The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry[†]

Georgios Zervas
School of Management
Boston University

Davide Proserpio, John W. Byers
Computer Science Department
Boston University

Last revised: February 11, 2015
First draft: December 14, 2013

Abstract

Spurred by technological advancement, a number of decentralized peer-to-peer markets, now colloquially known as the sharing economy, have emerged as alternative suppliers of goods and services traditionally provided by long-established industries. A central question surrounding the sharing economy regards its long-term impact: will platforms like Lyft and Airbnb materialize as viable mainstream alternatives to traditional providers, or will they languish as niche markets? In this paper, we study Airbnb, a sharing economy pioneer offering short-term accommodation. Combining data from Airbnb and the Texas hotel industry, we estimate the impact of Airbnb's entry into the Texas market on hotel room revenue, and study the market response of hotels. To identify Airbnb's causal impact on hotel room revenue, we use a difference-in-differences empirical strategy that exploits the significant spatiotemporal variation in the patterns of Airbnb adoption across city-level markets. We estimate that each 10% increase in Airbnb supply results in a 0.35% decrease in monthly hotel room revenue, translating to an impact exceeding 13% of revenue in Austin, where Airbnb supply is highest. We find that Airbnb's impact is non-uniformly distributed, with lower-priced hotels, and hotels not catering to business travel being the most affected segments. Finally, we find that affected hotels have responded by reducing prices, an impact that benefits all consumers, not just participants in the sharing economy. Our work provides empirical evidence that the sharing economy is making inroads by successfully competing with, and acquiring market share from, incumbent firms.

1 Introduction[†]

Much ado has recently been made of the "sharing economy", whereby the emergence of multi-sided technology platforms managed by firms such as Airbnb, Uber, Lyft and TaskRabbit have enabled individuals to collaboratively make use of under-utilized inventory via fee-based sharing. On the supply side, these platforms have enabled individuals to list their spare rooms for rent, become drivers for the day, or rent themselves out by the hour to perform small tasks and errands. On the demand side, consumers using the platforms benefit foremost from increased competition: each incremental Airbnb listing is another viable accommodation option, each additional Lyft driver lessens (sometimes artificial) scarcity of taxicabs; but also consumers benefit from availability of new services: arguably, before TaskRabbit, there was no inexpensive way to have your shoes picked up from the cobbler, other than to do it yourself.

While sharing of goods and services is an age-old phenomenon, the sharing economy is different, and goes beyond renting goods at lower costs, or with lower transactional overhead, than buying or renting through a traditional provider. The key differentiating enabler is that technological advances enable the networked platforms that underlie the sharing economy to operate at scale. Consider the case of Airbnb, the pioneer in short-term shared accommodations and rentals: for this firm, maintaining a worldwide database of millions of property listings with photographs and user reviews, and facilitating the hundreds of thousands of daily transactions on their platform, is now (but only quite recently) an entirely tractable data management problem for a start-up firm to undertake. Operating at scale also works to these firms' benefit, since, as the literature on platform economics predicts (as we review in Section 1.2), positive network effects create benefits to market participants as scale increases. As a result, suppliers with niche or remnant inventory that was formerly difficult to bring to market may now find greatly reduced barriers to entry using a platform: they can list their supply and obtain global reach, and incur costs that are typically only a fixed fraction of their revenue[1]. Moreover, these suppliers may take their supply on and off the market on a moment's notice, and experiment with pricing as they see fit. As a result of this convenience and flexibility offered to suppliers, successful platforms can list a much wider

[†]The authors thank the participants and organizers of SCECR'13 (http://scecr.org/scecr2013/), WISE'13 (http://wiseconf.org), and the seminar participants at Telefonica I+D Research, Barcelona and Technicolor Research, Paris for their helpful feedback on earlier drafts of this work. We thank Smith Travel Research (STR) for sharing data with us. We are also indebted to Flavio Esposito for motivating us to investigate Airbnb and for his contributions to our earlier research on the topic.

[1]Most sharing economy platforms, including those mentioned above, charge no up-front costs to their suppliers, but instead take a percentage cut for each transaction. Thus the only fixed cost for a supplier to list their supply on a platform is typically the cost of creating a webpage.

variety of supplied inventory than was previously possible; moreover, they can exploit search and personalization technology to increase ease of use and optimize the user experience on the demand side (Fradkin 2013). To date, we view the rise of platforms as having two distinct phases. In the first phase, web sites like eBay and Craigslist enabled the *selling* of goods through the use of searchable listings, allowing suppliers to reach broader audiences at minimal cost, and for purchasers to browse a massive online inventory from their desktop. In the second phase, platforms such as Airbnb have addressed the more challenging and less well-circumscribed problem of the *sharing* of goods and services, where unlike selling a good, retention of value, post-transaction, is an essential consideration for the supplier.

Along with the rise of the sharing economy comes a host of open research questions. Many of these, including those motivating our work, center on a debate about the sharing economy involving divergent views on its staying power and potency in the long run, especially as regards its contribution to social welfare. Clearly, there exist beneficial transactions enabled through the sharing economy, such as the cobbler example, that provide positive utility, have no negative externalities, and thus manifest themselves as a net gain in social welfare. Proponents of the sharing economy opine that a vast number of hypothetical transactions are of this form, and thus, the ability to unlock the sharing economy could provide a great boon to social welfare (see, *e.g.,* Botsman (2012)). On the other hand, unlocking these transactions have already come at some cost, for example due to tensions between taxi drivers losing business to Uber and Lyft, and anecdotes of neighbors disgruntled by the whims of short-term renters using Airbnb.[2] But more pointedly, detractors of the sharing economy at large view welfare gains in the long run as likely to be insubstantial[3]; and others, such as certain hoteliers, view upstarts such as Airbnb as a non-threat.[4]

Our work seeks to contribute to the debate regarding the impact of the sharing economy on incumbent firms; we contribute empirical evidence to the set of existing opinion pieces. The central question we study is, to what extent (if any) Airbnb has impacted the hospitality industry at large. If it is the case that transactions within this segment of the sharing economy are either complementary to pre-existing economic activity, or comprise a transaction volume that is inconsequential, then our analysis should find no effect. On the other hand, to the extent that an impact is identified, the magnitude of the effect both informs this debate and has managerial implications for hoteliers. In this case, we seek to

[2]See: "Is Airbnb Bothering the Neighbors?", September 2012, `http://betabeat.com/2012/01/is-airbnb-bothering-the-neighbors/`.

[3]See: "The Rise Of The Renting And Sharing Economy Could Have Catastrophic Ripple Effects", August 2013, `http://www.businessinsider.com/rise-of-the-renting-and-sharing-economy-2013-8?op=1`

[4]See: "Marriotts CEO Just Made a Pretty Good Sales Pitch for...Airbnb?", July 2014, `http://time.com/money/2964290/marriott-airbnb-arne-sorenson-neighborhoods/`

investigate what mechanisms are at work to produce these impacts, and from a managerial standpoint, we ask what factors make incumbent hotels more or less vulnerable to adoption of Airbnb in their area. We view our study as timely, as the market for short-term rentals provided by Airbnb has now reached a scale, tens of millions of room-nights, that certainly has the potential to meaningfully and measurably impact revenue at the most vulnerable hotels, if indeed it ever will, and thus even a finding of no effect would constitute a useful contribution.

1.1 Summary of Methodology and Findings

We empirically examine the impact of Airbnb on the hospitality industry in the state of Texas. Using data we collected from Airbnb over the six year period from 2008 to 2014, and monthly hotel room revenue data from approximately 4,000 hotels in Texas dating back to 2003, we show that Airbnb penetration has negatively impacted hotel room revenue, and that lower-end hotels incur most of the financial impact. Our empirical strategy identifies the causal relationship between reduced hotel room revenue and increases in Airbnb supply using difference-in-differences (DD). Specifically, due to the significant variability in both the temporal rate and the spatial density of Airbnb adoption, as well as the geographic specificity of both our hotel and Airbnb datasets, we are able to treat Airbnb market entry as a variable intervention in space and time against the hotel room revenue data. Our DD strategy identifies the Airbnb treatment effect by comparing differences in revenue for hotels in cities affected by Airbnb before and after Airbnb's entry against a baseline of differences in revenue for hotel in cities unaffected by Airbnb over the same period of time. Using this DD specification we find that, in Texas, each additional 10% increase in the size of the Airbnb market resulted in a 0.35% decrease in hotel room revenue. To calibrate the economic significance of this result, it is worth pointing out that in certain Texas municipalities (notably, Austin), Airbnb has grown exponentially over the past few years, resulting in an estimated revenue impact of over 13% for the most vulnerable hotels in our data.

Our DD specification allows for both time-invariant differences in revenue between hotels as well as common time-varying shocks to revenue across hotels. The key threat to identification potentially arises in the form of unobservable (to the researcher) city-specific, time-varying factors that differentially affect hotel room revenue depending on the intensity of Airbnb adoption within each city at a given point in time. To test the robustness of our DD estimate, we perform a series of checks. First, we show that our estimate for Airbnb's impact on hotel room revenue is robust to a number of covariates that vary by location and over time (*e.g.*, country-specific population, unemployment rate, and total hotel room

supply measured at the city level). Second, following common DD practice, we allow for flexible city-specific trends (*e.g.*, linear, or quadratic), which can parametrically control for unobserved endogenous trends that vary by city (the same level at which we observe variation in Airbnb supply). Third, we show that the basic set of controls included in our DD specification (*i.e.*, hotel fixed effects and temporal trends) explain approximately 93% of the variation in Airbnb supply. Therefore, little variation in Airbnb supply remains unexplained by our model, and could potentially be driven by unobserved factors that also affect hotel room revenue. Fourth, we check whether Airbnb adoption is driven by hotel performance, which would be a case of our confusing cause and effect. To the contrary, we find that a wide range of pre-Airbnb demographic and market characteristics – including, for example, hotel room prices, occupancy rates, and hotel room supply per city – that are significant predictors of post-Airbnb hotel room revenue, are not correlated with the patterns of Airbnb adoption we see in our data. Finally, in a separate analysis, we combine DD with coarsened exact matching (Iacus et al. 2012) to further reduce endogeneity concerns. Specifically, we first match each hotel affected by Airbnb to unaffected hotels belonging to same price-tier and sharing the same affiliation, discarding hotels that remain unmatched. The intuition behind matching is that similar hotels (*e.g.*, an upscale Hilton in Austin where Airbnb adoption is high, and an upscale Hilton in Dallas where Airbnb penetration is low) are less likely to differ in unobserved ways. We find that our CEM DD estimate is similar to our main DD analysis. Taken together, these robustness checks provide significant support for the assumptions underlying our DD analysis.

We then move to investigate both the mechanisms behind Airbnb's impact on hotel room revenue and the market response to Airbnb entry. With respect to mechanisms, given the nature of rentals on Airbnb today, which typically provide fewer amenities and services than many hotels, we expect those hotels providing more differentiated services to be less affected. We start by examining two such cases: high-end hotels and hotels catering to business travelers, both of which provide amenities that a typical Airbnb stay does not. First, segmenting hotels in five industry-standard price tiers (Budget, Economy, Midprice, Upscale, and Luxury) we find the impact of Airbnb is gradually magnified as we move down the price tiers. Then, through a similar analysis, using conference and meeting room space as a proxy for the extent to which a hotel caters to business travel, we find that the impact of Airbnb also falls disproportionately on those hotels lacking conference facilities. Finally, we examine Airbnb's differential impact on chain versus independent hotels, expecting that chain hotels will be less affected, for reasons ranging from larger marketing budgets and stronger brands to providing predictably consistent service. In contrast, independent hotels exhibit more variability and perhaps more inconsistency, as we would also expect with Airbnb prop-

erties. Indeed, our analysis confirms that the impact of Airbnb on independent hotels is disproportionately larger. Finally, with respect to market response, we study the extent to which affected hotels react to Airbnb's market entry. By reflecting on hotel industry performance metrics, notably occupancy and price, we find no statistically significant evidence of a change in occupancy rates, while instead, we observe almost all of the revenue impact being shouldered as a (negative) price response. Here, lower hotel room prices in areas with an established Airbnb presence benefit all consumers, whether they use Airbnb or not.

1.2 Related Work

Relatively few empirical papers have yet studied the sharing economy and its interplay with incumbent firms offering similar goods or servces. A handful of studies have examined the adoption and effects of car-sharing, for example, two studies have used survey analysis methods to find that car-sharing is associated with significant decreases in miles traveled, gasoline consumption, and car ownership (Cervero et al. 2007, Martin et al. 2010). As for accommodation sharing, we find a large number of opinion pieces in the popular press and on blogs, but little in the way of academic literature. Our closest comparison point is a set of short studies, commissioned by Airbnb, which claim that the Airbnb business model is complementary to the hotel industry, but primarily focus on arguing for and quantifying the substantial net economic benefit to cities that Airbnb travelers provide.[5] While our work is related to these studies, we apply a more sophisticated identification strategy, methodology, and segmentation analysis, resulting in conclusions that are both different and more nuanced.

Our work contributes to the growing literature on multi-sided platform competition, as Airbnb exemplifies a two-sided platform. Much of this literature establishes the economic theory of two-sided markets, for example through structural models that establish theories of price structure and usage (Rochet and Tirole 2003, Rysman 2009, Weyl 2010), and models which connect innovations in product design to network effects (Parker and Van Alstyne 2005). Other work, more closely related to our own, contributes empirical results to the literature that seek to explain the behavior of firms and individuals in two-sided markets (Jin and Rysman 2012), including the role of multihoming Landsman and Stremersch (2011), modeling response to regulation Carbó Valverde et al. (2010), and understanding the supply-side labor market Hall and Krueger (2015). Our work, in contrast to these, empirically studies a setting where a two-sided platform offers a substitute for consumer services supplied by traditional firms.

It is in this latter context that our work contributes to literature on substitution be-

[5]See: `https://www.airbnb.com/economic-impact/`

tween online and offline markets, as firms like Airbnb can be viewed as providing enabling technology that facilitates suppliers of niche inventory to bring their products to market. In contrast to offline markets, Airbnb provides sufficiently low cost of revenue for individuals to profitably list remnant inventory online; moreover, Airbnb provides enhanced reach by reducing consumer search costs (Bakos 1997). As such, our study can be viewed as investigating the consequences of an online platform lowering the barrier to entry for suppliers. Related work has studied similar examples in other domains. For example, a number of recent studies have focused on the impact of Craigslist – a website featuring free online classified ads – on the newspaper industry. Seamans and Zhu (2013) estimate the effect of Craigslist's market entry on several newspaper performance metrics. They find that in the face of increasing competition by Craigslist, newspapers with greater reliance on classified ad revenue responded by reducing their ad rates, and by increasing their subscription prices more than newspapers whose revenue were less reliant on advertising. Kroft and Pope (2014) estimate that Craigslist's entry resulted in a 7% reduction in the volume of classified ads appearing in newspapers during the period between January 2005 and April 2007. Further, they estimate that Craiglist's entry caused a decrease in the rental vacancy rate by approximately 1%. Our work shares a methodological trait with these studies: all of them rely on the temporal and geographic variation in Craigslist's entry to identify its effect. We exploit similar variation in the patterns of Airbnb adoption to measure its impact on hotel room revenue.

Finally, our work contributes to the literature studying the impact from external shocks on the tourism and the hospitality industry. Much of the prior work though, has centered on demand shocks. For example, O'Connor et al. (2008) study the impact of terrorism on tourism in Ireland; Baker and Coulter (2007) estimate the impact of the 2002 and 2005 terrorist attacks in Bali on the islands' vendors. Similarly, Kosová and Enz (2012) examine the adverse effects of the 9/11 attack and the 2008 financial crisis on hotel performance.

2 Data and the Airbnb Platform

For our study, we collect and combine data from various sources including the Airbnb website, the Texas Comptroller Office, Smith Travel Research (STR), county demographics from the U.S. Census Bureau, and the Current Population Survey (CPS) from the U.S. Bureau of Labor Statistics (BLS).

2.1 The Airbnb Platform

Much of the data used in our study is collected directly from the Airbnb website. Airbnb defines itself as *"a social website that connects people who have space to spare with those who are looking for a place to stay"*, and exemplifies a peer-to-peer marketplace in the sharing economy. Prospective hosts list their spare rooms or apartments on the Airbnb platform, establish their own nightly, weekly or monthly price, and offer accommodation to guests. Airbnb derives revenue from both guests and hosts for this service: guests pay a $9 - 12\%$ service fee for each reservation they make, depending on the length of their stay, and hosts pay a 3% service fee to cover the cost of processing payments. Since its launch in 2008, the Airbnb online marketplace has experienced very rapid growth, with more than one million properties worldwide and 30 million guests that used the service by the end of 2014 (18M of which was in the past year).[6]

Airbnb's business model currently operates with minimal regulatory controls in most locations, and as a result, hosts and guests both have incentives to use signalling mechanisms to build trust and maximize the likelihood of a successful booking. To reinforce this behavior, Airbnb has built an online reputation system that enables and encourages participants to rate and review each completed stay. Guests use star ratings to rate features of their stay, *e.g.*, cleanliness, location, and communication, while both guests and hosts are encouraged to post public reviews of each stay on the platform.

2.2 Airbnb Data: Listings and Market Entry

To estimate the extent of Airbnb's market entry, we periodically collected consumer-facing information from `airbnb.com` on the complete set of users who had listed their properties in the state of Texas for rental on Airbnb.

We refer to these users as *hosts*, and their properties as their *listings*. Each host is associated with a set of attributes including a photo, a personal statement, their listings, guest reviews of their properties, and Airbnb-certified contact information. Similarly, each listing displays attributes including location, price, a brief textual description, photos, capacity, availability, check-in and check-out times, cleaning fees, and security deposits. Figures 1 and 2 display a typical Airbnb listing and a typical Airbnb user profile, respectively. Our collected dataset contains detailed information on $10,555$ distinct hosts and $13,935$ distinct listings spanning a period from 2008 to August 2014.

We quantify Airbnb supply over time at the granularity of individual cities as follows: for a given city and date, we count the number of distinct listings that have (cumulatively)

[6] See http://www.cnbc.com/id/102389442.

appeared on Airbnb in that city prior to that date. We approximate the unobservable entry date of individual listings by using the prominently displayed date their owners became Airbnb members. We note that instantaneous supply fluctuates continuously, as some Airbnb hosts take properties on and off the market. Nevertheless, due to Airbnb's exponential growth, at any given point in time, cumulative supply strongly correlates with instantaneous supply.

While the presence of Airbnb listings in a city clearly does not by itself impact hotels, regressing hotel room revenue on Airbnb supply produces a meaningful coefficient estimate. We interpret a statistically significant negative coefficient on Airbnb supply as indicating that Airbnb listings lead to Airbnb bookings that substitute for hotel stays and impact hotel room revenue. We interpret a coefficient that is not statistically significantly different from zero as Airbnb listings having no effect on hotels. We interpret a positive coefficient, though implausible, as Airbnb listings benefiting hotels.

Separately, we must choose an appropriate level of geographic aggregation. Here, our data is suitably granular (with location accuracy to roughly 100 meters) to permit analysis at many different scales. Our choice of city-level granularity is driven by the observation that a city is the largest geographic unit within which we reasonably expect to see significant substitution patterns between hotels and Airbnb properties.

2.3 Hotel Data: Revenue, Prices, and Occupancy Rates

The main dependent variable we use in our analysis is monthly hotel room revenue, which we obtained from public records furnished by the Texas Comptroller's office, in their capacity as auditors of state tax collection.[7] In addition to monthly hotel room revenue, the dataset includes basic information including hotel name, address, and capacity. The raw dataset spans the period between Jan. 2003 and Aug. 2014 and contains approximately $735,000$ records.

Interestingly, according to Texas law, "a hotel is considered to be any building in which members of the public rent sleeping accommodations for $15 or more per day."[8] For this reason, revenue from Airbnb properties (as well as various other vacation rental options) complying with the Texas tax code is also reported in this dataset. This is evident from Figure 3, which plots the number of unique tax-paying properties in Austin broken down by capacity, *i.e.*, maximum occupancy. We conjecture that the rapid increase in low capacity properties starting in 2008 is related to Airbnb's entry into the Texas market at the same time. To exclude non-hotel properties from our analysis of impact on hotels, we cross-

[7] Available at `http://aixtcp.cpa.state.tx.us/hotel/hotel_qtr_all_srch.php`

[8] See `http://www.window.state.tx.us/taxinfo/hotel/faqhotel.html`

reference the Texas Comptroller dataset with the U.S. hotel census data provided to us by STR. The STR census includes all U.S. hotels and contains a rich attribute set for each hotel, including its opening date, price segment, capacity, operation type (chain vs. independent), and geographic location. In total, the STR dataset contains information on 4,673 hotels in the state of Texas. After linking the STR census dataset with the Texas tax dataset, we obtain high-confidence matches for a panel of 3,908 properties (83.6% of the Texas hotels in the STR census, which account for over 90% of the revenue in our data), and we exclude the rest. Given that we conduct a within-hotel analysis of revenue, unmatched hotels should not bias our estimation. This operation leaves us with about $370,000$ observations.

Airbnb can affect hotel room revenue through lower occupancy rates, decreased hotel room prices, or a combination of these two factors, conventionally reported within the hotel and hospitality industry as RevPAR (revenue per available room), which is the product of average room price and occupancy. Because the data we obtained from the Comptroller's office does not report either occupancy rates or hotel room prices, we obtain additional data on these quantities for a subset of Texas hotels from STR. The room price (also referred to as average daily rate, or ADR in the industry) and occupancy rate data from STR covers a subset of $3,225$ matched hotels in Texas who chose to report this information to STR over the same time period: Jan. 2003 to Aug. 2014.

Finally, we assemble a set of control variables derived from publicly available sources. We obtain monthly unemployment data at the city level, and annual demographic information at the county level from the BLS at `bls.gov` and the U.S. Census Bureau at `census.gov`.

3 Empirical Strategy and Results

Airbnb has seen widely varying degrees of traction within different local, regional and international markets, both with respect to initial market entry and the rate at which it has been adopted within markets. For example, consider Figure 5, which depicts the current extent of market penetration both of Airbnb properties and hotels within the state of Texas (top panels), and within the county encompassing the state capital, Austin (bottom panels). Unlike hotels, which have coverage throughout the state, and pockets of local density, such as in downtown Austin, Airbnb has spotty coverage at best throughout the state, but broader coverage across metro areas, including suburbs and exurbs. Table 1 reveals that patterns of Airbnb adoption, over the past eight years in the ten most populous cities in Texas, are themselves diverse, with several cities experiencing early adoption and rapid growth, while others experienced minimal Airbnb adoption. Our empirical strategy exploits this variability to identify the impact of Airbnb's rise on hotel room revenue using a differences in differences

(DD) identification strategy. Specifically, we estimate Airbnb's impact on hotel room revenue by comparing changes in hotel room revenue before and after Airbnb enters a specific city, against a baseline of changes in hotel room revenue in cities with no Airbnb presence over the same period time.

The key identification assumption we have to make to support a causal interpretation of this DD estimate is that there are no unobserved, time-varying, city-specific factor that are correlated with *both* Airbnb entry and hotel room revenue. Stated differently, we assume that unobserved factors that could potentially jointly affect both Airbnb adoption *and* hotel room revenue do not systematically vary *both* between different cities *and* over time. For instance, the following unobserved factors are accounted for in our estimate and do not bias our estimates: 1) city-specific time-invariant differences in adoption rates (*e.g.,* consumers in Austin overall being more likely to adopt Airbnb than consumers in Dallas); 2) factors that vary arbitrarily over time but do not vary across cities (*e.g.,* a generally increasing awareness of Airbnb shared across all consumers in Texas over time), and, 3) city-specific trends, which allow for unobserved confounders that vary both between cities and over time according to a pre-specified functional form (linear, quadratic, or cubic). Our DD specification takes the following form:

$$\log \text{Hotel Revenue}_{ikt} = \beta \log \text{Airbnb Supply}_{kt} + h_i + \tau_t + X'_{ikt}\gamma + \epsilon_{ikt}. \tag{1}$$

The dependent variable is the log of monthly room revenue of hotel i in city k at time t. To implement the DD strategy, our model includes hotel fixed effects h_i, and time (year-month) fixed effects τ_t. The first difference is taken using the hotel fixed effects, which allow for time-invariant differences in hotel room revenue between treated hotels (*i.e.*, hotels in cities with an Airbnb presence) and non-treated hotels (*i.e.,* hotels in cities with no Airbnb presence). The second difference in our DD specification is taken over time using year-month fixed effects τ_t which allow for unobserved time-varying revenue differences that are common across different cities. The coefficient of interest is β, which has the usual DD interpretation: it is an estimate of the percentage change in hotel room revenue in Airbnb-adopting cities subsequent to Airbnb's entry compared against a baseline of changes in hotel room revenue over the same time period in cities where Airbnb does not have a presence.

We now discuss and motivate the specific form of the specification and the controls we use, as well as the other best-practice methodologies from the literature that we employ, in carrying out this empirical identification strategy.

As we explained earlier, our DD estimate will be biased if there exist unobserved factors that vary across cities and over time, and which jointly influence Airbnb entry and hotel

room revenue. To further guard against this possibility, we allow for quadratic city-specific trends as a control in X_{ikt}. A concern with the inclusion of city-specific time-trends is that they can be confounded with hotels' response to Airbnb (Wolfers 2006). Fortunately, our dataset covers a long pre-Airbnb period from 2003 to 2008, allowing us to estimate these trends on a large sample of pre-treatment observations. As middle ground between flexibility and model complexity, we chose to include quadratic trends, but our results are not sensitive to different functional form specifications that we tested, including linear and cubic trends.

Another identification challenge we face is that increased demand for accommodation is likely correlated with increases in both Airbnb supply and hotel room supply. Concretely, it is plausible that over our decade-long observation period, hotel firms have been strategically developing new properties in areas of anticipated high demand. This pattern of competition could bias our estimation, because city-specific increases in hotel room supply can also drive per-hotel room revenue down, while at the same time correlate with increased Airbnb adoption. To guard against this type of concern, we construct a control variable $Hotel\ Room\ Supply_{-ikt}$, which measures the total supply of hotel rooms in the same city as hotel i (but excluding hotel i itself, thus the $-i$ in the subscript), for each time t. This control, which we also incorporate in X_{ikt}, allows for increases in the supply of hotel rooms provided by competitors to impact the room revenue of each hotel in our data, much as we hypothesize an increase in Airbnb rooms does.

A final issue we have to deal with is that the unit of analysis is hotel monthly room revenue, but the treatment, Airbnb adoption, occurs at the city level. As is well known, this mismatch in the level at which we measure our dependent variable compared to the treatment variable can result in understating the standard error of the estimate of Airbnb's impact, because it is likely that hotel room revenue is serially correlated over time within a city. We correct for this mismatch by clustering standard errors at the city level, which lets us account for possible serial correlation in hotel room revenue. In so doing, we follow the standard practice in the literature for analyzing panel data in a DD setting (see *e.g.*, the treatment recommended by Bertrand et al. (2004), as well as Donald and Lang (2007)). We report standard errors clustered at the city level for all subsequent regressions.

3.1 Identification checks

Before proceeding with estimation, we conduct a series of identification checks to assess whether our proposed empirical strategy can recover Airbnb's causal impact on hotel room revenue. Our DD identification strategy relies on randomness in Airbnb adoption with respect to unobserved city-specific time-varying factors (ϵ_{ikt}) that are also correlated with

changes in hotel room revenue (conditional on the control variables we include). As with any study relying on observational data, there is no conclusive test of this assumption. However, we can exploit the richness of our data to check if this assumption is likely to hold in practice. Similar to Akerman et al. (2013), we perform two checks that support the basis for our key identification assumption.

First, we show that most variation in Airbnb adoption is explained by time-invariant city-specific factors, time fixed effects, and city-specific trends – all of which are part of the DD model. To implement this test, we regress (the log of) Airbnb supply on city and time fixed effects as well as city-specific quadratic time-trends. These factors explain 93% of variation in Airbnb adoption, suggesting that our modeling assumption has a sound basis in practice. Next, we repeat this regression with the addition of city-specific time-varying observables that could potentially be correlated with hotel room revenue: population, unemployment rate, and employment in the accommodation sector. Together, these factors only account for an additional .02% of variation in Airbnb supply. These analyses suggests that little variation in Airbnb supply remains unexplained, and thus could potentially be correlated with the error term in our DD regression.

Second, we check whether pre-treatment city characteristics predict future Airbnb supply, where the time of treatment is taken to be 2008, when Airbnb entered the Texas market. The idea behind this test is that assuming Airbnb adoption is exogenous (with respect to hotel performance), it should not be correlated with pre-treatment factors. To perform this identification check, for each city, we compute its most recent pre-treatment (2007) population, unemployment rate, employment in the accommodation sector, hotel room supply, hotel room prices, and hotel occupancy rates. We then interact these pre-determined factors ($Z_{k,2007}$) with a vector of post-treatment year-month fixed effects (τ_t), and regress them on Airbnb supply. Concretely, with the units of analysis being post-2007 city-months, we estimate:

$$\log \text{Airbnb Supply}_{kt} = \text{City}_k + (\tau_t \times Z_{k,2007})' \theta + e_{kt}. \tag{2}$$

Each coefficient in the vector of coefficients θ is interpreted as a correlation between a specific pre-treatment characteristic and Airbnb adoption in each post-treatment period (from January 2008 onwards). Figure 4 presents the estimated coefficients θ for each characteristic together with their 95% confidence intervals. The only significant association we find is between pre-Airbnb population and subsequent Airbnb adoption, and, for this reason, we include population as a control in all our specifications. Visually, there also appears to be a weak correlation with pre-Airbnb unemployment rate, possibly driven by the help Airbnb

can provide to struggling or unemployed homeowners in paying their mortgage,[9] though nearly all individual correlation coefficients making up this trend are not statistically different from zero. Regardless, we also include county-level unemployment rates as a control in Equation 1. Beyond these associations, we find no other discernible trend in the remaining coefficients (whose 95% confidence intervals always include the zero point, or, no effect). It is especially reassuring that the pre-treatment hotel industry structure – as captured by hotel room supply, occupancy rates, room prices, and accommodation sector employment in 2007 – do not predict Airbnb supply from 2008 onwards.

As mentioned earlier, one cannot entirely rule out endogeneity concerns in a study using observational data. Here, we have shown that various factors potentially affecting hotel room revenue, including demographic trends, as well as the structure and performance of the hospitality industry across different cities, are not correlated with local patterns of Airbnb adoption. These checks increase our confidence that the identification assumptions needed to estimate Airbnb's causal impact on hotel room revenue hold in our data.

3.2 Results and Economic Significance

We report the results of estimating Equation 1 in the first column of Table 2. We estimate the coefficient $\beta = -0.035$, or equivalently, a 10% increase in Airbnb listings is associated with a statistically significant 0.35% ($p < 0.01$) decrease in monthly hotel room revenue. As we stated earlier regarding interpretation of a negative coefficient β, this estimate indicates that Airbnb listings result in some Airbnb stays that are substitutes for hotel stays in cities with an established Airbnb presence. We interpret the magnitude of this estimate shortly. In column 2 of Table 2 we incorporate the control variables discussed in the previous section. While the estimated coefficients for these controls have the signs and magnitudes we would expect (*e.g.*, increased hotel room supply and unemployment are both associated with decreased hotel room revenue), our estimates for the impact of Airbnb remain unchanged.

The economic significance of our estimates is best understood in the context of Airbnb's growth. For instance, in Austin, the city in Texas with the highest Airbnb penetration, we estimate that the present impact of Airbnb is 13% of hotel room revenue (based on present cumulative Airbnb supply, which exceeds 8,500 listings – see Table 1). Considering the high fixed costs associated with operating a hotel, this figure could represent a significant fraction of hotel profits. These currently substantial impacts are all the more striking in light of the fact that Airbnb continues to grow rapidly, including in cities like Austin, where it already has an established presence. Furthermore, larger markets in Texas such as Houston

[9] See "How Airbnb helps users save their homes", August 2012, `http://fortune.com/2012/08/16/how-airbnb-helps-users-save-their-homes/`

and Dallas appear to have ample room for Airbnb growth, due to their large population, and relatively low Airbnb penetration to date. Therefore, our results suggest the risk to incumbent hotels from Airbnb as a market entrant is both measurable and increasing.

An alternative way to assess the economic significance of these results is by comparing the estimated coefficients for the impact of increased Airbnb supply and hotel room supply on hotel room revenue. Our results in the second column of Table 2 show that a 10% increase in the supply of hotel rooms in Texas is associated with a roughly 0.41% decrease in Texas hotel room revenue, while a corresponding increase in Airbnb supply is associated with a smaller 0.35% decrease. It makes intuitive sense that increasing Airbnb supply has a smaller impact than an equal increase in hotel room supply, as we do not expect all Airbnb stays to substitute for a hotel room stay. Nevertheless, the two effects are surprisingly comparable in size, suggesting that Texas Airbnb inventory is not that different from incremental hotel inventory. We explore substitution between Airbnb and specific hotel types in much more detail in Section 4, where we seek to understand the mechanisms behind Airbnb's impact.

3.3 Robustness checks

To further reinforce the causal interpretation of our DD estimate, in this subsection we perform two additional checks: a matching method, which we use as a more stringent alternative in defining (otherwise similar) treated and untreated properties, and a specification test using an alternative functional form of Airbnb supply.

Since Airbnb adoption is clearly not random by design, to provide evidence in support of the DD identification assumptions, we showed that observed pre-treatment demographic and market characteristics do not correlate with the patterns of Airbnb adoption we observe in our data, which is what we would expect with exogenous Airbnb entry. Here, we combine DD with matching to further limit the potential for unobserved confounders biasing our estimates. To explain the matching approach, first recall our source of identification: roughly speaking, for each "treated" hotel, *i.e.*, a hotel affected by Airbnb competition, our DD analysis constructs a counterfactual outcome using a set of "untreated" hotels, *i.e.*, hotels unaffected by Airbnb. The intuition behind matching is that the more similar treated and untreated hotels are in their observed characteristics, the less likely they are to differ in unobserved ways, including bias-inducing factors. Matching methods aim to reduce endogeneity concerns by ensuring comparability between treated and untreated units (Heckman and Navarro-Lozano 2004). While various matching methods exist, here we use the Coarsened Exact Matching (CEM) procedure (Iacus et al. 2012), because it is intuitive and works well with categorical data (like most hotel characteristics).

CEM takes places in two steps. First, hotels are stratified based on observed characteristics; we use price segment (Budget to Luxury), operation (independent or chain), and hotel chain affiliation (*e.g.*, Hilton, or Marriott), if any. After this first step, each stratum contains hotels that are identical on the basis of these characteristics. For instance, a single stratum contains all Upscale Marriott hotels, some of which are eventually treated and some of which are not. In a setting with a binary treatment indicator, it is clear which units are eventually treated. In our case, where treatment intensity varies, we make the distinction between treated and untreated hotels by defining hotels in cities which see no Airbnb penetration by the end of our observation period as untreated, and the remaining hotels as treated. One could argue that this definition of treatment is too permissive; while we do not present these results for brevity, we found our CEM analysis to be robust to alternative definitions of treated units, such as hotels in cities that eventually have at least 100 Airbnb listings. In the second step of CEM, we discard strata containing only treated or untreated hotels, and re-normalize weights of observations in the remaining strata to place equal weight on treated and untreated units in each stratum. Applying CEM to our data leaves us with $3,001$ hotels.[10] Finally, we re-estimate the DD specification in Equation 1 on the subset of matched hotels using the CEM weights. Conceptually, DD on the CEM sample estimates a treatment effect within each stratum of comparable treated and untreated hotels, then averages these treatment effects to arrive at a final estimate. We report this estimate in the third column of Table 2. We find that the effect of Airbnb on hotel room revenue is robust to CEM, attaining a magnitude ($-0.037, p < .01$) that is nearly identical to our main analysis.

The second robustness check we perform guards against a functional specification concern in Equation 1: regressing the log of Airbnb supply on the log of hotel room revenue implicitly assumes a constant elasticity relationship between the two quantities. While this might be a reasonable assumption in data with limited variation in Airbnb supply, the constant elasticity assumption is likely violated in our setting, as it is implausible that doubling Airbnb supply from 1 to 2 units will have the same effect on hotel room revenue as doubling Airbnb supply from 100 to 200 units. To ensure that our results are not driven by this modeling choice, we model Airbnb supply non-parametrically using a categorical variable, which takes on one of the following (roughly log-binned) values: 0 Airbnb units, 1-99 Airbnb units, 100-999 Airbnb units, 1000+ Airbnb units. This model allows for the effect of Airbnb to vary depending on the number of Airbnb listings present in each city during a given period. We estimate this model by replacing Airbnb supply with this new categorical variable in Equation 1 using zero

[10] CEM entails a trade-off between matching granularity, and the number of discarded observations. We chose our matching criteria to strike a reasonable balance between ensuring units within each stratum are similar, and discarding too many observations. Our results our robust to alternate matching criteria.

Airbnb units as the reference level. We present our results in the fourth column of Table 2. We find that increasing levels of Airbnb penetration have proportionally larger impacts on hotel room revenue, as we would expect. For example, at Airbnb adoption rates exceeding 1000 rooms, the estimate $(-0.087, p < .01)$, indicates (since we are now working with a log-level specification) an average impact of 8.7% on hotel room revenue. Moreover, it is also reassuring that we find no statistically significant effect at low levels of Airbnb supply. This robustness check suggests that we are not identifying the Airbnb treatment effect from variation at low rates of Airbnb supply, which one would expect to have a negligible impact on hotel room revenue.

3.4 How Affected Hotels Respond to Airbnb

We now turn to the question of responses by incumbent hotels to Airbnb market entry. Patterns of response and novel response mechanisms by traditional incumbents to entrants facilitated by online technology is of increasing focus Seamans and Zhu (2013), Kroft and Pope (2014); we add to this literature.

We investigate whether hotels actively respond to Airbnb market entry through a price response. Recall that hotel room revenue is the product of two quantities: average occupancy rate within a given time period, and average daily room price (ADR) during that same period of time. We now separate the impact from Airbnb on hotel room revenue into two components: the effect due to reductions in occupancy and the effect due to pricing, as captured by changes in the average daily rate. A hotel that exerts no response to a supply shock would exhibit a reduction in occupancy, whereas alternatively, a manager could maintain occupancy levels via a price response. A key difference between the two responses is that the latter, reduced prices, is a net benefit for all consumers, whether they use Airbnb or not.

To estimate these component-wise effects, we re-estimate the DD specification in Equation 1, substituting the dependent variable first with occupancy rate, and then with (the log of) ADR. Similar to the room revenue analysis, these two quantities vary by hotel and by month. We report these results in the final two columns of Table 2. As reported in the fifth column of this table, we find no statistically significant connection between increased Airbnb listings and occupancy rate. However, in column 6, when we regress against ADR, we find that a 10% increase in Airbnb supply is associated with a statistically significant ($p < 0.01$) price decrease of 0.2%. This suggests that affected hotels actively responded to Airbnb entry by lowering their prices. To understand the economic significance of these results, we note that at present Airbnb supply levels, our estimates imply an 8% decrease in Austin hotel

prices, where competition by Airbnb penetration is most intense.[11]

4 Mechanisms and Market Response

4.1 Which hotels are most affected and why?

Previously, we provided evidence that, when treating hotels homogeneously, Airbnb has a negative impact on hotel room revenue in Texas. In this section, we investigate various mechanisms, through which Airbnb could exhibit heterogeneous impacts across different types of hotels. To motivate this analysis, we observe that while Airbnb can surely *sometimes* provide an alternative to hotels, one can hardly expect it to be a perfect substitute for all travel needs. As Airbnb has its roots in casual stays, including those involving shared accommodations, we expect it to be a less attractive option for those who are not on a budget. Specifically, business travelers whose hotel expenses will be reimbursed, or vacationers who frequent high-end hotels, are two examples of consumers we view as much less likely to substitute a hotel stay with an Airbnb stay. Following this logic, we further isolate the impact of Airbnb on hotel room revenue by partitioning hotels in two different ways, each dividing hotels into one class that is arguably less vulnerable to Airbnb's entry and another class that is more vulnerable, then estimating this additional interaction effect in our original DD specification. First, we segment hotels by price tier. Recall that the STR hotel census divides hotels into five price tiers: Budget, Economy, Midprice, Upscale, Luxury. Second, we differentiate hotels by their target customers. With most business travel traditionally arranged through employers, an Airbnb stay may be more difficult to substitute for these type of stays. Moreover, business travelers make greater use of those business-related hotel amenities not typically provided by Airbnb properties.

To estimate heterogeneous treatment effects, we estimate a new specification that adds an interaction effect between hotel types and Airbnb supply to the DD specification in Equation 1:

$$\begin{aligned} \log \text{Hotel Revenue}_{ikt} = & \beta_1 \log \text{ Airbnb Supply}_{kt} \\ & + \beta_2 \log \text{ Airbnb Supply}_{kt} \times \text{Hotel Type}_i \\ & + X'_{ikt}\gamma + \alpha_i + \tau_t + \epsilon_{ikt}. \end{aligned} \tag{3}$$

The coefficients of interest are β_2, which captures the differential impact of Airbnb on the

[11] An 8% price impact is not inconsistent with the 13% impact on room revenue we estimated in Section 3.2 for several reasons, e.g., the latter estimate is based on the subset of hotels reporting data to STR.

various segmentations by hotel type that we investigate. Specifically, following the segmentations described above, we first define Hotel Type$_i$ as a categorical variable identifying each one of the hotel price segments used by STR. In the second analysis we use a binary indicator of whether hotel i has conference or meeting space.

The results of these analyses appear in first two columns of Table 3. We start with price segmentation, presented in the first column. We estimate Equation 3, interacting hotel price segments with Airbnb supply. Based both on price comparison and the fact that upmarket hotels provide amenities (*e.g.,* pools, conference rooms, concierge) to travelers that typical Airbnb rentals do not, motivates our use of Luxury hotels as a reference category (though this choice does not affect our results.) We find the negative impact of Airbnb increasing as we step down price tiers, with statistically significant interaction coefficient estimates at the 1% level for each of the three lowest tiers (Midprice, Economy, and Budget). In contrast, we find only a small negative effect, not statistically significant, for the Upscale and Luxury segment (the latter being the reference level, and hence being captured by the main effect). From a managerial standpoint, this result has direct import: even though lower-end hotels in Texas account for a disproportionately small amount of room revenue as compared with upmarket hotels, they nevertheless bear the brunt of the impact of the market entry of Airbnb. Our evidence suggests that consumers are increasingly substituting Airbnb stays for lower-end hotels in Texas, possibly identifying the former as offering better value at a similar price point. While this increased competition affords consumers greater choice, it also places lower-end hotels in regions with high Airbnb penetration at greater risk.

In column two of Table 3 we report the results of the segmentation of hotels catering to business travelers. We use those hotels having conference and meeting space as the reference category. The estimated coefficient β_2 for the interaction between Airbnb supply and the absence of meeting space indicator is negative and statistically significant ($-0.016, p < .01$), suggesting that hotels lacking business facilities are more affected by Airbnb. These results are consistent with Airbnb's marketing strategy thus far, which has primarily targeted vacation travel. However, seeing a growth opportunity in the business travel segment, Airbnb recently launched an initiative to attract more business travelers.[12] An interesting open question going forward is the extent to which business travel will continue to differentiate the impact of Airbnb on hotels.

A separate distinction that we explore, relating to hotel operation rather than consumer behavior, is between chain hotels (including franchises) and independent hotels. Unlike independent hotels, chain hotels allocate large marketing budgets to advertising, brand building, guest loyalty programs, and other tactics which should make them less vulnerable to com-

[12]See: `http://bits.blogs.nytimes.com/2014/07/28/airbnb-expands-into-business-travel/`

petition. In addition, chains provide a more predictable standard of service, which further differentiates them from both Airbnb and independent hotels. We present this analysis in the third column of Table 3, using chain hotels as a reference level. The overall effect due to Airbnb remains negative and statistically significant ($-0.031, p < .01$), suggesting that hotels of all operation structures were affected. However, the estimated interaction coefficient for independent hotels ($-0.013, p < .01$) is also negative and statistically significant, suggesting that Airbnb has indeed had a larger impact on independent hotels.

Overall, we find that independent hotels, hotels that do not cater to business travelers, and lower-end hotels are all more heavily affected by Airbnb than our respective reference categories, hotels which without these characteristics. While these results help us better understand the most vulnerable hotel segments, and are certainly of importance to hoteliers, they also serve as a robustness check in that the heterogeneous substitution effects they reveal align with the effects we hypothesized based on the value proposition to consumers that Airbnb offers.

5 Discussion and conclusion

The sharing economy has recently emerged as a viable alternative to fulfilling a variety of consumer needs, ranging from prepared meals to cars to overnight accommodations, that were previously provided primarily by firms. As the size of the sharing economy has grown, so has the magnitude of its economic impacts. Our work is among the first to provide empirical evidence that the sharing economy is significantly *changing* consumption patterns, as opposed to generating purely incremental economic activity, as argued in prior work. Focusing on the case of Airbnb, a pioneer in shared accommodations, we estimate that its entry into the Texas market has had a quantifiable negative impact on local hotel room revenue. The substitution patterns we observe strongly suggest that Airbnb provides a viable, but imperfect, alternative for certain traditional types of overnight accommodation. Our results pinpoint lower-end hotels, and hotels not catering to business travelers, as those that are most vulnerable to increased competition from rentals enabled by firms like Airbnb. Hotels in areas where Airbnb has an established presence have responded to increased competition by lowering their prices, which harms their revenue, but benefits travelers, even those who do not use Airbnb.

Our work has some limitations which could be addressed in future work. First, one must recognize that our findings are representative of the state of Texas; directly generalizing them to other markets may not be appropriate given the varying of dynamics of supply and

demand for accommodation across different regional markets.[13] Additional studies which model the impact of Airbnb across these markets could be a useful contribution. A second limitation of work is that we analyze properties listed only on Airbnb, but not properties available through related vacation rental platforms like HomeAway and VRBO. We do not believe that our results are significantly affected by these competitors, since these firms primarily serve the smaller vacation rental market; moreover, they have not experienced the extremely rapid growth of Airbnb. Nevertheless, one could investigate the impact of all of these firms in aggregate, or individually. A final limitation of our study pertains to the precise characterization of hotels' response: here we have analyzed two metrics, price and occupancy rate, that managers can invoke as a response in the short-term. On longer time scales, hotels have other ways of responding to Airbnb, including promotions, advertising, and even re-positioning to provide more personalized Airbnb-like services. Mapping out the shape of hotels' response remains an interesting open question.

5.1 Conclusion

Returning to the thesis that the sharing economy has the potential to transformatively increase social welfare, as evangelized by Botsman (2012) and others, we assert that a large population of individuals worldwide have indeed benefited from Airbnb: not only hosts that derive incremental income by renting properties through Airbnb, and guests who select an Airbnb rental as an alternative to a hotel stay, but also those consumers who benefit from lower prices and increased competition in the accommodation industry. More broadly, our results should be viewed from outside the confines of the accommodation industry. This more encompassing viewpoint can weigh the positive change the sharing economy can bring about not only by providing imperfect substitutes for existing products, but also, through an application of Say's Law, by generating demand that did not previously exist through the supply of new products and services. Harkening back to arguments Airbnb has made, supply of inexpensive accommodations can increase travel and tourism spend overall, and thus the sharing economy could be a net producer of new jobs. However, these positives must be evaluated against various costs, including those estimated in this paper. Our study represents an empirical first step into understanding the complex set of issues surrounding the sharing economy. With the projected rapid growth of the sharing economy, a host of related studies will be needed to fully understand and reap its benefits.

[13]Following our study, others have arrived at similar estimates for Airbnb's impact in different markets. For example, Credit Suisse analysts used STR data to estimate that, in New York City, Jan. 2015 revenue per hotel room was 18.6% lower than a year ago. See "New York City hotel rooms are getting cheaper thanks to Airbnb" at http://qz.com/341292/new-york-city-hotel-rooms-are-getting-cheaper-thanks-to-airbnb/.

References

Akerman, Anders, Ingvil Gaarder, Magne Mogstad. 2013. The Skill Complementarity of Broadband Internet. Tech. rep., IZA Discussion Paper.

Baker, Kathleen, Alex Coulter. 2007. Terrorism and Tourism: The Vulnerability of Beach Vendors' Livelihoods in Bali. *Journal of Sustainable Tourism* **15**(3) 249–266.

Bakos, J Yannis. 1997. Reducing buyer search costs: Implications for Electronic Marketplaces. *Management Science* **43**(12) 1676–1692.

Bertrand, Marianne, Esther Duflo, Sendhil Mullainathan. 2004. How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics* **119**(1) 249–275.

Botsman, Rachel. 2012. The Currency of the New Economy is Trust. [Video file]. `http://www.ted.com/talks/rachel_botsman_the_currency_of_the_new_economy_is_trust.html`. [Online].

Carbó Valverde, Santiago, Sujit Chakravorti, Francisco Rodriguez Fernandez. 2010. Regulating Two-Sided Markets: An Empirical Investigation. Tech. rep., Working Paper, Federal Reserve Bank of Chicago.

Cervero, Robert, Aaron Golub, Brendan Nee. 2007. City CarShare: Longer-Term Travel Demand and Car Ownership Impacts. *Transportation Research Record: Journal of the Transportation Research Board* **1992**(1) 70–80.

Donald, Stephen G, Kevin Lang. 2007. Inference with Difference-in-Differences and Other Panel Data. *The Review of Economics and Statistics* **89**(2) 221–233.

Fradkin, Andrey. 2013. Search Frictions and the Design of Online Marketplaces. Tech. rep., Working paper, Stanford University.

Hall, Jonathan V., Alan B. Krueger. 2015. An Analysis of the Labor Market for Uber's Driver-Partners in the United States. `http://arks.princeton.edu/ark:/88435/dsp010z708z67d`.

Heckman, James, Salvador Navarro-Lozano. 2004. Using Matching, Instrumental Variables and Control Functions to Estimate Economic Choice Models. *Review of Economics and Statistics* **86**(1) 30–57.

Iacus, Stefano M, Gary King, Giuseppe Porro. 2012. Causal Inference without Balance Checking: Coarsened Exact Matching. *Political analysis* **20**(1) 1–24.

Jin, Ginger Zhe, Marc Rysman. 2012. Platform Pricing at Sports Card Conventions. Tech. rep., National Bureau of Economic Research.

Kosová, Renáta, Cathy A Enz. 2012. The Terrorist Attacks of 9/11 and the Financial Crisis of 2008 The Impact of External Shocks on US Hotel Performance. *Cornell Hospitality Quarterly* **53**(4) 308–325.

Kroft, Kory, Devin G Pope. 2014. Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist. *Journal of Labor Economics* **32**(2) 259–303.

Landsman, Vardit, Stefan Stremersch. 2011. Multihoming in Two-Sided Markets: An Empirical Inquiry in the Video Game Console Industry. *Journal of Marketing* **75**(6) 39–54.

Martin, Elliot, Susan A. Shaheen, Jeffrey Lidicker. 2010. Impact of Carsharing on Household Vehicle Holdings. *Transportation Research Record: Journal of the Transportation Research Board* **2143**(1) 150–158.

O'Connor, Noëlle, Mary Rose Stafford, Gerry Gallagher. 2008. The Impact of Global Terrorism on Ireland's Tourism Industry: an Industry Perspective. *Tourism and Hospitality Research* **8**(4) 351–363.

Parker, Geoffrey G, Marshall W Van Alstyne. 2005. Two-Sided Network Effects: A Theory of Information Product Design. *Management Science* **51**(10) 1494–1504.

Rochet, Jean-Charles, Jean Tirole. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* **1**(4) 990–1029.

Rysman, Marc. 2009. The Economics of Two-Sided Markets. *Journal of Economic Perspectives* **23**(3) 125–43.

Seamans, Robert, Feng Zhu. 2013. Responses to Entry in Multi-Sided Markets: The Impact of Craigslist on Local Newspapers. *Management Science* **60**(2) 476–493.

Weyl, Eric Glen. 2010. A Price Theory of Multi-sided Platforms. *American Economic Review* **100**(4) 1642–72.

Wolfers, Justin. 2006. Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results. *The American Economic Review* **96**(5) 1802–1820.

Table 1: Airbnb's spatial and temporal penetration. Cumulative counts of Airbnb listings per year in the ten most populous Texas cities.

(Pop.)	Houston 2.16M	San Antonio 1.38M	Dallas 1.24M	Austin 0.84M	Ft. Worth 0.78M	El Paso 0.67M	Arlington 0.38M	Corpus Christi 0.31M	Plano 0.27M	Laredo 0.24M
2008	1	9	0	25	0	0	0	0	0	0
2009	6	13	7	146	2	0	1	0	0	0
2010	39	22	23	468	10	0	3	0	1	0
2011	169	72	109	1,862	34	3	19	7	5	1
2012	425	171	271	5,158	68	8	27	24	20	1
2013	695	271	422	7,489	93	23	36	49	33	1
2014	891	346	526	8,575	114	31	52	60	44	2

Table 2: Difference-in-differences estimates of the impact of Airbnb on hotel room revenue, prices, and occupancy rates. The first four columns report estimates using data from the Texas Comptroller's office; the last two from STR.

	(1) Revenue	(2) Revenue	(3) Revenue	(4) Revenue	(5) Occupancy rate	(6) Room price
log Airbnb Supply	−0.035***	−0.035***	−0.037***		−0.004	−0.020***
	(−3.70)	(−4.09)	(−4.16)		(−1.29)	(−3.40)
Airbnb Supply (ref. zero listings)						
1 to 99 Listings				−0.024		
				(−1.50)		
100 to 999 Listings				−0.056**		
				(−2.46)		
1000+ Listings				−0.087***		
				(−2.95)		
log Hotel Room Supply		−0.041***	−0.041***	−0.032***	−0.017***	−0.004***
		(−5.62)	(−5.58)	(−3.30)	(−8.96)	(−2.86)
Unemployment Rate		−0.050***	−0.050***	−0.046***	−0.017***	−0.008*
		(−4.87)	(−4.76)	(−3.97)	(−3.83)	(−1.91)
log Population		−0.029	0.022	−0.032	−0.076	0.122**
		(−0.21)	(0.14)	(−0.26)	(−1.54)	(2.23)
CEM Sample	No	No	Yes	No	No	No
N	369640	369640	283550	369640	305950	305950
R^2 within	0.23	0.23	0.23	0.24	0.30	0.37

Note: The dependent variable is log Hotel Revenue$_{ikt}$ in columns 1-4, Occupancy rate$_{ikt}$ in column 5 and log Hotel Room Price$_{ikt}$ in column 6. Cluster-robust t-statistics (at the city level) are shown in parentheses. All specifications include hotel and time fixed effects, and a city-specific quadratic time trend.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Difference-in-differences estimates of heterogeneity in Airbnb's impact on hotel room revenue.

	(1) Price segment	(2) Business travel	(3) Operation
log Airbnb Supply	−0.012	−0.028***	−0.031***
	(−1.30)	(−3.13)	(−3.72)
Price segment (ref. Luxury)			
Budget × log Airbnb Supply	−0.039***		
	(−4.95)		
Economy × log Airbnb Supply	−0.033***		
	(−10.01)		
Midprice × log Airbnb Supply	−0.017***		
	(−4.43)		
Upscale × log Airbnb Supply	−0.006		
	(−1.06)		
Business travel (ref. hotel has meeting space)			
w/o Meeting Space × log Airbnb Supply		−0.016***	
		(−3.66)	
Operation (ref. chain)			
Independent × log Airbnb Supply			−0.013***
			(−5.61)
log Hotel Room Supply	−0.041***	−0.041***	−0.041***
	(−5.64)	(−5.63)	(−5.62)
Unemployment Rate	−0.050***	−0.050***	−0.050***
	(−4.86)	(−4.85)	(−4.85)
log Population	−0.007	−0.023	−0.031
	(−0.05)	(−0.17)	(−0.22)
N	369640	369640	369640
R^2 within	0.24	0.24	0.24

Note: The dependent variable is log Hotel Revenue$_{ikt}$. Cluster-robust t-statistics (at the city level) are shown in parentheses. All specifications include hotel and time fixed effects, and a city-specific quadratic time trend.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

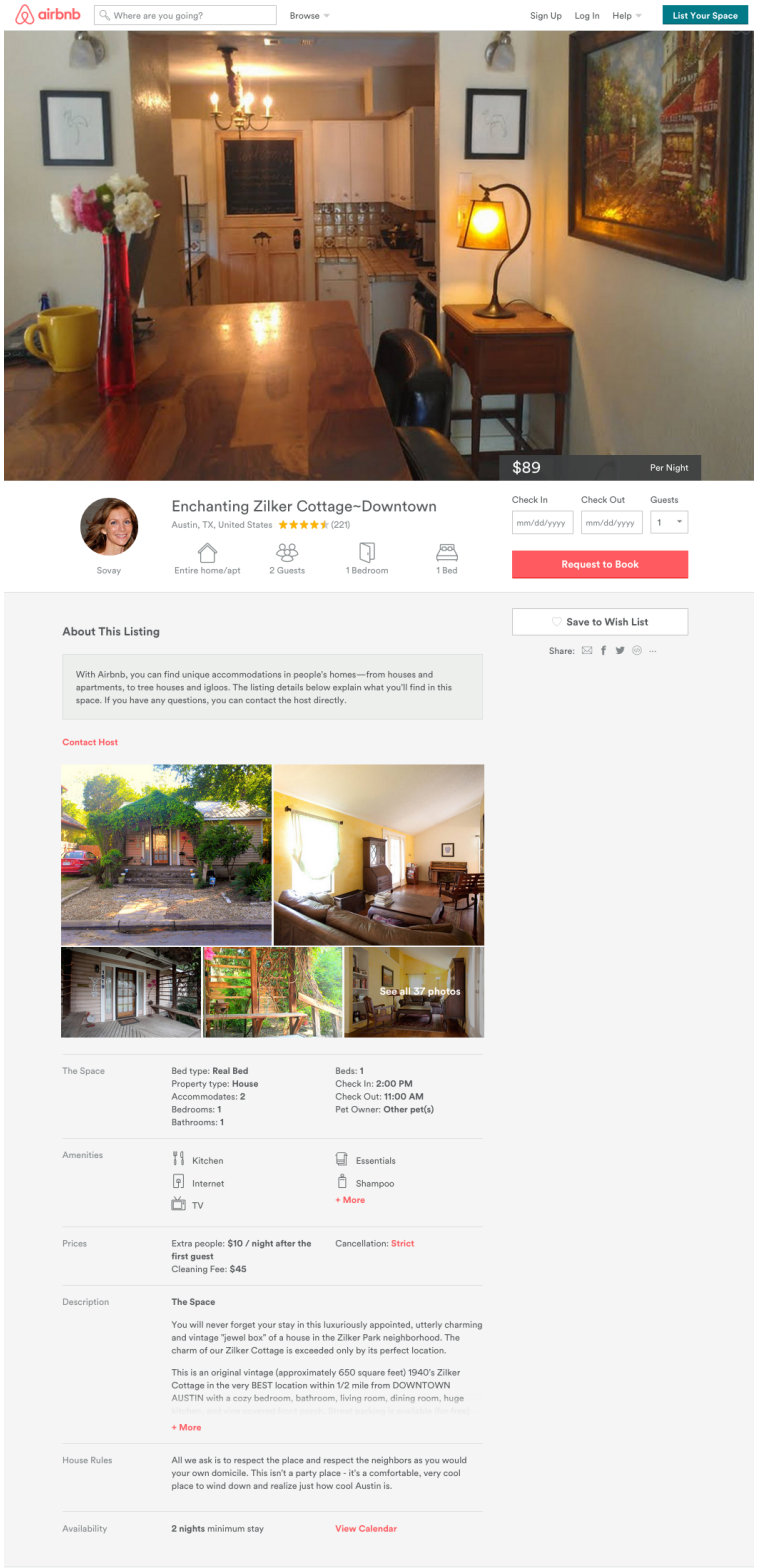

Figure 1: A typical Airbnb listing.

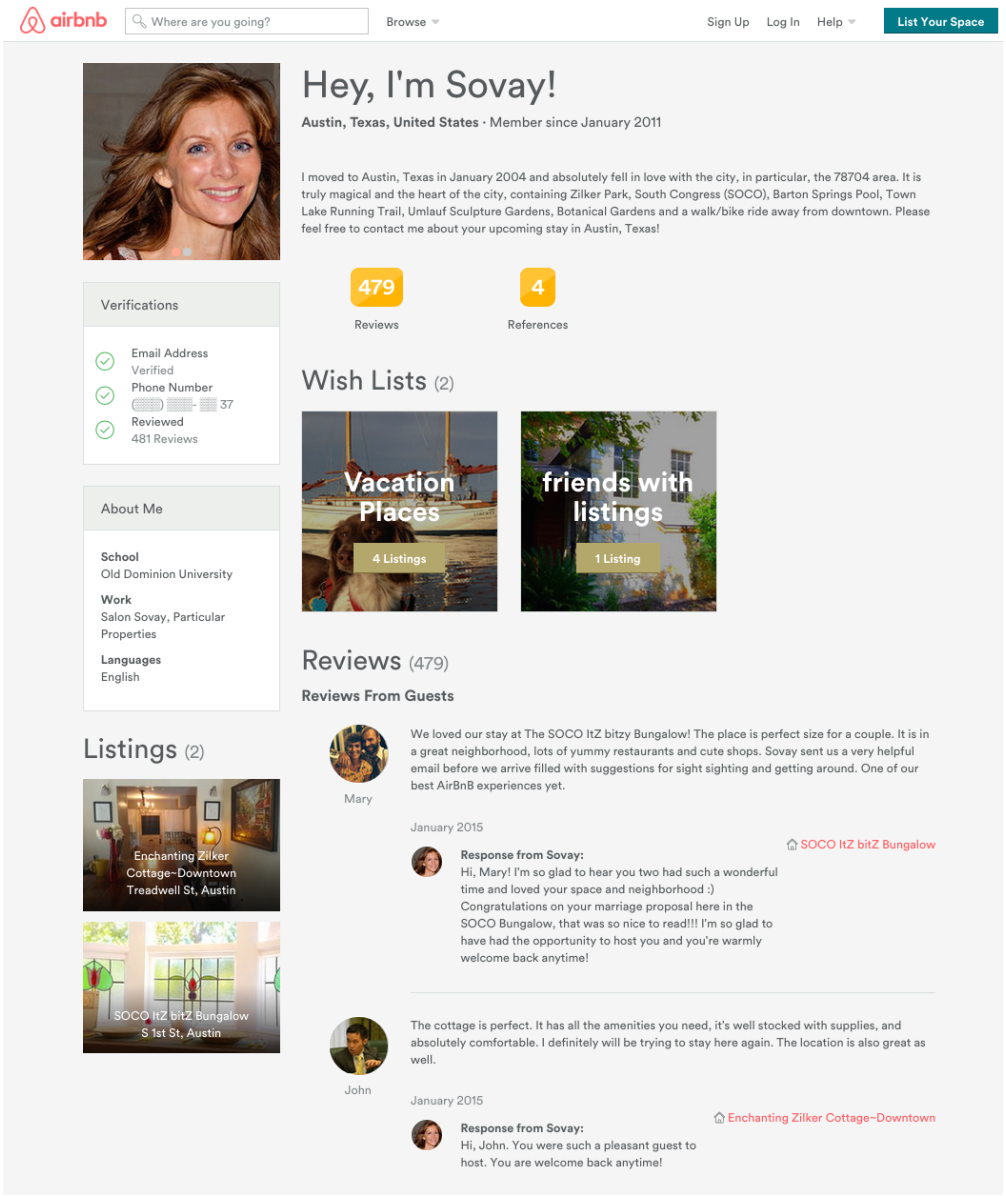

Figure 2: A typical Airbnb user profile.

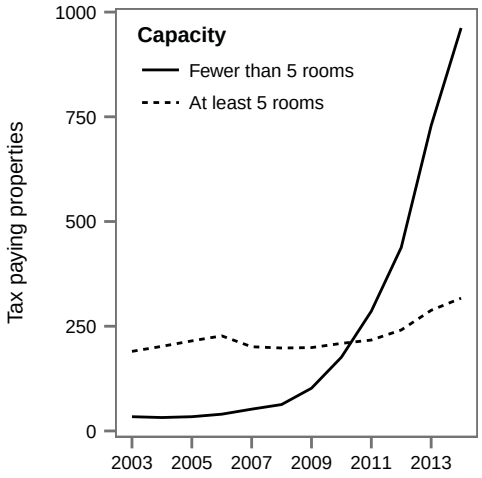

Figure 3: Annual counts of Austin properties that pay hotel occupancy tax, broken down by capacity.

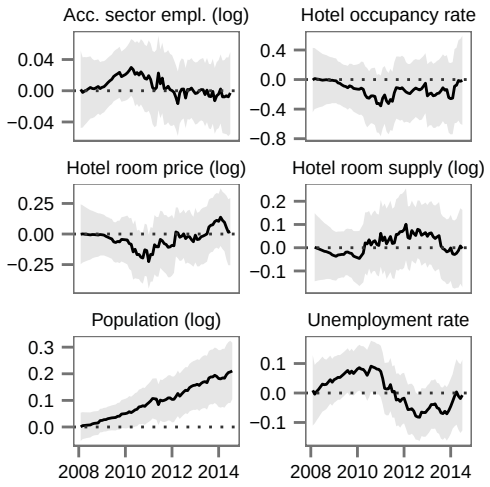

Figure 4: Correlation between Airbnb supply and pre-Airbnb (year 2007) city characteristics, with 95% confidence intervals.

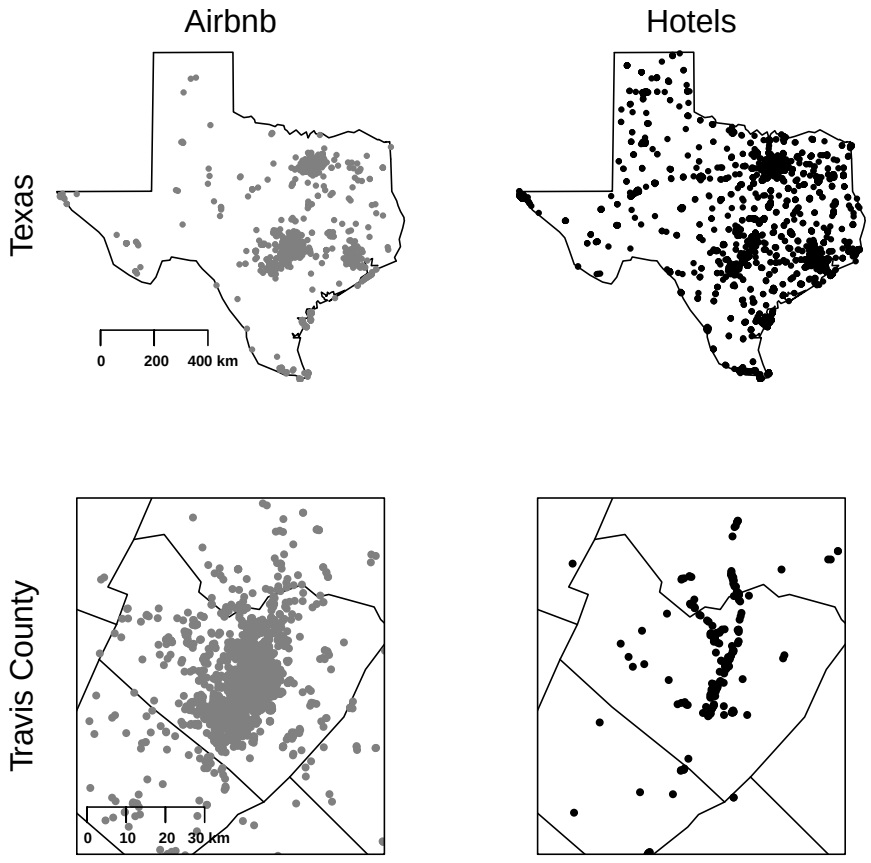

Figure 5: Geographical distribution of hotels and Airbnb listings in the state of Texas (top) and in Travis County, TX (bottom) in 2013.