University of California
Santa Barbara

Essays in Applied Spatial Economics

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Daniel J. Cullen

Committee in charge:

Professor Douglas G. Steigerwald, Chair
Professor Shelly Lundberg
Professor Andrew Plantinga
Professor Clément de Chaisemartin

September 2020

The Dissertation of Daniel J. Cullen is approved.

Shelly Lundberg

Andrew Plantinga

Clément de Chaisemartin

Douglas G. Steigerwald, Committee Chair

July 2020

Essays in Applied Spatial Economics

by

Daniel J. Cullen

To everyone who helped me along the way.

Acknowledgements

I would like to thank everyone who has supported me throughout my graduate studies. I would like to thank my parents Jim and Celesta, my brothers Matthew and Ryan, and my sister Allison. I am also thankful for the support of my friends and colleagues especially my roommates Renato, Dan, and Travis and my cohort members, James, Andrew, Juliana, Sumeyye, Nicole, Sahaab, and Shamlan. None of this would have been possible without the support of my family and friends.

I would like to thank my advisor Doug Steigerwald for his guidance throughout my entire PhD career. His advice and support helped me become the researcher I am today. I would also like to thank my committee, Andrew Plantinga, Clément de Chaisemartin, and Shelly Lundberg for valuable feedback and advice.

In addition to my committee, I would like to thank Peter Rupert and everyone at the UCSB Economic Forecast Project. I would also like to thank all of the members of the Econometrics Research Group for all of the feedback and discussion.

Curriculum Vitæ

Daniel J. Cullen

Education

2020	Ph.D. in Economics
	University of California, Santa Barbara
2014	M.A. in Economics
	University of Colorado Denver
2011	B.A. in Economics
	University of Florida

Research Interests

Applied Econometrics, Econometrics, Spatial Econometrics, Environmental Economics, Urban Economics, Health Economics

Research Experience

2018 – 2020	UCSB, Economic Forecast Project
2015 – 2016	Research Assistant – Heather Royer

Teaching Experience

Instructor – University of California, Santa Barbara

Datahack	Economics Department
	Summer 2018, Summer 2019
Data Analysis with R	Broom Center for Demography
	Fall 2019

Head Teaching Assistant – University of California, Santa Barbara

ECON 140A	Intro to Econometrics
	Winter 2017, Spring 2017, Winter 2018, Spring 2018

Teaching Assistant – University of California, Santa Barbara

ECON 140A	Intro to Econometrics
	Spring 2016, Summer 2016, Fall 2016, Summer 2017
ECON 140B	Intro to Econometrics II
	Fall 2017
ECON 1	Principles of Microeconomics
	Fall 2015, Winter 2016

Teaching Assistant – University of Colorado Denver

ECON 4071	Intermediate Microeconomic Theory Spring 2014
ECON 5073	Microeconomic Theory (Graduate Level) Fall 2013
ECON 2022	Principles of Microeconomics Spring 2013

Honors & Awards

Academic Senate Travel Grant, Spring 2019
White Graduate Fellowship, Summer 2018
Outstanding Teaching Assistant, Fall 2016
Andron Fellowship, 2014 – 2015
National Merit Scholarship, 2007 – 2011
The University of Florida Honors Program, 2007 – 2011

Abstract

Essays in Applied Spatial Economics

by

Daniel J. Cullen

This dissertation consists of three works which consider estimation of economic variables when a spatial component exists. Each essay utilizes different techniques and methodology for working with data that can be grouped into spatial clusters.

In the first essay, I estimate the impact of air pollution events caused by wildfire smoke on respiratory and circulatory health outcomes. Utilizing a combination of California health data and NOAA wildfire smoke data I can estimate the impact of exposure to wildfire smoke on health outcomes for all individuals in California. Using inpatient data I am able to construct a measure of exposure to wildfire smoke prior to the hospital visit, this allows for the identification of the impact of wildfire smoke exposure on different health outcomes. I find that an additional day of smoke exposure in a month leads to on average 11.38 additional hospital admissions for respiratory diagnoses and an additional 3 hospital admissions for circulatory diagnoses. This translates to an annual cost of wildfire smoke exposure in California due to respiratory and circulatory hospital admissions of $192,316,498.

The second essay, joint with Travis Cyronek, asks the question: How does the sharing economy affect traditional lodging markets? The advent of platforms such as *Airbnb* in 2008 has introduced a new channel of market interaction between those with space and those who seek it. This allows for transactions of lodging services that might otherwise be underutilized. This paper develops a framework to help think about how peer-to-peer transactions interact with traditional rental markets,

and what this means for property managers and tenants. Specifically, we examine how the introduction of sharing platforms (e.g. *Airbnb*) affect the listing decisions of vacant property managers and the lodging choices of dwelling seekers. The model features landlords who choose where to list vacant properties and renters who search for lodging. Renters can be either short or long-term, referencing how long they wish to occupy the property. Sharing platforms give landlords the option of accessing these short-term renters whom would otherwise occupy hotels, affecting traditional, long-term renters. We find that *Airbnbs* decrease hotel prices by about $24 while they increase average rents by $39 per month.

In the third essay, joint with Douglas G. Steigerwald, we study the behavior of cluster-robust test statistics in models with instrumental variables when cluster heterogeneity is present. Inference in a large number of papers using two-stage least squares regressions published in American Economics Association journals are driven by the presence of one or two influential clusters. We link a measure of cluster heterogeneity, the feasible effective number of clusters, to measures of influence. Using simulations, we demonstrate that high levels of cluster heterogeneity lead to coverage of less than 95% for 95% confidence intervals when using instrumental variables with panel data or with data that can be grouped into clusters. Using data from papers with two-stage least squares regressions published in American Economic Association journals, we show that the feasible effective number of clusters can be used as a pre-test to the sensitivity of two-stage least squares inference to influential clusters. We further show that when the feasible effective number of clusters is small, even when the number of clusters is large, the distribution of the test statistic in non-normal. When this severe cluster heterogeneity is present, the restricted wild cluster bootstrap can be used to return coverage to the appropriate level.

Contents

Chapter 1

The Health Cost of Wildfire Smoke

1.1 Introduction

The annual area burned by wildfires in the western United States has increased substantially in recent decades due to the increased frequency and size of wildfires (Westerling (2016)). In addition to the direct damages of wildfires, they also produce large amounts of smoke that can spread over thousands of miles and persists for days or weeks. The smoke plumes produced by wildfires contain large amounts of particulate matter (PM) and other toxic gases. The increase in wildfire smoke may be a contributing factor in the recent increase in the levels of air pollution in the United States. Clay and Muller (2019) find that wildfires may account for some of the observed increase in $PM_{2.5}$[1] from 2016 to 2018. The increase in air pollution exposure, specifically $PM_{2.5}$, increases mortality in vulnerable populations such as infants and the elderly (Chay and Greenstone (2003); Currie and Neidell (2005); Deryugina et al. (2019)) and increases hospitalization costs for respiratory and heart-related admis-

[1] $PM_{2.5}$ refers to fine atmospheric particles that have a diameter of less than 2.5 micrometers. The Environmental Protection Agency (EPA) sets and reviews national air quality standards for PM as part of the Clean Air Act.

sions (Schlenker and Walker (2016)).

Wildfire smoke not only releases $PM_{2.5}$, but also a wide range of pollutants including greenhouse gases, photochemically reactive compounds, and coarse particulate matter (Urbanski et al. (2008)). Past research on the impact of wildfire smoke has mainly focused on individual wildfire events and the communities nearby these fires, such as Haikerwal et al. (2015) who look at the 2006 wildfire in Victoria, Australia. Haikerwal et al. (2015) finds that $PM_{2.5}$ exposure was associated with increased risk of out-of-hospital cardiac arrests and coronary heart disease (CHD). Their results suggest that increased levels of $PM_{2.5}$ as a result of wildfires may act as a triggering factor for acute coronary events. Miller et al. (2017) were the first nationwide study of the impact of wildfire smoke on health. Using Medicare data and data on wildfire smoke plumes, Miller et al. (2017) found that exposure to wildfire smoke significantly increases mortality risk in the elderly.

In this paper, I utilize data from California's Office of Statewide Health Planning and Development (OSHPD) to look at the impact of wildfire smoke exposure on hospital admittance for respiratory and circulatory illnesses. Unlike previous studies that have utilized Medicare data, these data allow for the study of health outcomes for a broader demographic group. In addition, looking at hospital visits, not just mortality, allows for a more complete view of the impacts of wildfire smoke on health. This paper also utilizes data through 2018 which includes three of the largest fires in California history: the Mendocino Complex fire (2018), the Thomas fire (2017), and the Carr fire (2018).

To study the relationship between wildfire smoke exposure and hospital admissions, I link the hospital admittance data with satellite imagery data produced by The National Oceanic and Atmospheric Administration's (NOAA) Hazard Mapping System. These data track the location and movement of all wildfire smoke plumes in the

United States. I use this to derive the daily wildfire smoke exposure of every county in California from 2012 through 2018. Similar to the identification strategy in Miller, Molitor, and Zou (2017), I am able to utilize the quasi-random variation in exposure to wildfire smoke due to the nature of the drift in the wildfire smoke plumes. Using this year to year variation in whether a specific area was exposed to wildfire smoke at a specific point in time, I am able to identify the causal impact of wildfire smoke exposure. I estimate that each additional day of smoke exposure in a month leads to on average 5.56 additional hospital admissions that month for respiratory diagnoses and an additional 1.84 hospital admissions for circulatory diagnoses. In addition, each additional day of smoke exposure in the previous month leads to on average 3.78 additional hospital admissions for respiratory diagnoses and an additional 1.19 hospital admissions for circulatory diagnoses in that county. The total cumulative effect of an additional day of smoke exposure in a month leads to on average 11.38 additional hospital admissions for respiratory diagnoses which equates to a 2.78% increase in the average month. The cumulative effect on circulatory diagnoses is 3 hospital admissions which equates to a 0.41% increase in the average month in that county.

This paper also contributes to the literature on the costs of climate change. Under climate change scenarios it is predicted that wildfires will increase in area burned and fire intensity (Flannigan et al. (2000)). Not only will these fires have direct impacts on the surrounding communities, the smoke plumes produced by these fires will affect individuals far from the wildfire. Using estimates of the cost of illness (COI) for a hospital admission for respiratory or cardiovascular illness produced by The United States Environmental Protection Agency (EPA) I calculate that an additional day of wildfire smoke exposure leads to about $188,767 additional medical expenditures for respiratory and circulatory hospital admissions. This equates to a total annual cost of wildfire smoke exposure in California due to respiratory and circulatory hospital admissions of $193,903,107, or approximately 0.066% of California's annual health care spending.[2]

Section 1.2 provides background on wildfire smoke and the impacts of wildfire smoke and air pollution on health and productivity. Section 1.3 describes the data used and other potential data sources. Section 1.4 explains my empirical strategy and main results. Section 1.5 estimates the economics costs of wildfire smoke due to hospital admissions. Section 1.6 concludes.

1.2 Background

1.2.1 Wildfire Smoke

Wildfires can release large amounts of particulate matter (PM) and toxic gases including carbon monoxide (CO), nitrogen oxides (NOx) and non-methane organic compounds (NMOC) into the atmosphere. These compounds can then spread over hundreds of miles and persist for days or weeks in smoke plumes(Reisen et al. (2015)). Figure 1.1 shows an example of this spread during the Thomas Fire, a fire that began in Ventura County, California on December 4, 2017. This fire spread into neighboring Santa Barbara County and burned approximately 280,000 acres before being fully contained on January 12, 2018. The grey shaded regions shows the extent of wildfire smoke on December 9, 2017 during the first week of the fire. I produce these representations using the data produced the National Oceanic and Atmospheric Administration's (NOAA) Hazard Mapping System (HMS) data set. The shading of the plumes represents the relative density of the plume in the atmosphere, but does not reflect the concentration at ground level. The counties shaded red are the areas that

[2]California annual health care spending comes from the U.S. Centers for Medicare & Medicaid Services.

are exposed to any wildfire smoke on that day. In addition to the counties along the Pacific Coast exposed to the Thomas Fire smoke plume, we can also see smaller wildfire events in other areas of the U.S. on the same day. The smoke plume off the West Coast of the United States shows the large extent a single wildfire event can have. However, this wildfire event is one of the largest wildfires in U.S. history and most wildfire events are not on this scale. The majority of wildfires are more similar to the other wildfires seen around the U.S. in the Figure 1.1, only burning around 100 acres on average[3].

Figure 1.1: Map of smoke plume for December 9, 2017 during the Thomas Fire with smoke affected counties colored in red. The fire burned approximately 280,000 acres before being fully contained over a month later.

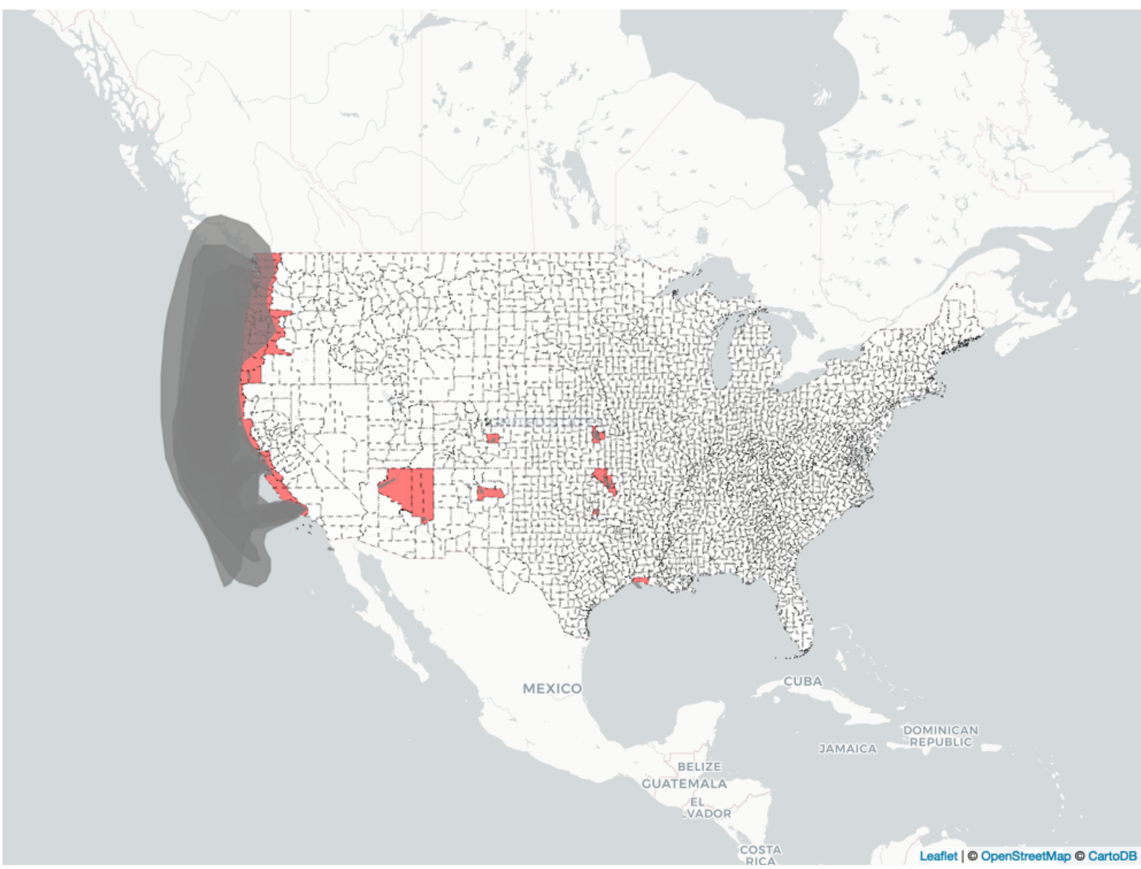

[3]Source: National Interagency Fire Center (NIFC)

While the location of wildfires is not exogenous to health outcomes, the spread of the smoke plume is mainly driven by the direction of the wind. Figure 1.2 also exemplifies the nature of the spread of wildfire smoke plumes. Exposure to the smoke is driven by shifts in wind patterns not on distance to the fire. Comparing the panels in Figure 1.2 we can see that exposure to wildfire smoke from one fire can vary on a day to day basis and that the distance to a fire is not the only important factor. This characteristic, the combination of fire ignition location and day to day wind patterns, creates exogenous variation in smoke exposure. It also addresses concerns about selection as predicting the location of wildfires and the direction of the spread of the smoke plume in advance seems unlikely.

The existing literature has found mixed results when looking at the relationship between wildfire smoke and morbidity. Liu et al. (2015) review fourteen studies that investigate the link between wildfire smoke and cardiovascular morbidity. Of those fourteen, only six reported a positive relationship. The strongest evidence for the negative effects of wildfire smoke on health come from Miller et al. (2017). Their nationwide study using Medicare administrative data shows that wildfire smoke significantly increases the mortality risk of the elderly.

Figure 1.2: The spread of wildfire smoke during the first 6 days of the Thomas fire which began on December 4, 2017 in Ventura County, California and lasted until January 12, 2018.

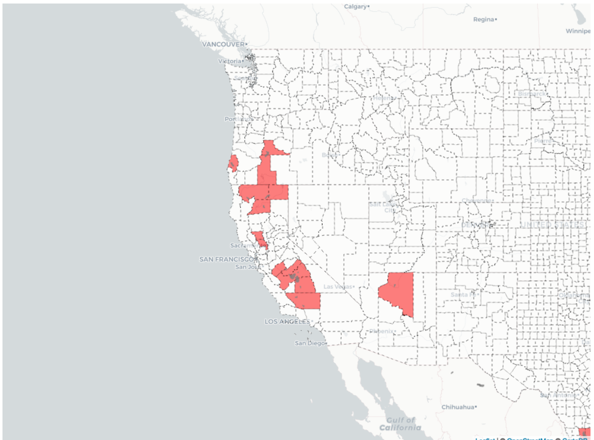

(a) December 4, 2017

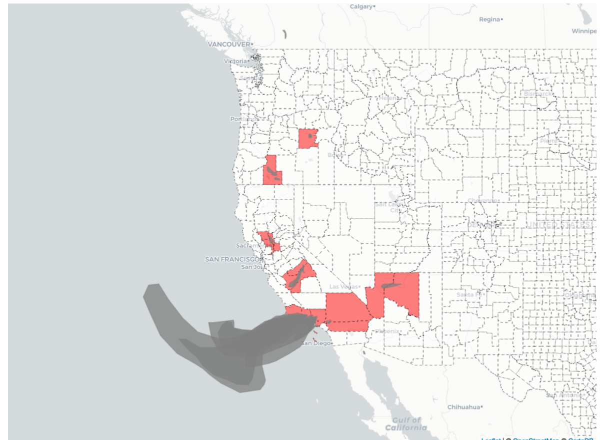

(b) December 5, 2017

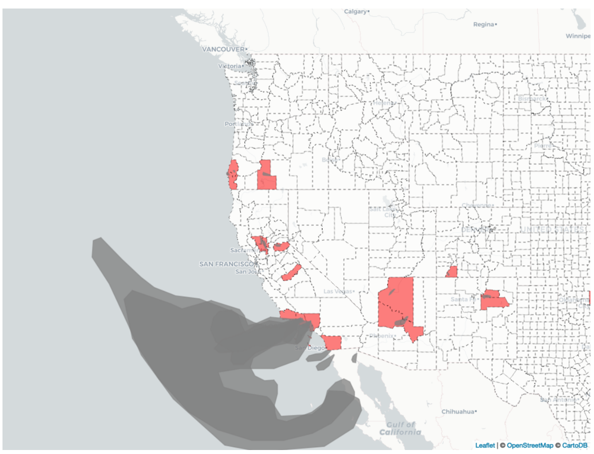

(c) December 6, 2017

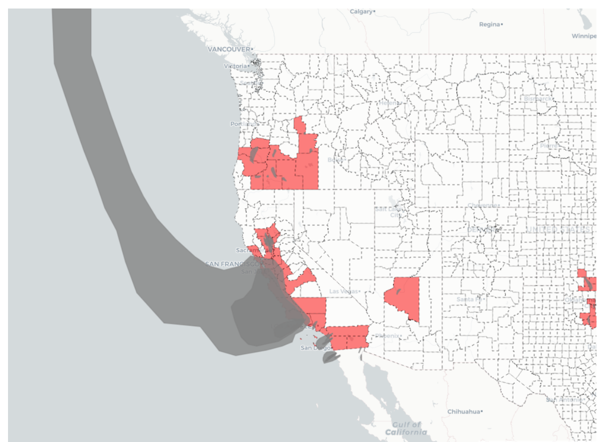

(d) December 7, 2017

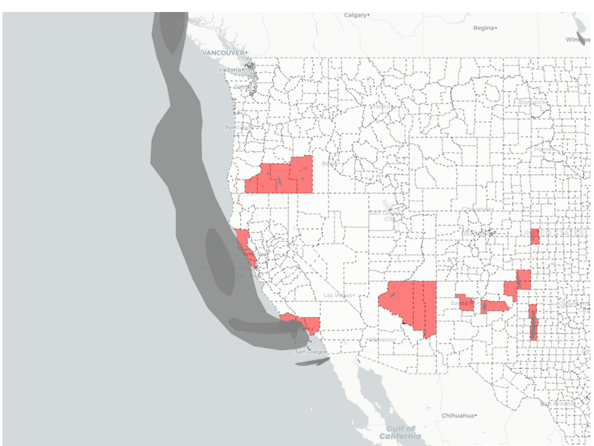

(e) December 8, 2017

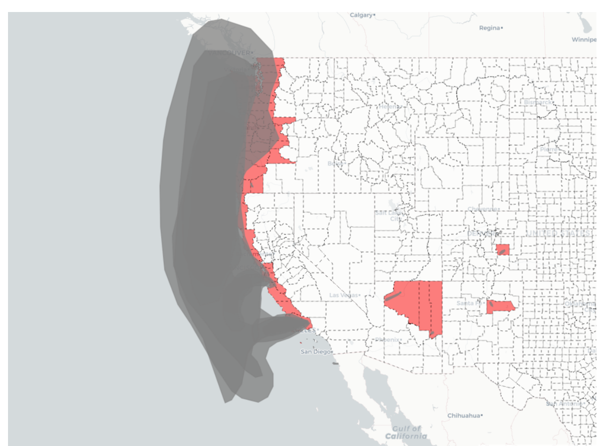

(f) December 9, 2017

1.2.2 Impact of Air Pollution on Health

While the number of studies investigating the impact of wildfire smoke exposure on health is relatively few, the relationship between air pollution and health continues to be a highly studied and debated topic. The impact of pollution on health is difficult to identify because pollution is not randomly assigned. The health of individuals who live in areas of high pollution may be different for reasons unrelated to pollution such as a preference for healthy activities. Many studies have looked at the impact of air pollution on infant health (Chay and Greenstone (2003);Currie and Neidell (2005)). Deryugina et al. (2019) were the first to conduct a large-scale, quasi-experimental investigation of the effects of short-term $PM_{2.5}$ exposure on elderly mortality, health care use, and medical costs. Schlenker and Walker (2016) develop a framework for estimating the contemporaneous effect of air pollution on health using variation in local air pollution driven by airport runway congestion. They find that an increase in the daily level of pollution leads to an increase in hospitalization costs for respiratory and heart-related admissions.

Lelieveld et al. (2020) estimate that global air pollution caused an additional 8.8 million premature deaths in 2015. These deaths represent an average shortening of life expectancy of nearly three years for all persons worldwide.

1.2.3 Impact of Air Pollution on Labor Productivity

In addition to the health effects, the impact of air pollution on productivity has been well documented. For example, Hanna and Oliva (2015) show the reduction in hours worked due to changes in air pollution in Mexico City. Chang et al. (2016) study the negative effect of $PM_{2.5}$ on productivity of pear packers.

Borgschulte et al. (2018) estimate the labor market impacts of ambient air pollution due to wildfire smoke exposure. They find that smoke exposure reduces earnings in the year of exposure and the following year, lowers labor force participation, and increases Social Security payments.

1.2.4 Identification

Identification in papers looking at the impact of air pollution on health utilize the variation created from ‘’natural experiments” to determine the impact of air pollution on health outcomes. Chay and Greenstone (2003) use the variation created by the implementation of the Clean Air Act of 1970. The Clean Air Act set limits to maximum allowable concentration of total suspended particulates (TSPs) that every county was required to meet[4]. The Clean Air Act led to larger reduction in TSPs in some counties than others, creating variation in the change in air pollution. This variation arose because industrial emitters of TSPs in counties that were above the legislated maximum were subject to stricter regulations.

Chay and Greenstone show that TSPs did decline in the early 1970s and the decline was entirely in nonattainment counties. They then use nonattainment status as an instrumental variable for changes in TSPs in 1971 to 1972 to identify the impact of TSPs on infant mortality. Nonattainment status may be a valid instrument because it is unlikely that a federal limit on emissions would impact county level infant mortality rates except through its impact on air pollution. They find that a one percent reduction in TSPs leads to a 0.5 percent decline in the infant mortality rate, with most of the reduction in the first month after birth.

[4]Total suspended particulates were defined by the EPA to include all particles with diameters less than or equal to 100 micrometers. Federal regulations that focused on PM_{10} and $PM_{2.5}$ were not established until 1987 and 1997, respectively.

Currie and Neidell (2005) use individual level data linked to weekly ZIP Code level pollution measures to identify the impact of air pollution on infant health in California in the 1990s. This detailed individual level data allows the authors to control for both postnatal and prenatal pollution exposure, the age of the child, observable characteristics of the mother and child. The authors also include month, year, and zip code fixed effects so the impacts of pollution are identified using variation within zip code, month, year level cells. They find that reductions in carbon monoxide during the 1990s saved the lives of approximately 1,000 infants in California.

Deryugina et al. (2019) use administrative data on Medicare beneficiaries combined with daily pollution data to evaluate the impact of $PM_{2.5}$ on health by utilizing variation in pollution due to changes in daily wind direction. The assumption that the authors make is that, after controlling for fixed effects and climate variables, variation in a county's daily wind direction does not affect a county's mortality or health care use except through a change in the level of air pollution. They estimate that an increase in exposure to $PM_{2.5}$ by 1 microgram per cubic meter leads to 0.69 additional deaths per million elderly individuals the three days following the increase.

Looking more specifically to studies evaluating the effects of wildfire smoke, a recent paper utilized a difference-in-differences model to study the impact of wildfire smoke exposure on birthweight and the probability of low birthweight in Colorado. Mccoy and Zhao (2020) use this approach to compare birthweights of infants born to mothers inside an area exposed to wildfire smoke plumes during pregnancy to infants located outside of the smoke plume. The authors estimate that infants exposed to wildfire smoke have a 0.034 increase in the probability of low birthweight.

Similar to Deryugina et al. (2019), Miller et al. (2017) utilize Medicare beneficiaries data but similar to this paper investigate the impact of exposure to wildfire smoke on health outcomes. The authors utilize variation in air quality as a results of drifting smoke plumes, specifically using year-to-year variation in smoke coverage of a particular area at a specific time of year. The authors include ZIP Code and week of year fixed effects to ensure the comparison is done within each ZIP Code in the same week across years with different levels of smoke exposure. The authors also include day of week indicator variables to account for differences in hospital admissions between weekends and weekdays. They estimate that the mortality rate increases by 0.522 deaths per million Medicare beneficiaries on the day of exposure and 1.204 deaths per million in a three day period following the exposure.

Similar to Miller et al. (2017), I utilize variation in year-to-year exposure to wildfire smoke plumes in a particular area. However, using Medicare data Miller et al. (2017) can only look at mortality rate of the elderly while I utilize data for all hospital admissions in California. The available admissions data is aggregated to the monthly level, so including day of week fixed effects is not necessary but month fixed effects are utilized to control for seasonality of both wildfire smoke and health outcomes. I also include county level fixed effects so that the wildfire smoke exposure is being compared in the same area across years. In addition to fixed effects, I include yearly level county demographic controls that impact healthcare utilization such as the unemployment rate, the percent of the population in poverty, and the percent of the population with health insurance. These controls capture similar information as year by County fixed effects[5].

[5]Including year by County fixed effects do not change the estimated coefficients in Section 4.

1.3 Data Sources

1.3.1 Patient Discharge Data

Patient discharge data come from California's Office of Statewide Health Planning and Development (OSHPD). These data provide the number of cases for several major diagnostic categories, Circulatory System, Respiratory System, Infections, Injuries and Poisonings, and All Other[6] at the monthly level. These data are aggregated to the county level for larger counties and grouped into multiple county units within an air basin[7] for smaller counties due to privacy restrictions[8]. The groupings are displayed in Figure 1.4. Including controls for seasonality will be important, Figure 1.3 demonstrates that respiratory system cases tend to peak in the winter months.

1.3.2 Hazard Mapping System Fire and Smoke Product

The smoke coverage data comes from the National Oceanic and Atmospheric Administration's (NOAA) Hazard Mapping System (HMS). These data are created using a combination of satellite imagery data and smoke analysts to determine the size, shape, and concentration of significant smoke plumes over the United States. Using these data, I am able to construct a measure of the exposure of each county to wildfire smoke at the daily level. For each day if a smoke plume intersects with any area of the county, then the county is counted as having a day of wildfire smoke exposure. It is important to note that this is an upper bound on the number of days an individual is exposed to wildfire smoke as a county unit is counted as exposed to wildfire smoke if

[6] Categorizations based on the Major Diagnostic Category ICD-9/ICD-10 coding system. Details can be found in Appendix Table A.1

[7] Air basin data comes from The California Air Resources Board https://ww3.arb.ca.gov/ei/maps/2017statemap/abmap.htm

[8] Even with these counties being combined there are still some observations that are masked due to the number of monthly admissions for a category is less than 11.

Figure 1.3: Monthly Totals of Hospital Admissions for Respiratory Illnesses

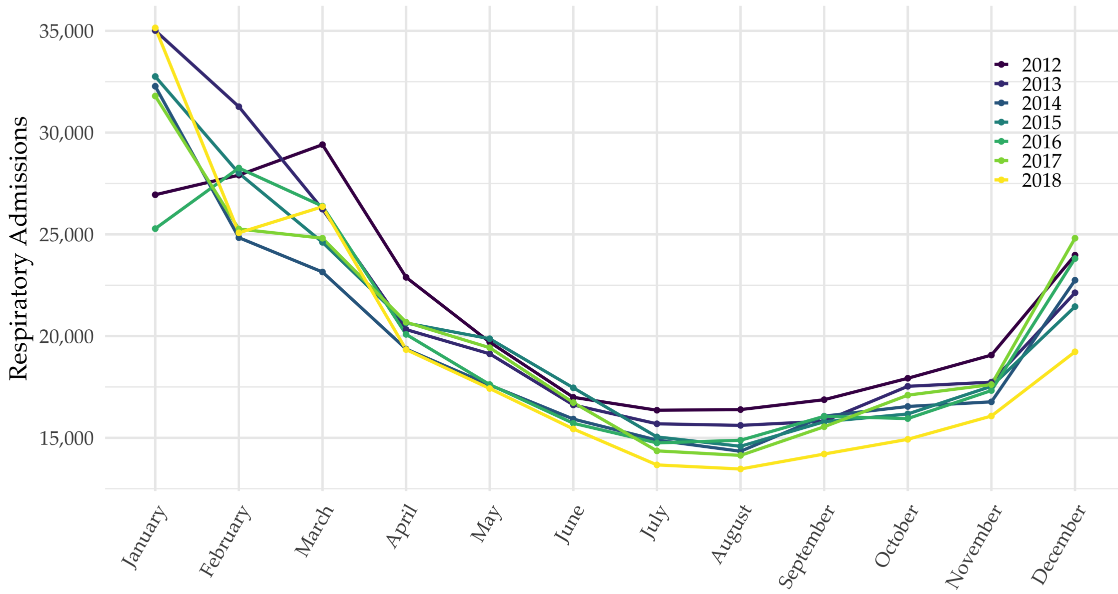

any of the county intersects a wildfire smoke plume on a given day. I then aggregate this daily county data to county units and to the monthly level to match the unit of observation in the healthcare data.

Figure 1.5 displays the distribution of the monthly number of smoke days[9], while Figure 1.6 shows a map of the average monthly number of smoke days across all months. The most frequent value Smoke Days takes is zero and the majority of months have less than three Smoke Days. Looking at the map in Figure 1.6 we can see the median number of Smoke Days in each county.

Similar to respiratory hospital admissions, the number of Smoke Days also follows seasonal trends. However, unlike respiratory hospital admissions which peak in the winter months, the number of Smoke Days tends to peak in the late summer months as can be seen in Figure 1.7.

[9]A boxplot representing the distribution of Smoke Days for every county can be found in Appendix Figure A.1

Figure 1.4: Units of analysis are outlined in black, divisions of counties within a unit of analysis are outlined with a dashed white line. All units of observation are displayed below the map.

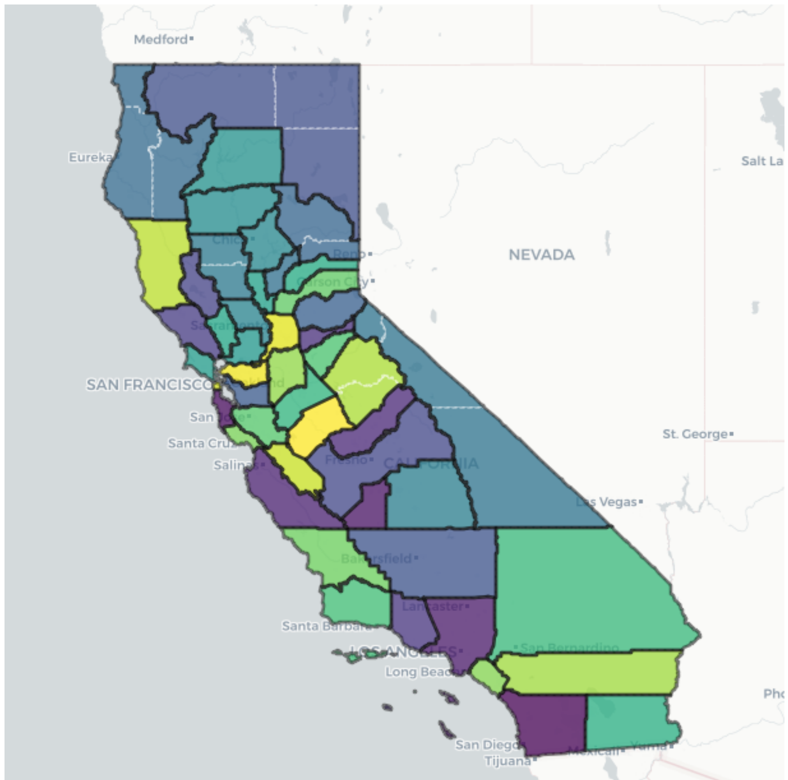

Alameda	Alpine—Inyo—Mono	Amador
Butte	Calaveras	Colusa—Glenn
Contra Costa	Del Norte—Humboldt—Trinity	El Dorado
Fresno	Imperial	Kern
Kings	Lake	Lassen—Modoc—Siskiyou
Los Angeles	Madera	Marin
Mariposa—Tuolumne	Mendocino	Merced
Monterey	Napa	Nevada
Orange	Placer	Plumas—Sierra
Riverside	Sacramento	San Benito
San Bernardino	San Diego	San Francisco
San Joaquin	San Luis Obispo	San Mateo
Santa Barbara	Santa Clara	Santa Cruz
Shasta	Solano	Sonoma
Stanislaus	Sutter	Tehama
Tulare	Ventura	Yolo
Yuba		

Figure 1.5: Distribution of the number of Smoke Days 2012 – 2018

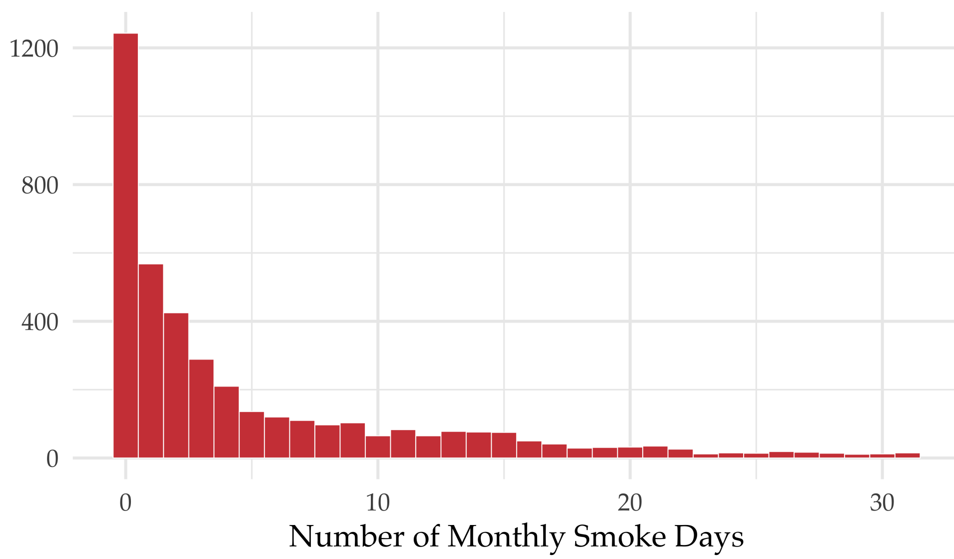

Figure 1.6: Monthly Median Number of Smoke Days 2012 – 2018

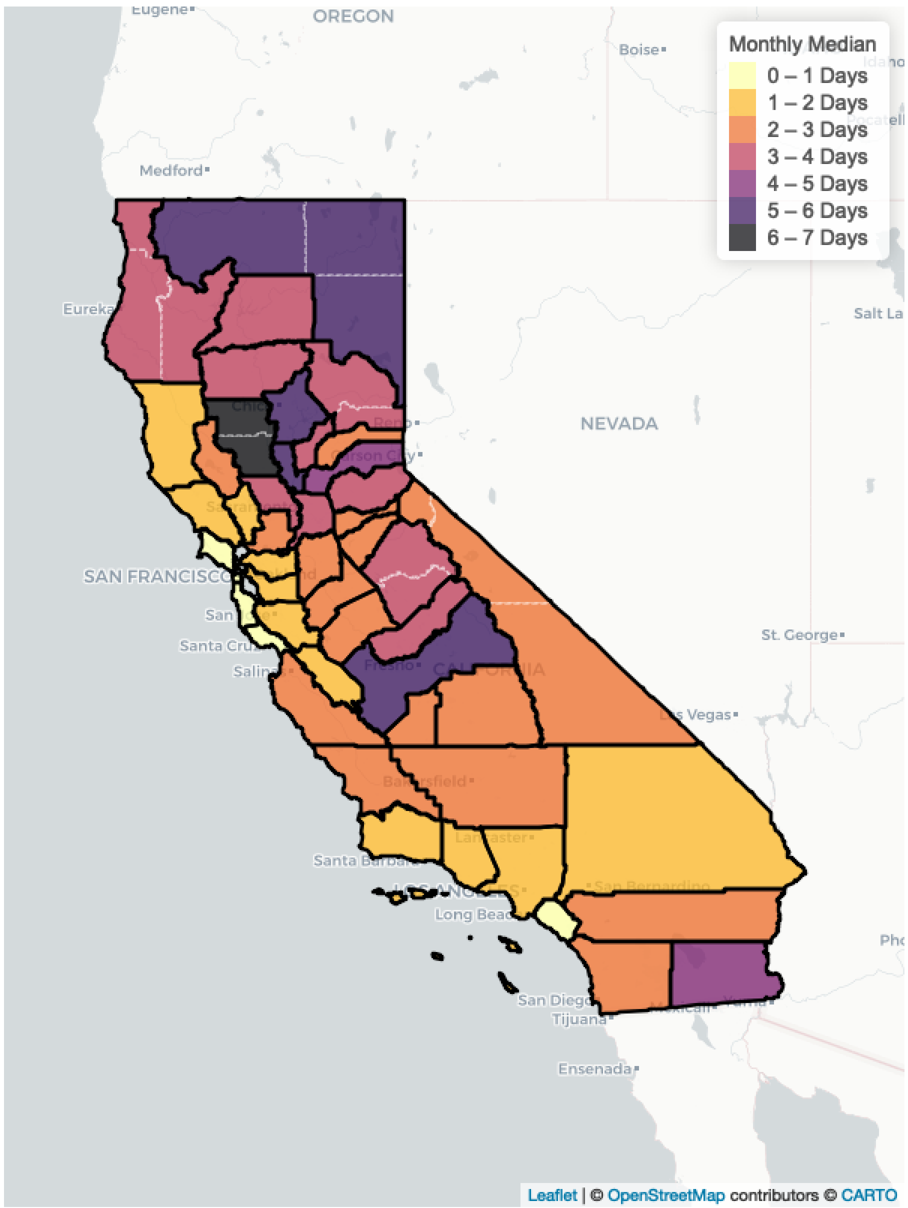

Figure 1.7: Monthly Average Smoke Days

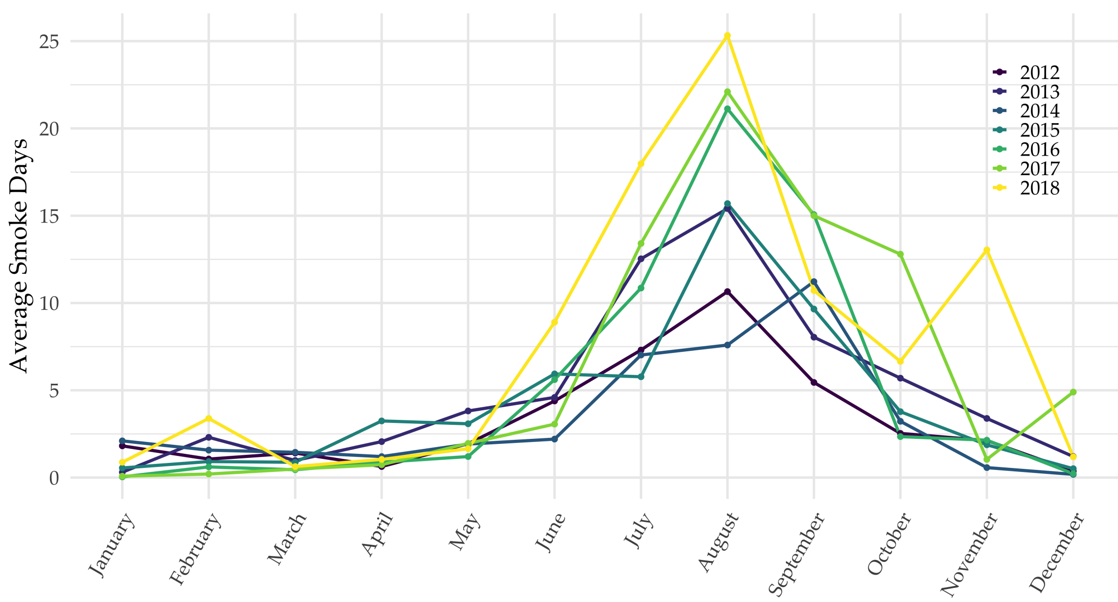

1.4 Empirical Strategy and Results

1.4.1 Hospital Admissions

While the impact of $PM_{2.5}$ on health has been well documented, wildfire smoke consists of more harmful substances in addition to $PM_{2.5}$. To evaluate the impact of exposure to wildfire smoke on health, I look at California hospital admissions data specifically for diseases and disorders of the respiratory and circulatory systems. This data is available at the monthly level so I utilize variation in the monthly number of wildfire smoke days at the county level to identify the impact of smoke exposure. Using monthly data from OSHPD, I construct a monthly level measure of smoke exposure $Smoke_{ct}$, by aggregating the number of smoke days for each month at the county level. I then use the following regression model:

$$Y_{ct} = \sum_{m=0}^{M} \beta_m Smoke_{c(t-m)} + X_{ct}\alpha + \gamma_{ct} + \eta_t + \tau_t + \varepsilon_{ct} \tag{1.1}$$

where Y_{ct} is a health outcome of interest in month t in county c.[10] I include M number of lags to account for previous smoke exposure, using M equal to 3 in the main results.[11] To control for differences at the county level, I include county by year fixed effects γ_{ct}. X_{ct} includes time-varying county level demographic controls, such as unemployment rate, poverty rate, health insurance coverage rate, population,[12] and weather controls, temperature and precipitation.[13] η_t is a year fixed effect to account for year to year differences that impact all counties, such as a year with a bad flu season. τ_t is a month by year fixed effect to account for seasonality in health outcomes.

[10] Counts are used rather than admission rates per capita because the counts are small for some counties and the variance of the outcome is very large.
[11] The results including up to 9 lags are presented in the Appendix.
[12] This data comes from the American Community Survey (ACS) 5 year estimates.
[13] Derived from NOAA's Climate Divisional Database (nCLIMDIV)

β_0, β_1, ..., β_m are the coefficients of interest, they can be interpreted as the effect of one additional day of wildfire smoke this month, last month, ..., m months ago on the number hospital admissions for a given category.

Table 1.1 shows the impact of an additional smoke day on hospital admissions for respiratory and circulatory diagnoses. Comparing columns 1 and 2, we can see the sign of the coefficient flip once the seasonality of smoke days and respiratory admissions is controlled for using weather variable controls and month fixed effects. As can be seen in column 4, an additional day of smoke exposure in a month leads to on average 5.56 additional hospital admissions that month and 3.78 additional hospital admissions the following month for respiratory diagnoses. Column 8 presents the results for circulatory admissions, an additional day of smoke exposure in a month leads to an additional 1.84 hospital admissions for circulatory diagnoses on average and about 1.19 additional admissions the following month.

Looking to the lagged impacts in columns 4 and 8, the positive coefficients for both respiratory admissions and circulatory admissions shows that the results are not just picking up a shift in timing. That is to say, the results are not being driven by wildfire smoke exposure causing admissions that would have occurred at a later date (a negative coefficient would be consistent with that effect). These results also match the mortality patterns found by Miller, Molitor, and Zou (2017). They found that the mortality effects tended to increase as the post wildfire exposure time period increases. This could be suggestive evidence that prolonged exposure has an even larger effect on hospital admissions.

Table 1.1: Impact of number of smoke days on hospital admissions in California.

	Respiratory Admissions				Circulatory Admissions			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cumulative Effect	-7.36^{***}	4.41	11.72^{*}	11.38^{*}	-0.30	-9.26^{**}	2.91^{**}	3.00^{**}
	(2.19)	(3.29)	(6.38)	(6.09)	(0.68)	(4.60)	(1.37)	(1.45)
Smoke Days	-4.74^{***}	4.06^{**}	5.61^{*}	5.56^{**}	-0.31	-1.73	1.97^{**}	1.84^{**}
	(1.29)	(1.96)	(2.84)	(2.66)	(0.32)	(1.60)	(0.93)	(0.81)
Smoke Days 1 Month Previous	-0.79^{***}	2.63^{*}	3.68^{**}	3.78^{**}	-0.06	-1.22	0.98^{***}	1.19^{**}
	(0.29)	(1.50)	(1.70)	(1.88)	(0.15)	(0.85)	(0.33)	(0.46)
Smoke Days 2 Months Previous	-3.50^{***}	1.08	2.51^{**}	2.25^{**}	-1.18^{***}	-1.44^{*}	0.67^{**}	0.69^{*}
	(1.06)	(0.67)	(1.24)	(1.01)	(0.39)	(0.84)	(0.28)	(0.37)
Smoke Days 3 Months Previous	1.66^{***}	-3.36^{**}	-0.08	-0.21	1.24^{***}	-4.88^{***}	-0.71	-0.72
	(0.43)	(1.35)	(1.40)	(1.42)	(0.34)	(1.56)	(0.51)	(0.48)
Weather Controls		✓	✓	✓		✓	✓	✓
Month FE		✓	✓	✓		✓	✓	✓
County-by-Year FE			✓	✓			✓	✓
Year FE				✓				✓
Num. obs.	4,059	4,059	4,059	4,059	4,054	4,054	4,054	4,054
Num. clusters	49	49	49	49	49	49	49	49
FENC	43.67	36.08	35.47	36.44	43.62	36.15	35.85	35.6

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
The feasible effective number of clusters is reported for the Smoke Days regressor.

Tables A.2 and A.3 in the Appendix present the results of the full fixed effects specifications in column 4 and 8 of Table 1.1 with an increasing number of lags. Figure 1.8 shows the robustness of the estimated coefficients for the contemporaneous effect and one and two month lagged effects of the number of Smoke Days on Respiratory Admissions. We can see that adding additional lags slightly changes the estimated coefficients, but the 95% confidence intervals for all of the regression specifications overlap all of the estimated coefficients. Figure 1.9 similarly shows the robustness of the estimated coefficients for the contemporaneous effect and one and two month lagged effects of the number of Smoke Days on Circulatory Admissions. The results are very similar for Circulatory Admissions as for Respiratory Admissions, although the difference in estimated coefficients is even smaller between models.

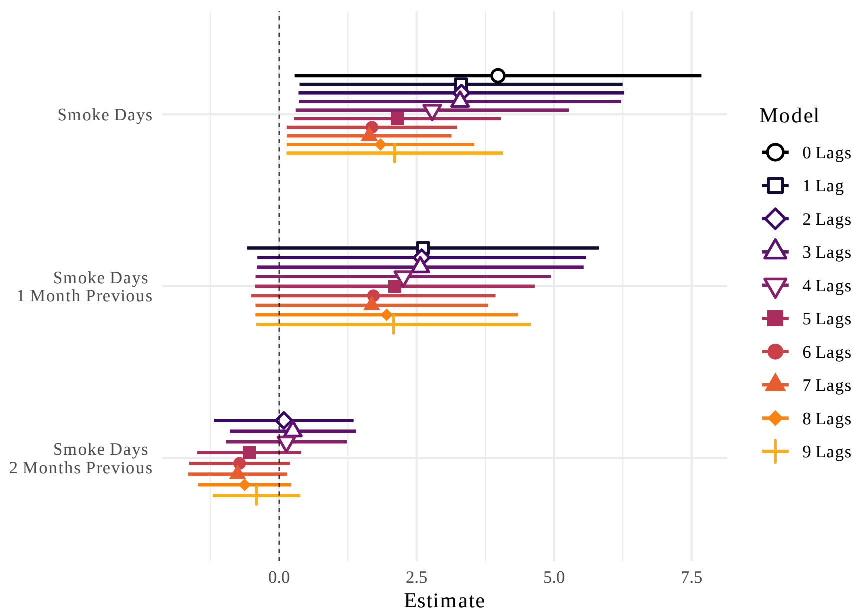

Figure 1.8: Coefficient Plot for the impact of number of smoke days on respiratory hospital admissions.

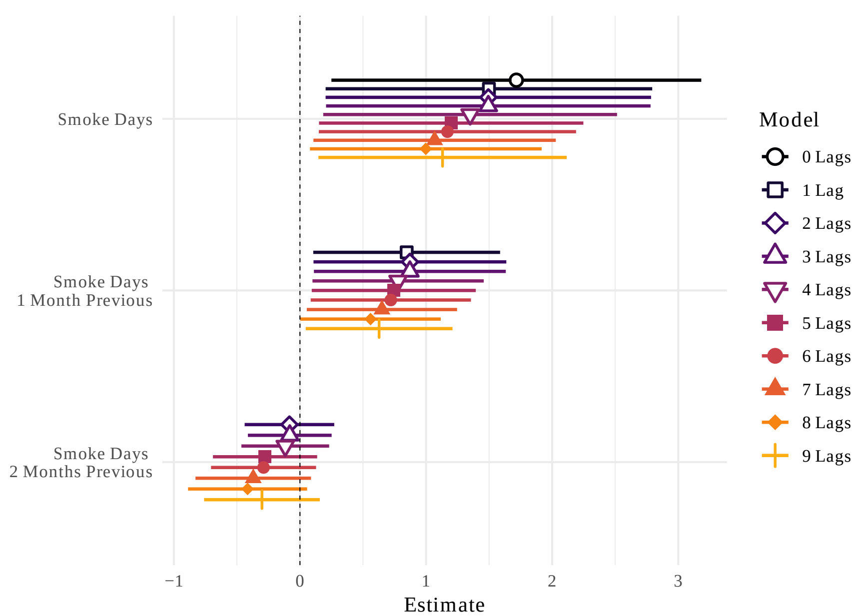

Figure 1.9: Coefficient Plot for the impact of number of smoke days on circulatory hospital admissions.

Table 1.2 presents the results for admissions for infections and injuries and poisonings, two admissions diagnoses that would not be expected to be impacted by exposure to wildfire smoke. Columns 4 and 8 show that there is no contemporaneous effect of additional smoke days on admissions for those two categories and even a slight decrease in the number of hospital admissions for infections. These results suggest that there is not some other factor increasing the number of hospital admissions when there is increased exposure to wildfire smoke.

Table 1.2: Impact of number of smoke days on hospital admissions in California.

	Infections Admissions				Injuries & Poisonings Admissions			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Smoke Days	-0.53	-2.71	0.10	0.14	1.25^{***}	-3.58^{**}	-0.92^{**}	-0.91^{**}
	(0.39)	(1.87)	(0.26)	(0.23)	(0.34)	(1.52)	(0.45)	(0.45)
Smoke Days 1 Month Previous	0.35	-1.34	0.58	0.49	0.04	-2.17^{***}	-0.63^{***}	-0.68^{***}
	(0.23)	(1.06)	(0.40)	(0.38)	(0.10)	(0.76)	(0.20)	(0.23)
Smoke Days 2 Months Previous	-1.28^{***}	-1.48	0.52^{**}	0.46^{**}	0.42^{***}	-1.58^{**}	-0.23	-0.22
	(0.44)	(0.90)	(0.21)	(0.20)	(0.13)	(0.62)	(0.19)	(0.19)
Smoke Days 3 Months Previous	1.59^{***}	-4.15^{**}	-0.12	-0.17	0.21	-2.69^{**}	-0.20	-0.19
	(0.46)	(1.79)	(0.33)	(0.32)	(0.19)	(1.09)	(0.19)	(0.20)
Weather Controls		✓	✓	✓		✓	✓	✓
Month FE		✓	✓	✓		✓	✓	✓
County-by-Year FE			✓	✓			✓	✓
Year FE				✓				✓
Num. obs.	3,962	3,962	3,962	3,962	4,057	4,057	4,057	4,057
Num. clusters	49	49	49	49	49	49	49	49
FENC	42.28	34.6	33.35	34.03	43.81	36.38	35.31	36.06

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
The feasible effective number of clusters is reported for the Smoke Days regressor.

1.4.2 Nonlinear Effect

To investigate whether the impact of the number of smoke days has a nonlinear impact on health, I create indicator variables for bins of the number of smoke days. It has been seen in the literature that the impact of temperature on health has a Ushaped relationship. Creating these binned variables will allow me to investigate if this relationship also exists for wildfire smoke exposure. I create $SmokeBin_{bct}$, a set of B indicator variables that are equal to 1 if a county's number of wildfire smoke days is within that range and 0 otherwise. I also create lags of the binned variable for 1 and 2 months. Table 1.3 displays the binned categories and the frequency of each bin by month. The bins for the largest number of smoke days occurs mostly in the summer months while the lowest bin is most prominent in the winter and spring. Using bins of 5 days allows for a flexible design while still providing fairly precise estimates, decreasing the bin size would allow for more flexibility at the expense of precision of the estimates. Section A.5 presents an alternative binning strategy constructing the bins using quantiles to ensure that the bins contain roughly equal numbers of observations. However, the quantile approach is not able to capture the nonlinearities for the larger number of Smoke Days due to the fewer number of months with these larger numbers.

Table 1.3: Total number of observations in each category by month 2012 – 2018

Month	<5	6 – 10	11 – 15	16 – 20	21 – 25	>25
January	329	12	2	0	0	0
February	320	18	5	0	0	0
March	338	5	0	0	0	0
April	324	19	0	0	0	0
May	325	17	1	0	0	0
June	213	88	34	8	0	0
July	76	97	102	41	23	4
August	53	49	49	55	58	79
September	80	90	102	48	18	5
October	219	55	45	22	2	0
November	273	26	36	7	1	0
December	321	19	1	2	0	0

To estimate the nonlinear effects of wildfire smoke exposure on health, I use the following regression model:

$$Y_{ct} = \sum_{m=0}^{M} \sum_{b=1}^{B} \beta_{mb} SmokeBin_{bc(t-m)} + X_{ct}\alpha + \gamma_{ct} + \eta_t + \tau_t + \varepsilon_{ct} \tag{1.2}$$

where Y_{ct} is a health outcome of interest in month t in county c. I include M number of lags to account for previous smoke exposure and bin smoke exposure into B bins. The results of the regression with the binned variables can be seen in Tables 1.4, 1.5, and 1.6. Table 1.4 shows the cumulative contemporaneous and lagged effect of wildfire smoke exposure on circulatory and respiratory hospital admissions. For both respiratory and circulatory admissions the largest impact is seen in months with greater than 15 days of wildfire smoke exposure.

Looking at the full results for respiratory admissions in Table 1.5 column 4, the impacts of exposure to wildfire smoke seem to only occur when the number of smoke days is high. The impacts are concentrated in the coefficients for 16 to 20 Smoke Days, 21 to 25 Smoke Days, and Over 25 Smoke Days. The 16 to 20 Smoke Days coefficient represents an increase of about 200 hospital admissions for respiratory illnesses compared to if the county experienced 5 or less smoke days.

Now looking to the results for circulatory admissions in Table 1.6 column 4, similar to the results for respiratory admissions the bins of wildfire smoke exposure with larger number of days lead to a significant increase in the number of hospital admissions for circulatory illnesses compared to the base case of below 5 smoke days. Three categories 16 to 20, 21 to 25, and over 25 all have large significant estimated coefficients, which means, for example, a month with between 21 and 25 smoke days leads to over 100 additional hospital admissions for circulatory diagnoses compared to if there were less than 5 days of wildfire smoke exposure.

Table 1.4: Cumulative impact of number of smoke days on hospital admissions in California.

	Respiratory Admissions				Circulatory Admissions			
Cumulative Effect	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
6 to 10 Smoke Days	-194.82^{***}	-33.51	-8.11	-19.66	-75.85^{***}	-64.93	-1.02	-0.99
	(70.90)	(25.31)	(29.20)	(32.24)	(26.69)	(44.06)	(6.80)	(7.95)
11 to 15 Smoke Days	-145.18^{***}	92.15	122.19	117.64	3.60	-26.77	38.16	38.37
	(32.84)	(72.87)	(97.78)	(89.39)	(10.73)	(40.97)	(23.32)	(24.94)
16 to 20 Smoke Days	-108.48^{***}	123.05^{**}	217.29^{**}	213.15^{**}	-5.32	-96.03	64.86^{***}	66.28^{***}
	(39.95)	(60.57)	(104.17)	(96.14)	(24.80)	(58.48)	(24.74)	(25.42)
21 to 25 Smoke Days	-54.28^{*}	196.58	303.81^{*}	305.04^{*}	-3.12	-129.76	105.14^{**}	108.83^{**}
	(30.08)	(122.25)	(167.27)	(164.41)	(24.81)	(93.34)	(45.61)	(48.10)
Over 25 Smoke Days	-33.89	210.05^{*}	299.62^{*}	296.90^{*}	13.18	-142.06^{*}	82.69^{**}	85.28^{*}
	(24.20)	(117.56)	(159.99)	(158.70)	(22.85)	(80.91)	(41.47)	(45.54)
Weather Controls		✓	✓	✓		✓	✓	✓
Month FE		✓	✓	✓		✓	✓	✓
County-by-Year FE			✓	✓			✓	✓
Year FE				✓				✓
Num. obs.	4,059	4,059	4,059	4,059	4,054	4,054	4,054	4,054
Num. clusters	49	49	49	49	49	49	49	49
FENC	33.26	37.38	21.37	24.11	32.55	36.89	2.13	0

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
The feasible effective number of clusters is reported for the 21 to 25 regressor.

Table 1.5: Impact of number of smoke days on hospital admissions in California.

	Respiratory Admissions			
	(1)	(2)	(3)	(4)
6 to 10 Smoke Days	-86.52^{***}	-2.10	6.64	3.07
	(26.59)	(7.94)	(9.52)	(9.43)
11 to 15 Smoke Days	-77.74^{***}	40.93	43.37	48.27
	(18.21)	(24.91)	(29.76)	(32.25)
16 to 20 Smoke Days	-64.30^{***}	60.36^{**}	81.25^{**}	82.58^{**}
	(19.84)	(26.03)	(35.30)	(33.70)
21 to 25 Smoke Days	-45.99^{***}	93.14^{*}	125.84^{*}	125.48^{*}
	(14.34)	(51.11)	(70.63)	(66.87)
Over 25 Smoke Days	-46.39^{***}	110.43^{*}	139.03^{*}	133.39^{*}
	(13.61)	(56.94)	(82.08)	(74.35)
6 to 10 Smoke Days	-62.51^{**}	-10.94	-5.51	-7.88
1 Month Previous	(23.89)	(8.95)	(12.24)	(12.84)
11 to 15 Smoke Days	-36.09^{***}	37.30	41.45	38.26
1 Month Previous	(10.54)	(33.53)	(38.15)	(33.98)
16 to 20 Smoke Days	-20.57^{**}	43.19^{*}	66.23^{**}	63.20^{**}
1 Month Previous	(8.54)	(24.22)	(31.96)	(29.53)
21 to 25 Smoke Days	1.43	61.81	86.54^{*}	86.91^{*}
1 Month Previous	(14.39)	(42.19)	(47.90)	(47.94)
Over 25 Smoke Days	11.94	60.14	72.82^{**}	75.48^{*}
1 Month Previous	(11.48)	(37.25)	(35.93)	(41.27)
6 to 10 Smoke Days	-45.79^{**}	-20.46	-9.23	-14.85
2 Months Previous	(21.32)	(12.34)	(11.65)	(13.75)
11 to 15 Smoke Days	-31.35^{***}	13.92	37.37	31.11
2 Months Previous	(9.57)	(16.50)	(31.75)	(25.18)
16 to 20 Smoke Days	-23.61^{*}	19.51	69.81^{*}	67.37^{*}
2 Months Previous	(14.08)	(14.23)	(39.54)	(35.93)
21 to 25 Smoke Days	-9.72	41.62	91.43^{*}	92.65^{*}
2 Months Previous	(7.75)	(31.03)	(51.36)	(52.22)
Over 25 Smoke Days	0.56	39.48	87.77^{*}	88.03^{*}
2 Months Previous	(9.03)	(26.64)	(45.90)	(46.72)
Weather Controls		✓	✓	✓
Month FE		✓	✓	✓
County-by-Year FE			✓	✓
Year FE				✓
Num. obs.	4,059	4,059	4,059	4,059
Num. clusters	49	49	49	49
FENC	33.26	37.38	21.37	24.11

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
The feasible effective number of clusters is reported for the 21 to 25 regressor.

Table 1.6: Impact of number of smoke days on hospital admissions in California.

	Circulatory Admissions			
	(1)	(2)	(3)	(4)
6 to 10 Smoke Days	-27.66^{***}	-13.29	8.61^{**}	8.29^{**}
	(9.21)	(15.02)	(3.70)	(3.60)
11 to 15 Smoke Days	2.09	1.30	26.25^{*}	24.36^{**}
	(3.77)	(17.69)	(13.61)	(11.98)
16 to 20 Smoke Days	-2.58	-27.42	31.60^{**}	31.28^{***}
	(10.12)	(24.85)	(12.05)	(11.03)
21 to 25 Smoke Days	2.45	-40.01	47.40^{**}	45.75^{**}
	(9.55)	(36.80)	(21.08)	(18.72)
Over 25 Smoke Days	12.79	-51.19	38.25^{*}	36.22^{**}
	(8.95)	(35.49)	(19.37)	(17.40)
6 to 10 Smoke Days	-26.89^{***}	-25.26^{*}	-7.40^{*}	-7.39
1 Month Previous	(9.72)	(14.75)	(4.01)	(4.81)
11 to 15 Smoke Days	0.02	-9.09	1.39	2.84
1 Month Previous	(5.67)	(12.66)	(6.87)	(7.45)
16 to 20 Smoke Days	-5.26	-23.57	14.63^{**}	18.07^{*}
1 Month Previous	(7.11)	(14.77)	(6.83)	(9.33)
21 to 25 Smoke Days	-10.58	-31.19	28.22^{**}	32.55^{*}
1 Month Previous	(7.82)	(26.19)	(13.63)	(16.96)
Over 25 Smoke Days	-10.15	-16.61	28.53^{*}	32.93
1 Month Previous	(6.65)	(20.71)	(15.05)	(20.16)
6 to 10 Smoke Days	-21.29^{**}	-26.38^{*}	-2.22	-1.89
2 Months Previous	(8.54)	(15.56)	(4.20)	(4.43)
11 to 15 Smoke Days	1.49	-18.97	10.52^{*}	11.17
2 Months Previous	(3.85)	(13.42)	(5.94)	(8.08)
16 to 20 Smoke Days	2.52	-45.04^{**}	18.63^{**}	16.93^{**}
2 Months Previous	(9.07)	(20.87)	(8.11)	(7.63)
21 to 25 Smoke Days	5.00	-58.56^{*}	29.52^{**}	30.53^{**}
2 Months Previous	(9.07)	(32.45)	(12.58)	(14.45)
Over 25 Smoke Days	10.54	-74.26^{**}	15.92	16.13
2 Months Previous	(9.47)	(29.89)	(9.97)	(11.27)
Weather Controls		✓	✓	✓
Month FE		✓	✓	✓
County-by-Year FE			✓	✓
Year FE				✓
Num. obs.	4,054	4,054	4,054	4,054
Num. clusters	49	49	49	49
FENC	32.55	2.13	0	0

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
The feasible effective number of clusters is reported for the 21 to 25 regressor.

Looking to a graph of the coefficients in Figure 1.10 we can see that these bins corresponding to larger numbers of smoke exposure have larger standard errors. The large estimated standard errors for these bins are not unexpected. Looking back to Table 1.3 the number of observations for these bins is much lower than the number of observations in other bins. In addition many months have 0 observations in the bins for 16 to 20, 21 to 25, and over 25 Smoke Days. We can also see that the results for Circulatory Admissions are more precisely estimated than the results for Respiratory Admissions.

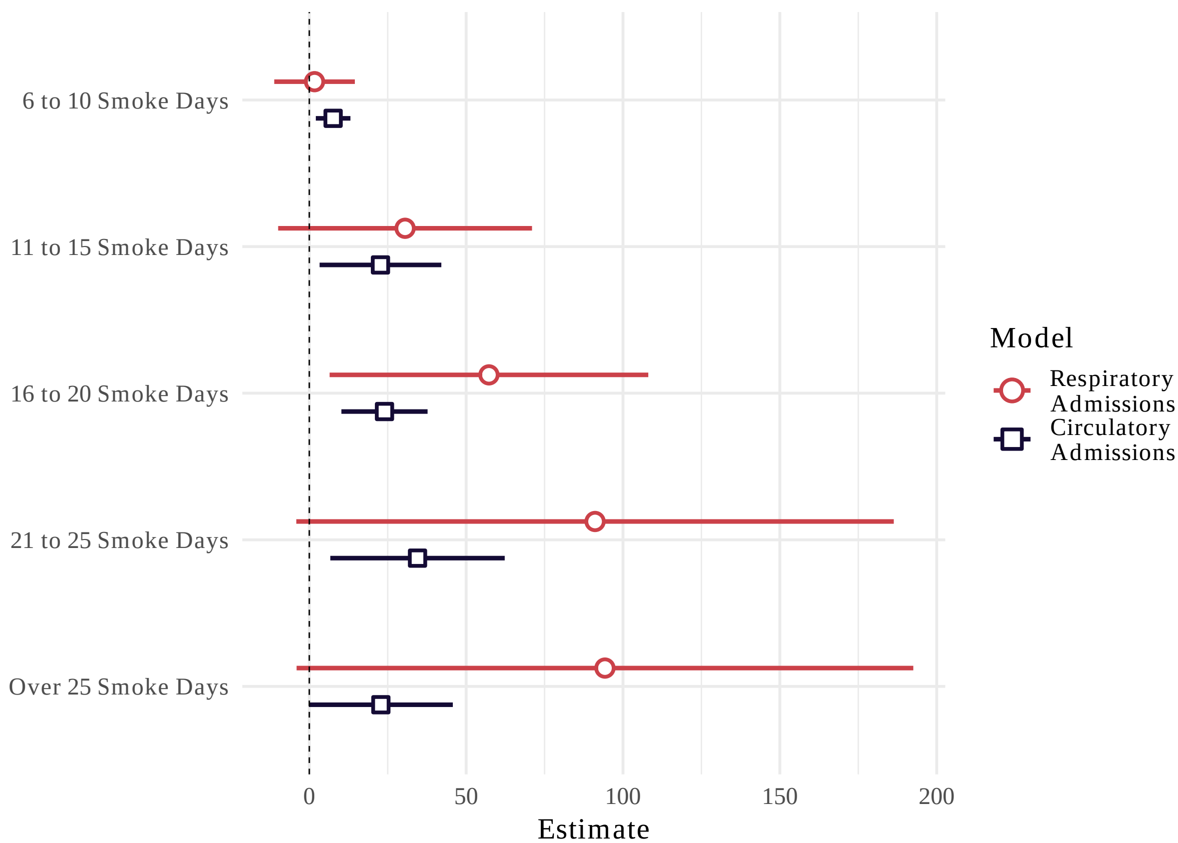

Figure 1.10: Coefficient Plot for the impact of number of smoke days on hospital admissions.

1.5 Economic Cost

Much of the research evaluating the economic cost of pollution looks at mortality as the outcome of interest. To evaluate the economic cost of mortality two methods are commonly used, multiplying the estimated number of lost lives by the value of a statistical life or estimate life-years lost and then multiple this by the value of a statistical life-year. However, when working with other health outcomes the economic valuation can be thought of as the amount that an individual is willing to pay to avoid the illness.

The United States Environmental Protection Agency (EPA) produced estimates of the cost of illness (COI) for a hospital admission for respiratory or cardiovascular illness as part of a series of reports on the benefits and costs of the Clean Air Act (1999). For this valuation, the EPA calculate the avoided medical costs and use this as an estimate of the value of avoiding the health effects causing the admission. The EPA estimates respiratory hospital admissions have a mean cost of $6,900 per case (1990 dollars) and cardiovascular hospital admissions have a mean cost of $9,500 per case (1990 dollars). These estimates can be thought of as a lower bound for the mean as they do not take into account all costs of illness only the medical costs.

To calculate the cost of an additional day of wildfire smoke exposure, I convert the EPA estimates to 2020 dollars using the Consumer Price Index (CPI) produced by the Bureau of Labor Statistics. I then multiply these per case values by the coefficients for exposure to wildfire smoke estimated in columns 4 and 8 of Table 1.1 and column 4 in Tables 1.6 and 1.5.

1.5.1 Economic Cost With Linear Effects

The regression specification in equation 1.1 produces coefficients that can be used to produce marginal effects. This allows for the interpretation as the impact of an addition day of wildfire smoke exposure. The cost of an additional day of wildfire smoke exposure can be calculated using the estimated values from the EPA. These calculations are displayed in Table 1.7. Each additional day of wildfire smoke exposure leads to about \$188,767 of additional medical expenditures on average.

Table 1.7: Estimated Average Cost per Additional Smoke Day

Contemporaneous Costs		
	Cases $\times$ Cost per case	
Respiratory Admission	$5.56 \times \$13,971.74 =$	\$77,678.98
Circulatory Admissions	$1.84 \times \$19,236.46 =$	\$35,417.03
	Total contemporaneous cost =	\$113,096.01
Lagged Costs		
	Cases $\times$ Cost per case	
Respiratory Admission	$3.78 \times \$13,971.74 =$	\$52,764.06
Circulatory Admissions	$1.19 \times \$19,236.46 =$	\$22,907.35
	Total lagged cost =	\$75,671.41
	Total cost =	\$188,767.42

Dollar values reported in 2020 dollars.

The total cost of wildfire smoke exposure for both respiratory admissions and circulatory admissions during the study period is presented in Table 1.8. From 2012 to 2018, there were a total of 20,574 days of wildfire smoke exposure in California counties. This leads to a total cost of wildfire smoke exposure in California due to respiratory and circulatory hospital admissions from 2012 to 2018 equal to \$2,326,837,289.

Table 1.8: Estimated Yearly Cost of Wildfire Smoke in California Due to Respiratory and Circulatory Hospital Admissions

Year	Smoke Days	Respiratory Admissions	Circulatory Admissions	Total
2012	1,944	$151,007,943	$68,850,698	$219,858,641
2013	2,956	$229,619,074	$104,692,728	$334,311,803
2014	1,973	$153,260,634	$69,877,792	$223,138,426
2015	2,543	$197,537,654	$90,065,497	$287,603,151
2016	2,967	$230,473,543	$105,082,316	$335,555,859
2017	3,713	$288,422,064	$131,503,417	$419,925,481
2018	4,478	$347,846,486	$158,597,442	$506,443,928

1.5.2 Economic Cost With Nonlinear Effects

The cost calculations using the linear model implicitly assume that the cost of each additional day of wildfire smoke exposure is constant. As we saw in Tables 1.5 and 1.6 the effect appears to be nonlinear, with larger effects when the number of Smoke Days is higher.

Table 1.9 presents similar cost calculations to Table 1.8, however the interpretation of these results are different. Using the estimated coefficients from the regression in equation 1.2 implicitly assumes a baseline month has between 0 and 5 days of wildfire smoke exposure while using the estimated coefficients from the regression in equation 1.1 implicitly assumes a linear effect of an addition smoke day. Comparing the average annual cost of wildfire smoke exposure using the results in Table 1.8 which gives $332,405,327 the results using the bin coefficient estimates give $192,316,498. The results from the binned regression give results approximately two-thirds the size of the linear regression results. However, the two estimated cost values are not found to be statistically different. Looking at the economic costs using either calculation we can see that the largest costs were in 2017 and 2018. These two years included the two largest wildfires in California history – the Thomas Fire (2017) and the Mendocino

Complex Fire (2018).

Table 1.9: Estimated Yearly Cost of Wildfire Smoke in California Due to Respiratory and Circulatory Hospital Admissions

Year	6 to 10	11 to 15	16 to 20	21 - 25	Over 25	Respiratory	Circulatory	Total
2012	49	36	11	1	6	$52,006,044	$36,361,563	$88,367,607
2013	98	57	22	10	7	$98,603,634	$69,250,232	$167,853,866
2014	66	29	10	9	2	$53,431,388	$39,443,199	$92,874,587
2015	110	33	17	12	4	$75,078,822	$56,577,385	$131,656,207
2016	45	63	36	18	14	$143,602,620	$83,952,559	$227,555,179
2017	63	68	47	26	19	$183,782,126	$106,307,112	$290,089,237
2018	64	91	40	26	36	$222,941,215	$124,877,587	$347,818,802

1.6 Conclusion

This paper provides estimates of the relationship between wildfire smoke exposure and respiratory and circulatory health. By combining data on hospital admissions and wildfire smoke plumes in California from 2012 to 2018, I am able to estimate the impact of increased wildfire smoke exposure on hospital admissions. I find that wildfire smoke exposure increases hospital admissions for respiratory and circulatory illnesses in the month of smoke exposure as well as in the following month. In addition, I find that increases in respiratory admissions mainly occur when the number of monthly wildfire smoke days per month is above 16 days. While a significant increase in circulatory admissions occurs for any level compared to five or less days with larger impacts occurring above 16 days.

While the absolute number of cases is not a large number compared to the total number of hospital admissions, each hospital admission does carry a large cost. I further estimate that the average annual cost of wildfire smoke exposure due to respiratory and circulatory health in California to be around $192,300,000 and the total cost during the study period to be around $1,346,200,000.

Chapter 2

The Sharing Economy and Rental Markets

with Travis Cyronek

2.1 Introduction

The peer-to-peer rental market has seen rapid growth since the introduction of *Airbnb* in 2008 and *Uber* in 2009. These platforms allow individuals to share and use goods and services that might have otherwise been underutilized. In the *Airbnb* example, entire apartments, houses, or individual rooms can be rented on a short-term basis. This increasingly prominent way to interact in the economy has led to regulatory battles throughout the United States as housing affordability has become a major political issue. Much of the discussion on how to regulate the short-term rental market has centered specifically on *Airbnb*. Before regulations are implemented it is important to understand the effects these markets have on rental and housing markets, as well as the impact on local residents. Proponents of these peer-to-peer markets argue that users of these services will see many benefits, including additional income, more efficient resource allocation, and the creation of new economic activity,[1] while opponents argue that these markets avoid regulations and increase rents for local renters.

The present research contributes to this discussion by studying the effect of peer-to-peer housing technologies on traditional, *impermanent* markets for accommodation. That is, we study how rental properties affect the availability and price of hotels and long-term (annually leased) rentals. To guide our work we endeavor to address two questions, one positive and the other normative. (1) How are the number of *Airbnb* listings in an area related to the average price paid for rentals / hotels and (2) What is the optimal way to regulate the market for short-run accommodation? In this effort we construct and use a novel dataset for the Santa Barbara, California housing market. We combine hotel price and vacancy data for hotels (*Visit Santa Barbara*) with information on rental properties readily available from the U.S. Census Bureau's *American Community Survey (ACS)*. Finally, we rely on scraped *Airbnb* data collected by *Inside Airbnb* and Tom Slee.

Though the data are rich in many respects, purely reduced form analyses of such markets may suffer from an inability to isolate exogenous variation in the key covariates and identify causal relationships. Even using fixed effects and controlling for amenity heterogeneity using proxy variables, identifying causal estimates without bias is implausible. Indeed, this difficulty is inherent to identification in the housing literature because amenity values are imperfectly measured. To circumvent some of these issues we develop a structural search and matching model where property managers post vacant rooms (and their prices) and tenants direct their search to these postings. By fully defining and describing the agents, their actions, and the equilibrium,

[1] https://blog.atairbnb.com/economic-impacts-los-angeles/

we bypass the need to directly address amenity heterogeneity and instead can use observables and the model's structure to disentangle the mechanisms at play[2]. Discrete agent types and the contracts they make with one another define three separate–but endogenously related–markets for lodging. Hotels are accessible by innkeepers and visitors, short-term rentals by visitors and landlords, and long-term rentals by landlords and residents. The key feature is that, since multiple markets are available to some agents, behavior in one market may influence the outcomes in the others. For example, a landlord's decision to list a property in the short-term market negatively affects residents who are seeking long-term accommodation.

In our calibration exercise, we target average (median) prices for hotels, *Airbnb*s, and rental units, sizes (i.e. number of) of these markets, and average length of stays for visitors and residents. After calibrating the model we find that *Airbnb*s decrease nightly hotel prices by about $24 while they increase average rents by $39 per room, per month. The added choice afforded to visitors, though, increases their flow utility by about 3%. This is offset quantitatively to losses in welfare of residents, who have fewer rentals to search for and higher prices. We ultimately find that, with limited entry, aggregate welfare is *lower* with *Airbnb*. Search decisions by visitors and landlords do not internalize the costs to innkeepers and residents. As a result, government policy can improve efficiency. We find that the optimal policy is to set a high transient occupancy tax on short-term rentals as the lost utility to residents is quantitatively dominant.

This paper relates to a limited yet growing literature on the relationship between short-term, peer-to-peer rental markets and traditional housing and rental markets. A majority of this scholarship is case studies of individual cities. These case studies

[2]To put this in a slightly different context, the model allows us to disentangle simultaneous equations that would, in a reduced form, introduce bias.

provide anecdotal and descriptive analysis on the relationship between the growth of *Airbnb* and housing and rental markets in a number of cities. For example, Lee (2016) suggests that *Airbnb* listings are limiting the supply of rentals for long-term use and pushing up rents in the Los Angeles housing market. He goes on to recommend a set of regulations and taxes that could help lead to more affordable housing. Quattrone et al. (2016) investigate when and where *Airbnb* listing are offered in London and the socio-economic conditions of the areas with concentrated *Airbnb* usage. They find that *Airbnb* listings tend to be in areas that are accessible to public transit, and have residents who are young, employed, and born outside the UK.

Others have explored the heterogeneous impact–both within and between cities—of these peer-to-peer technologies. Coles et al. (2017) explore the usage of *Airbnb* across neighborhoods in New York City to look at this differential relationship. Using matched census tract level data from *Airbnb* with neighborhood rent data produced by *Zillow*, they find that *Airbnb* listings have become more geographically dispersed over time. They also find that short-term rentals appear most profitable relative to long-term rentals in outlying, middle-income neighborhoods. Coyle and Yeung (2016) provide an overview of *Airbnb* in fourteen European cities. They find that the presence of *Airbnb* in a market has a negative relationship with hotel occupancy rates, but a positive relationship with average daily hotel rates. They also find an ambiguous relationship on the rental market, suggesting that the relationship between *Airbnb* and rental markets may depend on specific characteristics of the rental market.

Other studies have pursued identifying the causal effect of *Airbnb* on rental and housing markets. Horn and Merante (2017) utilize data from online apartment and *Airbnb* listings to evaluate the growth of *Airbnb* on asking rents in Boston. Using a fixed effects model, they show that a one standard deviation increase in *Airbnb* listings is associated with a 0.4% increase in asking rents in Boston. Garcia-López et al.

(2019) study the effect of *Airbnb* listings on rental rate in Barcelona, Spain. Using multiple econometric specifications, they find that a neighborhood with the average amount of *Airbnb* activity saw rents increase by 1.9%, while neighborhoods in the top 10% percent of *Airbnb* activity, saw increased rents by 7%. Looking to French cities, Ayouba et al. (2020) show that increase in *Airbnb* rentals is associated with increased rents in Lyon, Montpellier, and Paris, however *Airbnb* has no significant effect in other cities. Understanding how this heterogeneity in effect arises is an important characteristic for policy makers to study and understand.

Barron et al. (2018) offer the most complete look at the impact of *Airbnb* listings on rent and house prices across the United States. Using an instrumental variable approach they estimate that a 1% increase in *Airbnb* listings leads to a 0.018% increase in rental prices and a 0.026% increase in home prices. Doing a quick back-of-the-envelope calculation, this corresponds to a $9 increase in monthly rent and $1,800 increase in house prices. In addition they find that *Airbnb* does not impact the total supply of housing but does decrease the supply of long-term rentals. The model we write down is informed by this finding insofar as we take the supply of rental properties as given, but endogenously allow the fraction of vacancies posted in one market or another to depend on market conditions.

A major concern when estimating causal effects in rental and housing markets is the bias introduced by the exclusion of amenity values in an area. The utilization of area specific fixed effects can control for some of this bias by estimating parameters using within area variation. However, a time varying component of amenities may still be biasing the estimates. Using an instrument may also be able to reduce the bias, as the first stage relationship may reduce the correlation with the error term. However, as detailed in the next section, omitted variable bias is still a concern. Due to these doubts about producing unbiased causal estimates, we turn our attention to the theoretical literature on search and matching in housing markets.

Dating back at least as early as Wheaton (1990), models with search frictions have been used in the context of buying and selling property. First, they offer a realistic and intuitive reason for vacant properties to exist in equilibrium by taking seriously the idea that markets may clear through prices *and* time. Genesove and Han (2012) use a random search framework to study buyer and seller contact rates and time on the market. Second, models with search frictions have also had varied success in describing and explaining price dispersion (Albrecht et al. (2016)). Maury and Tripier (2014) and Moen et al. (2014) think about search strategies (e.g. where to direct search, when to make an offer, etc.) and their role in the dynamism of housing markets.

Indeed, models of directed search are particularly attractive theoretically inasmuch as they seem to align with what happens in observed housing markets. Zhu et al. (2017) develop and formalize multiple models of housing with price posting. We contribute to the above literature by applying the insights and tools of search theoretic models to *impermanent* lodging markets. More specifically, hotels and rental properties. In this effort we also seek to connect the often disjointed approach of looking at hotels and rentals separately. Finally, by explicitly modeling the endogenous relationships of key acting agents in these markets, we can therefore think in terms of normative assessments of optimal policy and address, perhaps incompletely, the discussions and debates about how peer-to-peer technologies should be governed.

In the next section we present and explore statistical facts about *Airbnb*, hotels, and rental properties in various regions. We also present and discuss some of the confounds and shortfalls of interpreting these results in the context of identifying the *effect* of *Airbnb* on these markets. We then develop a rich-yet-simple model of rental markets that allows us to circumvent these shortfalls and enables us to study the highly interrelated markets for lodging. We calibrate this model to a novel dataset that we construct for Santa Barbara, California, which we also use to assess questions of optimal policy regarding how to tax the various agents to maximize welfare.

2.2 Empirical Regularities

In this section we present the main sources of data and establish the statistical relationship between *Airbnb* listings and the price of apartment rentals. Our empirical strategy likely does not identify the causal effect of these listings on the price of rental properties. Its purpose is to motivate the key mechanisms in our structural model of impermanent housing markets. We begin by examining the *Airbnb* data and the distribution of prices and the growth of listings in major American metro areas since 2017. From *Zillow* and the *American Community Survey* we merge the features of long-term (i.e. "traditional") rental markets: prices, stocks, and vacancies. These characteristics are then later used to calibrate the structural model.

Using these data we estimate the statistical relationship between the number of *Airbnb* listings and the median rental price using a simple fixed effects model. The implied effect is then used to compare to the results generated by the structural model. As this is the first paper to use a search and matching model in this setting, this exercise allows us to compare our results to the empirical literature.

2.2.1 Description of *Airbnb* Data

The *Airbnb* data for major United States metro areas come from the free, publicly available data collected and hosted by *Inside Airbnb*.[3]. Data for Santa Barbara was

[3]The data was sourced from publicly available information from the *Airbnb* site and cleaned and aggregated by *Inside Airbnb* The is available under a Creative Commons CC0 1.0 Universal (CC0 1.0) "Public Domain Dedication" license at `http://insideairbnb.com/get-the-data.html`

collected and provided publicly by Tom Slee[4] on his website. This data consists of information about the room type, price, number of reviews, and exact location of each listing. The data also consists of the availability calendar for the next year into the future. The calendar for a listing gives a price for dates that are available to book, but nights that are unavailable to book cannot be differentiated from nights that have already been booked. However we aggregate the calendar for each listing from the daily to monthly level. This means that if a listing is available for at least one day in a month, the listing is considered active. We then calculate the price of a listing for the month by taking the median of its listed prices for the month.

Data are collected at *roughly* a monthly frequency, therefore we can observe many overlapping calendars for the same listing. In other words, for a listing collected in January 2017 we observe available nights for January 2017 to January 2018. When this listings data is collected again in February 2017 we observe February 2017 to February 2018, therefore we observe availability and prices for February 2017 to January 2018 twice in these two observations. That is, we can observe data for a single month for a listing up to twelve times.[5] The monthly price for a listing is calculated by finding the median monthly price in each observation then taking the maximum value across up to twelve monthly median price observations.[6]

Our dataset contains listings for 12 cities: Boston, Chicago, Denver, Los Angeles, Nashville, New York City, Portland, San Diego, San Francisco, Santa Barbara, Seattle,

[4] http://tomslee.net/airbnb-data-collection-get-the-data

[5] Because of this fact, even popular listings are likely to appear as available *at least* once even if they become fully booked.

[6] The method we use to determine the price of the listing will not impact the empirical analysis as we are only looking the relationship of the number of *Airbnb* listings on rent prices. This method will be used to determine the value used in calibration, however doing this exercise by taking the mean or median of the observed prices does little to affect the analysis because the data is aggregated to the ZIP code level for each month and only the median monthly price for the entire ZIP code is used. So, the price we calculate can be thought of as the maximum price a landlord can receive per night for their listing.

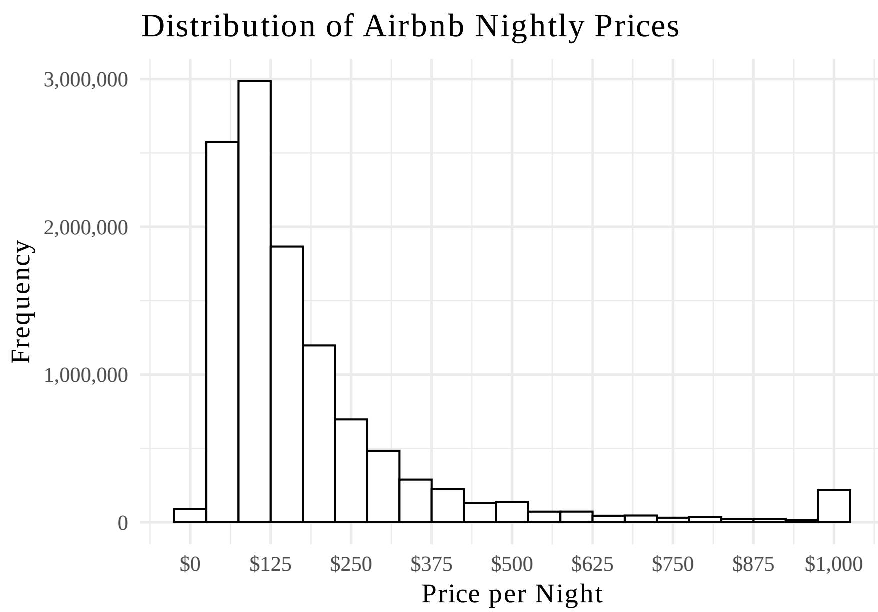

Figure 2.1: Distribution in the price of *Airbnb* listings for the entire dataset with price top coded at $1,000.

and Washington DC. This data contains characteristics of 477,314 unique listings for the years 2016 to 2020.[7] The distribution of *Airbnb* listing prices across the full sample can be seen in Figure 2.1. The median price per night of an *Airbnb* listing is $125 and the majority of listings have a price between $75 and $200 per night. The median *Airbnb* nightly price is a key calibration target.

Since its introduction *Airbnb* has seen heterogeneous growth with some markets growing extremely quickly seeing growth of several hundred percent in only a few years while others have seen relatively slow growth. Figure 2.2 presents the differences in growth between several large United States metro areas. This graph presents the time trend of the indexed number of rooms listed on *Airbnb* across from January 2017 to January 2020. We can see some cities, such as Denver, are growing extremely

[7]While, we use a full sample of cities for the empirical section which gives us a larger sample size and longer observation period, we calibrate the model using just data from Santa Barbara because we have richer data on the housing and hotel market available to us for that region.

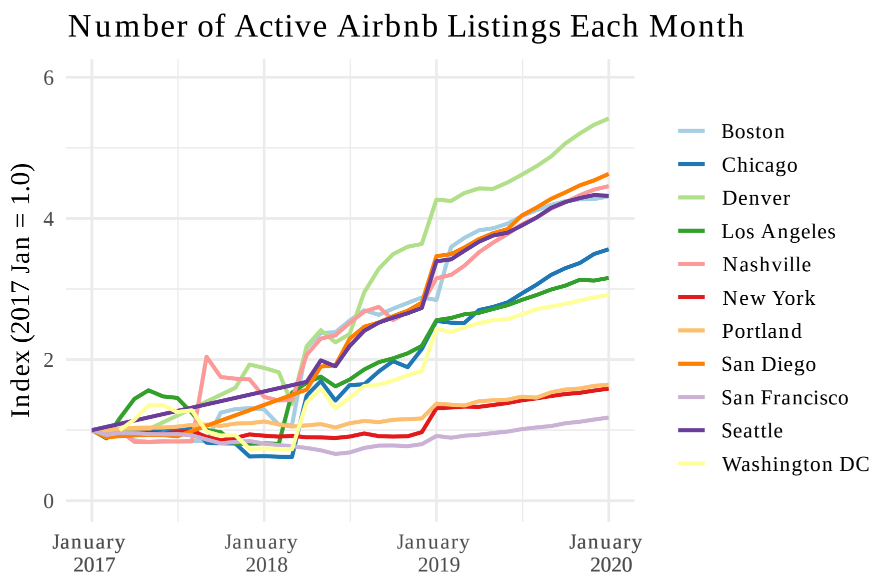

Figure 2.2: Monthly *Airbnb* listings across all room types.

quickly while other cities, especially those that already had a large number of *Airbnb* listings by January 2017, have seen much less growth in listings. There is also significant heterogeneity in the growth of *Airbnb* within metro areas. Figure 2.3 shows the spatial heterogeneity in the growth in the number of rooms listed by ZIP code from January 2017 to January 2020 in the Los Angeles area. We can see that some ZIP codes have seen much faster growth in the number of *Airbnb* listings than others. This spatial heterogeneity is going to be utilized in our regression analysis.

2.2.2 Description of Rental Market Data

Rental market price data comes from *Zillow.com*, an online real estate and rental marketplace company. *Zillow* maintains an online real estate database of over 110 million U.S. homes and estimates housing and rental prices across the United States. Because *Zillow* is used for finding houses and apartments listed for sale or rent, the price and rental rates represent the conditions in the long-term housing market. From

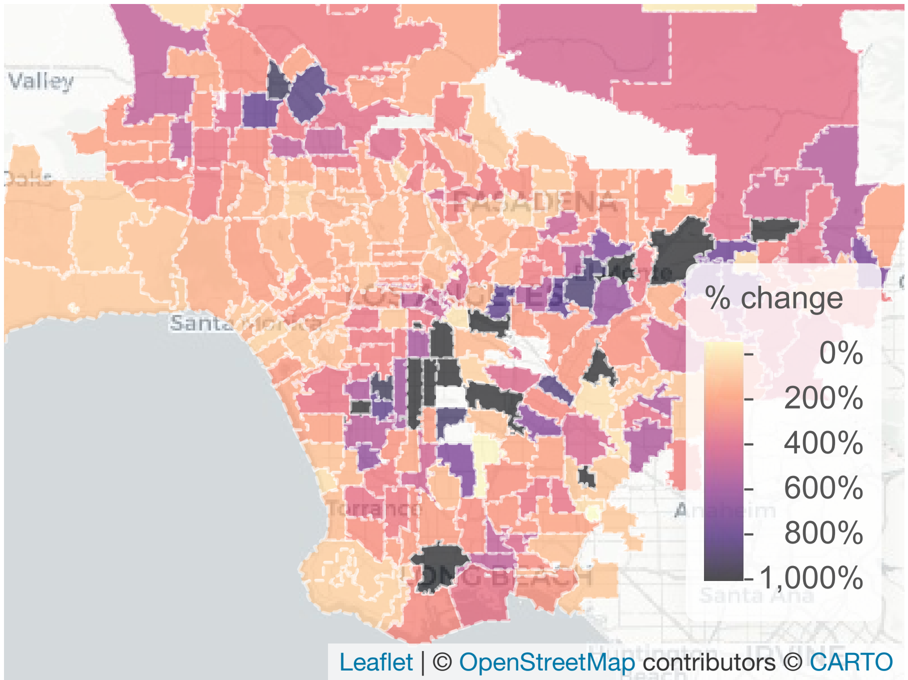

Figure 2.3: Change in the number of *Airbnb* listings from January 2017 to January 2020 by ZIP code in Los Angeles.

Zillow, we use data on the median rental price of apartments of various sizes at the zip code level. In Figure 2.4 we see differences in the changes in the estimated price to rent a 1-bedroom apartment from January 2017 through December 2019. We can also see the heterogeneity of changes within Los Angeles in Figure 2.5. While there isn't large change in the median rent price for a 1-bedroom apartment over this time period in Los Angeles, some ZIP codes in the area saw large increases of up to 40% while others saw very little change.[8]

In addition to *Zillow* housing data, additional housing data and socioeconomic variables come from the *American Community Survey* (*ACS*). From the *ACS*, we use the number of housing units, the number of occupied and vacant units, and the number of renter occupied housing units broken down by number of bedrooms at the ZIP

[8]This observation may be partly driven by area-specific rent controls.

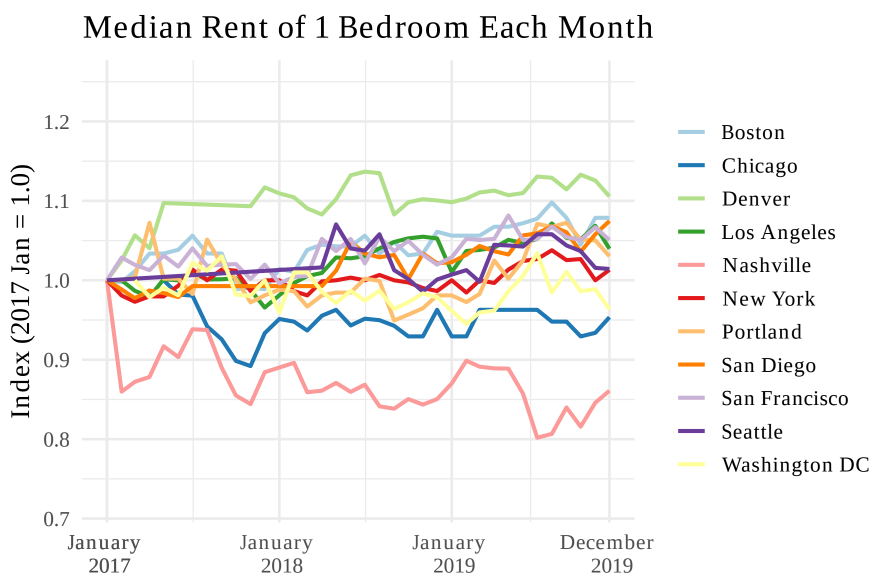

Figure 2.4: Monthly price to rent a 1-bedroom apartment.

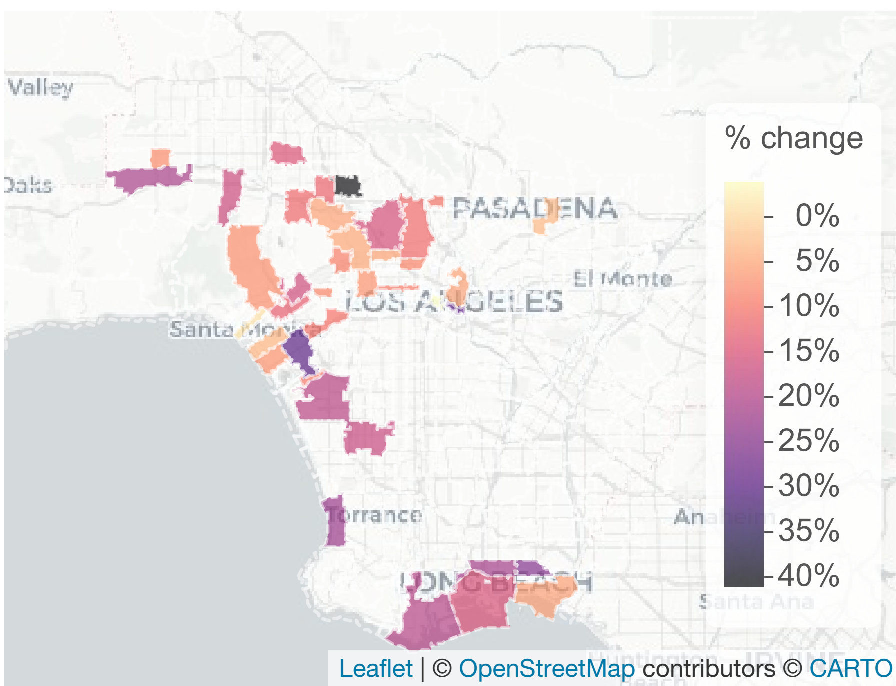

Figure 2.5: Change in the median price to rent a 1-bedroom apartment from January 2017 through December 2019 by ZIP code in Los Angeles.

	min	p25	median	p75	max
December 2016					
Airbnb Listings	1.00	13.00	45.00	156.00	3,145.00
Housing Units	0.00	7,928.75	12,821.00	18,007.25	48,196.00
Rental Units	0.00	2,916.75	5,785.50	10,207.25	32,060.00
Vacant Units	0.00	443.25	778.50	1,253.25	9,173.00
Airbnb Price	20.00	70.94	92.73	133.06	5,000.00
Rent 1 Bd.	1,112.00	1,600.00	1,885.00	2,561.38	4,425.00
December 2018					
Airbnb Listings	1.00	21.50	81.00	250.00	2,925.00
Housing Units	0.00	8,058.75	13,120.50	18,396.25	48,359.00
Rental Units	0.00	2,715.25	5,740.00	10,192.00	32,900.00
Vacant Units	0.00	430.00	784.50	1,368.50	7,799.00
Airbnb Price	18.03	74.65	95.98	143.20	1,162.10
Rent 1 Bd.	750.00	1,525.00	1,800.00	2,405.38	4,995.00

Table 2.1: ZIP Code Level Summary Statistics

code level for each year.[9] These data are derived from the *ACS* 5-year estimates for the years 2014 to 2018. We use the 5-year estimates because this data offers better precision when working within geographic areas with smaller populations such as ZIP codes(Decisions, 2008). Because the *ACS* data are calculated at a yearly frequency, we compute the monthly ZIP code level characteristics using a cubic spline.[10] Summary statistics for the data from *Airbnb*, *Zillow* and the *ACS* are in Table 2.1.

[9]The *American Community Survey* is a continuous survey conducted by the United States Census Bureau. The Census Bureau randomly samples and mails questionnaires to approximately 295,000 addresses every month. The sample does not target individuals but rather targets specific addresses to ensure good geographic coverage. The *ACS* defines a housing unit as occupied if it is the current place of residence of the person or group of people living in it at the time of interview, or if the occupants are only temporarily absent from the residence for two months or less. A vacant housing unit is defined as a housing unit that has no occupant living in it at the time of interview or the unit is occupied at the time of interview entirely by persons who are staying two months or less and who have a more permanent residence elsewhere. An occupied housing unit which is not occupied by the owner is considered as renter-occupied.

[10]Additional details can be found in section B.1

2.2.3 Relationship between *Airbnb* and Rental Markets

As seen in Figure 2.2, the number of *Airbnb* listings has seen large growth in many metro areas from 2017 to 2020. Visually, looking at Figure 2.2 and Figure 2.4 we can see that the metro area that saw the largest growth in the number of *Airbnb* listings, Denver, also saw the largest growth in the median rent of a 1-bedroom apartment. In fact, we can see that three of the five metro areas that saw the largest growth in *Airbnb* listings in this time frame (Denver, San Diego, and Boston) were the metros that saw the highest increases in the median rent of a 1-bedroom apartment. This limited evidence suggests that there may exist a relationship between the number of *Airbnb* listings and rental prices.

Looking to Table 2.1, we can see that the 25^{th} percentile, median, and 75^{th} percentile number of renter occupied units all decreased in this time span. This fact taken along with the increasing number of *Airbnb* listings and increasing number of vacant units (units occupied entirely by individuals who are staying two months or less are classified as vacant) could suggest housing units previously used as long-term rental units may have been converted to units used as short-term rental units. However, these aggregate trends may not reveal the true relationship.

As we saw in Figure 2.3 and Figure 2.5, not only is there heterogeneity between cities there is heterogeneity at the ZIP code level within cities. To formally evaluate the relationship between the number of *Airbnb* listings and the median rental price at the ZIP code level we use the following fixed effect specification:

$$\ln\left(y_{zmt}\right) = \beta \ln\left(Airbnb_{zmt}\right) + X_{zmt}\gamma + \eta_z + \tau_t + \mu_{month} + \varepsilon_{zmt}, \tag{2.1}$$

where y_{zmt} is the median rental price of a 1-bedroom apartment in ZIP code z, metro

m, and time t. $Airbnb_{zmt}$ is the number of rooms listed on *Airbnb* in a ZIP code in period t. X_{zmt} is a vector of observed ZIP code level characteristics including population, the number of housing units, the number of vacant housing units, and the unemployment rate. We include ZIP code level fixed effects, η_z, to control for time invariant ZIP code level characteristics and year fixed effects τ_t to control for aggregate trends. We also include month fixed effects, μ_{month}, to control for seasonality. Including these fixed effects means we are comparing the rents within a ZIP code in the same month of the year across years with different levels of *Airbnb* listings. Table 2.2 presents results for the regressions. Looking at column (4) we can see that the a one percent increase in the number of *Airbnb* corresponds to an increase in rent prices by 0.01%. The median ZIP code saw approximately a 36.20% year to year increase in *Airbnb* listings, which corresponds to a 0.31% increase in the price of rent. This equals a $5.80 increase in monthly rent at the median rent in the data.

Our findings are similar to the results found in Barron et al. (2018) which translate to a $9 increase in monthly rent. To separate differences in effects by the size of location, we also conduct this analysis for the rental price of two bedroom and three bedroom rentals[11]. The results are presented in the second and third panel of Table 2.2. While these figures are similar for two bedroom rentals, they are about half the size (and estimated with less precision) for three bedroom rentals.

It is important to note that the above analysis only demonstrates the statistical relationship between the number of *Airbnb* listings in an area and the median rental price. With this fairly naive approach, we do not believe that our econometric estimation leads to a causal interpretation of the results. When investigating the impact of *Airbnb* listings on rental market or housing prices, a major issue with identification

[11]The number of observations and clusters are different between each set of regressions because the *Zillow* data is available for different time frames and ZIP codes for the different variables.

	(1)	(2)	(3)	(4)
1 Bedroom				
ln(*Airbnb count*)	0.09***	0.02***	0.01***	0.01***
	(0.0099)	(0.0026)	(0.0022)	(0.0022)
Observations	10,818	10,818	10,818	10,818
Clusters	432	432	432	432
2 Bedroom				
ln(*Airbnb count*)	0.11***	0.02***	0.02***	0.01***
	(0.0088)	(0.0027)	(0.0023)	(0.0023)
Observations	12,127	12,127	12,127	12,127
Clusters	505	505	505	505
3 Bedroom				
ln(*Airbnb count*)	0.11***	0.01***	0.01***	0.005*
	(0.0119)	(0.0032)	(0.0025)	(0.0025)
Observations	5,877	5,877	5,877	5,877
Clusters	254	254	254	254
ZIP code FE		✓	✓	✓
Year FE			✓	✓
Month FE				✓

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

Table 2.2: Relationship between the Number of *Airbnb* Listings and the Median Rental Price

stems from isolating the impact on rental price *apart* from another factor driving both the demand for *Airbnb* and the demand for rental housing. That is, we may only be identifying changing desirability or amenities in an area. Reverse causality is also a concern. As rents rise in an area it may be more likely that individuals choose to rent a room in their apartment or if they own an apartment they may choose to relocate to another area to live and rent that apartment at a higher price.

Past research controls for this endogeneity by controlling for characteristics correlated with desirability such as crime rates, new building permits, and amenities such as nearby restaurants. While these characteristics may be a proxy for amenity values in an area and controlling for them removes some of the omitted variable bias, they are not a perfect measure of amenities and the estimated coefficient may still be biased. Barron et al. create a shift-share instrument interacting the popularity of *Airbnb*, measured by the *Google* search index, with an area's attractiveness to tourists in 2010, measured by the number of establishments in the food service and accommodation industries. The argument for this approach is that an exogenous time trend (i.e. *Google* searches for *Airbnb*) will differentially impact areas based on exposure ("touristiness"). Again, while this may remove some of the concern about omitted variable bias, disentangling the impact of *Airbnb* growth and differential recoveries from the Great Recession based on amenity levels is still a concern.

In an ideal setting, one would observe two identical housing markets and have one allow *Airbnb* rentals while the other limits (i.e. bans) them in some way. Observing a natural experiment where this occurs seems unlikely as places that have passed restrictions on *Airbnb*s tend to be places worried about increasing rents and affordable housing. Without such a natural experiment, identifying the impact of *Airbnb* listings on rental prices requires a factor (instrument) that shifts the number of these listings without directly impacting rental prices. Because the decision to choose to rent a unit through *Airbnb*, as opposed to the long-term market, is one determined by profit incentives, it would seem that any factor that changes the profitability of an *Airbnb* listing might also change it in the long-term housing market.

The use of a model circumvents some of these issues by explicitly adding structure to the endogenous relationships that may confound reduced-form analyses. In particular, we specify the objectives (e.g. profit, utility, or welfare optimization), decisions (e.g. price setting, vacancy posting, where to search), and trade-offs faced by the key agents in these markets. While we abstract away from many specifics about the realities of rental markets, we carefully consider those elements to be of first-order concern, motivating further study of the nuances introduced by extensions once the groundwork has been established.

2.3 A Model of Rental Markets

In this section we formulate a model of rental markets and study its behavior. This model is then extended to evaluate the public policy and welfare implications. It features three distinct decentralized markets distinguished by the types of agents that interact within them: hotels, short-term rentals, and long-term rentals. We think of long-term renters as annual leasers. Peer-to-peer technologies allow some property managers to compete with hotels for those looking for short-term stays. Property managers publish prices and potential tenants direct their search to these postings. Prices and queue times are determined endogenously with market utilities taken as given by these managers, making the equilibrium competitive. Importantly, the three markets are endogenously linked which allows us to investigate policy spillovers.

2.3.1 Environment

Time is continuous, infinite, and agents discount the future at a rate r. To find, purchase, and sell lodging services, property managers and tenants interact in three distinct, frictional dwelling markets.[12] Both property managers and tenants are one of two types, and these types exist in fixed measures. That is, we assume that the number of accommodation-seeking individuals and the number of available properties are fixed.[13] Property managers are either innkeepers ($\mathcal{I}$) or landlords ($\mathcal{L}$) and are endowed with a single dwelling unit that can be vacant and searching for a tenant, or occupied and receiving a flow payment p. Tenants are either visitors ($\mathcal{V}$) or residents ($\mathcal{R}$). If the tenant is accommodated, she receives flow utility $w - p$, otherwise they search for lodging and receive flow utility b, where w and b may vary by type of tenant. Further, an agent's type affects which of the three dwelling markets are available to them. Hotels (H) are available to innkeepers and visitors; short-term rentals (S) are available to landlords and visitors; and long-term rentals (L) are available to landlords and residents. The key feature in the above structure is that residents and innkeepers may only participate in one market, but landlords and visitors may participate in multiple, allowing the behavior in one market to influence outcomes in the others.[14]

Within each dwelling market (hotels, short-term, and long-term; indexed by i)

[12] "Frictions" here arise from the time it takes for a tenant and property manager to coordinate on a contract for accommodation.

[13] Regarding the fixed supply of properties, we argue that this is a reasonable assumption given that development is relatively slow and we concern ourselves with the short-run effect of *Airbnb* on the housing market. Accounting for entry is an interesting problem, and unlikely of first-order concern to our question and therefore left for future study.

[14] Importantly, these markets are *not* necessarily distinguished by location. Rather, we define markets by the types of agents that interact within them. One can think of hotels and rental properties as being spatially distinct and separate, but we abstract from any quality differences that may exist between hotels and short-term markets (and within hotels and rental properties more generally). This is done *not* without loss of generality, but to make the model as simple as possible to highlight the first-order effects of the peer-to-peer rental economy on the existing markets.

there is a continuum of sub-markets differentiated by price (indexed by j).[15] Each agent may only participate in one of these sub-markets, which are separate in the sense that search in the ij^{th} sub-market can only produce matches with other agents in that sub-market. Search is assumed to be directed as in Shimer (1996) and Moen (1997). Tenants observe all prices and choose where to search, but within a sub-market search is random. These stochastic, bilateral meetings are governed by a technology that maps the measures of unaccommodated tenants and vacant dwellings into matches: $m^i(u^{ij}, v^{ij})$. m^i is assumed to be increasing and concave in both arguments, has continuous derivatives, and satisfy constant returns. We allow the function (namely its parameterization) to vary by market. Further, let $m^i(u^{ij}, v^{ij})/v^{ij} = m^i(1/\theta^{ij}, 1) \equiv \lambda^i(\theta^{ij}) \equiv \lambda^{ij}$ denote the rate at which a property manager meets an unaccommodated tenant, where $\theta^{ij} \equiv v^{ij}/u^{ij}$ is the "tightness" of sub-market ij, and that $\lim_{\theta \to 0} \lambda(\theta) = \infty$ and $\lim_{\theta \to \infty} \lambda(\theta) = 0$. From the perspective of an unaccommodated tenant, the rate at which she finds a vacant dwelling is given by $m^i(u^{ij}, v^{ij})/u^{ij} = m^i(\theta^{ij}, 1) = \theta^{ij}\lambda^i(\theta^{ij})$.

We follow the literature in assuming that sub-markets are formed by a market maker who posts p^{ij} for each sub-market. Then, both property managers and tenants choose which sub-market to search in. Any sub-market that fails to attract tenants or managers is assumed to be costlessly shut down. As noted in Rogerson et al. (2005), the assumption of a market maker is a convenience and isomorphic to assuming that one side of the market posts prices and the other side directs to these postings. This is appropriate in the context of rental markets as tenants typically search for listings

[15] The assumption of a continuum is a mathematical convenience and will allow us to later write the problem faced by property managers as a well-defined optimization problem in a continuous domain. That is, it enables us to take derivatives. Further, it is important to note that this continuum is a feature of the environment, *not* the equilibrium. We later show that all search activity occurs in a single sub-market for each market. In other words, only "one" sub-market within each market will attract a positive measure of searchers and vacant listings in equilibrium.

and take the price as given.[16] To post a vacancy in any sub-market j within market i, property managers must pay a flow cost κ^i. In the abstract we interpret these costs as reflecting all technologies (e.g. physical, digital, or legal) that allows certain land to be sold to individuals for (temporary) residence. More concretely, we interpret differences in this cost by market as capturing differences in the regulatory structure associated with allowing a property to be sold to a tenant. Matches, i.e. tenant-manager pairs, are assumed to dissolve at a rates $\delta^{\mathcal{V}}$ and $\delta^{\mathcal{R}}$ for visitors and residents, respectively. Capturing differences in preferences between the two groups, it is assumed that $\delta^{\mathcal{V}} > \delta^{\mathcal{R}}$. A visual schematic of the environment is given in Figure 2.6.

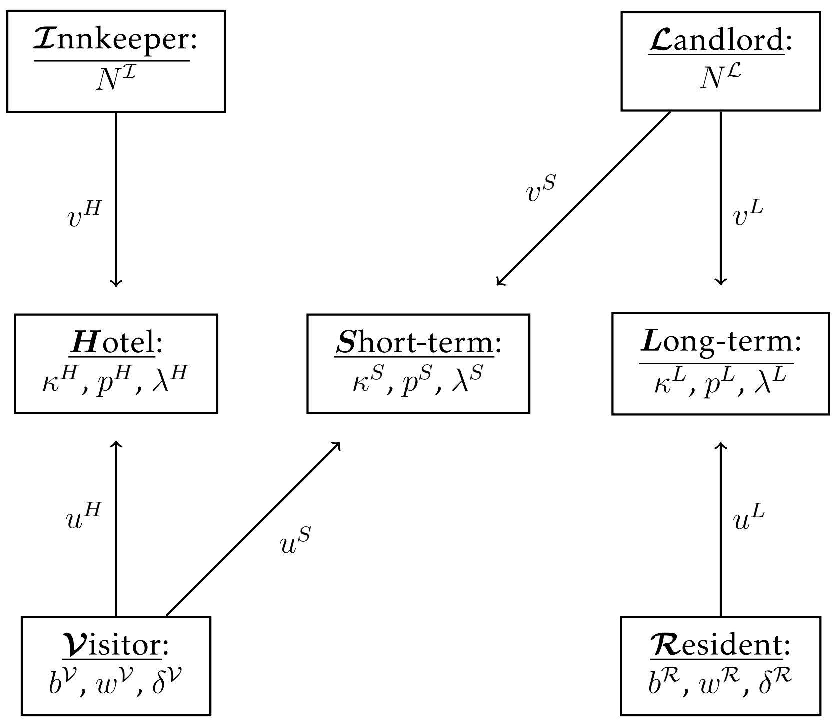

Figure 2.6: Schematic summary of the model environment.

[16] In contrast, this assumption would be less desirable when considering the more general housing market, where alternating offers and negotiations are more prevalent.

2.3.2 Value Functions

In this section we recursively formulate the value functions for each agent type $(\mathcal{R}, \mathcal{V}, \mathcal{I}, \mathcal{L})$ in sub-market ij depending on whether or not they are currently matched $\{0, 1\}$. First consider a resident searching for accommodation. Residents may only search for long-term rentals, but choose which sub-market j to search in. While searching she receives flow utility $b^{\mathcal{R}}$ plus the expected gain from locating a dwelling.

$$r\mathcal{R}_0^j = b^{\mathcal{R}} + \theta^{Lj}\lambda^{Lj}\left[\mathcal{R}_1^j(p^{Lj}) - \mathcal{R}_0^j\right] \tag{2.2}$$

An accommodated resident receives flow utility $w^{\mathcal{R}} - p^{Lj}$ plus the expected loss from separating.

$$r\mathcal{R}_1^j(p^{Lj}) = w^{\mathcal{R}} - p^{Lj} + \delta^{\mathcal{R}}\left[\mathcal{R}_0^j - \mathcal{R}_1^j(p^{Lj})\right] \tag{2.3}$$

It will prove useful to substitute Equation 2.3 for $\mathcal{R}_1^j(p^{Lj})$ into Equation 2.2 and simplify.

$$r\mathcal{R}_0^j = \frac{b^{\mathcal{R}}(r + \delta^{\mathcal{R}}) + \theta^{Lj}\lambda^{Lj}(w^{\mathcal{R}} - p^{Lj})}{r + \delta^{\mathcal{R}} + \theta^{Lj}\lambda^{Lj}} \tag{2.4}$$

for $p^{Lj} \leq w^{\mathcal{R}} - b^{\mathcal{R}}$. If the price is higher than the gain from finding lodging, the resident does not search and receives expected utility $b^{\mathcal{R}}/r$.

Visitors may search in either the hotel or short-term rental market, $i \in \{H, S\}$. Given this choice, they also choose which sub-market j to search for. The value functions for unaccommodated and accommodated visitors are given by the following.

$$r\mathcal{V}_0^{ij} = b^{\mathcal{V}} + \theta^{ij}\lambda^{ij}\left[\mathcal{V}_1^{ij}(p^{ij}) - \mathcal{V}_0^{ij}\right] \tag{2.5}$$

$$r\mathcal{V}_1^{ij}(p^{ij}) = w^{\mathcal{V}} - p^{ij} + \delta^{\mathcal{V}}\left[\mathcal{V}_0^{ij} - \mathcal{V}_1^{ij}(p^{ij})\right]. \tag{2.6}$$

As with residents, if no price induces visitors to search, they receive utility $b^{\mathcal{V}}/r$. Substituting out the value of being accommodated produces

$$r\mathcal{V}_0^{ij} = \frac{b^{\mathcal{V}}(r+\delta^{\mathcal{V}}) + \theta^{ij}\lambda^{ij}(w^{\mathcal{V}} - p^{ij})}{r+\delta^{\mathcal{V}}+\theta^{ij}\lambda^{ij}} \tag{2.7}$$

for $p^{ij} \leq w^{\mathcal{V}} - b^{\mathcal{V}}$.

Innkeepers manage property in the hotel market. When searching for a tenant, they must incur a flow cost κ^H but receive an expected gain when the vacancy is filled.

$$r\mathcal{I}_0^j = -\kappa^H + \lambda^{Hj}\Big[\mathcal{I}_1^j(p^{Hj}) - \mathcal{I}_0^j\Big] \tag{2.8}$$

When occupied, the innkeeper receives p^{Hj} plus the expected loss from separation.

$$r\mathcal{I}_1^j(p^{Hj}) = p^{Hj} + \delta^{\mathcal{V}}\Big[\mathcal{I}_0^j - \mathcal{I}_1^j(p^{Hj})\Big] \tag{2.9}$$

for $p^{Hj} \geq 0$. If the price is negative we assume that the dwelling remains indefinitely vacant. Combining the above we have

$$r\mathcal{I}_0^j = \frac{-\kappa^H(r+\delta^{\mathcal{V}}) + \lambda^{Hj}p^{Hj}}{r+\delta^{\mathcal{V}}+\lambda^{Hj}}. \tag{2.10}$$

Finally, landlords may participate in either the short-term or long-term rental markets, $i \in \{S, L\}$, and then additionally choose a sub-market j.

$$r\mathcal{L}_0^{ij} = -\kappa^i + \lambda^{ij}\Big[\mathcal{L}_1^{ij}(p^{ij}) - \mathcal{L}_0^{ij}\Big] \tag{2.11}$$

$$r\mathcal{L}_1^{ij}(p^{ij}) = p^{ij} + \delta^i\Big[\mathcal{L}_0^{ij} - \mathcal{L}_1^{ij}(p^{ij})\Big] \tag{2.12}$$

for $p^{ij} \geq 0$ and where $\delta^i = \delta^{\mathcal{V}}$ if $i = S$ and $\delta^i = \delta^{\mathcal{R}}$ if $i = L$. Eliminating Equation 2.12

we have

$$r\mathcal{L}_0^{ij} = \frac{-\kappa^i(r+\delta^i) + \lambda^i p^{ij}}{r+\delta^i+\lambda^{ij}}. \tag{2.13}$$

2.3.3 Equilibrium

In this section we establish and characterize the model's equilibrium, focusing in particular on the notion of *competitive search equilibria*.[17] One can solve for such an equilibrium by maximizing property managers' profits subject to tenants receiving some fixed level of utility. Since tenants are homogeneous within type, any sub-market that a positive measure of tenants searches in must pay them the same utility. This level of utility received by market participants is called their *market utility*. Agents in the economy take this level of utility as given, but it is determined endogenously in equilibrium.

Consider a resident searching for accommodation and denote her market utility as $\mathcal{R}_0$ where it must be that $r\mathcal{R}_0 \geq b^{\mathcal{R}}$. Plugging this into Equation 2.4 and rearranging gives us the following expression for the relationship she faces between accommodation finding and price for some given level of utility.

$$\theta^{Lj}\lambda^{Lj} = \frac{(r+\delta^{\mathcal{R}})(r\mathcal{R}_0 - b^{\mathcal{R}})}{w^{\mathcal{R}} - p^{Lj} - r\mathcal{R}_0} \tag{2.14}$$

From the above we can see that a searching resident must pay a high price to achieve a high finding rate and receive the market utility $\mathcal{R}_0$. In other words, Equation 2.14 describes her indifference curve. Further, the RHS is continuous and strictly increasing in both p^{Lj} and $\mathcal{R}_0$ on $p^{Lj} \in (-\infty, w^{\mathcal{R}} - r\mathcal{R}_0)$. As the price approaches $w^{\mathcal{R}} - r\mathcal{R}_0$, the

[17]The term "competitive search" equilibrium comes from Moen (1997) and, as explained by Rogerson et al. (2005), can be thought of as the combination of directed search and price posting. As noted earlier, posting with directed search can be made outcome-equivalent to assuming a third type of agent (or fifth in this paper), a market maker, sets up the sub-markets to attract both property managers and tenants.

gain from finding accommodation goes to zero, and the sub-market tightness goes to infinity. If the price is above $w^{\mathcal{R}} - r\mathcal{R}_0$, no resident searches and the sub-market shuts down. A similar argument with similar conditions can be made for a searching visitor. Here, though, her market utility is not only equal for all j within a given market, but also between markets $i \in \{H, S\}$.

$$\theta^{ij}\lambda^{ij} = \frac{(r+\delta^{\mathcal{V}})(r\mathcal{V}_0 - b^{\mathcal{V}})}{w^{\mathcal{V}} - p^{ij} - r\mathcal{V}_0} \tag{2.15}$$

We can thus think of the problems faced by property managers as a choice of sub-market with price p and tightness θ such that the (p,θ) relationships of Equation 2.14 and Equation 2.15 deliver searching tenants their market utility, where this market utility is taken as given. Letting $\theta^H(p^H;\mathcal{V}_0)$ describe this relationship for the hotel market, $\theta^S(p^S;\mathcal{V}_0)$ for the short-term market, and $\theta^L(p^L;\mathcal{R}_0)$ for the long-term market, the problems of property managers can be written as a (profit) maximization problem in θ *or* p given this market utility. In addition to visitors receiving $\mathcal{V}_0$ in both the hotel and short-term markets, landlords must also be indifferent to posting vacancies in the short and long-term markets.

$$\max_{p^H} \mathcal{I}_0\big(p^H, \theta^H(p^H;\mathcal{V}_0)\big) \tag{2.16}$$

$$\max_{p^S} \mathcal{L}_0\big(p^S, \theta^S(p^S;\mathcal{V}_0)\big) = \max_{p^L} \mathcal{L}_0\big(p^L, \theta^L(p^L;\mathcal{R}_0)\big) \tag{2.17}$$

The following lemma establishes that the above is well-defined.

Lemma 1 *Let* $\widetilde{\mathcal{I}}_0 \equiv \sup_{p^H} \mathcal{I}_0\big(p^H, \theta^H(p^H;\mathcal{V}_0)\big)$, $\widetilde{\mathcal{L}}_0^S \equiv \sup_{p^S} \mathcal{L}_0\big(p^S, \theta^S(p^S;\mathcal{V}_0)\big)$, *and* $\widetilde{\mathcal{L}}_0^L \equiv \sup_{p^L} \mathcal{L}_0\big(p^L, \theta^L(p^L;\mathcal{R}_0)\big)$, *where* $\widetilde{\mathcal{L}}_0^S = \widetilde{\mathcal{L}}_0^L = \widetilde{\mathcal{L}}_0$. *Further, assume that* $\widetilde{\mathcal{I}}_0 \geq 0$ *and* $\widetilde{\mathcal{L}}_0 \geq 0$. *Then the property managers' problems are well defined and the argmax in the price domain is achieved in* $[0, w^{\mathcal{V}} - r\mathcal{V}_0)$ *for the hotel and short-term markets, and in* $[0, w^{\mathcal{R}} - r\mathcal{R}_0)$ *for* *the long-term market.*

Proof: *See Appendix B.2.* ■

Notably, the solutions to the problems defined by Equation 2.16, Equation 2.17 are not necessarily unique. Put differently, many combinations of prices and finding rates may deliver tenants their market utility and maximize managers' profits. Given our assumptions on the matching technology, though, the following lemma establishes that there is no price dispersion *within* a market.

Lemma 2 *All property managers within market $i \in \{H, S, L\}$ choose the same price, and this price is a weighted average of each agent's gain from market participation.*

$$p^H = \eta_H(\theta^H)\big(w^{\mathcal{V}} - r\mathcal{V}_0\big) + \big(1 - \eta_H(\theta^H)\big)r\mathcal{I}_0 \tag{2.18}$$

$$p^S = \eta_S(\theta^S)\big(w^{\mathcal{V}} - r\mathcal{V}_0\big) + \big(1 - \eta_S(\theta^S)\big)r\mathcal{L}_0 \tag{2.19}$$

$$p^L = \eta_L(\theta^L)\big(w^{\mathcal{R}} - r\mathcal{R}_0\big) + \big(1 - \eta_L(\theta^L)\big)r\mathcal{L}_0, \tag{2.20}$$

where $\theta\frac{d\lambda}{d\theta}/\lambda \equiv \eta(\theta) - 1$ is the elasticity of the filling rate with respect to θ (and is a number between 0 and 1). Equivalently, $\eta(\theta)$ is the elasticity of the finding rate with respect to θ.

Proof: *See Appendix B.3.* ■

The equilibrium pricing equations Equation 2.18, Equation 2.19, and Equation 2.20 make clear the endogenous relationship between the three dwelling markets. The existence of a technology allowing landlords to compete with innkeepers makes visitors weakly better off (an increase $\mathcal{V}_0$). Given that innkeepers may lose customers, prices and profits in the hotel market will decline. Residents, too, are affected by this technology. With this additional renting channel, profits for landlords are weakly higher and may induce more of the fixed stock of rental units to be posted for short-term stays. This has upward pressure on prices in the long-term market. The introduction of this peer-to-peer technology has unclear welfare effects: though innkeepers and residents are worse off, landlords and visitors are better off. Resolving whether the aggregate welfare effect is positive or negative is therefore a quantitative exercise.

Turning to solve the model, we start by expressing equilibrium market tightnesses as implicit functions of a vacancy's value. Put differently, we derive the demand for vacancies per searcher as functions of their *cost*–i.e. the expected profits that a vacancy commands. Below, we formalize that this relationship is decreasing.

Lemma 3 *Let $\theta^H = \zeta_H(\mathcal{I}_0)$, $\theta^S = \zeta_S(\mathcal{L}_0)$, and $\theta^L = \zeta_L(\mathcal{L}_0)$ be functions that map the expected profits of a vacant dwelling into market tightnesses. In equilibrium, we have that*

$$\frac{d\zeta_H}{d\mathcal{I}_0} < 0, \qquad \frac{d\zeta_S}{d\mathcal{L}_0} < 0, \qquad \frac{d\zeta_L}{d\mathcal{L}_0} < 0.$$

Proof: See Appendix B.4. ■

To close the model we consider the steady state: the inflows into accommodation equal the outflows from it. Let $u^{\mathcal{V}}$ and $u^{\mathcal{R}}$ be the positive, exogenous measures of visitors and residents, respectively. Further, let $N^{\mathcal{I}}$ and $N^{\mathcal{L}}$ be the positive, exogenous measures of hotels and rentals, respectively. Starting with the hotel market, the measure of innkeeper-managed properties equals the sum of all vacant properties and those accommodating visitors: $N^{\mathcal{I}} = v^{\mathcal{I}} + a^{\mathcal{I}}$. In the steady state, the flows into and out of accommodation must be equal. That is $u^{\mathcal{I}}\theta^H\lambda^H = \delta^{\mathcal{V}}a^{\mathcal{I}}$. For ease of notation, define $\widetilde{\lambda} \equiv \theta\lambda$ and let $\chi = u^H/u^{\mathcal{V}}$ be the fraction of visitors searching in the hotel market. We then have

$$\begin{aligned} & N^{\mathcal{I}} = v^{\mathcal{I}} + \frac{\chi u^{\mathcal{V}} \widetilde{\lambda}^{H}}{\delta^{\mathcal{V}}} \\ \iff \quad & N^{\mathcal{I}} = \chi u^{\mathcal{V}} \Big[\zeta_H(\mathcal{I}_0) + \frac{\widetilde{\lambda} \circ \zeta_H(\mathcal{I}_0)}{\delta^{\mathcal{V}}} \Big], \end{aligned} \tag{2.21}$$

noting the substitution of $\theta^H = \zeta_H(\mathcal{I}_0)$.

In the rental market the measure of landlord-managed properties must equal the sum of all vacant properties and those accommodating visitors *and* residents: $N^{\mathcal{L}} = v^{\mathcal{L}} + a^{\mathcal{L}}$. Use the steady state conditions for both short and long-term markets and substituting for θ^S and θ^L.

$$\begin{aligned} & N^{\mathcal{L}} = v^{\mathcal{L}} + \frac{(1-\chi) u^{\mathcal{V}} \widetilde{\lambda}^{S}}{\delta^{\mathcal{V}}} + \frac{u^{\mathcal{R}} \widetilde{\lambda}^{L}}{\delta^{\mathcal{R}}} \\ \iff \quad & N^{\mathcal{L}} = (1-\chi) u^{\mathcal{V}} \theta^S + u^{\mathcal{R}} \theta^L - \frac{(1-\chi) u^{\mathcal{V}} \widetilde{\lambda}^{S}}{\delta^{\mathcal{V}}} + \frac{u^{\mathcal{R}} \widetilde{\lambda}^{L}}{\delta^{\mathcal{R}}} \\ \iff \quad & N^{\mathcal{L}} = (1-\chi) u^{\mathcal{V}} \Big[\zeta_S(\mathcal{L}_0) + \frac{\widetilde{\lambda} \circ \zeta_S(\mathcal{L}_0)}{\delta^{\mathcal{V}}} \Big] + u^{\mathcal{R}} \Big[\zeta_L(\mathcal{L}_0) + \frac{\widetilde{\lambda} \circ \zeta_L(\mathcal{L}_0)}{\delta^{\mathcal{R}}} \Big] \end{aligned} \tag{2.22}$$

$$\tag{2.23}$$

The above two conditions describe the steady state equilibrium conditions for properties managed by innkeepers and landlords. These two equations, though, are functions of three endogenous variables: $\mathcal{I}_0$, $\mathcal{L}_0$, and χ. Recalling that χ is the share of searching visitors in the hotel market, we pin down its value with the indifference condition of visitors–i.e. that visitors are indifferent between search in the hotel and short-term markets. To do so, separately rearrange Equation 2.13 for p^S and p^L. Plugging these into Equation 2.7 and equating them between markets, we have

$$\frac{b^{\mathcal{V}}(r+\delta^{\mathcal{V}})+\theta^{H}\lambda^{H}(w^{\mathcal{V}}-r\mathcal{I}_0)-\theta^{H}(r+\delta^{\mathcal{V}})(r\mathcal{I}_0+\kappa^{H})}{r+\delta^{\mathcal{V}}+\theta^{H}\lambda^{H}} = \frac{b^{\mathcal{V}}(r+\delta^{\mathcal{V}})+\theta^{S}\lambda^{S}(w^{\mathcal{V}}-r\mathcal{L}_0)-\theta^{S}(r+\delta^{\mathcal{V}})(r\mathcal{L}_0+\kappa^{S})}{r+\delta^{\mathcal{V}}+\theta^{S}\lambda^{S}}.$$

This describes an implicit relationship between market θs that we write as $\xi_H(\theta^H;\mathcal{I}_0) = \xi_S(\theta^S;\mathcal{L}_0)$. Rewriting θ^H in terms of known quantities and χ and substituting for θ^S,

$$\xi_H \circ \frac{\delta^{\mathcal{V}} N^{\mathcal{I}} - \chi u^{\mathcal{V}}\widetilde{\lambda} \circ \zeta_H(\mathcal{I}_0)}{\delta^{\mathcal{V}}\chi u^{\mathcal{V}}} = \xi_S \circ \zeta_S(\mathcal{L}_0). \tag{2.24}$$

Definition 1 *A steady state, competitive search equilibrium is a set of values* $\{\mathcal{V}_0, \mathcal{R}_0, \mathcal{I}_0, \mathcal{L}_0\}$, *prices* $\{p^H, p^S, p^L\}$, *and quantities* $\{\theta^H, \theta^S, \theta^L, \chi\}$ *that solve the following equations.*

$$N^{\mathcal{I}} = \chi u^{\mathcal{V}}\Big[\theta^H + \frac{\theta^H\lambda^H}{\delta^{\mathcal{V}}}\Big] \tag{2.25}$$

$$N^{\mathcal{L}} = (1-\chi)u^{\mathcal{V}}\Big[\theta^S + \frac{\theta^S\lambda^S}{\delta^{\mathcal{V}}}\Big] + u^{\mathcal{R}}\Big[\theta^L + \frac{\theta^L\lambda^L}{\delta^{\mathcal{R}}}\Big] \tag{2.26}$$

$$r\mathcal{V}_0 = \xi_H \circ \frac{\delta^{\mathcal{V}}N^{\mathcal{I}} - \chi u^{\mathcal{V}}\theta^H\lambda^H}{\delta^{\mathcal{V}}\chi u^{\mathcal{V}}} = \xi_S\big(\theta^S\big) \tag{2.27}$$

$$r\mathcal{R}_0 = \frac{b^{\mathcal{R}}(r+\delta^{\mathcal{R}}) + \theta^L\lambda^L(w^{\mathcal{R}} - p^L)}{r+\delta^{\mathcal{R}}+\theta^L\lambda^L} \tag{2.28}$$

$$\theta^H = \zeta_H(\mathcal{I}_0) \tag{2.29}$$

$$\theta^S = \zeta_S(\mathcal{L}_0) \tag{2.30}$$

$$\theta^L = \zeta_L(\mathcal{L}_0) \tag{2.31}$$

$$p^H = \eta_H(\theta^H)\big(w^{\mathcal{V}} - r\mathcal{V}_0\big) + \big(1-\eta_H(\theta^H)\big)r\mathcal{I}_0 \tag{2.32}$$

$$p^S = \eta_S(\theta^S)\big(w^{\mathcal{V}} - r\mathcal{V}_0\big) + \big(1-\eta_S(\theta^S)\big)r\mathcal{L}_0 \tag{2.33}$$

$$p^L = \eta_L(\theta^L)\big(w^{\mathcal{R}} - r\mathcal{R}_0\big) + \big(1-\eta_L(\theta^L)\big)r\mathcal{L}_0 \tag{2.34}$$

A graphical representation of the equilibrium is presented in Figure 2.7. In the center column we describe the indifference relation of visitors (top) and residents (bottom). Tenants receive their market utility, paying relatively low prices and finding accommodation slowly, or high prices and finding it quickly. The equilibrium lies along these indifference curves where property managers maximize the expected profits of a vacancy. For innkeepers this is straightforward. For landlords there is the added condition that the expected profits in both short and long-term markets is equal. This highlights the interconnectedness of the three markets. For example, changes that affect residents therefore alter the problems faced by landlords. This affects profit maximization in the short-term markets, and therefore visitors and innkeepers.

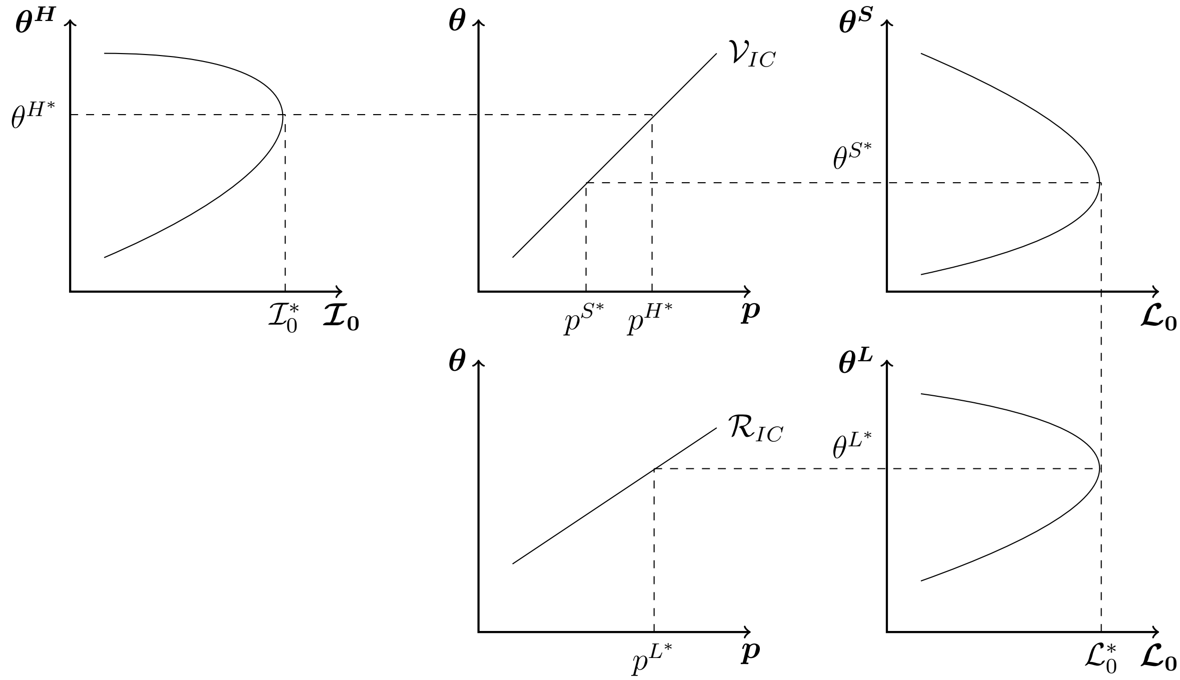

Figure 2.7: Graphical representation of the model's equilibrium.

2.3.4 Comparative Statics

We next present and discuss several exercises moving towards understanding the model laid out above. Though the model is reasonably simple, the interrelatedness of the three markets makes analytic comparative static exercises difficult, if not impossible. We thus rely on the computer to solve and disentangle it. Since the primary focus of this paper surrounds the existence (effect) of peer-to-peer technologies on traditional, lodging markets, we highlight the vacancy posting costs as convenient levers with which to pull. Namely, we can think of taking $\kappa^S \to \infty$ as reflecting the case when peer-to-peer meetings are impossible (or, rather, negligibly rare). When finite, we will later interpret the κ's as the *choices* of a government agent with its own objective function. For now we look at the effects of changes in κ^i on the model's endogenous variables holding all κ^j, $j \neq i$, constant. We report the results of this exercise in Table 2.3.

	$\mathcal{V}_0$	$\mathcal{R}_0$	$\mathcal{I}_0$	$\mathcal{L}_0$	p^H	p^S	p^L	θ^H	θ^S	θ^L
$\uparrow \kappa^H$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$-$	$+$	$+$
$\uparrow \kappa^S$	$-$	$+$	$+$	$-$	$+$	$+$	$-$	$-$	$-$	$+$
$\uparrow \kappa^L$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$+$	$+$	$-$

Table 2.3: Comparative Statics

First consider raising the posting cost of innkeepers, κ^H. This lowers the value of hotel vacancies and, recalling that the equilibrium price is an increasing function of $\mathcal{I}_0$, puts downward pressure on p^H. More visitors are inclined to search for hotels, decreasing the vacancy-to-searcher ratio in H (and increasing it in S). For landlords, the higher market utility enjoyed by visitors hurts them insofar as they must deliver tenants a combination of lower prices and higher finding rates. In response to the lowered profitability in the short-term market, more landlords post in the long-term market (partially undoing the increased tightness in S). Residents thus benefit as they more easily find accommodation at lower prices.

Next, assume that the cost of posting vacancies for short-term rentals, κ^S, increases. The value of unoccupied rentals declines and leads landlords to post more vacancies in the long-term market. This makes residents better off, as there are more vacancies vying for their business at lower prices. Visitors, on the other hand, are made worse off. More are pushed into the hotel market where innkeepers can raise prices alongside filling rates, increasing the value of a vacant hotel room. The fall in market utility for visitors is found to be large enough such that prices in the short-term market actually *increase*. Recalling the equilibrium pricing equation Equation 2.19, the fall in market utility makes the gain from accommodation higher. Though the value of a vacancy drops, the net effect is that visitors must pay more and find accommodation more slowly.

Last, consider raising costs for long-term rentals, κ^L. Profits for landlords are reduced and more prefer to list their vacancies in the short-term market. Because these markets are competitive, p^S and p^L fall. This unambiguously makes visitors better off who enjoy lower prices and faster finding. For residents, the effect is slightly less clear. Accommodation is harder to find, but prices are lower. Though, because accommodation finding is relatively fast, we find that the lowered prices are quantitatively dominant and result in raised resident market utilities.[18] Finally, the value of unoccupied hotel rooms falls as innkeepers must deliver visitors a higher market utility.

Overall, the above exercises demonstrate the importance of modeling all three markets. In models with only two of the three markets, much can be lost when failing to consider the spillovers associated with affecting any one type of agents' decisions. Further, these considerations may also impact notions of optimal policy concerning how short-term accommodation is governed. For example, thinking of changes in κ^H as a government's transient occupancy tax (TOT) policy, the above suggests that increases in this rate *could* benefit residents through multiple channels. Increased TOT revenues may be distributed directly, while indirectly benefiting them by reducing prices, raising finding rates, and lifting market utilities. This of course comes at the cost of property managers (both innkeepers and landlords). κ^S can similarly be thought of the fees charged to *Airbnb*. A lot of discussion has centered around whether or not these peer-to-peer websites should be allowed to operate in certain areas. A "ban" would correspond to $\kappa^S \to \infty$. What the optimal fees should be in each

[18]This result holds for a large portion of the parameter space, and all regions where this model makes sense qualitatively and quantitatively.

market, what the funds are used for or given to, and what the government's objective function is are all explored in the next section.

At this point it is important to discuss the assumption of a perfectly inelastic supply of dwellings. Some obvious detractions are that we know that properties are being developed for housing accommodation over time, and that decisions to develop are inherently tied to their profitability. Notwithstanding, we argue that the largest hold-up for new buildings centers around issues of permitting rather than, say, small changes in a TOT. In this respect, the model should be thought of in a static, short-run context. *Static* because we look at steady states, and *short-run* because of time-to-build restrictions on the construction of new lodging. Put differently, the results concerning the model's policy implications are conditional on there being no entry (or exit) response. Using a previous example, the identifying assumption requires that changes in the TOT do not affect the supply of hotels or rentals.

2.4 Calibration

We calibrate the model using Santa Barbara, California data. We do this for several reasons. The first is that we have detailed data on prices for hotels, Airbnb's, and rentals for the region. *Visit Santa Barbara*[19] provides data on hotel prices and vacancies. They report a sample of 75% of the rooms across their jurisdiction (Santa Barbara, Goleta, Montecito and Summerland), and therefore are estimates with a slight margin of error. These data are provided by "STR" and do not include hostels, vacation rentals or long-term rentals. Also, because these only show hotel room consumption, they do not represent any indicator of total visitor volume (it doesn't include day visitors from our surrounding area). We plot time series of monthly hotel demand

[19] https://santabarbaraca.com/

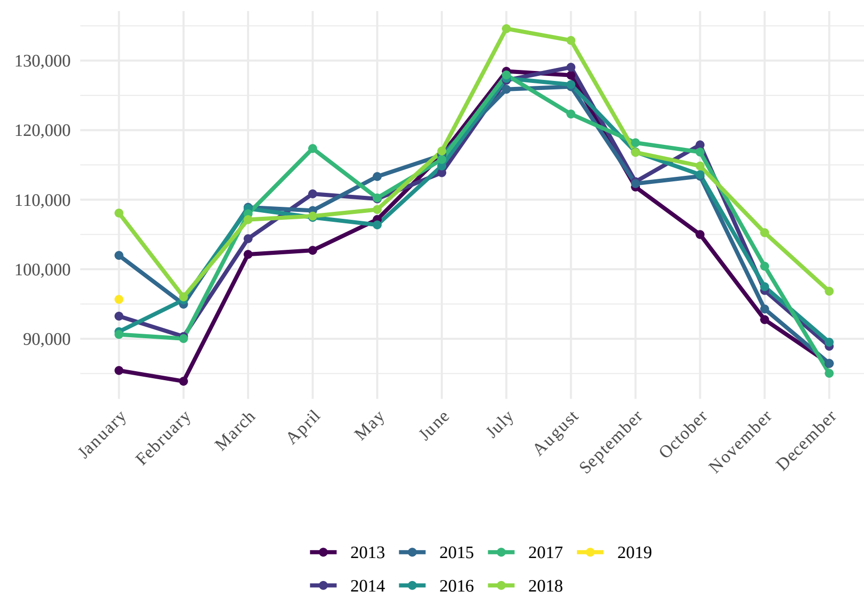

Figure 2.8: Monthly Demand for Hotels in Santa Barbara

in Figure 2.8. We rely on data from *Zillow* for information on rental properties. For comparison with hotel demand, we plot trends in the median rental prices in Figure 2.9. We use *Inside Airbnb* and Tom Slee for data on Airbnb listings. We plot the change in median prices of these listed short-term rentals for the Santa Barbara area in Figure 2.10.

The second reason we calibrate to this region is that the Santa Barbara Coast is fairly isolated along the central coast of California, with very limited expansion potential. Inland mountains prevent building away from the coast, while the coastal commission (paired with what one may call NIMBY sentiments) greatly hinders vertical construction. Since entry is impossible in the model, we view this as a near-ideal scenario to study and assess the policy and welfare implications of peer-to-peer technologies on communities. Indeed, the concern for affordable housing is an important topic for Santa Barbara residents and is a key topic for local politicians. We hope that the following exercises will provide insights into notions of optimal policy regarding

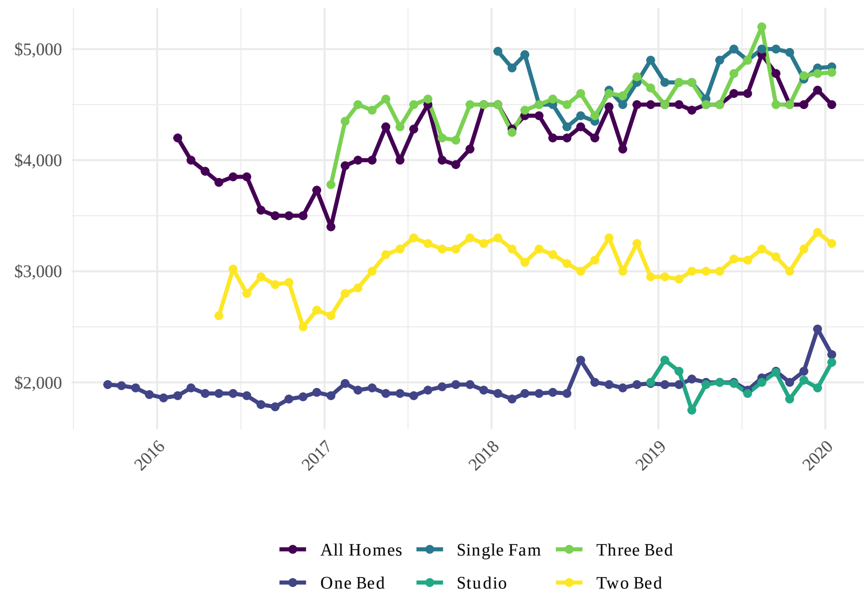

Figure 2.9: Median Monthly Rent in Santa Barbara

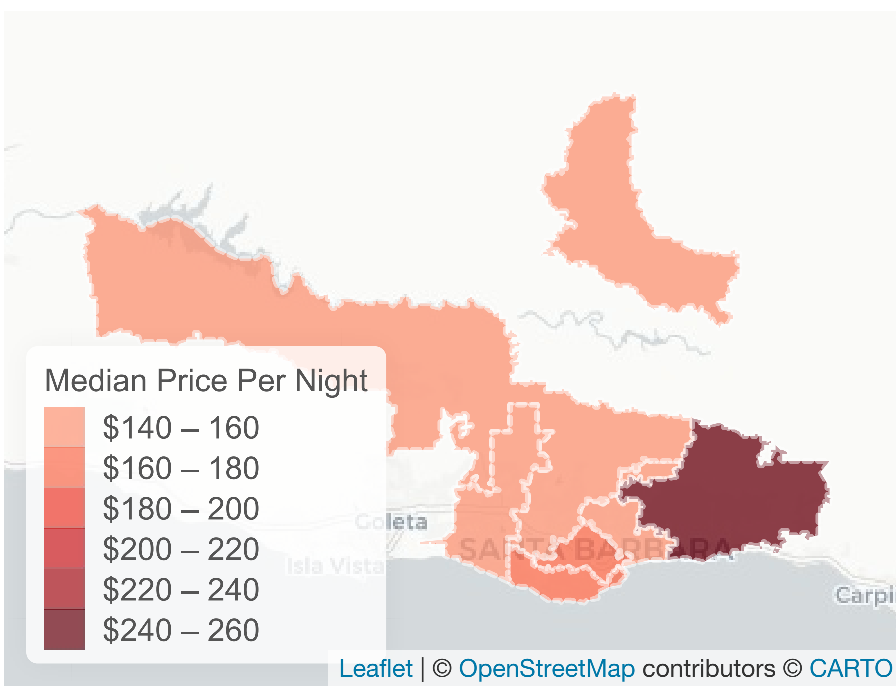

Figure 2.10: Median Price of Airbnb listings from January 2017 to July 2017 by ZIP code in Santa Barbara.

the effects of Airbnb and traditional lodging markets.

To calibrate the model we must make some functional form assumptions on the matching function. We assume that m^i is isoelastic and parameterized by $\eta^{\mathcal{I}}$ and $\eta^{\mathcal{L}}$ for each respective property manager. This assumption is attractive for two reasons. First, it reflects the idea that similar increases in the vacancy-to-searcher ratio may differentially impact vacancy filling for innkeepers or landlords. Second, it is quantitatively necessary to generate the price dispersion observed in data. Recalling the equilibrium pricing equations Equation 2.18 - Equation 2.20, a structurally rigid assumption of uniform elasticities *across* markets can produce only small differences in market prices faced by visitors. In total there are 16 parameters, 10 we set directly and 6 we jointly calibrate. We group these parameters into those related to preferences (r, $w^{\mathcal{V}}$, $w^{\mathcal{R}}$, $b^{\mathcal{V}}$, $b^{\mathcal{R}}$, $\delta^{\mathcal{V}}$, $\delta^{\mathcal{R}}$), search ($\eta^{\mathcal{I}}$, $\eta^{\mathcal{L}}$, $u^{\mathcal{V}}$, $u^{\mathcal{R}}$, κ^H, κ^S, κ^L), and the stock of lodgings ($N^{\mathcal{I}}$, $N^{\mathcal{L}}$).

To begin with the preference parameters, we calibrate the model to the daily frequency with discount rate r of 0.00013, corresponding to an annual discount rate of 5%. $\delta^{\mathcal{V}}$ and $\delta^{\mathcal{R}}$ are set so that the average stay for a visitor and residents, respectively, match what is observed in the data. According to *Visit Santa Barbara*, the average length-of-stay for tourists in 2017 was 2.8 days. The associated daily separation probability is therefore $1/2.8$. Converting this to a rate, we set $\delta^{\mathcal{V}} = 0.442$. For residents we assume annual leases, implying a separation rate of $\delta^{\mathcal{R}} = 0.0027$. Flow utilities for unaccommodation, $b^{\mathcal{V}}$ and $b^{\mathcal{R}}$, are unidentified and therefore normalized to zero. Those for accommodation, $w^{\mathcal{V}}$ and $w^{\mathcal{R}}$, are jointly calibrated such that prices paid by visitors and residents match the data. Utilizing January 2017 through July 2017 data, the median (nominal) price for an Airbnb in Santa Barbara is \$158.83 (*Tom Slee*). For the same time period, the median price for a hotel room was \$245.14 (Visit SB). Finally, Using Zillow listings for two bedroom apartments, the median per-room rental price is \$49.29.

Moving on to the search parameters, the matching function is characterized by $\eta^{\mathcal{I}}$ and $\eta^{\mathcal{L}}$. Unlike in the labor market context, there is a dearth of scholarship on the finding and filling rates of rental properties (or their elasticities w.r.t. market tightnesses). As a result, we reduce the number of parameters to calibrate by letting $\eta^{\mathcal{I}} = 0.5 + \eta$ and $\eta^{\mathcal{L}} = 0.5 - \eta$ so that we capture the *spread* in elasticities with a single parameter. Since these are free parameters (here centered at 0.5), we later evaluate our results' sensitivity to them and find that they are quantitatively robust for reasonable values.[20] Because *both* hotels and Airbnbs pay transient occupancy taxes, we calibrate κ^H and κ^S so that they correspond to the 12% TOT in Santa Barbara. Since TOTs are paid as a rate for the whole stay–and in the model they are paid in flow prior to the stay–we thus calculate the expected cost of search relative to the expected gain from filling a vacancy. In the model the total cost of a vacancy is κ/λ and the flow revenue is given by p/δ. Thus we have $\kappa\delta/\lambda p = 0.12$ for both the H and S markets. Rental properties do not pay any tax (other than property taxes), so we set $\kappa^L = 0$.

The final four parameters ($u^{\mathcal{V}}$, $u^{\mathcal{R}}$, $N^{\mathcal{I}}$, $N^{\mathcal{L}}$) are calibrated to match observed populations of rental market participation observed in data. We normalize the total measure of dwellings to equal 100. The fraction of innkeepers and landlords is then directly calculated using *ACS* and *Visit Santa Barbara* estimates on the total number of rental and hotel properties, respectively. From January through July 2017, there were 4,657 hotel rooms and 50,874 rental units, so $N^{\mathcal{I}} = 8.39$ and $N^{\mathcal{L}} = 91.61$. $u^{\mathcal{V}}$ and $u^{\mathcal{R}}$ are jointly calibrated to target the average measures of hotel and rental vacancies, respectively. Using the same data sources as before, there are an average of 1,075 unoccupied hotels and 4,601 vacant rental units ($v^H = 1.94$ and $v^{\mathcal{L}} = 8.29$). The care-

[20] By "reasonable" we tested values from $(0.20, 0.80)$ with negligible quantitative impact on the equilibrium values.

ful reader will have noticed that there is one more moment than jointly calibrated parameters above. To make this exercise exactly identified, we lastly use the number of Airbnb vacancies from Tom Slee's data, which find an average of 394 listings ($v^S = 0.71$). A summary of the calibration, and its results, are presented in Table 2.4.

Parameter	Value	Description	Target
Preferences			
r	0.00013	discount rate	5% annual discount rate
$w^{\mathcal{V}}$	1,345.72	visitor's flow utility of accomm.	price of hotels / Airbnbs ($\star$)
$w^{\mathcal{R}}$	8,201.30	resident's flow utility of accomm.	price of long-term rentals ($\star$)
$b^{\mathcal{V}}$	0.0	visitor's flow utility of search	unidentified, normalization
$b^{\mathcal{R}}$	0.0	resident's flow utility of search	unidentified, normalization
$\delta^{\mathcal{V}}$	0.442	visitor's separation rate	average stay of 2.8 days
$\delta^{\mathcal{R}}$	0.0027	resident's separation rate	annual lease
Search & Matching			
η	0.072	spread in matching elasticities	price dispersion ($\star$)
$u^{\mathcal{V}}$	0.34	measure of searching visitors	1,075 hotel vacancies ($\star$)
$u^{\mathcal{R}}$	0.03	measure of searching residents	4,601 rental vacancies ($\star$)
κ^H	17.14	hotel vacancy posting cost	12% TOT ($\star$)
κ^S	8.69	short-term vacancy posting cost	12% TOT ($\star$)
κ^L	0.0	long-term vacancy posting cost	No TOT equivalent
Stock of Lodging			
$N^{\mathcal{I}}$	8.39	measure of hotel units	4,657 hotels
$N^{\mathcal{L}}$	91.61	measure of rental units	50,874 rental units

Moment	Data	Model
average per day hotel price	\$209.50	\$209.50
average per day Airbnb price	\$158.83	\$158.83
average annual rent (per day)	\$49.29	\$49.29
hotel TOT	0.12	0.12
Airbnb TOT	0.12	0.12
number of Airbnb listings	0.709 (394 listings)	0.709
average long-term vacancy filling rate	0.011 (3 months)	0.062

Table 2.4: Results of the calibration. The top panel displays the parameters and the bottom reports the moments targeted in the joint exercise. Jointly calibrated parameters are "starred" in the top panel.

2.5 The Effect of Airbnb on Rental Markets

In this section we use the calibrated model to explore the effect of peer-to-peer rentals, namely *Airbnb*, on rental markets. To do so we compare two economies, the calibrated economy from the previous section and one where we let $\kappa^S \to \infty$. As discussed earlier, the case when κ^S is high can be thought of as a situation where *Airbnb*s are too costly to operate or, equivalently, that peer-to-peer technologies are not yet feasible. In the context of equilibrium quantities in the model, as we let $\kappa^S \to \infty$, χ becomes arbitrarily close to 1. We summarize the steady state equilibria in both models in Table 2.5.[21] We report prices, the share of visitors in the hotel market (χ), the measure of vacancies that landlords post in the long-term market ($\equiv \gamma$), market utilities, and aggregate welfare measures (for specifics about the precise welfare function we use, see Section 2.6).

From the benchmark model, note that the prices are the same (i.e. reproduced from) the calibration exercise. In this regime 87% of visitors search for hotels (13% for *Airbnb*s), and 84% of vacant rental properties are listed in the long-term market. As we squeeze the short-term market into nonexistence, we find some intuitive qualitative results. The presence of *Airbnb* depresses hotel prices as innkeepers must compete with them. Further, prices for long-term rentals increase as landlords must be adequately compensated for not listing in the short-term market. Quantitatively we find modest effects on prices. The average price for a hotel is about \$24.00 (per night) *less* expensive with *Airbnb*. The average room in a rental property is \$1.28 *more* expensive per day (about \$39 more per month). For visitors, added choice in search and lower prices make them better off by about 3% with *Airbnb* competition. Residents, however, are worse off. Property manager vacancy profits mimic these results:

[21]The presented results are found to be quantitatively robust to $\pm$ 5% perturbations to the calibrated parameters around the steady state.

	Prices			Search	
	p^H	p^S	p^L	χ	γ
Benchmark	\$209.50	\$158.83	\$49.29	0.87	0.84
No Airbnb	\$233.25	∞	\$48.11	1.00	1.00

	Values				Welfare	
	$r\mathcal{V}_0$	$r\mathcal{R}_0$	$r\mathcal{I}_0$	$r\mathcal{L}_0$	$r\mathcal{G}$	$r\mathcal{W} \times E5$
Benchmark	1,037.4	8,149.12	45.98	47.14	92.04	7.17
No Airbnb	1,006.7	8,150.33	77.33	45.98	83.48	7.24

Table 2.5: Equilibrium outcomes in the same economy with and without peer-to-peer rentals.

landlords are better off, innkeepers worse off.

To make these numbers comparable to the empirical literature, we use the model to "translate" the above results. In particular, we convert the model's results in terms of an elasticity: "a percent change in the number of *Airbnb* listings is associated with an $X\%$ change in Y." One difficulty in directly making this calculation, though, is that the number of *Airbnb*s in the model is endogenous, so directly manipulating the number of listed short-term rentals is difficult (read *impossible*). Instead we vary κ^S and solve for the model's equilibrium each time. We then use this collection of equilibria to construct a mapping from v^S to equilibrium outcomes. Interpolation of this discrete mapping allows us to uncover the desired statistics. It should be noted, though, that because v^S is varied *through* κ^S for this exercise, we cannot comment on the effect of changes in *Airbnb* listings on the price of *Airbnb*.[22] Results are displayed in Figure 2.11. In the left plot we display the effects on prices; values are displayed on the right. We graphically report the effect for -10% to 10% changes in posted vacancies.

[22]Instead, the exercise produces the effect of changes in the posting cost of *Airbnb*s on *Airbnb* prices.

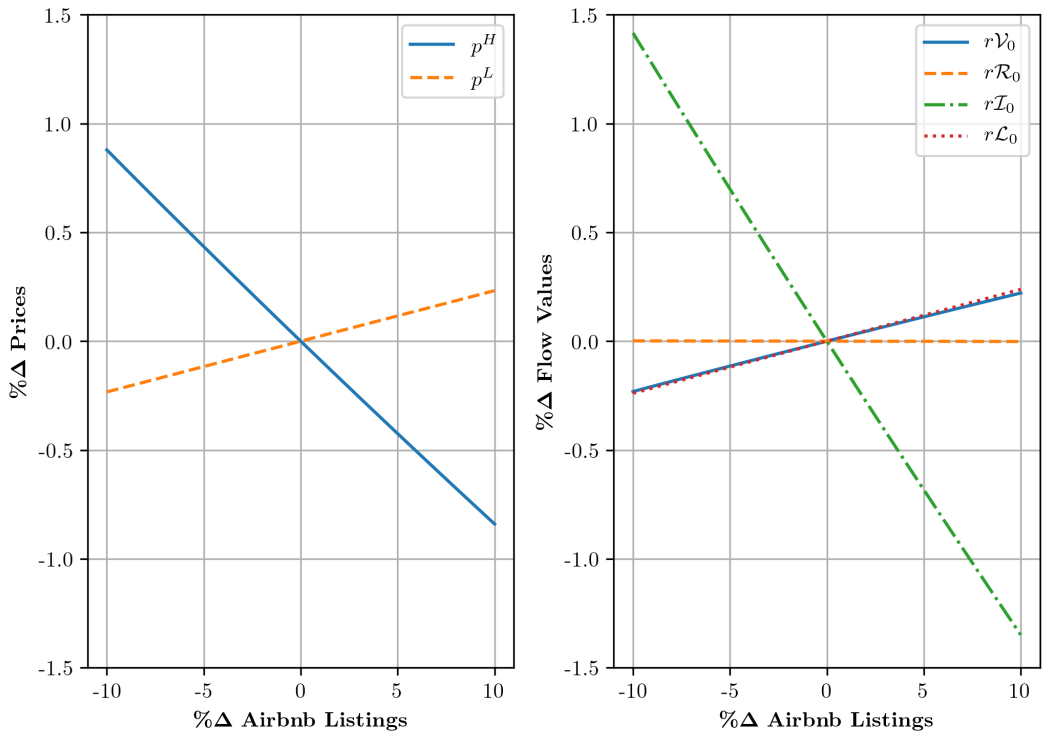

Figure 2.11: Percentage change in equilibrium values from a percentage change in the number of *Airbnb* listings.

A 1% increase in the number of *Airbnb* listings . . .

increases	rents	0.023%
	visitor utilities	0.022%
	the value of a vacant rental	0.024%
decreases	hotel prices	0.086%
	resident utilities	0.0001%
	the value of a vacant hotel	0.138%.

The model numbers are larger than what are found in our reduced-form analyses: the regression results found an effect of 0.01% on the price of rentals. However, they are closer to estimates found in Barron et al. (2018) which suggest an effect of 0.018%. Their paper utilizes data for the entire United States, so quantitative differences may stem from Santa Barbara's relatively unique isolation and development constraints.

More broadly this raises questions about the importance of building constraints in generating heterogeneous effects of *Airbnb*. This important extension is left for future research.

Notwithstanding, the model generally produces effects larger than what is found in empirical studies that we feel the need to comment on. While our structural approach avoids issues of measurement error and reverse-causation, for example, it does abstract in some key ways that might be important. Most notably there is no heterogeneity in quality; differences in price reflect differences in finding / filling rates and market-level search frictions. In reality we know that there are at least some key differences in most short and long-term rentals. Traditional (long-term) rentals are typically unfurnished, while *Airbnb*s are. If we think that *Airbnb*s attract furnished (or some notion of higher quality) properties, we *a priori* predict that that the above price effects would be smaller. Further, we might expect peer-to-peer markets to cater towards medium length–as opposed to short–stays. We have modeled both short-term rentals and hotels to be perfect compliments. Relaxing this might keep hotel prices from dropping as much as we vary the availability of *Airbnb*.

We argue that these results offer a starting point to exploring the general equilibrium effects of peer-to-peer technologies on existing, traditional markets. In addition, we find a relatively surprising normative result regarding *aggregate* welfare: the economy is worse off with *Airbnb*. Phrased differently, what are the sources of inefficiency that can generate lower welfare with *more* choice? In the next section we more formally define the welfare problem and think of this question in the context of optimal policy.

2.6 Policy and Welfare

Now we extend the model above to address questions of public policy and welfare. We interpret the κ's as elements in the choice set of governing agent $\mathcal{G}$. Let $\kappa^i = \kappa + \widetilde{\tau}^i$, $i \in \{H, S, L\}$, where κ is a flow cost charged to properties in all markets.[23]. Because of linear production and utility, we normalize κ to zero. $\widetilde{\tau}^i$ are fees paid to operate in market $i \in \{H, S, L\}$. $\widetilde{\tau}^H$ and τ^S model TOT fees, while τ^L are hypothetical fees placed on traditional rentals.

The government's objective is to maximize aggregate welfare using one of several policies. The different policies limit which markets the government can tax. For example, in one policy we suppose that the government cannot tax traditional rentals, but is free to tax hotels and *Airbnb*s. Deviations away from the unlimited policy (i.e. can levy fees on all markets) are *constrained optimal* and interpreted as "politically feasible" options. Letting $\boldsymbol{\tau} = [\widetilde{\tau}^H, \widetilde{\tau}^S, \widetilde{\tau}^L]'$, the government's objective is to maximize

$$\begin{aligned} \mathcal{W} = \max_{\boldsymbol{\tau}} \Bigg\{ & u^{\mathcal{V}} \mathcal{V}_0(\boldsymbol{\tau}) + u^{\mathcal{R}} \mathcal{R}_0(\boldsymbol{\tau}) + v^{\mathcal{I}}(\boldsymbol{\tau}) \mathcal{I}_0(\boldsymbol{\tau}) + v^{\mathcal{L}}(\boldsymbol{\tau}) \mathcal{L}_0(\boldsymbol{\tau}) \\ & + \frac{\chi(\boldsymbol{\tau}) u^{\mathcal{V}} \widetilde{\lambda}^H(\boldsymbol{\tau})}{\delta^{\mathcal{V}}} \Big[\mathcal{V}_1(p^H; \boldsymbol{\tau}) + \mathcal{I}_1(p^H; \boldsymbol{\tau}) \Big] \\ & + \frac{(1 - \chi(\boldsymbol{\tau})) u^{\mathcal{V}} \widetilde{\lambda}^S(\boldsymbol{\tau})}{\delta^{\mathcal{V}}} \Big[\mathcal{V}_1(p^S; \boldsymbol{\tau}) + \mathcal{L}_1(p^S; \boldsymbol{\tau}) \Big] \\ & + \frac{u^{\mathcal{R}} \widetilde{\lambda}^L(\boldsymbol{\tau})}{\delta^{\mathcal{R}}} \Big[\mathcal{R}_1(p^L; \boldsymbol{\tau}) + \mathcal{L}_1(p^L; \boldsymbol{\tau}) \Big] \\ & + \mathcal{G}(\boldsymbol{\tau}) \Bigg\}, \end{aligned} \tag{2.35}$$

where $r\mathcal{G}(\boldsymbol{\tau}) = v^H(\boldsymbol{\tau})\widetilde{\tau}^H + v^S(\boldsymbol{\tau})\widetilde{\tau}^S + v^L(\boldsymbol{\tau})\widetilde{\tau}^L$. In the above we make clear the dependence of the model's endogenous variables upon $\boldsymbol{\tau}$ through prices, values, and search.

[23]This, for example, includes property taxes which are paid on all types of property

In total there are eight welfare exercises whose outcomes are summarized in Table 2.6 along with the benchmark economy from the calibration. For ease of comparison, we report the policies as tax rates of total revenues ($\tau = \kappa\delta/\lambda p$) and summarize the steady state search behavior with the share of visitors searching for hotels (χ) and the share of vacant rental properties posted in the long-term market (γ). We group and order the exercises by how many markets are allowed to be taxed, beginning with a no-tax case. In this environment there are no government revenues, prices paid by visitors are lower, and prices paid by residents are higher relative to the calibrated benchmark. The short-term market is more attractive for both visitors and landlords, and so a larger proportion redirect their search away from the hotel and long-term market, respectively. A notable theme that will arise in the results to follow is that aggregate welfare can be improved by government intervention. Put differently, the no-tax case does *not* produce a socially optimal allocation of search effort and vacancy posting. Briefly, though search is directed and competitive, barriers to entry result in positive profits for posted vacancies. This opens up the possibility for agent decisions (to search in another market) to not fully internalize their effects on other agents.

Consider the government having access to tax revenues from each market separately. That is, one-by-one we set $\tau^j = 0.0$ for all markets $j \neq i$. When taxing hotels the government sets a high tax rate (49%) to maximize aggregate welfare. Since hotels must compete with each other *and* landlords, hotel prices fall, more visitors search in the hotel market, and market utilities rise. Indeed, they also rise for residents who enjoy lower prices and faster finding rates as landlords shift some vacancies to the long-term market. Here, visitor decisions to search for *Airbnb*s do not adequately compensate innkeepers for the lower filling rates. The government can improve aggregate welfare by taxing hotels to lower prices and induce these visitors back into the hotel market, where tax revenues can be distributed to hotels to compensate for the lost revenues. In the absence of these taxes, individual innkeepers take market utilities as given and so do not have incentives to lower prices and induce more visitors to search in the hotel market. In contrast the governing agent *can* affect market utilities. In this sense one may think of this inefficiency as a coordination problem in the price / market utility space. An atomistic innkeeper cannot improve visitor utility to induce enough short-term searchers to change their search behavior.

A similar intuition also follows when the government can only tax short-term rentals–though through a slightly different channel. Noting from above that visitors are inefficiently searching for *Airbnb*s, landlords also inefficiently re-direct vacancy postings from residents to this market. Since residents have no outside occupancy options (and because they are a large portion of the population) lost utility from lower finding rates and higher prices add up quickly and have large effects on aggregate welfare. Increasing fees in the short-term market can kill both of these birds (with one stone). Higher fees reduces landlords profits and leads more to advertise their vacant rentals to residents. Since both markets are available to landlords, in order to keep at least some posting in the short-term market, p^S must *increase*.[24] These higher prices lead more visitors to search for hotels. Since this policy addresses two sources of inefficiency, aggregate welfare is higher than when only taxing innkeepers. Further, due to the *Airbnb* market being relatively small, this welfare improvement is achieved with very little redistribution.

The last, single-market tax exercise is the long-term market. Very straightforwardly, we find that the optimal fee to place is 0.00 as one of the main sources of inefficiency involves not enough vacancies for residents. This can also be seen in the two-market exercise where the governing agent may tax hotels and long-term rentals.

[24]This differs from the case above because innkeepers do not have an outside option to lodge non-visitors.

We find no tax should be levied in L, and the high, 49% tax should be imposed in H to reallocate search effort. When taxing the two visitor markets, we find a slightly more "balanced" optimal policy wherein *Airbnb* taxes are slightly lower (though virtually identical after rounding) and hotel taxes are present, but small. Finally, the welfare maximizing policy can be achieved by taxing *only* landlords. Here, the government sets taxes on *Airbnb*s so high (71%) that they are effectively nonexistent, bringing back the economy to one where peer-to-peer rentals do not exist. In addition, since all agents have limited choice sets (i.e. no outside option for visitors or landlords), the government can levy a small, redistributive tax on long-term rentals à la the single-market hotel tax case. Because the search externalities can be corrected primarily with the high tax in the short-term market, there is no need to tax the hotel market.

$\mathcal{G}$ Choice Set			Policy			Prices			Search		Welfare	
H	S	L	τ^H	τ^S	τ^L	p^H	p^S	p^L	χ	γ	$r\mathcal{G}$	$r\mathcal{W}$ ($\times E5$)
X	X	X	0.00	0.00	0.00	\$204.21	\$154.14	\$50.56	0.77	0.71	0.00	7.09
✓	X	X	0.49	0.00	0.00	\$184.48	\$150.95	\$48.74	0.94	0.91	311.48	7.20
X	✓	X	0.00	0.33	0.00	\$239.64	\$172.71	\$48.40	0.96	0.96	5.21	7.22
X	X	✓	0.00	0.00	0.00	\$204.21	\$154.14	\$50.56	0.77	0.71	0.00	7.09
X	✓	✓	0.00	0.71	0.03	\$246.57	\$164.40	\$27.20	1.00	1.00	82.60	7.24
✓	X	✓	0.49	0.00	0.00	\$184.48	\$150.95	\$48.74	0.94	0.91	311.48	7.20
✓	✓	X	0.06	0.33	0.00	\$236.71	\$172.38	\$48.18	0.99	0.99	52.72	7.23
✓	✓	✓	0.00	0.71	0.03	\$246.57	\$164.40	\$27.20	1.00	1.00	82.60	7.24
Benchmark			0.12	0.12	0.00	\$209.50	\$158.83	\$49.29	0.87	0.84	92.04	7.17

Table 2.6: Policy experiment results.

2.7 Conclusion

In this paper, we examine the impact of the presence of *Airbnb* listings on the price of long-term rentals and hotels. We do this by developing a structural search and matching model where property managers post vacant rooms and tenants direct their search to these postings. In our model we have three separate but interconnected markets, the hotel market accessible to innkeepers and visitors, the short-term rental market accessible to landlords and visitors, and the long-term rental market accessible to landlords and residents. We then apply our model using a novel dataset for the Santa Barbara, California, housing and hotel markets combining data from several sources including *Visit Santa Barbara*, the *American Community Survey*, *Zillow*, and scraped *Airbnb* listings.

Our results suggest that *Airbnb*s decrease hotel prices by about $24 per night while increasing average rents by $39 per month. While the presence of *Airbnb* creates added choice in accommodation for visitors, increases their flow utility by about 3%, this welfare gain is more than offset by the reduction in welfare for residents due to fewer rentals to search for and higher prices. Overall, we find that with limited entry, aggregate welfare is reduced by the presence of *Airbnb*. As a result, a government policy to set a high transient occupancy tax on short-term rentals would increase aggregate welfare.

While this paper addresses the impact of *Airbnb*s on renters, there are other impacts of *Airbnb* on housing markets that are not accounted for, such as the effect on the price of owning a home. Furthermore, we do not explore the impacts of allowing for the development of new properties.

Chapter 3

Effective Number of Clusters and Inference with Instrumental Variables

with Douglas G. Steigerwald

3.1 Introduction

Empirical studies with clustered data have relied on the number of clusters to determine if critical values from the normal distribution are appropriate for inference. We show that the number of clusters is not the appropriate value for guiding inference when cluster heterogeneity is present. As cluster heterogeneity grows, inference using cluster-robust standard errors and normal critical values can suffer from false rejection of the null hypothesis. This result is more pronounced when using instrumental variables. We develop a measure for cluster heterogeneity when using instrumental variables and using simulations and empirical examples show how this measure can be used to guide inference.

When conducting inference with panel data or with data that can be grouped into clusters, failure to control for within cluster error correlation can lead to downward biased standard errors. Recent work examines the effect of cluster heterogeneity on the ability to control for within cluster correlation. Consistency of cluster-robust standard errors for models without instrumental variables is established in Carter, Schnepel, and Steigerwald (2017), who also provide a measure of the severity of the heterogeneity in a given sample. When cluster heterogeneity grows, inference using cluster-robust standard errors and normal critical values can suffer from false rejection of the null hypothesis. They show that substantial cluster heterogeneity requires the use of critical values that are larger than the normal critical values to produce accurate inference. This result may become more pronounced when using instrumental variables. We may expect the effect of cluster heterogeneity on inference documented for OLS estimates to be even more pronounced for two-stage least squares (2SLS) estimators for several reasons. One, the use of instrumental variables results in fitted regressor values that have less variation then the original regressor. Because the effective number of clusters is a function of the heterogeneity of the covariate matrix for each cluster, this reduction in variation can lead to a smaller effective number of clusters. Weak instruments can further reduce the variation in the fitted regressor and thus the effective number of clusters. A low effective number of clusters can then lead to an increase in the bias in the standard errors and to non-normal asymptotic distributions for the coefficient estimates. We show that the effective number of clusters can be an informative measure for determining the normality of the test statistic.

We use simulations to study the relationship between cluster heterogeneity and the distribution of the test statistic. Utilizing simulations allows for the study of the distribution of the test statistic. We find that when using instrumental variables the effective number of clusters is much lower than the actual number of clusters even when the relationship between the first and second stage error terms is zero, meaning even when the regressor of interest is exogenous using an instrumental variable greatly reduces the effective number of clusters when compared to OLS. We further show that when the effective number of clusters is low the test statistic is not distributed normally. This indicates that using critical values from the normal distribution will lead to rejecting the null hypothesis too frequently. We show that the restricted wild clustered bootstrap can be used in instances when the feasible effective number of clusters is low to return the size of the hypothesis test to the desired level.

We then look to published papers to evaluate how the choice of critical values impacts inference. Using data from empirical studies published in American Economic Association journals we show that inference with instrumental variables and clustered data is sensitive to cluster heterogeneity. Our results verify the observation of Young (2017) who shows using robustness checks that standard instrumental variable inference systematically understates confidence intervals leading to rejection rates greater than nominal size, using 2SLS regressions from published papers. We further show that the effective number of clusters can be used as an indicator to when over-rejection of the null hypothesis may be occurring.

In Section 3.2, we review the current instrumental variable methodology for working with clustered data and discuss the impacts of clustered data on causal inference. In Section **??**, we extend the results of Carter, Schnepel, and Steigerwald (2017) to a model using IV and 2SLS estimators. In Section 3.4, we use simulations to demonstrate the effect of cluster heterogeneity on the rejection rate of hypothesis tests in models with IV and 2SLS estimators and show that the restricted wild clustered bootstrap can be used for more accurate inference. In Section 3.5, we utilize published empirical results to demonstrate how to utilize our results to determine appropriate critical values for inference. In Section 3.6, we conclude and offer suggestions for the application of our results to future empirical research.

3.2 Setup

Data with clustered error structures is widely utilized in applied econometric research. Cameron and Miller (2015) provide a comprehensive overview of clustered-robust methods. Asymptotic theory for clustered based inference when clusters are equally sized is established in multiple papers including Liang and Zeger (1986) and Hansen (2007). Carter, Schnepel, and Steigerwald (2017) and Hansen and Lee (2019) establish asymptotic theory for clustered samples with heterogeneous clusters.

Cluster Robust Variance Estimators often perform well when conducting inference. However, Bertrand et al. (2004) demonstrate that if the number of clusters is small over rejection of the null hypothesis can occur. Carter, Schnepel, and Steigerwald (2017) further show that over rejection not only occurs when the number of clusters is few but can also occur when the number of clusters is large if substantial cluster heterogeneity is present in the data. When the number of clusters is small, bootstrap methods can be used to improve the accuracy of inference. Cameron et al. (2008), MacKinnon and Webb (2017), Djogbenou et al. (2019) demonstrate the improved properties of bootstrap methods with clustered data. Finlay and Magnusson (2019) show how bootstrap methods can improve inference when the number of clusters is small and instrumental variable models are used.

In applied work, researchers often rely on the number of clusters to infer the normality of the test statistic. A popular rule is that if there are 40 or more clusters, then the test statistic is approximately normal. Yet this "rule" relies on homogeneity of the clusters. If clusters are heterogeneous, for example if the number of observations in each cluster differs - then a sample with more than 40 clusters can yield a test statistic that is severely non-normal. This arises because heterogeneity across clusters leads to a larger variance for the estimated standard errors, which in turn leads to non-normality of the test statistic.

3.2.1 Asymptotic Behavior and the Effective Number of Clusters

There are n observations from the linear model

$$y = X\beta + u, \tag{3.1}$$

where the covariate matrix X consists of k linearly independent columns. There are two key features of the model. The first is that $X = [x : W]$ where only W is mean independent of the error u (that is, only W is exogenous). There exists a valid instrument z that is correlated with x and for which $Z = [z : W]$ is mean independent of u. The second key feature is that the observations can be split into G clusters, where errors are independent between clusters. Because x is also independent across clusters, this implies that the covariance matrix of u, given Z, can be written as a block diagonal matrix, where the diagonal element Ω_g is the covariance matrix for cluster g.

In this setting, the estimator of β is the two-stage least squares estimator

$$\widehat{\beta}_{TS} = \left(X^{\mathrm{T}} P X\right)^{-1} X^{\mathrm{T}} P y,$$

where P is the projection matrix $Z(Z^{\mathrm{T}}Z)^{-1}Z^{\mathrm{T}}$. The variance of the two-stage least squares estimator is

$$V := \mathrm{Var}\left[(X^{\mathrm{T}}PX)^{-1}X^{\mathrm{T}}Pu \mid Z\right].$$

Because Ω is block diagonal, this variance can be written as a function of the sum of the variance components for each cluster

$$V = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\left(\sum_{g=1}^{G}\widehat{\mathbf{X}}_g^{\mathrm{T}}v_g v_g^{\mathrm{T}}\widehat{\mathbf{X}}_g\right)(\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1},$$

where $\widehat{\mathbf{X}} = PX$, $v = Pu$, and the subscript *g* denotes the observations for cluster *g*.

Interest centers on hypotheses regarding the coefficients in (3.1), the general form of which is $H_0 : a^{\mathrm{T}}\beta = a^{\mathrm{T}}\beta_0$, where a^{T} is a vector selecting the coefficients under test. Because any factor that multiplies the selection vector cancels out of the test, we can simplify calculations by assuming that $||a|| = 1$. We focus on the cluster robust test statistic of the following form

$$t = \frac{a^{\mathrm{T}}(\widehat{\beta}_{TS} - \beta_0)}{\sqrt{\widehat{\mathrm{Var}}(a^{\mathrm{T}}\widehat{\beta}_{TS})}}$$

where $\widehat{\mathrm{Var}}(a^{\mathrm{T}}\widehat{\beta}_{TS}) = a^{\mathrm{T}}\widehat{V}a$ and $\widehat{V}$ is the cluster-robust variance estimator. The cluster-robust estimator is the sample analog of V

$$\widehat{V} = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\left(\sum_{g=1}^{G}\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{v}_g \widehat{v}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g\right)(\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1},$$

where for each cluster g the observed residuals $\widehat{v}_g$ replace the errors v_g.

The variance, and so the estimator of the variance, depend only on the within-cluster variation, that is, the variance of the cluster-specific estimators of β. As a result, it is the number of clusters, not the total number of observations, that governs the asymptotic behavior of the variance estimator and of the test statistic. To show this, we establish that both V and $\widehat{V}$ can be written as a weighted sum of the variances for $\widehat{\beta}_{TS\,g}$, which is the estimator based only on the observations for cluster g.

Lemma 4 *Under the assumption that $\{u_g\}$ is a sequence of independent random vectors:*

$$V = \sum_{g=1}^{G} C_g \operatorname{Var}(\widehat{\beta}_{TS\,g} \mid Z)\, C_g^{\mathrm{T}},$$

and

$$\widehat{V} = \sum_{g=1}^{G} C_g(\widehat{\beta}_{TS\,g} - \widehat{\beta}_{TS})(\widehat{\beta}_{TS\,g} - \widehat{\beta}_{TS})^{\mathrm{T}} C_g^{\mathrm{T}},$$

where $C_g = (\widehat{X}^{\mathrm{T}}\widehat{X})^{-1}\widehat{X}_g^{\mathrm{T}}\widehat{X}_g$.

We provide a proof in the appendix.

To allow for a wide range of behavior in cluster sizes we study $\widehat{V}_a/V_a$, where $\widehat{V}_a = a^{\mathrm{T}}\widehat{V}a$, rather than $\widehat{V}_a - V_a$. Following Carter, Schnepel, and Steigerwald (2017) (CSS), the main contribution to the mean-squared error of $\widehat{V}_a$ is governed by

$$G^* = \frac{G}{1 + \Gamma^*(\Omega, Z)}. \tag{3.2}$$

The quantity G^* is the effective number of clusters, which adjusts the number of clusters downward to account for the increased variation in $\widehat{V}$ brought about by heterogeneity in the clusters. The heterogeneity is captured by

$$\Gamma^*(\Omega, Z) = \frac{\frac{1}{G}\sum_{g=1}^{G}(\gamma_g - \overline{\gamma^*})^2}{\left(\overline{\gamma^*}\right)^2}, \tag{3.3}$$

where $\overline{\gamma^*} = \frac{1}{G}\sum_{i=1}^{G}\gamma_g^*$ and

$$\gamma_g^* = a^{\mathrm{T}}(\widehat{X}^{\mathrm{T}}\widehat{X})^{-1}\widehat{X}_g^{\mathrm{T}}\Omega_g\widehat{X}_g(\widehat{X}^{\mathrm{T}}\widehat{X})^{-1}a.$$

The quantity $\Gamma^*(\Omega, Z)$ plays an important role in understanding both the asymptotic behavior of the test statistic and the accuracy of the asymptotic approximation in finite samples. To establish an asymptotic distribution we need four conditions.

Assumption 1

i) The error vectors u_g are independent over g with bounded fourth cumulants.

ii) As $n \to \infty$, the number of clusters G grows without bound.

iii) As $G \to \infty$,

$$\frac{\mathbb{E}(\Gamma^*(\Omega, Z))}{G} \to 0.$$

iv) As $n \to \infty$,

$$\frac{1}{V_a} a^{\mathrm{T}} \sum_{g=1}^{G} \left((C_g - \frac{1}{G} I) V (C_g - \frac{1}{G} I) \right)^{\mathrm{T}} a \xrightarrow{\mathbb{P}} 0.$$

Two of the conditions are worthy of discussion. The third condition states no single cluster can have a dominant effect on the overall variation, in that the heterogeneity arising from that cluster will be asymptotically negligible. Of course, in finite samples, there may be a single, or a few, clusters that dominate the heterogeneity. The fourth condition also restricts the dominance of any single cluster, but directly captures the effect of heterogeneity in the explanatory variables weighted by their contribution to the overall variance. With this level of control on the asymptotic heterogeneity we can establish a central limit theorem.

Theorem 1 *Under Assumption 1:*

$\widehat{V}_a$ is a consistent estimator of V_a and, under H_0,

$$t \rightsquigarrow \mathcal{N}(0, 1),$$

where $\rightsquigarrow$ denotes convergence in distribution.

We provide a proof in the Appendix.

Any quantity that is a function of the latent Ω is infeasible and is denoted with a $*$ superscript. To implement these results we follow CSS and construct a feasible quantity by replacing Ω_g with the $n_g \times n_g$ unit matrix:

$$\begin{aligned}
\widetilde{\gamma}_g &= a^{\mathrm{T}}(\widehat{X}^{\mathrm{T}}\widehat{X})^{-1}\left(\sum_{g=1}^{G}\widehat{X}_g^{\mathrm{T}}\iota_g\,\iota_g^{\mathrm{T}}\widehat{X}_g\right)(\widehat{X}^{\mathrm{T}}\widehat{X})^{-1}a,\\
\widetilde{G} &= \frac{G}{1+\widetilde{\Gamma}(\iota_g\,\iota_g^{\mathrm{T}},Z)},
\end{aligned}$$

where ι_g is the $n_g \times 1$ vector of 1's and $\widetilde{\Gamma}(\iota_g\,\iota_g^{\mathrm{T}}, Z)$ is computed as in (3.3) with $\widetilde{\gamma}_g$ in place of γ_g^*. This generally delivers an upper bound for Γ^* and so a lower bound for G^*.

3.3 Behavior of the Effective Number of Clusters

To illustrate how cluster variation affects hypothesis testing in applied settings we turn to simulations. Because earlier work has confirmed the widespread belief that the effect of heterogeneity is most pronounced for tests of regressors that are highly correlated within clusters, we focus on cluster-level explanatory variables. The baseline model for individual *i* in cluster *g* is

$$y_{ig} = \beta_0 + \beta_1 x_g + \beta_2 w_{ig} + u_{ig}, \tag{3.4}$$

with $w_{ig} \sim \mathcal{N}(0,1)$ and $\mathbb{E}(u_{ig} \mid w_{ig}) = 0$. For the endogenous regressor x_g:

$$x_g = \alpha_0 + \alpha_1 z_g + \eta_g, \tag{3.5}$$

where $\mathrm{Cov}(z_g, u_{ig}) = 0$ and $\mathrm{Cov}(x_g, z_g) \neq 0$. There are two settings: x_g is a binary random variable and $z_g \sim$ Bernoulli(0.5); and x_g is a continuous random variable and $z_g \sim \mathcal{N}(0,1)$. The endogeneity arises through the construction of the error term

$$u_{ig} = \rho\eta_g + (1-\rho^2)^{\frac{1}{2}}\nu_{ig}, \tag{3.6}$$

where $\{\eta_g\}$ and $\{\nu_{ig}\}$ are each sequences of i.i.d. $\mathcal{N}(0,1)$ random variables and $\{\eta_g\}$ is independent of $\{\nu_{ig}\}$. Therefore $u_{ig} \sim \mathcal{N}(0,1)$ and $\rho = Corr(u_{ig}, \eta_g)$ measures the strength of the endogeneity. If $\rho = 0$, then x_g is exogenous.

The simulations are based on a sample size of 2,500, divided into 100 clusters. Cluster heterogeneity is controlled by varying the cluster sizes. In the baseline case, all clusters contain 25 observations. The first heterogeneous design has one cluster of 50 observations, with all other clusters having 24 observations. Each additional design increases the size of the large cluster by 25 observations until the large cluster contains half of the sample, which is the most extreme degree of heterogeneity that we consider.

For each cluster-size design, we have 33 simulated settings. Each setting corresponds to a value of ρ and a draw of the exogenous variables, z_g and w_{ig}. We allow ρ to take the values $\{0.0, 0.1, \ldots, 1.0\}$ and, for each value of ρ, we randomly draw the exogenous variables three times. For each of the 33 simulated settings we perform 1000 Monte Carlo simulations by randomly generating the error terms, η_g and ν_{ig}.

3.3.1 Feasible Effective Number of Clusters

In the case where all the explanatory variables are exogenous and the cluster level variable is Bernoulli, Carter, Schnepel, and Steigerwald (2017) show that $\widetilde{G} = G$ with equally sized designs and, as the fraction of the sample in the largest cluster grows, that $\widetilde{G}$ declines exponentially. Figure 3.1 reproduces these results with our designs (we set $x_g = z_g$ to mirror an exogenous covariate). The design in which the clusters are equally sized, so one percent of the observations are in the "largest" cluster, has $\widetilde{G}$ equal to, or nearly equal to, G in all cases. Because the exogenous variables are randomly drawn, they are not perfectly balanced over clusters and variation in the value can lead to a slight reduction in the feasible effective number of clusters.

The use of instruments may alter these results. It is well known that passing the endogenous explanatory variable through the first-stage filter reduces the variation in the variable, which could lead to less heterogeneity and a larger value of $\widetilde{G}$ than is reported in Figure 3.1. We report the results in Figure 3.2. The broad pattern is very similar between these two figures, with a sharp reduction in $\widetilde{G}$ as the fraction of the sample in the largest cluster grows. With 5 percent of the observations in the largest cluster $\widetilde{G}$ is half of G. As the size of the largest cluster approaches 10 percent of the sample, $\widetilde{G}$ falls below 10. The range of values for $\widetilde{G}$ for a given cluster-size design is larger for the endogeonous regressor case because of the additional variation arising from different values of ρ.

Figure 3.3 shows the median value of the feasible effective number of clusters across all simulations for each cluster-size design. The scaling on the horizontal axis is now the coefficient of variation in cluster sizes, which is the standard deviation of the cluster sizes divided by the average cluster size. This measure can be used in empirical settings where there are a number of clusters of unequal size. We report the corresponding value of the coefficient of variation for three empirical papers that are studied in more detail below. In addition to the binary endogenous covariate, a continuous covariate is also included. For the continuous covariate, the effective number of clusters is much lower for equal cluster sizes and declines at a slower rate. The low values for equal or nearly equal cluster sizes are the result of the fact that z_g and w_{ig}

Figure 3.1: Cluster size variation and the feasible effective number of clusters: exogenous binary covariate

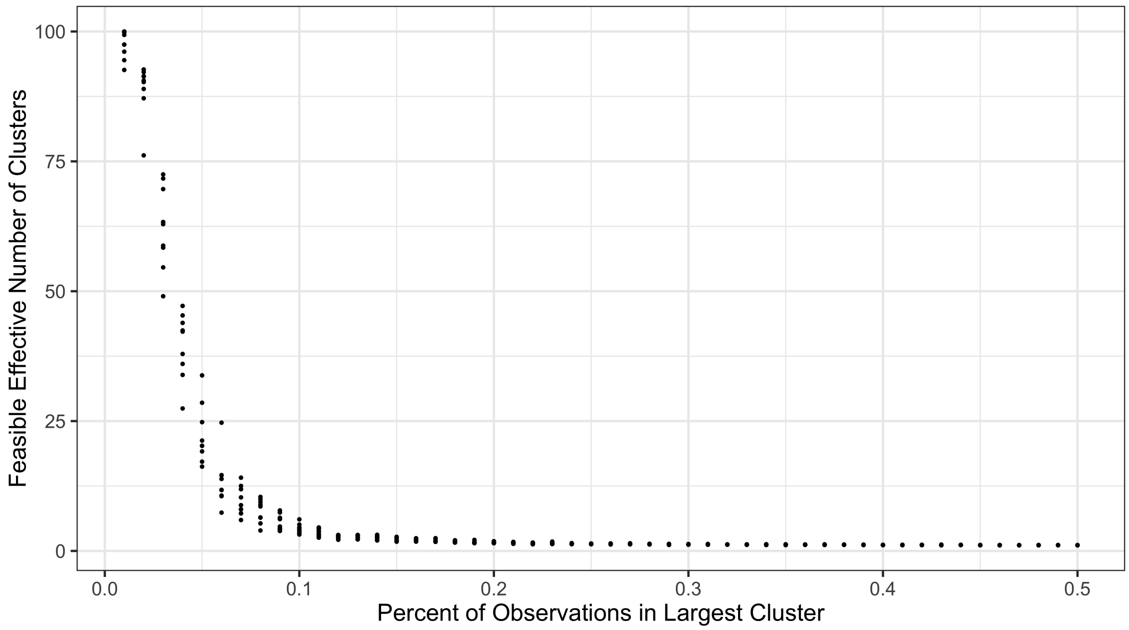

Figure 3.2: Cluster size variation and the feasible effective number of clusters: endogenous binary covariate

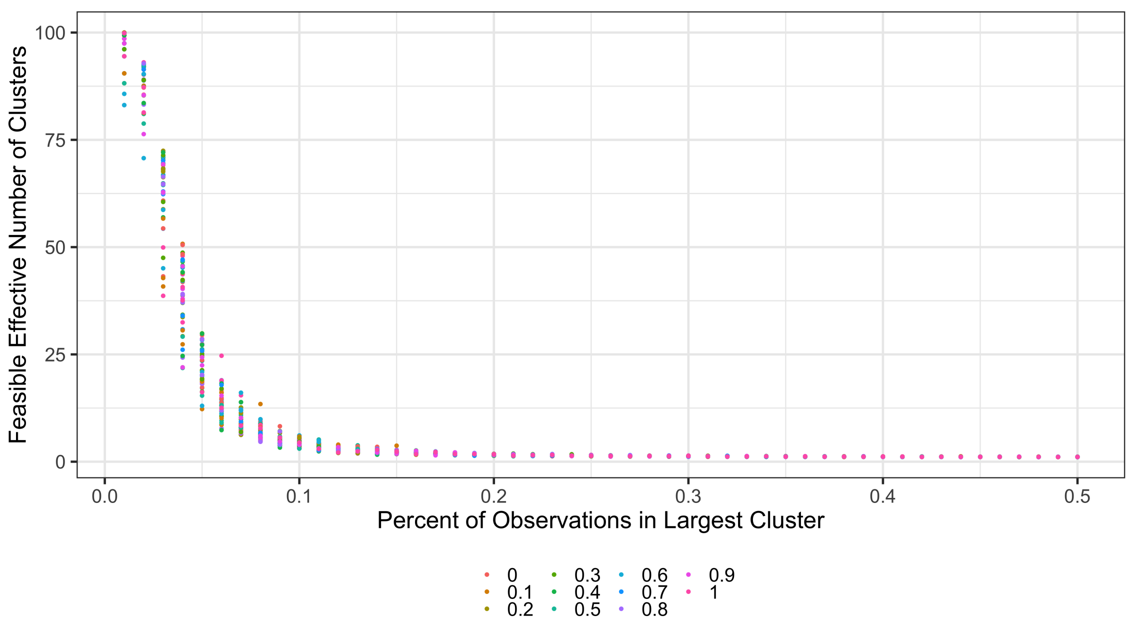

Figure 3.3: The feasible effective number of clusters for a continuous and a binary regressor for varying cluster size coefficient of variation.

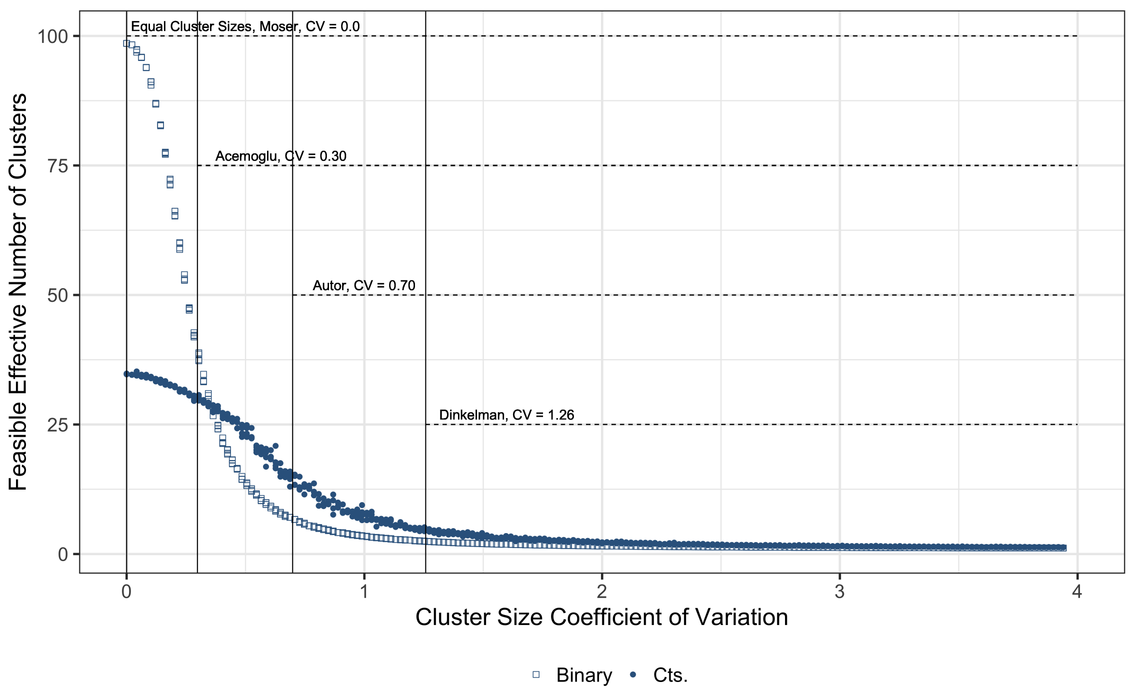

vary more over clusters in the continuous case than in the binary case.

3.3.2 Empirical Test Size

Of interest to empirical researchers is the relationship between feasible effective number of clusters and accuracy of inference. CSS demonstrated that with OLS, when the feasible effective number of clusters is small the empirical test size can be much greater than the nominal size. With 2SLS, for low values of $\widetilde{G}$ the empirical test size can lie far above the nominal size of 5 percent. Interestingly, even with a moderately high level of endogeneity, $\rho = 0.6$, the empirical test size is inflated even if $\widetilde{G}$ is as large as 40, which corresponds to the largest cluster having only 2-3 percent of the sample (the baseline is 1 percent). The binary pattern appears to be different, with size distortions that are much more responsive to ρ.

Figures 3.4 and 3.5 display the empirical test size as a function of the feasible effective number of clusters for a binary and continuous regressor. The cluster-robust IV test statistic of interest, t, is given by $\widehat{\beta} - \beta_0$ divided by the cluster-robust standard error. From these figures it is clear that for small values of ρ the empirical test size is below the nominal size while for larger values of ρ result in an empirical test size greater than nominal size. However, all values of ρ appear to have an increase in empirical test size when the feasible effective number of clusters drops below about 10.

When looking at the average empirical test size when split into low, medium, and high feasible effective number of clusters groups, we can see in Table 3.1 that for both continuous and binary regressors the empirical test size is only much larger than the desired 5 percent for the low feasible effective number of clusters group. To investigate why this may be occurring we look to the distribution of the test statistics of the simulated data.

Table 3.1: Average empirical test size split by low feasible effective number of clusters ($FENC < 15$), medium feasible effective number of clusters ($15 < FENC < 30$), and high feasible effective number of clusters ($FENC > 30$)

	Average Empirical Test Size		
	Low FENC	Medium FENC	High FENC
Continuous	0.098	0.047	0.047
Binary	0.097	0.053	0.049

In Figure 3.6, we can see that with equal cluster sizes for all values of ρ the simulated distributions are all close to the Standard Normal distribution and the feasible effective number of clusters is around 30 for all values of ρ. Once cluster heterogeneity is introduced, as seen in Figure 3.7, the feasible effective number of clusters drops and the distribution of the test statistic becomes non-normal. The distribution when the feasible effective number of clusters is 15.78 is slightly non-normal, while

Figure 3.4: Empirical test size for varying values of the feasible effective number of clusters when the regressor of interest is binary.

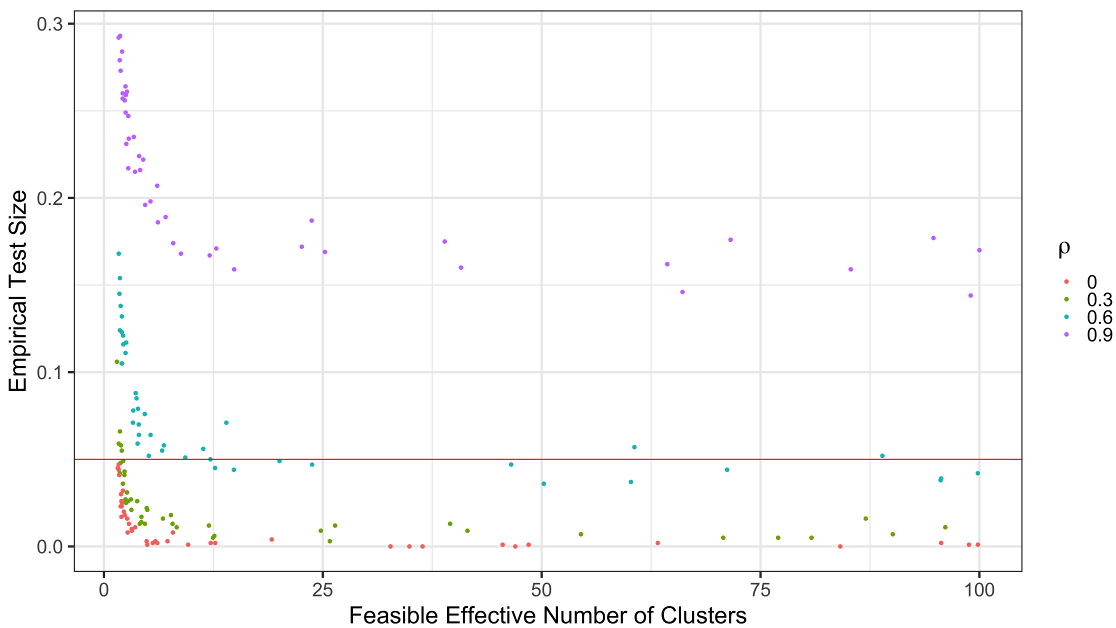

Figure 3.5: Empirical test size for varying values of the feasible effective number of clusters when the regressor of interest is continuous.

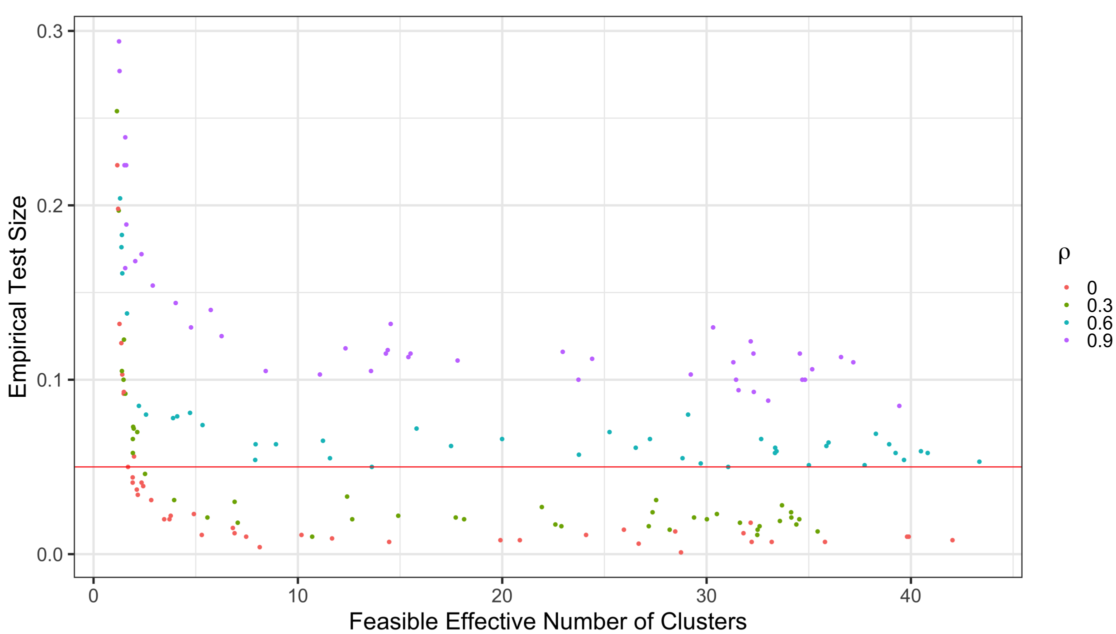

the distributions when the feasible effective number of clusters is near 1 are highly non-normal.

Figure 3.6: Distribution generated from 1,000 simulations for each specification with equal cluster sizes.

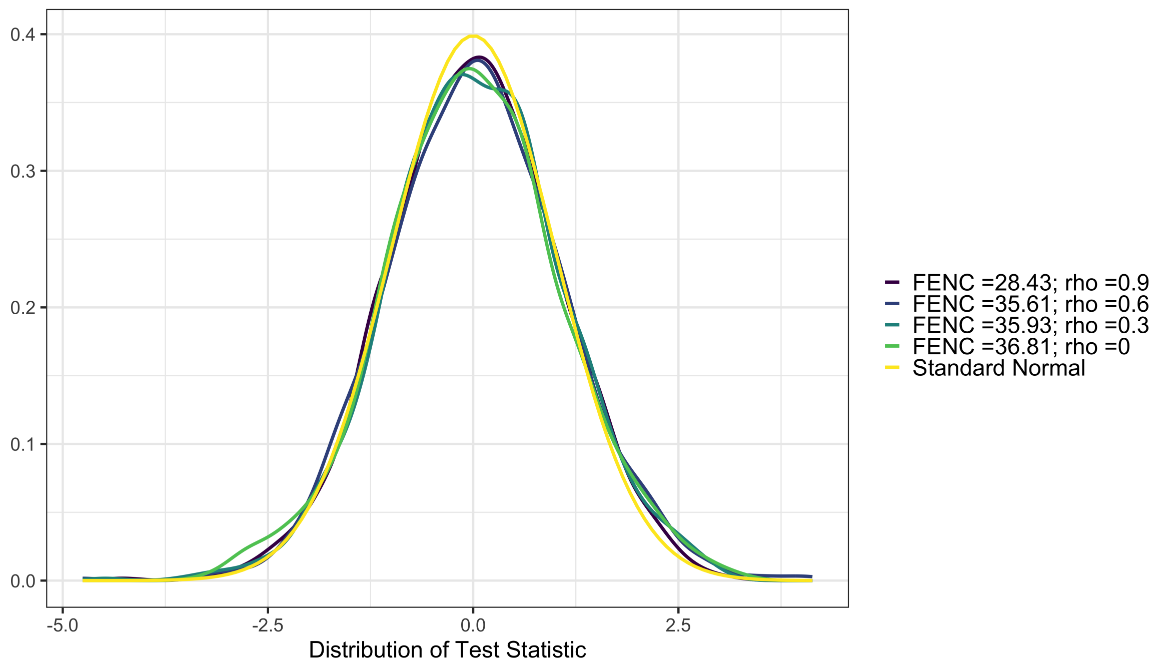

Figure 3.7: Distribution generated from 1,000 simulations for each specification with unequal cluster sizes.

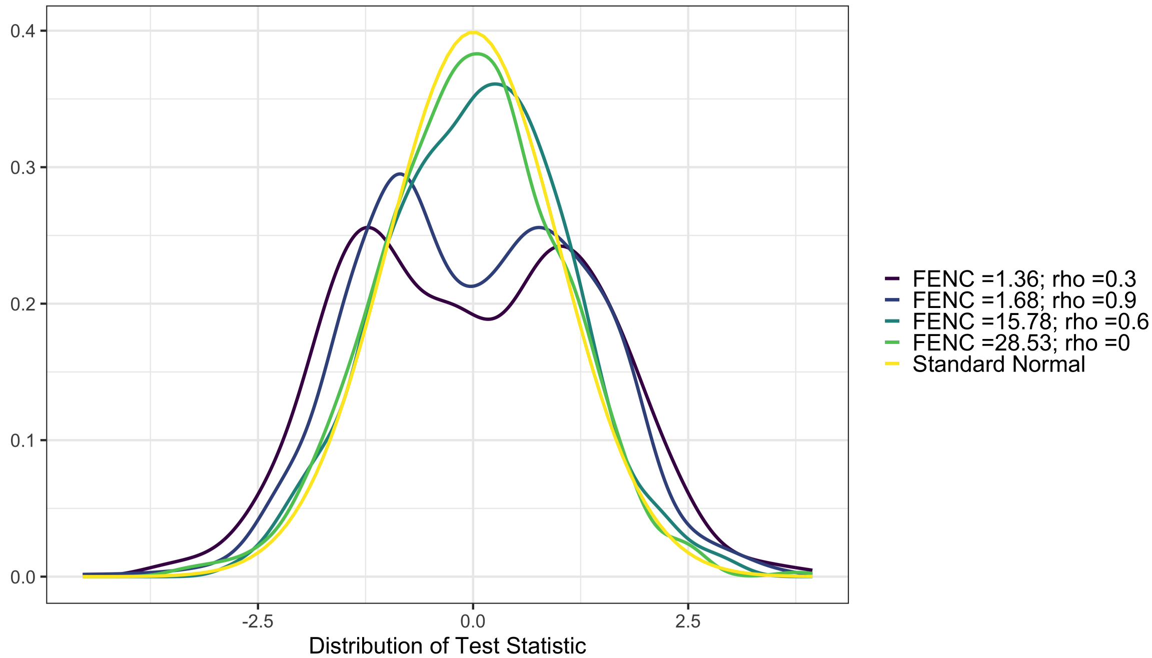

3.4 Non-normality of the Test Statistic

3.4.1 Bootstrap Approximation to the Distribution of the Test Statistic

We have seen from simulations that when using 2SLS with clustered data the empirical test size often does not match the nominal size. We have also seen that the distribution of the test statistic is non-normal. When conducting inference with test statistics of unknown sample distributions, the bootstrap can be used to obtain appropriate critical values. The "wild" bootstrap is useful when we believe that the variance of the error term depends on the value of the covariates. Instrumental variable estimation allows for consistent estimation, however can result in biased estimates. To bootstrap in the IV case we need to estimate the parameters in both the first and second stage regressions. We utilize a procedure adapted from the Simultaneous Equations section of MacKinnon (2019).

Consider the two-equation model

$$y_{ig} = \beta_1 + \beta_2 x_g + \beta_3 w_{ig} + u_{ig}$$

$$x_g = \alpha_1 + \alpha_2 z_g + \alpha_3 w_{ig} + \eta_g$$

where y_{ig} and x_g are endogenous variables and w_{ig} and z_g are exogenous variables.

1. Obtain estimates of the first and second stage parameters.
   (a) Estimate the restricted estimates of the structural equation, $\widetilde{\beta}_1$ and $\widetilde{\beta}_3$, and residuals, $\widetilde{u}_{ig}$, by imposing the null value β_0 and regressing $y_{ig} - \beta_0 x_g$ on w_{ig}.
   (b) Estimate the restricted estimates of the reduced-form equation by regressing x_g on w_{ig}, z_g, and $\widetilde{u}_{ig}$. The estimated restricted residuals, $\widetilde{u}_{ig}$, are included to estimate $\widetilde{\alpha}_1$ and $\widetilde{\alpha}_2$ because this yields more efficient estimates because u_{ig} and η_g are correlated.
2. Calculate the cluster-robust t statistic
3. For each of B bootstrap replications,
   (a) Generate a set of bootstrap disturbances for both stages by multiplying $\widetilde{u}_{ig}$ and $\widetilde{\eta}_g$ by v_g which takes the values -1 and 1 with equal probability.
   (b) Generate x_g^* followed by y_{ig}^* using the bootstrap disturbances and estimated restricted parameters such that

$$x_g^* = \widetilde{\alpha}_1 + \widetilde{\alpha}_2 z_g + \widetilde{\alpha}_3 w_{ig} + v_g \widetilde{\eta}_g$$

$$y_{ig}^* = \widetilde{\beta}_1 + \beta_0 x_g^* + \widetilde{\beta}_3 w_{ig} + v_g \widetilde{u}_{ig}$$

(c) Obtain bootstrap parameter estimates and calculate the bootstrap t statistic.

4. From the distribution of the B bootstrap t statistics calculate the critical values.

Table 3.2: 1000 simulations with the bootstrap performed 1000 times for each specification.

Homogeneous Clusters Sizes				
FENC	36.81	35.93	35.61	28.43
Coverage Probability Robust	0.92	0.94	0.93	0.94
Coverage Probability Bootstrap	0.94	0.95	0.94	0.96
Heterogeneous Cluster Sizes				
FENC	28.53	15.78	1.68	1.36
Coverage Probability Robust	0.94	0.94	0.90	0.86
Coverage Probability Bootstrap	0.96	0.95	0.95	0.95

Using simulations we study the behavior of the bootstrap test statistic for a variety of specifications, the coverage probabilities for a 95% hypothesis test are in Table 3.2. We can see that for all values of the feasible effective number of clusters the coverage probability is approximately 0.95 even for the low values of the feasible effective number of clusters when the coverage probability using normal critical values falls to below 0.90.

3.5 Empirical Examples

Empirical research papers that employ both instrumental variables and cluster robust inference have been spotlighted recently. Young (2017) notes that in a number of these papers there are one or two influential clusters. Indeed, he reports that for a large sample of published papers that report two-stage least squares estimates with statistical significance at 1% size, nearly half are "insignificant" when the most influential cluster is removed. While a test based on the minimal t-statistic - obtained by sequentially dropping one cluster - does not have a $\mathcal{N}(0,1)$ distribution, and so one cannot easily determine if a minimal t-statistic is significantly different from zero, the finding does draw attention to the importance of heterogeneity across clusters.

The effective number of clusters is closely tied to the concept of influence through the measure of leverage. To formally relate leverage to the effective number of clusters, consider a univariate model in deviations-from-means form, $y_{ig} = \beta x_g + u_{ig}$. The leverage of an observation is determined by the relative magnitude of the explanatory variable for that observation and in the two-stage least squares context, is given by

$$\frac{\widehat{x}_{ig}^2}{\sum_{g=1}^{G}\sum_{i=1}^{n_g}\widehat{x}_{ig}^2}.$$

From (3.3), the contribution of a single observation to the effective number of clusters is

$$\frac{\widehat{x}_{ig}^2\omega_{ig}^2}{\sum_{g=1}^{G}\sum_{i=1}^{n_g}\widehat{x}_{ig}^2},$$

which is the leverage weighted by the variance of the error. An influential observation is one for which the leverage and the magnitude of the error are high. Because a larger error variance generally leads to larger magnitude errors, the effective number of clusters is a natural measure of influence.

Unsurprisingly, there is a close link between the size of the largest cluster and the feasible effective number of clusters. Figure 3.8 presents the relation for the simulation model of Section 3.3.

To demonstrate the effectiveness of using the feasible effective number of cluster as an indicator for non-normal distributions of the test statistic we use data from

Figure 3.8: Relationship between size of largest cluster and total leverage of the largest cluster with OLS.

papers published in *AEA* journals that use instrumental variables. We first evaluate the sensitivity of the estimated coefficients and standard errors to influential clusters. We can further investigate the sensitivity of the inference by comparing the critical values generated from using the wild cluster bootstrap to the standard critical values and to the estimated test statistic.

3.5.1 Empirical Paper #1

The first empirical paper we look as is *The Long-Run Effect of Mexican Immigration on Crime in US Cities: Evidence from Variation in Mexican Fertility Rates* published in the *American Economic Review* (Chalfin (2015)). This paper uses instrumental variables to measure the effects of Mexican immigration on crime rates reported to police. We see in Table 3.3 that the total number of groups in this paper is 92. This value would appear to be large enough to use critical values from the Normal distribution for hypothesis testing. However, for the 2SLS regression specification in the paper the feasible effective number of clusters is equal to 11. This means the number we should use to determine if we should use a Normal distribution for the critical values is 11 not 92. From our simulations, we saw that a feasible effective number of clusters of 11 would be an indicator that using the bootstrap to determine the critical values would be appropriate. When leaving out just one cluster the estimated t-statistic changes from 2.05 to 1.29 demonstrating that the estimated coefficient is very sensitive to one cluster. We can also see in Table 3.3 that the largest change in test statistic corresponds to a 21% change in the estimated β and a 26% change in the estimated standard error.

Table 3.3: Leave-One-Out Estimation

	Chalfin
G	92
$\lvert\hat{t}\rvert$	2.05
Max Leverage	22.3 %
FENC	11
Leave-One-Out(Cross Validation)	
min. $\lvert\hat{t}\rvert$	1.29
% Change $\hat{\beta}$	21%
% Change $\hat{se}$	26%

Figure 3.9 show the distribution in the percentage change for the removal of each cluster from the regression estimation. We can see that the change in the estimated β ranges from -20% to 20% while the change in the estimated standard error ranges from about -18% to 26%. These large percentage changes from the removal of a single cluster shows that both the estimated β and standard error are sensitive to the removal of one cluster.

Figure 3.9: Distributions of estimated β and se with the removal of one cluster.

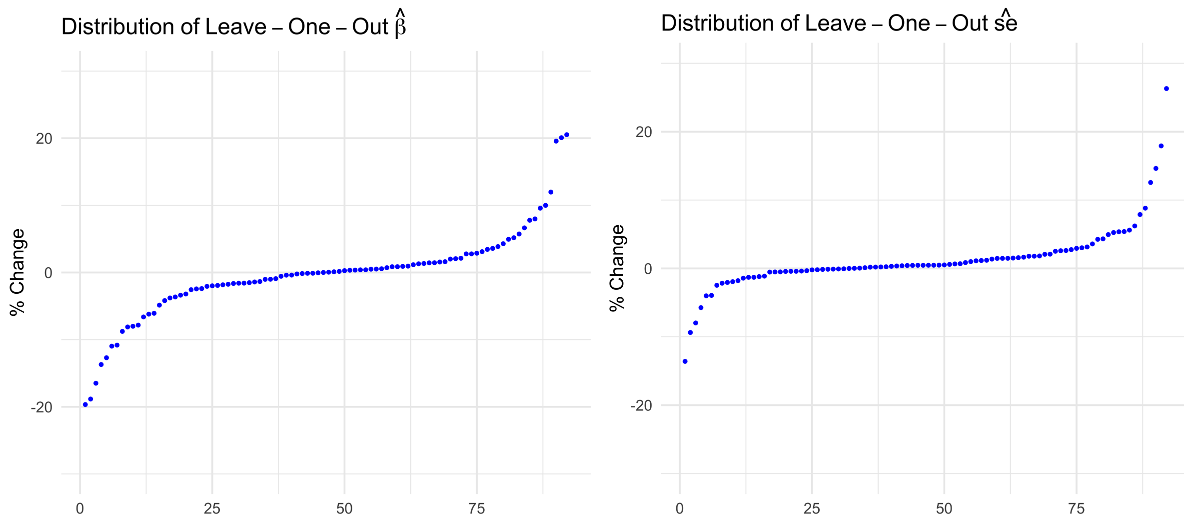

3.5.2 Empirical Paper #2

The second empirical paper we look at is *Income and Democracy* published in the *American Economic Review* (Acemoglu, Johnson, Robinson, and Yared (2008)). This paper uses instrumental variables to measure the impact of economic growth on measures of democracy. The data consists of 747 observations clustered at the country level with 114 countries. Comparing the evaluation of Acemoglu et al. and Chalfin, we can see that while the two data sets have a similar number of groups, 114 and 92, the feasible effective number of clusters for Acemoglu et al. is equal to 37 while for Chalfin it was only equal to 11. Referring back to Figure 3.5, the empirical test size seems to begin increasing when the effective number of clusters is below 10 to 15. We can see that the feasible effective number of clusters of the Acemoglu et al. data is above this threshold and thus would not require the use of a bootstrap procedure. The results of the leave-one-out estimation do not cause the results to become insignificant. As demonstrated in both of these papers, a *low* feasible effective number of clusters occurs when the minimum leave-one-out test statistic is *below* the Normal critical value while a *high* feasible effective number of clusters occurs when the minimum leave-one-out test statistic is *above* the Normal critical value

Table 3.4: Leave-One-Out Estimation

	Acemoglu et al.	*Chalfin*
G	114	92
$\lvert\hat{t}\rvert$	2.53	2.05
Max Leverage	14.9%	22.3 %
FENC	37	11
Leave-One-Out(Cross Validation)		
min. $\lvert\hat{t}\rvert$	2.39	1.29
% Change $\hat{\beta}$	-2%	21%
% Change $\hat{se}$	8%	26%

In Figure 3.10 we can see that the range of changes produced by the leave-one-out estimation is much smaller compared to Empirical Paper #1. This again shows the relationship between the sensitivity to leave-one-out estimation and the feasible effective number of clusters. Looking at the values in Figure 3.10 compared to the values in Table 3.4 we can see that the minimum test statistic does not correspond to the largest percent change in $\hat{\beta}$ or $\hat{se}$.

Figure 3.10: Distributions of estimated β and se with the removal of one cluster.

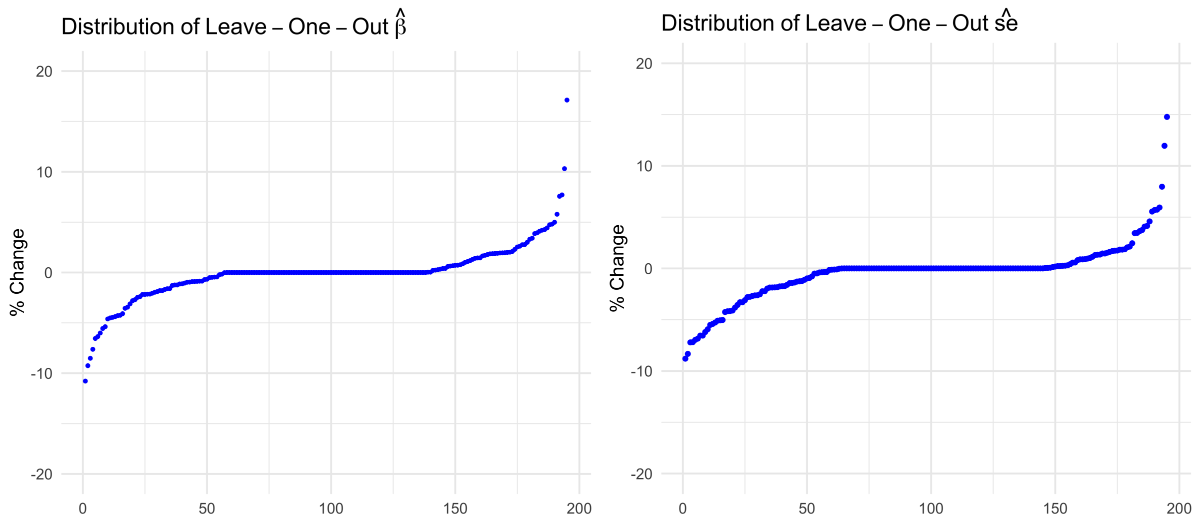

3.5.3 Empirical Paper #3

The third empirical paper we look at is *The China Syndrome: Local Labor Market Effects of Import Competition in the United States* published in the *American Economic Review* (Autor et al. (2013)). This paper analyzes the effect of rising Chinese import competition on US labor markets. They instrument for growth in Chinese imports to the United States using the contemporaneous composition and growth of Chinese imports in eight other developed countries.

Table 3.5: Leave-One-Out Estimation

	Acemoglu et al.	*Chalfin*	*Autor et al.*
G	114	92	48
$\lvert\hat{t}\rvert$	2.53	2.05	5.94
Max Leverage	14.9%	22.3 %	9.7 %
FENC	37	11	14.5
Leave-One-Out(Cross Validation)			
min. $\lvert\hat{t}\rvert$	2.39	1.29	5.39
% Change $\hat{\beta}$	-2%	21%	11.5%
% Change $\hat{se}$	8%	26%	-2.4%

While this paper has only 48 clusters in the data, about half the number of clusters as Empirical Paper #1, the feasible effective number of clusters is larger than that paper with FENC equal to 14.5 compared to 11. We also see in Figure 3.11 that these estimates are much less sensitive to both Empirical Paper #1 and Empirical Paper #2.

Figure 3.11: Distributions of estimated β and se with the removal of one cluster.

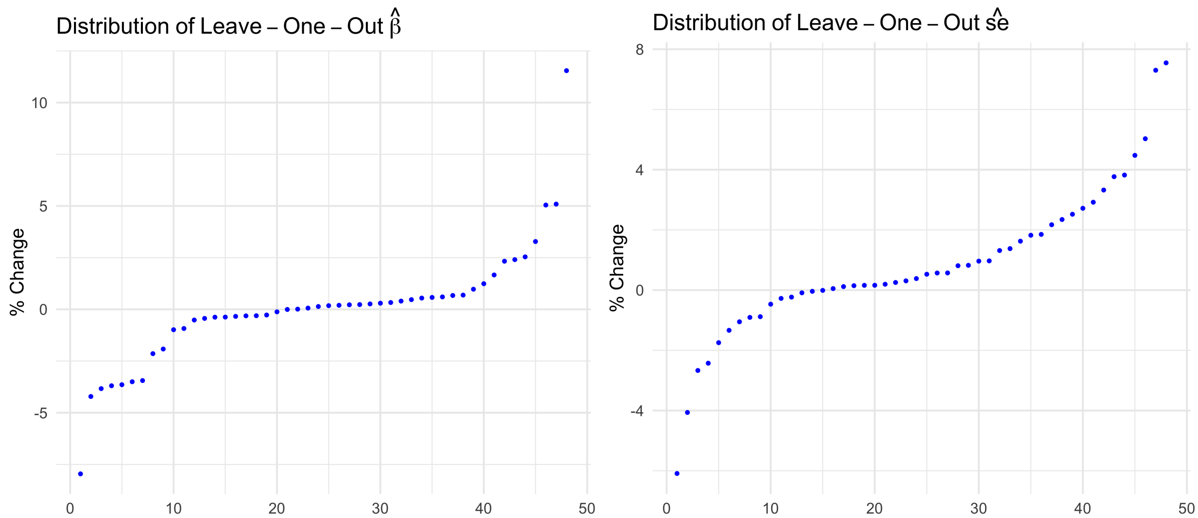

3.5.4 Empirical Paper #4

The fourth empirical paper we look at is *The Effects of Rural Electrification on Employment: New Evidence from South Africa* published in the *American Economic Review* (Dinkelman (2011)). This paper analyzes the impact of electrification on employment growth in South Africa instrumenting for program placement using the average community land gradient.

In Table 3.6 we see that this paper has 293 clusters the largest of any paper we study. But the feasible effective number of clusters is only 18.7 even with close to three hundred clusters. As expected with son many clusters, the maximum leverage of any one cluster is much smaller than the maximum leverage in any of the other

Table 3.6: Leave-One-Out Estimation

	Dinkelman	*Acemoglu et al.*	*Chalfin*	*Autor et al.*
G	293	114	92	48
$\lvert\hat{t}\rvert$	1.64	2.53	2.05	5.94
Max Leverage	3.75%	14.9%	22.3 %	9.7 %
FENC	18.7	37	11	14.5
Leave-One-Out(Cross Validation)				
min. $\lvert\hat{t}\rvert$	1.23	2.39	1.29	5.39
% Change $\hat{\beta}$	-30%	-2%	21%	11.5%
% Change $\hat{se}$	-6.8%	8%	26%	-2.4%

empirical papers looked at so far.

Figure 3.12: Distributions of estimated β and se with the removal of one cluster.

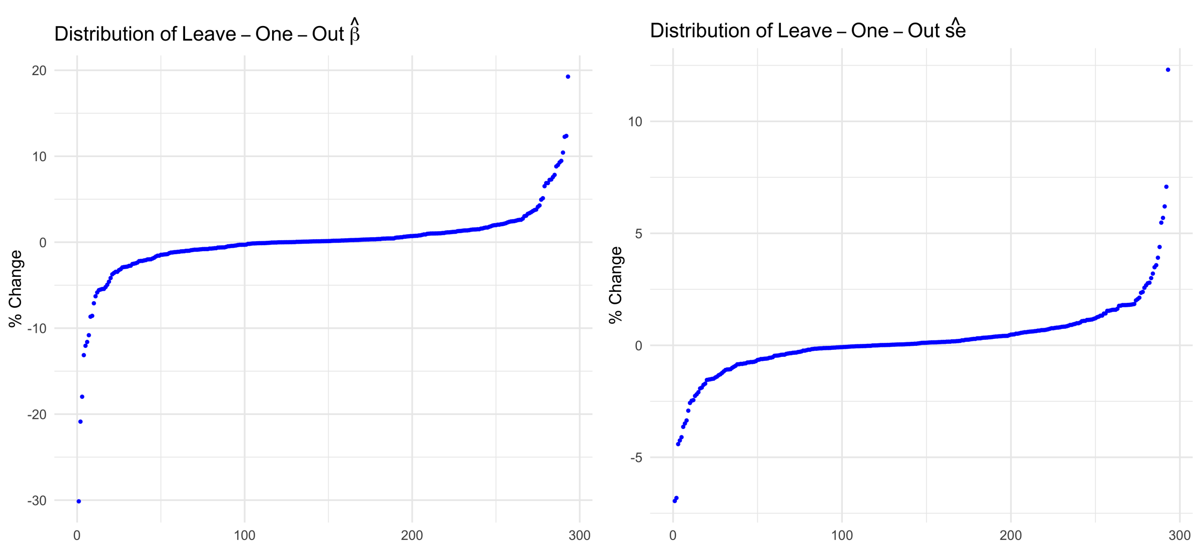

3.5.5 Empirical Paper #5

The final empirical paper we look at is *German Jewish Émigrés and US Invention* published in the *American Economic Review* (Moser et al. (2014)). This paper analyzes the impact of Jewish émigrés' from Nazi German on chemical innovation in the United States.

Table 3.7: Leave-One-Out Estimation

	Dinkelman	*Moser et al.*	*Acemoglu et al.*	*Chalfin*	*Autor et al.*
G	293	166	114	92	48
$\lvert\hat{t}\rvert$	1.64	2.48	2.53	2.05	5.94
Max Leverage	3.75%	1.47%	14.9%	22.3 %	9.7 %
FENC	18.7	1	37	11	14.5
Leave-One-Out(Cross Validation)					
min. $\lvert\hat{t}\rvert$	1.23	2.16	2.39	1.29	5.39
% Change $\hat{\beta}$	-30%	-14.1%	-2%	21%	11.5%
% Change $\hat{se}$	-6.8%	-1.38%	8%	26%	-2.4%

We can see in Table 3.7 that this paper has the lowest maximum leverage of any paper we study and the leave-one-out estimation does not produce a minimum test statistic below the Normal critical value of 1.96. However, the extremely low feasible effective number of clusters is would indicate to us the Normal critical value should not be used for inference with this data.

Figure 3.13: Distributions of estimated β and se with the removal of one cluster.

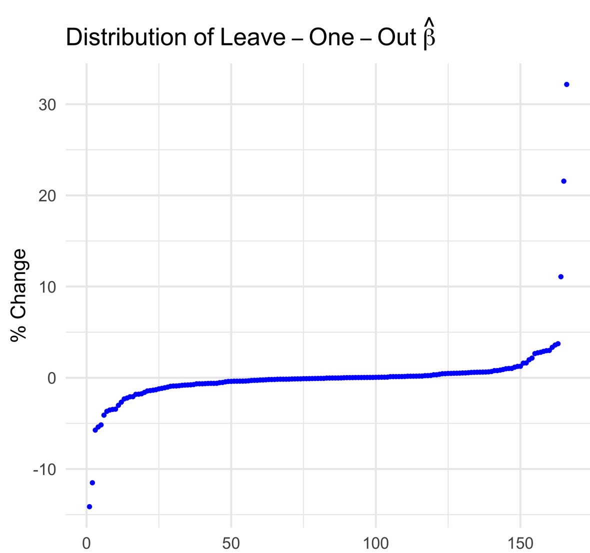

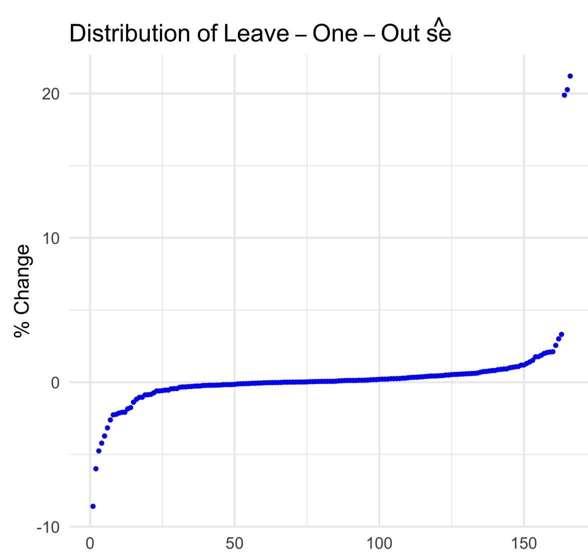

3.5.6 Analysis of Empirical Studies

Now that we have summarized the characteristics of the data and estimation in all five Empirical Papers we look at what these results mean for inference. Table 3.8 summarizes the use of Normal critical values and the use of bootstrapped critical values to conduct hypothesis testing. In the top panel we see G, the number of clusters, which traditionally has been used to guide inference would suggest using the Normal distribution to conduct inference for all five regressions. These results suggest failing to reject significance in only the regression in column 1. The leave-one-out sensitivity test in the middle panel finds that one regression (regression 4) is sensitive to an influential cluster, and would suggest failing to reject significance in both regression 1 and regression 4. In the bottom panel, the feasible effective number of clusters measure also identifies regression 4 as potentially having sensitive inference as well as identifying two other regressions (regressions 2 and 5). We see that with bootstrap critical values, the statistically significant results in regressions 2 and 4 disappear once the bootstrap is used. Using the feasible effective number of clusters to guide inference captures the sensitivity of the regressions captured by the leave-one-out estimation, as well as, identifying an additional regression that was not captured by the leave-one-out procedure.

Table 3.8: Evaluation of Inference in Five Empirical Papers

	Dinkelman (1)	*Moser et al.* (2)	*Acemoglu et al.* (3)	*Chalfin* (4)	*Autor et al.* (5)
G	293	166	114	92	48
Distribution	$\mathcal{N}$	$\mathcal{N}$	$\mathcal{N}$	$\mathcal{N}$	$\mathcal{N}$
$\lvert\hat{t}\rvert$	1.64	2.48	2.53	2.05	5.94
Conclude	Fail to Reject	Reject	Reject	Reject	Reject
Leave-One-Out(Cross Validation)					
min. $\lvert\hat{t}\rvert$	1.23	2.16	2.39	1.29	5.39
Conclude	Fail to Reject	Reject	Reject	Fail to Reject	Reject
Feasible Effective Number of Clusters					
FENC	**18.7**	**1.0**	**35**	**11**	**14.5**
Distribution	$\mathcal{N}$	Bootstrap	$\mathcal{N}$	Bootstrap	Bootstrap
Critical Value $\mathcal{N}$	1.96	1.96	1.96	1.96	1.96
Critical Value Bootstrap	-	2.84	-	2.20	2.52
Conclude	Fail to Reject	Fail to Reject	Reject	Fail to Reject	Reject

3.6 Conclusion

The use of G, the number of clusters in the data, to guide inference when using data that can be divided into clusters has become common among applied econometricians. We show that when cluster heterogeneity is present the effective number of clusters, not G, is the appropriate measure to use to determine critical values for hypothesis testing. When the feasible effective number of clusters is low, simulations show that the distribution of the test statistic is non-normal even when the number of clusters is large. This non-normality will lead to incorrect inference if the critical values are determined using the normal distribution. If the feasible effective number of clusters is low then critical values should be obtained from the wild cluster bootstrap.

Appendix A

The Health Cost of Wildfire Smoke

A.1 Health Data

The hospital admissions data is grouped by International Classification of Diseases Clinical Modification 9^{th} Revision (ICD-9-CM) and 10^{th} Revision (ICD-10-CM) produced by the Centers for Medicare and Medicaid Services (CMS) and the National Center for Health Statistics (NCHS). These codes are a morbidity classification for classifying diagnoses and reason for visits in American health care settings. These codings are based on the statistical classification of disease published by the World Health Organization (WHO).

Table A.1: Groupings based on principal diagnosis codes

	ICD-9 2012-2015Q3	ICD-10 2015Q4-2018
Circulatory System	390 to 460	I00 to J00
Infections	001 to 140	A00 to C00
Injuries Poisonings	800 to 999	S00 to T89
Respiratory System	460 to 520	J00 to K00

A.2 Smoke Data

Figure A.1 displays the distribution of monthly Smoke Days by County grouping. We can see the the Colusa and Glenn Counties unit had the highest median number of monthly Smoke Days with 7 days, while Marin, San Francisco, San Mateo, and Orange all have the lowest median number of 0 monthly Smoke Days.

Figure A.1: Distribution of Smoke Days 2012 – 2018

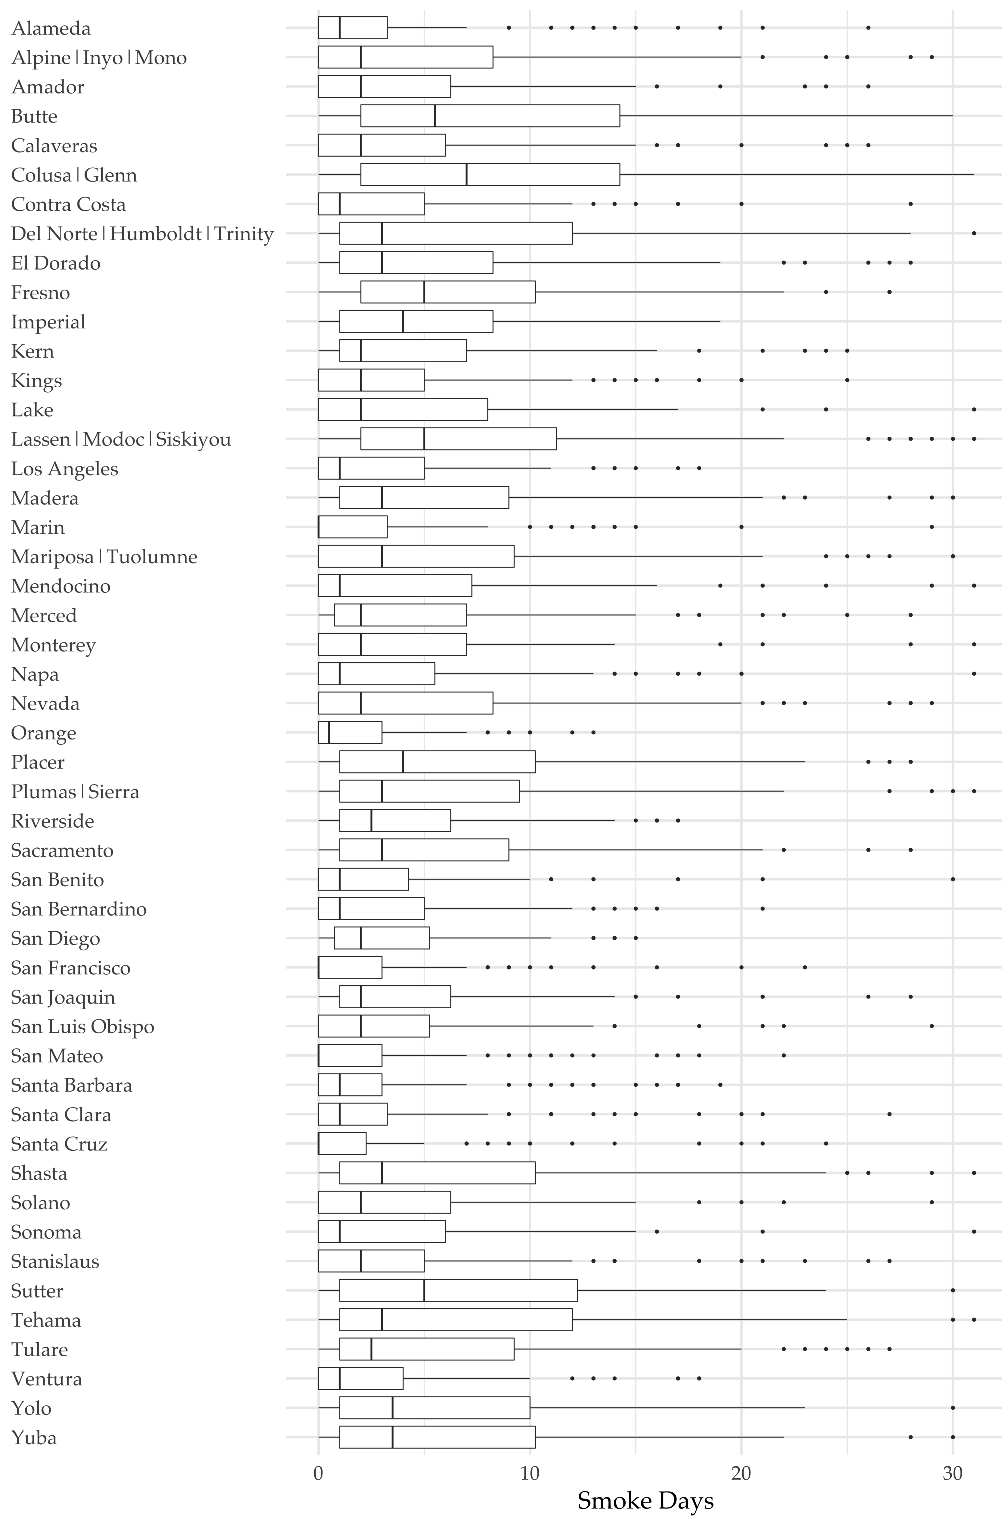

A.3 Dynamic Effects of Wildfire Smoke on Respiratory Health

Table A.2: Impact of number of smoke days on hospital admissions in California.

	Respiratory Admissions									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Smoke Days	6.57**	5.32**	5.58**	5.56**	5.05**	3.47*	3.03*	3.18*	3.48*	3.92*
	(3.11)	(2.54)	(2.64)	(2.66)	(2.40)	(1.77)	(1.62)	(1.65)	(1.81)	(1.96)
Smoke Days 1 Month Previous		4.52**	3.79**	3.78**	3.40*	3.00*	2.25*	2.31*	2.63*	2.82*
		(2.22)	(1.83)	(1.88)	(1.75)	(1.54)	(1.23)	(1.22)	(1.37)	(1.43)
Smoke Days 2 Months Previous			2.19*	2.25**	1.61*	0.60	0.33	0.53	0.71	1.06**
			(1.29)	(1.01)	(0.81)	(0.65)	(0.59)	(0.50)	(0.47)	(0.48)
Smoke Days 3 Months Previous				−0.21	2.01	0.98	0.41	0.50	0.91	1.07
				(1.42)	(1.95)	(1.58)	(1.43)	(1.42)	(1.53)	(1.53)
Smoke Days 4 Months Previous					−7.43**	−4.04*	−4.75*	−4.59*	−4.40*	−4.00*
					(3.63)	(2.18)	(2.48)	(2.43)	(2.30)	(2.16)
Smoke Days 5 Months Previous						−12.19**	−10.40**	−10.19**	−9.81**	−9.53**
						(5.46)	(4.64)	(4.59)	(4.37)	(4.27)
Smoke Days 6 Months Previous							−6.61**	−6.96**	−6.50**	−6.12**
							(3.11)	(3.17)	(2.92)	(2.80)
Smoke Days 7 Months Previous								1.39	0.62	1.02
								(0.99)	(0.84)	(0.90)
Smoke Days 8 Months Previous									3.03	2.36
									(1.86)	(1.63)
Smoke Days 9 Months Previous										2.77**
										(1.26)
Num. obs.	4,059	4,059	4,059	4,059	4,059	4,059	4,059	4,059	4,059	4,059
Num. clusters	49	49	49	49	49	49	49	49	49	49
FENC	32.3	35.69	36.9	36.44	35.29	35.71	35.25	38.25	35.94	34.56

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
All regressions include weather controls and month, county, and year fixed effects.

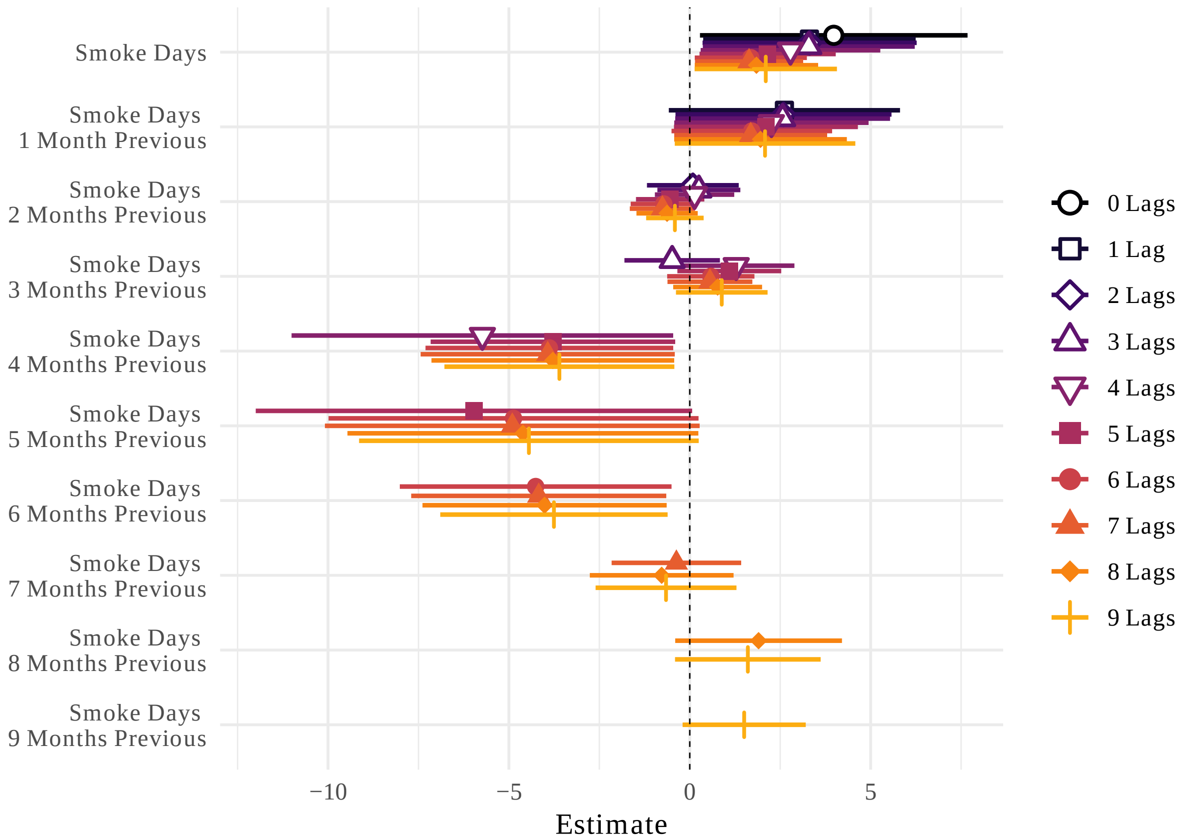

Figure A.2: Coefficient Plot for the impact of number of smoke days on respiratory hospital admissions.

A.4 Dynamic Effects of Wildfire Smoke on Circulatory Health

Table A.3: Impact of number of smoke days on hospital admissions in California.

	Circulatory Admissions									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Smoke Days	2.23^{**}	1.84^{**}	1.90^{**}	1.84^{**}	1.70^{**}	1.43^{**}	1.40^{**}	1.27^{**}	1.13^{**}	1.24^{**}
	(0.94)	(0.80)	(0.82)	(0.81)	(0.74)	(0.64)	(0.62)	(0.57)	(0.52)	(0.57)
Smoke Days 1 Month Previous		1.40^{**}	1.25^{**}	1.19^{**}	1.08^{**}	1.02^{**}	0.97^{**}	0.91^{**}	0.78^{**}	0.82^{**}
		(0.55)	(0.47)	(0.46)	(0.42)	(0.39)	(0.37)	(0.35)	(0.32)	(0.33)
Smoke Days 2 Months Previous			0.46	0.69^{*}	0.51	0.34	0.33	0.16	0.08	0.16
			(0.38)	(0.37)	(0.31)	(0.28)	(0.27)	(0.24)	(0.24)	(0.24)
Smoke Days 3 Months Previous				−0.72	−0.11	−0.29	−0.33	−0.40	−0.58	−0.54
				(0.48)	(0.45)	(0.43)	(0.44)	(0.44)	(0.45)	(0.44)
Smoke Days 4 Months Previous					-2.08^{*}	-1.50^{*}	-1.55^{*}	-1.67^{*}	-1.76^{*}	-1.66^{*}
					(1.07)	(0.86)	(0.88)	(0.92)	(0.94)	(0.91)
Smoke Days 5 Months Previous						-2.08^{**}	-1.96^{**}	-2.13^{**}	-2.30^{**}	-2.23^{**}
						(0.92)	(0.87)	(0.94)	(0.99)	(0.96)
Smoke Days 6 Months Previous							−0.47	−0.16	−0.36	−0.27
							(0.29)	(0.23)	(0.25)	(0.24)
Smoke Days 7 Months Previous								-1.16^{**}	−0.82	−0.72
								(0.56)	(0.49)	(0.49)
Smoke Days 8 Months Previous									-1.34^{**}	-1.50^{**}
									(0.52)	(0.58)
Smoke Days 9 Months Previous										0.68
										(0.44)
Num. obs.	4,054	4,054	4,054	4,054	4,054	4,054	4,054	4,054	4,054	4,054
Num. clusters	49	49	49	49	49	49	49	49	49	49
FENC	28.96	34.38	35.48	35.6	35.46	36.17	36.52	36.03	34.52	36.49

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.
All regressions include weather controls and month, county, and year fixed effects.

A.5 Nonlinear Effects

An alternate approach to creating bins of 5 days is to split the data based on the quintile of the number of Smoke Days. The advantage of this method is that the bins contain roughly equal number of observations. However, because the effect of wildfire smoke is only seen for the bins with the most days of smoke expose and these are also the least frequently observed the highest bin contains the largest range of values. As expected with this approach, only the coefficient for the highest number of Smoke Days quintile has a significant positive results. Therefore, this approach is not able to capture the variation in the impact of Smoke Days that is only present in the months that have a high number of Smoke Days.

Figure A.3: Count of Observations in Each Smoke Days Bin

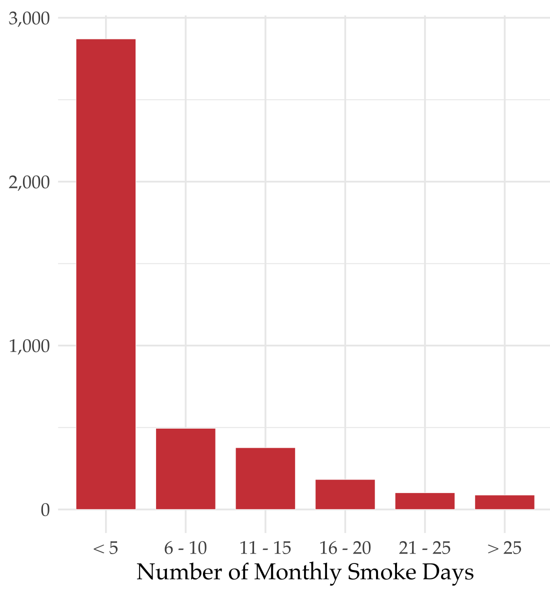

(a) 5 Day Bins

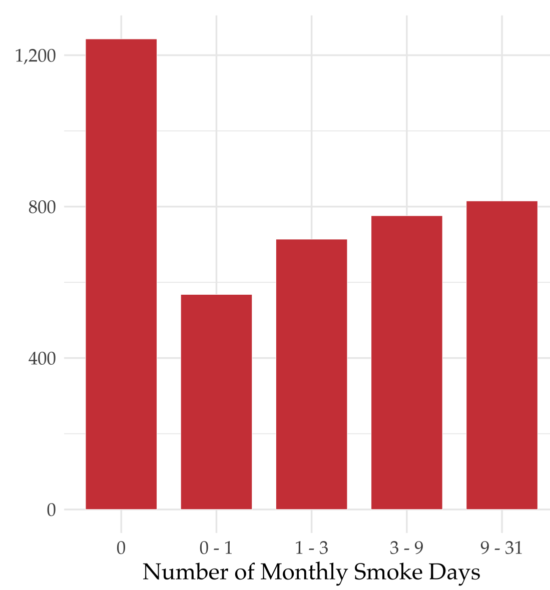

(b) Quintile Bins

Table A.4: Impact of number of smoke days on hospital admissions in California.

	Respiratory Admissions			
	(1)	(2)	(3)	(4)
0 to 1 Smoke Days	-45.80^{*}	-31.97	-32.43	-31.14
	(25.57)	(24.48)	(31.48)	(31.00)
1 to 3 Smoke Days	-60.19^{**}	-12.80	-4.58	-1.71
	(27.36)	(20.12)	(24.50)	(22.95)
3 to 9 Smoke Days	-95.26^{***}	-9.14	-8.98	-9.16
	(32.70)	(18.56)	(23.23)	(24.01)
9 to 31 Smoke Days	-89.86^{***}	45.57^{*}	45.67^{*}	50.13^{**}
	(24.57)	(23.88)	(23.08)	(23.88)
0 to 1 Smoke Days	-51.57^{*}	-37.55	-36.82	-43.19
1 Month Previous	(27.63)	(24.75)	(31.66)	(35.99)
1 to 3 Smoke Days	-60.89^{**}	-27.29	-22.91	-26.09
1 Month Previous	(28.87)	(21.46)	(29.93)	(30.66)
3 ot 9 Smoke Days	-77.14^{**}	-25.88	-30.50	-32.20
1 Month Previous	(34.70)	(22.35)	(32.82)	(32.72)
9 to 31 Smoke Days	-55.73^{***}	30.58	35.58^{*}	30.64^{*}
1 Month Previous	(16.61)	(19.84)	(17.92)	(18.09)
0 to 1 Smoke Days	-14.97^{*}	-34.55^{**}	-40.46^{*}	-40.83^{*}
2 Months Previous	(7.93)	(13.88)	(23.22)	(23.65)
1 to 3 Smoke Days	-22.32^{*}	-39.74^{**}	-42.33	-43.16
2 Months Previous	(11.98)	(17.14)	(28.76)	(28.09)
3 to 9 Smoke Days	-50.58^{**}	-60.10^{**}	-66.04	-74.70
2 Months Previous	(24.02)	(27.49)	(39.47)	(45.22)
9 to 31 Smoke Days	-32.76^{**}	-28.90^{**}	-14.45	-22.12
2 Months Previous	(14.14)	(12.45)	(17.93)	(23.57)
Weather Controls		✓	✓	✓
Month FE		✓	✓	✓
County-by-Year FE			✓	✓
Year FE				✓
Num. obs.	4059	4059	4059	4059
Num. clusters	49	49	49	49

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.

Table A.5: Impact of number of smoke days on hospital admissions in California.

	Circulatory Admissions			
	(1)	(2)	(3)	(4)
0 to 1 Smoke Days	-4.51	-28.89^{***}	-7.42	-8.36
	(4.84)	(10.69)	(5.92)	(5.97)
1 to 3 Smoke Days	-12.97^{*}	-25.28^{*}	5.55	2.96
	(6.48)	(13.55)	(4.01)	(3.48)
3 to 9 Smoke Days	-26.11^{***}	-25.11	0.99	-0.66
	(9.63)	(20.45)	(5.04)	(5.98)
9 to 31 Smoke Days	-6.89	-26.23	29.20^{**}	27.13^{**}
	(4.16)	(28.86)	(12.82)	(11.81)
0 to 1 Smoke Days	-2.66	-31.84^{***}	-10.30	-11.05
1 Month Previous	(4.08)	(11.68)	(6.89)	(7.62)
1 to 3 Smoke Days	-11.67	-35.19^{**}	-10.56	-9.30
1 Month Previous	(7.25)	(17.02)	(14.04)	(12.31)
3 ot 9 Smoke Days	-15.17	-36.79	-15.15	-14.26
1 Month Previous	(9.68)	(22.57)	(15.16)	(15.25)
9 to 31 Smoke Days	-3.09	-47.72^{**}	-0.33	1.53
1 Month Previous	(4.28)	(22.06)	(8.15)	(7.91)
0 to 1 Smoke Days	-3.10	-36.12^{***}	-14.39^{**}	-15.10^{**}
2 Months Previous	(5.81)	(11.85)	(6.45)	(7.38)
1 to 3 Smoke Days	-7.73	-39.38^{***}	-10.59	-10.53
2 Months Previous	(5.73)	(14.05)	(7.22)	(7.04)
3 to 9 Smoke Days	-17.56^{**}	-49.69^{**}	-20.56	-20.47
2 Months Previous	(7.38)	(21.80)	(12.83)	(12.48)
9 to 31 Smoke Days	-6.58	-65.75^{**}	-6.76	-5.86
2 Months Previous	(4.23)	(25.87)	(8.63)	(6.45)
Weather Controls		✓	✓	✓
Month FE		✓	✓	✓
County-by-Year FE			✓	✓
Year FE				✓
Num. obs.	4054	4054	4054	4054
Num. clusters	49	49	49	49

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$
Standard errors are clustered at the county level.

A.6 Feasible Effective Number of Clusters

Table A.6 presents the feasible effective number of clusters (fenc) calculated following Carter, Schnepel, and Steigerwald (2017). Each column corresponds to the columns in the other regression tables with columns 1 and 5 using no controls or fixed effects and columns 4 and 8 using month, year, and County fixed effects and weather controls. A point to note is that we would expect the fenc to be the same for a given covariate and and specification regardless of outcome variable (for example for Smoke Days column 1 is 43.67 and and 5 is 43.62), however they are slightly different because a few observations are dropped from the regression due to missing data differences between the data for respiratory and circulatory admissions.

Table A.6: The Feasible Effective Number of Clusters for Regression Coefficients

	Respiratory Admissions				Circulatory Admissions			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Smoke Days	43.67	36.08	35.47	36.44	43.62	36.15	35.85	35.6
Smoke Days 1 Month Previous	44.64	37.37	20.94	33.67	44.56	37.31	36.64	38.42
Smoke Days 2 Months Previous	44.26	37.54	34.88	35.48	44.07	37.29	37.34	21.66
6 to 10 Smoke Days	44.57	41.93	42.64	42.02	44.97	42.28	41.06	9.27
11 to 15 Smoke Days	44.56	40.47	0.41	39.56	44.82	40.76	0.28	0
16 to 20 Smoke Days	36.31	37.05	5.02	35.73	35.96	36.78	0.16	0
21 to 25 Smoke Days	33.26	37.38	21.37	24.11	32.55	36.89	2.13	0
Over 25 Smoke Days	25.99	27.36	0.21	17.29	27.03	27.86	4.94	0
Num. clusters	49	49	49	49	49	49	49	49

The Feasible Effective Number of Clusters calculated by using the formula derived in Carter, Schnepel, & Steigerwald, 2017

Appendix B

The Sharing Economy and Rental Markets

B.1 Interpolation of American Community Survey Data

In our empirical analysis we wish to control for trends in housing and demographic using data taken from the American Community Survey (ACS). The ACS collects demographic and housing data on a continuous basis from a national sample. Due to the nature of the collection of the data, the ACS estimates describe conditions over the time period during which the data was collected. Using the 5-year estimates means that about four-fifths of the data for one year overlaps with the data of the following year. That means comparing estimates from one year 5-year dataset to the next will not allow you to isolate the differences in the two estimates. The 5-year estimates however are useful for representing long run trends in the data.

The Airbnb listings data and the Zillow housing data are both available at the monthly ZIP code level, however the ACS data is only at the yearly ZIP code level.

Therefore, we wish to interpolate the ACS data to the monthly level. To do this we assign the month of December[1] to each reported value from the ACS then use a cubic spline to fill in the data for the remaining month between the observed years. The interpolated data for select ZIP codes can be seen in Figure B.1.

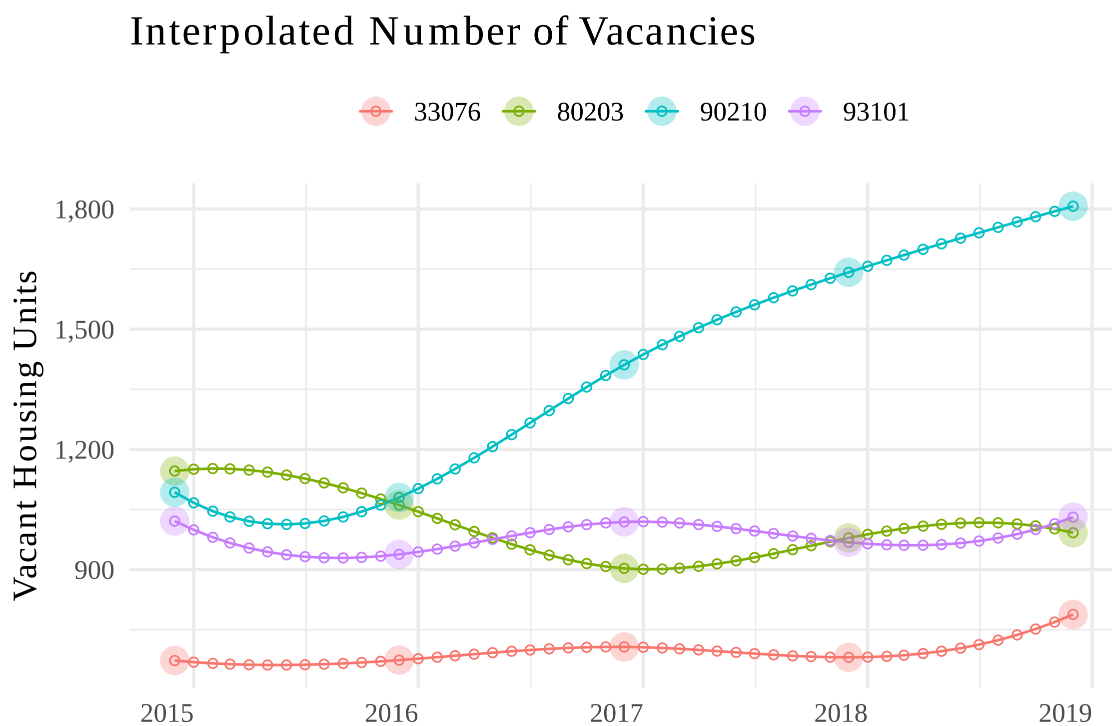

Figure B.1: Interpolated number of vacancies for select ZIP codes. Large circles represent ACS data points and small circles represent interpolated points.

B.2 Proof of Lemma 1

Proof:

Consider the problem faced by innkeepers and define $\widehat{\mathcal{I}}_0(p^H, \mathcal{V}_0) \equiv \mathcal{I}_0\big(p^H, \theta^H(p^H; \mathcal{V}_0)\big)$ for $0 \leq p^H < w^{\mathcal{V}} - r\mathcal{V}_0$ and $\widehat{\mathcal{I}}_0(w^{\mathcal{V}} - r\mathcal{V}_0, \mathcal{V}_0) = -\kappa^H$. It is easy to see that Equation 2.10

[1] Varying the month of the year the ACS estimate is assigned has little impact of the results of the analysis.

is continuous in p^H at $w^{\mathcal{V}} - r\mathcal{V}_0$ and so $\widehat{\mathcal{I}}_0$ is continuous across its domain. Noting also that $\widehat{\mathcal{I}}_0(0, \mathcal{V}_0) = -\kappa^H$, the innkeeper's problem is well defined and must achieve a maximum on the interval $[0, w^{\mathcal{V}} - r\mathcal{V}_0)$ since we have assumed $\widetilde{\mathcal{I}}_0 \geq 0$.

The problem faced by landlords in the short and long-term markets is structurally similar and, for brevity, not included. Since we have assumed that participation is weakly profitable, the short-term market's argmax, like that of the hotel market, must be in the interval $[0, w^{\mathcal{V}} - r\mathcal{V}_0)$, and the long-term market's in the interval $[0, w^{\mathcal{R}} - r\mathcal{R}_0)$.

■

(Back to Model)

B.3 Proof of Lemma 2

Proof:

It is sufficient to show that the first order conditions of property managers have unique solutions. First, consider an innkeeper who has entered the hotel market and is choosing which price to post. She maximizes Equation 2.10 subject to Equation 2.15. Rearranging the constraint for p^H, we have

$$p^H = w^{\mathcal{V}} - r\mathcal{V}_0 - \frac{(r + \delta^{\mathcal{V}})(r\mathcal{V}_0 - b^{\mathcal{V}})}{\theta^H \lambda^H}. \tag{B.1}$$

Substituting the above into Equation 2.16, we can write the problem of innkeepers as a choice of θ^H.

$$\max_{\theta^H} \left[-\kappa^H + \lambda^H \left(\frac{w^{\mathcal{V}} - r\mathcal{V}_0 - r\mathcal{I}_0}{r + \delta^{\mathcal{V}}} \right) - \frac{1}{\theta^H} \left(r\mathcal{V}_0 - b^{\mathcal{V}} \right) \right] \tag{B.2}$$

The first order condition is given by

$$-(\theta^H)^2 \frac{d\lambda^H}{d\theta^H}\big(w^{\mathcal{V}} - r\mathcal{V}_0 - r\mathcal{I}_0\big) = (r+\delta^{\mathcal{V}})(r\mathcal{V}_0 - b^{\mathcal{V}}) \tag{B.3}$$

which, given our assumptions on the matching function, has a unique solution. Thus, all innkeepers choose to search in a sub-market with the same market tightness, and because of the one-to-one relationship, the same price. By plugging Equation B.3 into Equation B.1 and simplifying, we uncover a classic competitive search result that the total surplus is split according to the elasticity of matching with respect to their participation. The problem in the short-term and long-term markets is structurally similar. The first order conditions are

$$-(\theta^S)^2 \frac{d\lambda^S}{d\theta^S}\big(w^{\mathcal{V}} - r\mathcal{V}_0 - r\mathcal{L}_0\big) = (r+\delta^{\mathcal{V}})(r\mathcal{V}_0 - b^{\mathcal{V}}) \tag{B.4}$$

$$-(\theta^L)^2 \frac{d\lambda^L}{d\theta^L}\big(w^{\mathcal{R}} - r\mathcal{R}_0 - r\mathcal{L}_0\big) = (r+\delta^{\mathcal{R}})(r\mathcal{R}_0 - b^{\mathcal{R}}). \tag{B.5}$$

■

(Back to Model)

B.4 Proof of Lemma 3

Proof:

To begin we derive the (implicit) demand functions, starting with the hotel market. First rearrange Equation 2.10 for p^H.

$$p^H = r\mathcal{I}_0 + \frac{(r+\delta^{\mathcal{V}})(r\mathcal{I}_0 + \kappa^H)}{\lambda^H}. \tag{B.6}$$

Combining this with Equation 2.7, Equation B.3, and simplifying we have

$$
\begin{aligned}
& -(\theta^H)^2 \frac{d\lambda^H}{d\theta^H}\left[w^{\mathcal{V}} - r\mathcal{I}_0 - \frac{b^{\mathcal{V}}(r+\delta^{\mathcal{V}}) + \theta^H\lambda^H(w^{\mathcal{V}} - r\mathcal{I}_0) - \theta^H(r+\delta^{\mathcal{V}})(r\mathcal{I}_0+\kappa^H)}{r+\delta^{\mathcal{V}}+\theta^H\lambda^H}\right] \\
& \qquad = (r+\delta^{\mathcal{V}})\left[\frac{b^{\mathcal{V}}(r+\delta^{\mathcal{V}}) + \theta^H\lambda^H(w^{\mathcal{V}} - r\mathcal{I}_0) - \theta^H(r+\delta^{\mathcal{V}})(r\mathcal{I}_0+\kappa^H)}{r+\delta^{\mathcal{V}}+\theta^H\lambda^H} - b^{\mathcal{V}}\right] \\
\iff \quad & -(\theta^H)^2 \frac{d\lambda^H}{d\theta^H}\left[\frac{w^{\mathcal{V}} - r\mathcal{I}_0 - b^{\mathcal{V}} + \theta^H(r\mathcal{I}_0+\kappa^H)}{\lambda^H}\right] \\
& \qquad = \frac{\lambda^H(w^{\mathcal{V}} - b^{\mathcal{V}} - r\mathcal{I}_0) - (r+\delta^{\mathcal{V}})(r\mathcal{I}_0+\kappa^H)}{\lambda^H} \\
\iff \quad & \left[r+\delta^{\mathcal{V}} + \big(1-\eta_H(\theta^H)\big)\theta^H\right](r\mathcal{I}_0+\kappa^H) = \eta_H(\theta^H)\big(w^{\mathcal{V}} - b^{\mathcal{V}} - r\mathcal{I}_0\big). \qquad \text{(B.7)}
\end{aligned}
$$

Equation B.7 describes an implicit function for the equilibrium demand for vacancies, θ^H, in terms of their cost, $\mathcal{I}_0$, which we write as $\theta^H = \zeta_H(\mathcal{I}_0)$. Next, differentiate w.r.t. $\mathcal{I}_0$.

$$
\frac{d\zeta^H}{d\mathcal{I}_0} = \frac{-r\left[r+\delta^{\mathcal{V}} + \eta_H(\theta^H) + \big(1-\eta_H(\theta^H)\big)\theta^H\right]}{(r\mathcal{I}_0+\kappa^H)\left[1-\eta_H(\theta^H) - \theta^H\frac{d\eta_H(\theta^H)}{d\theta^H}\right] - (w^{\mathcal{V}} - b^{\mathcal{V}} - r\mathcal{I}_0)\frac{d\eta_H(\theta^H)}{d\theta^H}} \tag{B.8}
$$

The above is strictly negative iff

$$
\frac{d\eta_H(\theta^H)}{d\theta^H} < \frac{\big(1-\eta_H(\theta^H)\big)(r\mathcal{I}_0+\kappa^H)}{(w^{\mathcal{V}} - b^{\mathcal{V}} - r\mathcal{I}_0) + \theta^H(r\mathcal{I}_0+\kappa^H)}. \tag{B.9}
$$

That is, if the marginal effect of market tightness on the filling rate elasticity is not *too* high, the demand for hotel vacancies is declining in $\mathcal{I}_0$. Given our standard assumptions on the matching function, $\frac{d\eta(\theta)}{d\theta} \leq 0$ so this condition is necessarily met. Under an isoelastic function, i.e. a Cobb-Douglas matching function,

$$
\frac{d\zeta^H}{d\theta^H} = \frac{-r\left[r+\delta^{\mathcal{V}} + \eta_H + (1-\eta_H)\theta^H\right]}{(1-\eta_H)(r\mathcal{I}_0+\kappa^H)} < 0. \tag{B.10}
$$

A similar set of steps establishes this result for the short and long-term markets. The implicit demand curves $\theta^S = \zeta_S(\mathcal{L}_0)$ and $\theta^L = \zeta_L(\mathcal{L}_0)$ are reproduced below.

$$\left[r + \delta^{\mathcal{V}} + \left(1 - \eta_S(\theta^S)\right)\theta^S\right](r\mathcal{L}_0 + \kappa^S) = \eta_S(\theta^S)\left(w^{\mathcal{V}} - b^{\mathcal{V}} - r\mathcal{L}_0\right) \tag{B.11}$$

$$\left[r + \delta^{\mathcal{R}} + \left(1 - \eta_L(\theta^L)\right)\theta^L\right](r\mathcal{L}_0 + \kappa^L) = \eta_L(\theta^L)\left(w^{\mathcal{R}} - b^{\mathcal{R}} - r\mathcal{L}_0\right) \tag{B.12}$$

■

(Back to Model)

B.5 More Comparative Static Results

In Table B.1 we report more comparative static results for completeness. These parameters do not as easily map into policy choices, but also provide some interesting model insights. First consider the flow values of unaccommodation. As they are increased, residents and visitors are made directly better off when searching for accommodation. Because property managers must deliver higher market utilities, they are made worse off. Increases in $b^{\mathcal{V}}$, ceteris paribus, increase finding rates for residents, while similar increases in $b^{\mathcal{R}}$, increase finding rates for visitors as landlords adjust vacancy posting strategies.

Increases in the flow value of being accommodated has similar effects. By making accommodation more attractive, market utility increases and prices rise. Increases in $w^{\mathcal{V}}$ hurts residents in terms of value and finding rates as landlords increase posting in the short-term market. The opposite holds when $w^{\mathcal{R}}$ increases. Interestingly, increases in $b^{\mathcal{V}}$ do negatively affect $\mathcal{R}_0$ like increases in $w^{\mathcal{V}}$ do (and the mirrored scenario). The key distinction is that increases in the flow value of searching effectively amount to better outside options. This pushes some landlords to post in the long-term market in the case of $b^{\mathcal{V}}$ increasing (and the short-term market in the when $b^{\mathcal{R}}$ increases). In other words, increases in one type's b directly increases their utility, while it indirectly improves the other by incentivizing landlords to the other market.

When δ increases, more properties are vacant in the steady state, increasing market tightnesses across the board and lowering vacancy values. If $\delta^{\mathcal{R}}$ increases, all prices fall with the value of the vacancies. In contrast, if $\delta^{\mathcal{V}}$ increases landlords can mitigate lost values by posting more in the short-term market and raising prices (which feeds through to the hotel market). Last, the results for increasing the number of searchers and properties are reported in the bottom of the table. Briefly, more tenants benefits property managers, and more properties benefit tenants.

	$\mathcal{V}_0$	$\mathcal{R}_0$	$\mathcal{I}_0$	$\mathcal{L}_0$	p^H	p^S	p^L	θ^H	θ^S	θ^L
$\uparrow b^{\mathcal{V}}$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$-$	$-$	$+$
$\uparrow b^{\mathcal{R}}$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$+$	$+$	$-$
$\uparrow w^{\mathcal{V}}$	$+$	$-$	$+$	$+$	$+$	$+$	$+$	$+$	$+$	$-$
$\uparrow w^{\mathcal{R}}$	$-$	$-$	$+$	$+$	$+$	$+$	$+$	$-$	$-$	$+$
$\uparrow \delta^{\mathcal{V}}$	$-$	$+$	$-$	$-$	$+$	$+$	$-$	$+$	$+$	$+$
$\uparrow \delta^{\mathcal{R}}$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$+$	$+$	$+$
$\uparrow u^{\mathcal{V}}$	$-$	$-$	$+$	$+$	$+$	$+$	$+$	$-$	$-$	$-$
$\uparrow u^{\mathcal{R}}$	$-$	$-$	$+$	$+$	$+$	$+$	$+$	$-$	$-$	$-$
$\uparrow N^{\mathcal{I}}$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$+$	$+$	$+$
$\uparrow N^{\mathcal{L}}$	$+$	$+$	$-$	$-$	$-$	$-$	$-$	$+$	$+$	$+$

Table B.1: More Comparative Statics

Appendix C

Effective Number of Clusters and Inference with Instrumental Variables

Proof of Lemma 1: Let $\widehat{\mathbf{X}}_g^*$ be the $n \times k$ matrix of explanatory variables with all rows that do not correspond to cluster g set to zero.

Part a: The cluster-specific estimator $\widehat{\beta}_{TS\,g}$ is constructed with a generalized inverse to allow both for cluster-invariant explanatory variables and for clusters with $n_g < k$. Because $\widehat{\mathbf{X}}^{\mathrm{T}}\mathbf{y} = \sum_{g=1}^{G} \widehat{\mathbf{X}}_g^{*\mathrm{T}}\mathbf{y}$ and $\widehat{\mathbf{X}}_g^{*\mathrm{T}}\mathbf{y} = \widehat{\mathbf{X}}_g^{\mathrm{T}}\mathbf{y}_g$,

$$\begin{aligned}
\widehat{\beta}_{TS} &= (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}^{\mathrm{T}}\mathbf{y} \\
&= (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\sum_{g=1}^{G}\widehat{\mathbf{X}}_g^{*\mathrm{T}}\mathbf{y} \\
&= \sum_{g=1}^{G}(\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g(\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g)^{-}\widehat{\mathbf{X}}_g^{*\mathrm{T}}\mathbf{y} \\
&= \sum_{g=1}^{G}C_g(\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g)^{-}\widehat{\mathbf{X}}_g^{\mathrm{T}}\mathbf{y}_g \\
&= \sum_{g=1}^{G}C_g\widehat{\beta}_{TS\,g},
\end{aligned}$$

where $C_g = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g$ and $(\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g)^{-}$ is a generalized inverse. Thus,

$$\begin{aligned}
V &= Var\left[\widehat{\beta}_{TS}|Z\right] \\
&= Var\left[\sum_{g=1}^{G}C_g\widehat{\beta}_{TS\,g}|Z\right] \\
&= \sum_{g=1}^{G}C_g Var\left[\widehat{\beta}_{TS\,g}|Z\right]C_g^{\mathrm{T}}.
\end{aligned}$$

Part b: Using the property $\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g = \widehat{\mathbf{X}}_g^{*\mathrm{T}}\widehat{\mathbf{X}}_g^{*} = \widehat{\mathbf{X}}_g^{*\mathrm{T}}\widehat{\mathbf{X}}$

$$\begin{aligned}
\widehat{\mathbf{X}}_g^{*\mathrm{T}}(\mathbf{y}-\widehat{\mathbf{X}}\widehat{\beta}) &= [\widehat{\mathbf{X}}_g^{*\mathrm{T}} - \widehat{\mathbf{X}}_g^{*\mathrm{T}}\widehat{\mathbf{X}}(\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}^{\mathrm{T}}]\mathbf{y} \\
&= \widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g[(\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g)^{-}\widehat{\mathbf{X}}_g^{*\mathrm{T}} - (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}^{\mathrm{T}}]\mathbf{y} \\
&= \widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{\mathbf{X}}_g(\widehat{\beta}_{TS\,g} - \widehat{\beta}_{TS}).
\end{aligned}$$

Therefore

$$C_g(\widehat{\beta}_{TS\,g} - \widehat{\beta}_{TS}) = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}_g^{*\mathrm{T}}(\mathbf{y} - \widehat{\mathbf{X}}\widehat{\beta}_{TS}) = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{v}_g, \tag{C.1}$$

where the second equality follows because $(y - \widehat{\mathbf{X}}\widehat{\beta}_{TS}) = \widehat{v}$ and $\widehat{\mathbf{X}}_g^{*\mathrm{T}}\widehat{v} = \widehat{\mathbf{X}}_g^{\mathrm{T}}\widehat{v}_g$. The cluster representation for $\widehat{V}$ follows directly from (C.1). ■

Proof of Theorem 1: We want to show that under H_0 the test statistic converges in distribution to $\mathcal{N}(0,1)$.

The first step is to establish that

$$T = \frac{a^{\mathrm{T}}(\widehat{\beta}_{TS} - \beta_0)}{\sqrt{V_a}} \rightsquigarrow \mathcal{N}(0,1),$$

where $V_a = a^{\mathrm{T}}Va$. Observe that

$$T = \sum_{g=1}^{G} D_g,$$

where the cluster-level components are $D_g := a^{\mathrm{T}}C_g(\widehat{\beta}_{TS\,g} - \beta_0)/\sqrt{V_a}$. Conditional on Z the elements of $\{D_g\}$ are asymptotically independent, $\mathbb{E}(D_g \mid Z)$ asymptotically approaches zero, and

$$\mathrm{Var}(D_g \mid Z) = a^{\mathrm{T}}C_g\,\mathrm{Var}(\widehat{\beta}_{TS\,g} \mid Z)C_g^{\mathrm{T}}a.$$

Let $s_G^2 := \sum_{g=1}^{G} Var(D_g \mid Z) = V_a$. Under Assumption 1(ii), $\mathbb{E}(D_g^4) < \infty$, so there exists a $\delta > 0$ for which

$$\lim_{G\to\infty} \frac{1}{s_G^{2+\delta}} \sum_{g=1}^{G} \mathbb{E}\left[|D_g \mid Z|^{2+\delta}\right] = 0.$$

Therefore, by the Lyapunov Central Limit Theorem the distribution of $D_g|Z$ converges to a standard normal distribution. The convergence is almost surely over Z, so

$$T \rightsquigarrow \mathcal{N}(0,1).$$

The test statistic

$$t = \frac{a^{\mathrm{T}}(\widehat{\beta}_{TS} - \beta_0)}{\sqrt{\widehat{\mathrm{Var}}(a^{\mathrm{T}}\widehat{\beta}_{TS})}},$$

which can be written as

$$t = T\left(\frac{V_a}{\widehat{V}_a}\right)^{\frac{1}{2}},$$

will converge in distribution to T if $\frac{\widehat{V}_a}{V_a} \xrightarrow{\mathbb{P}} 1$, by Slutsky's lemma.

Let

$$\widetilde{V} = (\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1}\sum_{g=1}^{G}\widehat{\mathbf{X}}_g^{\mathrm{T}}v_g v_g^{\mathrm{T}}\widehat{\mathbf{X}}(\widehat{\mathbf{X}}^{\mathrm{T}}\widehat{\mathbf{X}})^{-1},$$

be an unbiased estimator of V, with $\widetilde{V}_a = a^{\mathrm{T}}\widetilde{V}a$. Now

$$\frac{\widehat{V}_a}{V_a} - 1 = \frac{\widehat{V}_a - V_a}{V_a} = \frac{\widetilde{V}_a - V_a + \widehat{V}_a - \widetilde{V}_a}{V_a},$$

so $\frac{\widehat{V}_a}{V_a} \xrightarrow{\mathbb{P}} 1$ if

$$\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| \quad \text{and} \quad \left|\frac{\widehat{V}_a - \widetilde{V}_a}{V_a}\right|$$

are each $o_{\mathbb{P}}(1)$. This will follow if

$$\begin{aligned} \mathbb{P}\left\{\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| > \varepsilon\right\} \to 0, \\ \mathbb{P}\left\{\left|\frac{\widehat{V}_a - \widetilde{V}_a}{V_a}\right| > \varepsilon\right\} \to 0. \end{aligned} \tag{C.2}$$

To do so we show that

$$\begin{aligned} \mathbb{P}\left\{\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| > \varepsilon \middle| Z\right\} \to 0, \\ \mathbb{P}\left\{\left|\frac{\widehat{V}_a - \widetilde{V}_a}{V_a}\right| > \varepsilon \middle| Z\right\} \to 0. \end{aligned} \tag{C.3}$$

Because the probabilities in (C.3) are bounded, they are uniformly integrable functions of Z. Under uniform integrability, convergence in probability implies convergence in expectation so, for example,

$$\mathbb{P}\left\{\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| > \varepsilon\right\} = \mathbb{E}\left(\mathbb{P}\left\{\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| > \varepsilon \middle| Z\right\}\right) \to 0.$$

Adapting Lemma 1 from Carter, Schnepel, and Steigerwald (2017), Chebyshev's inequality implies

$$\mathbb{P}\left\{\left|\frac{\widetilde{V}_a - V_a}{V_a}\right| > \varepsilon \middle| Z\right\} \le \frac{2}{\varepsilon^2} \frac{1 + \Gamma^*(\Omega, Z)}{G}. \tag{C.4}$$

Under Assumption 1(ii)-(iii) the expected value of this bound goes to zero.

Adapting Lemma 2 from Carter, Schnepel, and Steigerwald (2017), Markov's inequality implies

$$\mathbb{P}\left\{\left|\frac{\widehat{V}_a-\widetilde{V}_a}{V_a}\right|>\varepsilon\middle|Z\right\}\leq\frac{1}{\varepsilon}\left(\frac{1}{G}+\frac{1}{V_a}a^{\mathrm{T}}Ma+2(\frac{1}{V_a}a^{\mathrm{T}}Ma)^{1/2}\right), \tag{C.5}$$

where $M=\sum_{g=1}^{G}(A_g-\frac{1}{G}I)V(A_g-\frac{1}{G}I)^{\mathrm{T}}$. Under Assumption 1(ii) and (iv) this bound goes to zero in probability. Hence

$$lim_{n\to\infty}\mathbb{P}\left\{\left|\frac{\widehat{V}_a}{V_a}-1\right|>\varepsilon\right\}=0,$$

which implies convergence in distribution. ■

Bibliography

Daron Acemoglu, Simon Johnson, James A Robinson, and Pierre Yared. Income and democracy. *American Economic Review*, 98(3):808–42, 2008.

James Albrecht, Pieter A Gautier, and Susan Vroman. Directed search in the housing market. *Review of Economic Dynamics*, 19:218–231, 2016.

David H. Autor, David Dorn, and Gordon H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68, October 2013. doi: 10.1257/aer.103.6.2121. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2121.

Kassoum Ayouba, Marie-Laure Breuillé, Camille Grivault, and Julie Le Gallo. Does airbnb disrupt the private rental market? an empirical analysis for french cities. *International Regional Science Review*, 43(1-2):76–104, 2020.

Kyle Barron, Edward Kung, and Davide Proserpio. The effect of home-sharing on house prices and rents: Evidence from airbnb. *Available at SSRN 3006832*, 2018.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119 (1):249–275, 2004.

Mark Borgschulte, David Molitor, and Eric Zou. Air pollution and the labor market: Evidence from wildfire smoke. Technical report, Working Paper, 2018.

A Colin Cameron and Douglas L Miller. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372, 2015.

A Colin Cameron, Jonah B Gelbach, and Douglas L Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90 (3):414–427, 2008.

Andrew V Carter, Kevin T Schnepel, and Douglas G Steigerwald. Asymptotic behavior of a t-test robust to cluster heterogeneity. *Review of Economics and Statistics*, 99(4): 698–709, 2017.

Aaron Chalfin. The long-run effect of mexican immigration on crime in us cities: Evidence from variation in mexican fertility rates. *American Economic Review*, 105 (5):220–25, 2015.

Tom Chang, Joshua Graff Zivin, Tal Gross, and Matthew Neidell. Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy*, 8(3):141–69, 2016.

Kenneth Y Chay and Michael Greenstone. Air quality, infant mortality, and the clean air act of 1970. Technical report, National Bureau of Economic Research, 2003.

Karen Clay and Nicholas Z Muller. Recent increases in air pollution: Evidence and implications for mortality. Technical report, National Bureau of Economic Research, 2019.

Peter A Coles, Michael Egesdal, Ingrid Gould Ellen, Xiaodi Li, and Arun Sundararajan. Airbnb usage across new york city neighborhoods: Geographic patterns and regulatory implications. *Forthcoming, Cambridge Handbook on the Law of the Sharing Economy*, 2017.

Diane Coyle and Timothy Yeung. Understanding airbnb in fourteen european cities. *The Jean-Jacques Laffont Digital Chair Working Papers*, 7088:1–33, 2016.

Janet Currie and Matthew Neidell. Air pollution and infant health: what can we learn from california's recent experience? *The Quarterly Journal of Economics*, 120(3): 1003–1030, 2005.

Helping You Make Informed Decisions. A compass for understanding and using american community survey data. 2008.

Tatyana Deryugina, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif. The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, 109(12):4178–4219, 2019.

Taryn Dinkelman. The effects of rural electrification on employment: New evidence from south africa. *American Economic Review*, 101(7):3078–3108, December 2011. doi: 10.1257/aer.101.7.3078. URL https://www.aeaweb.org/articles?id=10.1257/aer.101.7.3078.

Antoine A Djogbenou, James G MacKinnon, and Morten Ørregaard Nielsen. Asymptotic theory and wild bootstrap inference with clustered errors. *Journal of Econometrics*, 2019.

Keith Finlay and Leandro M Magnusson. Two applications of wild bootstrap methods to improve inference in cluster-iv models.'. *Journal of Applied Econometrics*, 34, 2019.

Michael D Flannigan, Brian J Stocks, and B Mike Wotton. Climate change and forest fires. *Science of the total environment*, 262(3):221–229, 2000.

Miquel-Àngel Garcia-López, Jordi Jofre-Monseny, Rodrigo Martínez Mazza, and Mariona Segú. Do short-term rental platforms affect housing markets? evidence from airbnb in barcelona. 2019.

David Genesove and Lu Han. Search and matching in the housing market. *Journal of Urban Economics*, 72(1):31–45, 2012.

Anjali Haikerwal, Muhammad Akram, Anthony Del Monaco, Karen Smith, Malcolm R Sim, Mick Meyer, Andrew M Tonkin, Michael J Abramson, and Martine Dennekamp. Impact of fine particulate matter (pm 2.5) exposure during wildfires on cardiovascular health outcomes. *Journal of the American Heart Association*, 4(7): e001653, 2015.

Rema Hanna and Paulina Oliva. The effect of pollution on labor supply: Evidence from a natural experiment in mexico city. *Journal of Public Economics*, 122:68–79, 2015.

Bruce E Hansen and Seojeong Lee. Asymptotic theory for clustered samples. *Journal of Econometrics*, 2019.

Christian B Hansen. Asymptotic properties of a robust variance matrix estimator for panel data when t is large. *Journal of Econometrics*, 141(2):597–620, 2007.

Keren Horn and Mark Merante. Is home sharing driving up rents? evidence from airbnb in boston. *Journal of Housing Economics*, 38:14–24, 2017.

Dayne Lee. How airbnb short-term rentals exacerbate los angeles's affordable housing crisis: Analysis and policy recommendations. *Harv. L. & Pol'y Rev.*, 10:229, 2016.

Jos Lelieveld, Andrea Pozzer, Ulrich Pöschl, Mohammed Fnais, Andy Haines, and Thomas Münzel. Loss of life expectancy from air pollution compared to other risk factors: a worldwide perspective. *Cardiovascular Research*, 2020.

Kung-Yee Liang and Scott L Zeger. Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1):13–22, 1986.

Jia C Liu, Gavin Pereira, Sarah A Uhl, Mercedes A Bravo, and Michelle L Bell. A systematic review of the physical health impacts from non-occupational exposure to wildfire smoke. *Environmental research*, 136:120–132, 2015.

James G MacKinnon. How cluster-robust inference is changing applied econometrics. *Canadian Journal of Economics/Revue canadienne d'économique*, 2019.

James G MacKinnon and Matthew D Webb. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254, 2017.

Tristan-Pierre Maury and Fabien Tripier. Search strategies on the housing market and their implications on price dispersion. *Journal of Housing Economics*, 26:55–80, 2014.

Shawn J. Mccoy and Xiaoxi Zhao. Wildfire and infant health: a geospatial approach to estimating the health impacts of wildfire smoke exposure. *Applied Economics Letters*, page 1–6, 2020. doi: 10.1080/13504851.2020.1730747.

Nolan Miller, David Molitor, and Eric Zou. Blowing smoke: Health impacts of wildfire plume dynamics., 2017.

Espen R Moen. Competitive search equilibrium. *Journal of political Economy*, 105(2): 385–411, 1997.

Espen R Moen, Plamen T Nenov, and Florian Sniekers. Buying first or selling first in housing markets. *Journal of the European Economic Association*, 2014.

Petra Moser, Alessandra Voena, and Fabian Waldinger. German jewish Émigrés and us invention. *American Economic Review*, 104(10):3222–55, October 2014. doi: 10.1257/aer.104.10.3222. URL https://www.aeaweb.org/articles?id=10.1257/aer.104.10.3222.

Giovanni Quattrone, Davide Proserpio, Daniele Quercia, Licia Capra, and Mirco Musolesi. Who benefits from the sharing economy of airbnb? In *Proceedings of the 25th International Conference on World Wide Web*, pages 1385–1394. International World Wide Web Conferences Steering Committee, 2016.

Fabienne Reisen, Sandra M Duran, Mike Flannigan, Catherine Elliott, and Karen Rideout. Wildfire smoke and public health risk. *International Journal of Wildland Fire*, 24(8):1029–1044, 2015.

Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of economic literature*, 43(4):959–988, 2005.

Wolfram Schlenker and W Reed Walker. Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, 83(2):768–809, 2016.

Robert Shimer. *Essays in search theory*. PhD thesis, Massachusetts Institute of Technology, 1996.

Shawn P Urbanski, Wei Min Hao, and Stephen Baker. Chemical composition of wildland fire emissions. *Developments in Environmental Science*, 8:79–107, 2008.

Anthony LeRoy Westerling. Increasing western us forest wildfire activity: sensitivity to changes in the timing of spring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1696):20150178, 2016.

William C Wheaton. Vacancy, search, and prices in a housing market matching model. *Journal of political Economy*, 98(6):1270–1292, 1990.

Alwyn Young. Consistency without inference: Instrumental variables in practical application. 2017.

Yu Zhu, Randall Wright, Damien Gaumont, et al. Modeling house prices. In *2017 Meeting Papers*, number 744. Society for Economic Dynamics, 2017.