Identifying Tax Compliance from Variation in Tax Policy: Theory and Empirics*

Andrew Bibler
andrew.bibler@unlv.edu
University of Nevada Las Vegas

Laura Grigolon
grigolon@uni-mannheim.de
University of Mannheim; CEPR

Keith F. Teltser
kteltser@gsu.edu
Georgia State University

Mark J. Tremblay
mark.tremblay@unlv.edu
University of Nevada Las Vegas; CESifo

August 21, 2025

Abstract

Governments increasingly use changes in tax rules to combat evasion. We develop a general approach to point-identify tax compliance along with supply and demand elasticities. Identification requires data on prices and quantities, variation in tax enforcement, and a demand or supply shifter. We illustrate our approach using data on Airbnb collection agreements, where taxes are enforced by shifting the statutory burden away from hosts and onto renters via the platform. We find that taxes are paid on roughly zero to 1.6 percent of Airbnb transactions prior to enforcement.

Keywords: taxation, tax compliance, evasion, enforcement, statutory incidence, remittance rules, online markets, sharing economy platforms, Airbnb

JEL Classifications: H20, H22, H26, L10

*We thank David Agrawal, Florian Englmaier, Mohammed Mardan, Martin Peitz, André Stenzel, Frank Verboven, and Jay Wilson for their helpful comments, along with seminar participants at the Mannheim Centre for Competition and Innovation 2024 Annual Conference, the Meeting of the Committee for Industrial Economics (German Economic Association), and the University of Mannheim. Laura Grigolon gratefully acknowledges support from the German Research Foundation (DFG) through CRC TR 224 (Project A02).

1 Introduction

In some markets, tax obligations are ambiguous or difficult to enforce, leading to less tax revenue and a competitive advantage for agents who can more easily evade. For example, in online marketplaces such as Amazon and Airbnb, cooperation between tax authorities and online platforms to increase compliance is now commonplace.[1] Changes in tax policies, such as remittance or enforcement rules, can increase compliance by, for example, changing the method of reporting or collecting, improving tracking of market transactions, shifting collection to the side of the market or platform with higher compliance, or otherwise increasing oversight. Identifying the level of compliance is crucial to determining the value of tax policy efforts. As entering into negotiations with platforms is costly for tax authorities, only a credible and precise quantification of the benefits of tax policy efforts can effectively inform these decisions.[2] We focus on changes in tax policy, including remittance rules, that effectively alter *enforceability* and *compliance* by changing incentives or mechanisms for tax collection.

We start by outlining a framework to estimate elasticities and compliance rates from changes in tax policy, where the level of tax compliance is unknown prior to the change. While the introduction of a fully enforced tax identifies both demand and supply elasticities (Zoutman et al., 2018), a model of tax enforcement (through policy changes) with non-compliance includes an additional parameter, tax compliance before enforcement, and the tax policy change is not sufficient on its own to point identify pre-enforcement compliance (Bibler et al., 2021). We show how to estimate the *pre-enforcement* tax compliance rate, along with demand and supply elasticities, when a change in tax policy leads to a potential change in compliance. We then extend this framework to encompass the identification of

[1] Partial compliance is particularly relevant for online markets. Those markets constitute a significant and growing portion of economic activity. The value of US online transactions is expected to exceed $3 trillion by 2024, approximately 10% of GDP (Statistica Digital Payments).

[2] In addition, even after enforcement is approved, agreements may be terminated or preempted by higher legislative bodies. Two examples of conflicts between local and state legislative bodies applied to Airbnb can be found in Ohio and Florida.

post-enforcement compliance rates as well.

Our framework includes all four possible combinations of shifts in the statutory burden under tax policy changes: (i) the statutory burden shifts from the supply to the demand side; (ii) the statutory burden shifts from the demand to the supply side; (iii) the statutory burden remains on the demand side before and after the change; and (iv) the statutory burden remains on the supply side. We adopt an operational definition of statutory tax burden based on observable pricing structures: the burden resides on the supply side under tax-inclusive pricing, where suppliers remit taxes; the burden resides on the demand side under tax-exclusive pricing, where consumers (or their agents, i.e., platforms) fulfill remittance obligations. This approach reflects the *de facto* allocation of compliance responsibilities that shape enforcement efficacy; our results do not hinge on formal legal liability determinations. Crucially, our identification strategy recovers compliance levels and elasticities from tax policy variation alone, independent of statutory tax rate changes.[3]

When tax compliance is not an issue, two exclusion restrictions are required to identify demand and supply elasticities from variations in the tax rate. In contrast, identifying tax compliance along with demand and supply elasticities necessitates three exclusion restrictions, requiring an additional supply or demand shifter. We outline the identification results, including the necessary exclusion restrictions, in all four remittance scenarios.

In general, our approach requires that either the pre- or post-enforcement compliance rate is known but can be between zero and one. However, when tax policy changes alter the remittance structure by shifting the statutory burden between market sides, our method point-identifies the pre- or post-enforcement compliance rate *without any knowledge* of the level of compliance at any point in time.[4] Finally, if, in conjunction with tax policy changes,

[3]Variation in the statutory tax rate is helpful to identify all the model parameters (compliance, demand, and supply elasticities) in conjunction with tax policy changes, as discussed below.

[4]This result emerges because the magnitudes of the supply and demand shifts each depend on one of the two compliance parameters, with each shift dependent on a different compliance parameter. As shown in Case A of Tables 1 and 2, when the supply shift is written as $\lambda_1 T$ and the demand shift as $\lambda_2 T$, this separation enables recovery of λ_1 independently from λ_2 using a single shifter. The compliance parameters can therefore be expressed *solely* as functions of reduced-form parameters, as demonstrated in the Results columns of Tables 1 and 2.

variation in the tax rate can be used as a demand or supply shifter, we show that all parameters — demand and supply elasticities, as well as pre- and post-enforcement compliance rates — are identified. In particular, the level of compliance is identified at any point in time and under any remittance scenario.

Our identification results significantly advance the literature on estimating tax compliance rates. While Bibler et al. (2021) discuss the absence of point identification under non-compliance in the Airbnb setting and propose a straightforward bounding strategy for pre-enforcement compliance, they neither formalize the conditions for point identification nor consider other remittance structures. We close this gap by: (i) formally stating the assumptions required under any remittance structure; (ii) providing point identification arguments where those assumptions hold; and (iii) extending the analysis to post-enforcement compliance. The extended arguments nest the bounding strategy proposed by Bibler et al. (2021) and imply analogous bounding strategies for situations in which an additional demand or supply shifter is unavailable.

Because our framework is adaptable to any change in tax policy, it provides a solution for estimating compliance rates, demand and supply elasticities in various settings, including e-commerce sales taxes, taxes on firms (local or trade tariffs), and taxes within the supply chain. This flexibility is crucial as tax policies can vary while still fitting into one of the four cases we cover. For instance, Airbnb has entered into numerous Voluntary Collection Agreements (VCAs) with state and local governments worldwide, whereby Airbnb collects taxes on applicable transactions (primarily transient occupancy taxes, sales taxes, and other similar lodging taxes) and remits them to the tax jurisdiction on behalf of the renters rather than relying on individual hosts to collect and remit. Similarly, Amazon is now required to collect sales taxes at checkout, rather than rely on consumer-based compliance. The model flexibility also extends to different types of variation in enforcement (temporal or cross-sectional). While we cast our framework in terms of temporal variation (before and after a change in enforcement), our approach generalizes to cases with cross-sectional variation in enforcement as well. For example, enforcement or monitoring efforts can vary across firms based on size (Almunia and Lopez-Rodriguez, 2018; Bachas and Soto, 2021), and auditing efforts may vary across individuals (Kleven et al., 2011).

We empirically illustrate our model using the tax collection agreements between Airbnb and several state and local governments. These agreements result in a switch from an unenforced period to full enforcement and a shift in the statutory incidence from the supply (hosts) to the demand side (renters) via the platform. In this case, the necessary restrictions comprise two restrictions resembling the Ramsey Exclusion Restriction (RER), one on the demand side and one on the supply side (Ramsey, 1927; Zoutman et al., 2018), plus one Standard Exclusion Restriction (SER) based on an additional demand shifter. These three restrictions identify the elasticity of supply, the elasticity of demand, and the pre-enforcement rate of tax compliance.

When Airbnb enters a VCA, prices shift from tax-inclusive (host-remitted) to tax-exclusive (platform-remitted on behalf of renters). Under our framework, this reassigns the statutory burden from the supply side to the demand side. This variation in tax policy constitutes a change in remittance rules rather than traditional enforcement mechanisms like penalties or audits (Slemrod, 2019). Unlike voluntary compliance by individual hosts, which may be endogenously determined by market conditions (Allingham and Sandmo, 1972), VCAs leverage Airbnb's centralized platform for streamlined collection, third-party reporting, and legal liability. This remittance shift thus effectively enhances enforceability by enabling stricter oversight and reducing evasion opportunities.

We use data on the Airbnb accommodation market, including prices and bookings during the pre- and post-enforcement periods in 24 metropolitan areas in the US. In addition, we construct three alternative variables that act as plausibly exogenous demand shifters: (i) the number of incoming flight passengers; (ii) the monthly search volume for hotels from Google Trends in a given metro; and (iii) the monthly search volume from Google Trends for Airbnb rooms. Identification rests on a difference-in-differences design that exploits the staggered roll-out of the agreements across cities, months, and tax rates. Although Airbnb's decision to sign a VCA could, in principle, reflect unobserved cost or market conditions, the timing of agreements is negotiated by the platform with local tax authorities rather than chosen unilaterally by the hosts. Arguably, this staggered, externally driven adoption provides a quasi-exogenous source of variation distinct from host-level compliance decisions. Event-study estimates show pre-treatment coefficients statistically indistinguishable from zero, indicating no differential pre-trends and supporting the validity of our identification strategy.

The estimated coefficients from the main Poisson specifications result in a market-level demand elasticity ranging between -0.41 and -0.66 and a supply elasticity between 1.42 and 2.17.[5] Taxes are paid on up to 1.6 percent of Airbnb transactions before enforcement.[6] All demand shifters yield similar estimates, and all specifications reject a 20% compliance rate at the 10% level. Using our approach to test for heterogeneity in compliance rates is straightforward. We illustrate this by distinguishing between listings operated by individual and professional hosts and find that pre-enforcement compliance rates for listings managed by professional hosts are significantly above zero.

Our results suggest that the Airbnb tax collection agreements addressed a substantial tax evasion issue, as pre-enforcement compliance is virtually null for the majority of listings. Our most conservative estimated compliance rate of 3.5% is substantially lower than the upper bound of 24% estimated by Bibler et al. (2021), demonstrating the practical value of point-identification of compliance and determining the associated benefits of enforcement. By predicting counterfactual full-tax-compliance prices and bookings in the pre-enforcement periods, we estimate that the lack of pre-enforcement compliance resulted in lost tax revenue of $178 per property-year (roughly $1,827,000 per jurisdiction-year) on average among the

[5]Bibler et al. (2021) obtain similar estimates of market-level demand elasticity; Bibler et al. (2021) and Farronato and Fradkin (2018) estimate similar supply elasticities.

[6]The most conservative estimates, using a linear specification, show a level of compliance up to 3.5%

treated jurisdictions in our estimation sample.[7]

Related Literature Our work contributes to the literature focused on tax evasion, particularly the research studying compliance in the presence of changes in remittance and enforcement regimes: Kopczuk et al. (2016), Baugh et al. (2018), Bibler et al. (2021), Fox et al. (2022), Agrawal and Shybalkina (2023), Waseem (2023), and Carrillo et al. (2023).[8] We directly build on the work of Bibler et al. (2021); the authors use reduced-form techniques to infer an upper bound of 24% on pre-enforcement tax compliance. We advance this literature by providing a framework that leverages a tax enforcement change along with an additional exogenous shifter to point-identify compliance. The expanded framework embeds the bounding results of Bibler et al. (2021) and extends the point-identification argument across various settings. Precise identification of compliance is fundamental to gaining a sense of the value of tax enforcement efforts and the plausibility of ex-ante counterfactual evaluations for jurisdictions that have yet to enter a tax agreement.[9]

More generally, the proposed framework advances the literature on using tax variation to identify demand and supply elasticities. Zoutman et al. (2018) demonstrate how variation in tax rates point-identifies both the supply and demand elasticities in a competitive model with full compliance. Dearing (2022) generalizes Zoutman et al. (2018) to markets with imperfect competition while maintaining the assumption of full tax compliance. We focus on modeling variation in tax enforcement, including remittance rules, in the presence of potential non-compliance in competitive markets, which we also extend to markets with imperfect competition.

Our empirical application to Airbnb also contributes to the growing literature on regu-

[7]For reference, the average predicted nightly booking price is $108.67, the average predicted nights booked is 1.39 per property month or 14,296 per jurisdiction-month, and the average pre-enforcement combined tax rate was 9.8%.

[8]Slemrod (2019) provides an overview of economic research in tax compliance and enforcement. Most studies focus on income taxes, showing substantial tax evasion when income is not subject to "enforcement" in the form of third-party reporting (e.g., Pissarides and Weber, 1989; Feldman and Slemrod, 2007).

[9]For example, Farronato and Fradkin (2022) implicitly assume that hosts do not pay lodging taxes when simulating the impact of tax regime changes on the Airbnb market. Our paper effectively validates their assumption.

lating the market for short-term rentals: Jia and Wagman (2020), Bekkerman et al. (2023), Chen et al. (2023), Jin et al. (2023). Our results suggest that tax evasion was rampant before the introduction of regulation and that collection agreements effectively closed the gap in tax treatment between Airbnb and brick-and-mortar hotels.

2 Background

Changes in remittance and enforcement rules can have profound effects on tax compliance (Slemrod, 2019). We show how these changes can be exploited to identify compliance. In this section, we discuss examples of all four possible combinations of statutory-incidence shifts that we cover in our theoretical framework and outline in Table 1. For platform-mediated transactions, we define the side of the market ultimately bearing the statutory burden as the side on behalf of which the platform collects and remits. In general, the statutory burden (or statutory incidence) falls on the party legally obligated to remit a tax, distinct from the economic burden determined by elasticities. The legal responsibility, in practice, aligns with pricing conventions: tax-inclusive prices mean that suppliers remit the tax, so the statutory burden lies on the supply side; when prices are tax-exclusive, the tax is added at checkout and the consumer (often via the platform) is the liable party, so the burden lies on the demand side. A shift from tax-inclusive to tax-exclusive pricing, such as through policy changes or platform-level agreements, reassigns the statutory burden from the supply side to the demand side.

The first example is the "Airbnb case" (Case A in Table 1), which is also the subject of our empirical application. Since 2014, the platform has entered into collection agreements with local jurisdictions, which shift the remittance obligation from the property host (supply) to the renter (demand) via the platform and directly affect the enforceability of taxation. We can safely assume that, after the change in the remittance rule, compliance is practically full as Airbnb takes measures to avoid off-platform transactions.[10] In addition, substitution

[10]For example, guests and hosts cannot exchange contact information prior to booking.

to alternative peer-to-peer home-sharing platforms is likely negligible as their market share is small with respect to Airbnb, which can offer significant network effects to hosts and renters.[11]

The second example, the "Amazon case" (Case B in Table 1), falls within the scope of regulating the taxation of online retail sales, which developed in three waves (Einav et al., 2014; Fox et al., 2022; Agrawal and Shybalkina, 2023). First, between 2011 and 2015, several state legislatures started to enforce the collection of sales tax on Amazon, the largest online retailer, at checkout (the Amazon Tax). Then, the 2018 *Wayfair* decision eliminated the physical presence nexus standard, ruling that the economic presence in a state is enough to subject a seller to a state's sales tax collection requirement. However, as sellers' compliance with the use tax was low due to limited enforcement capacity (Manzi, 2012; Agrawal and Mardan, 2019), a third wave of legislation, the Marketplace Facilitator laws, required all platforms that host a large number of smaller sellers to collect sales tax on all transactions on the platform. Empirical evidence from Fox et al. (2022) suggests that compliance is full or nearly full following legislation. Following Baugh et al. (2018), we treat the Amazon tax, which enforced consumer-based compliance via the platform without changing the statutory incidence, as an example of increased enforcement in which the burden remains on the customers (the demand side). Although the Amazon case is atypical because the platform also acts as the seller, this representation is consistent with our framework: the statutory burden is defined by the side for which the platform remits. Because the posted price is tax-exclusive both before and after the tax policy change, we view this as an example of a pure demand-side shift, with no corresponding change on the supply side.

Another significant application of our method pertains to estimating non-compliance in trade tariffs. The existing literature typically relies on reported import and export data, along with changes in trade tariffs, to infer evasion (as developed by Fisman and Wei, 2004). This approach only captures changes in evasion relative to tariff adjustments; it does not

[11] Bibler et al. (2021) do not find a significant impact on the number of properties listed following the introduction of the collection agreements.

provide a direct measure of evasion levels. In contrast, if countries enforce tariffs on suppliers digitally, as studied in Kitsios et al. (2020), our method can be used to measure the true level of trade tariff evasion. A change in tariff enforcement through digitalization does not affect the statutory burden, which remains on the supply side (Case C in Table 1).[12]

Finally, Kopczuk et al. (2016) leverage a change in the statutory incidence of diesel taxes, showing that diesel taxes statutorily levied on wholesalers and distributors raise more revenues than equivalent taxes on retailers. Shifting the statutory incidence up the supply chain, akin to a shift from the demand side (retailers) to the supply side (wholesalers and distributors) of the market (Case D in Table 1), directly affects the enforceability of taxation and compliance.

3 The Conceptual Framework

In this section, we start by presenting the standard model for estimating demand and supply elasticities, which assumes full compliance, and then develop our framework for identifying compliance when unobserved tax evasion exists. We do so by first considering the case that applies to our empirical example from Airbnb and then discussing how the method generalizes to all other cases where statutory burdens or enforcement policies differ. Lastly, we present other generalizations, including partially salient taxes, imperfect competition, and incorporating variation in tax rates.

3.1 The Standard Model: Full Compliance

We start by outlining the standard model developed by Zoutman et al. (2018) which assumes full compliance. Assume that we have equilibrium price and quantity panel data for a good.

[12]Case C in Table 1, where remittance and enforcement remain on the supply side, also applies to cases in which firms may evade local taxes. For example, Waseem (2023) leverages a tax reform in Pakistan to study VAT evasion via ghost firms. While the remittance rule and the statutory incidence do not change, remaining on the supply side, the reform dramatically reduced the tax liability for certain goods, effectively modifying the evasion incentives and, as a consequence, the level of compliance after the reform.

The index i can indicate a region, a firm, or an individual, and the index t denotes time. The following structural equations characterize demand and supply, respectively:

$$
\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \gamma^d T_{it} + v_{it}^d, \\
y_{it} &= \varepsilon^s p_{it} + \gamma^s T_{it} + v_{it}^s,
\end{aligned}
$$

where y_{it} denotes the logged quantity and p_{it} the logged price; thus, the price coefficients (ε^d and ε^s) represent elasticities. The demand and supply disturbances are denoted by v_{it}^d and v_{it}^s.[13] The term $T_{it} = f(\tau_{it})$ is a function of the ad-valorem tax rate, τ_{it}, such that y_{it} is linear in T_{it}. The tax rate τ_{it} is assumed to be exogenous (possibly after controlling for a vector of covariates in the empirical application). The demand and supply equations result in the following reduced-form equations for quantity and price:

$$
\begin{aligned}
y_{it} &= \frac{\gamma^d \varepsilon^s - \gamma^s \varepsilon^d}{\varepsilon^s - \varepsilon^d} T_{it} + \zeta_{it}^y, \\
p_{it} &= \frac{\gamma^d - \gamma^s}{\varepsilon^s - \varepsilon^d} T_{it} + \zeta_{it}^p.
\end{aligned}
$$

Let π_{Ty} and π_{Tp} denote the reduced-form coefficients. The relationship between reduced-form and structural coefficients can be represented as follows:

$$
\begin{aligned}
\pi_{Ty} &= \frac{\gamma^d \varepsilon^s - \gamma^s \varepsilon^d}{\varepsilon^s - \varepsilon^d}, \\
\pi_{Tp} &= \frac{\gamma^d - \gamma^s}{\varepsilon^s - \varepsilon^d}.
\end{aligned}
$$

To identify the structural demand and supply elasticities, Zoutman et al. (2018) make two assumptions. First, the Standard Exclusion Restriction (SER) states that the tax is levied on the demand side: $\gamma^s = 0$. Second, the Ramsey Exclusion Restriction (RER) states that demand depends only on the price after taxation: $\gamma^d = \varepsilon^d$. Imposing SER and RER

[13] This follows the specification adopted by Zoutman et al. (2018). The x_{it} terms included in Zoutman et al. (2018) are omitted for simplicity.

generates a system of two equations with two unknowns. Solving for ε^d and ε^s yields:

$$\varepsilon^s = \frac{\pi_{Ty}}{\pi_{Tp}}, \tag{1}$$

$$\varepsilon^d = \frac{\pi_{Ty}}{1+\pi_{Tp}}. \tag{2}$$

To illustrate how tax evasion impacts the identification of the structural parameters, Figure 1 presents two cases of tax enforcement in the presence of evasion: home-sharing (Airbnb) and online retail (Amazon). Figure 1a represents the Airbnb case, in which tax enforcement changes the statutory burden from the supply side, where the fraction of tax-compliant transactions before enforcement is denoted by λ_1, to the demand side, where the fraction of tax-compliant transactions is denoted by λ_2. Figure 1b represents the Amazon case, where the statutory burden initially falls on consumers, and a share λ_1 pays taxes. After Amazon enforces sales taxes at checkout, the burden remains on the demand side where all consumers ($\lambda_2 = 1$) now pay the tax. The fundamental divergence from the model of a tax introduction with full compliance is that the magnitude of the shift of one function (either supply or demand) resulting from an enforcement change depends on the pre-enforcement compliance parameter (λ_1).

3.2 Identification of Compliance: The Airbnb Case

We extend the framework proposed by Zoutman et al. (2018) to account for tax evasion. For simplicity, we begin our discussion by focusing on the example of tax collection agreements applied in Airbnb markets. Similar intuition carries through to all other possible remittance structures, as well as the identification of post-enforcement compliance rates (as we discuss in the next subsection).

We have a two-period framework (before and after a tax policy change). In the first period, the level of tax compliance (λ_1) is unknown. In the second period, the change in enforcement goes into effect, and the level of tax compliance post-change (λ_2) is known.

In practice, we assume that $\lambda_2 = 1$ because all renters pay taxes at the point of sale (we will revisit this assumption in Section 3.3). We show how to identify the unknown level of tax compliance before the change in enforcement, as well as the elasticities of supply and demand.

To start, consider the following updated system of demand and supply:

$$
\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \gamma^d D_t \cdot T_{it} + \rho^d Z_{it} + v^d_{it}, \\
y_{it} &= \varepsilon^s p_{it} + \gamma^s D_t \cdot T_{it} + \rho^s Z_{it} + v^s_{it},
\end{aligned}
$$

where D_t is a dummy variable equal to one if enforcement exists in period t and Z_{it} denotes an additional variable acting as a demand or supply shifter. Following Zoutman et al. (2018), we assume that the structural equations are written in logarithms; thus, price coefficients are the structural demand and supply elasticities. We represent the demand-supply system in the following reduced-form equations for quantity and price:

$$
\begin{aligned}
y_{it} &= \frac{\gamma^d \varepsilon^s - \gamma^s \varepsilon^d}{\varepsilon^s - \varepsilon^d} D_t \cdot T_{it} + \frac{\rho^d \varepsilon^s - \rho^s \varepsilon^d}{\varepsilon^s - \varepsilon^d} Z_{it} + \zeta^y_{it}, \\
p_{it} &= \frac{\gamma^d - \gamma^s}{\varepsilon^s - \varepsilon^d} D_t \cdot T_{it} + \frac{\rho^d - \rho^s}{\varepsilon^s - \varepsilon^d} Z_{it} + \zeta^p_{it}.
\end{aligned}
$$

Let π_{Ty}, π_{Tp}, π_{Zy}, and π_{Zp} capture the four reduced-form coefficients in the two equations above. The relationship between reduced-form and structural coefficients can be represented as follows:

$$
\begin{aligned}
\pi_{Ty} &= \frac{\gamma^d \varepsilon^s - \gamma^s \varepsilon^d}{\varepsilon^s - \varepsilon^d}, \\
\pi_{Tp} &= \frac{\gamma^d - \gamma^s}{\varepsilon^s - \varepsilon^d}, \\
\pi_{Zy} &= \frac{\rho^d \varepsilon^s - \rho^s \varepsilon^d}{\varepsilon^s - \varepsilon^d}, \\
\pi_{Zp} &= \frac{\rho^d - \rho^s}{\varepsilon^s - \varepsilon^d}.
\end{aligned}
$$

The assumptions necessary to identify elasticities and measures of tax evasion depend on which side of the market bears the statutory tax burden in the pre- and post-enforcement periods. In the case of the collection agreements applied in Airbnb markets, the statutory burden falls on the supply side pre-enforcement and shifts to the demand side after an agreement is in place (Case A in Table 1). To identify the parameters for Case A, we make the following assumptions:

Assumption 1. *Standard Exclusion Restriction (SER2). The variable Z_{it} is a demand shifter and does not appear in the structural supply equation: $\rho^s = 0$.*

Assumption 2. *Ramsey Exclusion Restriction (RER). Demand depends only on the price after taxation: $\gamma^d = \varepsilon^d$.*

Assumption 1 (SER2) is a standard exclusion restriction implying that Z_{it} acts as a demand shifter. Our SER2 exclusion restriction differs from the Standard Exclusion Restriction used in Zoutman et al. (2018), which states that, if the tax is levied on the demand side, $\gamma^s = 0$. In our case, we cannot rely on such an exclusion restriction on the tax because the change in enforcement is accompanied by a shift in the statutory burden from one side of the market to the other. Assumption 2 is the Ramsey Exclusion Restriction used in Zoutman et al. (2018); it states that demand depends on the after-tax price.

Under Assumptions 1 and 2, we express the supply and demand elasticities as follows:

$$\varepsilon^s = \frac{\pi_{Zy}}{\pi_{Zp}}, \tag{3}$$

$$\varepsilon^d = \frac{\pi_{Ty}}{1 + \pi_{Tp}}. \tag{4}$$

In addition, if the SER2 and RER hold, we can solve for γ^s as follows:

$$\gamma^s = \pi_{Ty} - \varepsilon^s \pi_{Tp}. \tag{5}$$

To identify the level of compliance prior to enforcement (λ_1) separately from the elasticity of supply (ε^s), we make a third assumption:

Assumption 3. *Ramsey Exclusion Restriction λ_1 (RERλ_1). The magnitude of the supply shift due to the tax can be represented as follows: $\gamma^s = \lambda_1 \varepsilon^s$, where $\lambda_1 \in [0,1]$ captures the tax compliance rate in the market before the change in enforcement.*

Similar to the RER assumption for demand, Assumption 3 relates the magnitude of the supply-side response from tax enforcement to the supply elasticity, ε^s, which must be scaled by the tax compliance rate, λ_1. Combining Assumption 3 with Equation (5), we solve for the tax compliance rate in the first period:

$$\lambda_1 = -\pi_{Tp} + \frac{\pi_{Ty}}{\varepsilon^s}. \tag{6}$$

The expression relates λ_1 and ε^s and describes the identification problem. To separately identify the level of compliance (λ_1) and the elasticity of supply (ε^s), we need a restriction on ρ^s; that is, we need a demand shifter Z_{it} to identify the elasticity of supply (ε^s).[14] The demand shifter is necessary to determine the portion of the enforcement-induced price change attributable to a shift in the supply curve rather than a movement along the supply curve. With an estimated elasticity of supply identified by the variation in a demand shifter, we can solve for the level of compliance (λ_1).[15]

In Equation (6), λ_1 has two components. The first one is the price change resulting from the tax enforcement, which represents an upper bound on the level of compliance; in

[14] When supply is perfectly inelastic ($\varepsilon^s = 0$), pre-enforcement compliance is not identified because supply is not a function of the enforced tax rate. Identification requires that tax enforcement affects both sides of the market.

[15] An alternative presentation of the identification argument can be made using the conditional expectation of y_{it}. Consider the case with an initial supply-side remittance that transitions to the demand side with full compliance after an enforcement change (as in the Airbnb case). Focusing on the supply equation with endogenous prices and assuming that $E[v^s_{it}|T_{it}, Z_{it}, D_t] = 0$:

$$E[y_{it}|T_{it}, Z_{it}, D_t] = \varepsilon^s E[p_{it}|T_{it}, Z_{it}, D_t] + \lambda_1 \varepsilon^s T_{it} D_t.$$

As demonstrated in Section 3.2, independent variation in Z_{it} and $T_{it}D_t$ separately identifies ε^s and λ_1. Further, variation in any two of T_{it}, Z_{it}, or D_t can be used to identify both parameters. We revisit the possibility of relying on variation in T_{it} and D_t for identification in Section 3.4.4.

the extreme case in which supply is perfectly elastic, the price change is solely due to a shift of a horizontal supply function. The second component describes the amount of the price change that can be attributed to a movement along the supply curve. Intuitively, the difference in the total change and the change explained by a movement along the supply curve is attributed to the shift in the supply function due to the alleviation of the statutory burden among the fraction of compliant suppliers. Compliance can be estimated using the estimates for each component of Equation (6).

3.3 General Scope

Our framework encompasses all four possible combinations of shifts in the statutory burden determined by an enforcement policy: (A) the statutory burden shifts from the supply to the demand side after the change; (B) the statutory burden remains on the demand side; (C) the statutory burden remains on the supply side; and (D) the statutory burden shifts from the demand to the supply side. In addition, the framework is adaptable to allow for the identification of post-enforcement compliance in cases where it is unknown.

First, we focus on the identification of pre-enforcement compliance (λ_1). We summarize the necessary assumptions and identification results for each remittance structure in Table 1. The table presents the general case in which post-enforcement compliance can be less than full ($\lambda_2 \neq 1$). Appendix A provides additional information on the necessary assumptions reported in Table 1 for Cases B to D that are not formally treated in this Section. As outlined in Table 1, the identification assumptions and results depend on the statutory incidence before and after the change in enforcement. Under Cases A and C, the magnitude of the enforcement-induced demand shift is known (either no shift or related to the size of the tax), but Z_{it} must act as a demand shifter to disentangle compliance and the elasticity of supply; hence the assumption that $\rho^s = 0$. In contrast, under Cases B and D, the size of the supply shift with enforcement is known, but Z_{it} must act as a supply shifter to disentangle compliance and the elasticity of demand; hence the assumption that $\rho^d = 0$. Using an additional indicator denoting which side of the market bears the statutory burden pre- and post-enforcement, we present a more parsimonious version of our conceptual framework, from which the four cases can be derived, in Appendix A, Section A.2.

Second, we illustrate the identification of post-enforcement compliance (λ_2). Table 2 summarizes the necessary assumptions and identification results for each remittance structure. This table presents the general case in which pre-enforcement compliance can be less than full ($\lambda_1 \neq 1$). As before, the identification assumptions and results depend on the statutory incidence before and after the change in enforcement. For example, in Case A, if λ_1 is known, the magnitude of the supply shift caused by enforcement is also known; an exogenous supply shifter is needed to separately identify the elasticity of demand and λ_2.

Finally, in Cases A and D of Table 1, it is worth highlighting that the identification of the tax compliance parameter pre-enforcement (λ_1) does *not* require knowing the tax compliance parameter post-enforcement (λ_2).[16] This symmetrically holds for Cases A and D of Table 2 as well: the identification of the tax compliance parameter post-enforcement (λ_2) does not require knowing the tax compliance parameter pre-enforcement (λ_1).[17] This feature renders our method fully general to estimate compliance when enforcement shifts the statutory burden from one side of the market to the other.

3.4 Extensions

Our method of identifying compliance can apply when the tax is not fully salient, competition is imperfect, and variation in the tax rate can be used as a demand or supply shifter together with changes in tax enforcement. In the remainder of this section, we outline precisely how.

[16]Without knowledge of the tax compliance parameter post-enforcement (λ_2), the demand elasticity is not point-identified in Case A, but is bounded. In Case D, the supply elasticity is not point-identified, but is bounded.

[17]In Cases B and C of Table 1 and Table 2, the change in compliance ($\lambda_2 - \lambda_1$) is identified without knowledge of the level of compliance at any point.

3.4.1 Salience

Online prices may not be fully salient to customers (Chetty et al., 2009; Blake et al., 2021). This is naturally a concern in our setting once tax enforcement occurs, especially in our application to the Airbnb market. Partial salience can take on various forms across the market settings considered in Table 1. In general, our approach cannot separately identify salience and compliance. Point identification of salience, along with tax compliance, would require an additional exclusion restriction, which depends on the remittance rules.[18] In this section, we show that the tax compliance parameter derived above (Case A in Table 1) is unaffected by incomplete salience after enforcement. Note that the same assertion is symmetrically valid for Case D in Table 1, as well as Cases A and D when identifying post-enforcement compliance as shown in Table 2.

Salience affects the Ramsey Exclusion Restriction. Assumptions 1 (SER2) and 3 (RERλ_1) remain unaltered. We modify Assumption 2 as follows:

Assumption 2′. *The Ramsey Exclusion Restriction under imperfect Salience (RERS): Demand depends on the total salient price after taxation so that $\gamma^d = \varphi \cdot \varepsilon^d$, where $\varphi \in [0, 1]$ denotes the degree of salience of the tax to consumers.*

The expressions for π_{Zy} and π_{Zp} are unchanged; however, π_{Ty} and π_{Tp} become:

$$\pi_{Ty} = \frac{\varphi\varepsilon^d \cdot \varepsilon^s - \gamma^s\varepsilon^d}{\varepsilon^s - \varepsilon^d},$$

$$\pi_{Tp} = \frac{\varphi\varepsilon^d - \gamma^s}{\varepsilon^s - \varepsilon^d}.$$

[18] The point is highlighted by Dearing (2022) in Appendix B. Before enforcement, salience and non-compliance can coexist as hosts may be non-compliant or simply unaware of their tax obligations (the tax is not salient). The distinction is irrelevant in the Airbnb setting, as the lack of salience would lead to non-compliance, which is our object of interest. In other settings, this may not be true. For instance, in Chetty et al. (2009), non-compliance with the sales tax by stores and non-salience of the sales tax to the customer are observationally equivalent. A change in enforcement would identify either non-compliance or salience, depending on the remittance rules.

Similarly, the expressions for ε^s and γ^s are unaltered, but the elasticity of demand becomes:

$$\varepsilon^d = \frac{\pi_{Ty}}{\varphi + \pi_{Tp}}.$$

The equation for λ_1 follows Equation (6). Differentiating with respect to the salience parameter reveals that the implied level of pre-enforcement compliance is unaffected by changes in salience:

$$\frac{d\lambda_1}{d\varphi} = -\frac{d\pi_{Tp}}{d\varphi} + \frac{d\pi_{Ty}}{d\varphi} \cdot \frac{1}{\varepsilon^s} = -\frac{\varepsilon^d}{\varepsilon^s - \varepsilon^d} + \frac{\varepsilon^d \cdot \varepsilon^s}{\varepsilon^s - \varepsilon^d} \cdot \frac{1}{\varepsilon^s} = 0.$$

Intuitively, incomplete salience impacts the estimated effect of tax enforcement on prices and quantities; these effects offset each other when calculating the level of pre-enforcement compliance, λ_1. Thus, changes in salience do not impact the identification of pre-enforcement tax compliance. Incomplete salience after enforcement, if present, would affect the estimated market elasticity of demand but not the compliance rate.[19]

3.4.2 Related Goods

In Section 3.3, we follow Zoutman et al. (2018)'s presentation, focusing on one-good markets. In multi-product settings, the presence of related goods introduces cross-price elasticities, which may affect the identification of key parameters such as demand elasticity. For example, in the Airbnb case, one may ask whether demand-side substitution between Airbnb and hotels may affect our identification argument. In our baseline model (The Airbnb Case), omitting the prices of related goods from the demand equation would lead to omitted variable bias in the OLS estimators for the reduced-form coefficients (π_{Ty} and π_{Tp}), leading to a biased demand elasticity estimator (ε^d). Importantly, if a related-good price p_{it}^k enters demand but is excludable in the supply equation, any bias in the reduced-form coefficients π_{Ty} and π_{Tp} cancels out in the ratio that delivers λ_1.[20] Similarly to the identification argument of

[19] If salience is less than one, the estimated demand elasticity is attenuated.

[20] As discussed in footnote 14, this is also apparent from the identification argument for ε^s and λ_1 through the supply equation, which does not rely on the estimation of ε^d.

the compliance parameter when prices are not fully salient (Section 3.4.1), a modification of the demand equation does not impact the identification of the supply elasticity and the compliance parameter when tax policy changes alter the remittance structure by shifting the statutory burden between market sides. In conclusion, the identification of the compliance parameter λ_1 in Case A in Table 1 does *not* hinge on consistent estimation of the demand elasticity.[21] By contrast, if one seeks unbiased demand elasticities, then cross-price elasticities with respect to related goods must be negligible.[22]

3.4.3 Imperfect Competition

We extend our framework to allow for imperfect competition in the presence of tax evasion, showing that the compliance parameter is identified under imperfect competition as well. In practice, we show that the demand equation and the first-order condition of a profit-maximizing firm under various forms of conduct (represented by a conduct parameter θ) are analogous to the base model. We establish that the conduct parameter does not interact with prices, taxes, or quantities; hence, the level of conduct does not affect the estimation of compliance.

With differentiated products, a source of market power arises. When data covers all the products in the relevant market, our straightforward framework can still identify market-level compliance under imperfect competition without estimating market power or product-level elasticities; firm- or product-level compliance cannot be separately identified without additional assumptions about the demand system.

Section A.3 in Appendix A provides the details. Table A.1 provides the solutions and assumptions across all four cases.

[21] As before, the same assertion is symmetrically valid for Case D in Table 1, as well as Cases A and D when identifying post-enforcement compliance as shown in Table 2.

[22] In our Airbnb setting, Farronato and Fradkin (2022) show in Appendix Table E9 that cross-price elasticities between Airbnb and hotels are extremely low. Our tests in Section 4.4 further corroborate the result.

3.4.4 Tax Variation as a Shifter

Variation in the tax rate is a source of identification for both the supply and demand elasticities, as discussed in Zoutman et al. (2018). In conjunction with a change in tax enforcement, variation in the tax rate allows for the identification of all structural parameters: demand and supply elasticities, as well as pre- and post-enforcement compliance rates. In the context of the Airbnb case, where the statutory burden shifts from the supply to the demand side, one can simply modify the system of demand and supply equations as follows:

$$\begin{aligned} y_{it} &= \varepsilon^d p_{it} + \gamma^d D_t \cdot T_{it} + \gamma^d Z_{it} + v_{it}^d, \\ y_{it} &= \varepsilon^s p_{it} + \gamma^s D_t \cdot T_{it} + v_{it}^s, \end{aligned}$$

where Z_{it} denotes a demand-side tax acting as a demand shifter after enforcement; hence, it does not appear in the supply equation.

Under Assumption 2, whereby $\gamma^d = \lambda_2 \varepsilon^d$, and Assumption 3, whereby $\gamma^s = \lambda_1 \varepsilon^s$, we obtain:

$$\begin{aligned} \varepsilon^s &= \frac{\pi_{Zy}}{\pi_{Z_p}}, \\ \varepsilon^d &= \frac{\pi_{Ty}}{\lambda_2 + \pi_{Tp}}, \\ \lambda_1 &= -\pi_{Tp} + \frac{\pi_{Ty}}{\varepsilon_s}, \\ \lambda_2 &= \frac{\pi_{Zy}\pi_{Tp} - \pi_{Zp}\pi_{Ty}}{\pi_{Ty} - \pi_{Zy}}. \end{aligned}$$

The analogous results apply to all remittance scenarios presented in Table 1. In the context of data presenting enough variation in the tax rate together with a change in tax enforcement, all structural parameters are identified without any knowledge of the level of compliance at any point in time and under any remittance scenario.[23]

[23] Unfortunately, this identification strategy is not applicable in our empirical illustration applied to the Airbnb market as we have insufficient variation in the tax rate in the post-enforcement sample.

4 An Application to the Airbnb Market

We present an application of the identification results outlined in the conceptual framework using the collection agreements stipulated between Airbnb and state and local governments. As discussed in Section 2, these agreements achieve full enforcement by shifting the tax burden away from hosts to renters via the platform. Using the results outlined in Section 3 (specifically in Section 3.2), we estimate the level of pre-enforcement compliance in Airbnb markets and the elasticity of supply and demand.

4.1 Data

We start with information derived from Airbnb.com on short-term rental listings, including daily price, daily availability, daily bookings, and date of booking. The data is collected by AirDNA, a third-party source that frequently scrapes property, availability, and host information from the website.

Our estimation sample covers 24 major metropolitan areas across the United States and includes 241,810 Airbnb listings active between August 2014 and September 2017 in 78 tax jurisdictions. We define tax jurisdictions as unique city, county, and state combinations.[24] To alleviate concerns about potential confounders, we follow Bibler et al. (2021) by excluding jurisdictions affected by confounding regulations. The sample includes listings that are relatively close substitutes to traditional short-term rental options.[25] Finally, we aggregate our property-day data to the property-month level for our analysis.

We augment the data with information on the timing of VCAs and the tax rate enforced by each VCA, which is summarized in Table B.1. This information is constructed using information published on the Airbnb website and from secondary sources such as news and government websites. We confirm the timing and the tax rates for the entire sample. En-

[24]Our sample includes a larger number of jurisdictions with respect to Bibler et al. (2021), which rely on within metro-year-month treatment variation.

[25]We remove shared room listings, properties with more than four bedrooms, listings allowing more than twelve guests, and listings with an average price in the bottom or top ten percent of the jurisdiction-specific price distribution.

forcement through VCAs varies across jurisdictions and within jurisdiction over time. Of the 78 tax jurisdictions, 45 are treated by a VCA, while the remaining 33 jurisdictions are never treated during the sample period.[26]

Finally, we construct three demand shifters. First, we use monthly data on the number of flight passengers by airport provided by Sabre Travel Solutions; we isolate incoming trips as part of a round trip from a different city and aggregate incoming passengers at the metro level to measure potential demand for accommodation (Farronato and Fradkin, 2022).[27] Our measure of incoming passengers proxies for demand fluctuations driven by area-specific seasonality, idiosyncratic shocks, and long-term trends in demand. In addition, we include two demand shifters based on Google Trends, which provides a normalized measure of search volume for a given query (Barron et al., 2021; Farronato and Fradkin, 2022). We use two queries: "hotels 'metro' " and "Airbnb 'metro' ", and extract monthly data series for each metro between June 2014 and November 2019. Google Trends series are standardized to equal 100 in the peak month over the search period and range from 0 to 100. Importantly, both measures reflect searches from all locations worldwide.

Table 3 presents the summary statistics.[28] The booking price is tax-inclusive before the implementation of a VCA and tax-exclusive after. The average booking price is roughly \$135 per night, while the average number of nights booked per property-month is 5.75. The number represents the number of nights booked in a property-month for future stays, which can exceed 31. The average tax rate enforced through the platform in treated jurisdictions is 10.9%, with modest variation across jurisdictions. Finally, the table includes summary statistics for the three demand shifters. The first one, *Arriving Passengers* (measured in

[26]While Airbnb began implementing tax collection agreements with jurisdictions in 2014, the process remains incomplete and continues to evolve. As of 2025, several states, including Alabama and Delaware, are still adopting new agreements, while others, like Ohio, maintain traditional lodging tax systems without platform collection. Local implementation varies considerably, particularly in rural areas, mirroring the gradual adoption pattern seen in international markets. This ongoing rollout creates continued variation in tax regimes across jurisdictions.

[27]We supplement this with official passenger statistics for the San Francisco International Airport (SFO), which is missing in our data.

[28]Table B.2 includes summary statistics by treatment status.

1000s), refers to the total number of arriving passengers at the metro-month level of passengers; the sample average is over 1.1 million passengers per month. The last two rows of Table 3 include the summary of the Google Trends variables. The sample average of the hotel trend is around 75; that is, the average search activity is equal to 75% of the peak month. Similarly, the sample average of the Airbnb trend is around 52, meaning that the average search activity is 52% of the peak month.[29]

4.2 Estimation

Our primary goal is to estimate pre-enforcement tax compliance, the elasticity of supply, and the elasticity of demand. To this end, we estimate the effects of VCAs on average booking prices and nights booked per property-month. VCAs shift tax remittance from individual hosts to Airbnb, which acts as a government intermediary, implementing changes at the tax jurisdiction level on a staggered basis. This remittance shift, distinct from potentially endogenous voluntary compliance by individual hosts as modeled in Allingham and Sandmo (1972), enhances enforcement through third-party reporting and legal liability (Slemrod, 2019), achieving full compliance. While Airbnb's decision to enter into VCAs may itself reflect unobserved factors, our difference-in-differences design exploits the staggered rollout across jurisdictions, supported by event-study estimates showing no differential pre-trends. Although Airbnb tax enforcement policies vary at the tax jurisdiction level, we use the property as our cross-sectional unit to control for property-specific heterogeneity. Our application aligns with the framework outlined in Section 3 and Case A in Table 1.

We estimate the quantity effects using the Poisson quasi-maximum-likelihood estimator with the following exponential mean function:

$$E[\text{Nights Booked}_{kjmt} \mid \tau_{jmt}, Z_{mt}, \delta_k, \delta_t] = \exp\left[\pi_{Ty} \ln(1+\tau_{jmt}) + \pi_{Zy} Z_{mt} + \delta_k + \delta_t\right] \quad (7)$$

[29] Figure B.1 displays the empirical distributions of the three demand shifters.

We estimate the price effects using the following difference-in-differences specification by OLS:

$$\ln(\text{Booking Price}_{kjmt}) = \pi_{Tp} \ln(1 + \tau_{jmt}) + \pi_{Zp} Z_{mt} + \delta_k + \delta_t + u^p_{kjmt}. \tag{8}$$

The outcome in Equation (7) is Nights Booked_{kjmt}, the number of nights booked for property k in tax jurisdiction j and metro m in month-year t. We report our main results from estimating the nights booked regressions via Poisson with two-way fixed effects (TWFE), as the outcome is weakly positive with a significant fraction of property-month observations with zero bookings (Chen and Roth, 2023). That said, we later show that using the transformed outcome, $\ln(1 + \text{Nights Booked}_{kjmt})$, and estimating the analogous two-way fixed effects equation by OLS yields very similar results.[30]

The outcome in Equation (8) is $\ln(\text{Booking Price}_{kjmt})$, the logarithm of the booking price for property k in tax jurisdiction j and metro m in month-year t. In both equations, the treatment variable is $\ln(1+\tau_{jmt})$; the term τ_{jmt} is the tax rate, in percentage terms, *enforced* in jurisdiction j, metro m, at time t. The parameters of interest, π_{Ty} and π_{Tp}, represent the percent change in quantity and prices associated with a one percent increase in $(1 + \tau_{jmt})$, which approximates a one percentage point increase in the tax rate enforced through the platform.[31]

We include a demand shifter at the metro and month-year level denoted Z_{mt}. The coefficient estimates associated with the demand shifter, π_{Zy} and π_{Zp}, are critical to disentangle the elasticity of supply from the pre-enforcement compliance rate. Each specification includes property fixed effects, δ_k, to control for time-invariant property-specific characteristics (number of bedrooms, number of bathrooms, maximum number of guests, or location), and month-year fixed effects, δ_t, to control for year and location-invariant monthly variation in the short-term rental market. The extensive set of fixed effects also ensures that the tax

[30] See Appendix Table B.3 and Figure B.3.

[31] One may (correctly) note that our estimation approach for this application does not account for network externalities that can arise in two-sided markets like Airbnb, which could be an alternate explanation for the observed effects. In earlier work, using a nearly-identical empirical approach and sample, Bibler et al. (2021) find virtually no effect of VCAs on listing entry or exit.

rates are plausibly exogenous as we base our inferences on within property and year-month variations. Finally, we cluster the standard errors by tax jurisdiction.

4.3 Results

4.3.1 Reduced Form Estimates

Before reporting the main estimates, we produce two sets of event studies based on binary treatment versions of Equations (7) and (8). We address the two primary concerns with estimating two-way fixed effects (TWFE) specifications in our setting: differential pre-trends between the treated and control groups, and the staggered adoption of the tax enforcement policies.

Parallel counterfactual trends are a necessary assumption for differences-in-differences estimators to deliver causal estimates. In addition, the TWFE estimator with staggered adoption delivers consistent estimates under the assumption of homogeneity in treatment effects across groups and time (Goodman-Bacon, 2021). We estimate an event study using the robust estimator introduced by Sun and Abraham (2021) along with the TWFE estimator. These event studies test the plausibility of the parallel trends assumption in our setting and the robustness of relaxing the treatment effect homogeneity assumptions.

Figure 2 presents the event study coefficients for quantity (Panel a) and price (Panel b). In both specifications, the estimates delivered by the TWFE estimator are close to the estimates obtained using the method proposed by Sun and Abraham (2021). This shows that the results are unlikely to be driven by issues related to treatment effect heterogeneity and negative weighting that arise from using staggered treatments. In addition, pre-treatment coefficients are close to zero and exhibit little to no evidence of differential pre-trends, while post-treatment coefficients are substantially larger in magnitude than any of the pre-treatment coefficients, which is consistent with the parallel trends assumption. To probe the issue further, we also include the sensitivity analysis suggested by Rambachan and Roth (2023) for our post-treatment estimates in Appendix Figure B.2.

Table 4 reports reduced-form estimates of the effect of a tax enforced through a VCA on the number of nights booked (Panel A) and the booking price (Panel B). The first column in the table shows results for the simplest specification, which includes property fixed effects and month-year fixed effects. We estimate that a 10 percentage point increase in the enforced tax rate decreases the number of nights booked by 4.8% and reduces booking price by 2.4%. Columns 2 to 4 of the table include estimates of Equations (7) and (8) using three different variables that act as demand shifters: (i) the number of incoming flight passengers (column 2); (ii) the search volume for hotels from Google Trends in a given metro (column 3); and (iii) the search volume for Airbnb rooms from Google Trends in a given metro (column 4).

We view (ii), hotel search volume from Google Trends, as our preferred instrument. Intuitively, hotel search volumes are unlikely to be driven by variations in hotel supply, given the fixed supply of hotels in the short run. Our exclusion restriction assumes that fluctuations in accommodation demand caused by holidays or special events and captured by hotel search volumes are correlated with fluctuations in Airbnb demand (not in Airbnb supply) after conditioning on property and month-year fixed effects. Importantly, we use hotel searches, not bookings, so the correlation with Airbnb prices and quantities is unlikely to be driven by a supply response to hotel capacity constraints. Regarding (i) and (iii), we argue it is unlikely that the availability of Airbnb listings drives tourists to travel or search for accommodations in a particular area. Farronato and Fradkin (2022) advance a similar defense for using search volumes from Google Trends and flight travelers as exogenous demand shifters, arguing that Airbnb bookings make up a small share of travel demand. However, we acknowledge that it is possible that Airbnb advertisements for particular destinations or attractive listings could result in more Google searches for Airbnb as well as flights to a particular destination. In any case, the similarity in the reduced-form coefficients and implied structural parameters across the three specifications reassuringly suggest that all three shifters act primarily on the demand side.

In all specifications, the estimated effects of the enforced tax rate on both nights booked and prices are similar. We also find that, intuitively, increased demand leads to higher quantities and prices. For example, a 10% increase in arriving passengers yields a statistically significant 4.7% increase in the number of nights booked (Panel A) and 3.3% increase in booking prices (Panel B). A one-point increase in the volume of Google hotel searches leads to a 0.8% increase in nights booked and an increase in the booking prices of 0.4%.[32]

Appendix Table B.3 and Figure B.3 present a robustness check on our main quantity estimates which are based on Poisson TWFE regression (Chen and Roth, 2023). In an alternative approach, we use the logarithmic transformation $\ln(1 + \text{Nights Booked}_{kjmt})$ and estimate the equation by OLS, obtaining very similar estimates for the reduced-form coefficients and implied structural parameters.[33]

4.3.2 Structural Estimates

Table 5 includes estimates for the market-level elasticity of demand, the elasticity of supply, and pre-enforcement tax compliance. The structural parameters are constructed using the reduced-form estimates. Case A of Table 1 displays the relationships between the reduced-form and the structural estimates in our application. As we assume that VCAs lead to full compliance, $\lambda_2 = 1$.

We focus the discussion of our results on specification (ii), hotel search volume from Google Trends as a demand shifter. The market-level elasticity of demand, $\varepsilon^d = \frac{\hat{\pi}_{Ty}}{1+\hat{\pi}_{Tp}}$, is equal to $\frac{-0.480}{1-0.244} = -0.635$. Using the same demand shifter, we obtain an elasticity of supply, $\varepsilon^s = \frac{\hat{\pi}_{Zy}}{\hat{\pi}_{Zp}}$, equal to $\frac{0.008}{0.004} = 2.010$. The estimated elasticities of demand and supply are consistent across the three specifications employing different demand shifters. The market-level elasticity of demand ranges between -0.41 and -0.66. These estimates are consistent with Bibler et al. (2021); they estimate a demand elasticity of -0.48 using a smaller sample and

[32] Google Trends are standardized to the peak month over the trend period, so a one-point change in the trend reflects a one-percentage-point change in the search interest for a given metro area.

[33] Note that using the logarithmic transformation $\ln(1 + \text{Nights Booked}_{kjmt})$ yields slightly higher estimated pre-enforcement compliance rates than the Poisson approach, such that we view the former as providing slightly more conservative empirical results in terms of implied evasion.

a different set of fixed effects. The elasticity of supply ranges between 1.42 and 2.17. These estimates align with the ones obtained by Farronato and Fradkin (2018) equal to 2.16, and the lower bound estimate of 1.5 obtained by Bibler et al. (2021).

Finally, we estimate pre-enforcement compliance, $\hat{\lambda}_1 = -\hat{\pi}_{Tp} + \frac{\hat{\pi}_{Ty}}{\hat{\varepsilon}^s}$. The estimated compliance rate has two components: the total price change and the price change that could be explained by a movement along the supply curve. Using, again, Google trends for hotel searches as the demand shifter, we obtain a pre-enforcement compliance rate equal to $0.244 - \frac{0.48}{2.01} = 0.005$; only 0.5% of transactions were compliant before enforcement. The confidence intervals are tight around the obtained values of pre-enforcement compliance. We test the hypotheses that $\lambda_1 > 0.1$ and $\lambda_1 > 0.2$; the $p-$values are 0.264 and 0.097, respectively. These results suggest that we can rule out even modest compliance rates.

Using different demand shifters, we obtain very similar estimates; the pre-enforcement compliance rate, λ_1, is between zero and 1.6 percent. Similarly, using the alternative linear estimation of the quantity effects produces estimates of λ_1 between 1.1% and 3.5%. In other words, the price change can be explained almost entirely by a movement along the supply function based on the estimated elasticity of supply and the change in the number of nights booked.

Our finding of low compliance before the tax collection agreements implies that enforcement leads to a considerable increase in tax revenues for tax jurisdictions (even after accounting for demand and supply equilibrium effects). Anecdotal evidence in the form of celebratory news articles attributing increased revenue to the implementation of collection agreements corroborates our findings: Los Angeles, Texas, Arizona, Tennessee, Florida, and US and Canada.

When calculating the structural parameters, we assume the tax is fully salient to renters ($\theta = 1$). At the end of Section 3, we demonstrate that while imperfect tax salience would attenuate the estimated market elasticity of demand, it does not affect the estimated compliance rate. Hence, our conclusions related to the pre-enforcement compliance rate remain unaltered in the event that the actual tax salience is less than one.

4.4 Additional Results

Testing for Spillovers Our empirical analysis relies on comparing listings in jurisdictions with and without VCAs in place. A VCA adopted in treated jurisdictions might influence outcomes in untreated jurisdictions, which are sometimes in close proximity. We therefore include three complementary robustness checks, which show that geographic spillovers do not appear to play any role in our estimates. First, we directly estimate the effects of VCAs on control jurisdictions. Using the $n = 33$ untreated jurisdictions, we regress outcomes (nights booked and prices) on two metro–level variables: (i) a dummy equal to one if any jurisdiction in the metro has a VCA and (ii) the natural logarithm of the highest VCA tax rate in the metro. Table B.4 shows that neither variable has a statistically significant effect on prices or nights booked in the control jurisdictions, suggesting that VCAs do not transmit across borders within the same metro area.

Second, we re-estimate our baseline specification on a subsample restricted to the single largest tax jurisdiction in each metro. Because these jurisdictions are the most likely to adopt VCAs and the least likely to be influenced by smaller neighbors, this exercise minimizes potential cross-border contamination. The results, reported in Table B.5, closely mirror the baseline; the highest estimated compliance rate is 0.1%.

Finally, we remove all control jurisdictions located in metros where any VCA is in force and re-estimate our main model. By excluding observations that could conceivably be exposed to spillovers, we obtain the estimates in Table B.6. These estimates remain consistent with our primary results, and all inferred compliance parameters are smaller than their counterpart baseline values in Table 5.

Accounting for Supplier Preferences The structure of the Airbnb market may raise a further concern related to time-varying supplier preferences. Hosts may reserve their properties for personal use, so within-supplier (seasonal) changes could be correlated with demand fluctuations. To address this, we augment the baseline specification with a proxy for these unobserved preferences: the number of nights each listing is booked or available in the observation month. Re-estimating the model with this control (Table B.7) leaves the main results essentially unchanged. The highest implied compliance rate, obtained when demand is instrumented with the Airbnb-search shifter, is 2.3%.

Heterogeneous effects We illustrate the usefulness of the identification strategy to test for heterogeneity in compliance rates, focusing on differential compliance rates between casual and professional hosts, where we define professional hosts as those with five or more Airbnb listings. Professional and casual hosts may differ in compliance rates for several reasons. For example, professional hosts may be more aware of tax obligations, face lower costs of complying, and/or have different levels of risk aversion or evasion incentives.

In Appendix Tables B.8 and B.9, we present results for both professional and non-professional hosts. Appendix Table B.8 reports reduced-form effects of tax-enforcement agreements on quantities and prices, using interacted specifications to estimate differential effects for professional and non-professional hosts. From Panel A, we find generally stronger effects on the number of nights booked among casual hosts, with elasticities up to –0.63 (compared with –0.49 in the pooled sample), whereas we find statistically insignificant changes for professional hosts. From Panel B, we estimate generally stronger effects on listing prices among professional hosts, with elasticities up to –0.44 (compared with –0.25 in the pooled sample). This pattern is consistent with higher compliance rates among professional hosts, as price declines are driven by the alleviation of tax obligations among tax-compliant hosts. Similarly, quantity effects are mitigated by higher levels of pre-enforcement compliance. In the extreme case, shifting the tax burden from the supply side to the demand side with both sides fully compliant does not generate quantity effects in the baseline model.

In addition to these differential reduced-form effects, we allow for differential effects of the demand shifters (Z) on nights booked and prices across host types. We estimate broadly comparable reduced-form effects for professional and casual hosts, although the effects are slightly larger for professional hosts in both quantities and prices.

Turning to the estimated structural parameters reported in Appendix Table B.9, we find significant differences between professional and casual hosts. Most notably, professional hosts, who constitute only about 16% of the sample (according to our definition), display pre-enforcement compliance rates above zero (37–59%), indicating higher compliance prior to enforcement.[34]

5 Conclusion

In this paper, we develop a simple theoretical approach to identify tax compliance from variation in enforcement. We present a fully general framework to estimate demand and supply elasticities along with tax compliance rates in settings where variation in tax enforcement leads to potential differences in tax compliance rates. Identification of tax compliance along with demand and supply elasticities requires incorporating an additional variable that acts as a supply or demand shifter, depending on which side of the market bears the statutory burden before and after the change in the tax policy.

Our approach is especially appealing to investigate tax compliance in online transactions, where tax obligations are particularly ambiguous or difficult to enforce. We illustrate the theoretical identification argument using Airbnb tax enforcement agreements with local jurisdictions. Exploiting the staggered introduction of these agreements, we use a difference-in-difference design to estimate the level of pre-enforcement compliance. We find that only zero to 4.6% of transactions were compliant before enforcement.

Our approach generalizes to settings in which policy changes elicit partial behavioral responses. First, consider information shocks like the drop in the 1099-K reporting thresholds

[34] Some caution in interpreting these results is warranted because, for professional hosts, strategic pricing behavior and/or the availability of close demand-side substitutes could affect the calculation of the structural parameters, as discussed in Sections 3.4.2 and 3.4.3.

studied by Garin et al. (2025) or digital monitoring systems, as in Boyer and d'Astous (2023). Second, the framework can be readily extended to recover the degree of salience (rather than compliance) when consumers face prices that either include or exclude fees, as in Blake et al. (2021); in that variant, the key parameter measures the share of consumers who perceive the full price, not the share who comply. Third, digital goods and services (e.g., software, streaming, non-fungible tokens) pose distinctive compliance challenges because their intangibility and the absence of a physical buyer address often place them outside Nexus/Wayfair sales tax requirements; substantial cross-state heterogeneity in taxability and definitions further creates scope for non-compliance. Platform-level rule changes or future policy reforms in these markets can generate quasi-experimental variation analogous to Airbnb's VCAs. For example, if a platform begins enforcing tax collection on digital goods, our framework can be employed to identify compliance effects. Finally, digitalization reforms aimed at curbing cross-border tax fraud, such as mandatory electronic payments for customs duties studied by Kitsios et al. (2020), raise enforcement intensity by improving transparency and limiting evasion opportunities. In our taxonomy, these map to enforcement-intensity cases (e.g., Cases B/C). Using trade data (import volumes and prices) together with exogenous shifters (trade seasonality or staggered port-level implementation timing), the framework recovers pre- and post-reform compliance and elasticities. Relatedly, variation in statutory tariff rates also identifies compliance; for example, exploiting tariff hikes during the recent trade wars would allow estimation of compliance among trading partners ex-ante and ex-post tariff variations.

References

Agrawal, David R. and Iuliia Shybalkina, "Online shopping can redistribute local tax revenue from urban to rural America," *Journal of Public Economics*, March 2023, *219*, 104818.

_ **and Mohammed Mardan**, "Will destination-based taxes be fully exploited when available? An application to the U.S. commodity tax system," *Journal of Public Economics*, January 2019, *169*, 128–143.

Allingham, Michael G. and Agnar Sandmo, "Income Tax Evasion: A Theoretical Analysis," *Journal of Public Economics*, 1972, *1* (3-4), 323–338.

Almunia, Miguel and David Lopez-Rodriguez, "Under the Radar: The Effects of Monitoring Firms on Tax Compliance," *American Economic Journal: Economic Policy*, February 2018, *10* (1), 1–38.

Bachas, Pierre and Mauricio Soto, "Corporate Taxation under Weak Enforcement," *American Economic Journal: Economic Policy*, November 2021, *13* (4), 36–71.

Barron, Kyle, Edward Kung, and David Proserpio, "The effect of home-sharing on house prices and rents: Evidence from Airbnb," *Marketing Science*, 2021, *40* (1), 23–47.

Baugh, Brian, Itzhak Ben-David, and Hoonsuk Park, "Can taxes shape an industry? Evidence from the implementation of the "Amazon tax"," *The Journal of Finance*, 2018, *73* (4), 1819–1855.

Bekkerman, Ron, Maxime C Cohen, Edward Kung, John Maiden, and Davide Proserpio, "The effect of short-term rentals on residential investment," *Marketing Science*, 2023, *42* (4), 819–834.

Bibler, Andrew J, Keith F Teltser, and Mark J Tremblay, "Inferring tax compliance from pass-through: Evidence from Airbnb tax enforcement agreements," *The Review of Economics and Statistics*, 2021, *103* (4), 636–651.

Blake, Tom, Sarah Moshary, Kane Sweeney, and Steve Tadelis, "Price salience and product choice," *Marketing Science*, 2021, *40* (4), 619–636.

Boyer, M. Martin and Philippe d'Astous, "Tax compliance and firm response to electronic sales monitoring," *Canadian Journal of Economics/Revue canadienne d'économique*, October 2023, *56* (4), 1430–1468.

Bresnahan, Timothy F., "The oligopoly solution concept is identified," *Economics Letters*, January 1982, *10* (1–2), 87–92.

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal, "Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador," *American Economic Review: Insights*, December 2023, *5* (4), 427–44.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, December 2023, *139* (2), 891–936.

Chen, Wei, Zaiyan Wei, and Karen Xie, "Regulating Professional Players in Peer-to-Peer Markets: Evidence from Airbnb," *Management Science*, May 2023, *69* (5), 2893–2918.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, September 2009, *99* (4), 1145–77.

Dearing, Adam, "Estimating structural demand and supply models using tax rates as instruments," *Journal of Public Economics*, 2022, *205*, 104561.

Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan, "Sales taxes and internet commerce," *American Economic Review*, 2014, *104* (1), 1–26.

Farronato, Chiara and Andrey Fradkin, "The Welfare Effects of Peer Entry in the Accommodation Market: The Case of Airbnb," *NBER Working Paper*, 2018.

_ **and** _ , "The welfare effects of peer entry: The case of Airbnb and the accommodation industry," *American Economic Review*, 2022, *112* (6), 1782–1817.

Feldman, Naomi E. and Joel Slemrod, "Estimating Tax Noncompliance with Evidence from Unaudited Tax Returns," *The Economic Journal*, March 2007, *117* (518), 327–352.

Fisman, Raymond and Shang-Jin Wei, "Tax rates and tax evasion: evidence from "missing imports" in China," *Journal of Political Economy*, 2004, *112* (2), 471–496.

Fox, William F., Enda Patrick Hargaden, and LeAnn Luna, "Statutory incidence and sales tax compliance: Evidence from Wayfair," *Journal of Public Economics*, September 2022, *213*, 104716.

Garin, Andrew, Emilie Jackson, Dmitri Koustas, and Alicia Miller, "The Impact of Third-Party Reporting on Tax Compliance: Evidence from Gig Workers," Technical Report 2025.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, December 2021, *225* (2), 254–277.

Jia, Jian and Liad Wagman, "Platform, Anonymity, and Illegal Actors: Evidence of Whac-a-Mole Enforcement from Airbnb," *The Journal of Law and Economics*, November 2020, *63* (4), 729–761.

Jin, Ginger Zhe, Liad Wagman, and Mengyi Zhong, "The Effects of Short-term Rental Regulation: Insights from Chicago," Technical Report 2023.

Kitsios, Emmanouil, João Tovar Jalles, and Genevieve Verdier, "Tax Evasion from Cross-Border Fraud: Does Digitalization Make a Difference?," *International Monetary Fund*, 2020, *245*.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod, "Does tax-collection invariance hold? Evasion and the pass-through of state diesel taxes," *American Economic Journal: Economic Policy*, 2016, *8* (2), 251–86.

Manzi, Nina, "Use tax collection on income tax returns in other states," Technical Report, Policy Brief, Research Department, Minnesota House of Representatives 2012.

Pissarides, Christopher A. and Guglielmo Weber, "An expenditure-based estimate of Britain's black economy," *Journal of Public Economics*, June 1989, *39* (1), 17–32.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, February 2023, *90* (5), 2555–2591.

Ramsey, Frank P, "A Contribution to the Theory of Taxation," *The Economic Journal*, 1927, *37* (145), 47–61.

Slemrod, Joel, "Tax Compliance and Enforcement," *Journal of Economic Literature*, December 2019, *57* (4), 904–954.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199. Themed Issue: Treatment Effect 1.

Waseem, Mazhar, "Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax," *Journal of Public Economics*, 2023, *218*, 104783.

Zoutman, Floris T., Evelina Gavrilova, and Arnt O. Hopland, "Estimating Both Supply and Demand Elasticities Using Variation in a Single Tax Rate," *Econometrica*, 2018, *86* (2), 763–771.

Figures and Tables

Figure 1: Examples of Tax Enforcement

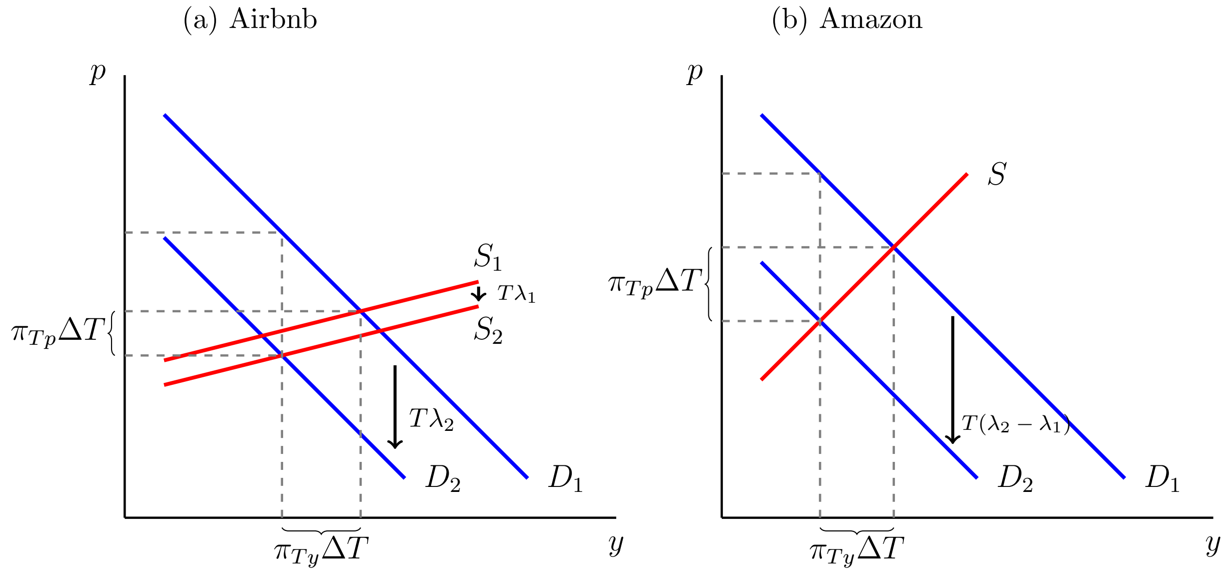

Figure 1 presents two cases of tax enforcement in the presence of evasion: home-sharing (Airbnb) and online retail (Amazon). Figure 1a represents the Airbnb case, in which tax enforcement changes the statutory burden from the supply side, where the fraction of tax-compliant transactions before enforcement is denoted by λ_1, to the demand side, where the fraction of tax-compliant transactions is denoted by λ_2. Figure 1b represents the Amazon case, where the statutory burden initially falls on consumers, and λ_1 is the share of transactions for which taxes are paid. After Amazon enforces sales taxes at checkout, the burden remains on the demand side where all consumers ($\lambda_2 = 1$) now pay the tax.

Figure 2: Event study estimators

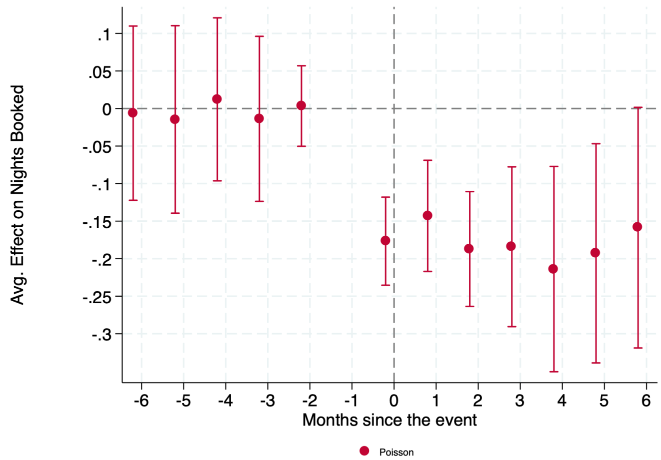

(a) Effect of VCAs on nights booked, Poisson

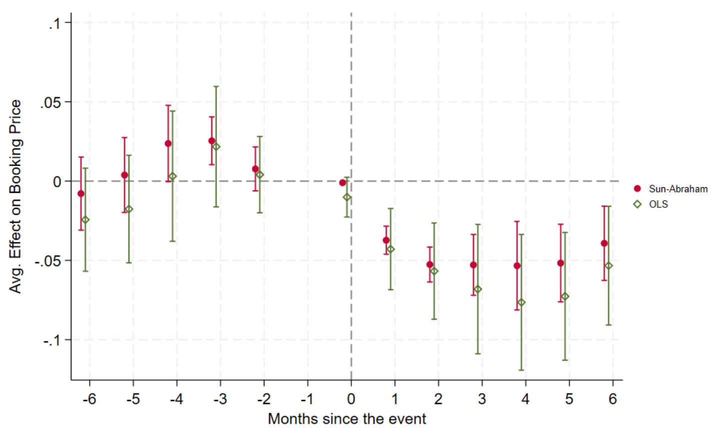

(b) Effect of VCAs on booking prices

The figures report: in Panel (a) a dynamic version of the TWFE model, Equation (7), estimated using Poisson regression. The outcome is Nights Booked_{kjmt}, the number of nights booked for property k in tax jurisdiction j and metro m in month-year t; in Panel (b) a dynamic version of the TWFE model, Equation (8), estimated using OLS and Sun and Abraham (2021). The outcome is ln(Booking Price_{kjmt}), the logarithm of the booking price for property k in tax jurisdiction j and metro m in month-year t. The figures display six pre-periods and six post-periods. The bars represent 95 percent confidence intervals. Standard errors are clustered at the tax-jurisdiction level.

Table 1: Summary of Results: Identifying Pre-Enforcement Compliance

Examples	Burden Pre	Burden Post	Assumptions	Results
Case A:	Supply	Demand	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{\lambda_2 + \pi_{Tp}}$
Airbnb			$\gamma^d = \lambda_2 \varepsilon^d$	$\varepsilon^s = \frac{\pi_{Zy}}{\pi_{Zp}}$
Bibler et al. (2021)			$\gamma^s = \lambda_1 \varepsilon^s$	$\lambda_1 = \frac{\pi_{Ty}}{\varepsilon^s} - \pi_{Tp}$
Case B:	Demand	Demand	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Amazon			$\gamma^d = (\lambda_2 - \lambda_1)\varepsilon^d$	$\varepsilon^s = \frac{\pi_{Ty}}{\pi_{Tp}}$
Baugh et al. (2018)			$\gamma^s = 0$	$\lambda_1 = \lambda_2 + \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^d}$
Case C:	Supply	Supply	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{\pi_{Tp}}$
Trade Tariffs			$\gamma^d = 0$	$\varepsilon^s = \frac{\pi_{Zy}}{\pi_{Zp}}$
Fisman and Wei (2004)			$\gamma^s = (\lambda_1 - \lambda_2)\varepsilon^s$	$\lambda_1 = \lambda_2 - \pi_{Tp} + \frac{\pi_{Ty}}{\varepsilon^s}$
Case D:	Demand	Supply	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Diesel Fuel			$\gamma^d = -\lambda_1 \varepsilon^d$	$\varepsilon^s = \frac{\pi_{Ty}}{\pi_{Tp} - \lambda_2}$
Kopczuk et al. (2016)			$\gamma^s = -\lambda_2 \varepsilon^s$	$\lambda_1 = \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^d}$

The table outlines the necessary assumptions and identification results for four possible combinations of shifts in the statutory burden determined by an enforcement policy: (A) enforcement shifts the statutory burden from the supply to the demand side; (B) the statutory burden remains on the demand side before and after the change in enforcement; (C) the statutory burden remains on the supply side before and after the change in enforcement; and (D) enforcement shifts the statutory burden from the demand to the supply side. We provide an in-depth discussion of each identification assumption and the results for four cases in Appendix A. *Burden Pre* and *Burden Post* refer to the side of the market that bears the statutory burden before and after the change in enforcement, respectively. The *Assumptions* column specifies the necessary assumptions, and the *Results* column includes the solutions for the structural parameters in terms of the reduced-form parameters. It is important to note the distinction between tax-inclusive versus tax-exclusive prices across the four cases. The price in the burden pre-stage is the market price observed in the data. In the burden post-enforcement, the price in Cases A and B is tax-exclusive (since demand shifts downward), and the price in Cases C and D is tax-inclusive (since supply shifts upward).

Table 2: Summary of Results: Identifying Post-Enforcement Compliance

Examples	Burden Pre	Burden Post	Assumptions	Results
Case A:	Supply	Demand	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Airbnb			$\gamma^d = \lambda_2 \varepsilon^d$	$\varepsilon^s = \frac{\pi_{Ty}}{\lambda_1 + \pi_{Tp}}$
Bibler et al. (2021)			$\gamma^s = \lambda_1 \varepsilon^s$	$\lambda_2 = \frac{\pi_{Ty}}{\varepsilon^d} - \pi_{Tp}$
Case B:	Demand	Demand	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Amazon			$\gamma^d = (\lambda_2 - \lambda_1)\varepsilon^d$	$\varepsilon^s = \frac{\pi_{Ty}}{\pi_{Tp}}$
Baugh et al. (2018)			$\gamma^s = 0$	$\lambda_2 = \lambda_1 - \pi_{Tp} + \frac{\pi_{Ty}}{\varepsilon^d}$
Case C:	Supply	Supply	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{\pi_{Tp}}$
Trade Tariffs			$\gamma^d = 0$	$\varepsilon^s = \frac{\pi_{Zy}}{\pi_{Zp}}$
Fisman and Wei (2004)			$\gamma^s = (\lambda_1 - \lambda_2)\varepsilon^s$	$\lambda_2 = \lambda_1 + \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^s}$
Case D:	Demand	Supply	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{-\lambda_1 + \pi_{Tp}}$
Diesel Fuel			$\gamma^d = -\lambda_1 \varepsilon^d$	$\varepsilon^s = \frac{\pi_{Zy}}{\pi_{Zp}}$
Kopczuk et al. (2016)			$\gamma^s = -\lambda_2 \varepsilon^s$	$\lambda_2 = \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^s}$

The table outlines the necessary assumptions and identification results for estimation of post-enforcement compliance (λ_2), including all four possible combinations of shifts in the statutory burden determined by an enforcement policy are included. *Burden Pre* and *Burden Post* refer to the side of the market that bears the statutory burden before and after the change in enforcement, respectively. The *Assumptions* column specifies the necessary assumptions, and the *Results* column includes the solutions for the structural parameters in terms of the reduced-form parameters. It is important to note the distinction between tax-inclusive versus tax-exclusive prices across the four cases. The price in the burden pre-stage is the market price observed in the data. In the burden post-enforcement, the price in Cases A and B is tax-exclusive (since demand shifts downward), and the price in Cases C and D is tax-inclusive (since supply shifts upward).

Table 3: Summary Statistics

	N	Mean	Std. Dev.	25th Percentile	Median	75th Percentile
Book Price	1,259,409	135.12	84.26	81.7	115	165
Nights Booked	3,592,522	5.75	11.66	0	0	7
Tax Rate	3,592,522	5.53	5.93	0	5	10.5
Tax Rate, with VCA	1,823,992	10.9	3.28	7.5	10.5	14
Arriving Passengers (1000s)	3,592,522	1152.49	718.99	565.24	955.19	1756.63
Hotel Search	3,592,522	74.8	14.03	65	75	86
Airbnb Search	3,592,522	52.47	19.13	39	51	66

The table reports summary statistics of the main variables. *Arriving Passengers* (in 1000s) refers to the number of passengers arriving in a metro area in a given month, excluding return flights. *Hotel Search* refers to the Google Trends search volume for the search *hotels 'metro'* in the month. and *Airbnb Search* refers to the Google Trends search volume for the search *Airbnb 'metro'* in the month. Google Trends series are standardized to the maximum search activity over the period June 2014 - November 2019.

Table 4: Reduced form Estimates, Poisson

			Google Searches	
			Hotels	Airbnb
Panel A: Nights Booked, Poisson TWFE				
$\ln(1+\tau_{jmt})$	-0.482*	-0.495*	-0.480*	-0.339
	(0.286)	(0.256)	(0.245)	(0.231)
ln(Arrivals)		0.468***		
		(0.065)		
Google Trends			0.008***	0.011***
			(0.001)	(0.001)
Observations	3,118,578	3,118,578	3,118,578	3,118,578
Panel B: ln(Nightly Booking Price)				
$\ln(1+\tau_{jmt})$	-0.237**	-0.252***	-0.244***	-0.172***
	(0.099)	(0.079)	(0.080)	(0.059)
ln(Arrivals)		0.330***		
		(0.045)		
Google Trends			0.004***	0.005***
			(0.001)	(0.001)
Observations	1,259,409	1,259,409	1,259,409	1,259,409
Property FE	x	x	x	x
Month-Year FE	x	x	x	x

Panel A reports the reduced-form estimates of the effect of tax collection agreement on nights booked using Poisson regression. Panel B reports the reduced-form estimates on booking price. The top row of each panel $\ln(1+\tau_{jmt})$ includes the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. The number of jurisdictions is 78. Standard errors, in parentheses, are clustered at the tax jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Structural Parameter Estimates, Poisson

	Demand Shifter		
	Passengers	Hotels Trend	Airbnb Trend
ε^d	-0.662	-0.635	-0.409
	(0.763)	(0.758)	(0.489)
ε^s	1.419	2.010	2.172
	(0.228)	(0.248)	(0.261)
λ_1	-0.098	0.005	0.016
	(0.203)	(0.150)	(0.125)
p-value, $H_0 : \lambda_1 > 0.1$	0.165	0.264	0.250
p-value, $H_0 : \lambda_1 > 0.2$	0.071	0.097	0.070

The table reports the structural parameters with standard errors (in parentheses below the estimates). Standard errors are computed using a bootstrap with 500 repetitions and random sampling at the tax-jurisdiction level. For each bootstrap repetition, we construct the structural parameters and report the standard deviation of the bootstrap distribution. The first column includes estimates using the incoming flight passengers variable. Columns 2 and 3 include estimates using the volume of searches reported in Google Trends for hotels and Airbnb. The $p-$values are calculated on the basis of the parameter estimates and their standard errors assuming normality.

Appendix A

A.1 Alternative Specifications: the Tax Burden

In Table 1, we outline the necessary assumptions and identification results in the four possible cases of increased tax enforcement, which depend on the side of the market bearing the statutory tax burden in the pre- and post-enforcement periods. Each of the results follows from using the solution concept outlined in Section 3.

For Case B, in which the tax burden is on the demand side of the market before and after enforcement, we require that Z_{it} acts as a supply shifter, so $\rho^d = 0$. We refer to this assumption as SER3. Without a supply shifter, the other two assumptions we make in this case will not separately identify pre-enforcement compliance from the elasticity of demand. The modified RER assumption in this case is $\gamma^d = (\lambda_2 - \lambda_1) \cdot \varepsilon^d$, which is adjusted for the magnitude of the demand shift due to the increase in tax enforcement. Because the burden is on the demand side in both periods, the magnitude of the enforcement-induced shift is mitigated to the extent that buyers are tax-compliant in the pre-enforcement period. Lastly, in this case the statutory burden falls on consumers in both periods, so we make the SER assumption that $\gamma^s = 0$. Intuitively, the change in tax enforcement does not lead to a shift in the supply function, as in the case presented by Zoutman et al. (2018).

In Case C, the tax burden is on the supply side of the market both before and after the change in tax enforcement. For this case, we require that Z_{it} acts as a demand shifter, so $\rho^s = 0$ (as in Case A), so that assumption SER2 holds. Otherwise, the elasticity of supply and pre-enforcement compliance can not be separately identified, as enforcement leads to a shift in supply that depends on λ_1 and ε^s. The modified RER assumption in this case is $\gamma^s = (\lambda_2 - \lambda_1) \cdot \varepsilon^s$, which is adjusted for the magnitude of the supply shift due to the increase in tax enforcement. Because the burden is on the supply side in both periods, the magnitude of the enforcement-induced shift is mitigated to the extent that sellers are tax-compliant in the pre-enforcement period. Analogous to Case B, where the burden does not change sides, the tax is always levied on supply, so we apply the standard SER assumption for a supply-side tax that $\gamma^d = 0$. In other words, tax enforcement does not lead to a shift in the demand function.

Case D refers to the situation in which the tax burden switches from the demand to the supply side with enforcement, which requires that Z_{it} acts as a supply shifter. That is, similar to Case B, we assume that $\rho^d = 0$ (SER3), to facilitate separate identification of the elasticity of demand and pre-enforcement compliance. The modified RER assumption in this case is that $\gamma^s = -\lambda_2 \varepsilon^s$, which describes the magnitude of the supply shift due to the increase in tax enforcement. Given that the supply side does not bear the pre-enforcement statutory burden, the enforcement-induced supply shift is scaled by the rate of post-enforcement compliance (λ_2) and does not depend on λ_1. Lastly, because the tax burden shifts from the demand to the supply side, the magnitude of the demand shift due to the tax does depend on pre-enforcement compliance from buyers and is given by $\gamma^d = -\lambda_1 \varepsilon^d$. This is analogous to the third assumption discussed in Section 3 for Case A and captures the shift in demand that follows from removing the statutory burden from the demand side.

A.2 A Parsimonious Formulation of the Model

We rewrite the four cases separately discussed in Table 1 in a parsimonious model. We start from the system of demand and supply represented in Equation (3):

$$\begin{aligned} y_{it} &= \varepsilon^d p_{it} + \gamma^d D_t \cdot T_{it} + \rho^d Z_{it} + v_{it}^d, \\ y_{it} &= \varepsilon^s p_{it} + \gamma^s D_t \cdot T_{it} + \rho^s Z_{it} + v_{it}^s, \end{aligned}$$

Let the indicator variable D_{pre} be equal to 1 ($D_{post} = 1$) if the supply side bears the statutory burden pre-(post-) enforcement. Analogously, $D_{pre} = 0$ ($D_{post} = 0$) corresponds to the demand bearing the statutory burden of the tax pre- (post-) enforcement. This allows us to reduce Assumptions 1, 2, and 3 to two main assumptions encompassing all four possible combinations of shifts in the statutory burden determined by an enforcement policy presented in Table 1. First, Assumption 1 can be rewritten as follows:

- $\rho^s = 0$ if $D_{pre} = 1$,
- $\rho^d = 0$ if $D_{pre} = 0$.

Second, Assumptions 2 and 3 can be rewritten as:

- $\gamma^d = (1 - D_{post})\lambda_2 \varepsilon^d - (1 - D_{pre})\lambda_1 \varepsilon^d$,
- $\gamma^s = D_{pre}\lambda_1 \varepsilon^s - D_{post}\lambda_2 \varepsilon^s$.

If the supply bears the statutory burden pre-enforcement ($D_{pre} = 1$), we obtain:

$$\begin{aligned} \rho^s &= 0, \\ \varepsilon^s &= \frac{\pi_{Zy}}{\pi_{Zp}}, \\ \varepsilon^d &= \frac{\pi_{Ty}}{(1 - D_{post})\lambda_2 + \pi_{Tp}}, \\ \lambda_1 &= D_{post}\lambda_2 + \frac{\pi_{Ty}}{\varepsilon^s} - \pi_{Tp}. \end{aligned}$$

If the demand bears the statutory burden pre-enforcement ($D_{pre} = 0$), we obtain:

$$\begin{aligned}
\rho^d &= 0, \\
\varepsilon^d &= \frac{\pi_{Zy}}{\pi_{Zp}}, \\
\varepsilon^s &= \frac{\pi_{Ty}}{\pi_{Tp} - D_{post}\lambda_2}, \\
\lambda_1 &= (1 - D_{post})\lambda_2 - \frac{\pi_{Ty}}{\varepsilon^d} + \pi_{Tp}.
\end{aligned}$$

A.3 Tax Compliance with Imperfect Competition

We present a model of tax compliance under imperfect competition, showing that the compliance parameter is also identified in this case. To maintain generality, we denote λ^d as the compliance rate on the demand side and λ^s as the compliance rate on the supply side. We rewrite the demand equation as follows:

$$\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \lambda^d \varepsilon^d T_{it} + \rho^d Z_{it} + v_{it}^d, \\
\ln(Q_{it}) &= \varepsilon^d \ln P_{it} + \lambda^d \varepsilon^d \ln(1 + \tau_{it}^d) + \rho^d Z_{it} + v_{it}^d, \\
Q_{it} &= ((1 + \tau_{it}^d)^{\lambda^d} P_{it})^{\varepsilon^d} \cdot e^{(\rho^d Z_{it} + v_{it}^d)}. \qquad \text{(A.1)}
\end{aligned}$$

To simplify notation, we suppress the good and market subscripts (it). We generalize the Cournot first-order condition with N symmetric firms to allow for various forms of conduct. Following Bresnahan (1982) "conduct parameter" approach, let $\theta = 1/N$ denote a conduct parameter nesting different forms of competition: perfect competition ($\theta = 0$), full collusion ($\theta = 1$), and Cournot ($\theta = 1/N$). We write the inverse demand curve as follows:

$$P(Q) = \frac{\left(e^{-(\rho^d Z + v^d)} \cdot Q\right)^{\frac{1}{\varepsilon^d}}}{(1 + \tau^d)^{\lambda^d}}.$$

A firm's profit is given by:

$$\Pi = (1 - \tau^s)^{\lambda^s} P(Q) q - c q^{\phi},$$

where $\phi > 1$ ensures profit maximization. Note that we distinguish between τ^d and τ^s as the interaction between enforcement and statutory burden will impact these terms. In Case A, $\tau^s = \tau$ and $\tau^d = 0$ prior to enforcement, and $\tau^s = 0$ and $\tau^d = \tau$ after enforcement. In Case B, $\tau^s = 0$ and $\tau^d = \tau$ before and after enforcement. In Case C, $\tau^s = \tau$ and $\tau^d = 0$ in both the pre- and post-enforcement market. Lastly, in Case D, $\tau^s = 0$ and $\tau^d = \tau$ prior to enforcement, and $\tau^s = \tau$ and $\tau^d = 0$ after enforcement.

The generalized first-order condition is therefore:

$$
\begin{aligned}
(1-\tau^s)^{\lambda^s} P - c\phi q^{\phi-1} + (1-\tau^s)^{\lambda^s} \frac{P}{Q\varepsilon^d} q &= 0, \\
(1-\tau^s)^{\lambda^s} P - c\phi(\theta Q)^{\phi-1} + (1-\tau^s)^{\lambda^s} \frac{P\theta}{\varepsilon^d} &= 0, \\
Q = \frac{1}{\theta}\left[\frac{1}{c\phi}(1-\tau^s)^{\lambda^s}\left(1+\frac{\theta}{\varepsilon^d}\right)P\right]^{\frac{1}{\phi-1}}&, \qquad \text{(A.2)}
\end{aligned}
$$

where the second equality uses the fact that $q = \theta Q$.

Writing Equation (A.1), the demand equation, and Equation (A.2), the first-order condition, in logarithmic form and reintroducing the good and market subscripts, we obtain:

$$
\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \varepsilon^d \ln\left[(1+\tau^d_{it})^{\lambda^d}\right] + \rho^d Z_{it} + v^d_{it}, \\
y_{it} &= \frac{1}{\phi-1} p_{it} + \frac{1}{\phi-1}\ln\left[(1-\tau^s_{it})^{\lambda^s}\right] - \frac{1}{\phi-1}\ln(c_{it}) + \left[\frac{1}{\phi-1}\ln\left(\frac{1}{\phi}\left(1+\frac{\theta}{\varepsilon^d}\right)\right) - \ln\theta\right],
\end{aligned}
$$

where $\ln(c_{it})$ can be written as a function of Z_{it} and an additive error term: $\ln(c_{it}) = f(Z_{it}, \Omega) + \mu_{it}$. Noting that the term $\left[\frac{1}{\phi-1}\ln\left(\frac{1}{\phi}\left(1+\frac{\theta}{\varepsilon^d}\right)\right) - \ln\theta\right]$ is constant, the above equations reduce to:

$$
\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \varepsilon^d \ln\left[(1+\tau^d_{it})^{\lambda^d}\right] + \rho^d Z_{it} + v^d_{it}, \\
y_{it} &= \frac{1}{\phi-1} p_{it} + \frac{1}{\phi-1}\ln\left[(1-\tau^s_{it})^{\lambda^s}\right] + \frac{1}{\phi-1} f(Z_{it}, \Omega) + v^s_{it},
\end{aligned}
$$

where v^s_{it} represents a composite error term.

In Case A, enforcement eliminates the tax responsibility for the sellers and places it on buyers so that $\lambda^s = -\lambda_1$ and $\lambda^d = \lambda_2$. In addition, we use the following approximation: $\ln\left[(1-\tau)^{\lambda}\right] \approx \ln\left[(1+\tau)^{-\lambda}\right]$ for $\lambda \in [0,1]$ and τ close to zero.[35] Applying this approximation to the system above implies the following with respect to enforcement:

$$
\begin{aligned}
y_{it} &= \varepsilon^d p_{it} + \lambda_2 \varepsilon^d T_{it} + \rho^d Z_{it} + v^d_{it}, \\
y_{it} &= \frac{1}{\phi-1} p_{it} + \lambda_1 \frac{1}{\phi-1} T_{it} + \frac{1}{\phi-1} f(Z_{it}, \Omega) + v^s_{it},
\end{aligned}
$$

which is identical to the base model with $\frac{1}{\phi-1}$ in place of ε^s.

Importantly, the conduct parameter, θ, does not interact with prices, taxes, or quantities;

[35]Using this approximation is only required in Cases A and D when the statutory burden of the tax changes with enforcement.

as a consequence, the level of conduct does not affect the level of compliance represented by the parameters λ^d and λ^s. The solutions and assumptions across all four cases are presented in Table A.1. The structural parameters are not functions of the conduct parameter θ in any of the four cases, which confirms the separation between conduct and compliance estimation.

Table A.1: Summary of Results: Identifying Pre-Enforcement Compliance with $0 < \theta < 1$

Examples	Burden Pre	Burden Post	Assumptions	Results
Case A:	Supply	Demand	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{\lambda_2 + \pi_{Tp}}$
Airbnb			$\lambda^d = \lambda_2$	$\frac{1}{1-\phi} = \frac{\pi_{Zy}}{\pi_{Zp}}$
Bibler et al. (2021)			$\lambda^s = -\lambda_1$	$\lambda_1 = (1-\phi)\pi_{Ty} - \pi_{Tp}$
Case B:	Demand	Demand	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Amazon			$\lambda^d = \lambda_2 - \lambda_1$	$\frac{1}{1-\phi} = \frac{\pi_{Ty}}{\pi_{Tp}}$
Baugh et al. (2018)			$\lambda^s = 0$	$\lambda_1 = \lambda_2 + \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^d}$
Case C:	Supply	Supply	$\rho^s = 0$	$\varepsilon^d = \frac{\pi_{Ty}}{\pi_{Tp}}$
Trade Tariffs			$\lambda^d = 0$	$\frac{1}{1-\phi} = \frac{\pi_{Zy}}{\pi_{Zp}}$
Fisman and Wei (2004)			$\lambda^s = \lambda_2 - \lambda_1$	$\lambda_1 = \lambda_2 - \pi_{Tp} + (1-\phi)\pi_{Ty}$
Case D:	Demand	Supply	$\rho^d = 0$	$\varepsilon^d = \frac{\pi_{Zy}}{\pi_{Zp}}$
Diesel Fuel			$\lambda^d = -\lambda_1$	$\frac{1}{1-\phi} = \frac{\pi_{Ty}}{\pi_{Tp} - \lambda_2}$
Kopczuk et al. (2016)			$\lambda^s = \lambda_2$	$\lambda_1 = \pi_{Tp} - \frac{\pi_{Ty}}{\varepsilon^d}$

This table extends Table 1 by outlining the necessary assumptions and identification results for the four possible cases in the model with imperfect competition. The ρ terms capture which shifter is used, and $\phi > 1$ captures the underlying cost structure of the firms.

Appendix B: Supplemental Tables and Figures

B.1 Additional Descriptive Statistics

Table B.1: Summary of Tax Introductions

Tax Intro	City	Metro	Initial Tax	Max. Tax
Feb. 2015	Washington	Washington DC	7.25	14.5
Apr. 2015	Malibu	Los Angeles	4.4	12
Jun. 2015	Charlotte	Charlotte	15.25	15.25
Jul 2015	Oakland	Oakland	14	14
	Phoenix	Phoenix	5.3	12.57
	San Diego	San Diego	5.76	10.5
Oct. 2015	Bellevue	Seattle	6.58	12.4
	Kirkland	Seattle	5.76	11
	Redmond	Seattle	5.76	11
	Santa Clara	San Jose	5.21	9.5
	Seattle	Seattle	5.26	10.5
	University Place	Seattle	6.25	12.1
	Vashon	Seattle	4.72	8.6
Nov. 2015	Jersey City	New York	6	6
	Delray Beach	Miami	6	7
	Four Corners	Orlando	7	7.5
	Four Corners	Orlando	7	7
	Kissimmee	Orlando	7	7.5
	Orlando	Orlando	6.5	12.5
	Sunny Isles Beach	Miami	7	13
	West Palm Beach	Miami	6	7
Jan. 2016	Evanston	Chicago	3.38	7.17
	Oak Park	Chicago	3.38	11.17
Apr. 2016	Cleveland Heights	Cleveland	5.5	5.5
	Lakewood	Cleveland	5.5	5.5
	Metairie	New Orleans	5	5
	New Orleans	New Orleans	5	9
Jun. 2016	Bethesda	Washington DC	7	7
	Silver Spring	Washington DC	7	7
Aug. 2016	Anchorage	Anchorage	12	12
	Los Angeles	Los Angeles	14	14
Sep. 2016	Golden	Denver	3	8.43
	Millcreek	Salt Lake City	11.6	11.6
	Salt Lake City	Salt Lake City	12.6	12.6
	Sandy	Salt Lake City	13.1	13.1
Jan. 2017	Mesa	Phoenix	14.02	14.02
	Scottsdale	Phoenix	13.92	13.92
	Tempe	Phoenix	14.07	14.07
Feb. 2017	Lakewood	Denver	5.43	5.43
May 2017	Austin	Austin	6	6
	Dallas	Dallas	6	6
	Fort Worth	Dallas	6	6
	Galveston	Houston	6	6
	Houston	Houston	6	6
Jun. 2017	Richmond	Oakland	10	10

Table B.2: Summary Statistics by Treatment Status

Treated	N	Mean	Std. Dev.	25th Percentile	Median	75th Percentile
Book Price	1,002,668	137	87	84	116	165
Nights Booked	2,878,807	6	12	0	0	6
Tax Rate	2,878,807	7	6	0	7	14
Tax Rate, with VCA	1,823,992	11	3	8	11	14
Arriving Passengers (1000s)	2,878,807	1156	726	550	973	1778
Hotel Search	2,878,807	75	14	64	75	86
Airbnb Search	2,878,807	52	19	37	50	65

Untreated	N	Mean	Std. Dev.	25th Percentile	Median	75th Percentile
Book Price	256,741	128	72	76	107	160
Nights Booked	713,715	6	12	0	0	7
Tax Rate	N/A					
Tax Rate, with VCA	N/A					
Arriving Passengers (1000s)	713,715	1139	688	680	940	1540
Hotel Search	713,715	75	12	68	75	85
Airbnb Search	713,715	56	20	42	56	74

The table reports summary statistics of the main variables by treatment status. The top panel inlcudes observations for treated jurisdictions. The lower panel includes observations for never treated jurisdictions. *Arriving Passengers* (in 1000s) refers to the number of passengers arriving in a metro area in a given month, excluding return flights. *Hotel Search* refers to the Google Trends search volume for the search *hotels 'metro'* in the month. and *Airbnb Search* refers to the Google Trends search volume for the search *Airbnb 'metro'* in the month. Google Trends series are standardized to the maximum search activity over the period June 2014 - November 2019.

Figure B.1: Demand Shifter Histograms

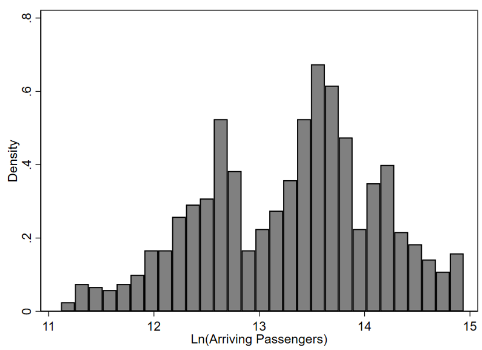

(a) Ln(Passengers)

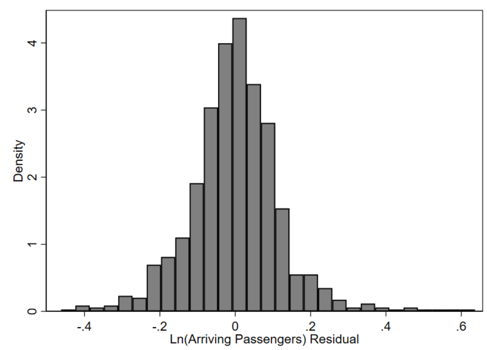

(b) Ln(Passengers) Residuals

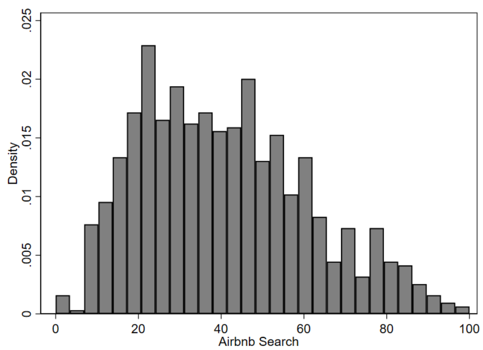

(c) Airbnb Search

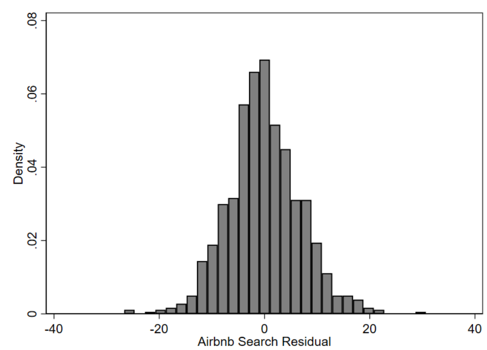

(d) Airbnb Search Residuals

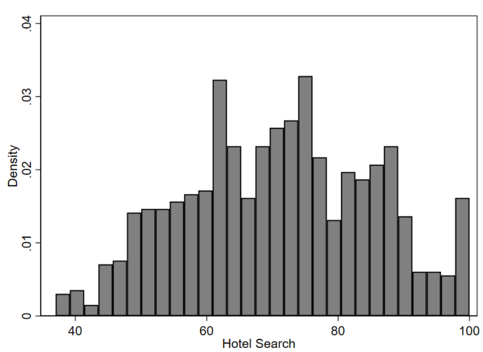

(e) Hotel Search

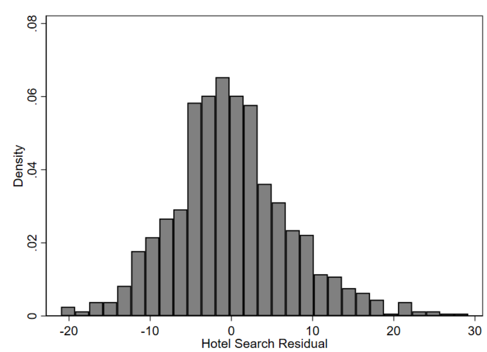

(f) Hotel Search Residuals

The figures report the distributions of the demand shifters. Each histogram displays the distribution of one of the shifters (Z_{mt}), using one observation per metropolitan area by month, which is the level of variation. The panels on the left side show the unconditional distribution, while the panels on the right hand side display the residualized analog. The residuals are obtained from a linear regression of (Z_{mt}) on metro and month fixed effects.

B.2 Assessing the Parallel Trend Assumption

We adopt the "honest approach" to parallel trends proposed by Rambachan and Roth (2023) to test the robustness of our findings to alternative assumptions about different trends in treated versus untreated tax jurisdictions. If we restrict the post-treatment violation of parallel trends to be no larger than the maximal pre-treatment violation of parallel trends, we obtain confidence sets that are slightly wider than the original ones but rule out a null effect on both prices and quantities. We also verify that the breakdown value for a null effect is around a violation that is twice as large as the maximal pre-treatment violation: see Figure B.2. We also construct robust confidence sets about how non-linear the difference in trends can be, allowing for linear violations of parallel trends and larger deviations from linearity. Our results are robust to linear violations and, up to the arbitrary amount $M \leq 0.03$, to nonlinear violations.

Figure B.2: Sensitivity estimates on nights and prices based on Rambachan and Roth (2023)

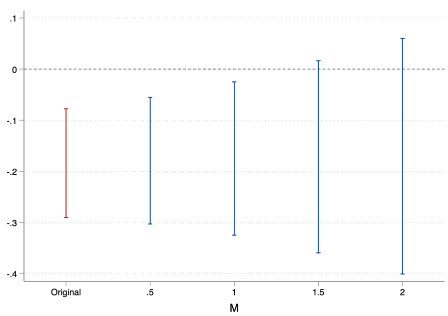

(a) Sensitivity on nights booked: Poisson TWFE

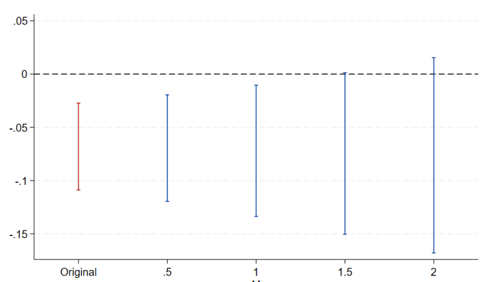

(b) Sensitivity on booking prices: TWFE

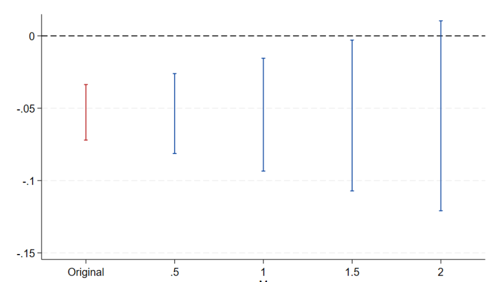

(c) Sensitivity on booking prices: Sun and Abraham (2021)

The figures report a sensitivity analysis of the estimated effects on nights (Panels a) and prices (Panels b and c) to potential violations of parallel trends per Rambachan and Roth (2023). The red bar in each panel represents the 95% confidence interval of the difference-in-difference estimate for $t = 4$ months after the introduction of a VCA agreement (baseline estimates). The blue bars represent the corresponding 95% confidence intervals permitting M deviations (x-axis) from the parallel trends assumption.

B.3 Robustness and Heterogeneity

Figure B.3: OLS TWFE Event Study and Sensitivity

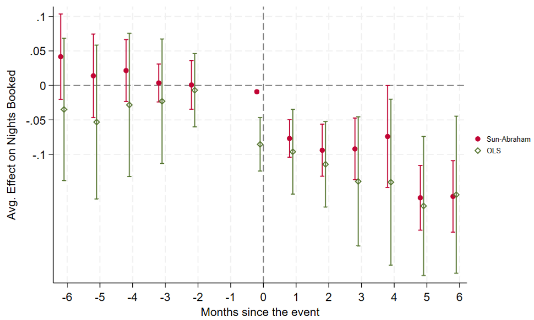

(a) Effect of VCAs on nights booked: OLS

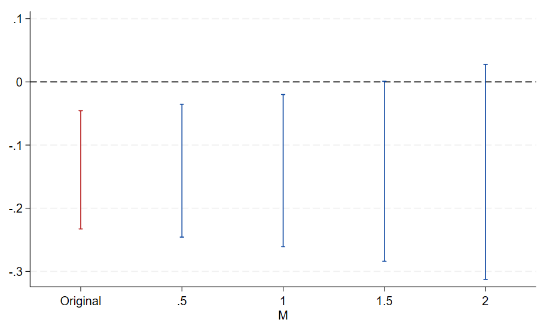

(b) Sensitivity on nights booked: TWFE

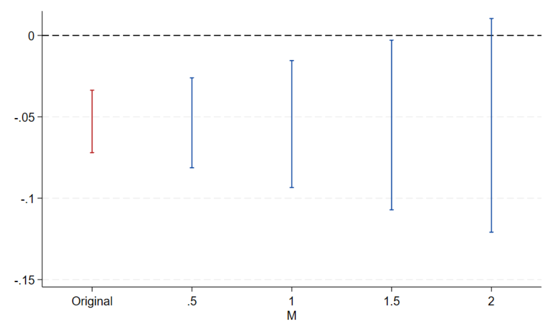

(c) Sensitivity on nights booked: Sun and Abraham (2021)

The figures report the (a) event studies and (b) and (c) sensitivity analysis (per Rambachan and Roth (2023)) of the estimated effects on nights booked using OLS regression with $\ln(1+Y)$ as the outcome variable rather than Poisson regression. The outcome is the logarithm of 1 plus the number of nights booked for property k in tax jurisdiction j and metro m in month-year t, $\ln(1 + Nights\ Booked_{kjmt})$, and the estimates are based on OLS regression that controls for property and month-year fixed effects. In Panels (b) and (c), the red bar represents the 95% confidence interval of the difference-in-difference estimate for $t = 4$ months after the introduction of a VCA agreement (baseline estimates). The blue bars represent the corresponding 95% confidence intervals permitting M deviations (x-axis) from the parallel trends assumption.

Table B.3: Reduced Form and Structural Estimates, OLS

Panel A: ln(1 + Nights Booked)			Google Searches: Hotels	Google Searches: Airbnb
$\ln(1+\tau_{jmt})$	-0.383**	-0.366**	-0.422***	-0.290
	(0.157)	(0.183)	(0.151)	(0.186)
ln(Arrivals)		0.539***		
		(0.058)		
Google Trends			0.008***	0.009***
			(0.001)	(0.001)
Observations	3,592,522	3,592,522	3,592,522	3,592,522

Panel B: Structural Parameter Estimates

	Demand Shifter: Passengers	Demand Shifter: Hotels Trend	Demand Shifter: Airbnb Trend
ε^d	-0.489	-0.558	-0.351
	(0.234)	(0.207)	(0.224)
ε^s	1.632	2.014	1.802
	(0.215)	(0.173)	(0.167)
λ_1	0.028	0.035	0.011
	(0.130)	(0.109)	(0.116)
p-value, $H_0 : \lambda_1 > 0.1$	0.287	0.274	0.221
p-value, $H_0 : \lambda_1 > 0.2$	0.091	0.064	0.052

Panel A reports the reduced-form estimates of the effect of tax collection agreement on nights booked (controlling for property and month-year fixed effects) using OLS with $ln(1+Y)$ as the outcome variable. The top row of Panel A $\ln(1+\tau_{jmt})$ includes the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. Panel B reports the corresponding structural parameter estimates, which are obtained using the Panel A OLS estimates for quantity along with the Table 4 price estimates. The number of jurisdictions is 78. Standard errors in Panel A, in parentheses, are clustered at the tax-jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1. Standard errors in panel B are computed from a non-parametric bootstrap with 500 repetitions and clustering at the tax-jurisdiction level. The $p-$values are calculated on the basis of the parameter estimates and their standard errors, assuming normality.

Table B.4: Effects of Within Metro Treatment on Control Units

	Booked	ln(Booking Price)	Booked	ln(Booking Price)
1[Tax in Metro]	-0.056	-0.004		
	(0.040)	(0.007)		
ln(1 + Max. Tax)			-0.338	-0.005
			(0.349)	(0.069)
Observations	629,214	256,741	629,214	256,741
Clusters	33	33	33	33

Estimated spillover effects of within-metro treatments on control units. The sample contains listings in jurisdictions with no VCA during our sample period. 1[Tax in Metro] is an indicator for having any VCA in the same metro in the observation month. ln(1 + Max. Tax) is the highest VCA enforced tax rate in the metro in the given month. All regressions include property fixed effects and month-year fixed effects. There are 33 jurisdictions in the sample.

Table B.5: Main Results, Including Largest Jurisdiction in Each Metro Only

			Google Searches	
			Hotels	Airbnb
Panel A: Nights Booked, Poisson TWFE				
$\ln(1+\tau_{jmt})$	-0.431	-0.424*	-0.416*	-0.289
	(0.286)	(0.251)	(0.225)	(0.232)
ln(Arrivals)		0.443***		
		(0.086)		
Google Trends			0.008***	0.010***
			(0.002)	(0.002)
Observations	2,411,942	2,411,942	2,411,942	2,411,942
Panel B: ln(Nightly Booking Price)				
$\ln(1+\tau_{jmt})$	-0.214*	-0.215**	-0.217**	-0.145**
	(0.115)	(0.087)	(0.089)	(0.065)
ln(Arrivals)		0.334***		
		(0.057)		
Google Trends			0.004***	0.005***
			(0.001)	(0.001)
Observations	976,112	976,112	976,112	976,112
Property FE	x	x	x	x
Month-Year FE	x	x	x	x
Panel C: Structural Parameter Estimates				
ε^d		-0.541	-0.532	-0.338
		(0.935)	(0.851)	(0.554)
ε^s		1.325	1.925	1.929
		(0.278)	(0.285)	(0.221)
λ_1		-0.105	0.001	-0.005
		(0.226)	(0.149)	(0.138)

Estimation of main results while including only the largest jurisdiction in each metro area. Panel A reports the reduced-form estimates of the effect of tax collection agreement on nights booked using Poisson regression. Panel B reports the reduced-form estimates on the booking price. The top row of each panel $\ln(1+\tau_{jmt})$ includes the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. Panel C includes the resulting estimates of the structural parameters using each demand shifter. The number of jurisdictions is 24. Standard errors, in parentheses, are clustered at the tax jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.6: Main Results, Dropping Controls in Treated Metros

			Google Searches	
			Hotels	Airbnb
Panel A: Nights Booked, Poisson TWFE				
$\ln(1+\tau_{jmt})$	-0.538*	-0.554**	-0.550**	-0.361
	(0.306)	(0.274)	(0.267)	(0.251)
ln(Arrivals)		0.468***		
		(0.065)		
Google Trends			0.009***	0.011***
			(0.002)	(0.002)
Observations	2,857,213	2,857,213	2,857,213	2,857,213
Panel B: ln(Nightly Booking Price)				
$\ln(1+\tau_{jmt})$	-0.217**	-0.235***	-0.233***	-0.134**
	(0.098)	(0.076)	(0.082)	(0.056)
ln(Arrivals)		0.332***		
		(0.046)		
Google Trends			0.004***	0.005***
			(0.001)	(0.001)
Observations	1,157,566	1,157,566	1,157,566	1,157,566
Property FE	x	x	x	x
Month-Year FE	x	x	x	x
Panel C: Structural Parameter Estimates				
ε^d		-0.725	-0.717	-0.416
		(0.747)	(0.788)	(0.475)
ε^s		1.409	1.975	2.118
		(0.209)	(0.240)	(0.242)
λ_1		-0.158	-0.045	-0.036
		(0.221)	(0.148)	(0.131)

Estimation of main results excluding control units in treated metros from the sample. Panel A reports the reduced-form estimates of the effect of tax collection agreement on nights booked using Poisson regression. Panel B reports the reduced-form estimates on booking price. The top row of each panel $\ln(1+\tau_{jmt})$ includes the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. Panel C includes the resulting estimates of the structural parameters using each demand shifter. The number of jurisdictions is 55. Standard errors, in parentheses, are clustered at the tax jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.7: Main Results, Conditional on Days Supplied

		Google Searches	
		Hotels	Airbnb
Panel A: Nights Booked, Poisson TWFE			
$\ln(1+\tau_{jmt})$	-0.484*	-0.466*	-0.325
	(0.261)	(0.249)	(0.238)
ln(Arrivals)	0.491***		
	(0.065)		
Google Trends		0.009***	0.011***
		(0.002)	(0.001)
Supply	0.022***	0.022***	0.022***
	(0.001)	(0.001)	(0.001)
Observations	3,118,578	3,118,578	3,118,578
Panel B: ln(Nightly Booking Price)			
$\ln(1+\tau_{jmt})$	-0.252***	-0.244***	-0.172***
	(0.079)	(0.080)	(0.059)
ln(Arrivals)	0.330***		
	(0.045)		
Google Trends		0.004***	0.005***
		(0.001)	(0.001)
Supply	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)
Observations	1,259,409	1,259,409	1,259,409
Property FE	x	x	x
Month-Year FE	x	x	x
Panel C: Structural Parameter Estimates			
ε^d	-0.647	-0.617	-0.392
	(0.779)	(0.762)	(0.498)
ε^s	1.488	2.089	2.185
	(0.242)	(0.275)	(0.269)
λ_1	-0.074	0.021	0.023
	(0.202)	(0.151)	(0.128)

Estimation of main results conditional on the number of days the property is available (nights in use plus nights available) in the given month. Panel A reports the reduced-form estimates of the effect of tax collection agreement on nights booked using Poisson regression. Panel B reports the reduced-form estimates on booking price. The top row of each panel $\ln(1+\tau_{jmt})$ includes the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. Panel C includes the resulting estimates of the structural parameters using each demand shifter. The number of jurisdictions is 78. Standard errors, in parentheses, are clustered at the tax jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.8: Reduced Form Estimates, Individual vs. Professional Hosts (Poisson)

			Google Searches	
			Hotels	Airbnb
Panel A: Nights Booked				
$\ln(1+\tau_{jmt}) \times < 5$	-0.619**	-0.631**	-0.619**	-0.383
	(0.313)	(0.278)	(0.269)	(0.248)
$\ln(1+\tau_{jmt}) \times \geq 5$	0.164	0.146	0.209	-0.075
	(0.255)	(0.259)	(0.240)	(0.277)
ln(Arrivals) $\times < 5$		0.472***		
		(0.065)		
ln(Arrivals) $\times \geq 5$		0.446***		
		(0.113)		
Google Trends $\times < 5$			0.009***	0.010***
			(0.001)	(0.001)
Google Trends $\times \geq 5$			0.007***	0.014***
			(0.002)	(0.002)
	3,118,578	3,118,578	3,118,578	3,118,578
Panel B: ln(Nightly Booking Price)				
$\ln(1+\tau_{jmt}) \times < 5$	-0.222*	-0.213**	-0.217**	-0.115
	(0.115)	(0.090)	(0.090)	(0.070)
$\ln(1+\tau_{jmt}) \times \geq 5$	-0.305***	-0.443***	-0.386***	-0.406***
	(0.059)	(0.071)	(0.073)	(0.083)
ln(Arrivals) $\times < 5$		0.302***		
		(0.043)		
ln(Arrivals) $\times \geq 5$		0.447***		
		(0.059)		
Google Trends $\times < 5$			0.004***	0.005***
			(0.001)	(0.001)
Google Trends $\times \geq 5$			0.005***	0.006***
			(0.001)	(0.001)
	1,259,371	1,259,371	1,259,371	1,259,371
Property FE	x	x	x	x
Month-Year FE	x	x	x	x

The table reports the reduced-form estimates of the effect of tax collection agreement on nights booked (Panel A) and booking price (Panel B) for two subsets of the sample: (i) listings from hosts with fewer than 5 listings ("Individual") and (ii) listings from hosts with 5 or more listings ("Professional"). The top two rows of each panel $\ln(1+\tau_{jmt})$ include the estimated effects of tax enforcement. The first column includes no additional demand shifter. Columns 2-4 include an additional demand shifter (Z_{mt}). Column 2 includes the logarithm of incoming flight passengers. Columns 3 and 4 include the volume of searches reported in Google Trends for hotels and Airbnb in the month. All estimates are from a single regression that includes interactions between the tax variable and indicators for hosts with fewer than 5 listings and hosts with 5 or more listings, and interactions between the demand shifter and indicators for hosts with fewer than 5 listings and hosts with 5 or more listings. The number of jurisdictions is 78. Standard errors, in parentheses, are clustered at the tax-jurisdiction level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.9: Structural Parameter Estimates, Individual vs. Professional Hosts (Poisson)

	Demand Shifter		
	Passengers	Hotels Trend	Airbnb Trend
Panel A: Hosts with < 5 Listings			
ε^d	-0.802	-0.791	-0.433
	(0.871)	(0.863)	(0.504)
ε^s	1.563	2.262	2.174
	(0.286)	(0.268)	(0.273)
λ_1	-0.191	-0.057	-0.062
	(0.191)	(0.137)	(0.134)
p-value, $H_0 : \lambda_1 > 0.1$	0.063	0.126	0.106
p-value, $H_0 : \lambda_1 > 0.2$	0.020	0.031	0.023
Panel B: Hosts with ≥ 5 Listings			
ε^d	0.262	0.340	-0.127
	(0.947)	(1.091)	(0.868)
ε^s	0.999	1.348	2.197
	(0.216)	(0.261)	(0.312)
λ_1	0.589	0.541	0.372
	(0.396)	(0.271)	(0.178)
p-value, $H_0 : \lambda_1 > 0$	0.068	0.023	0.018

The table reports the structural parameters with standard errors (in parentheses) for two subsets of the sample: (i) listings from hosts with fewer than 5 listings ("Individual") and (ii) listings from hosts with 5 or more listings ("Professional"). Structural parameter estimation based on the reduced-form results in Table B.8. Standard errors are computed from a bootstrap with 500 repetitions and clustering at the tax jurisdiction level. The first column includes estimates using the incoming flight passengers variable. Columns 2 and 3 include estimates using the volume of searches reported in Google Trends for hotels and Airbnb. The $p-$values are calculated on the basis of the parameter estimates and their standard errors, assuming normality.