are related, and that, in consequence, pass-through in markets with inattentive, or amateur suppliers may be different than in markets with traditional firms. Although there are various studies which already suggest that consumers face optimization frictions that affect their responsiveness to changes in tax rate or tax administration (e.g., Chetty et. al 2010; Goldin and Homonoff 2012; Homonoff 2016; Lockwood 2017), there is comparatively little evidence on whether similar optimization frictions also affect suppliers.

Relatedly, while tax incidence is traditionally exclusively determined by market-level factors such as the level of competition and supply and demand elasticities (see, e.g., Myles 1989; Weyl and Fabinger 2013), there is some empirical evidence that differences in firms' characteristics, such as managerial resources that affect price-setting strategies, can lead to variation in tax incidence within a market where some firms have market power. For example, small, independent firms are more likely to rely on simplified pricing rules, such as round-number heuristics, and may not fully incorporate tax changes into price-setting behavior (Harju, Kosonen, and Skans 2015). My empirical findings suggest that more sophisticated hosts pass on less of the tax burden resulting from the elimination of evasion opportunity, lending support to this hypothesis.

A caveat is warranted. Price changes provide direct evidence of the increased cost of maintaining consumption after the policy, and can also provide indirect insight into underlying market functions (Kopczuk et al. 2012; Stolper 2016). However, I proceed with caution in inferring that the policy changed the tax incidence, at least incidence in the way that it is conventionally defined—as the ratio of reductions of total consumer and producer surplus that results from imposition of a tax. I interpret my results to suggest that the tax burden on guests in these cities rose—they are now paying higher prices for identical products. Yet, if some hosts were previously evading their remittance obligation, as seems likely, their absolute tax burden rose as well (from zero). Therefore, the policy changed the tax incidence in the sense that it shifted the burden of the tax from taxpayers to the suppliers and consumers, rather than the relative burden of the tax as shared between consumers and suppliers.

Another limitation of my approach is that it relies in the main on data from a single platform firm, in a single industry. While I duly acknowledge that this inherently limits generalizability of my estimates, I maintain that two key features of this context expand the project beyond a case study: first, despite obvious difficulty in valuation, hosts have full autonomy in price-setting, and second, a large contingent of hosts on Airbnb are amateurs—lots of amateurs in the market characterizes other emerging platform or "market-maker" driven markets.

For clarity, I explicitly define key terms employed throughout the paper as follows. I understand *pass-through*—distinct from incidence—as the degree to which tax exclusive prices adjust to shift the economic burden of the tax to non-remitting parties to the taxable transaction. As is standard, I express pass-through as a percentage calibrated to the total tax liability.[1] I refer to individual suppliers, who list their property on the Airbnb as *hosts*, and consumers or short-term renters as *guests*, in keeping with Airbnb's nomenclature. I define an *amateur* (host) as a host who is a casual participant in the rental market—they did not secure their property interest

[1] However, I refer to pass through of the policy as if it constituted a new tax (rather than being partially constituted by a change in compliance costs).

for the purpose of short-term rental, and they lack the price setting acumen that accrues to professionals through intensive rental activity or centralized price-setting resources.

The remainder of this paper is organized as follows: Section 2 lays out a theoretical framework of the effects of shifting remittance duty which I use to motivate and interpret the empirical findings. Section 3 provides background on Airbnb rental markets and the natural experiment afforded by cities' VCA negotiations, while Section 4 introduces the data and characteristics of the sample. The next three sections explore empirical claims corresponding to the predictions of the model. Section 5 studies the effect of the remittance shift on tax-inclusive prices and collection of tax revenue. Section 6 considers supplier heterogeneity in pass-through and Section 7 asks whether the policy impacted market exit decisions. Section 8 summarizes the empirical results and discusses their implications for ongoing academic and policy dialogues about tax system design. Section 9 concludes.

1.2 Conceptual Framework

This section sets forth a simplified, partial equilibrium model of supplier behavior to develop intuition for how Airbnb's policy of remitting hotel taxes on behalf of consumers changes the distribution of prices and supplier composition in equilibrium. It also offers several predictions for how these changes differ based on the level of pre-policy compliance (e.g., how many hosts were remitting taxes voluntarily) and hosts' price setting sophistication.

In this model, hosts differ along two dimensions: honesty and attention. Honest hosts remit in full any known tax obligations; dishonest hosts remit nothing, or some fraction of their true liability. Hosts also vary in innate attention, which affects their price setting in two ways: inattentive hosts are less perceptive of demand for their listing—leading to errors in price-setting—and, in addition, are unaware of their remittance obligation, inhibiting optimal response when the remittance regime changes.

Comparative static analysis yields three predictions. First, in the absence of evasion, the policy will not affect tax inclusive prices. Second, in the presence of evasion, more attentive hosts pass through less of the effective tax increase. Finally, on the extensive margin, the policy may induce some non-compliant hosts to exit the market.

Like most economic models, this model makes several assumptions that are unlikely to hold in reality. I assume the prices of other goods do not enter explicitly into the host's price-setting. However, in a general equilibrium model with a representative consumer who optimizes with respect to all goods in the economy, this choice does not *necessarily* preclude other hosts' prices from having an *indirect* effect on the host's price-setting through the parameters of her demand function. To see this, consider a typical linear inverse demand function $p_i = a - bq_i$. If the relative price of the monopolist's good increases—say, because other goods that consumers purchase become less expensive—this will be expressed through a downward shift of the entire demand function (i.e., a decrease in the value of a, the demand curve intercept). Thus, changes in relative prices may shift the value of a, which in turn would affect the monopolist's optimal price. However, in my one period, partial equilibrium model, the parameter values of an individual host's demand function are static (i.e., the values of a and b are fixed).[2]

[2] What if, instead, Airbnb listings were imperfect substitutes for one another (i.e., moving from a strict monopolist to a model of imperfect competition)? Imagine that there are two hosts with imperfectly substitutable listings who differ along a single dimension: honesty. In a Nash-in-prices equilibrium, the optimal price of each listing takes the other listing's price as an argument. The policy triggers a series of strategic price interactions that shift the market to the new equilibrium. In contrast to the monopolist set-up, the post-policy difference in price response by type is muted. The honest host will lower her listing price, but not by the full amount of the tax, because her optimal price is a function of the dishonest host's price response. The basic logic of this simplified two host model flows through to the current model with three types of hosts. As in the example, the price responses of the attentive-honest and attentive-dishonest would be attenuated. The price response of inattentive hosts will depend on the extent to which the inattentive host accurately perceives how the policy changes the best response function of other hosts. Either the inattentive host is oblivious to the price changes of other hosts, which would appear to assume the result, or the

A related limitation is that the guest's decision of which listing to choose given prices and availability is omitted from the model. While it may be the case that demand for temporary lodging within a city is relatively inelastic in the short run, it is almost certainly more elastic in the long run. If this is the case, then the incidence of the hotel tax will be increasingly borne by hosts in the long run, as guests have their choice of visiting cities with and without remittance agreements.

1.2.1 Set-up

Each Airbnb host supplies a listing made unique by its location and amenities from those offered by competitors. For simplicity, each host i is assumed to be a monopolist facing a downward-sloping demand q_i in price p_i for her listing.[3] The price paid by guests is either the one set by the host if hosts are expected to remit ($\theta = 0$), or the one set by the host plus a specific tax t if guests are expected to remit ($\theta = 1$).[4] More succinctly, guests face a price of $p_i^c = p_i + t\theta$.

In the host's optimization, however, what the guest is *perceived to pay* is a function of the host's attentiveness, $a_i \in \{0,1\}$. An attentive host ($a_i = 1$) correctly perceives the price paid by

---

inattentive host perceives that her competitor's prices have changed, and responds by lowering her listing price—which (a) muddles the meaning of "inattentive" and (b) is not in line with the empirical finding that a sizeable number of hosts do not adjust at all to the policy. More generally, the strategic price interactions required for an imperfectly competitive setting seem incongruent with the limited experience and resources that characterize price-setting for the average Airbnb host.

[3] Though not included here for concision, I also consider a model of host price-setting under monopolistic competition, where demand, $q_i(p_i, p_{-i})$, is declining in the host's own price $\left(\frac{\partial q_i}{\partial p_i} < 0\right)$ and increasing in the prices set by all other hosts $\left(\frac{\partial q_i}{\partial p_{-i}} > 0\right)$. Assuming an equilibrium that is symmetric Nash in prices, the main results survive in sign, though the strategic price-setting attenuates their magnitude. My results also hold for a simpler model that incorporates limited strategic price-setting, in which a host's demand is partially dependent on the average price, over which an individual host has negligible influence. The results of this model fall between those of monopoly and monopolistic competition, however they impart limited additional intuition while greatly complicating exposition.

[4] For simplicity, this model assumes a two-sided market (omitting the platform firm as a potential third party to the transaction). In reality, Airbnb assumes the obligation of remitting on behalf of guests.

guests ($p_i + t\theta$), while an inattentive host ($a_i = 0$) is both unaware that guests might be required to remit a tax *and* misestimates demand due to the inclusion of a noise term ε_i. Therefore, hosts optimize with respect to a guest price of $p_i^c = p_i + ta_i\theta + (1 - a_i)\varepsilon_i$.

Each host chooses a price to maximize her *perceived* profit:

$$\pi_i^P = [p_i - ta_i(1-\theta)(1-\gamma_i) - c_i]q_i(p_i + ta_i\theta + (1-a_i)\varepsilon_i) - F \tag{1}$$

where c_i is an exogenous, host-specific marginal cost, and F is a uniform fixed cost. In contrast, a host's *actual* profit is:

$$\pi_i = [p_i - ta_i(1-\theta)(1-\gamma_i) - c_i]q_i(p_i + t\theta) - F \tag{2}$$

Relative to attentive hosts, the perceived profit of inattentive hosts differs in two ways that affect price-setting. First, as described earlier, inattentive hosts fail to account for the prices actually paid by guests, due to being unaware of guests' possible remittance obligations and due to demand forecast error. The decision to relate hosts' attentiveness to the tax environment and their ability to accurately forecast demand is intuitive: a lack of resources or experience could potentially explain both.

Second, inattentive hosts are unable to evade their remittance obligations because they are unaware that such obligations exist. In contrast, attentive hosts are further characterized along an additional dimension of heterogeneity: honesty. An attentive host who is honest ($\gamma_i = 0$) will remit the tax in full, while an attentive host who is dishonest ($\gamma_i = 1$) will not.[5] Notice that this distinction is irrelevant for hosts who are inattentive to the policy ($a_i = 0$). As a result, hosts can fail to comply with the tax in two ways: either through conscious non-compliance

[5] The purpose of this model is to generate comparative statics regarding price-setting and entry behavior by hosts in response to a change in remittance policy, rather than to predict changes in compliance behavior. As a result, evasion is modeled as a byproduct of hosts' exogenously determined honesty, rather than an endogenous choice, and no cost to evading is included in the model (nor does it appear, from anecdotal evidence, that many hosts who failed to comply were detected and punished).

(“evasion”), or by being unwitting non-compliers. However, neither γ_i nor a_i determines compliance with the tax when guests are obligated to remit ($\theta = 1$).[6]

1.2.2 Host's Decision

In effect, the model above describes three types of hosts: (1) attentive and honest, (2) attentive and dishonest, and (3) inattentive. All hosts choose an optimal price determined by a vector of host-specific and general parameters: $p_i^* = p_i^*(c_i, t, a_i, \theta, \gamma_i, \varepsilon_i)$. Actual profit is then determined by the fixed and variable costs, the parameters governing the tax policy, the host's attentiveness and honesty, and the quantity demanded given the true after-tax price faced by guests:

$$\pi_i^* = [p_i^* - ta_i(1-\theta)(1-\gamma_i) - c_i]q_i(p_i^* + t\theta) - F \tag{3}$$

Prior to the policy change, profit is weakly increasing in the level of honesty: inattentive hosts cannot consciously evade, while attentive hosts earn rents from being dishonest ($\frac{d\pi_i^*}{d\gamma_i} > 0$ if $a_i = 1$).[7] However, the relationship between profit and attentiveness is less straightforward. Profits are increasing in attentiveness for hosts who are dishonest, but may be decreasing for compliant hosts. Recall that inattentiveness is positively correlated with a host's tendency to misestimate the demand curve she faces. If the degree to which an inattentive host misperceives her demand is small relative to the cost of complying with her tax obligations once made aware of them, then being more attentive can actually lower profit.

---

[6] It is assumed that when guests are obligated to remit, they are unable to evade their obligation because Airbnb performs it on their behalf.

[7] Formally: $\frac{d\pi_i^*}{d\gamma_i} = \left(\frac{dp_i^*}{d\gamma_i} + a_i t(1-\theta)\right) q_i(p_i^* + t\theta) + (p_i^* - ta_i(1-\theta)(1-\gamma_i) - c_i)\frac{\partial q_i^*}{\partial p_i^*}\frac{dp_i^*}{d\gamma_i}$, where $-a_i t < \frac{dp_i^*}{d\gamma_i} < 0$ as monopolists shift only part of the perceived tax savings in evading to consumers.

1.2.3 The Policy

In response to a shift in remittance obligation from hosts to guests ($\theta \rightarrow 1$), the change in a host's optimal price will differ by her attentiveness and honesty.

Proposition 1: If a host is attentive and honest (previously complied with the tax), then shifting the remittance obligation from hosts to guests has no effect on the final prices faced by guests.[8]

A host who was attentive and honest ($a_i = 1, \gamma_i = 0$) prior to the policy change will lower her price by the full amount of the tax, $\frac{dp_i^*}{d\theta} = -t$, after the policy change goes into effect, thereby ensuring that the price faced by guests remains unchanged, $\frac{dp_i^c}{d\theta} = 0$. This is the classic result that incidence is independent of which side of the market remits a tax.

Proposition 2: If a host is attentive and dishonest or inattentive (previously non-compliant with the tax), then shifting the remittance obligation from hosts to guests increases the final prices faced by guests.

A host who was attentive and dishonest ($a_i = 1, \gamma_i = 1$) prior to the policy change will decrease her price by less than the full amount of the tax, $\frac{dp_i^*}{d\theta} > -t$, after the policy change goes into effect, thereby raising the price faced by guests, $\frac{dp_i^c}{d\theta} > 0$. Because this host was fully non-compliant with the tax previously, the policy acts as an effective tax increase by shifting the remittance obligation to a side of the market that is unable to evade it. As an effective tax increase, the host will respond by lowering her price by an amount reflective of guests' demand elasticity.

[8] This proposition depends on the market structure assumptions that I make in section 2.1: namely, that hosts are monopolists whose demand curves are independent of prices set by other hosts.

Finally, if a host was fully non-compliant with the tax previously due to being inattentive ($a_i = 0$), then she will be unaware of the policy change and therefore not adjust her price, $\frac{dp_i^*}{d\theta} = 0$. If a host does not adjust her price downward, then the policy, which mechanically adds the amount of the tax at check-out for guests, will increase guest prices by the full amount of the tax, $\frac{dp_i^c}{d\theta} = t$. Note that this is not equivalent to the incidence of the tax being fully borne by guests; failing to adjust her price downward will result in lost profits for the host.

Proposition 3. Marginal hosts who are inattentive or dishonest will exit the market.

Attentive and honest hosts ($a_i = 1, \gamma_i = 0$) in the market prior to the policy change will remain in the market after the policy change: after adjusting their prices to reflect the change in remittance obligation, their profits will remain unchanged (see Proposition 1). On the other hand, attentive and dishonest hosts ($a_i = 1, \gamma_i = 1$) and inattentive hosts ($a_i = 0$), who both previously did not comply with the tax, will see their profit decrease from what is an effective tax increase. Attentive and dishonest hosts, being aware of the policy change, will reduce prices accordingly (see Proposition 2); in doing so, they will no longer collect evasion "rents,"[9] and their profit will decrease: $\frac{d\pi_i^*}{d\theta} = \frac{dp_i^*}{d\theta} q_i(p_i^* + t\theta) + (p_i^* - c_i)\frac{dq}{dp}\left(\frac{\partial p_i^*}{\partial \theta} + t\right) < 0$, because $-t < \frac{dp_i^*}{d\theta} < 0$ and $\frac{dq}{dp} < 0$.[10] Inattentive hosts are, by definition, unaware that a policy change occurred and as a result do not adjust their prices, implicitly passing the tax through fully to guests who reduce the quantity they demand and, ultimately, hosts' profit: $\frac{d\pi_i^*}{d\theta} = t(p_i^* - c_i)\frac{dq}{dp} < 0$, because $\frac{dp_i^*}{d\theta} = 0$. From

[9] Evasion "rents" refer to the surplus captured by hosts from knowingly failing to remit their tax obligations. In this context, only attentive hosts can enjoy evasion rents. For examples of this terminology, see KMMS and X.

[10] Formally: $\frac{d\pi_i^*}{d\theta} = \left(\frac{dp_i^*}{d\theta} + ta_i(1-\gamma_i)\right) q_i(p_i^* + t\theta) + (p_i^* - c_i - ta_i(1-\theta)(1-\gamma_i))\frac{dq}{dp}\left(\frac{\partial p_i^*}{\partial \theta} + t\right)$.

among these hosts, some will have been only marginally profitable, and the policy change will cause their profits to decrease, forcing some to exit.

1.3 Background

This section provides three types of background information relevant for subsequent analysis. First, I describe characteristics of the emerging, platform-driven, short-term rental market, and Airbnb specifically. Next, I discuss a timeline of the Airbnb VCAs, which provide the plausibly exogenous policy variation needed for analysis. Finally, I provide details of Airbnb's implementation of agreements, including how and when the tax was displayed during booking in Airbnb's interface.

1.3.1 What is Airbnb?

Airbnb is the largest of several firms facilitating short-term, peer-to-peer residential space rentals through an online platform. Originally conceived as an online marketplace to connect couch surfers, Airbnb has experienced remarkable growth in recent years, expanding exponentially in popular tourism cities around the globe.[11] Hosts on Airbnb create listings for each of their properties. Each listing includes information about the space's characteristics, such as the number of beds, kitchen availability, and whether it is a private apartment or a shared space. Hosts can designate a listing's availability and set its price for each calendar day.

1.3.2 How are Airbnb Transactions Taxed?

In addition to consumer safety concerns, local governments expressed frustration with Airbnb hosts' avoidance of short-term rental taxes. In cities with significant tourism, the

[11] Paris is thought to have nearly 40,000 active Airbnb listings, the most of any city in the world.

estimated loss of occupancy tax revenue is significant. Initially, Airbnb's position was that its rentals were not subject to occupancy taxes because transactions were "peer-to-peer" rather than commercial in nature. In May 2014, the company officially retracted this view and announced that it believed its hosts were responsible for paying occupancy taxes to local governments. It also amended its "Terms of Service" agreement to inform hosts of their obligation to research and comply with applicable local taxes and regulations.[12]

On June 28, 2014, Airbnb announced that it had reached an agreement with the city of Portland, OR to collect an 11.5% occupancy tax on all reservations booked on its site, and to pay these taxes to the city at the end of each quarter. Crucially, the agreement explicitly prohibited Portland's city government from requiring Airbnb to disclose information related to taxable transaction that could individually identify hosts. As part of the exchange, the Portland City Council agreed to pass a code revision that would legalize short-term home rentals if residents obtained a $180 permit and installed fire alarms.

Between August 2014 and August 2015, similar agreements to collect and remit hotel sales taxes were signed with San Francisco, CA (14.5%), San Jose, CA (10%), Chicago, IL (4.5%), Washington, DC (14.5%), Philadelphia, PA (8.5%), Durham, NC (6%), San Diego, CA (10.5%), and Phoenix, AZ (3%), as well as several smaller municipalities. Typically, an agreement is announced two weeks before the date when Airbnb begins collecting taxes on all bookings in that jurisdiction. Airbnb notifies affected hosts of the policy change via email shortly after the announcement.

---

[12] Beginning May 1, 2014, Airbnb's Terms of Service includes the following paragraph:
YOU AS A HOST UNDERSTAND AND AGREE THAT YOU ARE SOLELY RESPONSIBLE FOR DETERMINING (I) YOUR APPLICABLE TAX REPORTING REQUIREMENTS, AND (II) THE TAXES THAT SHOULD BE INCLUDED, AND FOR INCLUDING TAXES TO BE COLLECTED OR OBLIGATIONS RELATING TO APPLICABLE TAXES IN LISTINGS. YOU ARE ALSO SOLELY RESPONSIBLE FOR REMITTING TO THE RELEVANT AUTHORITY ANY TAXES INCLUDED OR RECEIVED BY YOU. AIRBNB CANNOT AND DOES NOT OFFER TAX-RELATED ADVICE TO ANY MEMBERS.

1.3.3 How do Taxes Appear to Airbnb Guests?

When a guest searches for a rental on Airbnb, she is presented with a set of search results that includes an image, location, and tax-exclusive estimate of the nightly fee for each listing (Figure 1.1).[13] After a guest clicks on a listing, she is shown a more detailed accounting of the rental cost, including Airbnb's service fee and any occupancy tax. Figure 1.2 shows examples of listings from two jurisdictions: one that has a bilateral agreement with Airbnb (Chicago, IL), and one that does not (Evanston, IL). Notice that both listings appear among same set of search results. Without clicking on a listing, it is not evident whether an occupancy tax applies to it.

1.4 Data

My analysis makes use of multiple datasets. Below, I describe each dataset's source and features, and then discuss descriptive analysis of hosts' price-setting behavior.

1.4.1 Data Sources and Details of Key Variables

To measure the response of hosts to Airbnb's remittance agreements, I collect information on listings for selected U.S. cities and their surrounding areas between December 2014 and August 30, 2016.[14] My data collection focused on cities with large tourisms sectors and cities who had announced, but not yet implemented, occupancy tax remittance agreements with Airbnb. In total, 20 cities enactment agreements during the period of data collection. See Table

[13] The price shown in the search results is the average cost per night of the room, excluding taxes and Airbnb's service fee. For example, if a listing's rental prices for Friday, Saturday, and Sunday are \$90, \$100, and \$110, respectively, and the listing has a \$30 cleaning fee, then the price displayed in the search results will be \$110 (90+100+110+30 / 3 = 110).

[14] These data are collected using an automated script or "crawler" that systematically browses Airbnb.com and collects information on listings associated with a particular geographic search term (e.g., "New York, NY"). The script mimics the browsing experience of a potential guest by clicking through each listing in the search results and obtaining its characteristics.

1.1 for a list of enactment dates. I also collect data for five cities that do not enact agreements during the period—these cities serve as controls. In addition to listings within the city itself, I collect data on listings in metro areas (MSAs) to which the implementing cities belong. For each listing, I obtain its approximate geographic coordinates,[15] price, unit type (e.g., shared, private room, entire home), number of reviews, and whether it can be booked instantly. Listings and hosts are each identified by a unique ID, facilitating the tracking of listings over time.

Data are collected in multiple waves, based on the implementation dates of remittance agreements. To supplement these collection efforts, I purchased additional listing data from Airdna, a company that collects Airbnb listing data. My final analysis sample includes all listings in the city and greater metro areas[16] for all cities in the study between Dec. 2014 and Aug. 2016.

When a guest searches for listings in a given location, Airbnb's site returns information on the price and neighborhood of up to 18 listings per page. By clicking on a listing, the user gains additional information about its amenities, reviews, and availability. Availability is displayed using a calendar that the host controls, and where days can be designated as either available for booking or not. If designated available, the default price for that day is the listing price. However, hosts have the option of overriding the listing price for a particular day, such as for a major sporting event. In the analysis that follows, I distinguish between the 'listing price' and the 'booking price.' The latter is the final consumer price, equal to the listing price plus the Airbnb service fee, the cleaning fee, and the tax if an agreement is in place. Consumers review the booking price before the transaction is completed.

[15] Geographic coordinates are purposefully offset by a small distance from the street address registered by the host for privacy. Once a listing is booked, the guest is sent an email with the exact street address. Anecdotal evidence, based on discussions by hosts on internet forums, suggests that these offsets are small (less than 1/8 mile) and, importantly, according to Airbnb's website, offsets are done within neighborhoods.

[16] Listings are included based on the intersection of approximate longitude and latitude coordinates and the U.S. Census MSA boundary files.

1.4.2 Descriptive Statistics

Table 1.2 contains descriptive statistics for treatment and control cities. Col. 1 provides the number of unique listings in the entire metro area (both treatment and control), while Col. 2 contains the number of listings located within the municipal boundary. Col. 4-6 provide means of relevant variables for each metro.

Treated cities differ in the number of listings observed in un-treated, neighboring municipalities. For example, almost one third of listings in the Washington metro area are located in neighboring municipalities, compared to a relatively smaller fraction of listings in the Chicago metro area. Washington, D.C. is perhaps uniquely well-suited for the purpose of comparing treated host behavior to that of untreated, nearby controls: more than a third of the listings returned in a search for the city were located in Arlington, VA, Falls Church, VA, or Bethesda, MD, three municipalities that did not sign remittance agreements with Airbnb. Visual evidence of this is provided in Figure 1.4, which shows the spatial distribution of listings in Washington, Chicago, Oakland and Los Angeles.

Figure 1.3 displays the fraction of listings that change price at least once in three of the treatment cities in each week, limited to those listings appearing at least once in both the pre- and post-agreement periods. On average over the study period, approximately twenty percent of listings change price each week, while in San Diego, closer to a third of listings observed in any given week change prices at least once.

Finally, Figure 1.5 displays a histogram of prices across all listings under $250 in the data. It is evident that hosts employ a number of heuristic pricing strategies, such as choosing prices in increments of $10 or $5.

1.5 VCA Effect on Prices and Tax Collections

This section provides evidence on the first prediction of the model: in the absence of evasion, which side of the market is tasked with remitting the tax is irrelevant. To generate this in the context of cities adopting Airbnb VCAs, I answer two specific questions. First, what is the effect of shifting the hotel tax remittance duty from individual hosts to the platform firm itself on tax-inclusive prices? Second, does this policy affect revenue collection in a manner that is consistent with pre-policy evasion? I find that the policy increases both after-tax prices and the city's collections of hotel tax revenue.

1.5.1 What is the Effect of Shifting the Remittance Responsibility to the Firm on Consumer Prices?

To identify the effect on consumer prices of VCA adoption, I employ two complementary estimation techniques that rely on separate identifying assumptions. First, I exploit variation in when and where—both across and within metropolitan areas—VCAs were adopted to estimate a triple difference specification, as well as its event study analogue. I further refine both of these specifications by allowing the magnitude of treatment to scale with the local hotel tax rate.

Next, I leverage the substantial number of listings in my sample that are proximate to a municipal political border to implement a geographic regression discontinuity (RD) design. The validity of this design relies on the fact that, to consumers, listings within close proximity to a political border may otherwise be seen as close substitutes, yet those in an implementing jurisdiction may have sharply different tax liabilities and remittance obligations. To guard against a failure of the assumptions necessary for a geographic RD design to be valid, I also implement a difference-in-discontinuities, or "diff-in-disc," design.

1.5.1.1 Triple Difference and Event Study Estimates

Using data at the listing-date level, I estimate the following OLS equation on after-tax prices:

$$y_{imct} = \gamma_i + \gamma_t + \gamma_{mt} + \gamma_{ct} + \pi M_i C_i POST_{mt} \tau_m + \varepsilon_{imct} \tag{4}$$

where i is an individual listing in metro area m at time t. Each metro area fully contains the boundaries of a city, and each listing is located either within that city ($C_i = 1$) or outside of it but still within the metro area ($C_i = 0$). The coefficient of interest, π, captures differences in the outcome for listings that meet three conditions: they are (1) located within a metro area that contains a city that adopts a VCA ("treated metro") ($M_i = 1$), (2) located within a city ($C_i = 1$), and (3) observed after the VCA adopted by the city within the treated metro goes into effect ($POST_{mt} = 1$). Controls include fixed effects that absorb any time-invariant listing characteristics (γ_i), national time trends (γ_t), metro-specific time trends (γ_{mt}), and trends common to listings inside or outside of cities (γ_{ct}).

Conceptually, this specification compares listings within the cities of treated metros to listings outside of those cities in the same metros, and to listings in metros that are not treated, as well as to themselves before treatment starts. Because the magnitude of the policy's effect can be expected to vary in direct proportion to the prevailing hotel tax rate (τ_m), the intensity of treatment is allowed to vary with its dose.

I report estimates for this specification (4) on the log tax-inclusive price paid by consumers in Table 1.3 (Col 3). For example, for each one percentage point increase in the effective tax rate, the price paid by consumers rises by approximately 0.9 percent.

This price increase, in addition to violating statutory irrelevance, suggests that the burden of increased compliance falls heavily on consumers. The effects on the advertised, pre-tax price (Col 1), and on reservations (Col 2), have the opposite sign, as expected, but much more modest one tenth of one percent decrease. I also report the results of a traditional difference-in-differences specification, which restricts the sample to listings from treated metros. Estimates from the pooled diff and triple diff are appreciably similar, but diverge (in some cases, significantly) when estimated separately by metro.

Figure 1.6 visually displays the coefficients[17] of the analogue event study for Washington D.C., Chicago and San Diego. Notably, the difference in the control and treatment listings is negligible in the weeks preceding the policy, satisfying the parallel trends assumption.

1.5.1.2 Regression Discontinuity and Difference-in-Discontinuities

Using data at the individual listing level for all listings within two miles of a municipal border of a treated city, I first estimate the following parametric RD specification using OLS:

$$y_i = \beta X_i + \pi_1 R_i + J_i(\pi_2 + \pi_3 R_i) + \varepsilon_i \qquad (5)$$

where *i* is an individual listing located R_i miles inside the municipal border, $J_i = 1[R_i \geq 0]$ is an indicator for listings within the municipal border, and X_i is a vector of time-invariant listing

[17] The estimating equation is:
$y_{imct} = \gamma_i + \gamma_t + \gamma_{mt} + \gamma_{ct} + \sum_{j>-A}^{B-1} \delta^j D_{mt}^j M_{im} C_{ic} \tau_m + \delta^B D_{mt}^B M_{im} C_{ic} \tau_m + \varepsilon_{imct}$
where D_{mt}^j is an indicator for metro *m* in time period *t* being *j* periods away from the implementation of the policy.

characteristics. The coefficient of interest, π_2, captures differences in the outcome for listings just inside the municipal border relative to those just outside of it.

Validity in this context requires that any observed and unobserved listing characteristics must vary smoothly across the cutoff, while the only factor changing sharply at the border is the tax treatment. On the one hand, Airbnb's interface encourages consumers to treat geographic search areas as contiguous, showing results without being constrained to municipal boundaries. This would suggest that, from the consumer's perspective, two listings close to but on opposite sides of a municipal border are equally attractive, and the setting is an appropriate one for an RD. On the other hand, local amenities and property taxes can differ sharply across neighboring municipalities, and these differences could manifest themselves in the quality of the housing stock—and the critical identification assumption fails.[18]

To address this concern, I take advantage of the fact that enactment of a VCA might impact the magnitude of any discontinuity that predated it. Following the approach of Grembi, Tommaso, and Troiano (2012), I modify the RD estimating equation in (5) to measure the change in after-tax prices on either side of the border, before and after VCA enactment:

$$y_i = \beta X_i + \pi_1 R_i + J_i(\pi_2 + \pi_3 R_i) + T_t\big(\rho_1 R_i + J_i(\rho_2 + \rho_3 R_i)\big) + \varepsilon_i \tag{6}$$

This "diff-in-disc" specification shares much in common with (5), with the addition of an indicator for the post-treatment period, T_t. The coefficient of interest, ρ_2, captures any *additional* difference in the outcome for listings just inside a municipal border ($J_i = 1$) in the post-treatment period ($T_t = 1$). Estimates, reported in col. 3 and 4 of Table 1.4, confirm the existence of a price

[18] To empirically test whether the RD assumptions are met, I present in Table 1.4 estimates of (5) separately for a time period before and after policy enactment within select metropolitan areas containing treated cities. The results in columns 2 and 5 strongly suggest that a price discontinuity exists after the enactment of VCAs in these cities. However, in several cases—San Diego, Washington, D.C., and Phoenix—a discontinuity also exists *before* any VCA is enacted.

discontinuity.[19] For most cities, this discontinuity is largest when measured one month immediately before and after the policy, and diminishes when that window is broadened to two months before and after, suggesting that the competitive pressure from cross-border listings may (slowly) cause hosts inside the city to lower their prices.

Like the RD specification, the diff-in-disc specification provides a causal estimate of the effect of VCA enactment on outcomes for listings close to a municipal border. This estimate may or may not differ from that recovered by the DDD specification. For example, if the proximity of similar listings that are not (directly) affected by the VCA across a municipal border introduces additional competition that limits the degree of pass-through, then the effect of VCA enactment might vary with distance to the border. In that case, the diff-in-disc estimates will differ from the DDD estimates, which reflect the average price increase across all listings within the city. My estimates weakly support this intuition; for most metros, the diff-in-disc estimates are lower than the triple difference estimates (col. 0, repeated from Table 1.3, col. 3).

Figure 1.9 provides visual evidence of discontinuities in after-tax prices after—and in some cases before—enactment of VCAs in Washington D.C., Chicago, Oakland and Los Angeles. The points plotted here are average residuals of after-tax price after controlling for listing characteristics and a linear time trend, for listings at different distances from the border. The left panel plots listings' average residuals for the thirty days *before* the policy, overlaid with lines of best fit estimated separately on either side of the boundary. The right panel plots the same figure for the thirty days *after* the policy. Relative to any pre-policy discontinuity that

[19] The ten treatment jurisdictions shown are selected from the larger sample of treated cities because they have a sufficient mass of observations within a bandwidth around the municipal border. Each city's border is represented as a polygon in GIS software and divided into 0.1-mile segments, with each pair of segments connecting at a vertex. This creates between 400 and 800 vertices per city, and each listing's distance is calculated relative to its nearest vertex. Identifiers for vertices are included as fixed effects in (5), ensuring that listings on opposite sides of the same border segment are compared with one another. The regressions are estimated for both 30 and 60 days before and after policy implementation.

existed, the gap in after-tax prices between listings on either side of the border increases with the adoption of the policy in all four cities.

As was alluded to in the introduction, it is difficult to interpret from either set of estimates whether or how economic incidence was affected by this policy. My estimates show that the after-tax price rose significantly after remittance was reassigned, and, at least in the short term, there is no indication that the quality of rentals increased. Therefore, it seems reasonable to infer that consumer surplus declined. However, for previously non-remitting hosts, the tax increased their absolute tax burden (from zero) and weakly reduced demand, likely decreasing producer surplus. Without strong assumptions over underlying demand and supply elasticities, and pre-policy compliance, it is difficult to estimate the *comparative* reduction in surplus. Incidence in this context is further discussed in Section 6.

1.5.2 What Effect did the Policy Have on the Municipal Hotel Market and Hotel Tax Receipts?

In this section, I evaluate the effects of the policy on a city's hotel market and hotel tax receipts, using monthly data from STR, a market research firm, and tax collection data obtained from municipal governments via Freedom of Information Act requests. By re-assigning the duty to remit hotel taxes from hosts to Airbnb, and therefore making it more difficult for hosts to evade the tax, the policy could be expected to have at least two effects on a city's hotel market and its hotel tax receipts. First, it effectively increases the price of Airbnb listings, and may therefore increase demand for hotel rooms to the degree that those are seen as substitutes for short-term rentals. Second, even if demand for Airbnb rentals declines somewhat following the policy, it will likely increase a city's hotel tax receipts as the opportunities for evasion dwindle.

Using monthly hotel market and hotel tax receipt data from 2010 through October 2016 for four cities that enacted these policies and three that did not,[20] I estimate the following difference-in-differences specification:

$$y_{mt} = \gamma_m + \gamma_t + \pi Treat_m POST_{mt} + \varepsilon_{mt} \qquad (7)$$

where y_{mt} is the outcome of interest for municipality m in month t. Characteristics invariant to municipality or time period are captured by municipality and time fixed effects, respectively. The coefficient of interest, π, captures the difference in the outcome between municipalities that adopted the policy and those that did not, both before and after its enactment.

The hotel market data capture several monthly measures of a city's hotel market: the occupancy rate, revenue per available room, and total revenue. The occupancy rate is the number of rooms sold divided by the number of available rooms, while the revenue per available room is the total guest revenue divided by total number of available rooms. Table 1.6 reports results from equation (7) for log versions of these hotel market measures. These point estimates suggest that the enactment of the policy had almost no effect on the occupancy rate of hotels, though it did increase revenue per available room by 6.4 percent and total revenue by 1.3 percent; however, none of these estimates are statistically distinguishable from zero.

Table 1.6 also reports results from equation (7) for log hotel tax receipts. Enactment of the policy increased hotel tax receipts by 10 percent, though this estimate is only significant at the 10 percent threshold.

Taken together, these estimates suggest that enactment of the policy bolstered cities' hotel tax collection efforts, as evidenced by the increase in their tax receipts, but they do not provide conclusive evidence one way or the other on its effects on the local hotel market.

[20] Complete hotel market data (from STR) and hotel tax receipt data (from FOIA requests) were assembled for four cities that enacted the policy (San Diego, Palo Alto, Phoenix, Philadelphia) and three that did not (Houston, Austin, Dallas).

1.6 Heterogeneity in Pass-Through by Attention Correlates

In this section, I explore how much of the observed heterogeneity in the price effect can be explained by differences among individuals in characteristics that suggest their attention to price-setting. Concluding from the previous section that non-compliance was pervasive before the policy, I interpret this as heterogeneity in pass-through of an effective tax increase. I find that hosts which present as "more attentive" pass-through less of the effective tax increase to consumers. This finding may generalize to pass-through of actual tax rate changes by inattentive suppliers in the absence of differential evasion opportunities.

1.6.1 How Does the Effect on Prices Differ by Host Observables?

As discussed in section 1.1, hosts differ in their approach to setting prices. For example, variation in price setting sophistication may cause some hosts to respond to the policy by changing (i.e., lowering) their listing price because they anticipate that consumers will be less willing to book at higher prices.[21]

I test for heterogeneity in price response by host characteristics that may be associated with price setting sophistication: time series correlation between the host's pre-policy prices and the prices of hotel rooms; heuristic pricing, such as setting a price divisible by 5 or 10; and proxies for the intensity of rental activity, such as enabling the "instant booking" feature, listing multiple properties on Airbnb, or listing an entire unit (as opposed to a private room in what is likely an owner-occupied dwelling). For each binary host characteristic X_i, I estimate the following triple-difference specification:

[21] Assuming the listing price before the policy was optimal, and demand is not perfectly inelastic, the optimal listing price after the policy is lower.

$$y_{imct} = \gamma_i + \gamma_t + \gamma_{mt} + \gamma_{ct} + \pi M_i C_i POST_{mt} \tau_m X_i + \varepsilon_{imct} \quad (8)$$

The coefficient π represents the average percent difference in tax-inclusive listing prices for hosts with characteristic X_i, located within the major city ($C_i = 1$) in a "treated" metro ($M_i = 1$) after the policy is enacted ($POST_{mt} = 1$), for each percentage point of the hotel tax rate in that metro (τ_m).

Table 1.5 reports results of π estimated separately for each host characteristic. Taken in aggregate, hosts who are less likely to be sophisticated—who do not have instant booking turned on, who exhibit evidence of heuristic price-setting behavior, who do not rent out an entire unit, and who do not list multiple properties—usually have slightly higher tax-inclusive listing prices following the policy than hosts who are more likely to be sophisticated price-setters. For example, heuristic price-setting behavior is associated with a 0.3 or 0.4 percent higher price for every 1 percentage point of effective tax increase. Hosts who enable instant booking, on the other hand, have listing prices that are approximately 0.1 percent lower for every 1 percentage point of effective tax increase.

To further explore the relationship between "attention" and pass-through, I examine how price setting response is related to hosts' pre-policy price correlation with local hotel prices. In comparison to the previously discussed binary characteristics, this measure is continuous, and arguably more comprehensive than self-reported attributes like whether an entire unit is being rented out. To the extent that hotels and Airbnb rental properties are even imperfect substitutes, demand shocks to the hotel market should affect the Airbnb market as well. And there are a number of reasons why hotel price movements should be informative about the direction and magnitude of these shocks: hoteliers, particularly those affiliated or owned by large chains, likely set prices centrally, have extensive experience in doing so, and are pricing a largely standardized product. It is therefore likely that when hotel prices rise or fall, it is due to changes in the demand for short-term rentals that apply to Airbnb hosts as well.

Figure 1.8 (top panel) plots event study coefficients for hosts, estimated separately by whether hosts' pre-policy price correlations with hotel prices are above or below the median within their metro. Hosts whose prices correlated more closely with those of hotels are also more likely to adjust their prices upward by less after the policy, passing through less of the tax to consumers, at least initially. Figure 1.8 (bottom panel) plots event study coefficients estimated separately for hosts at different deciles of the host-hotel price correlation distribution; here, it is even more apparent that the more "sophisticated" hosts, whose prices tracked more closely to those of hotels, pass on less of the tax in the time shortly after the policy. The subsequent convergence of prices may suggest that these sophisticated hosts, upon learning more information about the resilience of consumer demand for Airbnb rentals, bring prices up and in line with those of hosts who were inattentive to the policy's impact on prices.

1.7 Effect of VCA on Entry and Exit

In addition to adjusting prices, hosts can respond to the policy on the extensive margin: by deciding whether and where to list their properties. Hosts whose costs exceed their listing price in the absence of "evasion rents" have two extensive margin responses. First, they may exit the market for short term rentals altogether. Second, they may continue evading the tax by listing on an alternative platform that does not remit tax on behalf of hosts. Similarly, some prospective hosts who would have entered the Airbnb market prior to the policy may choose not to in light of it, or may choose to enter the market through an untaxed platform. I refer to this behavior as "platform jumping." To the extent that some of the tax savings are reflected in lower prices on the untaxed platform, consumers will also have an incentive to search on that platform. I examine both extensive margin behaviors in the next two sub sections.

1.7.1 Airbnb Exit

One plausible margin of adjustment to the policy is a host's decision to exit the Airbnb market. This decision can appear in the data in one of two ways. First, a host can delete her account, which is indicated by her listing no longer being observed after the exit date. Second, a host can "effectively exit" the market by no longer actively making her unit available. (Airbnb is set up to require that hosts actively identify dates during which their units are listed on their calendars as available.) To determine what length of continuous inactivity likely signals an effective exit, I compare the likelihood that a host exhibits subsequent activity—by making the unit available or having it reserved—after inactivity spells of varying. On the basis of this analysis, I find that after 90 or more days of inactivity, hosts have a ten percent or smaller likelihood of becoming active again, and therefore use 90 days as a threshold for effective exit.

I then estimate the triple difference specification (4) where the dependent binary variable y_{imct} is equal to one if host i exited the market—either by deleting her account or effectively exiting— on or after time t.[22] Table 1.7 reports the results. On average, the policy increased the likelihood of exit by one third of one percentage point (Col 1). For comparison, the likelihood that a host in a control city leaves the market on any given day is approximately 1.2 percent, implying that the policy increased the likelihood of exit by 25 percent.

[22] To perform this analysis, the listing-date analysis dataset is extended so that each listing is observed through the end of the data window. This means that if a host deleted her account prior to the end of the data window, new records are created for which the host's listing is flagged as having exited the market.

1.7.2 Platform Jumping

VRBO, Airbnb's main competitor, is one such alternative (untaxed) platform. Platform jumping might be more prevalent between VRBO and Airbnb because the interfaces and requirements for the two sites are virtually identical.[23] While creating an account for the first time on any platform takes a modest amount of effort (Airbnb advertises that it takes less than an hour), the marginal cost of creating an additional listing profile on a similar platform is likely even lower.

I test for a decline in Airbnb entries and reservations (and a corresponding increase in VRBO entrants and bookings) by estimating the following difference-in-differences equation separately for each platform:

$$y_{mt} = \gamma_m + \gamma_t + \pi POST_{mt} + \varepsilon_{mt} \tag{9}$$

where the dependent variable measures entrants/bookings in metro *m* in month *t*. $POST_{mt}$ is equal to one in treated metros after the VCA implementation date. The identifying assumption is that parallel trends in entry exist between treated and untreated metros prior to implementation of the policy. While I find that entry into Airbnb declines and that VRBO entries increase, both effects are only marginally statistically significant.[24]

1.8 Results Summary and Discussion

In this section, I consider the relevance of my findings to broader academic and policy discussions on the role that statutory features play in tax compliance, the long-term collection

---

[24] Alternative DDD specification:

$$y_{pmt} = \gamma_p + \gamma_t + \gamma_m + \gamma_{pm} + \gamma_p POST_{mt} + \gamma_m POST_{mt} + \pi M_m P_p POST_{mt} + \varepsilon_{mpt}$$

efficacy of VCAs, and the welfare and distributional consequences of taxing markets substantially populated by unresponsive sellers.

In previous sections, I establish four main empirical findings:

(1) Shifting the remittance duty substantially increased tax-inclusive prices. I estimate this effect using both a triple differences and discontinuity-in-differences approach. Pooled triple difference estimates indicate the policy increased tax-inclusive prices by 0.9 percent for everyone percentage point of tax re-assigned to the platform, though estimates by metro vary.

(2) Shifting the remittance duty increased tax revenue collections. For every one percentage point re-assigned, tax revenues increase by 0.8 percent, scaled by Airbnb's market share, though it is not clear what portion of this effect is driven by an increase in traditional hotel prices.

(3) Extent to which tax-inclusive prices increased correlated with attention. I estimate the triple difference specification interacted with host characteristics likely associated with attention to price setting. Hosts whose prices closely correlate with traditional hotel prices pass on less of the effective tax increase.

(4) Shifting remittance duty induces exit of less attentive hosts. I find that host entry into the Airbnb market drops after VCA adoption, and, further, that entry into VRBO, Airbnb's closest competitor, increases after VCA adoption.

1.8.1 Location of the Duty to Remit Affects Compliance with, and Incidence of, Consumption Tax

I find that shifting the remittance duty from hosts to the Airbnb platform increased equilibrium prices by nine-tenths of a percent for each percentage point of a municipality's hotel tax rate. These findings are broadly consistent with an emerging strand of the consumption tax incidence literature which finds that statutory features—such as the identity of the tax remitter, the point of collection, or the direction of the tax change—can have first order effects on tax incidence.

I also find suggestive evidence that local government hotel tax receipts rise by 10 percent, on average, following policy enactment, lending support to the evasion channel hypothesis of KKMS, extended to a monopolistically competitive market structure. In other words, the change in tax incidence here may result, in part, from different evasion opportunities available to each side of the market. These different evasion opportunities mean that a tax levied on the demand side of the market may result in a different equilibrium outcome than a similar tax levied on the supply side of the market. Indeed, prior to the enactment of the policy, anecdotal evidence suggests very few hosts were complying with the law and remitting hotel taxes. Unlike textbook tax incidence examples, then, the policy may not merely shift tax incidence between the two sides of a market, but rather changes its overall magnitude as well.

A significant appeal of requiring firms to assess and remit taxes, such as payroll and income taxes in the U.S., is that it is more cost effective to administer given the small number of firms relative to taxpayers. However, when there are many small "firms" each responsible for remitting a small fraction of total tax revenue—as is the case with individual Airbnb hosts and local hotel taxes—it becomes costly to monitor compliance with the tax. This situation is likely to grow more prevalent as technology and business practices lower the barriers to individuals monetizing their time or possessions; not only will many more people be subject to new tax obligations—stretching tax authorities thin—but they may also be unaware of them. If tax receipts do not keep pace with tax obligations, it will not always be clear whether sellers are making a conscious decision to evade in light of a low probability of detection, or whether they lack information about the tax and their duty to pay it. Distinguishing between these two will be crucial to designing remedies to ensure greater compliance.

1.8.2 Welfare and Distributional Concerns in Taxing Unresponsive Sellers

The presence of unresponsive sellers in a market, as appears to be the case with Airbnb hosts, can have significant welfare consequences. The tax salience literature has shown that, when taxes are not salient, consumers will underreact to them (Chetty, Looney, Kroft 2009). As a result, the deadweight loss of imposing a sales tax is inversely proportional to how salient that sales tax is. Yet, in the context of Airbnb, if a *host* underreacts to a change in remittance obligation that functionally acts as an effective tax increase, she may end up passing through 100 percent of the tax to consumers. This, in turn, can have a large effect on consumer behavior and result in a greater deadweight loss than if the host was aware of, and responsive to, the tax. In the long run, this may be mitigated by the introduction of algorithms that assist hosts in setting prices, but in the short run, where pricing decisions are often the result of inertia or inattention, this remains a real concern.

1.8.3 The Promise and Peril of Government Reliance on VCAs

Voluntary compliance agreements (VCAs) are attractive tax collection tools for local governments for two reasons. First, in the U.S., most sales taxes are imposed by state and local governments that have limited power to compel “remote” or platform sellers to remit taxes; absent a federal solution, VCAs allow these governments to recoup some of this otherwise foregone tax revenue. Second, VCAs offer an alternative to information reporting for capacityconstrained states that may be unable to collect a tax even with identifying information about the seller.

However, the long-term effects of VCAs remain unclear, and may be potentially troubling. Platforms that negotiate VCAs with local governments often do so from a position of considerable market strength, and as a result they can secure significant concessions. For example, in exchange for remitting hotel taxes as a lump sum, Airbnb's VCAs with local governments do not require it to provide identifying information about hosts to the tax authorities. Not only does this prevent local tax authorities from recouping taxes owed on previous transactions from hosts directly, it also prevents them from monitoring their behavior on other platforms, including direct competitors to Airbnb, that have not signed VCAs. Put differently, VCAs can make local governments permanently dependent on the individual firm for significant revenues, and are signed when those firms have accrued sufficient market power to negotiate them on their terms.

1.9 Conclusion

In classical economic theory, the incidence of a consumption tax is exclusively determined by market-wide demand and supply elasticities. This paper contributes to an emerging empirical literature which suggests that other factors, such as assignment of the remittance obligation, may affect incidence in practice.

I find that shifting the legal obligation to remit hotel taxes from small, independent hosts to Airbnb increases after-tax prices paid by consumers. The magnitude of this effect differs by a number of host characteristics related to sophistication. While several rationalizations of my estimates are possible, this primary result is consistent with low levels of voluntary compliance among individual hosts prior to implementation of mandatory withholding, despite the existence of a paper trail and federal information reporting on Airbnb rental transactions. This finding has potentially important implications for understanding the potential revenue and distributional consequences of taxing non-employee service transactions facilitated by digital platforms.

Table 1.1: Airbnb Voluntary Collection Agreements

City	Tax Rate	Announcement Date	Implementation Date	Metropolitan Statistical Area
Treatment				
Boulder, CO	7.5		October 1, 2016	Boulder, CO Metro Area
Chicago, IL	4.5	February 1, 2015	February 15, 2015	Chicago-Naperville-Elgin, IL-IN-WI Metro Area
Cleveland, OH	3	June 20, 2015	July 1, 2016	Cleveland-Elyria, OH Metro Area
Washington D.C.	14.5	February 1, 2015	February 15, 2015	Washington-Arlington-Alexandria, DC-VA-MD-WV Metro Area
Golden, CO	7.5		November 1, 2016	Denver-Aurora-Lakewood, CO Metro Area
Kill Devil Hills, NC	6.75	May 23, 2015	June 1, 1915	Kill Devil Hills, NC
Jersey City, NJ	6	October 12, 2015	February 1, 2016	New York-Newark-Edison, NY-NJ-PA Metropolitan Statistical Area
Los Angeles, CA	14	July 18, 2016	August 1, 2016	Los Angeles-Long Beach-Anaheim, CA Metro Area
Malibu, CA	12		April 20, 2015	Oxnard-Thousand Oaks-Ventura, CA Metro Area
Newark, NJ	14.5	April 12, 2016	May 1, 2016	New York-Newark-Jersey City, NY-NJ-PA Metro Area
Oaks Island/Myrtle Beach	6.75	May 23, 2015	June 1, 2015	Myrtle Beach-Conway-North Myrtle Beach, SC Metropolitan
Oakland, CA	14	July 5, 2015	July 15, 2015	San Francisco-Oakland-Hayward, CA Metro Area
Palo Alto, CA	14	November 30, 2014	January 1, 2015	San Jose-Sunnyvale-Santa Clara, CA Metro Area
Philadelphia, PA	8.5	July 1, 2015	July 15, 2015	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD Metro Area
Portland, OR	11.5		July 1, 2014	Portland-Vancouver-Hillsboro, OR-WA Metro Area
Phoenix, AZ	5	July 1, 2015	July 1, 2015	Phoenix-Mesa-Scottsdale, AZ Metro Area
Santa Clara, CA	9.5		November 1, 2015	San Jose-Sunnyvale-Santa Clara, CA Metro Area
San Diego, CA	10.5	July 1, 2015	July 15, 2015	San Diego-Carlsbad, CA Metro Area
San Francisco, CA	16.5	August 1, 2014	October 1, 2014	San Francisco-Oakland-Hayward, CA Metro Area
San Jose, CA	10	January 1, 2015	February 1, 2015	San Jose-Sunnyvale-Santa Clara, CA Metro Area
Control				
Austin, TX	0	NA	NA	Austin-Round Rock, TX Metro Area
Dallas, TX	0	NA	NA	Dallas-Fort Worth-Arlington, TX Metro Area
Houston, TX	0	NA	NA	Houston-The Woodlands-Sugar Land, TX Metro Area
New Orleans, LA	0	NA	NA	New Orleans-Metairie, LA Metro Area
Savannah, GA	0	NA	NA	Savannah, GA Metro Area

Table 1.2: Sample Summary Statistics

City	N (Metro)	N (City)	N (Listing X Days)	Avg. Price	Avg. No. per Month per Host: Price Changes	Avg. No. per Month per Host: Reservations	Entire Apt. (%)
Treatment							
Boulder, CO	3,657	2,446	193,878	143	2.9	4.9	24.1%
Chicago, IL	21,453	18,786	682,957	139	3.6	4.6	54.0%
Cleveland, OH	4,065	1,926	128,975	467	1.8	2.4	34.3%
Washington D.C.	22,401	11,904	769,856	148	3.0	4.9	51.8%
Golden, CO	12,927	112	373,278	124	3.6	5.4	29.8%
Kill Devil Hills, NC	151,370	3,993	2,064,362	170	1.6	4.3	81.4%
Jersey City, NJ	2,967	850	125,898	207	3.3	4.5	16.6%
Los Angeles, CA	87,598	51,793	2,437,104	198	3.5	5.0	49.2%
Malibu, CA	89,913	685	510,603	711	3.5	4.1	31.7%
Newark, NJ	151,370	503	876,406	161	1.6	4.6	77.0%
Oak Island, NC	3,454	467	106,225	198	2.7	3.9	10.6%
Oakland, CA	21,669	4,815	1,849,500	178	1.8	5.2	34.4%
Palo Alto, CA	14,720	1,813	1,323,801	415	2.1	3.9	33.0%
Philadelphia, PA	17,664	13,979	1,847,512	491	1.1	2.6	35.3%
Portland, OR	10,727	7,810	437,326	123	4.0	6.2	24.0%
Phoenix, AZ	12,219	4,438	505,783	329	2.7	3.8	22.9%
Santa Clara, CA	7,696	1,729	711,486	467	2.0	3.6	33.4%
San Diego, CA	21,096	14,995	686,205	219	3.8	4.4	29.9%
San Francisco, CA	47,623	25,954	5,558,380	218	1.7	5.2	45.7%
San Jose, CA	12,907	5,211	1,056,904	454	2.0	3.6	33.3%
Control							
Austin, TX	21997	19,250	949,109	277	3.3	3.6	25.9%
Dallas, TX	7823	3,710	274,168	142	3.5	4.7	47.6%
Houston, TX	12726	8,497	409,408	239	2.8	3.0	37.1%
New Orleans, LA	10539	9,723	424,294	187	4.3	4.6	29.2%
Savannah, GA	1847	1,082	66,989	235	5.7	6.7	22.8%

Table 1.3: Triple Difference; Dependent Variable Log Prices

	All Listings			Fixed Listing Composition		
City	**Log Listing Price**	**Reservations**	**Log Price**	**Log Listing Price**	**Reservations**	**Log Price**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Pooled	-0.001***	-0.001***	0.009***	-0.001***	-0.000***	0.009***
	0.000	0.000	0.000	0.000	0.000	0.000
Boulder, CO	0.039***	-0.073***	0.111***	0.021***	-0.041***	0.093***
(t=7.5%)	(0.002)	(0.003)	(0.002)	(0.002)	(0.004)	(0.002)
Chicago, IL	-0.009***	0.002	0.035***	-0.006***	-0.002	0.038***
(t=4.5%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Cleveland, OH	0.032***	0.045***	0.062***	0.005	0.010***	0.035***
(t=3%)	(0.004)	(0.003)	(0.004)	(0.005)	(0.004)	(0.005)
Washington D.C.	0.001**	-0.017***	0.137***	0.002***	-0.028***	0.138***
(t=14.5%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Golden, CO	0.032***	-0.074***	0.104***	0.031***	-0.041***	0.103***
(t=7.5%)	(0.002)	(0.005)	(0.002)	(0.002)	(0.006)	(0.002)
Jersey City, NJ	0.009***	-0.017***	0.068***	0.014***	-0.005***	0.073***
(t=6.75%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Kill Devil Hills, NC	0.004	0.088***	0.069***	-0.056***	0.034***	0.010***
(t=6%)	(0.003)	(0.004)	(0.003)	(0.003)	(0.005)	(0.003)
Los Angeles, CA	0.002***	-0.023***	0.133***	-0.011***	-0.023***	0.120***
(t=14%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Malibu, CA	0.015***	-0.005***	0.128***	0.001	-0.013***	0.114***
(t=12%)	(0.001)	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)
Newark, NJ	-0.035***	-0.035***	0.023***	-0.034***	-0.019***	0.024***
(t=14.5%)	(0.001)	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)
Oak Islands	-0.035***	0.016***	0.030***	-0.094***	0.027***	-0.028***
(t=6.75%)	(0.002)	(0.003)	(0.002)	(0.004)	(0.005)	(0.004)
Oakland, CA	0.006***	-0.030***	0.137***	0.002***	-0.016***	0.133***
(t=14%)	(0.001)	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)
Palo Alto, CA	-0.019***	-0.008***	0.112***	-0.036***	0.017***	0.095***
(t=14%)	(0.001)	(0.003)	(0.001)	(0.001)	(0.003)	(0.001)
Philadelphia, PA	-0.012***	-0.030***	0.070***	-0.005***	-0.012***	0.077***
(t=8.5%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Phoenix, AZ	-0.070***	-0.019***	-0.021***	-0.028***	0.016***	0.021***
(t=5%)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Santa Clara, CA	0.045***	-0.096***	0.136***	0.037***	-0.034***	0.127***
(t=9.5%)	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)
San Diego, CA	-0.001	0.028***	0.099***	-0.018***	0.022***	0.082***
(t=10.5%)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
San Francisco, CA	0.028***	-0.088***	0.181***	0.026***	-0.044***	0.179***
(t=16.5%)	(0.001)	(0.002)	(0.001)	(0.001)	(0.003)	(0.001)
San Jose, CA	-0.017***	-0.021***	0.078***	-0.048***	0.006**	0.047***
(t=10%)	(0.001)	(0.002)	(0.001)	(0.001)	(0.003)	(0.001)

Table 1.4: Regression Discontinuity Estimate on Log Price at the Municipal Border

		30 Day Window			60 Day Window		
City	DDD*	Pre	Post	Diff-Disc	Pre	Post	Diff-Disc
		(1)	(2)	(3)	(4)	(5)	(6)
Los Angeles, CA	0.133***	0.005	0.099***	0.127***	0.012	0.095***	0.110***
(t=14%)	(0.001)	(0.02)	(0.02)	(0.004)	(0.01)	(0.01)	(0.002)
		373,888	390,840	878,352	748,978	806,280	1,768,612
San Diego, CA	0.099***	0.051***	0.182***	0.099***	0.056***	0.177***	0.101***
(t=10.5%)	(0.001)	(0.00)	(0.00)	(0.010)	(0.00)	(0.00)	(0.006)
		420,674	457,678	287,399	811,026	957,586	577,025
Palo Alto, CA	0.112***	0.046	0.065***	0.057**	0.026	0.043***	-0.003
(t=14%)	(0.001)	(0.03)	(0.02)	(0.027)	(0.02)	(0.01)	(0.016)
		12,992	29,058	42,050	25,185	60,365	85,550
San Jose, CA	0.078***	-0.008	0.090***	0.093***	0.008	0.101***	0.120***
(t=10%)	(0.001)	(0.01)	(0.01)	(0.017)	(0.01)	(0.01)	(0.010)
		142,158	145,241	57,507	284,778	292,247	115,281
Santa Clara, CA	0.136***	0.098*	0.281***	0.105***	0.043	0.283***	0.097***
(t=9.5%)	(0.002)	(0.05)	(0.04)	(0.013)	(0.03)	(0.03)	(0.007)
		28,275	29,232	93,970	57,581	57,700	190,713
Oakland, CA	0.137***	0.030	0.040	0.178***	-0.001	0.032*	0.151***
(t=14%)	(0.001)	(0.02)	(0.02)	(0.024)	(0.02)	(0.02)	(0.015)
		46,951	47,019	52,229	94,933	95,780	109,935
Chicago, IL	0.035***	0.041	0.122***	0.116***	0.021	0.057*	0.131***
(t=4.5%)	(0.001)	(0.04)	(0.04)	(0.006)	(0.03)	(0.03)	(0.003)
		24,215	28,014	211,555	46,521	63,414	424,835
Washington D.C.	0.137***	-0.029***	0.095***	0.029**	-0.027***	0.092***	0.033***
(t=14.5%)	(0.001)	(0.01)	(0.01)	(0.013)	(0.01)	(0.01)	(0.008)
		105,734	105,821	316,259	210,914	213,920	647,244
Phoenix, AZ	-0.021***	0.155***	0.280***	0.131***	0.160***	0.271***	0.115***
(t = 5%)	(0.001)	(0.01)	(0.01)	(0.010)	(0.01)	(0.01)	(0.006)
		107,995	108,313	216,308	219,713	220,783	440,496
Boulder, CO	0.111***	-0.006	0.088***	0.083***	0.008	0.110***	0.068***
(t=7.5%)	(0.002)	(0.02)	(0.02)	(0.022)	(0.02)	(0.01)	(0.012)
		58,638	60,755	121,829	116,746	123,779	250,440

Notes. RD coefficients are estimated separately for the thirty day and sixty day intervals around the policy, estimates are reported in columns (1), (2) and columns (4), (5) respectively. Difference and discontinuity estimates are reported in col. (3) and (6). All specifications include for listing characteristics and time fixed effects. To ensure that like listings are being compared, I calculate the closest border vertex for each listing, and include vertex fixed effects. The sample is limited to listings within two miles on either side of the municipal border.
*Triple difference estimates from Table 1.3, col (3) have been repeated for readers' convenience.

Table 1.5: Pooled Triple Difference Estimates on Log Price by Host Characteristic

Dependent Variable: Log Price	(1)	(2)	(3)	(4)
Instant Book Enabled?	-0.002***	-0.001***	-0.002***	-0.001***
	0.000	0.000	0.000	0.000
Base Divisible by 10	-0.001***	0.004***	-0.001***	0.004***
	0.000	0.000	0.000	0.000
Base Divisible by 5	-0.002***	0.003***	-0.002***	0.003***
	0.000	0.000	0.000	0.000
Entire Apartment	0.007***	0	0.008***	0
	0.000	0.000	0.000	0.000
Multiple Properties	0.002***	0	0.001***	-0.000***
	0.000	0.000	0.000	0.000
Host Fixed Effects?	No	Yes	No	Yes
Control Cities?	No	No	Yes	Yes

Table 1.6: Effect of VCA on Log of Hotel Tax Revenue

	Diff-in-Diff	Event Study
	(1)	(2)
Treat * Post* Airbnb Market Ratio	2.656*	
	(1.09)	
Pre Month 4		-0.973
		(2.32)
Pre Month 3		-2.328
		(2.53)
Pre Month 2		-1.769
		(2.59)
Post Month 0		0.599
		(3.19)
Post Month 1		-0.498
		(3.11)
Post Month 2		0.736
		(3.18)
Post Month 3		0.244
		(3.35)
Post Month 5		0.240
		(3.43)
Post Month 6		0.892
		(3.06)
Post Month 7		2.770
		(3.33)
Post Month 8		1.463
		(3.27)
Post Month 9		1.810
		(3.39)
Post Month 10		1.717
		(3.36)
Post Month 11		2.090
		(3.65)
Post Month 12		3.951
		(3.58)
Post Month>12		8.946***
		(2.40)
N	659	659

Notes. Dependent variable is log of city's monthly hotel tax revenues. Treatment is defined as the interaction between Ever Treated City*Post* Relative Airbnb Market Size at the time of treatment. Col. 1 reports difference and difference estimates, Col (2) reports event study estimates for the equivalent specification. Cities included in the sample are Palo Alto, Chicago, Washington D.C. and San Diego (treated); Austin, Dallas and Houston (Never treated). Standard errors are reported below coefficient estimates. All specifications include controls for seasonality.

Table 1.7: Effect of Policy on Entry and Exit (Difference-in-Differences)

	Airbnb Entry	Airbnb Entry (logs)
	(1)	**(2)**
Platform* Treat* Post	-120.743	-1.743
	(55..68)	(1.68)
N	1647	1647

Notes. Col.1 reports average effect of the policy on Airbnb hosts (absolute measure). Col. 2 reports the effect of airbnb entry in logs. Both specifications estimated with seasonal effects. Treatment is defined as the interaction between Platform*Ever Treated City*Post. Platform is equal to 1 if Platform is Airbnb. Standard errors are reported under coefficients.

Figure 1.1: Airbnb Search Results

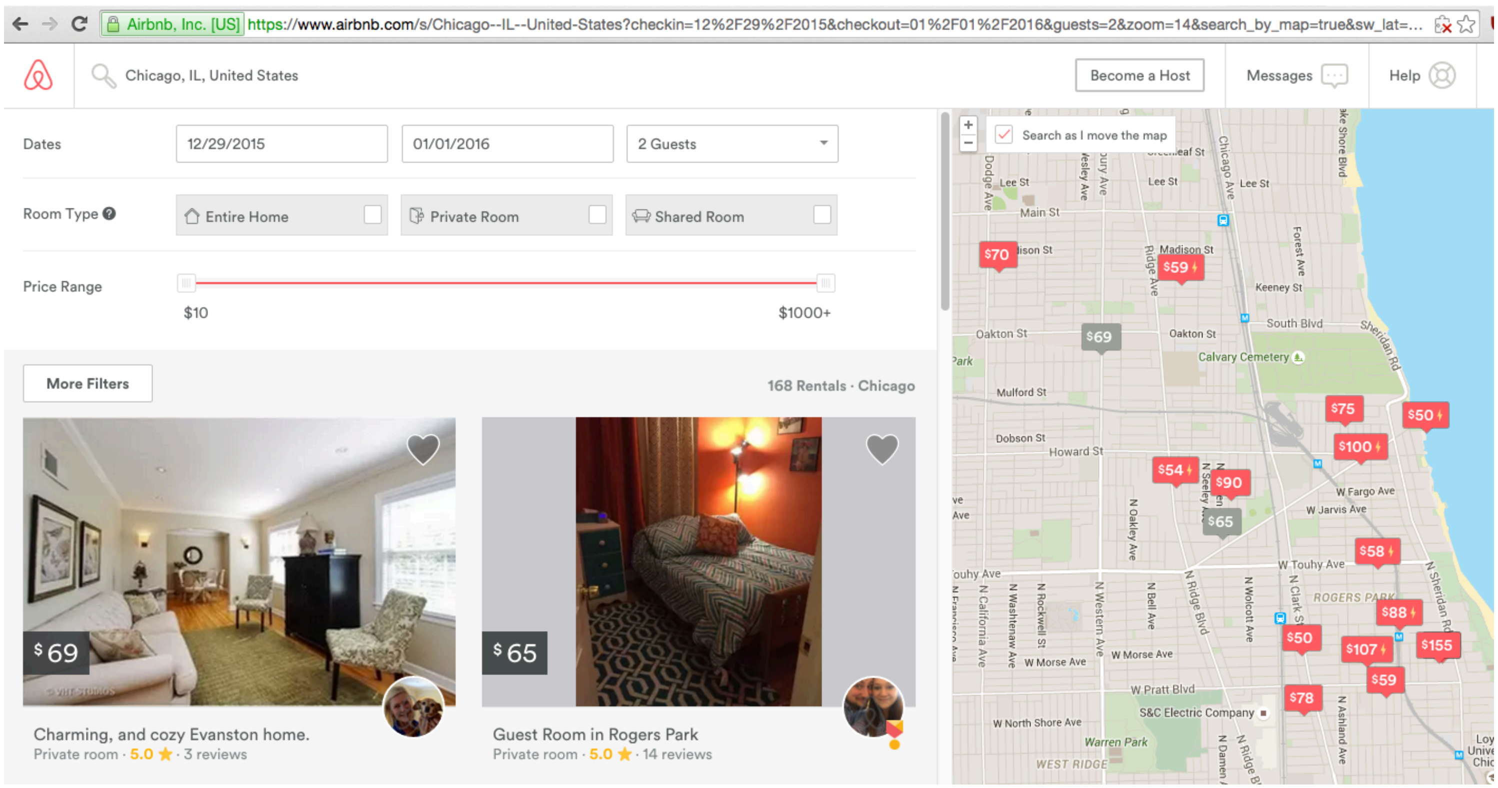

Figure 1.2: Airbnb Listing Details in Chicago and Evanston, IL

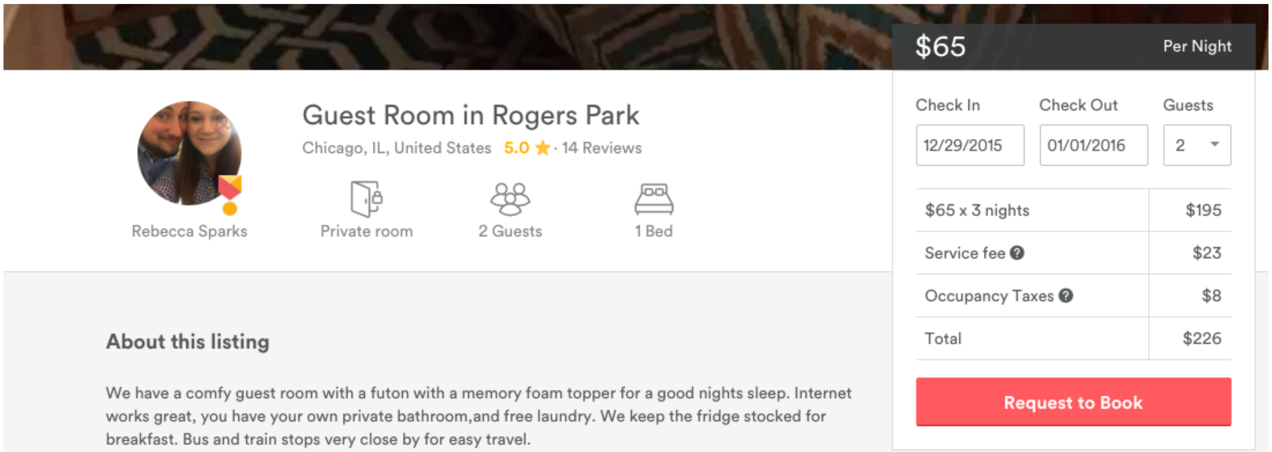

Chicago, IL

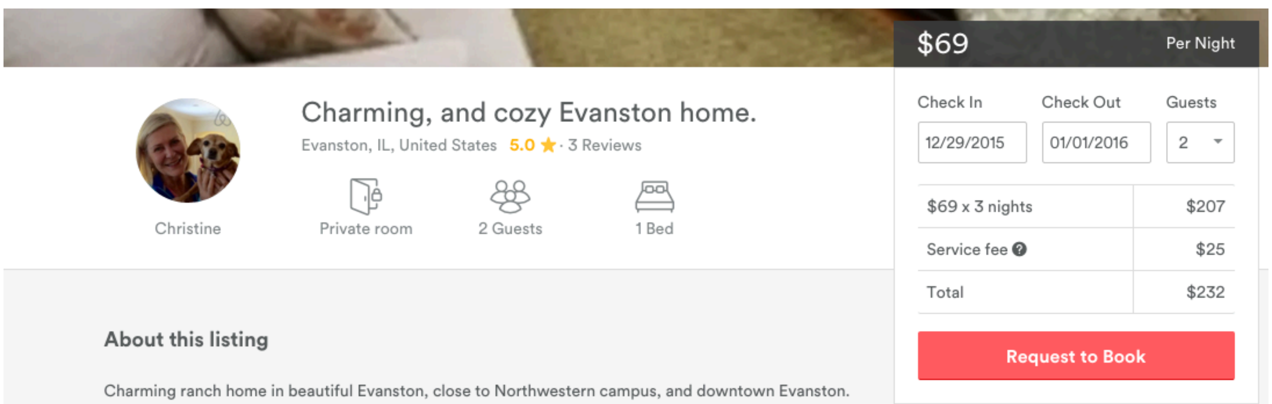

Evanston, IL

Figure 1.3. Percent of Listings with Price Changes (Weekly)

Washington DC

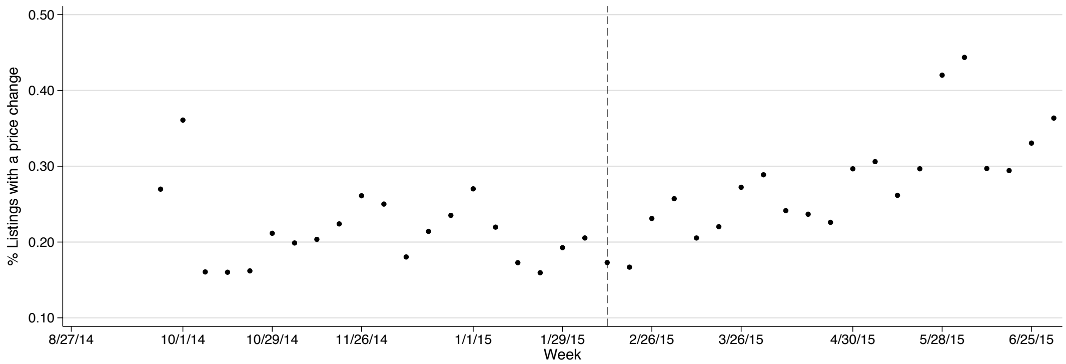

Chicago, IL

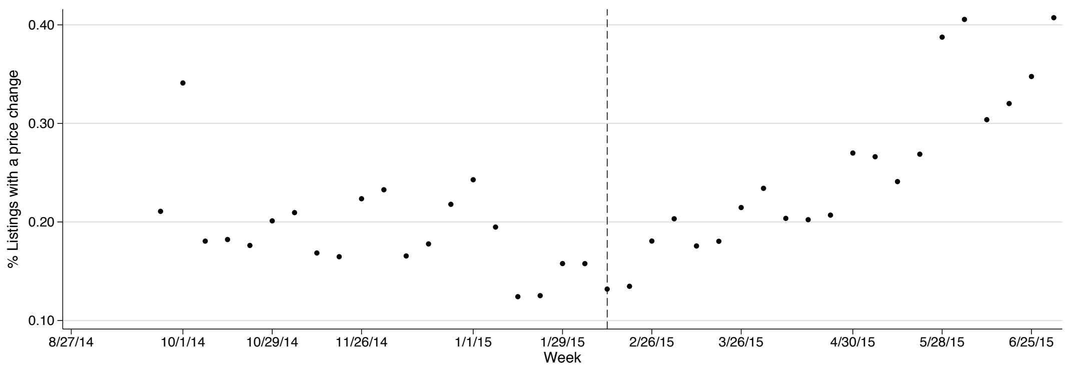

San Diego, CA

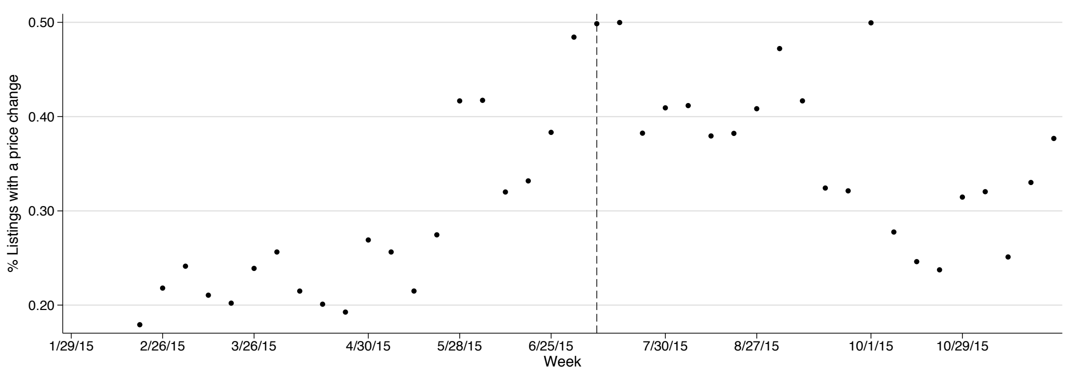

Figure 1.4. Spatial Distribution of Listings, Relative to City Boundaries (D.C., Los Angeles, Chicago, Oakland)
Note: Red dots represent Airbnb listings, Orange dots represent VRBO listings)

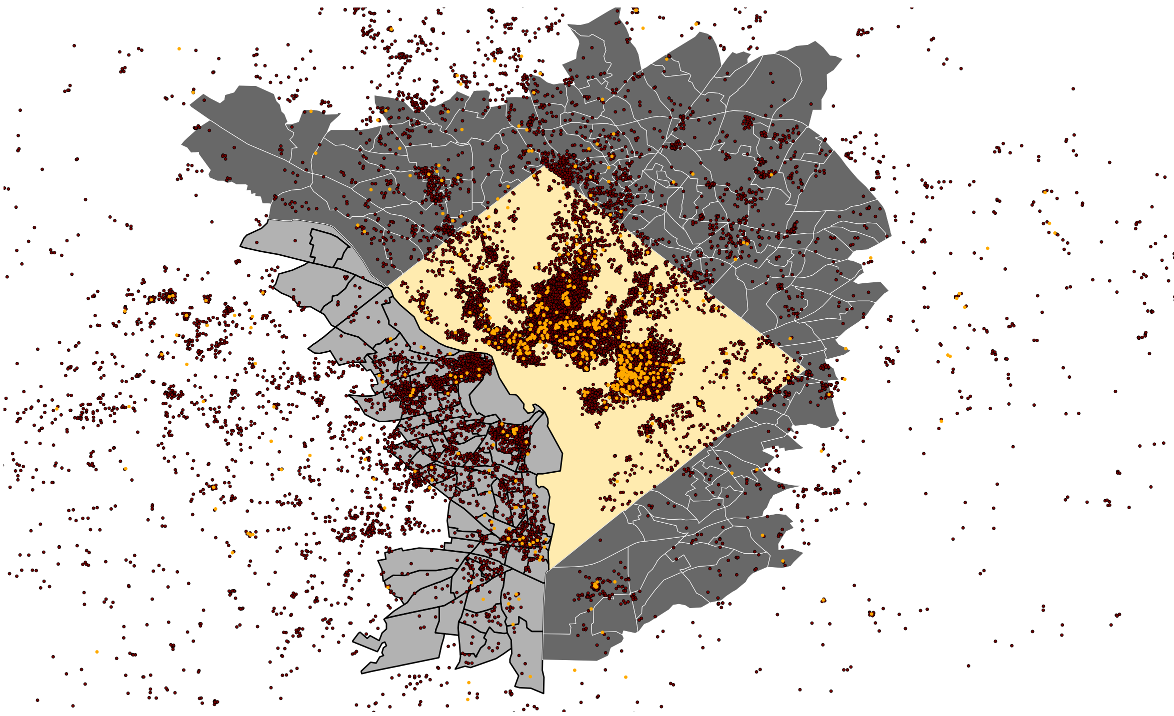

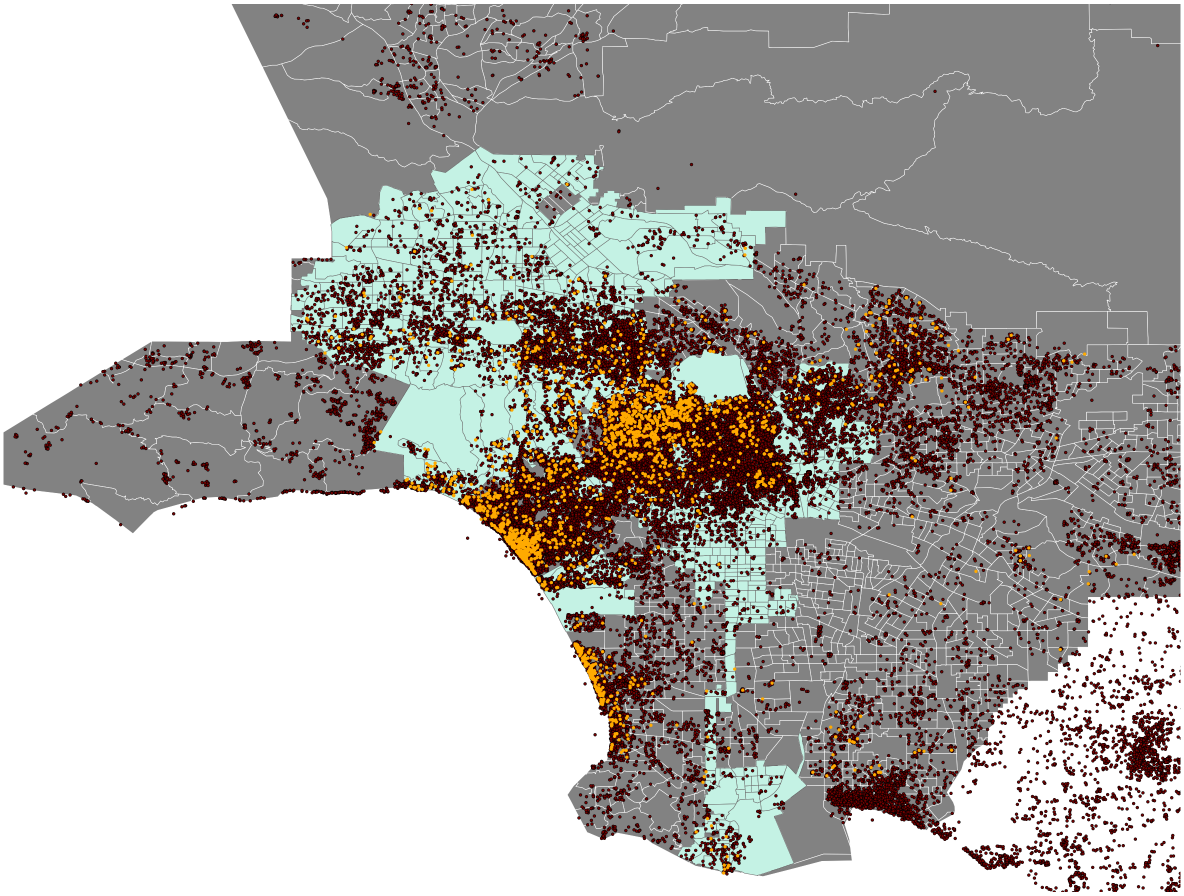

Figure 1.4 (continued): Chicago and Oakland

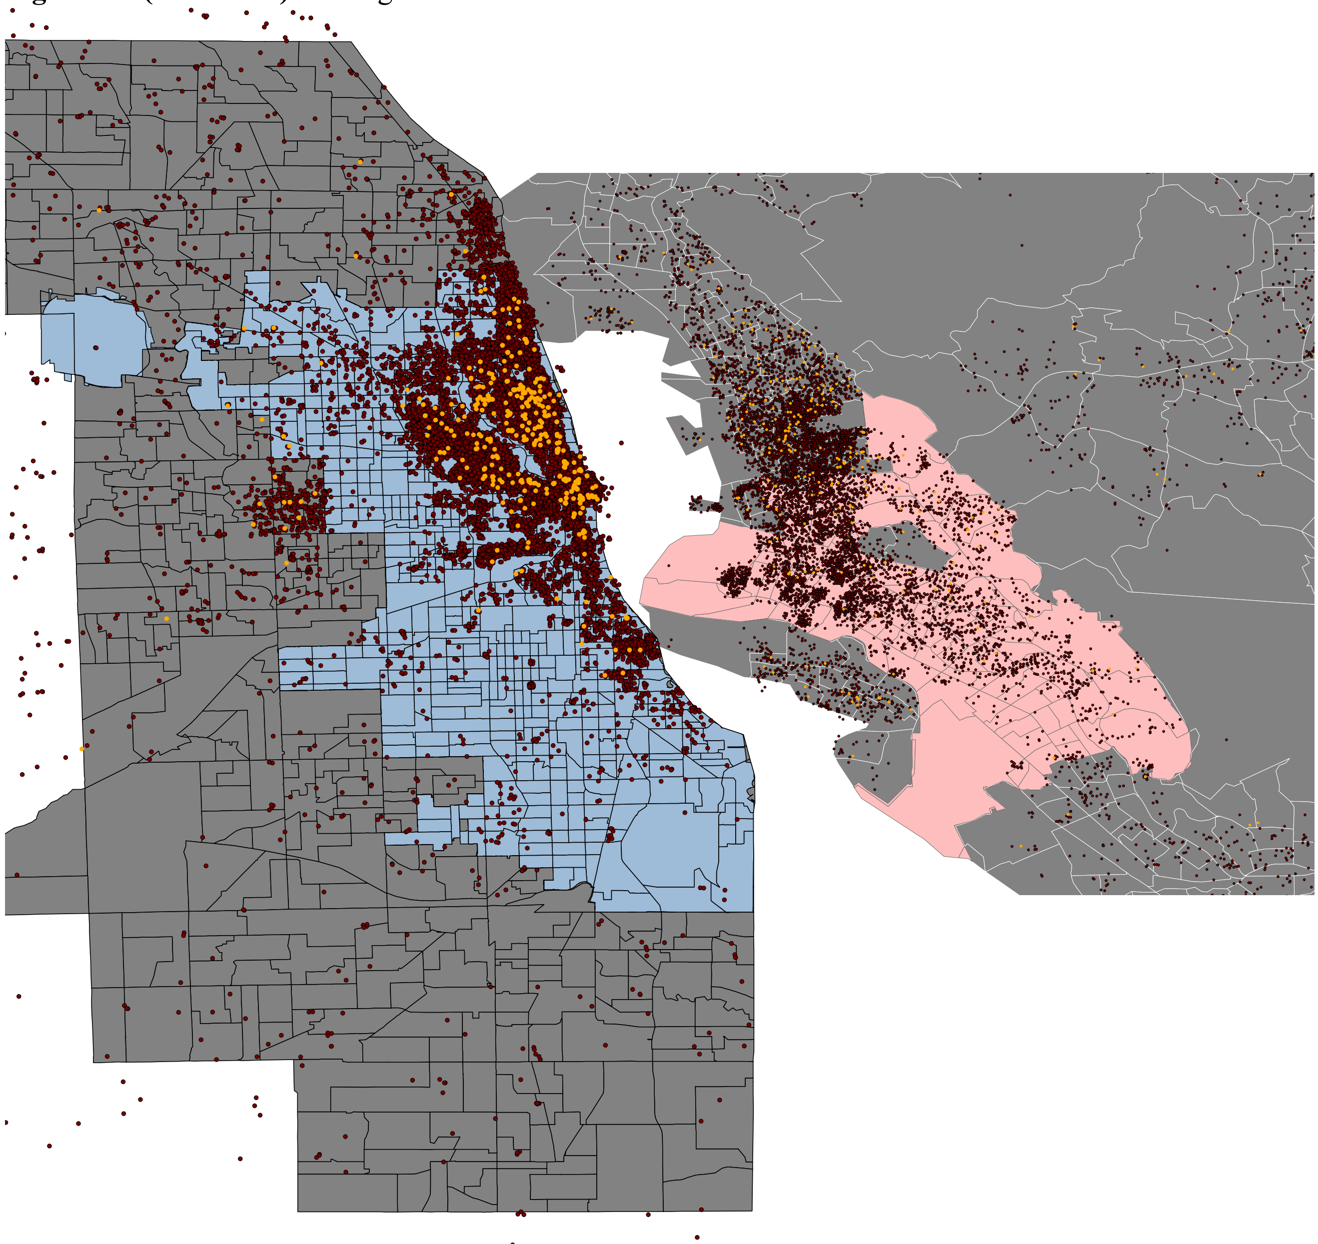

Figure 1.5: Histogram of Listing Prices

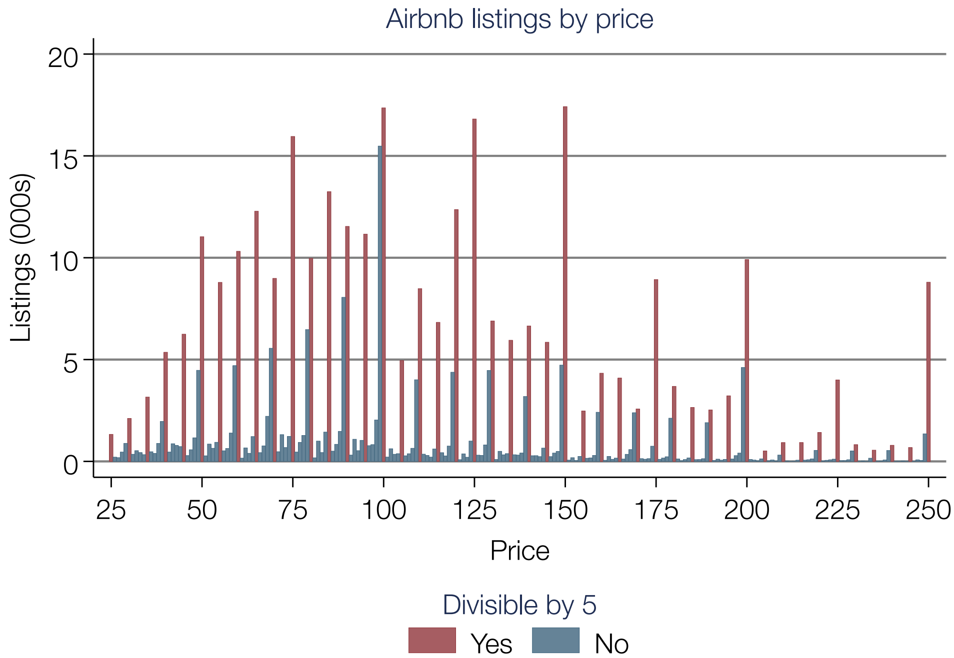

Notes: Figure displays the frequency of listings by price, for all observed listings priced under $250.

Figure 1.6. Event Study Estimates of Policy on Log of Booking Price

Washington DC

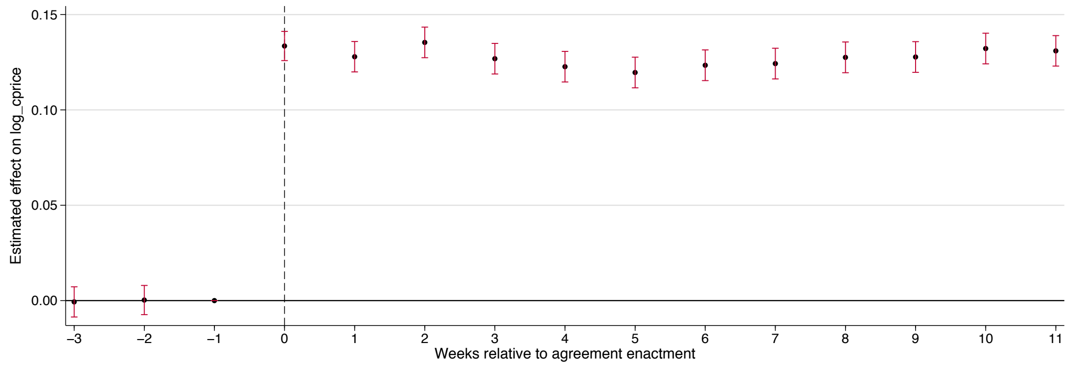

Chicago, IL

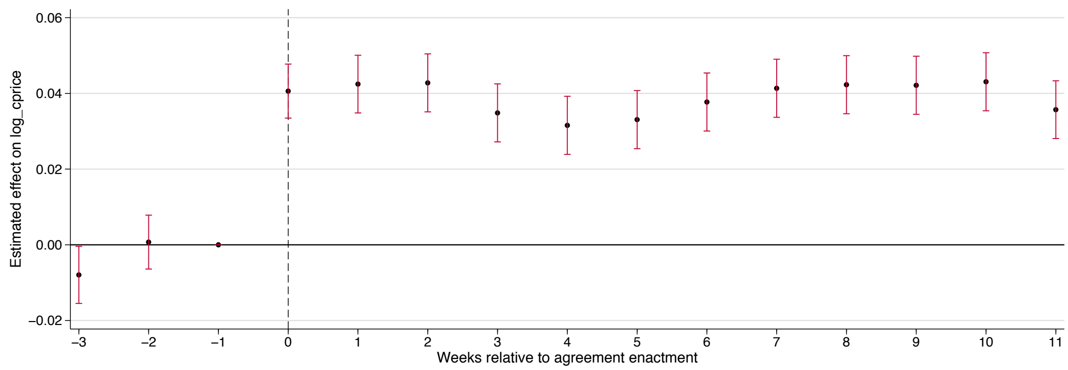

San Diego, CA

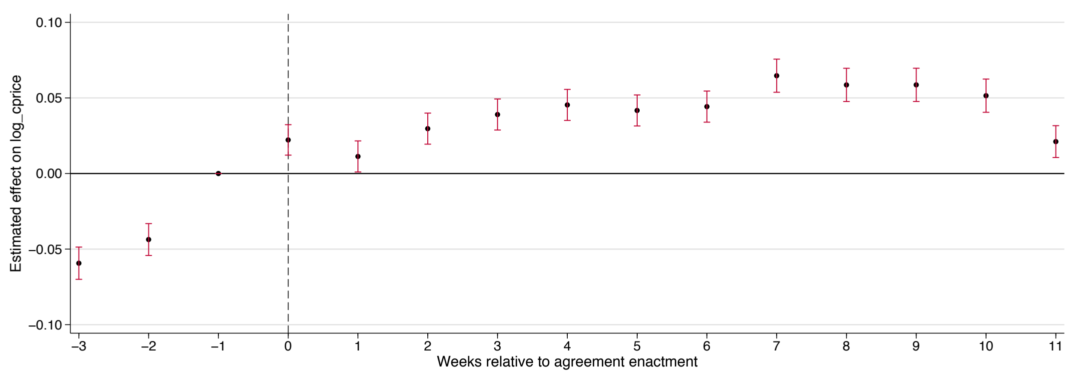

Figure 1.7. Event Study Estimates for Listings with Round Base (Divisible by 10)

Washington, D.C.

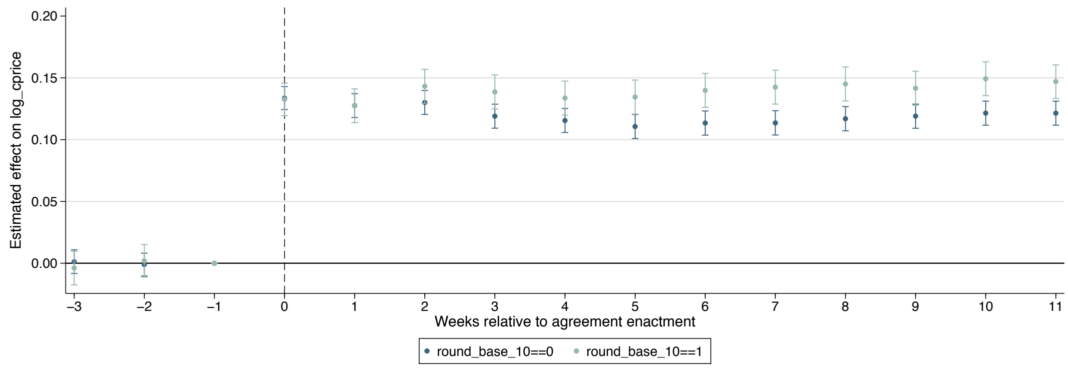

Chicago IL

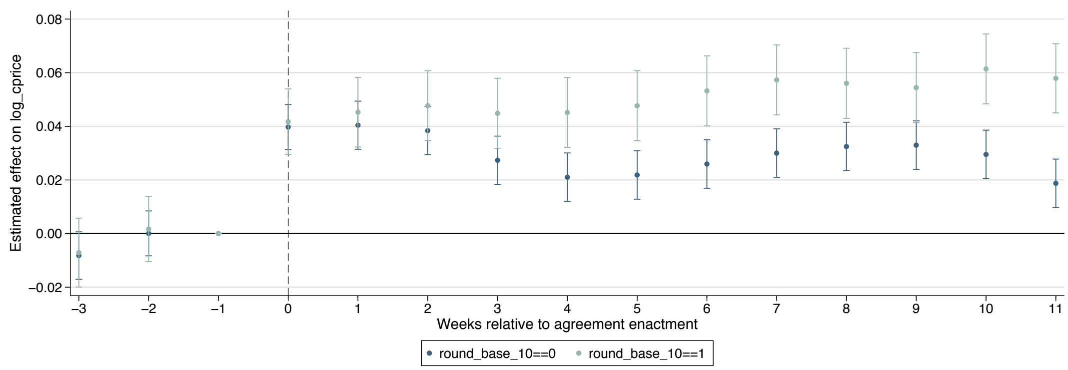

San Diego, CA

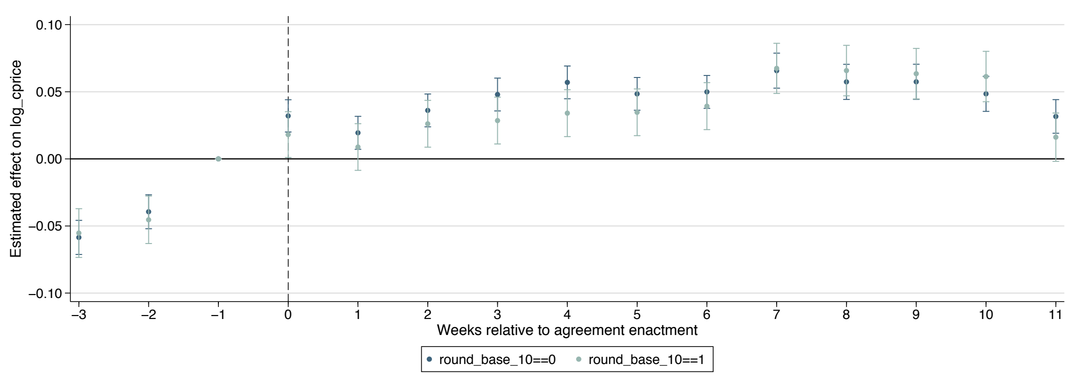

Figure 1.8: ES on Log Price by Pre-policy Correlation with Hotel Price

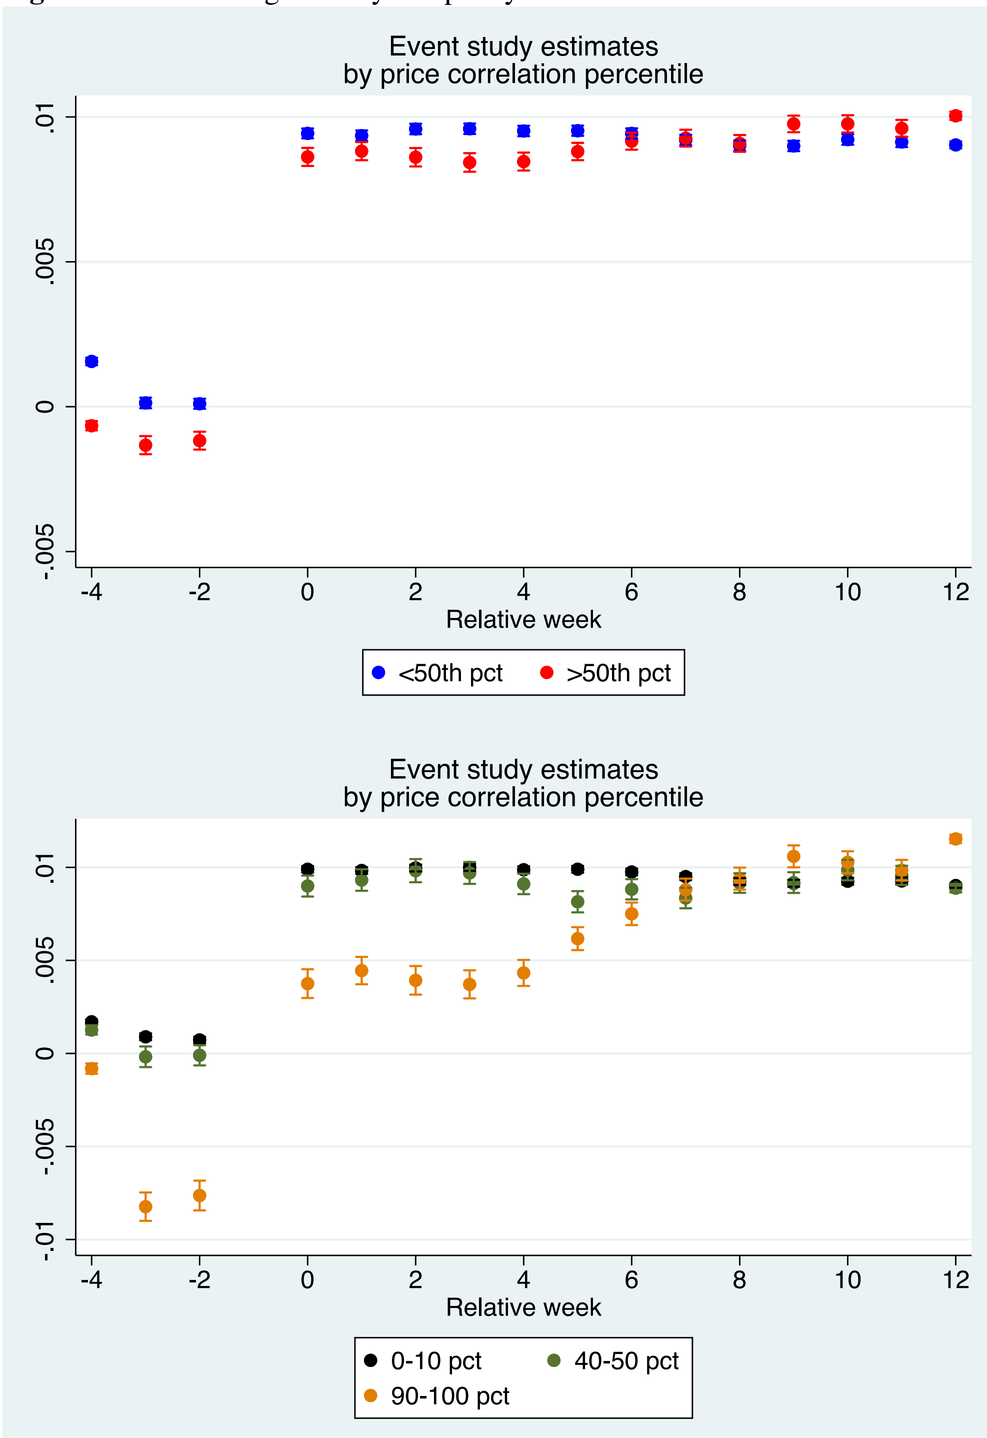

Figure 1.9: Difference in Discontinuity Residuals

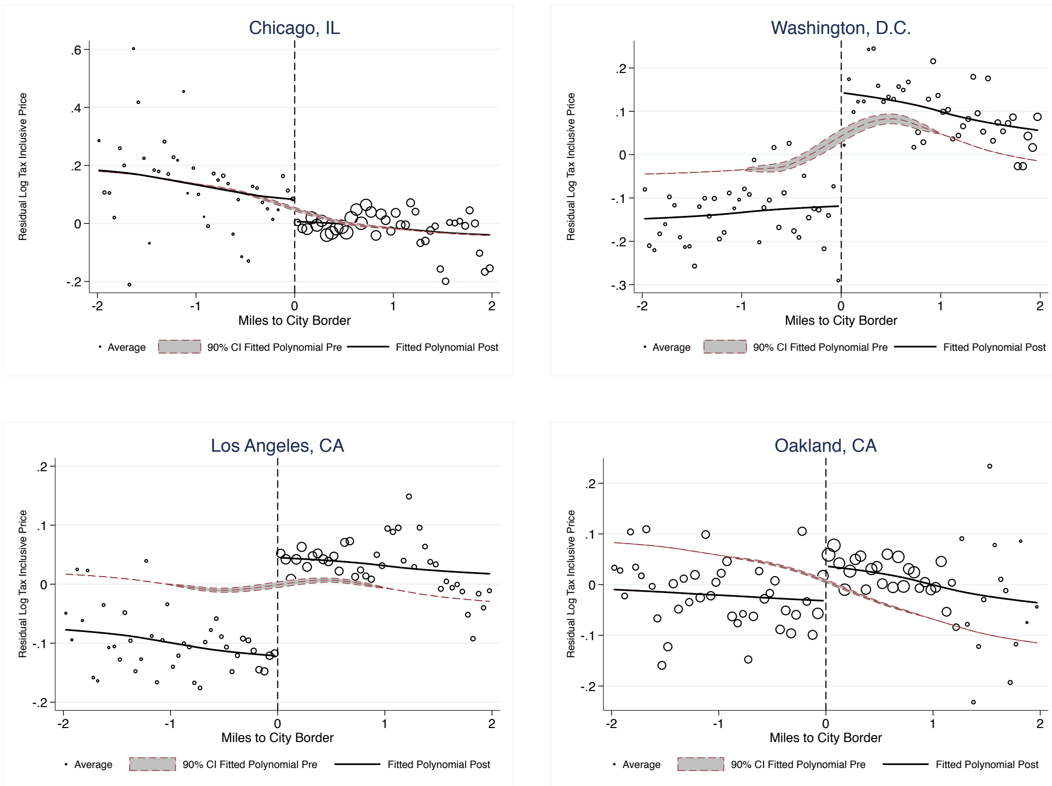

References

Agostini, Claudio A. and Claudia Martinez A. (2012), "Tax Credits Response to Tax Enforcement: Evidence from a Quasi-Experiment in Chile," working paper, Universidad Adolfo Ibanez and Universidad da Chile.

Alm, James, Edward Sennoga, and Mark Skidmore. (2009) "Perfect competition, urbanicity, and tax incidence in the retail gasoline market," Economic Inquiry 47:1, p.118-134.

Balke, Nathan S. and Grant W. Gardner. (1991) "Tax Collection Costs and the Size of Government," working paper, Southern Methodist University.

Barnett, Paul G., Theodore E. Keeler, and Teh-wei Hu. (1995) "Oligopoly structure and the incidence of cigarette excise taxes," Journal of Public Economics 57:3, p. 457-470.

Besley, T. (1989) "Commodity taxation and imperfect competition: A note on the effects of entry," Journal of Public Economics 40, p. 359-367.

Besley, Timothy J. and Harvey S. Rosen. (1999) "Sales taxes and prices: an empirical analysis," National Tax Journal 52, 157-178.

Chernick H. and A. Reschovsky. (1997) "Who pays the gasoline tax?" National Tax Journal 50, 157-178.

Chetty, Raj. 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy*, 1(2): 31–52.

Chouinard, Hayley and Jeffrey M. Perloff. (2004) "Incidence of federal and state gasoline taxes," Economics Letters 83, p. 55-60.

De Paula, Aureo and Jose A. Scheinkman. (2010) "Value-Added Taxes, Chain Effects, and Informality," American Economic Journal: Macroeconomics 2:4, p. 195-221.

Delipalla S, O'Donnell O. (2001) "Estimating tax incidence, market power and market conduct: The European cigarette industry," International Journal of Industrial Organisation 19, p. 885-908.

Doyle, Joe J. and Krislert Samphantharak. (2008) "$2.00 Gas! Studying the effects of a gas tax moratorium," Journal of Public Economics 92, p. 869-884.

Finkelstein, Amy. 2009. "EZ-Tax: Tax Salience and Tax Rates." *Quarterly Journal of Economics*, 124(3): 969-1010.

Fullerton, Don and Gilbert E. Metcalf, "Tax Incidence," in Alan J. Auerbach and Martin S. Feldstein, eds., Handbook of Public Economics, volume 4, New York: North-Holland, 2002, p. 1787-1872.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano. "Policy responses to fiscal restraints: A difference-in-discontinuities design." (2012).

Harju, Jarkko, Tuomas Kosonen, and Oskar Nordström Skans. 2015. "Strategists, generalists and firm-level price setting: Evidence from consumption tax reforms." Unpublished manuscript. http://www.sv.uio.no/econ/english/research/news-and-events/events/guest-lectures-seminars/ofs-seminar/2015/swefin_reforms_05032015.pdf

Kau, James B. and Paul H. Rubin. (1981) "The Size of Government," Public Choice 37:2, p. 261-274.

Kopczuk, Wojciech. (2005) "Tax bases, tax rates and the elasticity of reported income," Journal of Public Economics 89, p. 2093-2119.

Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod. Forthcoming. "Does Tax-Collection Invariance Hold? Evasion and the Pass-through of State Diesel Taxes." *American Economic Journal: Economic Policy*.

Myles, Gareth D. 1989. "Ramsey tax rules for economies with imperfect competition." *Journal of Public Economics*, 38(1): 95-115.

Pomeranz, Dina. (2011) "No Taxation without Information: Deterrence and Self- Enforcement in the Value Added Tax," working paper.

Poterba, James M. (1996) "Retail Price Reactions to Changes in State and Local Sales Taxes," National Tax Journal 49:2, p. 165-76.

Rothstein, Jesse. (2010) "Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence," American Economic Journal: Economic Policy 2:1, p. 177-208.

Saez, Emmanuel, Manos Matsaganis, Panos Tsakloglou. (2012) "Earnings Determination and Taxes: Evidence from a Cohort-Based Payroll Tax Reform in Greece," The Quarterly Journal of Economics 127, p. 493-533.

Sidhu, Nancy D. (1971) "The Effects of Changes in Sales Tax Rates on Retail Prices," Proceedings of the Sixty-Fourth Annual Conference on Taxation p. 720-33.

Slemrod, Joel. (2001) "A General Model of the Behavioral Response to Taxation," International Tax and Public Finance 8:2, p. 119-128.

Slemrod, Joel. (2008) "Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance," National Tax Journal 61:2, p.251-275.

Weyl, E. Glen, and Michal Fabinger. 2013. "Pass-through as an economic tool: Principles of incidence under imperfect competition." *Journal of Political Economy*, 121(3): 528-583.

CHAPTER II

Independent Contractor or Employee?
The Changing Relationship Between Firms and Their Workforce and Implications for the U.S. Income Tax

The number of U.S. workers classified as independent contractors has risen dramatically over the last two decades. While this trend, in part, reflects technological changes in how work is carried out, some of the increase may also reflect firms and workers taking advantage of the legal ambiguity between classifications to obtain preferential tax treatment or to avoid complying with regulations. To study this phenomenon, we exploit a sharp discontinuity in the marginal cost of hiring an employee, created by the 2010 Affordable Care Act. We use U.S. tax returns from 1997-2015 to link firm filings to associated employees and independent contractors. Here, we find preliminary evidence that firms substitute independent contractors for employees in order to qualify for size-based regulatory exemptions. We discuss empirical strategies to distinguish whether this substitution reflects re-organization in the production process (a real response) or misclassification (an evasion response).

2.1 Introduction

Which workers are employees, and how do firms choose? Firms purchase labor from individuals who can be classified as either employees or as independent contractors. Firms' labor and regulatory compliance costs, and workers' labor protections, benefits eligibility, and preferential tax treatment hinge on this classification. In theory, this determination is made according to a series of rules about the nature of the work being performed, and about the relationship between the firm and the worker. In practice, there is substantial legal ambiguity about which classification is appropriate in a given firm-worker arrangement, and enforcement is challenging.

The number of workers in the U.S. who are classified as independent contractors in some or all of their relationships with firms is rising. Some of this increase reflects technological changes in how work is carried out, most notably in the "gig economy" where independent contractors monetize their time, vehicles, and dwellings. Some of this increase may also reflect firms and workers taking advantage of the legal ambiguity between classifications to obtain preferential tax treatment or avoid complying with regulations.

We exploit a sharp discontinuity in the marginal cost of hiring an employee to explore related questions arising from the second category of motivations. First, do firms substitute independent contractors for employees to qualify for size-based regulatory exemptions? Does this substitution reflect re-organization in the production process (a real response), or misclassification (an evasion response)? And, second, to the extent that independent contractors at firms just below the threshold are more likely to be close substitutes with employees, what is the effect of worker status on key outcome variables, such as total compensation or reported income? To answer these questions, we construct a dataset linking firms to workers of both classifications from administrative information contained in U.S. tax returns.

As researchers in other contexts have noted (Schivandi and Torrini (2008); Braguinsky (2011); Garciano et al. (2016)), some firms respond to laws and regulations that bind once a firm reaches a certain size—usually defined by its number of employees—by reducing their size to remain under the compliance threshold. Firms have several ways of accomplishing this. One way is to actually reduce their size by lowering the number of employees. Another way is to misreport their size, particularly if the audit risk is low. But a third way that exists in the U.S. context is to either reclassify employees as independent contractors or otherwise substitute away from using workers of the former type and toward using workers of the latter type. This approach seems particularly attractive given the ambiguity about which classification is appropriate, the relatively low cost of changing classification, and the relatively high cost of compliance with certain laws and regulations.

To study this question, we focus on several laws that impose "size-based" requirements on firms: the Patient Protection and Affordable Care Act (ACA), the Family Medical Leave Act (FMLA), and minimum wage laws. These laws vary in their time of enactment, the jurisdiction (city, state, federal) to which they apply, and the firm size threshold at which they bind. By exploiting this variation, we show that some firms just beyond the thresholds of these laws manipulate their size to avoid complying with them, and that the degree of this manipulation varies with the expected cost of compliance. We also show that the form this manipulation takes is largely not that of firms reducing their overall size—employees and independent contractors. Instead, firms try to limit the impact of this manipulation on their productivity, either by reclassifying existing employees as independent contractors, or restructuring their operations to rely on independent contractors to a greater degree.

Ultimately, we plan to use this exogenous variation in the likelihood that a worker at a particular type of firm—one that is most likely to have taken steps to increase its reliance on independent contractors to avoid compliance—is an independent contractor to estimate the effect of this worker status on a variety of outcomes, including compensation and tax filing behavior.

2.2 Background

2.2.1 Worker Classification

The legal distinction between an employee and an independent contractor is not native to tax law: it originates in common law principles of vicarious liability. Consider the doctrine of *respondeat superior*, under which parties can be liable, or legally responsible, for acts of their agents. The underlying logic is that of incentive alignment—the doctrine extends a financial penalty for causing harm to the party with power to instruct and supervise the agent to ensure the work is done with sufficient care to avoid causing harm. In a similar vein, in the U.S. establishment of an employer-employee depends to a large extent to which the purchaser of services controls how the work is completed, such as whether or not the purchaser provides tools, dictates the timing of work, and to the extent to which the purchaser has financial control over the service provider.

These factors have been codified, with minor variations, into state law as a multi-factor balancing test that requires the arbiter to weigh the relative importance of several features of the relationship to determine whether the worker is an employee. This approach has been adopted by a number of state and federal agencies for regulatory purposes. However, enforcing the distinction in this context has proved problematic: evaluating the holistic nature of the relationship between a worker and her firm is factually intensive and requires significant commitment of agency audit resources. Consider the indeterminate guidance provided by the IRS to potential employers in deciding whether a worker is an independent contractor:

> Businesses must weigh all these factors when determining whether a worker is an employee or independent contractor. Some factors may indicate that the worker is an employee, while other factors indicate that the worker is an independent contractor. There is no "magic" or set number of factors that "makes" the worker an employee or an independent contractor, and no one factor stands alone in making this determination. Also, factors which are relevant in one situation may not be relevant in another.

Despite the cost and obvious difficulty of enforcing such a standard, several features of the tax code depend on worker classification. Independent contractors are entitled to claim "above the line" business cost deductions and are not subject to payroll or income tax withholding. In addition, preferential tax treatment of non-wage compensation is only available to employees.

2.2.2 Firm size-based regulation

Our design relies on laws that create a discontinuity in the cost of hiring an additional employee.[1] There are at least two substantial size-based laws enacted in years covered by the CDW. First, certain provisions of the 2010 Affordable Care Act apply only to firms with more than 50 employees. Second, the 1993 Family Medical Leave Act (FMLA) similarly mandated that firms with 50 or more employees provide unpaid, job-protected leave for specified family and medical events, raising the expected cost of hiring an additional employee above the

[1] These laws and regulations apply differentially to firms who are below or above an arbitrary threshold number of employees.

threshold. From 1999 to 2010, several states passed "binding" family leave statutes that extend these benefits to smaller firms.

One of the most significant size-based regulations in the U.S. is contained in the Affordable Care Act (signed in 2010 by President Obama). In one of several provisions designed to encourage enrollment in health insurance, the ACA's "employer mandate" requires all large employers to provide essential qualifying coverage. Whether or not a firm is "large" is based on its full-time-equivalent (FTE) employment in the calendar year prior to the coverage year, with 50 FTE employees triggering the mandate. Part-time employees contribute to size in proportion to their hours worked.[2] Employers who do not comply are subject to either a Section 4980H(a) penalty or a Section 4980H(b) penalty. The employer mandate was implemented in 2014 with a significant employer subsidy, i.e. "transition relief." No subsidy was provided in 2015.

For coverage year 2017, the 4980H(a) penalty is $2,265 per full-time employee (the first 30 full-time employees are exempt) on the payroll during the coverage year, prorated by month. The 4980H(b) penalty is $3,398 for each full-time employee that buys coverage on the exchanges, capped at $2,265 per full-time employee on the payroll during the coverage year. Neither penalty is deductible from the employer's business income tax, which makes it more expensive than the same dollar amount paid as employee salary. In total, the prior-year hire that triggers the large-employer designation–i.e. that raises the number of FTEs above 50—costs as much as $68,987 (in addition to employee's wage and benefits).

2.3 Empirical Approach

[2] Section 4980H(c)(2)(E) of the Internal Revenue Code, as amended by the ACA, says that the conversion factor from part-time employees to full-time employees is the ratio of the former group's monthly work hours to 120. For example, if February had exactly four work weeks, then every employee working 15 hours per week would count as one half of a full-time equivalent for the month of February.

This section describes our empirical approach for determining whether firms substitute independent contractors for employees to avoid size-based regulations; whether this substitution reflects a real or evasion response; and the effect of worker status on outcomes such as total compensation or reported income. The approach relies on being able to detect unusual bunching in the distribution of firm size, and on being able to link workers to firms.

2.3.1 Bunching and substitution toward independent contractors

Consider the histogram of firm size shown in Figure 2.1a(i).[3] As is commonly the case, firm size here follows a power law distribution, shown as the overlaid dashed line. In the absence of any regulations that create large distortions at points along this distribution, we would expect to see a pattern like this in the data.

Now consider the histogram of firm size shown in Figure 2.1a(ii). Here, a size-based regulation imposes a large cost to hiring the 50th employee. In response, some firms choose to manipulate their size in order to evade the regulation by changing their reported size to be less than 50 employees. This phenomenon creates excess mass just to the left of the threshold, shown in the figure by the light blue bars.

There are multiple avenues for a firm to reduce its reported size to below 50 employees, including reducing its total workforce by firing employees, or simply by misreporting. The channel we focus on, however, is a substitution toward independent contractors. Using data on workers—both employees and independent contractors—linked to firms, we can calculate the

[3] Firm size is measured in full-time equivalent (FTE) units, one that is commonly used in size-based laws and regulations.

share of a firm's workforce that is comprised of independent contractors. Suppose we compute the average value of this metric for each firm size bin, as shown in Figure 2.1b.

The increase in the independent contractor share of the total workforce among firms to the left of the threshold suggests that these firms are substituting away from employees and toward independent contractors as a means of evading the regulation.

2.3.2 Evasion or a real response

Without being able to observe how firms carry out this substitution, it is difficult to determine whether it takes the form of an evasion response like misclassification, or a real response like restructuring of the production process to rely on independent contractors to a greater degree.

One way of disentangling these two narratives is to look for direct evidence of workers having their status reclassified. In other words, relying on the panel nature of the data, construct a measure at the worker-level that reflects whether worker *i* in year *t* is an independent contractor with firm *f* after having been an employee in year *t-k*. Then, use OLS to estimate:

$$\bar{y}_{ft} = \sum_{s=1}^{S} \pi_k FTE_{ft}^{s} + \sum_{g=0}^{G} \rho_g D(FTE_{ft}, g) + \epsilon_{ft} \tag{1}$$

The first term in this expression, a polynomial of order *S* in firm size, captures any underlying relationship between firm size and the average rate at which employees are reclassified as independent contractors with the same firm each year. The second term in this expression includes a series of indicators for whether firm *f*'s size falls within the bin denoted by *g*.

Let $g = 0$ denote the bin immediately to the left of the threshold. As the bin where firms appear to be bunching the most and where the independent contractor share is highest, we estimate equation (1) and test whether $\widehat{\rho_0} > 0$; that is, whether that bin has a higher-than-expected share of employees who are reclassified. If so, this provides suggestive evidence that some firms in this section of the size distribution arrived there by misclassifying their former employees as independent contractors.

Another way to help distinguish these two narratives is to examine whether firm productivity varies significantly in the neighborhood around the firm size threshold. If firms are merely reclassifying workers without substantively changing their behavior, then productivity—as measured, for example, by revenue—should not be adversely affected by the policy. On the other hand, if firms are making substantive changes to their production process, then this should cause a reduction in productivity relative to the optimum in the absence of any size-based regulation.

2.3.3 *Effect of worker status*

Conditional on obtaining evidence that firms manipulate their reported size by reclassifying employees as independent contractors, we can use the variation in worker status that this creates to estimate the effect of status on worker outcomes of interest.

Let C_{ift} be an indicator for whether worker *i* in year *t* is an independent contractor with firm *f*. Then, use OLS to estimate with worker-level data:

$$C_{ift} = \beta_1 X_{it} + \sum_{s=1}^{S} \pi_k FTE_{ft}^{s} + \sum_{g=0}^{G} \rho_g D(FTE_{ft}, g) + \epsilon_{ift} \tag{2}$$

As in equation (1), we are interested in whether $\widehat{\rho_0} > 0$, or whether a worker at a firm in that bin has a higher-than-expected likelihood of having independent contractor status. If this proves to be the case, we can then estimate equation (3) using two-stage least squares:

$$y_{ift} = \beta_1 X_{it} + \sum_{s=1}^{S} \pi_k FTE_{ft}^{s} + \sum_{g=1}^{G} \rho_g D(FTE_{ft}, g) + \beta_2 \widehat{C_{ift}} + \epsilon_{ift} \qquad (3)$$

To interpret β_2 as the causal effect of independent contractor status on y, we must first explicitly state and justify the underlying assumption. As an example, suppose the size-based law in question is the ACA, which mandated that employers with over 50 FTE employees offer them health insurance. If firms close to the 50-employee threshold respond to the ACA by shifting below it and rely on reclassification to do so, then it is plausible that $\widehat{\rho_0} > 0$ in equation (2).

The primary assumption needed to justify the use of the bin indicator as an instrumental variable is that it only affects our outcomes of interest through its effect on C_{ift}, a worker's status. Put differently, among workers at firms in this bin, the variation in their status needs to be unrelated to these outcomes. Because we believe that firms, rather than workers, motivated these status changes, this seems plausible. Of course, it may be the case that worker outcomes may be correlated with working for the type of firm that would engage in such manipulation; to mitigate this, the vector of worker characteristics, X_{it}, can also include the industry in which they are employed.

The empirical design in equations (2) and (3) can be further strengthened by leveraging the panel aspects of the data, as well as information on compliance with the law or regulation prior to the enactment of a size-based threshold. For example, consider equations (2') and (3'):

$$C_{ift} = \beta_1 X_{it} + \sum_{s=1}^{S} \pi_k FTE_{ft}^{s} + \gamma_t + \beta_2 T_f + \beta_3 T_f POST_t + \sum_{g=0}^{G} D(FTE_{ft}, g)\,[\rho_g^{(1)} + \rho_g^{(2)} POST_t + \rho_g^{(3)} T_f + \rho_g^{(4)} T_f POST_t] + \epsilon_{ift} \tag{2'}$$

$$y_{ift} = \beta_1 X_{it} + \sum_{s=1}^{S} \pi_k FTE_{ft}^{s} + \gamma_t + \beta_2 T_f + \beta_3 T_f POST_t + \sum_{g=0}^{G} D(FTE_{ft}, g)\left[\rho_g^{(1)} + \rho_g^{(2)} POST_t + \rho_g^{(3)} T_f\right] + \sum_{g=1}^{G} \rho_g^{(4)} D(FTE_{ft}, g)\, T_f POST_t + \beta_4 \widehat{C_{ift}} + \epsilon_{ift} \tag{3'}$$

The indicator T_f captures whether firm *f* was non-compliant with the policy prior to its implementation. Now, the exogenous variation used to identify changes in contractor status comes from comparing workers in the bin immediately to the left of the threshold, for firms not previously in compliance (for whom the policy has the most "bite"), and in the periods after the policy took effect, to other workers.

2.4 Data

To conduct the analysis, we construct a dataset using the universe of digitized tax filings in the U.S. for tax years 1997-2015. We link individual income tax returns with information reports,[4] and link individuals to their employing firms. This new panel dataset allows us to observe traditional employees and contractors working for the same firm, linked with employment and compensation characteristics of the employing firm. Additionally, we use combinations of characteristics derived from income tax returns and information reports to define what constitutes

[4] In particular, we use income tax reports Form 1040, Schedule C and Schedule SE, and information reports W-2, 1099-MISC and 1099-K.

independent contractor labor for this paper. Following the work of Knittle et al. (2011), we use restrictions on the type and amount of expenses reported to differentiate small businesses from independent contractors. Other criteria, such as the number of firms from which an individual receives a 1099 and the share of individual income coming from 1099s, are used to create consistent definitions of independent contractors relative to other types of self-employment labor.

We expand on previous methods of identifying independent contractors in administrative data by including incorporated independent contractors (single owner LLCs) and test whether the trends we identify are consistent under more(less) inclusive criteria.

There are several notable advantages to using tax return data to study questions related to independent contractors. First, and most importantly, it is possible to link independent contractors to all firms that compensate them. Unlike census-based linked employer-employee datasets such as the LBD, which are aggregated from state unemployment insurance records, tax return data allows us to identify all individuals that a firm compensates for labor services through annual (mandatory) firm information reporting to the IRS. Employees are issued form W-2, while independent contractors are issued Form1099-MISC or, in some cases, Form 1099-K.

Second, we can link workers to their individual income tax returns (Form 1040). Making this link allows us to observe information relevant to the nature of the firm-worker relationship, such as the degree to which the worker relies on income from the firm, the length of time a worker has been associated with a specific firm, and whether or not the same worker has switched classification while working for the firm. Form 1040 also contains various outcome variables of interest: detailed information on deduction taking, total reported income, and some information on non-wage compensation.

Finally, we can access a rich set of firm characteristics by linking employing firms to their business income returns and other filings. These forms include information on firm deductions for employee benefits, which we use to identify the subset of firms that were unlikely affected by the regulation because they were already providing the benefit. It also provides variables other than labor input which should vary smoothly with the number of employees, such as total investment or revenue per employee.

There are also some important draw-backs to these data. We are unable to observe hours worked, which makes it more difficult to directly compare compensation across workers. Also, for regulations that are not enforced by the IRS, we do not have the exact firm size measure used by the regulatory agency to determine whether the firm is subject to the regulation (i.e. the FMLA is enforced by the Department of Labor, not the IRS).

2.5 Descriptive Analysis Results

2.5.1 Sample

The analysis sample, a repeated cross-section, is constructed by drawing a 2 percent random sample from the universe of firms with at least one employee, separately during each tax year from 2000-2015.[5] Firms are identified by their TINs, as they appear on Form W-2 (for employees) or Form 1099-MISC (for independent contractors). When collapsing from the worker-firm-year level to the firm-year level, information about the compensation distribution of the firm-year (e.g., median W-2 compensation, median 1099-MISC compensation) is preserved.

[5] This restriction is partly for convenience: all firms with at least on employee in the cleaned SOI Databank, which is already linked to W-2 workers. However, as our research question explores substitution between employees and ICs, this limitation is also a sensible one.

The sample contains 2,032,732 firm-year observations, or approximately 130,000 firms per tax year. Firm size, as measured by the number of employees, follows a power law distribution, with very few large firms in the population (See Figure 2.2). As a result, within each tax year, the sample contains approximately 3,000 firms with over 100 employees, or 2 percent of firms. In the future, we will consider drawing a stratified sample with a higher sampling rate for large firms to attain a sufficient number of them.

We define two broad measures of independent contractor (IC) usage by firms. Extensive margin IC usage is an indicator for whether a firm hired at least one IC in a given tax year. Intensive margin IC usage is one of several continuous measures representing the degree to which the firm made use of ICs in a given tax year, such as the ratio of ICs to total workers (employees and ICs), or the ratio of IC compensation (from 1099-MISCs) to total compensation (from W-2s and 1099-MISCs).

2.5.2 Firms are hiring more independent contractors.

Figure 2.3a shows the change in firms' extensive margin IC usage by quartiles of the firm's median wage.[6] Extensive margin IC usage rose gradually from 2000 to 2010, and then increased sharply in 2011. By 2015, extensive margin IC usage rose by almost 20 percent for most firms, and by almost 30 percent for firms with low median wages, since 2000. The increase in extensive margin IC usage was also shared across firms of different size, though was highest in small firms with four or fewer employees (Figure 2.3b).

[6] Firms are assigned to quartiles based on their median employee compensation within the tax year. Similar results obtain when assigning firms to quartiles based on their average, 25th percentile, 75th percentile, or 90th percentile employee compensation.

The growth in any IC usage by small and low wage firms masks the fact that larger and higher wage firms, which had higher levels of IC usage to begin with, grew in the intensity with which they used ICs. Figures 2.4a and 2.4b show analogous information but for changes in an intensive margin measure of IC usage: the ratio of ICs (1099-MISCs) to total workers (W-2s and 1099-MISCs). Higher-paying firms saw their IC usage grow by over 8 percent, compared to under 4 percent for lower-paying firms. Likewise, large firms saw their IC usage grow by over 10 percent between 2000 and 2015.

For both of these sets of trends, there is a distinct break after tax year 2009, presumably related to the recession. These trends are also likely understated, as we do not count individuals who received a 1099-K for labor services they provided as ICs. They are consistent with the pattern of individuals increasingly working as ICs found by Jackson, Looney and Ramnath in their OTA working paper, which also uses CDW data.

2.6 Evidence of bunching in the FTE distribution

This section provides evidence that some firms respond to the employer mandate by reducing their full-time employment below 50. We employ standard bunching analysis techniques to visually identify missing mass to the right of the threshold, possibly created by firms that shifted, or "bunched," just below 50 FTE.

2.6.1 Data from ACA related firm filings

In contrast to the 2% sample used in the descriptive analysis, this section uses the universe of firm tax filings for firms with fewer than 500 FTE.[7] Employers mandated by the Affordable Care Act (ACA) to provide their employees with health insurance or pay a penalty are called Applicable Large Employers (ALEs). An ALE is defined as an employer with an average of 50 or more full-time employees (FTE) in the previous tax year. In addition to imposing a potential penalty, the ACA also imposes specific information reporting requirements on ALEs. In the current tax year, ALEs must file one Form 1095-C for each employee, as well as a single Form 1094-C that accompanies the transmittal of the Forms 1095-C. Each Form 1095-C indicates whether or not the firm offered health insurance coverage to a given employee.

Form 1094-C contains additional identifying information about the firm, including the number of FTE it employed in each month of the current tax year, which is used to determine its ALE status for the next tax year. In addition, the 1094-C reports whether the firm is part of an aggregated ALE group, whether it provided minimum qualifying coverage, and its number of employees for the purpose of assessing penalties for failure to offer coverage,[8] all at the monthly level. Tax year 2016 was the first year in which the IRS assessed penalties for ALEs that did not offer mandatory coverage for all employees.

2.6.2 Analysis of the firm size distribution ("Bunching results")

Firm size tends to follow a power law distribution, with a long right tail of large firms (Figure 2.2). This appears to be true both internationally (Garicano, Lelarge, van Reenen 2016)

[7] FTE is a lower bound for total employment—all full-time or part-time employees that count towards ALE status are issued W-2s, but there are some W-2 workers that do not contribute to FTE (i.e., because they are qualify for the "transitioning" employee exemption).
[8] This number may differ from the FTE because employers are not penalized for not providing coverage to part time employees, or employees who work for the firm for fewer than 3 months of the tax year.

and in the U.S. as measured with census data. However, in the presence of size-based regulations (SBRs), this expected power law distribution may not hold if firm size is sufficiently elastic: in response to a cost or other penalty for being above a certain size, for example, some firms may bunch to the left of a size threshold. In this context, because only firms with 50 FTE or more are required to issue Forms 1095-C and 1094-C, we cannot directly observe the distribution of firms on either side of 50 FTE (Onji 2009). Nevertheless, we can still look for evidence of bunching in response to SBRs like the ACA to the right of 50 FTE.

In Figures 2.5a and 2.5b, we plot the number of firms reporting each integer value of FTE from 50 and above in 2015 and 2016, excluding firms that belong to aggregate groups. If the power law distribution held, we would expect the number of firms to be monotonically decreasing from 50 FTE onward; instead, we see a non-monotonic distribution, with fewer firms than expected between 50 and 60 FTE. (After 60 FTE, the series appears to resume the expected pattern.) This non-monotonicity is more pronounced for firms that pay lower wages, and were less likely to already be providing health insurance coverage to their employees prior to the ACA, than it is for firms that pay higher wages (Figures 2.6a and 2.6b).

2.6.3 Interpretation of bunching patterns

Though suggestive, these patterns are not definitive proof of bunching by firms in response to the ACA or other SBRs that anchor on the 50 FTE threshold. For example, the mandatory reporting obligation for firms at or above 50 FTE may result in firms near this threshold being uncertain about whether they are required to submit Forms 1095-C and 1094-C. If this uncertainty is increasing the closer a firm is to the threshold, and if some of these uncertain firms choose not to file, then we can expect to see a pattern similar to what we observe.

Despite this caveat, there are several reasons to think the missing mass we observe is indeed a response to the ACA. First, although some other SBRs bind at the 50-employee threshold (e.g., FMLA, ADA), no other federal SBR use the same definition of firm size as the ACA (i.e., no other SBR counts employees for the purposes of the exemption in the same manner that the ACA counts FTE). Though there are similarities across the definitions, the incentives created by these other SBRs should not create a discontinuity exactly at 50 FTE. However, we see the fewest firms at exactly 50 FTE (e.g. Figure 2.5b).

Second, the cost of violating these other SBRs is not expected to vary systematically with the cost of violating the ACA's firm mandate. Yet if we focus on firms that were more likely to provide their employees with insurance coverage prior to the ACA, such as firms paying above median wages, we notice that bunching behavior is significantly reduced.

In the canonical bunching papers in the public finance literature, the goal of detecting and measuring excess mass in the taxable income distribution is to arrive at an elasticity of taxable income (ETI) (see, e.g., Saez 2008; Kleeven and Waseem 2013; Chetty et al. 2012). The ETI estimate can then be used as a general parameter to estimate the revenue implications of *any* tax change.

In our context, the general parameter of interest that could be estimated from the bunching we observe is the regulatory cost elasticity of FTE. This parameter would allow us or other researchers to estimate the effects of regulations (or taxes) on firm size. Estimating this parameter in the context of the ACA requires us to define regulatory costs. Complying with the ACA varies with the specific circumstances of the firm: the cost of an essential qualifying insurance plan in the market in which the firm is located, the number of employees who would take-up the insurance if it was offered, etc. In theory, the expected penalty, based on federal maximums, should provide a lower bound on the regulatory cost faced by the firm: if it was less expensive to provide insurance than to pay the penalty, the firm would provide insurance. However, the penalty also varies, though we may be able to get traction on this issue in several ways. The maximum amount of the penalty faced by firms for each uninsured employee is set at a national level, without respect to the local cost of healthcare. Additionally, linking the information on Form 1095-C with the individual's tax return will allow us to determine how many of a firm's employees at the time of the regulation were receiving insurance through another member of their household.

2.7 Directions for Future Analysis

Building on the evidence described in Section 6 suggesting a potential response by firms to the ACA's SBR, we are interested in whether this response takes the specific form of substitution between ICs and employees, either via restructuring of a firm's operations ("real response") or misclassification ("reported response"). The challenge we face is that we are unable to observe firms' FTE for those firms with fewer than 50 FTE; all we can infer from the absence of a Form 1094-C filing is that a firm has fewer than 50 FTE. However, within this group, we can neither directly measure a firm's proximity to the 50 FTE threshold, nor can we order firms to create a relative ranking by FTE.

Though this is a serious limitation, we consider in this section two approaches to advancing the analysis of IC usage in response to the ACA's SBR. The first method relies on leveraging firm characteristics that we can observe for the full distribution to predict FTE, then using the predicted FTE to analyze patterns in IC usage that vary with proximity to the threshold.

The second method limits empirical analysis to firms to the right of the threshold for whom we observe reported FTE. We suggest two examples of testable hypotheses that only require data for firms to the right of the threshold.

2.7.1 Predicting FTE

This approach uses the relationship between pre-ACA observable characteristics and post-ACA FTE, for firms that report FTE and are unlikely to alter their reported size in response to the ACA, in order to predict FTE for all firms in the absence of the ACA. This analysis will proceed in three phases.

First, we will identify firms in the unaffected region of the observed FTE distribution to estimate the relationship with pre-ACA observable characteristics. In theory, firms located anywhere to the right of the 50 FTE threshold face an incentive to relocate to the left of the threshold. In practice, firms' willingness to act on this hinges on their elasticity of FTE with respect to regulatory costs. Larger firms, for whom a real response could be highly disruptive and a reporting response could significantly increase audit risk, are the least likely to respond to the ACA by altering their reported size. In keeping with current practice, we can identify the unaffected region of the FTE distribution visually, and then perform a series of robustness checks to ensure that our prediction results are not sensitive to this choice (see, e.g., Saez 2009, Kleven 2016). Based on Figure 2.5a**,** it appears that the FTE distribution beyond approximately 65 assumes the shape of a power law distribution that we would expect to see in the absence of distortionary SBRs.

Second, we will use machine learning (ML) methods to estimate the relationship between firm post-ACA FTE and pre-ACA observables for unaffected firms, and then predict FTE for all firms (including those who did not file a 1094-C). ML approaches, which seek to minimize out-of-sample prediction error, are ideally suited for this type of exercise.[9] Because we are concerned that certain contemporaneous observable firm characteristics (e.g., number of employees, ICs) may themselves change in response to the ACA, we estimate this relationship using only firms' pre-ACA characteristics. Finally, by estimating this relationship using only firms in the "unaffected" region, we seek to predict the FTE that a firm would report in the absence of the ACA, rather than the potentially distorted FTE a firm would report in response to the ACA.[10]

Finally, we will compare the IC usage of firms that have predicted FTE>50, separately by whether the firm did and did not file a 1094-C. The intuition underlying this comparison is that firms with predicted FTE>50 that did not file a 1094-C are those that, based on their pre-ACA observable characteristics, were predicted to exceed the reporting threshold for the 1094-C yet did not report, possibly by relocating below the threshold through IC and employee substitution. Formally, we can test this using the following specification:

$$y_i = \alpha\widehat{FTE}_i + 1\left[\widehat{FTE}_i > 50\right]\left(\beta_1 + \beta_2 1\left[\text{No 1094-C filing}\right]\right) + \varepsilon_i \tag{4}$$

where the dependent variable, y_i, is a measure of firm i's IC usage, such as the ratio of ICs to total workers or the ratio of IC compensation to total. The coefficient of interest, β_2, represents

[9] In general terms, most ML techniques use cross-validation, or training the model on several subsets of data and evaluating their performance on a complementary set, to avoid over-fitting.

[10] To appreciate the difference between these counterfactuals, one can think of FTE in the absence of the ACA as the number of full time equivalent employees that the firm would hire if there was no discontinuous change in regulatory costs anywhere in the firm size distribution. "FTE contemporaneous to the ACA," in contrast, is the number of FTE that firms would report if subject to the 1094-C reporting requirements, i.e., in a world where all firms were required to report FTE, regardless of the number of FTE. This measure is FTE taking into account firm responses to relocate (by either restructuring or misclassifying). These are both separate from a third counterfactual, "FTE upon audit," which would measure a firm's "true" FTE conditional on using the appropriate classification of employees and ICs.

the difference in IC usage for firms predicted to have FTE>50 but that nevertheless did not file 1094-C. If $\beta_2 > 0$, for example, then IC usage is higher for firms that did not report but were predicted to have FTE>50. Conditional on finding economically meaningful substitution towards ICs for these firms, we could potentially distinguish restructuring and misclassification by estimating the same specification using a productivity metric as the dependent variable (e.g., reported profit). If firms respond to the ACA by restructuring, their productivity should diminish, as otherwise that structure was available to them in the absence of the ACA. On the other hand, if firms simply misclassify employees, their productivity is less likely to be adversely affected.

The primary limitation of this approach is that the relationship between observed pre-ACA characteristics and FTE for firms in the unaffected region may differ systematically from that of firms at or below the threshold. However, we can try to learn about the accuracy of this type of out-of-sample prediction by applying it to other portions of the observed FTE distribution. For example, we could estimate the relationship using firms with FTE>100 and use it to predict FTE for firms that report values between 65 and 100.

We have taken initial steps to try to determine if this strategy is viable. We expect that the number of *any* employees (W-2s) will be a strong predictor of the number of full-time employees. We begin to explore this relationship in Figures 2.7a and 2.7b, which plot these two values against one another in 2015 and 2016. Unfortunately, these figures show the relationship between these two related employment measures is quite noisy. Presumably, exploiting additional information from firm filings will reduce this error considerably.

2.7.2 Limit analysis to firms that reported FTE above the threshold

A different approach is to exclude from the analysis firms that do not file a 1094-C and for which we are unable to observe a reported FTE. Instead, by focusing on firms that file a 1094-C and have a reported FTE, we may still be able to uncover evidence that will allow us to understand the nature of firms' response to the ACA. Specifically, we consider two hypotheses:

First, if firms close to the threshold on the right reclassify employees as ICs in order to move to the left, then the firms that remain on the right and file a 1094-C should be disproportionately those that have a lower elasticity of FTE with respect to regulatory costs. One reason why they may have a low elasticity is because they are in an industry in which it is more difficult to reclassify employees as ICs. This suggests a testable hypothesis concerning the industry composition of firms immediately to the right of the threshold, and whether those industries are disproportionately those in which reclassification is more difficult.

Second, given the frictions involved with firing employees, firms seeking to avoid complying with the ACA may instead opt to delay hiring. This strategy is likely to be more sensible for firms that are not growing rapidly toward a size far in excess of the threshold. As a result, we might expect to see higher average employment growth rates among firms immediately to the right of the threshold than we would otherwise, as firms that would otherwise have FTE>50 and be growing slowly would instead have slowed hiring to keep FTE<50.

2.8 Conclusion

Policymakers and economists are increasingly focused on the changing structure of the U.S. workforce, and in particular the perceived shift away from "traditional" employment relationships toward potentially "alternative" work arrangements, a broad category that includes temporary agency workers, on-call workers, contract workers, and independent contractors.

Recent empirical studies have documented increases in the prevalence of non-standard work arrangements over the last decade.[11] However, there is compelling evidence that use of survey data to identify a worker's legal classification is unreliable, and likely understates these trends. Additionally, there is much less evidence on which factors, on either the supply or demand side, might be behind this increase. Our paper aims to enhance our understanding of the changing relationship between firms and workers, and why this relationship is a critical component of the design of effective tax and regulatory policies (i.e. labor protections and social safety-net policies).

We make two empirical contributions. First, we provide evidence on whether, and to what extent, the use of IC arrangements has increased over the last decade. We find that firms are indeed hiring more ICs, a pattern that we decompose into both extensive margin measures (firm hiring of any ICs) and intensive margin measures (extent of firm hiring of ICs). Our findings, derived from high quality administrative data, are broadly consistent with the increase in IC arrangements measured using individual surveys.

Second, we provide the first evidence of how firms may respond to the ACA by altering their reported size. The employer mandate, which applies to firms with 50 or more full-time employees, appears to induce some firms to report having fewer than 50 FTE (by failing to file a 1094-C): our analysis reveals a substantial "missing mass" in the distribution of firms by FTE just above 50 in the years 2015 and 2016. We find a larger response among firms that pay lower wages and were less likely to have provided their employees with insurance prior to the ACA, for whom the regulation likely had more bite. With additional assumptions, we can translate the missing mass into a more general elasticity of full-time employment to regulatory costs.

[11] See Katz and Krueger (2016) for evidence based on survey data, and Jackson, Looney and Ramnath (2017) for evidence from administrative tax records.

Figure 2.1a. Firm size distribution without and with a size-based regulation

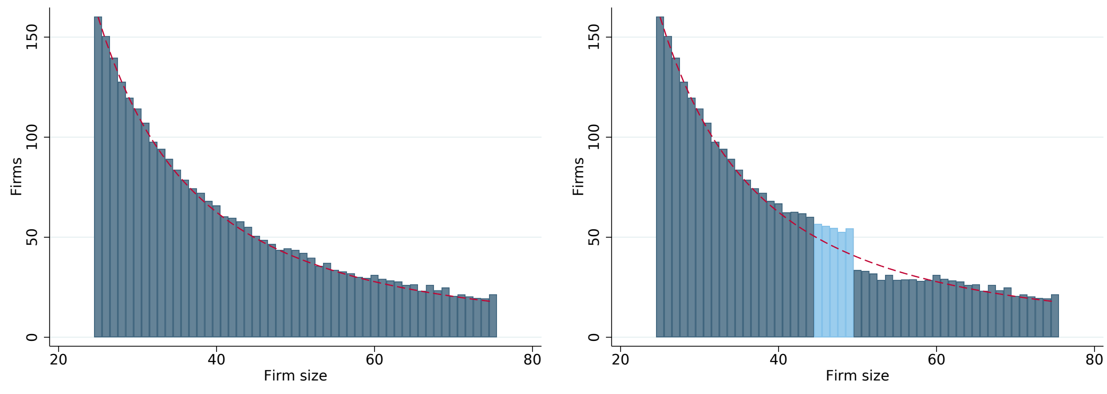

(i) (ii)

Figure 2.1b. Independent contractor share of total workforce

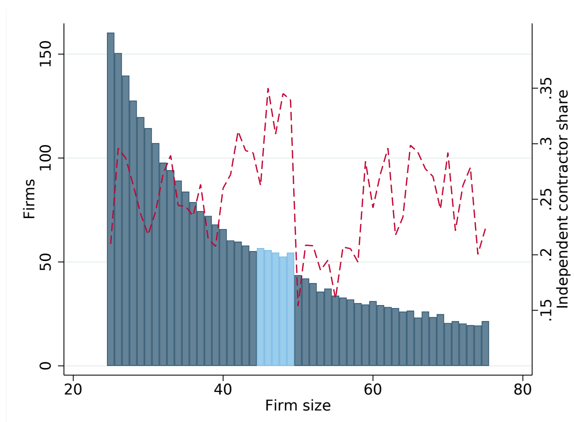

Figure 2.2. U.S. Firm Size Distribution (Total Employment) in 2014

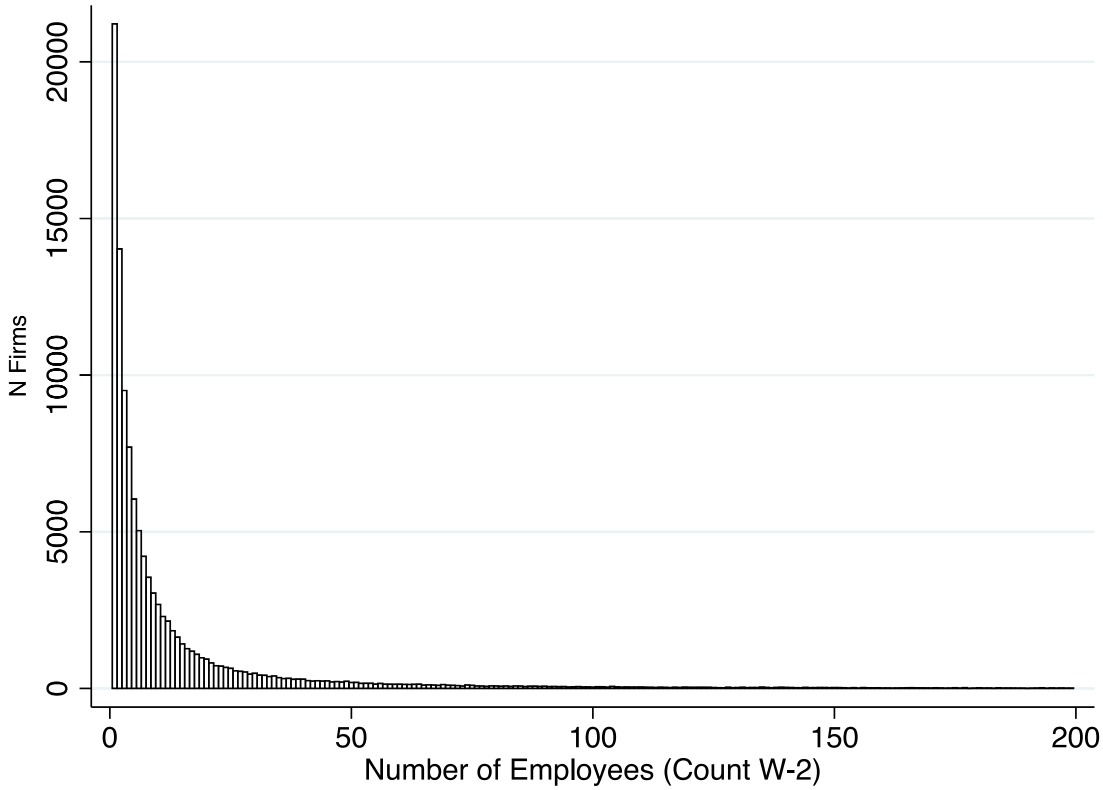

Figure 2.3a. Extensive margin IC use by quantile of firm median wage

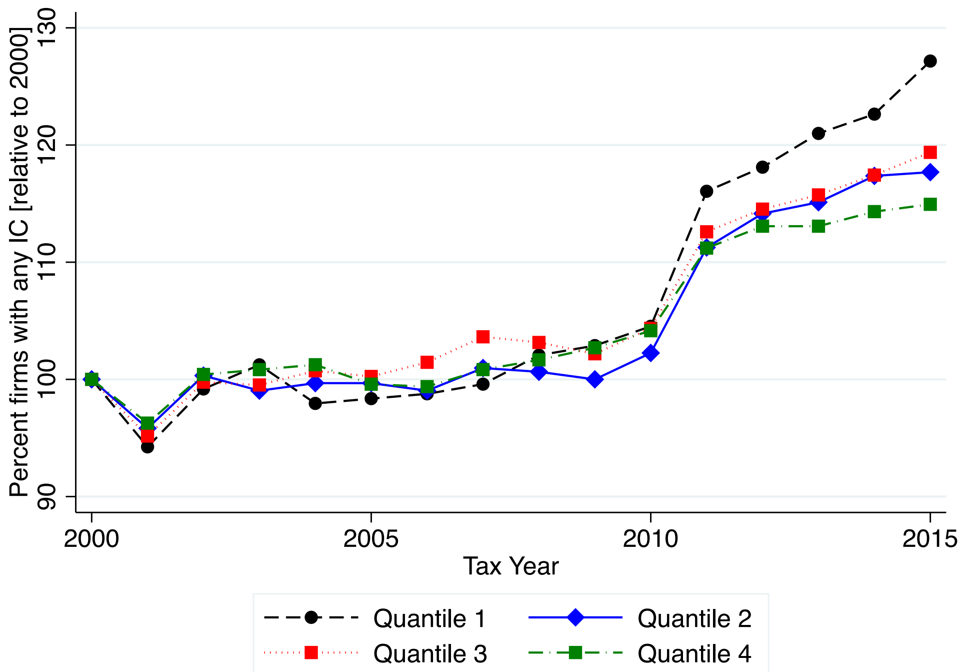

Figure 2.3b. Extensive margin IC use by number of employees

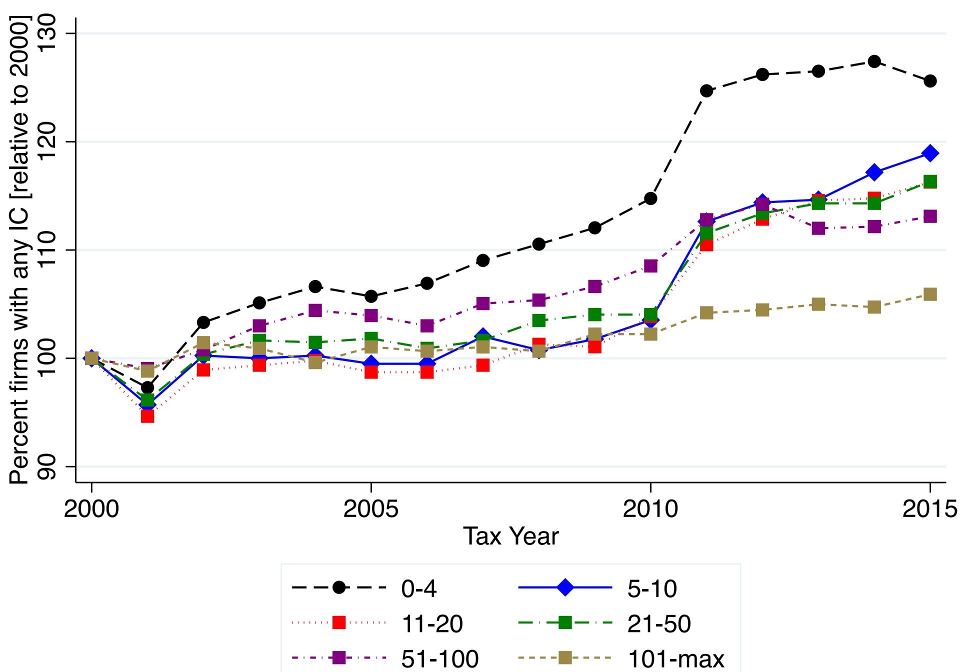

Figure 2.4a. Intensive IC usage by quantile of firm median wage

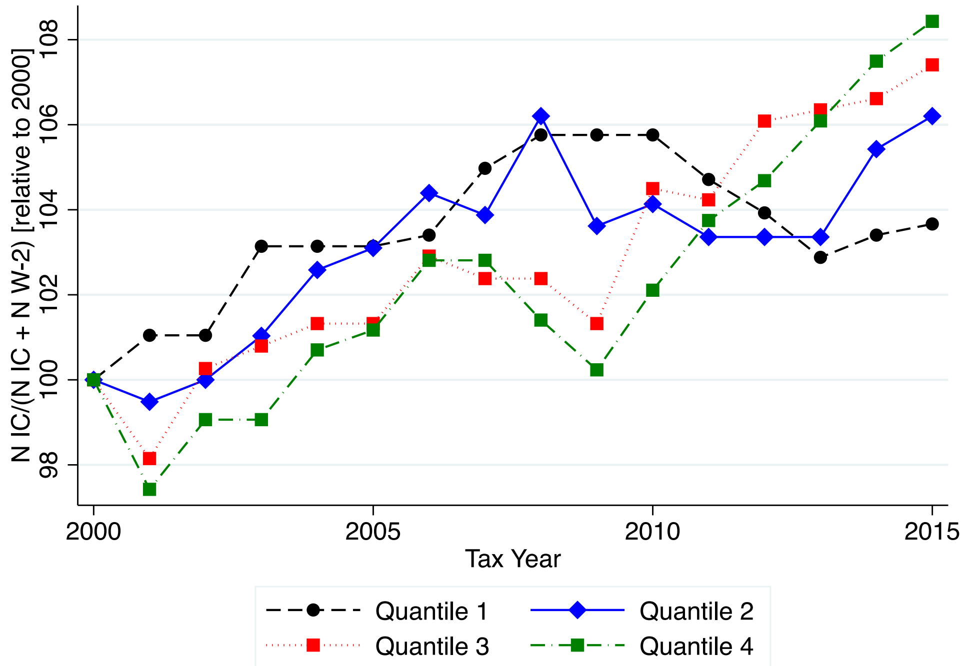

Figure 2.4b. Intensive IC usage by number of employees

Figure 2.5a. Number of firms by FTE (2015)

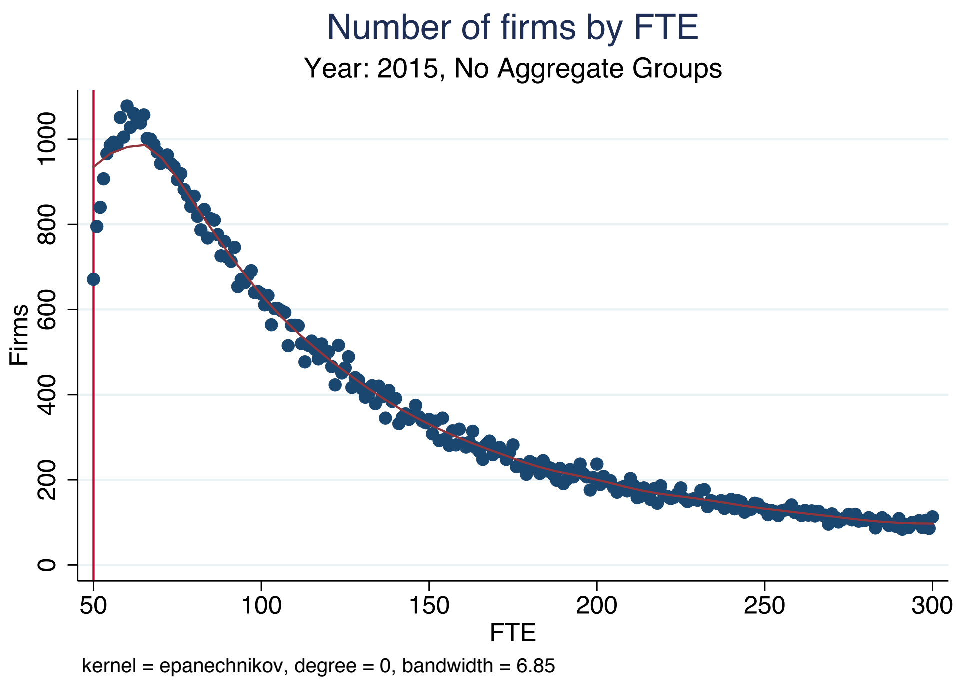

Figure 2.5b. Number of Firms by FTE (2016)

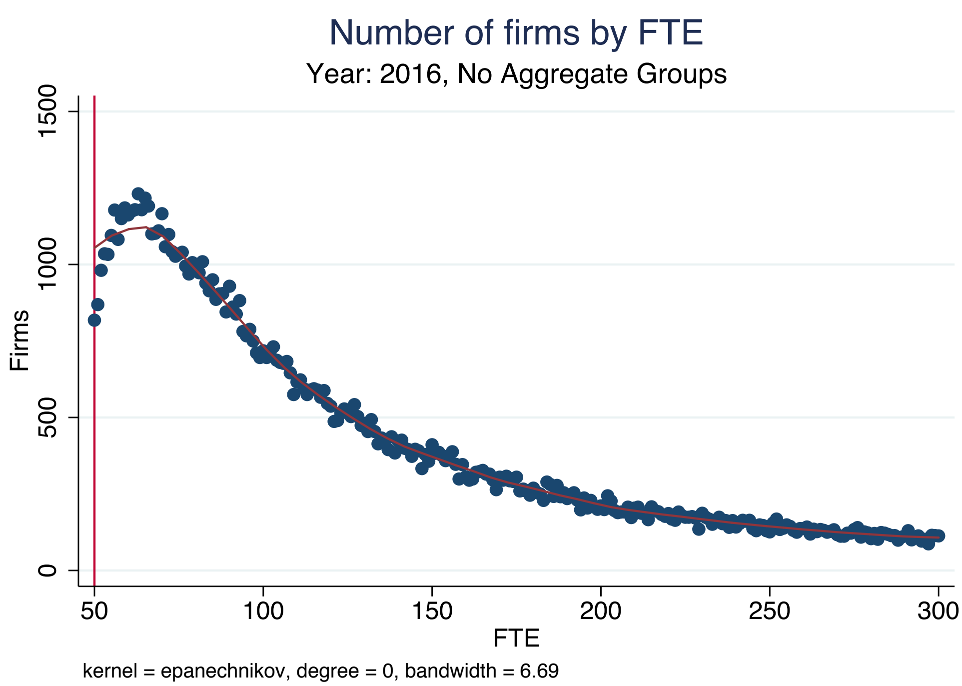

Figure 2.6a. Number of firms by FTE for firms with average wages below 25th percentile (2016, No Aggregate groups).

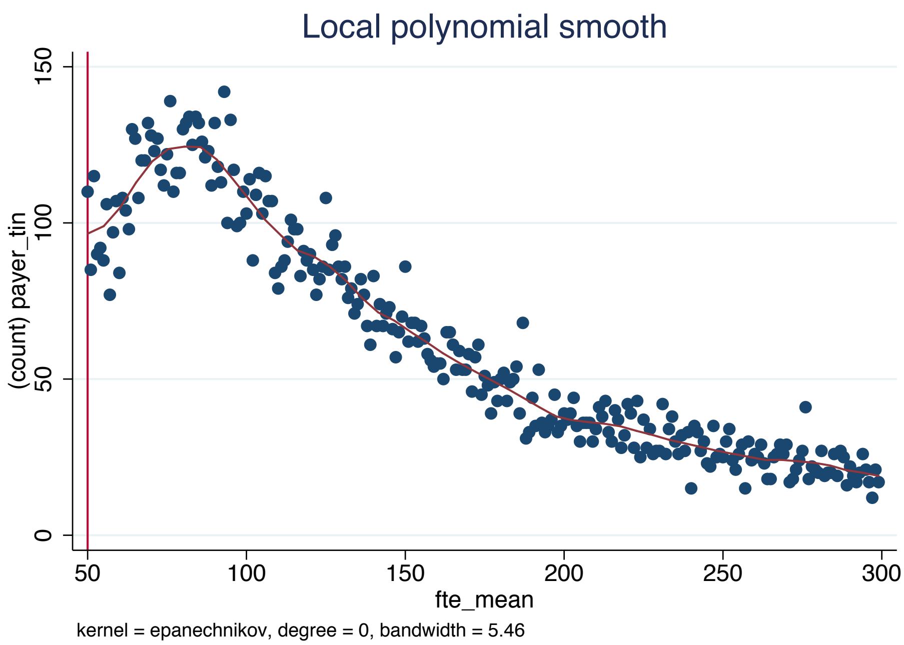

Figure 2.6b. Number of firms by FTE for firms with average wages above 75th percentile (2016, No Aggregate groups).

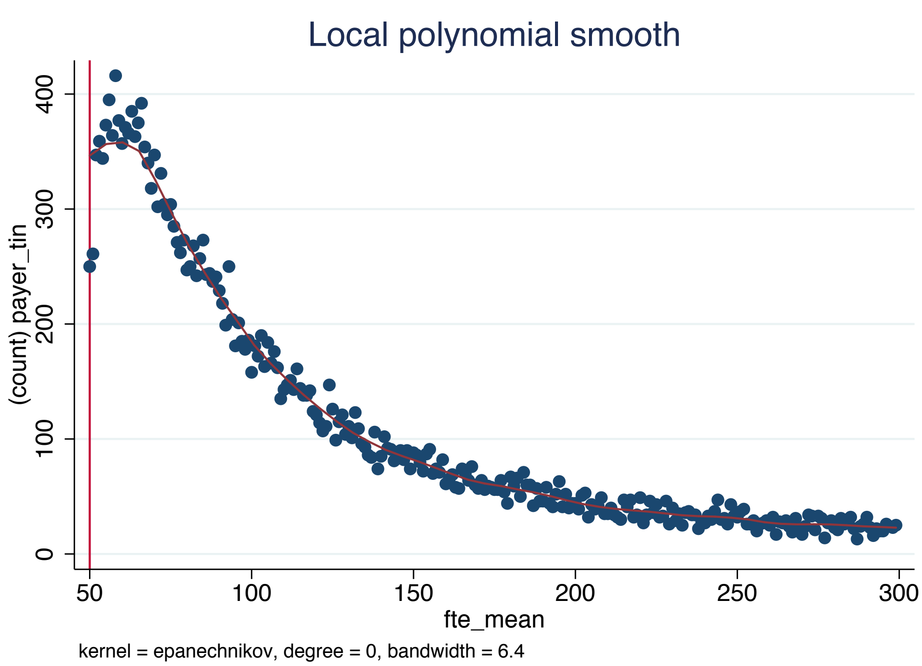

Figure 2.7a. FTE vs. Number of W-2s (2015)

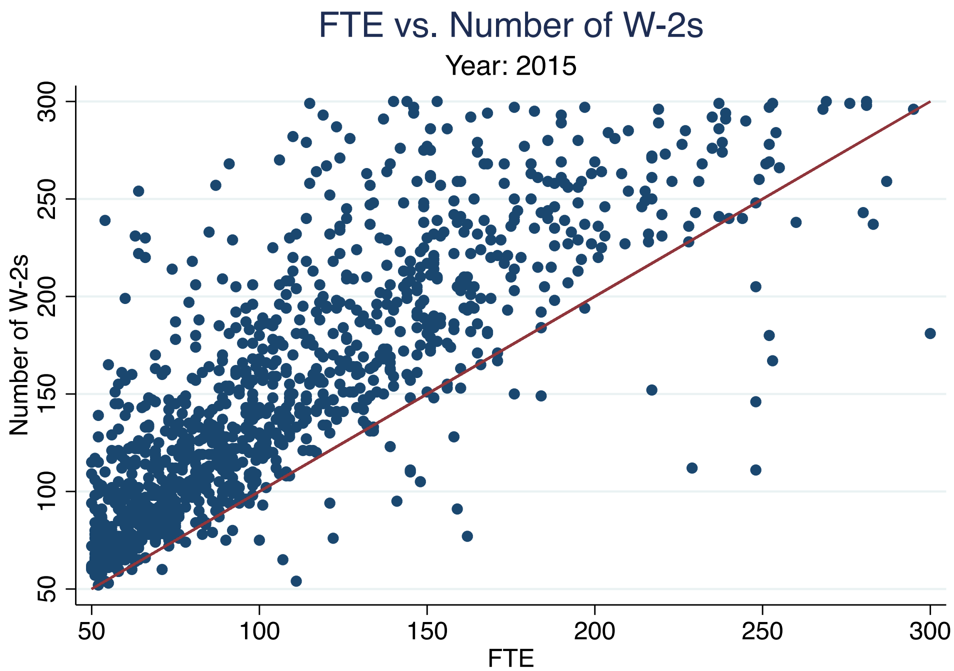

Figure 2.7b. FTE vs. Number of W-2s (2016)

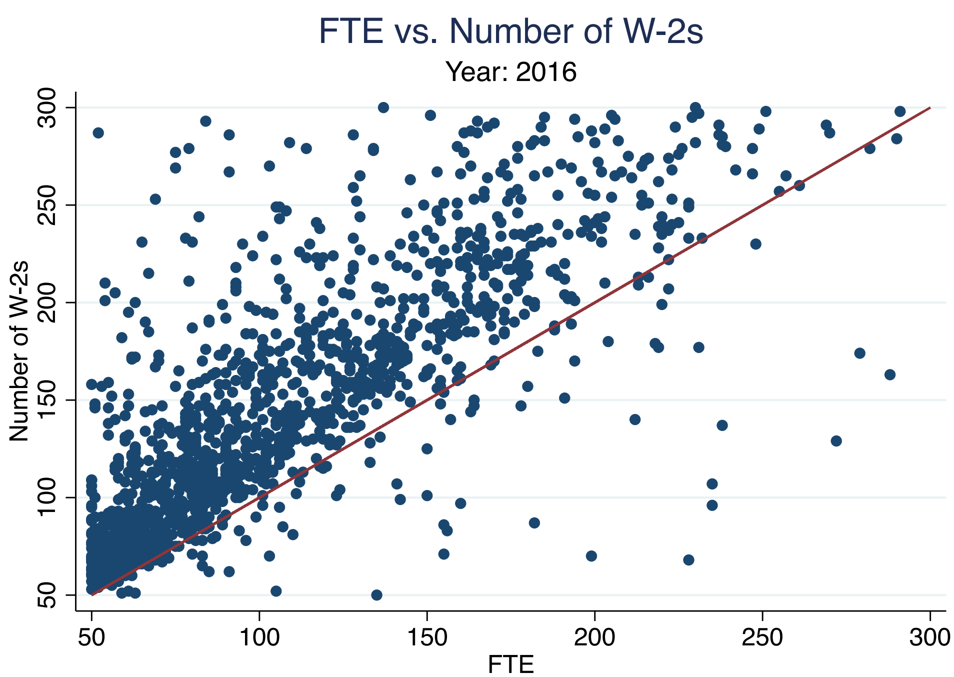

References

Abraham, Katherine, John C. Haltiwanger, Kristin Sandusky, and James R. Spletzer. "Measuring the Gig Economy." *Available at: www. sole-jole. org/16375. pdf* (2016); URL http://www.sole-jole.org/16375.pdf.

Bickley, James (2011). "Tax Gap: Misclassification of Employees as Independent Contractors." Congressional Research Service (March 10).

Goldschmidt, Deborah, and Johannes F. Schmieder. "The rise of domestic outsourcing and the evolution of the German wage structure." *The Quarterly Journal of Economics* 132, no. 3 (2017): 1165-1217.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw (1999). "Evaluating the Effect of an Antidiscrimination Law Using a Regression-Discontinuity Design." NBER Working Paper 7131.

Harris, Seth D., and Alan B. Krueger. *A Proposal for Modernizing Labor Laws for Twenty-First-Century Work: The" Independent Worker"*. Hamilton Project, Brookings, 2015.

Internal Revenue Service (1987). "Strategic Initiative on Withholding Noncompliance Study." Department of the Treasury, Pub 1415-E (Rev. 10-93).

General Accounting Office Contingent workforce: Size, characteristics, earnings, and bene ts. Technical Report GAO-15-168R, GAO, April 2015. URL http://www.gao.gov/products/GAO-15-168R.

Jackson, Emilie, Adam Looney, and Shanthi Ramnath. "The rise of alternative work arrangements: Evidence and implications for tax filing and benefit coverage." *Office of Tax Analysis Working Paper* 114 (2017).

Katz, Lawrence F., and Alan B. Krueger. *The rise and nature of alternative work arrangements in the United States, 1995-2015*. No. w22667. National Bureau of Economic Research, 2016.

Kleven, Henrik Jacobsen. "Bunching." *Annual Review of Economics* 8 (2016): 435-464.

Knittel, Mathew, Susan Nelson, Jason DeBacker, John Kitchen, James Pearce, and Richard Prsinzano (2011). "Methodology to Identify Small Businesses and Their Owners." *Office of Tax Analysis, Department of Treasury.*

Mas, Alexandre, and Amanda Pallais. "Valuing alternative work arrangements." *American Economic Review* 107, no. 12 (2017): 3722-59.

Onji, Kazuki (2009). "The Response of Firms to Eligibility Thresholds: Evidence from the Japanese Value-Added Tax." *Journal of Public Economics*, 93(5-6), June, pp. 766-775.

Schivardi, F., and Torrini, R. (2004). "Threshold Effects and Firm Size: The Case of Firing Costs." CEP Discussion Papers.

U.S. Department of Labor, Office of the Assistant Secretary for Policy (2012). "Independent Contractors and the Fair Labor Standards Act." March 2 2012.

CHAPTER III

Does the Elasticity of the Sales Tax Base Depend on Enforcement? Evidence from U.S. states' Voluntary Collection Agreements

In addition to taxpayer preferences, elasticity of taxable income has been shown to depend on parameters of the tax system—including the costs and expected penalty of tax evasion, and the costs of tax avoidance. However, less is known about how consumption elasticities change in response to enforcement. The theory of statutory neutrality predicts that structuring the as a 'use tax'—where the consumer remits—or, as a sales tax—under which the retailer remits, should have no effect on the fundamental parameters of the tax system. We test this in the context of U.S. state restructuring of the remittance regime governing online sales shipped to state residents. Using detailed purchase data from the Nielsen Consumer Panel and monthly, zip-code level information on local sales tax rates, we find that consumers reduce their online expenditure in response to Voluntary Collection Agreements (VCA). However, we do not find evidence of a large change in elasticity of the tax base with respect to tax changes. We conclude that shifting the remittance duty to the party with fewer evasion opportunities, akin to an enforcement increase, could affect the responsiveness of the tax base to future tax rate changes but that the effect of the enforcement on online retailers is too small to measure.

3.1 Introduction

In standard economic models, demand for a taxed good is solely a function of utility function parameters and the good's tax inclusive price. Implicitly, this assumes that a change in the log of the tax-exclusive price and changes in the tax rate have identical effects on behavior. However, recent literature casts doubt on this equivalence; for example, if the tax is less salient than the tax exclusive price at the point of decision (Chetty 2009, Finkelstein 2004), or if the tax increase can be mitigated by avoidance or evasion behavior in a way that a price increase could not (Slemrod and Yitzhaki 1996).

In the same vein, Slemrod and Kopczuk (2002) argue that the elasticity of taxable income is not a structural parameter. Rather, response is conditional on various parameters of the tax system – in particular, enforcement. In this paper, we explore whether the behavioral response of consumers and producers to a consumption tax change similarly varies with evasion opportunity in the context of U.S. states' voluntary tax collection agreements, a structural change in remittance assignment which substantially increased compliance with state sales taxes for online purchases. If consumers become less price-elastic as a result of this enforcement measure, states can potentially raise more revenue while lowering excess burden.

Existing literature has established that consumers purchase online to avoid sales taxes. Goolsbee (2001) was the first to suggest this channel of evasion. Einav, Knoepfle, Levin and Sundaresan (2014) find evidence of this evasion in consumers' online shopping response to taxes on the Ebay marketplace. Goolsbee, Lovenheim and Slemrod (2010) use data on state-level smoking rates and internet penetration from 1980 to 2005 to show that the price elasticity of cigarette sales rose as ability to purchase cigarettes online increased. Baker and Keung (2017) show using the Nielsen data that the internet is used as a means of evading the sales tax on a broader set of consumption goods, not just those subject to high sales and excise taxes. Policies that increase online sales tax compliance are thus a natural setting to study the effect of enforcement on consumption elasticities. We exploit time and geographic variation in adoption of Voluntary Collection Agreements (VCAs) which dramatically increased online sales tax compliance by shifting the duty to remit from the buyer to the seller. One recent paper shows that the VCA agreements had a measurable impact on consumers' shopping behavior on Amazon. Baugh et al (2017) show that this increase in tax on online purchases was salient to consumers, and that they reduced their Amazon purchases by about 9%. However, to our knowledge, this is the only paper to have so far examined the consequences of VCA agreements on the elasticity of the sales tax base.

To understand the likely effects of VCAs on consumption elasticity, we first build a simple theoretical framework to predict what might happen to the elasticity of the effective tax base when tax-exclusive prices remain fixed and consumers choose to either purchase a commodity online or at a brick-and-mortar store.

Next, we test the underlying assumptions and predictions of our framework using a large panel of household purchases from the Nielsen Consumer Survey. The rich information in the Nielsen data, which includes unique product identifiers, allows us to observe the elasticity of consumers' purchases with respect to tax changes at both online and brick-and-mortar retailers. We show that consumers reduce their online taxable expenditure in response to the VCA, while maintaining consumption of tax exempt items. We use an event-study design to test whether monthly online expenditure of households in states that enacted a VCA between 2010 and 2014 decreases following the VCA adoption relative to states which do not have a VCA. In line with the findings of Baugh, Ben-David and Park (2016) for online expenditure on Amazon, we find that expenditure at large online retailers fell in response to the VCA.

Next, we decompose this reduction in total expenditure into a change in reported tax-exclusive prices of online goods and a change in quantity demanded by consumers. The decrease in tax-exclusive expenditure online comes from consumers who continue to purchase online, but switch to cheaper varieties and cheaper commodities; and from consumers who simply stop shopping online-an extensive margin response. Households switch from purchasing the same products online to brick-and-mortar stores. Since online retailers typically price their goods for sale anywhere in the United States and the VCAs are implemented by state, it is reasonable that producers do not change their tax-exclusive price in response to the VCA and that any effective tax increase is passed through to the consumer.

Finally, we test whether the price elasticity of purchases at brick-and-mortar stores decreases because of the effective tax increase on online purchases resulting from enforcement. While these results are still very preliminary, we find limited evidence that enforcement significantly reduced the elasticity of the tax base.

3.2 Context

In this section, we discuss why states collect use taxes on online sales, and variation in state strategy to collect these taxes.

Forty-five U.S. states levy sales taxes on goods purchased for consumption within their physical borders, and require sellers, usually retailers, in these transactions to assess and remit the tax. To mitigate the tax arbitrage incentive to purchase products in low tax jurisdictions, states with general sales taxes often levy parallel "use taxes" on goods consumed in their states by their residents, but purchased outside the state or online.

Use tax provisions require residents to declare and self-assess the value of goods purchased elsewhere that would have been subject to sales tax if purchased in-state, and then to remit the equivalent sales tax amount to the state tax authority.[1] In theory, this minimizes revenue loss and distortion by equalizing after-tax prices. However, in practice, very few residents remit use taxes from either purchases made online, or those made in other states. In 2012, the percent of income tax returns reporting use tax (i.e. reporting tax liability on online purchases) ranged from 0.2 percent in Mississippi to 10 percent in Maine.[2]

3.2.1 Collecting Use Tax on Online Sales

States may impose a sales or use tax on purchases made by their residents, even if the retailer is out of state.[3] However, the state cannot legally impel the retailer to remit said tax unless there is a constitutionally sufficient relationship (a "nexus") between the retailer and the state.[4]

As internet sales have grown in volume, states have utilized a variety of strategies to recoup uncollected use taxes without running afoul of the constitution's nexus provision.

---

[1] States differ in their procedure for remitting use taxes. Several states require residents to report and remit use taxes annually, frequently via state income tax return. However, Vermont requires residents to report and remit each month (cite the form they use). Additionally, most states allow residents to deduct any sales tax that was paid in the source state, i.e. if Michigan has x% sales tax, and Michigan resident purchases a taxable item from Wisconsin and consumes it in Michigan, and pays a y% sales tax on her purchase, she need only pay the difference in use taxes to Michigan.

[2] See report published by Maine's tax authority: <http://www.house.leg.state.mn.us/hrd/pubs/usetax.pdf>

[3] The nexus requirement arises from two provisions in the U.S. Constitution: the Due Process Clause[3] and the Interstate Commerce Clause.[3] In the seminal case on this issue, *Quill v. North Dakota* (1992), the Supreme court held that a nexus exists only if the online retailer has a physical presence in the state (such as a store, office, warehouse or employees) or, if the retailer has purposefully solicited the state's residents.

[4] In addition to remittance, a state cannot impose any kind of "tax duty" (such as, requiring the retailer to report sales information to the state tax authority. Cite CO case.

Broadly, state actions can be divided into two categories; legislation, which tried to expand the definition of nexus to (large) online retailers in a manner consistent with *Quill*, and voluntary collection agreements (VCAs), essentially contracts between a single retailer and the tax authority in which the retailer agreed to remit future sales tax in return for some benefit. Although collectively referred to in popular parlance as “Amazon *Laws*,” this term is a misnomer; in most cases, states signed VCAs with Amazon and other large retailers either before or in conjunction with legislation.

Legislation, pioneered by New York and referred to as “click-through nexus,” imposes a duty to remit sales taxes on any retailer with in-state affiliate or associate that directs residents to the retailer’s website.[5] This extended the duty to remit to large retailers such as Amazon or BackCountry, unless they dropped all affiliated sellers in the state that sold through their platform. In several states, Amazon initially dropped affiliates to avoid nexus (CO, NC, TN), but in large states with hundreds of affiliates, Amazon acknowledged nexus and began remitting. In our study period, three states (CA, NJ, PA and VA) passed such legislation.

In contrast, fourteen states announced VCAs with Amazon during our study period. In general terms, a VCA is a non-standard contract between a business and a state or local tax authority in which the business “voluntarily” agrees to assess and remit taxes going forward, even if not legally required to do so. In the context of online sales, large retailers signed these agreements in exchange for concession by the state, such as release from back taxes, or a commitment by the state not to require the retailer to disclose individual buyer data. For example, in July 2012, Amazon signed a VCA with the state of Texas promising to remit future

[5] The language of the 2008 New York statute creates a rebuttable presumption of nexus “if the seller enters into an agreement with a resident of this state under which the resident, for a commission or other consideration, directly or indirectly refers potential customers, whether by a link on the Internet website or otherwise, to the seller.” N.Y. Tax Law

taxes and to increase capital investment in the state. In exchange, the Texas State Comptroller agreed not to pursue collection of the estimated $269 million in sales tax that Amazon had not collected between 2005-2009.

Our design relies on variation in state sales tax rate, variation in VCA adoption (See Figure 3.1), and variation in the tax base to which the VCA applies (i.e. exemptions). Several states have also enacted temporary exemptions "sales tax holidays" for specific product categories (e.g. school supplies), which we can potentially exploit for further variation. Consumers differ in propensity to purchase online; Figure 3.4 shows that the ratio of total expenditures online to total expenditures is rising in household consumption.

Sales taxes in the United States are set by states and local option sales taxes at the county or city level supplement these standard rates. Sales tax exemptions can vary by state. In addition, some goods are taxed at a special discounted or higher rate. Some goods like alcohol and tobacco are also subject to additional excise taxes. We focus on goods taxed at the standard sales tax rate and exempt goods only for now, excluding items taxed at a special rate.

3.3 Data

The Nielsen Consumer Panel is a nationally and regionally representative, stratified longitudinal panel of between 40,000 – 60,000 households from 2004 – 2014.[6] For this draft, we focus on the sample of households observed between 2010 and 2014, which is the period when most VCT agreements were signed, to keep the dataset of a manageable size. Households self-report their purchasing behavior to Nielsen through in-home scanners for a set of "Nielsen-tracked" products. These products include both food and non-food items purchased at any outlet,

[6] The sample was increased from 40,000 to 60,000 in 2007.

including purchases made online. Households record their purchases from each shopping trip, which includes information on total amount spent, retailer type, payment type, value of each item purchased and quantity of each item purchased. Items are identified by a unique product code (UPC) with details on brand variation, size, multipacks etc. This detailed product and quantity information allows us to more accurately measure the impact of the VCA on consumer purchase behavior. Unlike Baugh et al., we are able to separately analyze the response of taxable and exempt consumption. We also decompose the expenditure response into the price and quantity demanded response to the VCA.

The Nielsen-tracked product groups capture approximately 30 percent of total household consumption. Our estimates of consumption elasticity with respect to the tax rate therefore only reflects consumption elasticity of this subset of household consumption rather than total household consumption. Notably, Nielsen emphasizes fast-moving consumer goods over durables like washing machines or cars. Therefore, our price elasticity estimates are likely to be smaller since durables consumption is generally more elastic.

For tax rate changes, we use data on monthly sales tax rates at the state, county, and local (school district, etc.) level purchased from zip2tax. Table 3.1 shows the number of sales tax rate changes in our data at each administrative level. Most changes over this time period occur at the city level (2089). Figure 3.8 shows that the distribution of tax changes before and after the VCA are not very different.

We construct a measure of total tax exclusive expenditure at each household. Each shopping trip a household makes is assigned a retailer code and each retailer is assigned a "channel type". One of the channel categories is "Online Retailer", which allows us to distinguish online shopping trips[7]. We construct a measure of total monthly total online expenditure for each household by adding the reported item-level expenditure, which are exclusive of tax. Similarly, we measure total online taxable expenditure and exempt expenditure separately by adding up item-level expenditure of items within each category.

Our predictions about the effect of the VCA on elasticity of the tax base assumes that the effective tax increase due to the VCA is fully passed through to the consumer. We test the pass-through of the VCA to the consumer directly as the effect on tax-exclusive price at the UPC-level. We create unit-level price of each purchase as the total price after any coupons divided by the quantity recorded.

Next, we turn to whether the reduction in online expenditure as a result of the VCA agreements, also translates to lower sensitivity of the effective tax base to sales tax changes. Assuming that use tax compliance prior to the VCA is zero and 100 percent afterward, we define the "effective tax base" as brick-and-mortar expenditure prior to the VCA and the sum of brick-and-mortar expenditure and online expenditure after the VCA. This definition is intended to capture the expenditure that is likely reported to the tax authority.

3.4 Illustrative Model

In this section, we present an illustrative model of how VCA adoption affects online and offline consumption elasticities and the elasticity of the effective sales tax base. We first present an identity for tax elasticity of demand for taxable goods, and then move to the consumer's problem. We conclude with three predictions that we can take to the data.

[7] Although the identity of individual retailers is unknown, we can identify "large online retailers" through the volume, diversity and ubiquity of sales recorded on Nielsen. One retailer code is a generic "Other" category but we believe we can identify this retailer code.

3.4.1 *Tax Elasticity of Demand for Taxable Goods*

Total demand for taxable goods, D_x is the sum of demand for taxable online goods, D_x^o and taxable offline goods, D_x^b. So, the tax-elasticity of demand for taxable goods is:

$$\epsilon_{x,t} = \frac{\partial D_x}{\partial t}\frac{t}{D_x} = \left(\frac{\partial D_x^o}{\partial t} + \frac{\partial D_x^b}{\partial t}\right)\frac{t}{D_x} = \frac{\epsilon_{x^o,t} D_x^o + \epsilon_{x^b,t} D_x^b}{D_x} \tag{1}$$

$$= \epsilon_{x^o,t}\theta + \epsilon_{x^b,t}(1-\theta)$$

Where $\theta = \frac{D_x^o}{D_x}$, online demand for the product as a share of total demand. The smaller that θ is, the closer tax elasticity of total demand is to tax elasticity of demand for offline products.

θ is also directly affected by the tax rate, whether the VCA is in place, the relative price of the good online and offline, as well as consumers' relative preference for online and offline purchasing. We present a simple model below that illustrates how θ, $\epsilon_{x^o,t}$, and $\epsilon_{x^b,t}$ might change as a result of the VCA.

3.4.2 *Consumer's Problem*

Consumers vary in their preference for online versus brick-and-mortar purchasing, which is represented by the parameter $\psi_i \sim F(.)$ and affects their relative value of online versus offline purchases of identical goods. For example, a consumer that prefers to try on clothes before purchase might prefer an identical shirt in a brick and mortar store over its online equivalent: she would have a negative value of ψ_i. Conversely, a consumer who values the convenient delivery of online sales, perhaps because he does not own a car, would have a positive value of ψ_i. The consumer solves:

$$u(x^o, x^b, e^o, e^b) = (\psi_i x^o + x^b)^\alpha (\psi_i e^o + e^b)^{1-\alpha} \tag{2}$$

Such that,

$$q_x^o x^o + q_x^b x^b + q_e^o e^o + q_e^b e^b = I \tag{3}$$

where x^o and x^b represent composite taxable online and brick-and-mortar goods, respectively and e^o and e^b represent composite tax-exempt online and brick-and-mortar goods, q_j^i is the after-tax unit price of good j. This utility function reflects two key assumptions: (1) that offline and online versions of the taxable and exempt goods are perfect substitutes and (2) that consumers spend a constant share of their income on the taxable and exempt goods. The first should hold generally – any individual consumer probably does not purchase the same good both online and offline. The second is assumed for now for simplicity but may be relaxed later.

Each unit of the online good is ψ_i times as valuable as the offline good. $\psi_i > 1$ would imply that the price of x^b must be less than $\frac{1}{\psi_i}$ times the price of x^o for the consumer to prefer x^b. So, in the first step consumers choose between the online and offline version of the goods.

A consumer chooses $X = x^o$ if $\frac{q_x^o}{\psi_i} < q_x^b$. They choose $E = e^o$ if $\frac{q_e^o}{\psi_i} < q_e^b$.

The probability that a consumer chooses x^o is therefore:

$$P\left(\frac{q_x^o}{\psi_i} < q_x^b\right) = P\left(\frac{q_x^o}{q_x^b} < \psi_i\right) = 1 - F\left(\frac{q_x^o}{q_x^b}\right) \tag{4}$$

The consumer's problem then reduces to:

$$\max_{X,E} X^{\alpha} E^{1-\alpha} \tag{5}$$

$$s.t.\ C_x X + C_E E = I$$

where

$$X = I\left(\frac{q_x^o}{q_x^b} < \psi_i\right)\psi_i x^o + \left[1 - I\left(\frac{q_x^o}{q_x^b} < \psi_i\right)\right]x^b \tag{6}$$

$$E = I\left(\frac{q_e^o}{q_e^b} < \psi_i\right)\psi_i e^o + \left[1 - I\left(\frac{q_e^o}{q_e^b} < \psi_i\right)\right]e^b \tag{7}$$

$$C_x = I\left(\frac{q_x^o}{q_x^b} < \psi_i\right)q_x^o + \left[1 - I\left(\frac{q_x^o}{q_x^b} < \psi_i\right)\right]q_x^b \tag{8}$$

$$C_E = I\left(\frac{q_e^o}{q_e^b} < \psi_i\right)q_e^o + \left[1 - I\left(\frac{q_e^o}{q_e^b} < \psi_i\right)\right]q_e^b \tag{9}$$

For the tax-exempt goods, $q_e^i = p_e^i$ where p_e^i is the pre-tax price. Prior to the VCA, online sales were effectively treated as exempt, i.e. $q_x^o = p_x^o$; after the VCA, they were subject to sales tax $q_x^o = p_x^o(1 + t)$. For taxable goods in brick and mortar stores, $q_x^b = p_x^b(1 + t)$ both before and after the VCA. We also assume that the tax-exclusive prices are fixed and exogenously given, i.e. perfectly elastic supply curves, an assumption we will justify in the next section.

The Cobb-Douglas utility gives rise to the following individual demand functions: $X(C_x, I) = \frac{\alpha I}{C_x}$ and $E(C_E, I) = \frac{(1-\alpha)I}{C_E}$. Aggregate demand for taxable online goods is therefore[8]:

$$D_x^o = \left[1 - F\left(\frac{q_x^o}{q_x^b}\right)\right]\left[\frac{\alpha I}{q_x^o}\right] \tag{10}$$

Similarly, aggregate demand for taxable brick and mortar goods is:

$$D_x^b = F\left(\frac{q_x^o}{q_x^b}\right)\left[\frac{\alpha I}{q_x^b}\right] \tag{11}$$

3.4.3 Predictions

Comparative statics yield three testable predictions relevant to the effect of the policy on consumption elasticities:

Proposition 1: If the VCA substantially increased sales tax compliance for online purchases, the tax elasticity of online taxable goods changes sign, becoming negative.

Prior to VCA adoption, the tax elasticity of demand for taxable online goods is:

---

[8] Because,

$$\begin{aligned} D_x^o &= \int_{\Psi_i} P\left(\frac{q_x^o}{\psi_i} < q_x^b\right)[X(C_x, I)]\mathrm{dF}(\Psi_\mathrm{i}) \\ &= \int_{\Psi_i}\left[1 - F\left(\frac{q_x^o}{q_x^b}\right)\right]\left[\frac{\alpha I}{q_x^o}\right]\mathrm{dF}(\Psi_\mathrm{i}) \\ &= \left[1 - F\left(\frac{q_x^o}{q_x^b}\right)\right]\left[\frac{\alpha I}{q_x^o}\right] \end{aligned}$$

$$\epsilon^{pre}_{x^o,t} = \frac{\partial D^o_x}{\partial t}\frac{t}{D^o_x} = \frac{t}{1+t}\left(\frac{p^o_x}{q^b_x}\right)\frac{f\left(p^o_x/q^b_x\right)}{\left[1 - F\left(p^o_x/q^b_x\right)\right]} > 0 \tag{12}$$

This is intuitive. As the tax rate increases, making taxed brick and mortar goods comparatively more expensive, consumers with marginally lower values of ψ_i who previously purchased a product in a brick and mortar store will switch to purchasing online.

After VCA adoption, $q^o_x = p^o_x(1+t)$ and $\frac{\partial\left(q^o_x/q^b_x\right)}{\partial t} = 0$.

Therefore,

$$\epsilon^{post}_{x^o,t} = -\frac{t}{1+t} < 0 \tag{13}$$

After the VCA, a change in the tax rate does not change the relative price of online and offline products, and so it does not affect the share of online demand. The demand elasticity for online goods is the same as that of offline goods.

Proposition 2: Tax elasticity of brick and mortar taxable goods remains negative but becomes smaller in magnitude.

After the VCA,

$$\epsilon^{post}_{x^b,t} = \epsilon^{post}_{x^o,t} = -\frac{t}{1+t} > \epsilon^{pre}_{x^b,t} \tag{14}$$

This reflects the elasticity of demand coming from the increase in after tax price of x^b as well as the decrease in demand from consumers shifting from offline to online consumption. This is in accord with standard models: $\epsilon_{x^b,t}$ is generally negative prior to the VCA since an increase in the local sales tax rate would induce individuals to switch to purchasing online, or to demand less offline.

After the implementation of the VCA, $\epsilon_{x^o,t}$, should become negative since an increase in the tax rate would also increase the relative after tax price of online goods. $\epsilon_{x^b,t}$ will become smaller in magnitude as individuals will no longer switch from purchasing offline to online. How these changes in demand elasticity for online and offline products affects overall elasticity will depend on the relative importance of the online and offline demand for the product as well as the magnitude of the change in elasticity.

Proposition 3: The elasticity of the effective tax base, defined as the value of goods on which tax is remitted, becomes smaller in magnitude.

We define the effective tax base as purchases reported to the tax authority (and on which tax is remitted). Prior to the VCA, the base is simply the offline purchases as almost no online purchase is reported. After the VCA we assume full compliance on both online and offline purchases. Therefore, the effective tax base is now both online and offline purchases.
Post VCA the tax-elasticity of online and offline tax base is identical and the elasticity of the effective tax base is:

$$\epsilon_{x\ \ ,t}^{post} = -\frac{t}{1+t} \tag{15}$$

Prior to the VCA, the base is equal to offline expenditure and therefore the elasticity is:

$$\epsilon_{x\quad,t}^{pre} = \epsilon_{x^b,t}^{pre} = -\frac{t}{1+t}\left[\frac{f\left({p_x^o}\big/{q_x^b}\right)}{F\left({p_x^o}\big/{q_x^b}\right)}\left(\frac{p_x^o}{q_x^b}\right)+1\right] > \epsilon_{x\quad,t}^{post} \tag{16}$$

3.5 VCA Effect on Online Prices and Consumption

In this section, we establish several preliminary empirical facts implicitly assumed by our model. We find that the VCA had an effect on online purchasing by households, and that these effects are consistent with an after tax price increase in online goods. First, we evidence that VCAs substantially increased the number of online purchases on which sales taxes were collected by online retailers. Next, using two measures of consumer behavior, we show that consumers reacted to this change in remittance policy akin to a tax increase, suggesting that use tax compliance was low prior to VCAs.

3.5.1 More Online Foods Are Taxed at Point of Sale After the VCA

A prima facie question is “Did the Amazon Laws actually induce online retailers to collect and remit sales taxes?” Given the difficulty and expense state authorities face in enforcing remittance obligations against out of state retailers, we do not assume that VCA were efficacious. Instead, we establish that retailers began collecting tax on online purchases from the data. Nielsen records expenditure in two variables – *item-level* expenditure and *trip-level* expenditure. The *trip-level* expenditure is always tax-inclusive while the *item-level* is tax exclusive.[9] If no sales tax is collected at the point of transaction, the aggregate of all item expenditures for a given trip will equal the trip-level expenditure. If the VCA induced retailer remittance, we expect the fraction of online transactions where no sales tax was collected to fall.

[9] We investigate this crucial aspect of the data in detail. See the data appendix.

We visually test whether this is indeed the case. After restricting the data to trips in which only taxable items were purchased, we separately plot the share of trips where the sum of the *item-level* expenditure equals the *trip-level* expenditure for online and offline purchases, relative to the VCA adoption (see Figure 3.2). Prior to the VCA agreements, about 25 percent of online trips have no tax collected, whereas only about 12 percent of offline trips have no tax collected (or report *item-level* tax-inclusive expenditures). We see a sharp drop in this fraction for online trips, suggesting that online retailers began collecting sales taxes soon after implementation of the VCA.

Having established that online retailers remitted after the policy, we now turn to the consumer response. Classical tax theory, which assumes full salience and compliance, would predict that shifting the remittance duty from the consumer to the retailer should have no effect on equilibrium quantities and prices. However, if, as we suspect, compliance with use taxes was low, for most consumers the policy increased the tax inclusive price of online goods.[10]

3.5.2 Consumers Reduced Total Online Spending on Taxed Goods, Though Not Tax-Exempt Goods

We estimate the effect of the policy on online purchasing behavior by estimating the following difference in difference specification:

$$y_{hm} = \beta_0 + \beta_1 T_h * Post_{hm} + \beta_2 X_{hm} + \gamma_m + \delta_h + \varepsilon_{hm} \tag{17}$$

where y_{hm} is either (1) total online taxable expenditure or (2) total online exempt expenditure of household *h* in month *m*. β_1 is the parameter of interest where T_h indicates where household *h* is

[10] The exact amount that after tax prices increase depends on relative demand and supply elasticities, but, as most Nielsen tracked products are commodities, we think 0% pass through is unlikely. Also, NB, I haven't actually met anyone in the flesh who remits use taxes.

in a state that adopts the VCA between 2010 and 2014 and $Post_{hm}$ is an indicator for whether we are observing household h in a month m following adoption of VCA in that state. We also control for time fixed effects (γ_m) and household fixed effects (δ_h), as well as time-varying area-level characteristics (X_{hm}) such as a local cost of living index[11].

If the parallel trends assumption holds—that is, if the online purchasing habits of households in states that did not adopt VCAs are a suitable counterfactual for the purchasing habits of households in states that adopted VCAs—then this parameter represents the difference-in-differences estimator of the effect of VCA adoption on the extensive and intensive margin of online sales. We would expect that online expenditure on taxable items falls as a result of the VCA but that online expenditure on exempt items does not change.

We find that the introduction of the VCA reduced total monthly tax-exclusive expenditure online by about 20 cents on average, which represents an 8 percent decrease relative to the mean (Table 3.2, column 1). In contrast, we see no statistically significant effect of the VCA on online expenditure on exempt goods and the estimated magnitude is close to zero. These estimates control for household, year and month fixed effects and standard errors are clustered by state. Total monthly expenditure at brick and mortar stores increase by about one dollar but these effects are not statistically significant. Note however, that monthly expenditure at brick-and-mortar stores also include expenditure on goods that are never purchased online even prior to the VCAs. It is therefore possible that expenditure on goods that were previously purchased online increases while there is no change in other expenditure. Figure 3.3 shows, there is no anticipatory effect of the VCA in the quarter before and the parallel trends assumption holds.

[11] We create this measure following steps outlined in Baugh et al. (2017)

Figure 3.3 shows that online tax-exclusive expenditure on taxable goods falls in months following the VCA. In months before the VCA we do not see any anticipatory effects on online expenditure. Prior to the VCA, although we see fluctuation in expenditure from month to month, on average the difference in expenditure in these months relative to month just prior to the VCA is 0. After the VCA, we see that on average the expenditure is about 20 to 30 cents lower each month.

We also separately estimate the effect of the VCA on large versus small retailers. Nielsen lists a unique retailer code that identifies where each purchase was made. We define "large" retailer as the two retailer codes that together represent about 50 percent of all online purchases. We find that the expenditure decline comes largely from declines at these large online retailers. Expenditure at these retailers declines by nearly 20 percent (Table 3.2, column 2). Expenditure at small retailers on the other hand shows a small but statistically insignificant increase (Table 3.2, column 3). Similarly, we find small and statistically insignificant increases in taxable and exempt expenditure at brick-and-mortar stores (Table 3.2, columns 4 and 5).

In Figure 3.6 we see a small and possibly delayed effect on monthly online tax-exclusive expenditure of households on exempt goods. However, on average this is a statistically significant expenditure change. Again, there does not seem to be evidence of anticipatory effects or differential trends in online expenditure between households in states that do and do not adopt the VCA.

3.5.3 Online Retailers Do Not Adjust Tax-Exclusive Prices; Consumers Reduce Quantity Purchased

We decompose the change in tax-exclusive online expenditure into the change in the tax-exclusive price of goods and change in consumer demand. Our specification estimating the effect on tax-exclusive prices:

$$\log(p_{cmu}) = \beta_0 + \beta_1 T_c * Post_{cm} + \gamma_m + \upsilon_y + \delta_u + \alpha_c + \varepsilon_{smu} \tag{18}$$

where the coefficient of interest is again β_1, which represents the average percent change in the tax-exclusive price across all products due to the VCA.

Next, we test the effect on consumer demand (quantity purchased) within UPC using the following specification:

$$\log(q_{cmu}) = \beta_0 + \beta_1 T_c * Post_{cm} + \gamma_m + \upsilon_y + \delta_u + \alpha_c + \varepsilon_{smu} \tag{19}$$

where β_1 is the estimate of average percent change in quantity demanded for product, conditional on purchase (i.e. intensive margin effect on quantity). The drawback of this specification is that a null effect could be consistent with a couple of different interpretations: (1) Consumers do not reduce their quantity demanded on most goods, conditional on online purchase, as a result of the VCA, (2) Consumers reduce their quantity demanded of higher price goods and substitute to purchasing lower price goods (therefore increasing quantity demanded of these goods). On average, this would translate to no effect on quantity demanded. For example, if consumers switch from a higher priced variety of household cleaner to a lower priced variety, this would appear on average as no change in quantity demanded across UPC. Or, if consumers decide not to purchase an expensive kitchen appliance and instead spend more of their budget on other lower priced items – they would have decreased quantity demanded in one UPC but increased demand for another. (3) Consumers only respond on the intensive margin, i.e. they stop purchasing any amount of the product online.

To distinguish between (1) and (2), we examine the effect of the VCA on quantity interacted with the average price of each UPC across purchases from all states in 2011, a year in which no state introduced a VCA. This price is by definition, unaffected by the VCA. In this way, we can examine heterogeneous effects on demand due to the VCA across high and low-price commodities.

$$\log(q_{cmu}) = \beta_0 + \beta_1 T_c * Post_{cm} * PreVCAp_{cmu} + \beta_2 T_c * PreVCAp_{cmu} + \beta_3 T_c * Post_{cm} + \gamma_m + \upsilon_y + \delta_u + \alpha_c + \varepsilon_{smu} \qquad (20)$$

Now β_1 measures the average decrease in consumer demand across UPC, scaled by the price of each UPC. If consumers behave as described in (2), we would expect β_1 to be negative. On the other hand, if consumers behave as described in (1), we would expect β_1 to be zero.

Effect of VCA on tax-exclusive prices. Table 3.3 decomposes the effect on total online expenditure into the effect on prices and quantity separately. We do this analysis at the purchase level including fixed effects for each UPC, time and household/county. Columns 3 – 6 shows the effect on log of prices. The coefficient of interest should be interpreted as the percent change in prices due to the VCA. We find that the VCA reduced prices by 0.9 percent, but this reduction is coming mostly from purchases of video products. We find no evidence of a statistically significant change in the tax-exclusive price of most goods purchased online, suggesting that any effective tax increase due to the VCA was fully passed through to consumers. Therefore, the reduction in expenditure is coming from consumers reducing quantity demanded of goods online.

Quantity Purchased Online – Intensive Margin. Columns 1-3 in Table 3.3 show the effect of the VCA on the intensive margin of purchases. That is, conditional on observing a purchase of a particular UPC, how does quantity purchased of that UPC change as a result of the VCA? We find no evidence of an intensive margin effect on quantity on average. That is, conditional on an online purchase, we do not see a decrease in quantity on average across all commodities. However, this result could be consistent with a decrease in quantity purchased of some goods and an increase in quantity purchased of others. For example, if consumers substituted away from a more expensive to a less expensive variety, we would not find evidence of a decrease in quantity on average.

One way to test whether this happens would be to interact quantity effect with average pre-VCA price of each UPC. These results are presented in Table 3.4. We calculate the average price in 2011 for each UPC, a year in which there were no VCA adoptions, and interact the treatment effect with this price. Column 1 shows that quantity demanded decreases as a result of the VCA by more for higher price taxable goods, suggesting that consumers substitute away from higher price varieties to lower price varieties or lower price goods. A $1 increase in the average price of a UPC in 2011 translates to a 0.2 percentage point greater decrease in the quantity demanded of that UPC. This effect is robust to the inclusion of household fixed effects. For exempt goods, quantity demanded increases on average following the VCA but less so for higher price commodities. Overall, the VCA does not decrease quantity demanded of exempt goods on the intensive margin.

3.6 VCA Effect on Tax Elasticities

While the first empirical section evidenced the direct effects of the VCA on online purchasing, this second section will examine the effect of VCA adoption on fundamental parameters of the sales tax system: the elasticity of the tax base. We estimate the effect of the VCA on the "effective tax base", which we define as the expenditure that is reported to the tax authority. Prior to the VCA, this base is only expenditure at brick-and-mortar retailers. We assume that no online expenditure is reported, which we feel is reasonable given the near zero compliance rate on use taxes. After the VCA, the base is the sum of both online and offline expenditure.

We also test the effect of the VCA on a subset of this base – expenditure on UPC that are purchased online between 5 and 95 percent of the time. Since the change in elasticity is expected to come from consumers who no longer purchase the good online in response to a tax change, we would expect that the effect is strongest on goods that can be purchased both online and at brick-and-mortar retailers.

3.6.1 VCA Effect on the Elasticity of the Tax Base with Respect to the Tax Rate

The revenue consequence of an increase in the statutory tax rate are often divided into two countervailing effects: the "arithmetic" effect of a higher rate which increases revenue, and the "economic effect" of reducing the base by dis-incentivizing the taxed economic activity (See, e.g. Laffer 2004 for discussion). We are concerned with the later. In our context, a higher sales tax rate increases the cost of taxed goods, and can reduce the sales tax base through multiple channels. The sales tax base can shrink in if taxpayers substitute to tax exempt products (substitution effect) or reduce consumption of all products (income effect), both of which are determined by consumers' utility functions. In addition, the base might shrink if consumers respond to a higher tax rate by putting greater effort into sales tax avoidance via cross border shopping, black market purchases, or in the context at hand, by ordering items online. Unlike the substitution and income effects, the "avoidance effect" of increasing the tax rate on the tax base is determined by other features of the tax system, such as the enforcement regime, that determine the effort/costs the taxpayer must expend/incur to avoid their jurisdiction's tax rate. (See, e.g. Slemrod 2008).

Let the sales tax base in jurisdiction of household *h* at time *t* be defined as

$$\mathrm{B}_{ht} = \sum_{i}^{I} p_{iht}(\tau) X_{iht}(q_{iht}(p_{iht}(\tau), \tau) \tag{21}$$

where *I* is the set of all taxable goods in the jurisdiction of household h, and τ is the sales tax rate. The first term, $p_{iht}(\tau)$, denotes the tax exclusive price; X_{ict} is the aggregate demand for product *i* in jurisdiction *c* at time *t*, and is a function of the tax inclusive price $q_{iht} = p_{iht}(\tau)(1 + \tau)$. The effect of the VCA on tax rate elasticity of demand can be expressed as $\frac{d\mathrm{B}_{ht}}{dVCA\, d\tau} = \frac{d\mathrm{B}_{ht}}{d\tau}\Big|_{VCA=1} - \frac{d\mathrm{B}_{ht}}{d\tau}\Big|_{VCA=0}$; the difference in the derivative of tax base with respect to the tax rate when a VCA is in place.

We estimate this effect with the following OLS specification at the household-month level:

$$\Delta \log(e_{hcst}) = \beta_0 + \beta_1 \Delta\tau_{ct} + \beta_2 treat_s post_{st} + \beta_3 \Delta\tau_{ct} treat_s post_{st} + \pi X_{ct} + \gamma_h + \gamma_t + \gamma_s * t + \varepsilon_{hcst} \tag{22}$$

β_1 captures the relationship between the tax base and the sales tax rate in untreated states. β_2 represents the tax rate invariant effect of VCA adoption on expenditures. The coefficient of interest, β_3 , captures how the effect of VCA adoption varies with changes in the sales tax rate. X_{ct} is a vector of time-varying county-level controls, including the unemployment rate. Household and time fixed effects are included to control for any time-invariant household characteristics and time trends, respectively.

In Table 3.5, we estimate the above equation over three tax bases: first, the effective tax base, second the subset of the effective tax base that is purchased both online and offline, and finally, the brick-and-mortar tax base.

Specifying at the household-month level has two advantages: we can include household effects which absorb idiosyncratic variation in expenditures within a household, thus making our estimates considerably more precise; and it further mitigates omitted variable concerns by partially controlling for endogenous sorting of households into local tax jurisdictions. After transformation, the coefficient estimate for β_3 in Col. 3 suggests that households' taxed expenditures became somewhat less elastic to a tax rate change but that this change is not statistically significant.

In column 1, we estimate the elasticity of the tax base with respect to the tax rate for all goods that are subject to the standard sales tax rate over all time periods. This base excludes goods like soda, alcohol or cigarettes and other goods that may be taxed at special rates. It also excludes goods that are tax exempt. We focus on only positive tax rate changes. We find a very large estimated elasticity of the tax base of -3.8, which is even larger when we restrict the sample to just prior to the VCA (and to states that adopt the VCA) at -4 (column 2). The coefficient on β_3 in column 4 is positive, as we expect, suggesting that the VCA reduced the elasticity of the effective base.

In column 5, we estimate the same specification on the subset of goods that are purchased both online and offline. Surprisingly, we do not find the same effect of the VCA on this subset of goods. Columns 6 – 9 restrict the analysis to only brick-and-mortar purchases. The estimated effects are largely similar to what we see in the effects on the effective tax base. This is what we would expect since brick-and-mortar purchases form 99 percent of the effective base.

Although we control for hyper local market conditions, we recognize that our estimates may still be vulnerable to omitted variable bias. Future analysis will explore potential instruments for locality I's tax rate, for example, the lagged tax rate of similar counties or proximate counties, or fixing price to some period before the treatment window (See Case, Rosen, Hines (1993); spatial correlation in tax rate paper estimated on UK data).

3.7 Conclusion

With the share of consumer purchases made online expected to grow, policymakers are understandably focused on ways of ensuring that online retailers remit sales taxes. In this paper, we study the impact of states adopting VCAs with Amazon, the largest online retailer, on the prices and purchases of online goods. We are also interested in the effect of VCA adoption, which makes it more difficult for consumers to purchase products online from non-remitting retailers, on the sales tax elasticity. To investigate these questions, we exploit variation in the location and timing of VCA adoption by states between 2010 and 2014, and we use data from the Nielsen Consumer Panel.

First, we find that VCA adoption increases the share of online goods sold that are taxed at the point of sale. To establish this, we measure the percentage of taxable sales where the after-tax item price is equal to the pre-tax item price, implying that sales taxes were not being remitted by the online retailer. The proportion of online sales meeting this criterion falls by nearly half in response to the VCA, with the most pronounced changes at the largest retailers who are likeliest to comply; the analogous proportion for brick-and-mortar sales remains constant.

Second, we find that consumers respond to VCA adoption by reducing their online consumption. On average, households in VCA-adopting states reduce online purchases by 8%, similar to the findings of Baugh et al. that households reduce purchases on Amazon following VCA adoption by 9-12%. The discrepancy in these estimates is likely caused by the fact that, in our data, we capture all online expenditures, rather only those for Amazon, and many small retailers did not sign VCAs. This response suggests a sales tax elasticity of between -1.2 and -1.4, smaller but still in the range of elasticity estimates reported by Baker and Keung in their study of consumer response to local sales tax rate changes. Both of these findings call into question the view that sales taxes are not salient to buyers at the point of purchase.

Finally, we attempt to measure the impact of VCA adoption on the sales tax elasticity. Unfortunately, only two states in our sample change their tax rates after adopting a VCA, providing insufficient variation to reliably measure this estimate. In future work, we plan to use state sales tax holidays as an alternative source of tax rate variation to measure this impact. For example, if VCA adoption meaningfully limits consumers' ability to avoid paying sales taxes throughout the year, then we would expect consumer response to sales tax holidays to increase.

Figure 3.1. Date of Implementation of Amazon VCAs

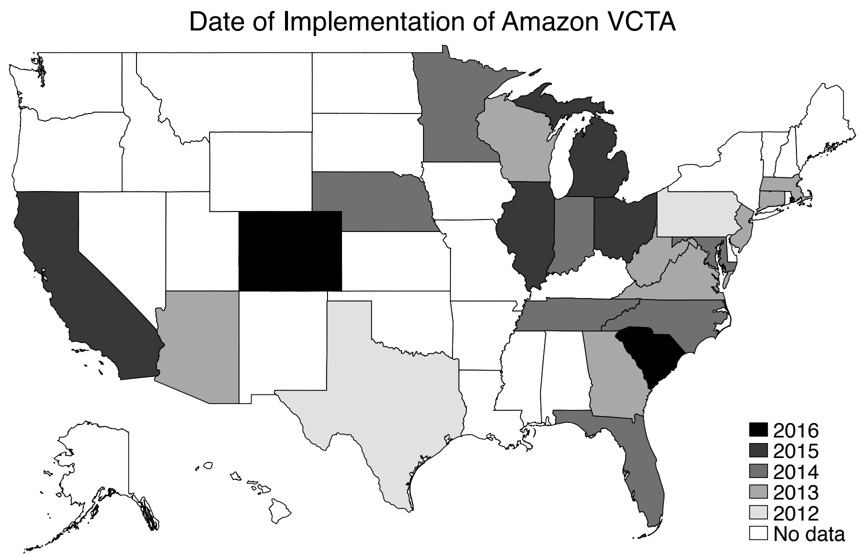

Figure 3.2. Fraction of Trips with Only Taxed Items that Paid No Sales Tax

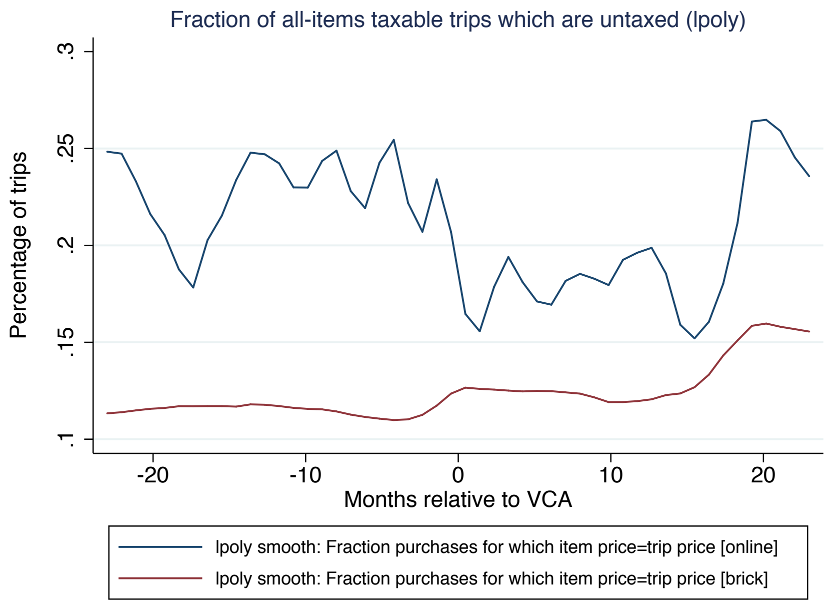

Notes: Observations are on(at) the state-month level, N=2080. Panel is unbalanced due to variation in treatment timing, which determines number of post treatment observation months. For computational reasons, this figure was generated on a random 10% sample of trips.

Figure 3.3. Change in Average Monthly Household Expenditure on Taxable Goods Online

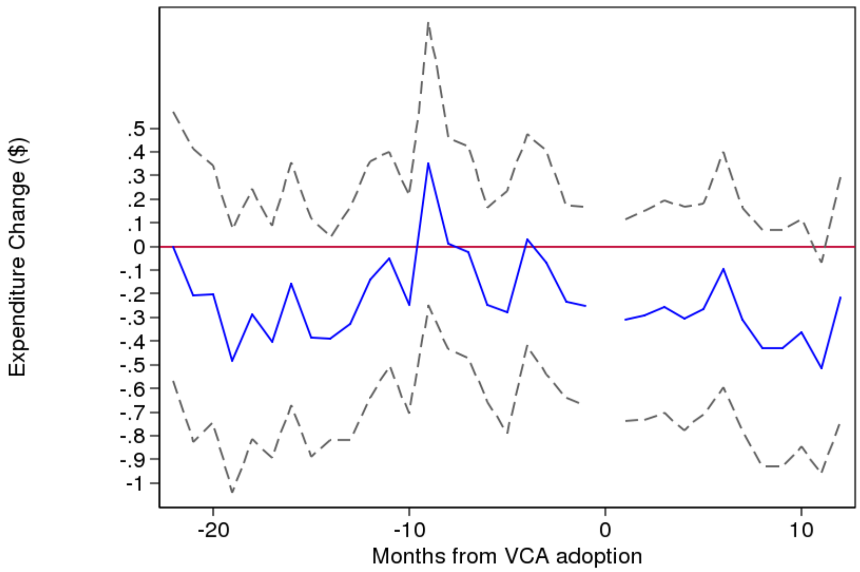

Regression coefficients from spec. with Year and Month FE.
Expenditure is tax-exclusive

Figure 3.4. Change in Average Monthly Household Expenditure on Taxable Goods Online at Large Retailers,

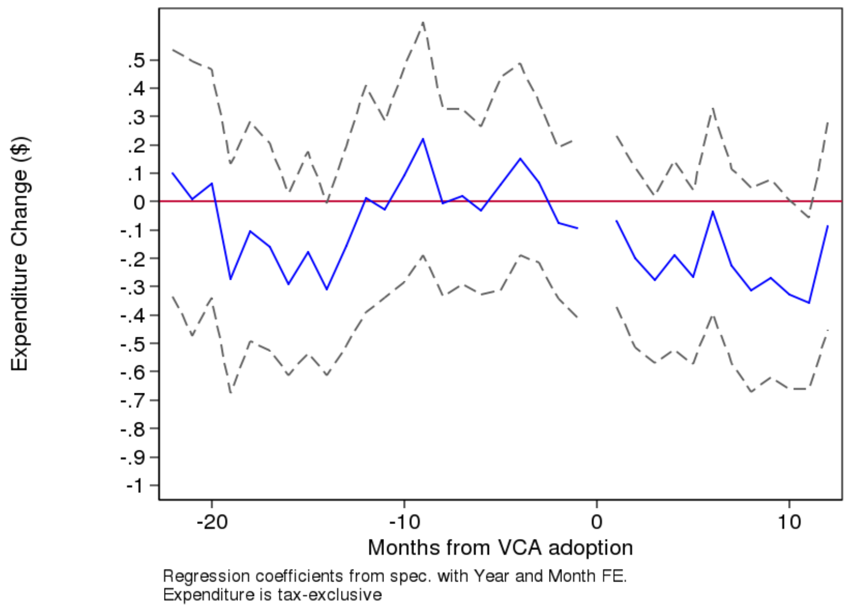

Figure 3.5. Change in Average Monthly Household Expenditure on Taxable Goods Online at Small Retailers

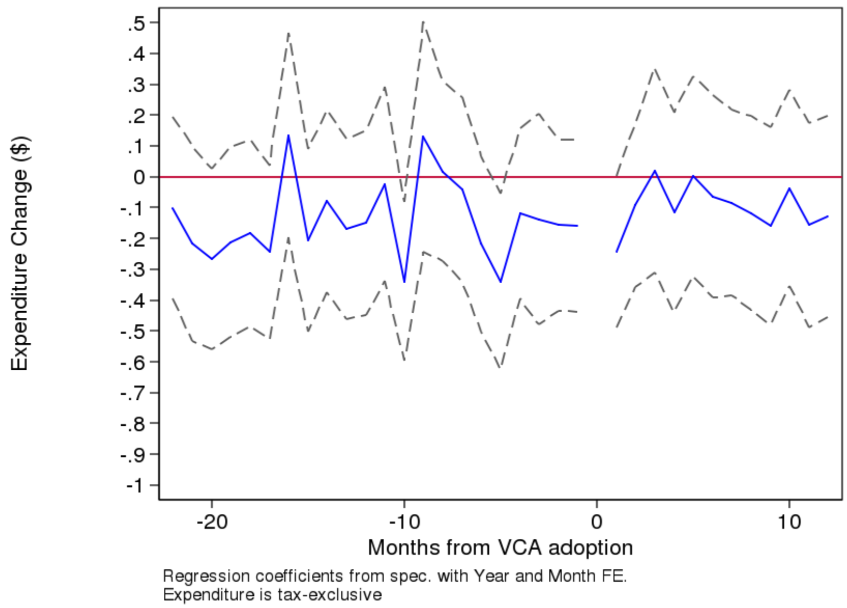

Figure 3.6. Change in Average Monthly Household Expenditure on Exempt Goods Online,

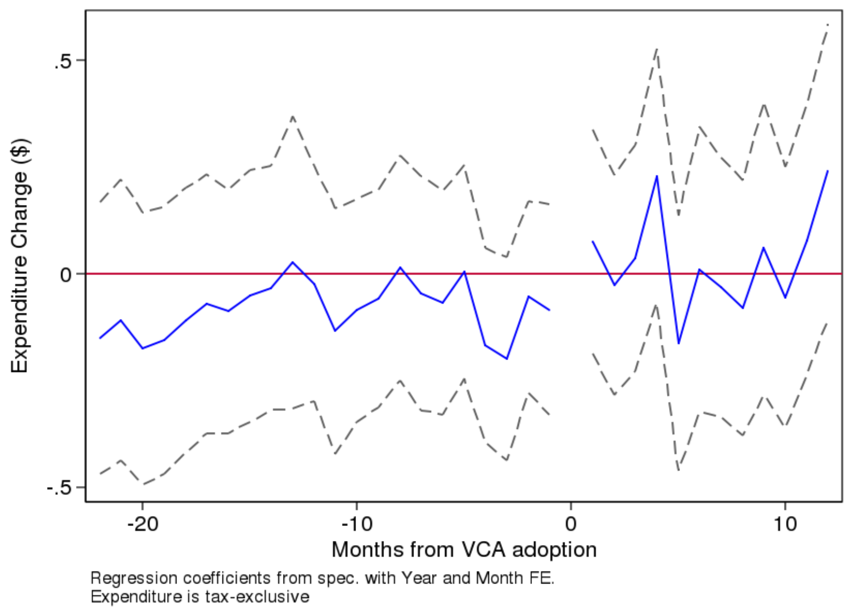

Figure 3.7. Change in Average Monthly Household Expenditure on Taxable Goods at Brick-and-Mortar Stores

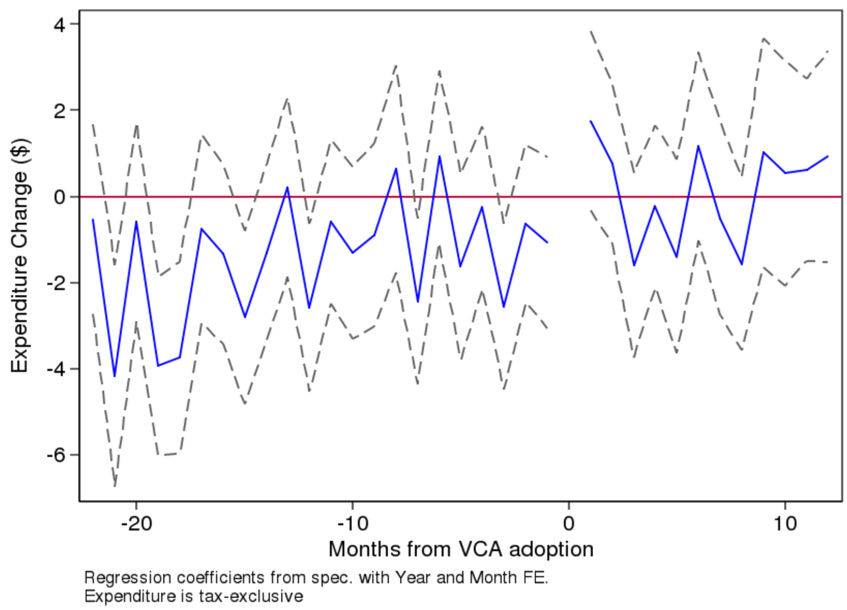

Figure 3.8. Local Sales Tax Rate Changes

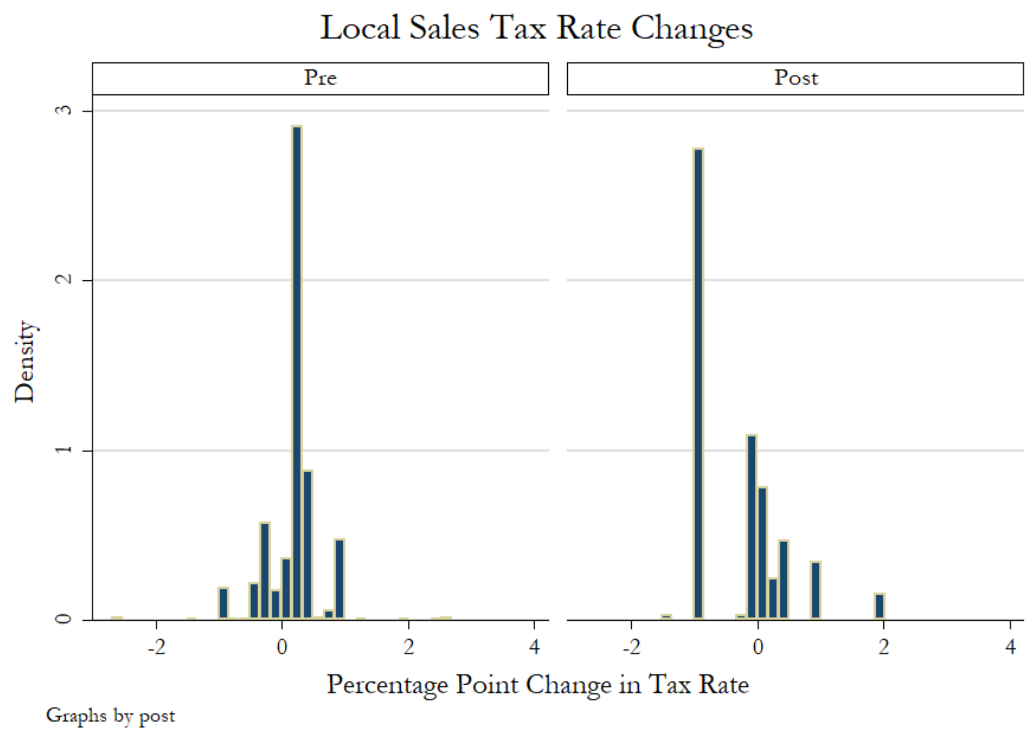

Table 3.1. State and Local Sales Tax by Year

State and Local Sales Tax by Year

	2010	2011	2012	2013	2014	Total
State Rate Changes	3	3	0	8	0	14
County Rate Changes	68	59	65	92	88	372
City Rate Changes	207	247	1109	271	255	2089
Total	278	309	1174	371	343	2475

Table 3.2. Effect of VCA on Change in Average Monthly Household Expenditure

	Dependent Variable: Average Monthly Household Expenditure in Category					
	(1)	(2)	(3)	(4)	(5)	(6)
	Online				Brick-and-Mortar	
	Taxable			Exempt	Taxable	Exempt
	Total	Large Retail	Small Retail			
Treat X Post	-0.231*	-0.247***	0.016	0.118	0.556	0.546
	(0.119)	(0.091)	(0.103)	(0.088)	(0.677)	(1.105)
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3591230	3591230	3591230	3591230	3591230	3591230

Note: Standard errors in parentheses, clustered by state. *p<0.1, **p<0.05, ***p<0.1. Dependent Variable is measured in dollars. Average online taxable expenditure per month is approximately $2.70 (of which $1.20 is at large retail), average online exempt expenditure online is approximately $1.40. Online taxable expenditure declines by 8.5 percent. Online taxable expenditure at large retailers declines by almost 20 percent.

Table 3.3. Pass-through of VCA to Prices, and Effect of VCA on Quantity Demanded Controlling for UPC-level Fixed Effects

	Dependent Variable: Log(Quantity)			Dependent Variable: Log(Tax-Exclusive Price)		
	(1)	(2)	(3)	(4)	(5)	(6)
	Taxable		Exempt	Taxable		Exempt
Treat X Post	-0.007	-0.007	-0.001	-0.009*	-0.001	0.002
	(0.006)	(0.008)	(0.009)	(0.005)	(0.005)	(0.007)
UPC FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs (Purchases)	884647	778420	888773	879392	773596	885008

Note: Standard errors in parentheses, clustered by state. *p<0.1, **p<0.05, ***p<0.1. Columns (2) and (5) exclude videos and CDs/ DVDs.

Table 3.4. Variation in Effect of VCA on Quantity Demanded by Ex-Ante Price of UPC

	Dependent Variable: Log(Quantity)			
	(1)	(2)	(3)	(4)
		Taxable Products		Exempt Products
Treat X Post X 2011 Price	-0.001**	-0.002**	-0.001**	-0.009***
	(0.001)	(0.001)	(0.001)	(0.003)
Treat X 2011 Price	-0.000	-0.000	-0.001*	0.002
	(0.000)	(0.000)	(0.000)	(0.002)
Treat X Post	-0.004	-0.004	-0.007	0.007
	(0.007)	(0.009)	(0.007)	(0.011)
UPC FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes		Yes
Household FE			Yes	
Obs (Purchases)	884647	778420	769024	888773

Note: Standard errors in parentheses, clustered by state. *p<0.1, **p<0.05, ***p<0.1. 2011 price is the average price recorded for each UPC in 2011 across all purchases in the Nielsen Homescan data. Column (2) excludes video products.

Table 3.5. Effect of VCA on Elasticity of the Tax Base.

	Δ Log(Effective Base)				Δ Log (Effective Base-select)	Δ Log(Brick-and-Mortar Base)			
	All	Pre	Pre	All	All	All	Pre	Pre	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ Tax Rate	-3.860***	-4.611***	-3.431*	-3.191	-5.757	-3.913***	-4.547***	-3.354*	-3.110
	(1.466)	(1.539)	(1.980)	(1.986)	(8.809)	(1.470)	(1.538)	(1.980)	(1.988)
Treat X Δ Tax Rate			-2.791	-2.769	3.752			-2.822	-2.786
			(3.182)	(3.154)	(11.864)			(3.180)	(3.151)
Post X Δ Tax Rate				7.828	-10.752				6.297
				(5.390)	(20.024)				(5.539)
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N (Household-Month)	3298724	2714777	2714777	3298724	542098	3296636	2714777	2714777	3296636

Note: Standard errors in parentheses. *p<0.1 **p<0.05 *** p<0.01. Sample restricted to time periods that experience either positive or no change in tax rate. Columns (1), (4), (5),(6) and (9) include all time periods. Columns (2),(3),(7) and (8) online include time periods prior to VCA and therefore also only states that implement the VCA between 2010 and 2014. The "Effective Base- select" includes only expenditure on UPC that are purchased online between 5 and 95 percent of the time. The "Effective Base" is defined as expenditure on brick-and-mortar purchases prior to the VCA (or in the absence of VCA) and expenditure on brick and mortar purchases as well as online expenditure after the VCA.

References

Baugh, Brian, Itzhak Ben-David, and Hoonsuk Park. *Can Taxes Shape an Industry? Evidence from the Implementation of the "Amazon Tax"*. No. w20052. National Bureau of Economic Research, 2014.

Chetty, Raj, Adam Looney, and Kory Kroft. "Salience and taxation: Theory and evidence." *The American economic review* 99.4 (2009): 1145-1177.

Goldin, Jacob, and Tatiana Homonoff. "Smoke gets in your eyes: cigarette tax salience and regressivity." *American Economic Journal: Economic Policy* 5.1 (2013): 302-336.

Goolsbee, Austin, Michael Lovenheim and Joel Slemrod. "Playing with Fire: Cigarettes, Taxes, and Competition from the Internet." *American Economic Journal: Economic Policy* 2.1 (2010):131-154

Slemrod, Joel and Wojchiech Kopczuk. "The Optimal Elasticity of Taxable Income." *Journal of Public Economics* 84.1 (2002): 91-112

APPENDIX

Measuring Tax-Exclusive and Tax-Inclusive Price in the Nielsen Data

This appendix describes our investigation of Nielsen's price data to determine the accuracy with which tax-exclusive and tax-inclusive prices are recorded.

A.1 Are Nielsen's Recorded Prices and Expenditure Tax-inclusive or Tax-exclusive?

The distinction between tax-inclusive and tax-exclusive price is crucial for an analysis of incidence or other impacts of taxation. Nielsen does not explicitly request consumers to enter the tax-exclusive price. Two variables provide information on expenditure. One is the *trip-level* total expenditure, the other is *item-level* expenditure given separately for each item purchased in the trip. Nielsen's documentation states that the *trip-level* total expenditure is tax inclusive but that the *item-level* expenditure is generally exclusive of tax. We test how often this is true by imputing our own measure of total *trip-level* tax inclusive expenditure from the *item-level* expenditure by adding up expenditure on each item, along with our measure of the applicable tax. If the *item-level* expenditure is always tax exclusive, and we are able to accurately impute the tax then the *imputed* measure of the *trip-level* expenditure should match the *actual trip-level* expenditure.

In the Nielsen documentation, they specify a number of reasons the *imputed trip-level* expenditure might not equal the *actual trip-level expenditure* ("total spent"). These include the trip price is generally tax inclusive, whereas the item prices are not; not all items in the trip are recorded by the panelist;[12] not all items purchased by the panelist are tracked by Nielsen (only "fast moving" goods tracked);[13] the scanner malfunctioned; and item price is censored (capped) at $999.99 for non-magnet items.

A.2 Analysis of discrepancy: Predicted vs. Actual Tax-inclusive Expenditures

Applicable tax rates on items are estimated using zip-code level information on local sales tax rate and the exemption status of products recorded in LexisNexis.[14] Any errors in *item-level* expenditure makes it more likely that there are discrepancies between *imputed* and *actual trip-level* expenditure in trips where more than one item was purchased, we separately analyze trips with one item versus multiple items (See Figure A.1 for the respective distributions of items per trip).

We generate two measures of discrepancies in tax inclusive expenditure. First, we calculate the difference between the *imputed trip-level* expenditure and the *actual trip-level* expenditure ("tax discrepancy"). I plot the densities of this measure separately for online and brick-and-mortar purchases. For both markets, there are mass points at common sales tax rates, suggesting an error in correctly applying the tax rather than an error in item price recording (See Figure A.2).

Next, because the *imputed tax-inclusive* expenditure may not have accurately assigned the tax rate, we restrict the sample to trips in which no exempt items were purchased and identify trips in this sample where *imputed* expenditure equals *actual* expenditure.

[12] Nielsen Documentation, p66. "The panelist didn't scan all products purchased. Some items never make it into the home to get scanned. Consider items purchased at a hardware store that might get stored in the garage rather than being brought into the home, or a candy bar that was purchased and eaten before the consumer got home."
[13] Nielsen Documentation, p66. "Some items aren't "coded" by Nielsen – Nielsen mostly tracks fast-moving consumer goods (e.g. not most apparel, electronics or home furnishings, etc.)."
[14] LexisNexis has the least intuitive interface in the history of search. Encyclopedia volumes are faster.

I collapse the total number of such purchases separately for online and BM retailers, from the trip level to the state – treatment month level (approximately 40 periods *50 states= 2080 observations), and plot weighted kernel smoothers for online and BM separately relative to VCA passage (See Figure 3.2). As expected, the number of online purchases with no sales tax is much higher than for brick purchases in the pre-treatment period, and fall sharply after VCA passage. However, the drop in online purchases without sales tax belies a minimal change in the levels: up to 30 months after a VCA, approximately 1 out of 4 purchases are untaxed compared to 1 out of 10 for brick purchases.

Figure A.1. Histogram of Number of Items Purchased per Trip

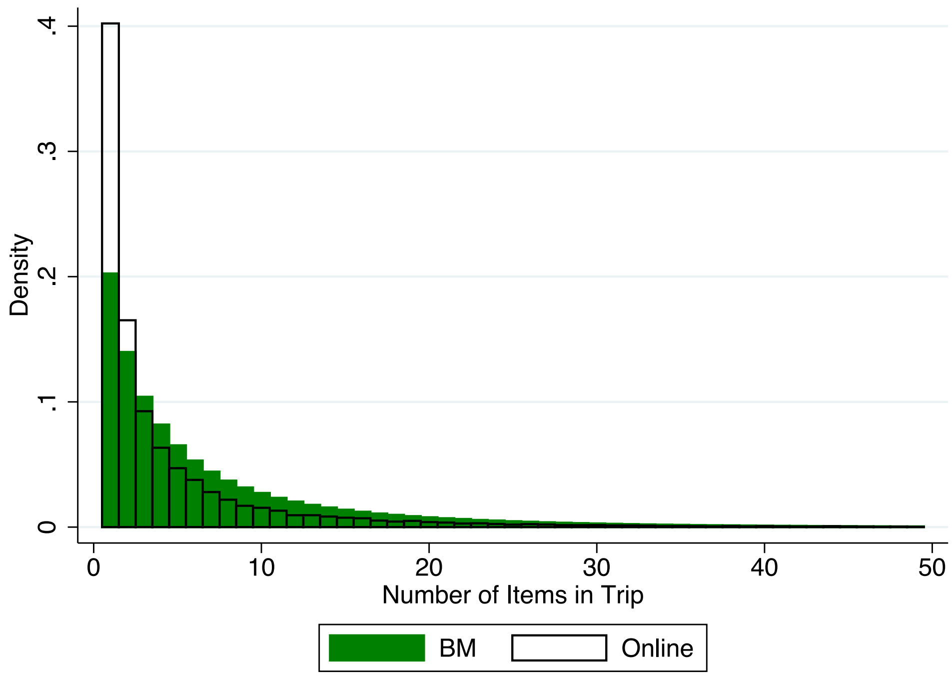

Figure A.2. Discrepancy between Computed and Observed Tax-inclusive Prices