Title

Essays in Applied Microeconometrics

Author

Taylor, Reid Barrett

Publication Date

2024

Peer reviewed|Thesis/dissertation

Essays in Applied Microeconometrics

By

REID BARRETT TAYLOR

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

David Rapson, Chair

James Bushnell

Erich Muehlegger

Committee in Charge

2024

To Léonie

Contents

Abstract

This dissertation analyzes the role of regulation and competition in a variety of markets.

In the first chapter, I examine how regulation and market structure can lead to market unraveling when firms face rapidly increasing risk due to climate change. I first present a simple model of an adversely selected disaster insurance market to investigate how price regulation and increasing climate costs impact private markets in the presence of an insurer of last resort. I then empirically study the California non-renewal moratoriums, a first-of-its-kind regulation aimed at stymieing the retreat of insurance companies from high wildfire risk areas by forcing insurers to supply insurance to current customers following disasters in 2019 and 2020. Using quasi-random geographic variation in regulatory borders and a difference-in-differences identification strategy, I find that the moratoriums successfully reduced company-initiated non-renewals in the short run. However, the effects only lasted for one year, with insurers resuming non renewals as soon as the moratorium lapsed. Additionally, the moratoriums had no discernible effect on participation in the State's insurer of last resort.

In the second chapter, I estimate the effect of nearby entry and exit on the pricing of incumbent firms using high frequency price data and the precise geographic location of all gas stations in California. Using a difference-in-differences design, I find that entry of a new station is associated with a two-cent decrease in prices at incumbent stores, which equates to a 5% decrease in estimated retail markups. The effects are immediate, persistent, and show no sign of deterrence or limit pricing behavior. However, nearby exit results in precisely estimated null effects on prices. Both results are robust to various specifications and market size definitions. This paper also contributes to the conversation on California's "Mystery Gas Surcharge", a divergence between California and nationwide gas prices after the February 2015 Torrance refinery fire that cannot be explained by differences in taxes, regulation, or input and spot prices (Borenstein, 2023). I show that the magnitude of entry effects was unchanged, providing evidence that retail competition was largely unaffected.

In the third chapter, I study the link between auto insurance and health insurance market. specifically, I study how the increase in health insurance status from the implementation of the Affordable Care Act (ACA) overlapped with automobile insurance. Using data on automobile insurance claims from the National Association of Insurance Commissioners, I identify a plausibly causal impact of the ACA on the frequency of automobile insurance claims. These increases are restricted to the private market portion of the ACA and are associated with a 4.7% increase in bodily injury liability claim frequency.

Acknowledgments

This dissertation has been made possible thanks to the help and support of many friends, family and colleagues. I would like to begin by thanking my committee members, David Rapson, Jim Bushnell, and Erich Muehlegger. Your support and guidance over the past years has been invaluable. I would also like to thank the many other faculty members at UC Davis and CU Denver who helped guide me along this journey. Lastly, thank you to Bruce Mayer, Dan Messerschmidt, and Ed DeClair from Lynchburg College for first showing me what research was.

To my wife, Katie, thank you for your sacrifice and support. Without you, none of this would have been possible and I will be eternally grateful. To my parents, thank you for your unwavering commitment to expecting the best from me and support all these years. To my daughter Léonie, you give me a sense of purpose every day, and your smiles and giggles do more than you will ever know.

To my GYS brothers, Patrick and Jamie Jacobs, your daily chats, debates, and laughter kept my sanity. To Dodd Mitchell, Thomas Anderson, Andew Gallagher, and countless others, thank you for being great friends and providing conversation, games, distraction, and support.

CHAPTER 1

Climate Change and the Regulation of a Crashing Insurance Market (With Madeline Turland and Joakim Weill)

1.1. Introduction

Natural disasters pose substantial financial risks to households, firms, and communities, highlighting the urgent need for well-functioning insurance markets. However, escalating costs associated with disasters, compounded by spatially correlated and potentially catastrophic losses, present new challenges for insurers. In a majority of US states, "insurers of last resort" provide coverage to households unable to purchase insurance elsewhere. Initially meant to temporarily insure only the riskiest properties, state-run insurers of last resort now hold substantial market share in many places across the United States. This presents a puzzle for insurance market regulators: why are insurers of last resort growing, and which policies can help prevent private insurance markets from unraveling?

In this paper, we examine how price regulation can result in the unraveling of the private insurance market, and thus full reliance on the insurer of last resort, when insurance firms face rapidly increasing risk due to climate change. We begin by constructing a conceptual model of a natural disaster insurance market featuring price regulation, a residual risk pool serviced by the insurer of last resort, and a private ("voluntary") market in which firms can use information about the marginal cost of consumers to choose which properties to insure. We show how the presence of a residual market is necessary to guarantee full market coverage when regulation constrains price below the expected average cost. When firms observe the underlying risk profile of consumers, they choose not to insure those that have expected costs above the regulated price, resulting in a bifurcated market as a portion of consumers receive coverage in the residual risk pool.

Our framework rationalizes the recent dynamics observed in the California homeowners insurance market, the largest homeowners insurance market in the country. Following consecutive record-setting wildfire seasons in 2017 and 2018, insurers in California refused to renew more than

200,000 homeowner's insurance policies primarily in areas of high wildfire risk (Bikales, 2020), a 61% increase from prior years, citing restrictive price-setting regulations (California Department of Insurance, 2021). At the same time, take-up of policies offered through the state's insurer of last resort, the California FAIR plan, spiked in the same areas. This is consistent with private firms reducing their exposure by ceding the highest risk policies to the insurer of last resort as cost increases outpaced increases in the regulated price. The model predicts that if prices are not allowed to reflect climate costs, then adjustments must occur through the insurer of last resort.

We then study a first-of-its-kind policy that California implemented to counter the explosive growth in its insurer of last resort: the non-renewal moratorium, which forces insurers to continue supplying insurance for at least one year to certain homeowners. The regulation impacts zip codes located near state-declared 'state of emergency' wildfires for one year following the emergency declaration, first in 2019 and then in 2020. We exploit the quasi-random variation generated by the regulatory boundaries of the moratoriums to study the causal impact of the policy on insurance market outcomes. Our treatment group includes zip codes adjacent to but not directly impacted by the wildfires. This avoids confounding the treatment effect of the moratorium with the direct effects of the fire, which are likely correlated with our outcomes of interest. Our control group comprises zip codes directly adjacent to those impacted by the regulation. These areas offer a credible counterfactual, as they closely resemble the treatment zip codes and were not subject to the moratorium: the quasi-random occurrence of wildfire ignition sites and zip code boundaries suggests their exemption from regulation was purely coincidental. In robustness tests, we show that results are similar with an alternative control group that uses nearest-neighbor matching between treatment zip codes and observably similar but untreated zip codes in the rest of the state.

We find that the moratoriums successfully increase insurance supply by decreasing company-initiated non-renewals while they are active, with no evidence that firms are able to avoid the regulation by forcing out customers using other methods. However, this effect is short-lived; firms increase non-renewals by 72% to 96% as soon as the year-long moratoriums ended. Additionally, we estimate that the moratorium had no discernible impact on slowing the transition of policies from the voluntary market into the FAIR plan. While the regulation restricted non-renewals of currently insured customers, it had no effect on firms refusing to insure new customers. These results highlight that the California moratorium only acted as a short-term band-aid, and that deeper changes of the rate-setting guidelines are required to avoid an unraveling of the market.

This paper first contributes to a growing literature on natural disaster insurance markets (Born and Klimaszewski-Blettner, 2013; Knowles and Kunreuther, 2014; Kousky, 2011, 2022; Kunreuther, 1996, 2001; Marcoux and Wagner, 2024; Oh et al., 2023). Due to both the historical importance of flood losses and the lack of publicly available data on other homeowner insurance policies, work in this area largely focuses on flood insurance (Bradt et al., 2021; Gallagher, 2014; Mulder, 2022; Wagner, 2022; Weill, 2023). Two recent exceptions are Sastry et al. (2023), who focus on wind damage in Florida, and Boomhower et al. (2023), who examine the impacts of wildfires on the Californian homeowner insurance market. Boomhower et al. (2023) investigate the issue of wildfire risk estimation for insurers, while our paper focuses on the interaction between private insurance markets and the state insurer of last resort, with an application to the California non-renewal moratoriums.

This paper also contributes to the literature on insurance regulation, both theoretically and empirically. We develop a model of adverse selection in a segmented market under price regulation and rapidly changing risk that extends the canonical model of Einav et al. (2010), which has been used broadly to study adverse selection in insurance and lending markets (Boyer et al., 2020; Cabral and Cullen, 2019; Spinnewijn, 2017).[1] A majority of states are now operating "residual markets" or an "insurer of last resort" (Kousky, 2011). We provide a simple framework to investigate how these markets interact with private markets. Our paper also shows how regulating both price and risk selection, as is the case with the California non-renewal moratoriums, can lead to long-run firm exit and unravelling of the private market. Our work expands on the natural disaster insurance regulation literature (Born and Klimaszewski-Blettner, 2013; Born and Viscusi, 2006; Oh et al., 2023) and the smaller set of studies focused on regulatory effectiveness and efficiency in California's insurance market (Liao et al., 2022).

Finally, this work fits into a broader literature on climate adaptation (Barreca et al., 2016; Botzen et al., 2019; Diaz and Moore, 2017; Kahn, 2021; Kousky, 2019; Sastry, 2021). We first contribute to a large literature that focuses on housing markets and climate change; numerous studies document how insurance pricing and access impact the real estate market (Issler et al.,

[1] See Einav and Finkelstein (2023) for a review of literature that makes use of the framework from Einav et al. (2010).

2024; Nyce et al., 2015) and how natural disaster risk can impact mortgage repayment (Biswas et al., 2023; Xudong et al., 2024). We also add to the small but growing literature on firm level adaptation to climate change (Bilal and Rossi-Hansberg, 2023; Castro-Vincenzi, 2022; Gu and Hale, 2022; Prankatz and Schiller, 2021).

This paper proceeds as follows. Section 2 provides background on the California Moratoriums while section 3 introduces our conceptual model. Section 4 presents the data used, section 5 introduces the econometric framework and section 6 presents the results. Section 7 concludes.

1.2. Institutional Background

1.2.1. Insurance Markets. The price of an insurance policy is set before any potential losses are incurred.[2] This implies that insurers' profitability depends on both accurate projections of expected losses, and premiums that reflect these projections. Theoretically, if firms were unconstrained in their ability to calculate and set policy premiums, they would be able to offer a price for all risks. However, regulations have emerged with the dual goals of protecting customers from rates that are unfairly discriminatory and unreasonable, and ensuring premiums are sufficient to guarantee solvency. In some cases, regulation can lead premiums to diverge from expected costs through both suppressing premium growth and limiting the firm's ability to accurately incorporate cost forecasts. We focus our discussion on the main distortions rate regulation introduces to natural disaster risk pricing in the California homeowner's insurance market, which generalize to a varying degree to other state markets.

Before new insurance rates can be implemented, insurers must obtain prior approval of rates with the state Department of Insurance. This administrative process is cumbersome, and frequently lasts more than 12 months (Oh et al., 2023). The specifics of the rate approval process vary widely between states. For instance, in California, regulators face three specific regulations which suppress premium growth in practice. First, overall rate increases of 7% or higher (calculated over the entire insurer portfolio) are subject to in-depth public scrutiny at the unrecoverable cost of the insurer (California Ballot Propositions and Initiatives, 1988). This regulation has resulted in an effective rate increase cap as most rate increase filings are below this threshold, with significant bunching

[2]Most property and casualty lines of insurance follow experience rating whereby premiums can be adjusted for losses incurred in previous contract periods. However, some policies use retrospective rating, which settles the final premium amount due at the end of the period and takes into account losses from that same period. Retrospective rating is generally reserved for worker's compensation and commercial policies.

at 6.9% (Boomhower et al., 2023).[3] Second, California regulation requires the overall rate for natural disasters, known as a catastrophe load, to be justified by historical averages of losses over at least the past 20 years (California Code of Regulations, 2024).[4] Until 2023, insurers in California were not allowed to incorporate forward-looking catastrophe models or other means of forecasting as justification for higher catastrophe loads, exacerbating premium inadequacy when past loss experience does not reflect future expectations. Finally, California restricts firms from passing reinsurance costs through to consumer premiums. Recent industry literature has highlighted greatly increasing reinsurance premiums as climate risks increase, with reinsurance companies not subject to the same regulatory oversight as consumer facing insurance companies. This further drives the difference between the cost firms incur and the premium they are able to charge the customer.

State regulation also routinely specifies which observable home and homeowner characteristics are permissible in the underwriting and rating processes. While the classic case of adverse selection relies on consumers having private information unobserved by the firm, adverse selection can also arise from regulation restricting the set of permissible characteristics used for pricing. Thus, even after conditioning on the permissible observables, consumers that are offered the same premium can still vary in their expected costs. Regulation in California in late 2022 made two changes to the underwriting process: first, firms were forced to incorporate defensible space characteristics into their rating plan, and second, any use of catastrophe or risk scoring to underwrite or create rate differentials had to be filed with the state. The setting of rate differentials is distinct from catastrophe load calculations, which, as previously stated, have never been permitted to integrate catastrophe models. The second regulation change presents a hurdle to firms as they were given the option of publicly filing proprietary and confidential models (some contracted through 3rd party companies) or to cease their use.

1.2.2. California Moratoriums. In response to large losses from record breaking wildfires in 2017 and 2018, insurance companies began to withdraw from high wildfire-risk areas. A standard, one-year insurance policy can typically only be cancelled mid-term by the insurer due to lack of payment or material fraud on behalf of the insured. However, an insurer is able to non-renew

[3]Because the 7% regulatory constraint is calculated over the full portfolio of the insurer, the premiums can still increase faster than 7% for some homeowners.

[4]Insurers can opt to weigh certain years more than others in the premium calculations. However, actuarial convention and the regulator can push back on nonuniform weighting schemes, especially if firms are over-weighting certain years to increase expected costs.

(not offer a subsequent contract) for a wider range of reasons, including changing beliefs about the probability of a claim. In an attempt to stymie the retreat of insurance companies from high-risk locations, the California legislature passed Senate Bill 824 in 2018. This bill prohibits insurance companies from non-renewing a policy because of wildfire risk in any zip code either directly impacted by, or adjacent to, a wildfire that was declared a state of emergency by the state government. The commissioner of the department of insurance cited the bill as giving, "millions of Californians breathing room and hits the pause button on insurance non-renewals while people recover."[5] The regulation impacts firms by limiting their ability to geographically diversify and to drop policies which are likely otherwise unprofitable given the firm's rating plan.

Each moratorium lasts one year from the date of disaster declaration. For the years examined in our study, the earliest start date for a moratorium is August 18 and the latest start date is November 18. We refer to the moratoriums by yearly cohorts. The collection of non-renewal moratoriums initiated following the 2019 fire season is the "2020 Moratorium", while those initiated after the 2020 fires is referred to as the "2021 Moratorium".

Due to the stochastic nature of wildfires, and specifically wildfire perimeters, zip codes located near each other can be differentially impacted by the moratorium despite being observably similar. Additionally, high risk areas in other parts of the state that have not yet experienced a fire post-legislation are not covered by the moratoriums, despite being similar. The quasi-random nature of the initial coverage of the moratorium, coupled with the lack of lead time and anticipation for firms, forms the basis of our identification strategy to identify the causal impacts of the moratoriums on various insurance and consumer outcomes.

1.2.3. Insurers of Last Resort. Residual markets, or "insurers of last resort", are state-run or state-sponsored plans that sell coverage for properties considered very high risk and unable to find insurance in the voluntary market. Insurers of last resort vary between states; Fair Access to Insurance Requirements (FAIR) plans were established in twenty-six states, the District of Columbia, and Puerto Rico in 1968 following the riots and bushfires of the 1960s (Dwyer (1978) offers an illuminating discussion of the establishment of FAIR Plans). Some states have "wind pools" or "beach plans", which cover specific perils with limited geographic coverage. These plans are typically more expensive than policies in the voluntary market and their coverage varies by state

[5]See https://www.insurance.ca.gov/01-consumers/140-catastrophes/MandatoryOneYearMoratoriumNonRenewals.cfm.

(see Kousky (2011) for a discussion of ten state-sponsored disaster insurance programs). Today, 34 states and the District of Columbia have an insurer of last resort (shown on Figure 1.1) – Colorado became the latest state to establish a FAIR Plan in 2023 in response to increasing wildfire losses.

Figure 1.1. Insurers of last resort in the United States

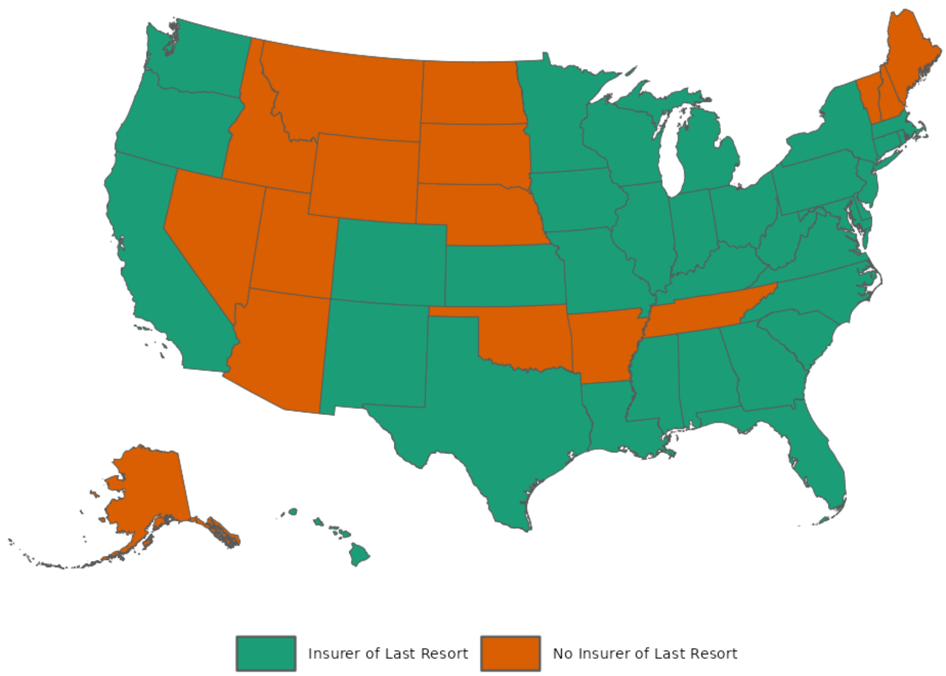

The California FAIR Plan issues policies on behalf of all companies licensed to write property and casualty insurance in California. Each member company participates in the profits, losses, and expenses in direct proportion to its market share in the state and thus are invested in the financial stability of the FAIR Plan. Its purpose is to provide temporary, basic fire insurance when traditional insurance is not available. It is designed to act as an emergency safety net while homeowners search for insurance in the traditional (or voluntary) market. California FAIR Plan policies are typically more expensive than the voluntary market, have a maximum coverage amount of $3 million, and require customers to obtain an additional Difference in Coverage (DIC) policy in order to replicate the coverages offered in a standard homeowner's insurance policy. Despite these differences, FAIR Plan and voluntary insurance policies are substitutable in the market because both can satisfy the requirements of mortgages lenders for homeowner's insurance.

1.3. Conceptual Model

This section presents a general, but simple, framework to clarify the role of the residual insurance market in a market with adverse selection. We begin by characterizing supply and demand in the presence of price regulation, and use a graphical approach to illustrate how the market equilibrium changes with increasing wildfire risk.

Our framework closely follows Einav et al. (2010); consumers make a discrete purchase decision for a homogeneous full-coverage insurance policy, which they buy at the lowest price available from profit-maximizing firms competing in the market. Consumers purchase their policy from either the private ("**voluntary**") market or the **residual** market.

We focus on two distinct pricing regulations. First, conditional on a set of permissible property characteristics $\{c_i\}$, the regulator sets a fixed price $\tilde{P}$ in the voluntary market. Second, in the residual market, prices can adjust freely but the regulator imposes a zero-profit condition.

The pricing constraint imposed in the voluntary market captures features found in state-level regulation of homeowner insurance rates.[6] In California, the Department of Insurance limits the characteristics $\{c_i\}$ that insurers can use to determine premiums. In particular, estimates of risk from catastrophe models cannot currently be used to set household-level insurance premiums (California Code of Regulations, 2023).[7] Insurers in California determine rates through a complex rate-filing process, in which requested premium increases averaged over the insurer portfolio must stay below 7% to avoid a costly public hearing.[8]

Most property-level characteristics that impact expected losses and that are observable to homeowners (such as location, building materials, number of floors, etc.) are readily observable to insurers. This stands in contrast to health or auto insurance markets, where consumers typically have private information about their expected losses. However, regulations that restrict the observable characteristics, c_i, permitted for rate-making prevent insurers from achieving perfect price discrimination, resulting in what Finkelstein and Poterba (2014) call "asymmetrically used information". Imperfect price discrimination manifests itself as consumers with the same permissible characteristics being charged the same price, *despite* otherwise observable differences in their expected

[6] For example, California, Hawaii, Maryland, Massachusetts, Michigan, Nevada, Oregon and Utah prohibit the use of credit scores to determine home insurance rates.

[7] Relaxing premium regulation to allow for catastrophe modeling is an active debate in California (State of California Department of Insurance, 2023; Watkins and Lee, 2022).

[8] Additional details regarding insurance rate regulation in California are discussed in Boomhower et al. (2023).

costs. This information asymmetry results in adverse selection, characterized by consumers with the highest expected costs also having the highest willingness to pay resulting in downward sloping marginal and average cost curves Einav et al. (2010). For simplicity, we assume that demand is higher than average cost at every point, implying that at actuarially fair prices, every consumer prefers to purchase insurance than to go without.[9]

Although firms do not have the freedom to set prices, they control which consumers to offer a contract at the regulated price, and thus quantities in the voluntary market. Insurers can observe the marginal cost curve and decide not to offer insurance contracts to certain properties. However, the California moratoriums implemented in 2020 and 2021 directly eliminated this decision-making variable for insurers.

1.3.1. Market Segmentation. In the graphical analysis that follows we depict one tranche of the market where all consumers have the same set of permissible rating characteristics and are charged the same premium, but vary in their expected losses. In Panel (a) of Figure 1.2, we consider the case where the regulator imposes an exogenous price $\hat{P}$ below the average cost curve at every point. While all consumers would opt to buy insurance at this price, insuring the entire market (Q^{max}) would lead to negative expected profits for firms.

In Panel (b), firms use their knowledge about the marginal cost curve to select which consumers they offer coverage. They choose to offer coverage only to consumers that are profitable, such that $\hat{P} \geq MC$. This results in only a portion of the market receiving insurance coverage from the voluntary market, consumers from Q^R to Q^{max}. The remainder of the market (Q^0 to Q^R) is forced to purchase from the residual market. The zero-profit condition imposed on the residual market results in price being set at the average costs over the customers that purchase from the residual insurer: $P^R = AC^R(Q^R)$. By construction, all customers ceded from the voluntary market will purchase a policy from the residual insurer because their willingness-to-pay is greater than the average cost curve at every point.[10] Positive profits (shown in blue) in the short run are possible in

[9]According to the National Association of Insurance Commissioners, about 90% of homeowners have insurance, largely due to the requirement to buy insurance to obtain a mortgage.

[10]In an alternate scenario, if the demand curve were steep enough, consumers that are marginally ceded from the voluntary market will not purchase from the residual market as the pooled price is higher than their willingness-to-pay.

the voluntary market because prices are fixed at the regulated level and the costs of market entry are non-trivial.[11]

This stylized example highlights the crucial role played by the residual market; it ensures that all consumers can purchase an insurance contract *despite* the price regulation. To see this, note than in the absence of the residual market, consumers with marginal costs above the regulated price cannot buy insurance *regardless of how much they are willing to pay* as no firm would be willing to insure them. In contrast, in the absence of regulation, competitive firms would perfectly discriminate and charge each consumer a price equal to their marginal cost, ensuring insurance availability. The residual market thus allows price-suppressing regulation to be sustainable in the voluntary market.

Relative to the perfect price discrimination benchmark, where the price for each customer is equal to their marginal cost, the scenario with a residual market and price regulation entails clear distributional consequences. All consumers in the voluntary market are charged more than their marginal costs, with the lowest risk consumers paying the highest markups. Consumers buying in the residual market are charged an average cost necessarily greater than both the regulated price of the voluntary market ($\hat{P}$) and the average cost pooled across the total market, as the risk pooling is concentrated on only the highest risk consumers. Within the residual market, the riskiest consumers are charged less than their marginal cost, while the least risky consumers are charged more than their marginal cost.

FIGURE 1.2. Baseline Market

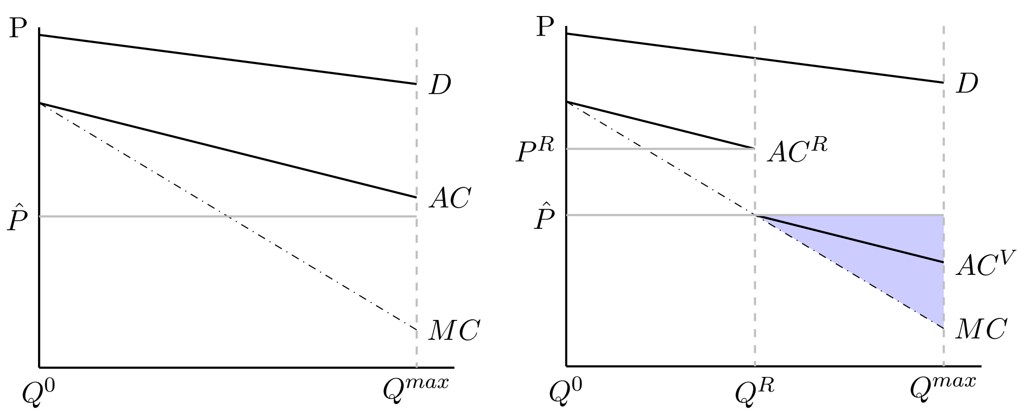

(a) No Segmentation (b) Presence of a Residual Market

[11]Characterization of the long-run, dynamic nature of rate requests and the role profits play in future negotiation with the regulators is beyond the scope of this paper.

1.3.2. Increasing perceptions of wildfire risk. Consider what happens when the industry experiences an extreme weather event (or series of events) that causes insurers to update their risk perceptions. We distinguish between actual risks (which we assume remain constant) and perceptions of risk (which are impacted by recent events). We assume that when perceptions change, they shift closer to the true risk levels. In Figure 1.3, MC' and AC' represent the increase in perceived insurance costs following a particularly bad series of extreme events. For simplicity we assume these curves are a parallel shift in expected costs for each consumer.

We focus on what happens when the regulated price adjusts more slowly than perceived costs increase, as is the case in our setting. That is, the regulated price increases from $\hat{P}$ to $\hat{P}'$, with $\hat{P}' - \hat{P} < AC' - AC$. Such situations can occur due to the length of the rate-filing process, or due to specific regulatory constraints. In contrast, because the price in the residual market is not constrained by the regulated price, price adjusts to the average cost over customers already insured in the residual market, $P^{R'} = AC^{R'}(Q^R)$, keeping profits equal to zero in the residual market.

Panel (a) of Figure 1.3 coincides with the market under the California non-renewal moratoriums when firms are not allowed to adjust which consumers they serve in the voluntary market (Q^R is held constant). As premiums increase, consumers suffer a reduction in wealth represented by the red rectangle in the residual market and the red and purple areas in the voluntary market. Firms in the voluntary market expect higher costs, shown by the grey and purple areas, which is only partially offset by the increase in the regulated price. In this situation, firms lose money on consumers with a new marginal cost greater than the new regulated price, and would choose not insure these properties absent the moratorium. If the perceived increase in costs is high enough that the average cost over the remaining consumers is above the regulated price, firms will exit the voluntary market in the long run and the market will collapse.

FIGURE 1.3. Market With Expected Cost Increase

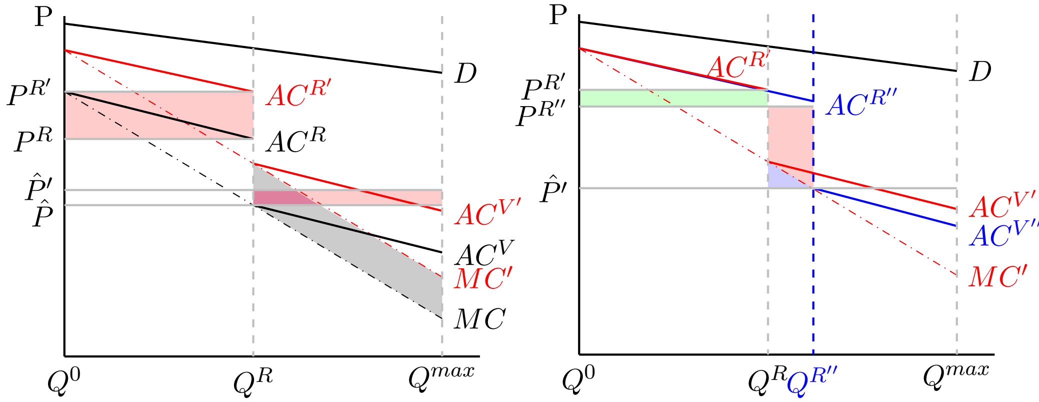

(a) No Residual Market Adjustment (b) Residual Market Adjustment

Panel (b) depicts the market if firms are able to adjust their portfolio given the new regulated price and cost curves, meaning that Q^R is allowed to adjust. In relation to our empirical study, this coincides with the market at the end of the non-renewal moratorium. Following the increase in risk perceptions, the voluntary market cedes any consumers who have a new marginal cost higher than the new regulated price, which are between Q^R and $Q^{R''}$, where $Q^{R''}$ is determined by the intersection of the new marginal cost curve and the new regulated price. Because we consider a demand curve that is always above the residual market's average cost curve, the consumers dropped from the voluntary market will purchase insurance in the residual market. Given the residual market operates as a non-profit, and that the consumers dropped from the voluntary market have lower marginal costs than those already participating in the residual market, the residual market price drops to $P^{F''}$.

The costs of allowing adjustment from Q^R to $Q^{R''}$ are entirely born by the group of consumers forced out of the voluntary market as a result of the adjustment. These consumers lose the red and blue areas in Figure 1.3 (b) due to the higher price $P^{F''}$. The firms capture the blue portion of the welfare loss due to the reduction in expected losses. Customers already in the residual market experience a benefit, shown by the green rectangle, because the addition of lower-risk consumers to the risk-sharing pool reduces the price.

In sum, this model generates four simple predictions in relation to our empirical study of non-renewal moratoriums under changing risks: (i) premiums increase in both the residual market and voluntary market, (ii) consumers in the voluntary market are not ceded to the residual market when the moratorium is active, leading to short-run losses for firms, (iii) the residual market share increases when the moratoriums become inactive, and (iv) holding costs constant, residual market price decreases when the moratoriums become inactive. We test these predictions in the follow section and assess how the characteristics of consumers in the residual market and voluntary market changed following the moratoriums.

1.4. Data

1.4.1. Insurance Data. We obtain homeowners insurance data from the California Department of Insurance (DOI). These data are a combination of three separate products: the Community Service Statement (CSS), the Personal Property Exposure (PPE), and the Residential Property Experience (RPE). The CSS contains information on earned exposures, earned premiums, number of policies, and average premium at the company-zip code-year level for all insurance companies licensed to operate in California from 2009 to 2022. Importantly, the California FAIR plan reports data in the CSS alongside companies in the voluntary market which allows us to calculate a zip code level FAIR plan market share. The PPE survey reports the amount of coverage provided, number of units insured, and deductible amounts at the company-zip code-year level from 2009 to 2021. All companies writing more than $5 million in insurance in California are required to report. Lastly, the RPE data set reports the number of new, renewed, and non-renewed policies at the zip code-year level. Importantly, we observe whether the decision to non-renew the policy was initiated by the insurer or by the customer. The RPE is reported yearly from 2015 to 2021.

1.4.2. Wildfire Risk. We use the Risk to Potential Structures (RPS) data from the US Forest Service to construct zip code level measures of wildfire risk.[12] The RPS relates both the probability of a fire occurring as well as the expected intensity of a potential fire, asking the question, "What would be the relative risk to a house being located on this pixel?" Thus, the measure does not rely on the current presence of a building in order to assess the risk. This allows for an insurance relevant wildfire risk measure to be calculated even in sparsely populated portions of the state, and comparison between currently inhabited and not yet inhabited locations. We calculate the zip code level average RPS by calculating the mean across the RPS values for each 30 meter pixel

[12]Formally, zip codes are not geographic in nature, but yet relate a collection of mail routes. The census thus created Zip Code Tabulation Areas (ZCTA) which are geographic representations of zip codes. We use ZCTAs to construct all geographic level data to match the level of observation of our insurance data, but use the more common term "zip code" in the rest of the text.

located within the boundary of the zip code. We also calculate the standard deviation of the RPS values within a zip code to capture the variability of fire risk within a zip code. The RPS is time invariant and represents a snapshot of wildfire conditions modeled in 2014. In reality, wildfire risk can change over time following drought conditions and recent wildfire activity. Additionally, the RPS data do not account for changes or variation in home construction types, which is an important way homeowners can try to manage wildfire risk.

1.4.3. Wildfire Boundaries. We use geolocated fire perimeters from the California Department of Forestry and Fire Protection (CalFire)'s Fire and Rescue Assessment Program (FRAP) to identify the location of wildfires during our sample period. The fire perimeters are developed by CalFire jointly with the US Forest Service, the Bureau of Land Management, and the National Park Service, and the Fish and Wildlife Service covering both public and private lands in California. Data on the location, area covered, cause of the fire, and the responding agency are available. We exclude prescribed fires from our dataset. Wildfires occur in both Northern and Southern California, largely concentrated in the foothill and mountainous areas along both the coastal and Sierra Nevada ranges.

1.4.4. Non-Renewal Moratorium Status. We identify zip codes subject to a non-renewal moratorium in 2020 or 2021 using data from the office of the California Insurance Commissioner. We classify zip codes into 4 categories. 'Fire' zip codes are those that were included in the moratorium because they directly experienced a fire that was declared a disaster. 'Treatment' zip codes are included in the moratorium by regulation due to being adjacent to zip codes which burned, but did not directly experience the fire causing the disaster declaration. These zip codes form the basis of our identification strategy discussed in the following sections. 'Adjacent' zip codes are zip codes that are not included in the moratorium but share a border with a zip code covered under a moratorium. 'Rest of State' encompasses all other unimpacted zip codes. The Department of Insurance reports all zip codes subject to a moratorium without distinction, thus we spatially merge the fire perimeter data to differentiate the 'Fire' and 'Treatment' zip codes. Figure 1.4 depicts the various moratorium classifications for the state of California in 2020 and 2021, separately.

1.4.5. Descriptive Statistics. Table 1.1 presents summary statistics for the dataset, grouping zip codes by the 2020 moratorium classifications. As expected, 'fire' zip codes were also the

Figure 1.4. Zip Code Classifications

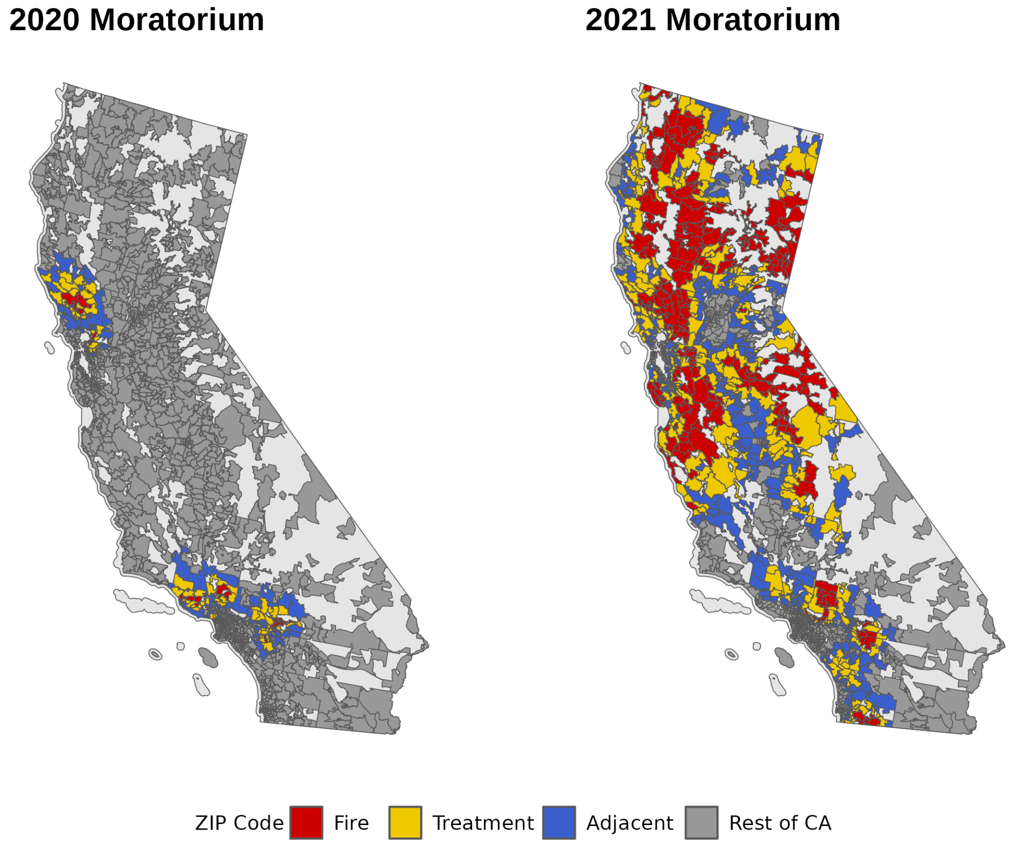

riskiest, *ex ante*, as measured by RPS, but indistinguishable from nearby treatment and adjacent zip codes while zip codes in the 'Rest of State' category have a notably lower wildfire risk level. This supports the assumption that while areas prone to wildfire are not random distributed, the precise location of wildfire events and the interaction with zip code boundaries is essentially random. While areas impacted by fires and the moratorium have higher FAIR Plan market shares, only 3% of the market is served by the FAIR plan on average in wildfire impacted zip codes. Over the entire sample period, there does not appear to be any trends associated with the number of new policies, renewals, or customer or company initiated non-renewals.

Figures 1.5 and 1.6 show the evolution of the number of company-initiated non-renewals, customer-initiated non-renewals, new policies, and renewals by zip code classification status separately for the 2020 moratorium and 2021 moratorium. Statistics are indexed relative to their level in 2015.

Table 1.1. Summary Statistics by Zip Code Classification (2020 Moratorium)

Zip code Classification	Fire		Treatment		Adjacent		Rest of State	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Premium (dollars)	1049.5	388.4	1099.2	579.8	1085.1	677.0	966.3	495.2
New Policies (count)	517.9	513.5	848.6	653.5	790.6	710.8	511.8	577.5
Renewals (count)	4262.1	3924.9	6309.0	4303.0	5672.2	4639.0	3995.7	4171.0
Customer Nonrenewals (count)	404.7	385.7	622.9	481.1	594.5	540.4	382.0	433.4
Company Nonrenewals (count)	101.2	100.3	168.0	128.3	143.5	123.1	95.3	102.9
FAIR Plan Market Share	0.03	0.05	0.02	0.04	0.02	0.05	0.02	0.04
RPS	0.5	0.6	0.5	0.6	0.4	0.5	0.2	0.4

For the 2020 moratorium, ‘fire’ and ‘treatment’ zip codes both saw a decrease in the number of company-initiated non-renewals in the year they were covered by the moratorium, consistent with the moratorium being a binding constraint on firm behavior. However, we also see a large reversal the year the moratorium was lifted for these zip codes. While only preliminary, this suggests that the moratorium was only successful at altering firm behavior in the short term. A potential concern in our setting is that firms are able to circumvent the moratorium by forcing-out customers through making their product less attractive in an effort to have them cancel, thereby disguising a company-initiated non-renewal as a customer-initiated non-renewal. This would result in the moratorium seeming more effective than it actually is. While we note here that customer-initiated non-renewals increased the most in ‘fire’ zip codes, we return to this question later in our results using a causal framework, showing that this is likely not a concern.

For both the 2020 and 2021 moratorium designations, we see a clear increase in both FAIR plan policies and the total number of policies in the year after the fire. This is consistent with evidence that large disasters can drive *ex post* demand for insurance, as well as the narrative that firms are accelerating their retreat from the highest risk zip codes.

Taken together, this descriptive evidence suggests that the moratorium may have had significant impacts on the market by reducing non-renewals in the short-term, but that these reductions were concentrated to just the year of the moratorium coverage. In our next section we formalize the assumptions needed to identify the causal impacts of the moratoriums.

Figure 1.5. Statistics by 2020 Moratorium Classification

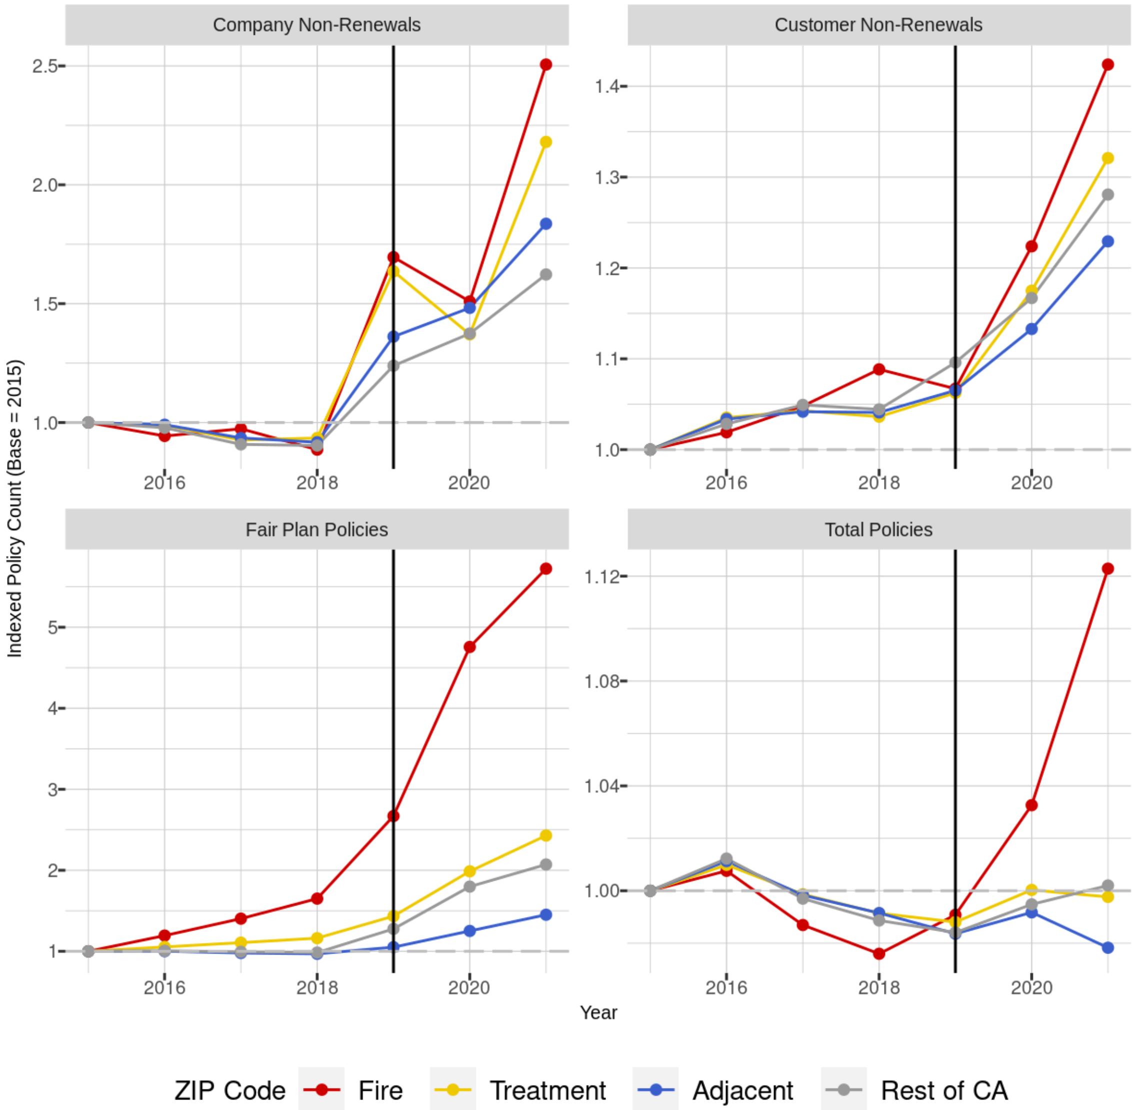

Notes: Zip codes are broken out by moratorium classifications. 'Fire' zip codes were directly impacted by a wildfire in 2019 and covered by the non-renewal moratorium in 2020. 'Treatment' zip codes were covered by the non-renewal moratorium in 2020 but did not experience a wildfire in 2019. 'Adjacent' zip codes share a border with zip codes covered by the non-renewal moratorium in 2020. 'Rest of State' zip codes are the remaining zip codes not covered by the moratorium.

FIGURE 1.6. Statistics by 2021 Moratorium Classification

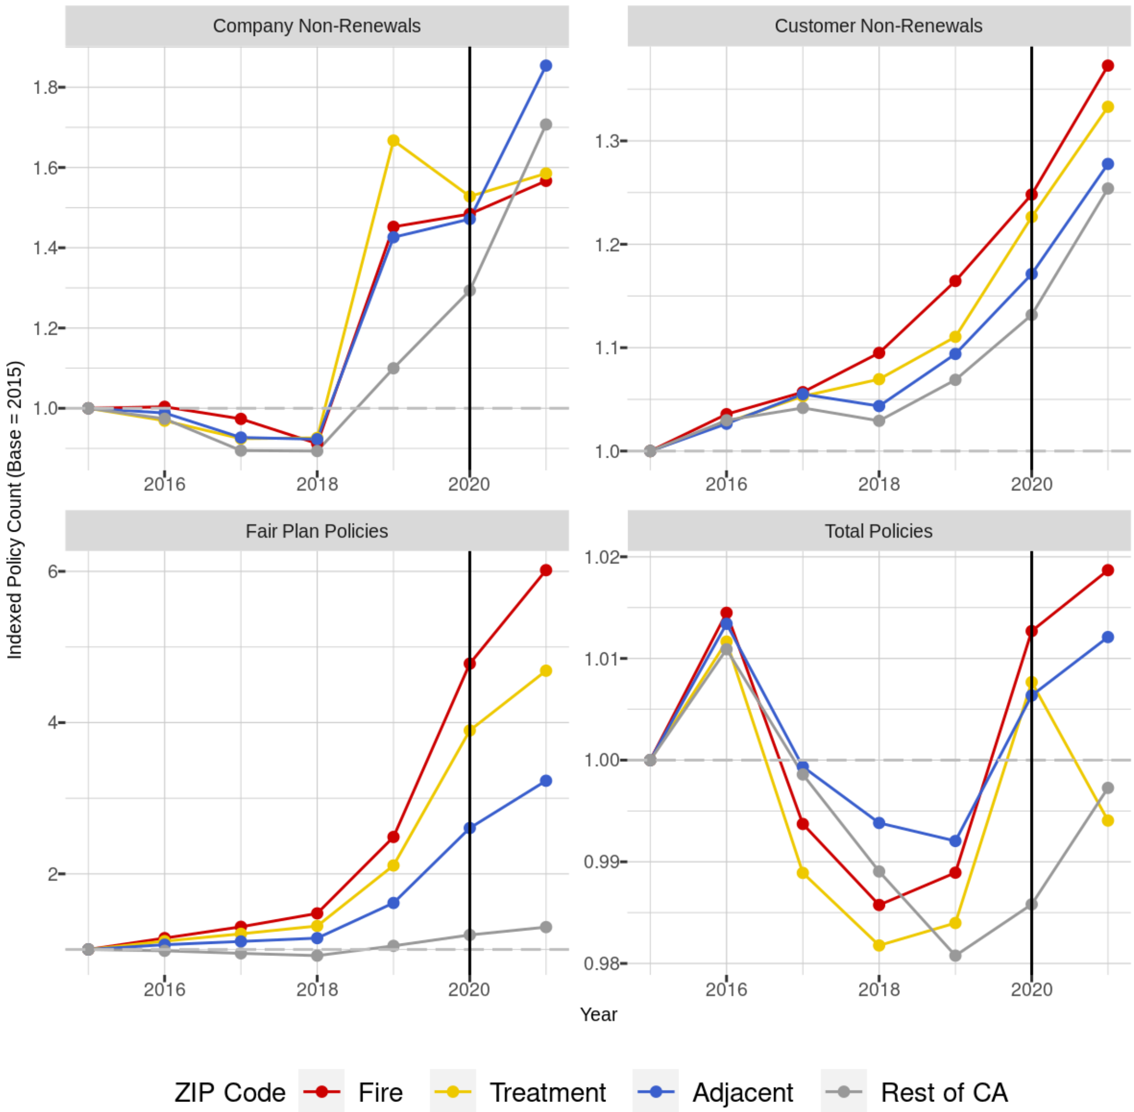

Notes: Zip codes are broken out by moratorium classifications. 'Fire' zip codes were directly impacted by a wildfire in 2020 and covered by the non-renewal moratorium in 2021. 'Treatment' zip codes were covered by the non-renewal moratorium in 2021 but did not experience a wildfire in 2020. 'Adjacent' zip codes share a border with zip codes covered by the non-renewal moratorium in 2020. 'Rest of State' zip codes are the remaining zip codes not covered by the moratorium.

1.5. Methods

The stochastic nature of wildfires and the unique geographic coverage of the California non-renewal moratorium allow for a difference-in-differences specification to recover causal estimates of the policy's impacts. We make use of the sharp geographic border discontinuity between neighboring zip codes, comparing zip codes covered by the moratorium to those zip codes located just outside the border of the moratorium, before and after the policy change.

The non-renewal moratorium covers policies in zip codes that experienced a state declared disaster fire and their immediate neighboring zip codes. Identification in our model requires that no other changes, contemporaneous with the policy, could explain the observed changes in the outcome variables. As such, we omit zip codes that are directly impacted by a disaster fire from treatment as they experience housing supply shocks, receive disaster relief funding, and are impacted by other unobserved factors perfectly correlated with the timing of the moratorium.

Our estimating equation is,

$$y_{zt} = \alpha + \sum_{j=0}^{1} \beta_j T_z D_j + \sigma_z + \delta_t + \varepsilon_{zt}, \tag{1.1}$$

where y_{zt} is the outcome of interest in zip code z in year t, T_z is the treatment indicator variable which takes a value of 1 if zip code z is impacted by a moratorium during the sample period, and D_t are post-period event time indicators taking a value of $D_j = 1$ for the year of the moratorium ($j = 0$) or the first year post treatment ($j = 1$). We include zip code fixed effects, σ_z, to control for time-invariant geographic heterogeneity correlated with wildfire risk and the insurance outcome variables, such as climate, elevation, slope, vegetation types, population density, and access to emergency services. We also include year fixed effects, δ_t, to account for common annual shocks across all units. This controls for unusually dry or hot seasons or macro-financial trends which impact the risk appetite of firms. We cluster model standard errors at the zip code level.

Identification of causal estimates from equation (1.1) relies on three main assumptions. First, the common trends assumption requires that outcome variables in both treatment and control areas should evolve along the same trend over time, and would have continued along a similar path absent the moratorium. We can provide supporting evidence for the assumption through estimation of the following event study analog of Equation (1.1):

$$y_{zt} = \alpha + \sum_{j=-6}^{1} \beta_j D_{j(gj)} + \sigma_z + \delta_t + \varepsilon_{zt}, \tag{1.2}$$

where y_{zt} is our outcome of interest and $D_{j(gt)}$ is matrix of indicator variables which take a value of one if the first year of the moratorium is j years away for zip codes in moratorium group g in year t. We would expect to estimate no statistical difference between the treatment and control groups in the pre-policy event-time coefficients.

Secondly, recent methodological advances show that the TWFE model, as shown in equation (1.1), only yields consistent causal estimates of the average treatment effect on the treated when the treatment effects are homogeneous across groups (Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfoeuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). The main concern is due to the staggered timing of the policy across zip codes, the two-way fixed-effects model uses all possible combinations of treatment and control comparisons, leading to earlier treated groups being used as controls for later treated groups, resulting in inconsistent estimates for the average treatment effect on the treated if effects are heterogeneous across cohorts.

There are two reasons why we would expect heterogeneous treatment effects in our setting. First, a non-renewal moratorium is novel to the California insurance market, meaning insurers may adapt over time in how they respond to the policy and the Department of Insurance may also adapt in their enforcement role. Second, due to the record-breaking 2020 wildfire season, the 2021 moratorium covered more territory than in 2020, leading to a potentially different response from the insurers as a larger share of their business was under compulsory supply.

We take several steps to address these concerns. Our main event-study results make use of the estimator from Sun and Abraham (2021), which delivers consistent estimates in the presence of heterogeneous effects and differential timing of treatment. Secondly, we estimate equation (1.1) separately for the 2020 and 2021 treatment cohorts selecting only never-treated control units from each cohort's neighboring zip codes, limiting the "forbidden comparisons".

Lastly, as we use geographic borders to designate treatment, identification requires that the populations on either side of the border are homogeneous and that there is no selection into treatment or other feature of zip codes correlated with treatment. In order to limit the inherent differences between treatment and control groups, and to account for unobserved heterogeneity, we restrict the control group to zip codes that immediately border a treated zip code, but do not experience the moratorium during our sample. We believe these zip codes represent a plausible counterfactual due to geographic proximity. Importantly, as zip codes are not administrative boundaries, such as city or county borders, we would expect the unobserved heterogeneity to be smooth across the border. We also unaware of any increased funding or interventions implemented by jurisdictions in response to the fires that follow zip code designations or would apply to zip codes which were not impacted by the fire perimeter.

In Figure 1.7 we report the coefficients from cross-sectional baseline regressions using demographic and housing characteristics from the 2018 Census American Community Survey 5-year estimates at the zip code level. Results show that the treatment and adjacent zip codes are observably similar in the pre-treatment period. After controlling for observable factors, it is by random chance that these zip codes were not included in the moratorium boundaries because the location and size of disaster fires is as good as random each year.

FIGURE 1.7. Baseline Regressions

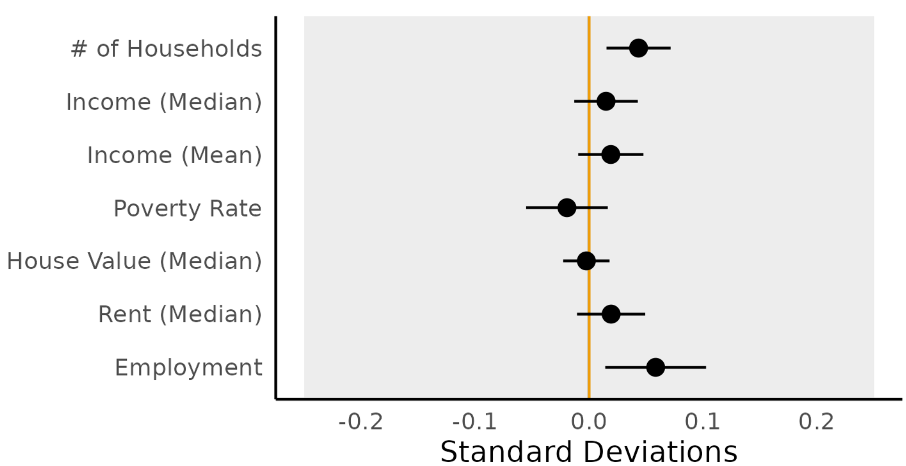

Notes: Plotted estimated coefficients and 95% confidence intervals are from cross sectional regressions at the zip code level of the respective outcome variable on an indicator for whether the zip code is treated during our sample. Observations are limited to our treatment and adjacent zip codes. Outcome variables have been standardized. Data is sourced from the 2018 Census American Community Survey's 5-year estimates.

It is plausible that for some areas the most geographically proximate zip codes follow different trends and that the most similar zip codes are actually located elsewhere in the state. For example, zip codes along the foothills of the central valley have very high fire risk, but often border nearby flat farm land which has near zero wildfire risk as measured in our data. In response, in addition to our main difference-in-difference specification, we refine the control group by using a nearest-neighbor matching approach, matching treated zip codes covered under the moratorium with zip codes from the 'rest-of-state' zip codes that do not border any treatment area, are never treated, nor experience a disaster fire during our sample period. Following the synthetic control literature, we match based on pre-treatment period trends in outcome variables as well as our time invariant measures of average and variance of wildfire risk (RPS) at the zip code level.

An additional benefit of using the matching approach is that by choosing zip codes that do not directly neighbor the moratorium zip codes, we are able to test whether there are spillover effects from the treatment areas to the neighboring zip codes. The imposition of the moratorium disrupts a firm's ability to balance the geographic concentration of their portfolio by forcing supply in the moratorium zip codes. This may lead to increased departure of firms from the closest zip codes to avoid being too highly concentrated, biasing the estimates from our difference-in-differences approach. Similar results between the sample using adjacent controls and our nearest-neighbor matched sample provides supporting evidence that there is no differential spillover between nearby and more distant zip codes.

1.6. Results

We begin our discussion of the impact of the non-renewal moratorium on insurance markets by estimating whether the regulation was in fact binding for firms. If the moratorium is binding, we would expect to see a sharp decrease in company-initiated non-renewals in treated zip codes compared to the control groups while the moratorium is in effect. We present event study results using the estimator from Sun and Abraham (2021) for the effect of the non-renewal moratorium on company-initiated non-renewals in Figure 1.8. Event time 0 represents the year the moratorium was in effect in the treatment zip code, while event time 1 represents the year after the moratorium has been lifted and the regulation is no longer in effect. Point estimates are shown along with 95% confidence intervals. Results from the specification using the adjacent neighboring zip codes as the control group are shown in panel (a). We find a large and statistically significant decrease in company-initiated non-renewals the year of treatment.

During the period of the moratorium, firms decrease their non-renewals by 15% compared to the control zip codes. Non-renewals do not completely disappear during the moratorium as firms are still able to non-renew policies for a variety of reasons, and only non-renewals that uniquely cite increased wildfire risk were restricted by the moratorium. The sharp decrease provides evidence that the regulation was binding and firms were not able to fully avoid the regulation.

However, the effect of the moratorium is short-lived as the decrease in non-renewals is quickly reversed the year after the moratorium is lifted. We estimate a large subsequent increase in non-renewals of 20% when compared to the adjacent control zip codes at event time 1. This is consistent with the narrative that firms not only simply delayed the non-renewal action to the following contract period, but accelerated their retreat from moratorium areas.

While the adjacent zip codes are located next to treatment zip codes covered by the moratorium, it is possible that these zip codes differ, not just in their levels, but also trends in insurance outcomes over time. Zip codes that burn may be systematically different than zip codes geographically close by, and therefore the best counterfactual setting may be located in other parts of the state. To address this, we also report estimates from a matched difference-in-differences model, where control zip codes are chosen through nearest-neighbor matching on average pre-treatment outcomes and zip code wildfire risk (RPS), shown in panel (b) of Figure 1.8. We restrict the pool of potential control zip codes to those that are not adjacent to treatment. We estimate a nearly identical decrease in non-renewals the year of the moratorium and a slightly larger post-policy increase the year the moratorium expires, with similar precision of estimates when compared to the specification using adjacent zip codes as control units.

Because the matched control group draws uniquely from non-adjacent zip codes, this specification allows for spatial spillover effects of treatment. Forcing firms to retain additional policies in treated areas that they would have otherwise non-renewed could lead firms to adjust their portfolio in adjacent zip codes in order to avoid being geographically concentrated in high risk areas. The similarity between the results using adjacent control zip codes and the matched control group suggests that spillovers are not a concern in our setting.

Figure 1.8. Effect of the Moratorium on Company-Initiated Non-Renewals

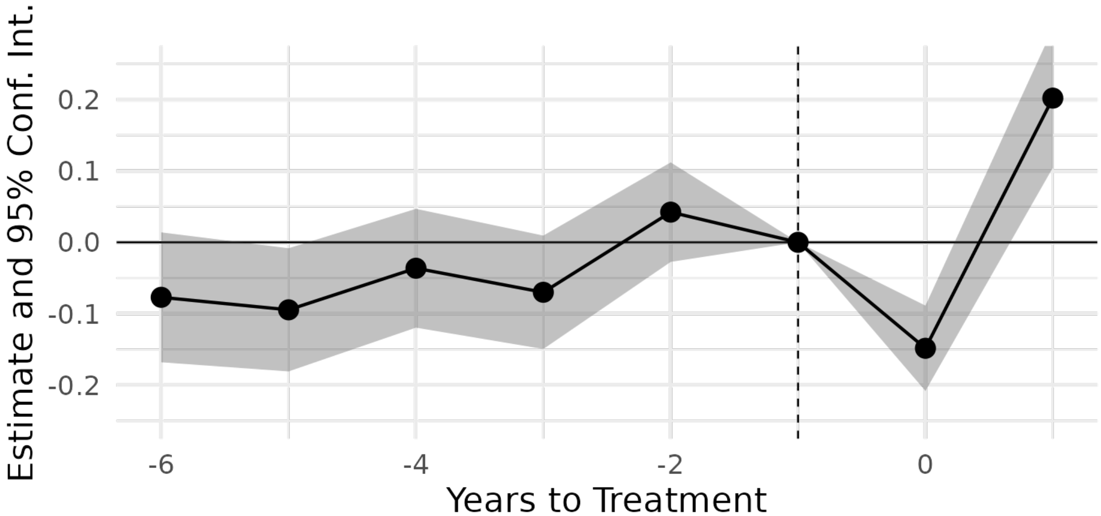

(a) Control group: Adjacent zip codes

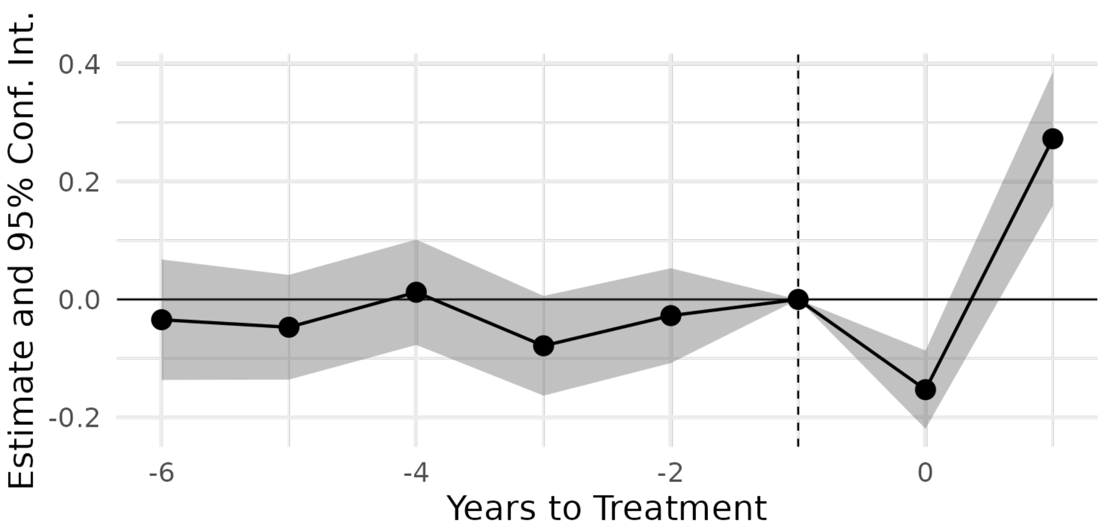

(b) Control group: NN-matched zip codes

Notes: Event study estimates for the effect of the California non-renewal moratorium on the log of company-initiated non renewals are shown with 95% confidence intervals. Results are separated by control group, with Panel (a) using adjacent non-impacted zip codes as controls and Panel (b) using the nearest neighbor matched control group from the rest of the state. Event time = 0 represents the year the moratorium was in place in the treatment zip codes.

Identification of causal effects in the event study framework relies on common pre-trends between treatment and control zip codes and that these trends would have continued in the absence of the policy. We can test the first part of this assumption using the pre-treatment estimated coefficients from the plots in Figure 1.8. For both control groups, pre-period coefficients include zero in the confidence-interval with no systematic trend, providing supporting evidence that the treatment and control groups were not on separate trends.

In Appendix Tables 1.11 and 1.12 we explore heterogeneity in the effect along the dimensions of wildfire risk and income. We estimate equation (1.2) separately by RPS and income quartile. Results show the effect was largely contained to the highest wildfire risk zip codes. However, there is little supporting evidence of heterogeneous effects by income quartile.

We now turn to the effect of the moratorium on non-renewals initiated by the customer. Figure 1.9 plots the estimated coefficients using both adjacent zip codes and the matched control zip codes in separate panels. Customer-initiated non-renewals were unaffected during the moratorium, but increased the year it was lifted. Coefficients on the pre-treatment years are precisely estimated and indistinguishable from zero, providing support that the common trends assumption holds in this setting. The lack of an effect the year of the moratorium provides further supporting evidence that the moratorium was binding for firms. While we cannot rule out that firms increased non-renewal activity for other reasons in response to the moratorium, this result shows that firms were not successful pushing away existing customers or manipulating their reports. If they had been able to do so, we would expect to see a positive coefficient, similar in magnitude, to the coefficient from the regression on company-initiated non-renewals in Figure 1.8 at event time 0.

We provide the following explanation for the post-moratorium positive coefficient. Once firms were able to non-renew policies again, this created a salient shopping trigger for customers. Non-renewals have to be delivered in writing to the policyholder at least 75 days in advance of the expiration date of the policy (California Code, Insurance Code – INS § 678.1). Thus, the process of the insurers initiating non-renewal action likely drove increased customer-initiated non-renewals as customers shopped for new policies ahead of their contract expiry. Thus, the true effect of the moratorium being lifted represents a combination of the customer and company-initiated non-renewals in event time 1. We can think of the addition of the two coefficients as the upper-bound estimate of the combined effect of the moratorium on firm driven non-renewal activity.

We next turn to the efficacy of the moratorium in slowing the transition of policies from being insured by the voluntary market to being covered by the FAIR plan. From the descriptive evidence presented in the data section, we know that the number of policies insured by the FAIR plan in

Figure 1.9. Effect of the Moratorium on Customer-Initiated Non-Renewals

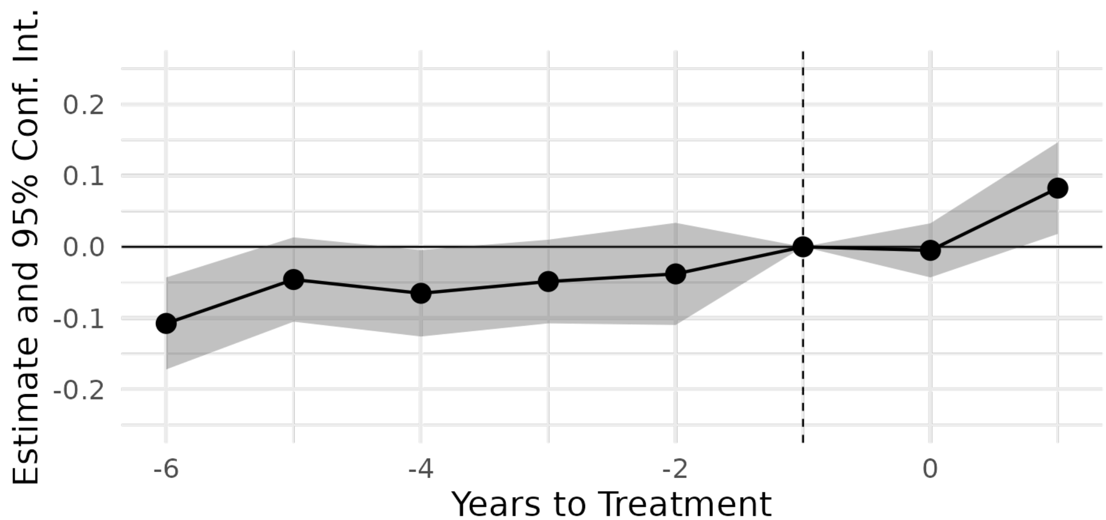

(a) Control group: Adjacent zip codes

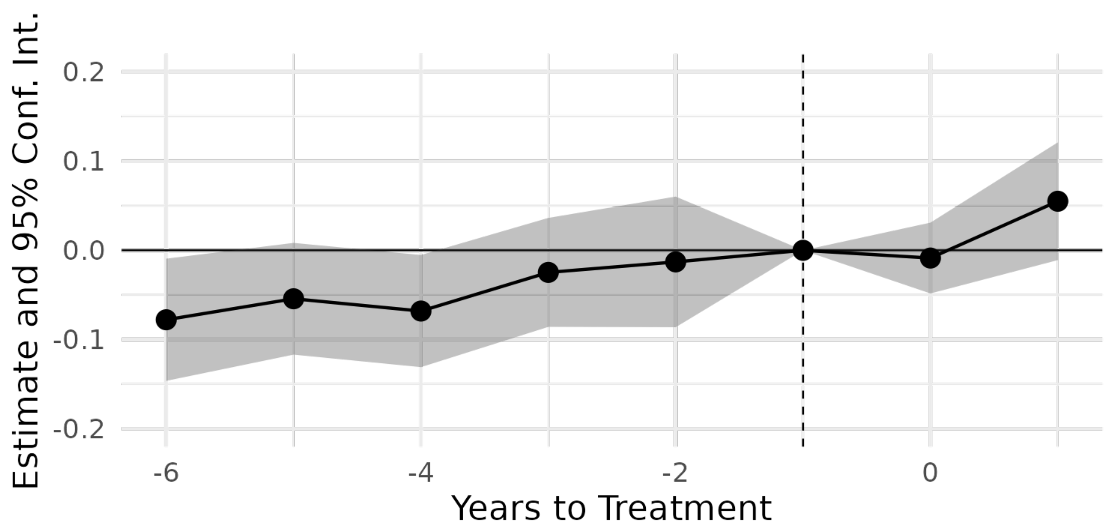

(b) Control group: NN-matched zip codes

Notes: Event study estimates for the effect of the California non-renewal moratorium on the log of customer-initiated non renewals are shown with 95% confidence intervals. Results are separated by control group, with Panel (a) using adjacent non-impacted zip codes as controls and Panel (b) using the nearest neighbor matched control group from the rest of the state. Event time = 0 represents the year the moratorium was in place in the treatment zip codes.

both treated and control areas is increasing during the period leading up to the moratorium as firms had already begun to limit their exposure in high risk areas. When using our econometric approach with FAIR plan policies as the outcome variable (shown in Figure 1.10), we do not detect any discernible impact of the moratorium on the market share of the FAIR Plan.

Figure 1.10. Effect of the Moratorium on FAIR Plan Policies

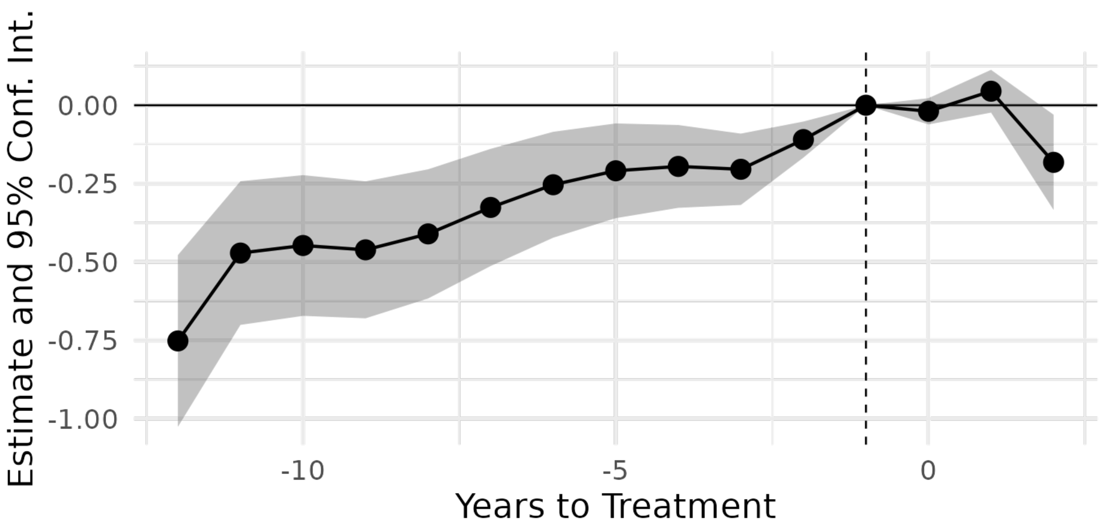

(a) Control group: Adjacent zip codes

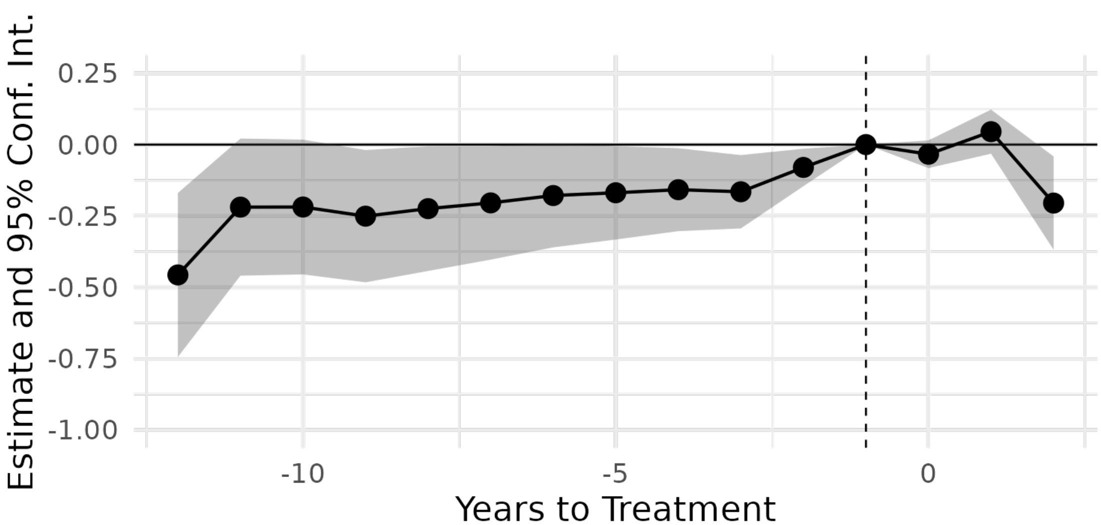

(b) Control group: NN-matched zip codes

Notes: Event study estimates for the effect of the California non-renewal moratorium on the log of FAIR plan market share are shown with 95% confidence intervals. Results are separated by control group, with Panel (a) using adjacent non-impacted zip codes as controls and Panel (b) using the nearest neighbor matched control group from the rest of the state. Event time = 0 represents the year the moratorium was in place in the treatment zip codes.

Results using the adjacent control zip codes exhibit significant diverging pre-trends, which violate the assumption needed for causal interpretation of the regression coefficients. The results are consistent with the story that FAIR plan market share was increasing in the treated zip codes faster than in control zip codes prior to the moratorium. Yet, upon employing matched control units, we notice considerably milder pre-trends between treatment and control units. Nevertheless, we continue to estimate a precise null effect of the moratorium.

The moratorium eliminated one channel through which firms could reduce their exposure in high-risk areas. However, our results show that customers were rejected from the voluntary market and found insurance in the FAIR Plan at similar rates in both areas covered by the moratorium and control areas. This provides evidence that the policy was an ineffective tool at slowing the retreat of insurers from high risk areas, despite providing temporary reprieve for select customers. Firms were still able to reduce their exposure through not writing new policies for current residents who had their policy cancelled by another firm for non-protected reasons, or homeowners new to the zip code who were also not protected by the moratorium.

1.7. Conclusion

As climate change increases the risk of large scale natural disasters, well-functioning insurance markets will be necessary for consumers who rely on them often as the sole method of risk transfer. Our paper highlights how regulation and market structures that has traditionally been designed to benefit consumers by suppressing price levels can have large distortionary effects and lead to the unraveling of the market as prices and risk diverge. The California non-renewal moratorium is a unique policy tool the government implemented in an attempt to maintain a stable supply of homeowners insurance in the face of rapidly increasing wildfire risk. The moratoriums were effective in achieving this goal, but only in the short-term, and the strong rebound effect suggests that this policy is not an effective long term solution to correct market failures.

There is still a need for a permanent solution to this problem, which should include measures to reduce wildfire risk faced by households, both by adapting to wildfire risk and discouraging migration to high risk areas exacerbated by artificially low homeowner's rates. Regulating the industry in a way that allows firms to react to increased risk and earn reasonable profits can reduces the incentive for firms to retreat from the market and results in a functioning private market, a stated goal of the Department of Insurance.

1.8. Appendix

Figure 1.11. Effect of Moratorium on Company-Initiated Nonrenewals by RPS Quartile

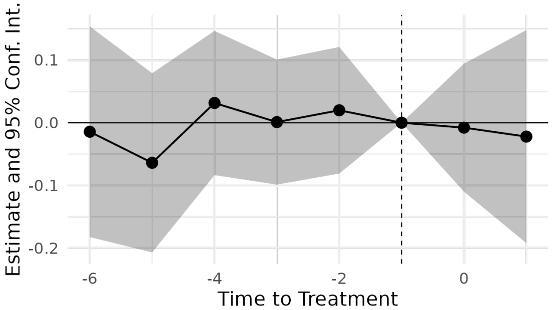

(a) 1st

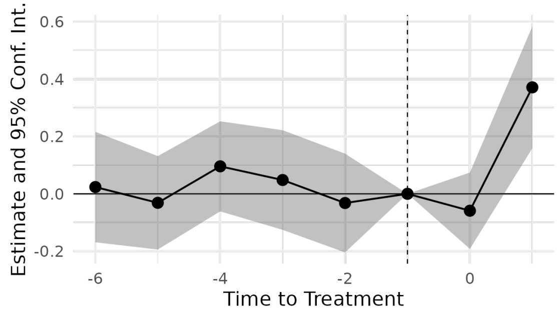

(b) 2nd

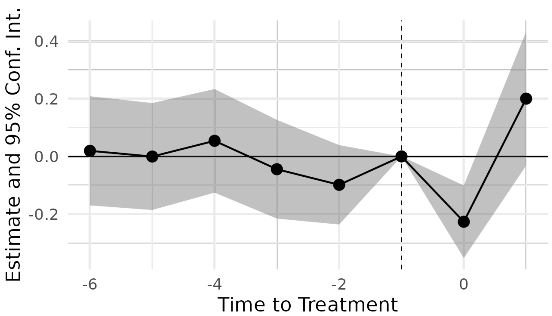

(c) 3rd

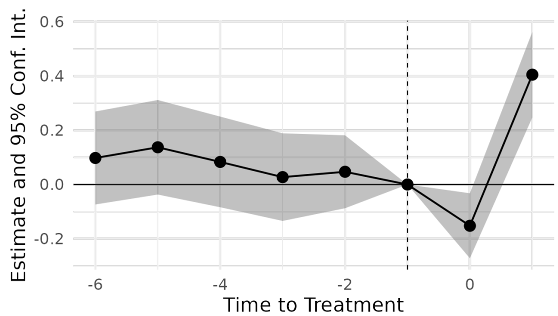

(d) 4th

Notes: Event study estimates for the effect of the California non-renewal moratorium on company-initiated non renewals are shown with 95% confidence intervals. Results are separated by RPS quartile. Event time = 0 represents the year the moratorium was in place in the treatment zip codes.

Figure 1.12. Effect of Moratorium on Company-Initiated Nonrenewals by Income Quartile

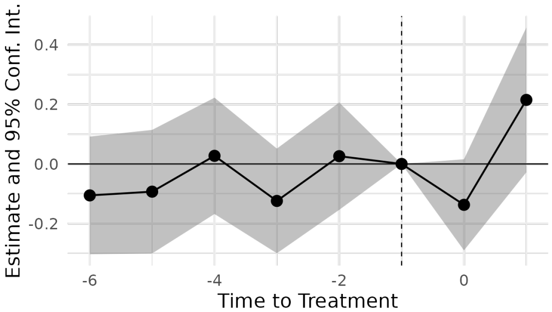

(a) 1st

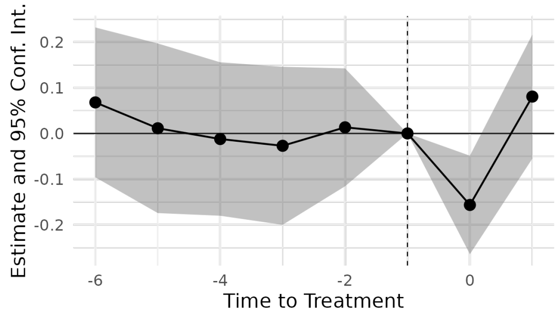

(b) 2nd

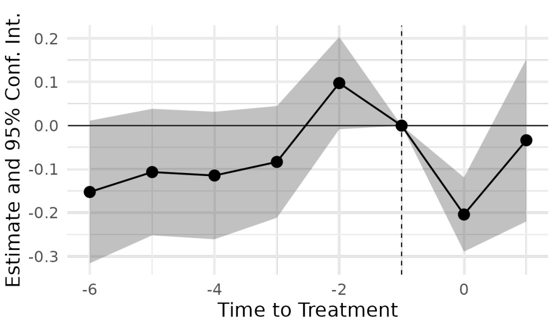

(c) 3rd

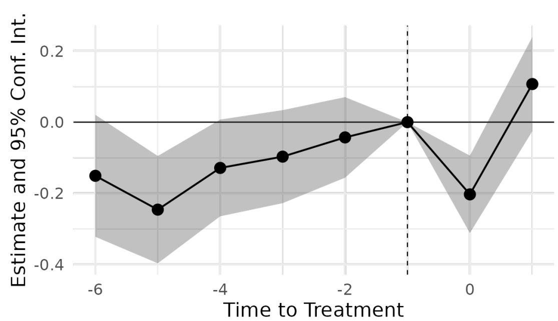

(d) 4th

Notes: Event study estimates for the effect of the California non-renewal moratorium on company-initiated non renewals are shown with 95% confidence intervals. Results are separated by income quartile. Event time = 0 represents the year the moratorium was in place in the treatment zip codes.

CHAPTER 2

The Impacts of Entry and Market Power in the California Retail Fuel Industry

2.1. Introduction

Understanding market power in the California gasoline market is of interest to policy makers and has been the subject of review by state agencies and the California Department of Justice for the past two decades due to prices that are consistently higher than the rest of the country.[1] Entry costs for new stations are high due to land values, environmental regulation, and local zoning laws. California blend rules result in an isolated market supplied nearly exclusively by the few in-state refineries. Further, these same refiners can have vertical relationships with downstream retail stores either through direct ownership of convenience stores or franchise contracts. Taken together, these factors suggest the potential for market power in oligopolistic local retail gasoline markets.

In this paper, I study competition in the California retail gasoline market by estimating the short-run effect of market size on incumbent firm pricing. Theory is ambiguous on the effect of market size and entry on prices (Barron et al., 2004). Empirically estimating the impact of market size and firm entry on prices has traditionally been a challenge due to the endogeneity of market structure. Profit maximizing firms are attracted to markets with higher prices and profit margins. This selection bias results in cross-sectional studies yielding biased estimates for the causal effect of market size on price. Estimation has been further constrained by the availability of sufficiently rich data that contain numerous exogenous changes in market structure and granular price data (Haucap et al., 2017). In light of price transparency regulations that have generated administrative data sets and the increased availability of third-party high-frequency data, modern panel data methods can be used to explicitly control for the inherent market structure.

To calculate the reduced-form causal effect of market size and the entry of nearby competitors, I use a panel of daily station-level prices and the precise geographic location for the universe

[1] https://oag.ca.gov/antitrust/gasoline

of California gas stations from 2014-2018 with geographic and temporal variation in exposure to changes in the number of nearby competitors. The resulting data set includes over 700 new station entry and station exit events and 35 million price observations. Using both difference-in-differences and event-study designs, I compare the difference in prices at incumbent stations before and after a station enters or exits the market, with unaffected stations serving as the control group. I am able to control for the endogenous location decision of entering and exiting firms by including station, city-specific linear time trends, and day-of-sample fixed effects. The coefficient is thus identified by within-station variation in the number of nearby competitors and relies on the assumption that the exact timing of entry is exogenous conditional on these market characteristics.

I find that increased market competition reduces prices, but that the effect of changes in the number of competitors is heterogeneous. Entry of a new gas station nearby is associated with a two-cent reduction in gas prices at incumbent stations that is immediate upon the timing of entry. The amount, while small in absolute terms, represents an average of 5% of station markups on gasoline during the sample period. Conversely, I estimate precise null effects of station exit on nearby incumbent station pricing. Event study results provide support for the key identifying assumption for causal interpretation of the estimates, markets with and without entry or exit were following observably parallel trends in the periods leading up to the respective event.

Focusing on the entry effect, I show that entry leads to a lower price equilibrium that is highly persistent in the market, lasting for years after the event, and is present across all three major blends of gasoline sold. Additionally, I show that the estimated effects are not attenuated by entry or exit of other stations nearby in the periods preceding, following, or simultaneous with entry and exit. Estimated effects are highly localized and decay as the market definition broadens, consistent with prior literature on the tight spatial nature of retail gasoline competition. There is no evidence that firms engaged in deterrence behavior or limit pricing in the periods leading up to entry, nor evidence of predatory pricing by incumbent firms leading to firm exit.

Next, I present results focused on market power in the retail gasoline markets. Recent work has shown that gas prices in California rose sharply in response to the Torrance refinery fire and the ensuing supply shock in February, 2015 (Borenstein, 2023). In the years since, an unexplained divergence between prices in California and the rest of the nation has emerged, averaging around 40 cents, resulting in $48 billion dollars of additional expenses for Californians as of 2023. This gap cannot be explained by differences in taxes or input prices and is concentrated after the point of wholesale distribution. I provide evidence that suggests the supply shock had little effect on competition among retail stores and that market power of retail stations may not have been affected by the event. I show that the estimated magnitude of entry effects for new stations was unaffected by the supply shock.

This work contributes to two strands of literature. First, this paper contributes to the growing literature on the effects of market size and entry in retail markets. Following Arcidiacono et al. (2020), similar studies using panel data methods emerged studying entry effects in gasoline markets focused on international markets where administrative price data was available. I add to this literature by presenting the first causal estimates of the effect market size and station entry on pricing in California. Results are similar in magnitude to other studies conducted in Mexico (Davis et al., 2023) and Germany (Fischer et al., 2023) also making use of high frequency, daily station data. Bernardo (2018) makes use of a liberalization of entry restrictions in Spain to study the impact of new entry on prices while González and Moral (2023) use an IV approach in the same setting.

Additionally, I contribute to the subset of the literature studying the unique California gasoline market. Stations in California, by regulation, must sell a special blend of gasoline (CARBOB), resulting in a distinct wholesale market from surrounding states. This has generated substantial literature documenting market power and pricing in the California market (Borenstein et al., 2004), and vertical mergers (Hastings, 2004; Taylor et al., 2010).

Understanding the effect of market structure in gasoline markets is also of importance to the conversation on the energy transition. As society moves towards reducing its reliance on fossil fuels, we need to understand how markets will be impacted. Supply-side restrictions on capital, either through reduced access to financial markets or through outright bans on fossil fuel extraction, transportation, and delivery have the effect of improving the position of legacy infrastructure. This is present in the California gasoline markets as several jurisdictions, starting with Petaluma in 2021, have banned the construction of new gas stations and restricted the expansion of existing stations. While gasoline demand is expected to decrease over time as the market penetration of electric vehicles increases, millions of gasoline-powered cars will still be operating on California's roads in the years to come and new gas stations are still in demand as shown by data. Banning entry of new stations results in counter-factually higher prices, increasing the profits of incumbent firms.

The paper proceeds as follows: Section 2 presents the data used in the empirical study, Section 3 discusses the empirical strategy, Section 4 presents the estimation results, as well as a series of robustness checks, Section 5 discusses the role of market power and Section 6 concludes.

2.2. Data

To calculate the effect of market size on incumbent pricing, I obtain a panel of daily retail station gasoline prices in California from the Oil Price Information Service (OPIS) for years 2014 through 2018. I observe daily prices by fuel blend with the precise geographic coordinates for each station and a unique site identifier. Importantly, a station is identified in the OPIS data by its geographic location. Thus, the unique site identifier remains unchanged when a store undergoes renovations, changes ownership, or changes store and fuel branding. This is important for the analysis as entry and exit will not be confounded with changes that do not alter the number of competitors in the market. This however comes at the expense of being able to account for owner specific characteristics. For convenience, in the rest of the paper I use the terminology "station" to refer to the unique site where a gas station is located. To account for erroneous data, I exclude stations with less than 14 total price observations over the 4-year sample period as well as the 4 stations without geographic information. This results in a panel of 9,539 stations with over 35 million price observations.

FIGURE 2.1. Observations by Station: 2014-2018

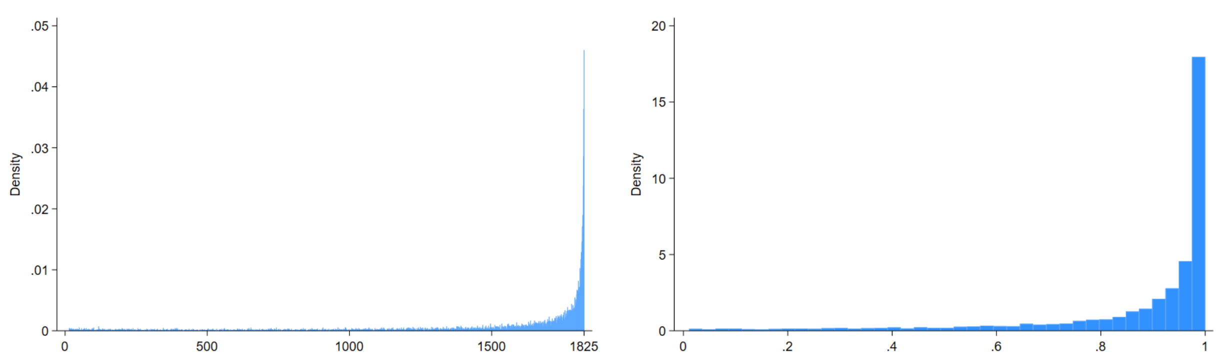

(a) Observation Count (b) % of Potential Observations

Notes: Data sourced from Oil Price Information Service (OPIS).

Prices in OPIS are collected through monitored fleet credit card transactions, customer reports, and direct feeds from stations. Only one price is reported per station per day per fuel type, however a price is not observed every day for each station or each fuel blend. OPIS reports prices for blends labeled as regular, mid-grade, and premium with the terms directly mapping to 87, 89, and 91 octanes. Separate from any convenience store branding for the station, the fuel dispensed can be branded or unbranded from the wholesale point. Branded fuels are associated with a company name and often include additives that are blended after the wholesale point. Both the store branding and the fuel branding are observed in the data.

Panel (a) of Figure 2.1 shows the distribution of observation counts by station for the sample period 2014-2018 (1,825 days) for regular gasoline. Data coverage is high, with the median store having 1,756 prices reported. In fact, 75% of stores report more than 1,500 price observations over the sample period. To account for differential start and stop dates by station during the sample period, Panel (b) of Figure 2.1 reports the number of price observations as a percent of potential reporting days for each station using as a denominator, the number of days spanned between the first and last observed price observations for the station. The median store has prices reported for 97% of its respective sample period, and 75% of stores have prices reported for more than 86% of potential days.

Figure 2.2. Month of First and Last Observations by Station: 2014-2018

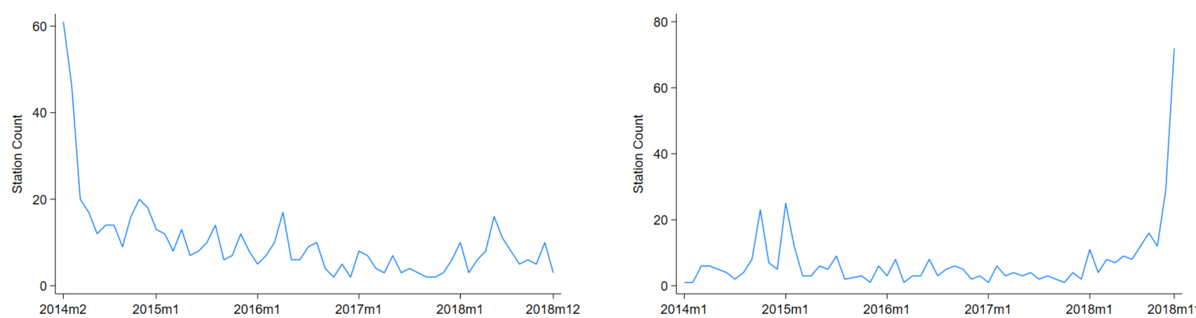

(a) First Observation

(b) Last Observation

Notes: Data sourced from Oil Price Information Service (OPIS)

The OPIS data forms the basis for the empirical analysis on entry and exit. For each station, I calculate the entry date as the date of the earliest price observation and exit as the date of the last observed price observation. Figure 2.2 reports firms by the date of their first and last price observation, grouped by months (January 2014 and December 2018 are excluded from the graph since the overwhelming amount of first and last observations fall in these two months). There are a disproportionate number of entries and exits in the extreme tails of the sample period. This is likely driven by the fact that stations do not have prices reported every day and could have price observations just before or just after the sample period. Considering these stations as entries and exits would result in biased estimate. Therefore, I consider any station that has an initial price observation in January-March 2014 not to be an entry and any station to have their final price observation in November-December 2018 not to be an exit, and assume that they were present through the beginning or ending of the sample period. Stations that go dormant but have subsequent prices during the sample period are not considered as station exits or subsequent entry. This results in 484 unique stations having an entry date during the sample period and 348 unique station exits.

FIGURE 2.3. Market Characteristics by Station: 2014-2018

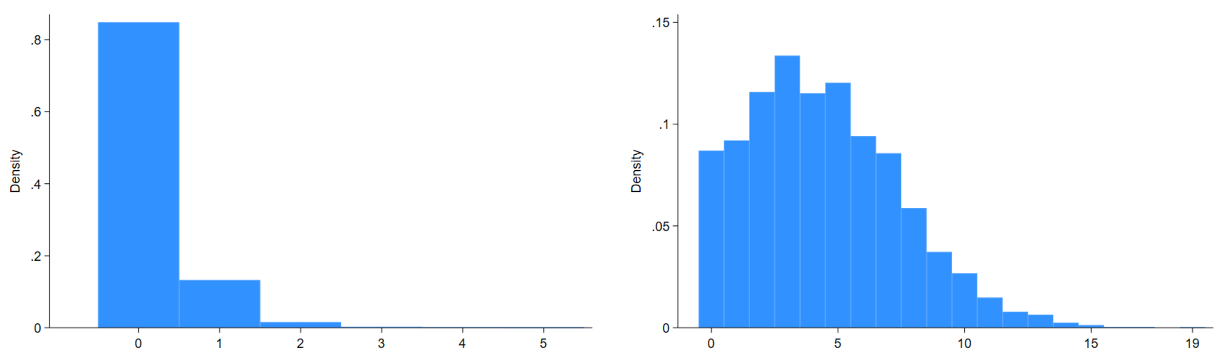

(a) Number of Entrants (b) Initial Market Size

Notes: Data sourced from Oil Price Information Service (OPIS)

Using the entry date, exit date, and geographic location of each station, I define the relevant market as the number of competitor stations operating within a 1 mile distance at the station-date level.[2] Panel (a) of Figure 2.3 shows the number of entry events within 1 mile experienced by incumbent stations. 85% of stations do not experience market entry during the sample period. Multiple entry is rare as 90% of stations that do experience entry only have 1 entrant during the sample period. As such, in the subsequent event study analysis, I focus on the first entry event observed during the sample period. Panel (b) of Figure 2.3 shows the distribution of initial market

[2]Results for other distances measures varying from .5 to 10 miles are included as a robustness check.

sizes for the full sample of stations. Markets contain relatively few stations, with the median store having four other gas stations within a 1-mile radius. There are stations which operate in monopoly markets and one station in downtown Los Angeles with 19 competitor stations within a 1 mi. radius. 10% of stations are located in highly concentrated markets with 8+ stations. Appendix Figure 2.10 reports the number of entrants experienced by incumbent stations by the initial market size. The majority of entries occur in markets with 3-8 competitors and are the only entrant during the sample period. Figure 2.4 below gives a visual representation of the variable construction, with gasoline stations in Sonoma county plotted in black and each station's 1-mile market shown in red.

FIGURE 2.4. Map of Gasoline Stations in Sonoma County: 2014-2018

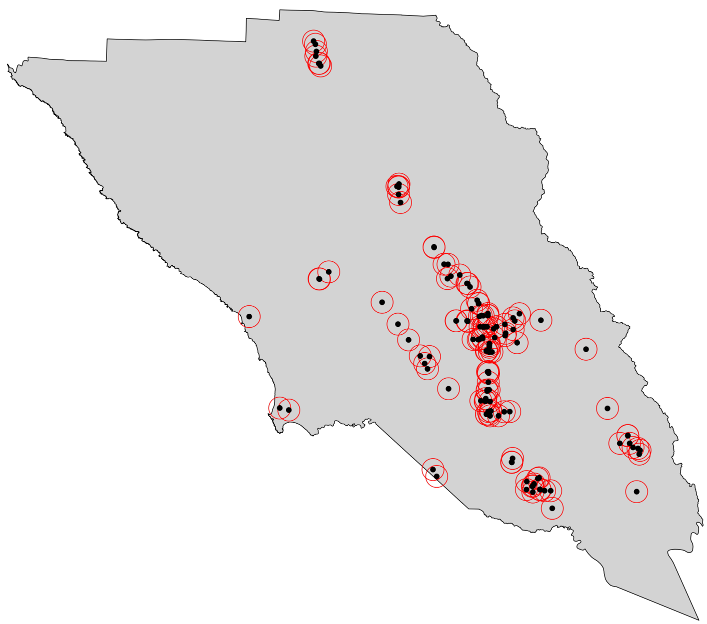

Notes: Gas stations located in Sonoma county, California are shown in black with the 1 mi. market definition shown in red. Data sourced from Oil Price Information Service (OPIS)

There does not exist an official public source of gas station operation dates for the state of California and, to the best of our knowledge, the OPIS data represents the best source of information for the research presented. Collecting business permit data or underground tank information can locate stations but lacks the pertinent temporal component of on-site business operations which is likely the relevant metric of entry and exit for competing firms. State tax data on gas stations is reported at the owner later, and thus does not contain store-level information across multiple stores under the same ownership.

Figure 2.5. Map of Gasoline Station Entrants: 2014-2018

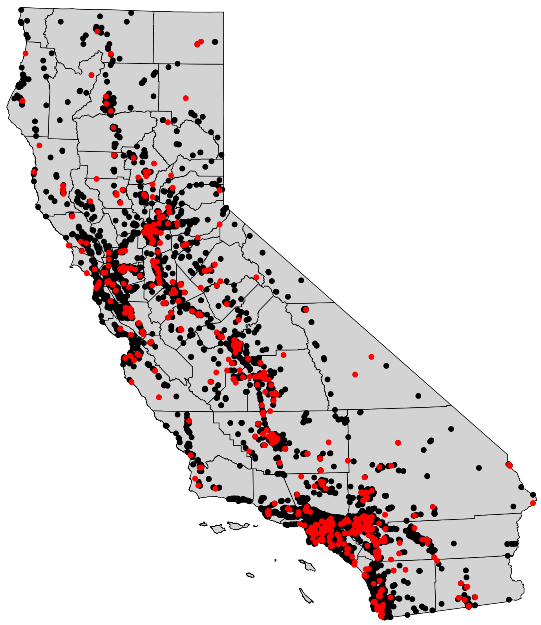

Notes: Entrant stations are shown by red dots, with incumbent stations represented by black dots. Data sourced from the Oil Price Information Service (OPIS).

There exists the potential for false positives in the OPIS entry and exit variable construction due to the limited window of the sample period and the fact that observations are when prices are reported, not official open or closing dates for the store. There also is the chance that a station goes unreported to OPIS. I validate the OPIS data by bench-marking the number of stations and change in stations over the sample period against two other measures for California: the Census Bureau's County Business Patterns data and the estimated number of stations calculated by the California Energy Commission. The County Business Patterns data reports the number of businesses as of the week of March 12th of the appropriate year by NAICS code. The data show a net increase of 228 gas stations in California from 2014-2019. The CEC also undertakes an effort to estimate the number of stations in California based on returns from the A15 survey and other government data sources. For the same period, the CEC estimates a net increase of 190 stations in the state and around 10,000 gas stations, thus the OPIS data is similar in magnitude and change. Neither source differentiates between entries or exits.

Figure 2.5 shows the location of the entrant gas stations in red and incumbent gas stations in black. Entry occurs in nearly all parts of the state, with concentration in the major urban areas and along the I-5 and CA-99 corridors through the central valley. Visually, entry appears to largely match the locations of existing stations and not concentrated in new markets. There is observed entry in remote and rural parts of the state, however these events will only contribute to identification of the entry parameter in the empirical analysis if they are located sufficiently close to an incumbent station.

The main concern with previous cross-sectional analyses is that entry and exit are likely to occur in locations that differ from markets that do not observe a structural change. To test for baseline differences along observable demographic characteristics, I use demographic data from the 2014 American Community Survey's 5-year estimates. I compare census tracts that experience at least one entry or exit event to census tracts which have gas stations but experience no entry/exit during the sample period using the following cross sectional regression specification:

$$Char_t = \alpha + \beta Event_t + \epsilon_t \tag{2.1}$$

which regresses the various demographics variables for tract t on a binary indicator variable for entry or exit in separate regressions. Regressions are weighted by tract population.

Table 2.1 reports the coefficients from estimation of equation (2.1). Census tracts that experience entry and exit on average have lower household income, higher poverty rates, lower housing values, and more households than tracts with stations, but no events during the sample period. Tracts with entry are less likely to have zero-vehicle households, more likely to commute via vehicle, but have similar commuting times. Exit occurs in more populated tracts with lower income and housing values, more zero-vehicle households, and shorter commute times.

The baseline demographic differences between locations that do and do not observe market composition changes highlight the need to account for the inherent market characteristics to address the endogenous market change and continuing operation decisions of stations. Additionally, to the extent that there are unobservable demographic characteristics that are correlated with both station entry and demand for gasoline, cross-sectional regressions of price on the number of competitors are likely to yield biased estimates of the effect.

Table 2.1. Baseline Demographic Differences By Census Tract

	Entry	Exit
Income (Median)	-7,664.6***	-7,525.8***
	(1,359.8)	(1,644.3)
Income (Mean)	-11,109.9***	-7,702.9***
	(1,757.9)	(2,129.3)
Poverty Rate	2.874***	3.065***
	(0.548)	(0.664)
Households	325.9***	121.0**
	(45.8)	(55.6)
House Value (Median)	-111,373.0***	-50,330.9***
	(10,628.6)	(12,794.4)
% No Vehicle	-0.979***	0.911**
	(0.323)	(0.391)
% Commuting by Vehicle	1.290***	-0.148
	(0.485)	(0.587)
Commute Time	0.122	-1.660***
	(0.281)	(0.339)

Notes: Coefficients from cross-sectional regressions are shown, with standard errors below in parentheses. Tract level demographic variables are regressed on an indicator for whether the census tract experienced entry, in column 1, or exit, in column 2. Regressions are weighted by tract population. Data are sourced from the 2014 American Community Survey 5-year estimates.

2.3. Empirical Strategy

The localized nature of gasoline station competition allows for the geographic and temporal variation in exposure to station entry and exit across firms to form the basis of a difference-in-differences estimation for the causal effect on incumbent pricing. Following prior work by Arcidiacono et al. (2020), I treat the exact timing of the entry or exit of a new gasoline station as a short-run exogenous shift in the market structure for incumbent firms after conditioning on the inherent market structure. In the preferred specifications, I define the relevant market as the 1 mile radius circle around the incumbent station. This market size is based on prior literature which ranges between 1 and 2 miles (Barron et al., 2004; Bernardo, 2018; Carranza et al., 2015; Davis et al., 2023; Fischer et al., 2023; Hastings, 2004; Lewis, 2015).

Importantly, I condition the model on a rich panel of fixed effects to account for unobserved variable bias inherent to the endogenous location decision of entering and exiting firms. By restricting the model to identification from within station variation in the number of nearby competitors, over time, the model accounts for factors important to the locating decision such as the overall price level in the market, local price elasticity of demand, local traffic patterns, and relevant customer characteristics.

The estimating equation is:

$$P_{st} = \alpha + \beta N_{st} + \sigma_s + \delta_t + \Phi_c(t) + \varepsilon_{st} \tag{2.2}$$

where the main outcome variable is the retail price at station s on day t and N_{st} is the count of stations within a 1-mile radius of station s on day t, increasing upon entry and decreasing with nearby exit. Separate regressions are run for the various blends of gasoline.

Station fixed effects (σ_s) are included to capture time-invariant differences between locations, such as station amenities and size, location effects, and distance to the wholesale terminal which largely drives differences in input costs. Day-of-sample fixed effects (δ_t) capture state-wide daily shocks to both input costs, such as oil-prices and refinery supply shocks, as well as common daily shocks to product demand. Lastly, in the preferred specifications, I include a city-specific linear time trend, $\Phi_c(t)$, which accounts for changes over times which may be correlated with price and station demand, flexibly across markets. Since the market definition varies from store to store, I cluster standard errors at the city level to account for common shocks across units.

The use of unit and time fixed effects as in the above equation is known as the two-way fixed effects (TWFE) estimator. In the setting, stations that do not experience a change to market size within 1 mile during the sample period and previously treated stations both serve as control units for stations that experience entry or exit. Recent literature has shown that in the presence of heterogeneous treatment effects across treatment timing, the use of previously treated units as control units can lead to biased estimates. I return to this point later in Section 4.3 and show that the results are largely unchanged in the robustness checks where I add increased flexibility to the TWFE equation as described in Wooldridge (2021) to account for potential heterogeneity. I report the main results using the TWFE specification from equation 2.2 for two main reasons: The highly localized nature of gas station competition makes homogeneous effects across treatment timing likely and the proposed solution estimators are computationally infeasible for the sample size with daily variation and requires significant aggregation.

Identification of the main coefficient of interest, β, as the causal effect of a change in the number of nearby stations on prices requires three main assumptions. First, the main identifying assumption requires entry and exit to be conditionally uncorrelated with the error term. Specifically, conditional on station fixed effects, time fixed effects, and time trends, the changes in station count are exogenous.

$$E[\varepsilon_s|\sigma_s, \delta_t, \Phi_c(t), N_{st}] = 0 \tag{2.3}$$

This also requires that there were no other factors perfectly correlated with the timing of the station entry that also impacted the pricing of nearby stations.

Secondly, the SUTVA assumption requires that there is no spillover of treatment onto control units outside of the impacted market. This is plausibly satisfied in this setting due to the local geographic nature of gas station competition, eliminating spillover price effects from treatment units to the larger pool of control observations. Lastly, identification in the differences-in-differences framework relies on the parallel trends assumptions. This requires that treatment and control units to be following common trends in the pre-treatment period and would have continued to do so in the absence of treatment.

In addition to the difference-in-differences model specified in equation 2.2, I estimate the following event study model for station entries, and the equivalent analog for exit events separately:

$$P_{st} = \alpha + \sum_{k=-24}^{24} \beta_k \mathbb{1}[Entry_{st} = k] + \sigma_s + \delta_t + \Phi_c(t) + \varepsilon_{st} \tag{2.4}$$

setting event time indicators for the number of months before and after the first nearby entry or exit observed at station s. End points are binned to include 24 or more months before/after the event.

This additional specification has two attractive properties. First, it provides support for the parallel trends assumption needed for causal estimates. The concern being that if markets that observe entry or exit were not just different in terms of levels, but also changing differently over time, and these changes were correlated with the endogenous entry decision, this would introduce bias into the estimates. Precisely estimated null effects in the time periods leading up to entry or exit event provide support that treatment and control markets were following common trends.

Secondly, I am able to decompose the effect of changes in the station count variable used in equation (2.2) by estimating separate regressions for entry and exit events.

The model specified above yields estimates of the average treatment on the treated (ATT) when the identification assumptions are satisfied. Given the baseline differences in treatment and control areas shown in Table 2.1, the estimated coefficient represents an internally valid estimate of the casual effect of entry or exit at locations where firms decide to enter or exit. The effect can be generalized to answer the question,"What will be the average effect of a station's planned entry or exit?"

2.4. Results

2.4.1. Station Count Results. Results for the impact of changes in the station count within 1 mile on incumbent prices are shown in Table 2.2. Column 1 contains results from the simple linear regression of price on the station count variable as an initial point of comparison. There is an economically small, but significant negative relationship between market size and prices. The addition of a station within a 1 mi. radius lowers prices at incumbent firms by .6 cents. The specification fails to account for the endogenous relationship between price and entry, leading the estimate to likely be biased.

Table 2.2. Effect of Station Entry Within One Mile on Incumbent Pricing

Dependent Variable	(1)	(2)	(3)	(4)	(5)
Price	-0.006***	-0.006***	-0.067***	-0.012***	-0.010***
	(0.002)	(0.002)	(0.021)	(0.003)	(0.003)
R Sq.	0.001	0.809	0.163	0.956	0.958
Obs.	14,745,166	14,745,166	14,745,166	14,745,166	14,745,166
Log Price	-0.002***	-0.002**	-0.021***	-0.004***	-0.004***
	(0.001)	(0.001)	(0.006)	(0.001)	(0.001)
R Sq.	0.001	0.802	0.165	0.951	0.953
Obs.	14,745,166	14,745,166	14,745,166	14,745,166	14,745,166
Station FE	No	No	Yes	Yes	Yes
Day of Sample FE	No	Yes	No	Yes	Yes
Linear Time Trend	No	No	No	No	Yes

Notes: OLS estimates for the effect of a change in the count of competitor stations within 1-mile on incumbent firm pricing for regular unleaded gasoline are shown. Column 1 reports estimates from a linear regression of the outcome on station count. Columns 2-4 report results using the same model, but with the addition of the fixed effects listed in the panel below. Row 1 presents results where price is the outcome variable while row 2 presents results using the log transformation of price as the outcome variable. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018. Model standard errors are reported in parentheses and clustered at the city level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Columns 2-5 report results with added layers of fixed effects. In Column 2, the addition of a day-of-sample fixed effects to the naive regression accounts for daily common shocks to input costs and demand such as crude oil prices or major weather events. While this accounts for a sizable portion of the variation in prices, as evidenced by the substantial increase in r-squared from Columns 1 to 2, the effect of entry largely remains unchanged. Column 3 accounts for inherent market characteristics that are endogenous to the firm's entry and exit decisions by restricting identification to within-station variation with the inclusion of the station fixed effects. An additional nearby competitor is now associated with a significant 6.7 cent decrease in prices. Columns 4 is the canonical two-way fixed effects model including both day-of-sample fixed effects and station fixed effects. Column 5 represents the preferred specification, showing results are robust to the inclusion of a city-specific linear time trend. An additional competitor within 1 mi. results in a 1 cent reduction in incumbent prices, which represents around a 2.5% reduction in firm markups during the sample period which average 40 cents.[3]

The evolution of coefficients across the table highlights the importance of addressing the endogenous entry decision of firms. Without the inclusion of station fixed effects, estimates rely on cross-sectional variation in station counts. Once controlling for the inherent market characteristics with store fixed effect, chiefly whether entry occurs in a high or low-price market, the entry effect increases and is statistically distinguishable from zero. The change in the results with the inclusion of time effects highlights the need to account for seasonality in prices and entry timing as well as demand shocks. Lastly, a concern in this setting is that markets that are selected for entry or exit vary not only in terms of price levels, but are on separate growth trends, and that these trends are correlated with gasoline demand. The robustness of the coefficient to the inclusion of a city-specific linear time trend between Columns 4 and 5 provides support that treatment and control markets also do not vary in terms of trends.

2.4.2. Event Study Results. We now turn to results from the event study specification shown in equation 2.4. Figure 2.6 presents results for the first entry event experienced by a given station, and Figure 2.7 reports results for the first exit event experienced during the sample period. An event study design is useful in this setting for multiple reasons. First, causal identification in a difference-in-differences model specification relies on the parallel trends assumption which requires

[3]California Energy Commission estimates of CA gasoline price breakdown and margins.

treatment and control groups to evolve along common trends in the pre-treatment period and would have continued in the absence of treatment. The first part of this assumption can be tested visually and statistically via the event study design.

FIGURE 2.6. Effect of Station Entry on Incumbent Pricing

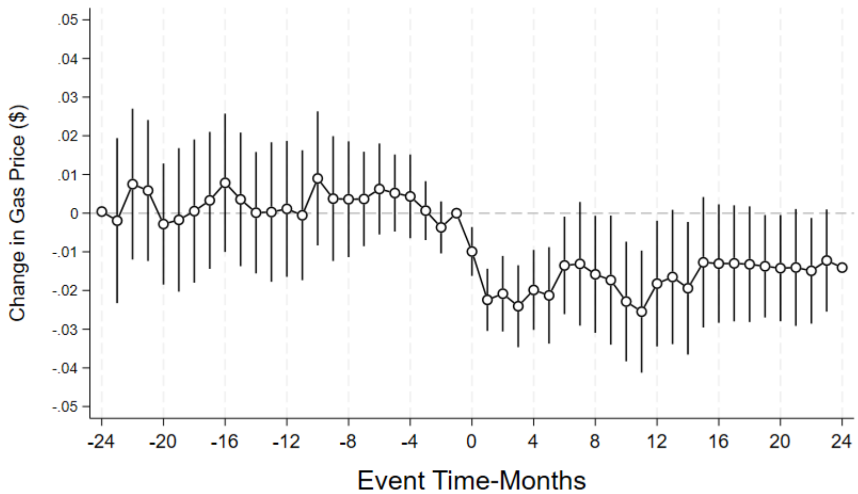

Notes: Coefficient estimates for the effect of station entry within 1 mile of an incumbent station on incumbent pricing of regular unleaded gasoline are shown. Event time t=-1 is the month prior to the arrival of the entrant. Coefficients are estimated from a linear regression of price on a panel of event time dummy variables, station fixed effects, day-of-sample fixed effects, and city-specific linear time trends. Standard errors are clustered at the city level. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018.

Point estimates for the event-time coefficients for the months prior to both entry and exit events are close to zero and all time periods include zero within the confidence intervals, supporting the identifying assumption. The null estimates in the pre-period also provide supporting evidence that firms do not engage in deterrence behavior or limit pricing when threatened by entry by dropping prices to make entry seem unappealing, nor do they engage in predatory pricing to force the exit of a nearby station. While I do not have data on construction or permitting dates for new stations, the lack of any significant change in prices for years prior to arrival is likely sufficient. Firms only drop prices at the start of new operations and accommodate entry.

Secondly, by breaking out the entry and exit effect separately, we can see that the prior results from the difference-in-differences estimator, which is identified off of variation from both entry and exit events using the station count variable as the regressand, attenuated the effect due to heterogeneous effects by event type. Figure 2.7 shows precisely estimated null effects for exit events, which stand in contract to the negative effects of entry.

Lastly, an event study design allows for estimates of the dynamics of the treatment over the duration of the post-period. Focusing on the entry effect in Figure 2.6, results show that entry of a station is associated with a sharp and immediate drop in the price of incumbent of around 2 cents. Effects are precisely estimated and persistent over the long run suggesting that the entry of a new station results in a quick shift to a new, lower price equilibrium.

FIGURE 2.7. Effect of Station Exit on Incumbent Pricing

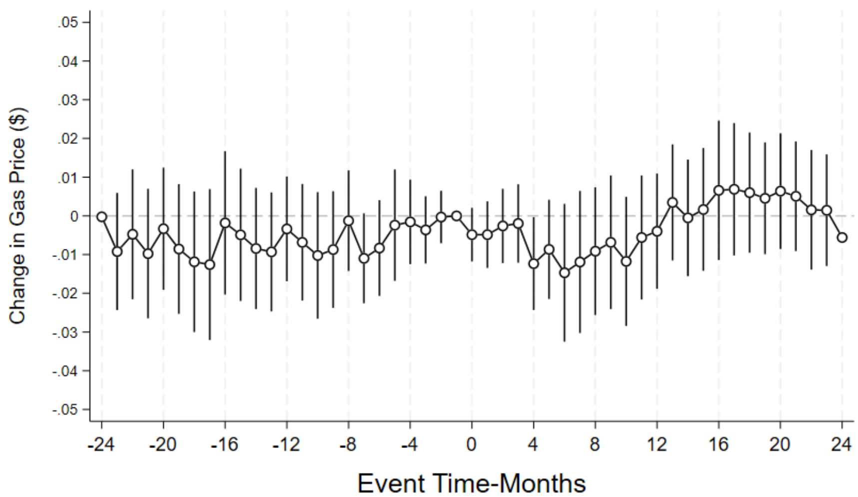

Notes: Coefficient estimates for the effect of station exit within 1 mile of an incumbent station on incumbent pricing of regular unleaded gasoline are shown. Event time t=-1 is the month prior to the departure event. Coefficients are estimated from a linear regression of price on a panel of event time dummy variables, day-of-sample fixed effects, and city-specific linear time trend. Standard errors are clustered at the city level. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018.

In both the difference-in-differences and event study models, results would be attenuated if the arrival of a new gas station were simply a replacement for a nearby exiting firm within the same market. In the extreme, perfectly timed entry and exit would not contribute to identification in the difference-in-differences model as the station count variable would remain constant across time. The event study design may also result in muted effects if entry and exit often accompany each other around the timing of the event. In spite of this, a negative relationship between entry and price could still exist in the presence of correlated entry and exit if the new station varies in product space with better site amenities or selling capacity compared to the exiting firm.

FIGURE 2.8. Effect of Station Entry and Exit on Market Size

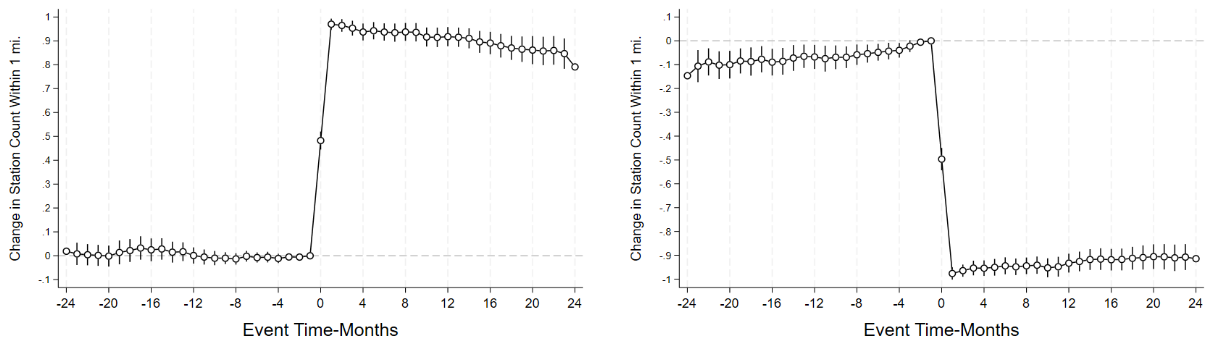

(a) Entry (b) Exit

Notes: Coefficient estimates for the effect of station entry (Panel (a)) and exit (Panel (b)) within 1 mile of an incumbent station on the number of competitors within 1 mile are shown. Event time t=-1 is the month prior to the event. Coefficients are estimated from a linear regression of station count on a panel of event time dummy variables, day-of-sample fixed effects, and city-specific linear time trend. Standard errors are clustered at the city level. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018.

In Figure 2.8, I report the estimated coefficients from an event study specification similar to equation 2.4, but using the station count within 1 mile as the outcome variable regressed on the event time dummies for entry events in Panel (a) and exit events in Panel (b). Results suggest that there is no evidence of exits preceding entry events, shown by precisely estimated zero for event time coefficients during the pre-period. If there were significant exit preceding entry events, we would expect a downward sloping line above zero leading up to event time zero. At the time of the entry, the station count increases by 1 eliminating simultaneous entry and exit as a source of bias. The change in market size is persistent over the post-period with minimal decay in the store count up to one year post event. This rules out exit in the post-period as a driving force in the price results. This also supports the earlier descriptive analysis which showed very few stations experienced multiple entry. Results for exit events are qualitatively similar, with only slight evidence of leading entries in the market. At the time of the exit, there is no evidence of simultaneous entry, and no evidence of lagging entry in the year after the station exit event.

2.4.3. Robustness Checks. In the presence of heterogeneous treatment effects and staggered treatment timing across units, recent work has shown that estimates of the average treatment on the treated from the canonical two-way fixed effects model, as outlined in equation 2.2, can be biased. This is due in part to the method's use of previously treated observations as control observations for later treated units, which is an incorrect counterfactual comparison in the presence of heterogeneous effects across treatment cohorts. While this is not likely to be a concern in my setting since the majority of observations come from never-treated stations, I confirm that results are robust to newer estimators below.

Many of the proposed alternative estimators rely on estimating treatment effects for each treatment cohort immediately before and after treatment non-parametrically to ensure appropriate control units are used in estimation. These approaches are computationally inefficient in the presence of many treated cohorts, which are defined by the exact date of entry in the current data specification. Additionally, to the extent that there is any mismeasurement of the exact date of entry or exit, these methods can yield biased estimates. Collapsing the data to higher levels of aggregation eases the computational constraint as cohorts can be defined by common month or quarter of entry, however this strips the model of the essential identifying variation of daily prices.

I follow Davis et al. (2023) in using the solution proposed in Wooldridge (2021) to show that results are robust to increased levels of flexibility to the TWFE model. The method fully saturates the TWFE model by estimating coefficients for treatment-cohort and time interaction terms in addition to the unit and time effects from the TWFE specification, parametrically estimating the heterogeneous treatment paths across cohorts. These group coefficients can then be aggregated across treatment cohorts to calculate the average treatment effect of the treated parameter.

Focusing first on the effect of entry on price in Row 1 of Table 2.3 below, Column 1 reports the coefficient from a version of the TWFE regression in equation 2.2 using a treatment indicator variable for the timing of the first entry event with station and quarter-of-sample fixed effects as a point of direct comparison for the later Wooldridge (2021) estimator. Entry is associated with a 2.2 cent decrease in incumbent pricing, similar to the results shown earlier from the event study. In Column 2, I show results are robust to the correction used in Davis et al. (2023) which increases flexibility in the model by including city-specific linear time trends and a treatment cohort-specific linear time trend, where treatment cohorts are defined by common quarter-of-sample for the entry event.

Table 2.3. Effect of Nearby Station Entry and Exit on Incumbent Pricing

Dependent Variable	(1)	(2)	(3)	(4)
Entry	-0.022***	-0.019***	-0.024***	-0.027***
	(0.004)	(0.004)	(0.003)	(0.004)
Exit	0.002	.002	-.001	-.001
	(0.004)	(0.004)	(0.003)	(.004)

Notes: Estimated coefficients on dummy variables for the first entry and exit experienced by incumbent firms are estimated in separate regression and shown by row. Column 1 reports the baseline TWFE estimates. Column 2 adds city-specific linear time trends and treatment cohort time trends. Column 3 implements the estimator from Wooldridge (2021) designating treatment cohorts at the quarter level, while Column 4 uses a monthly treatment cohort designation, but only a 5% sample of never-treated station. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018.

Next, in Columns 3 and 4 I report the estimated average treatment effect of the treated parameters from implementation of the full estimator from Wooldridge (2021). Designation of treatment cohorts below the quarter-of-entry level was computationally infeasible using the full sample of over 14 million observations. Column 3 reports the aggregated ATET using the quarter-of-entry to designate treatment cohorts. In Column 4, I report the aggregated ATT using month-of-entry to designate treatment cohorts using the full sample of treatment observations, but a restricted 5% random sample of never treated stations for control observations. Results obtained from the TWFE specifications are robust to the correction methods, and similar results are produced from both levels of aggregation in the Wooldridge (2021) estimator.

Turning towards the exit effects in the next row of Table 2.3, the estimated null effects from the TWFE model are also robust to both the use of the Wooldridge (2021) estimator and the level of cohort designation in Columns 3 and 4.

To this point results have focused on the most commonly sold blend of gasoline, regular unleaded 87 octane gasoline, which represents around 70% of motor gasoline sales in California. In table 2.4, I present results for the other two main blends of gasoline, mid-grade (89 octane) and premium (91 octane). Results closely mirror the outcomes for regular blend gasoline shown above. Nearby entry is associated with a 1.6 cent price decrease for mid-grade gasoline, while premium gasoline decreases by 1.5 cents. as with regular gasoline, exit events are not associated with statistically distinguishable changes to incumbent pricing.

Table 2.4. Effect of Entry and Exit by Fuel Blend

	Regular	Mid-Grade	Premium
Count	-0.010***	-0.009***	-0.008**
	(0.003)	(0.003)	(0.003)
Entry	-0.020***	-0.016***	-0.015***
	(0.004)	(0.004)	(0.004)
Exit	0.000	0.003	0.001
	(0.004)	(0.004)	(0.004)
Station FE	Yes	Yes	Yes
Day of Sample FE	Yes	Yes	Yes
Linear Time Trend	Yes	Yes	Yes

Notes: Coefficients for the station count variable, the first entry, and the first exit experienced by a station are estimated separately and shown by row. Columns indicate the blend of gasoline used for the outcome variable. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Turning focus back to the regular grade gasoline, I report results from varying the distance measure used to define the market size surrounding a station in Table 2.5. Column titles show the length of the radius, in miles, used to define the circular market. Results evolve across the distances as expected with the strongest effects calculated for the narrowest market definition. Results suggest that effects quickly dissipate with increasing market size.

Table 2.5. Effect of Entry and Exit by Market Definition

	.5	1	2	3	4	5	10
Count	-0.011***	-0.010***	-0.009***	-0.008***	-0.006***	-0.006***	-0.005***
	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Entry	-0.021***	-0.020***	-0.018***	-0.019***	-0.013***	-0.008*	-0.011*
	(0.006)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.006)
Exit	0.002	0.000	0.001	-0.001	0.000	0.001	0.006
	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)

Notes: Effects of station count, entry, and exit on incumbent prices are reported for various market definitions. Column titles indicate the size of the circular market definition around each station in miles. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

2.5. Discussion

2.5.1. Market Concentration. There is not a consensus in the literature on the effect of changes to market size on price levels. Canonical work by Bresnahan and Reiss (1991) on competition in homogeneous goods markets suggests that entry into smaller, consolidated markets results in larger competitive effects and that this effect dissipates as the number of competitors in a market increases. However, Barron et al. (2004) show that this result can reverse in search model with price dispersion depending on the search costs involved. I test for heterogeneous effects of entry across market sizes in my setting by estimating equation (2.2) separately by initial market size.

Results using the preferred specification with day-of-sample and station fixed effects, and city specific time trends are shown in Figure 2.9. Effects dissipate towards zero, mostly monotonically, as market size increases. Effects are concentrated among markets with less than 5 competitors within a 1-mile radius. This in fact aligns with the findings in Bresnahan and Reiss (1991) which found entry above 5 firms led to no changes the competitive conduct of incumbent firms in their setting. Given the lack of data on station characteristics, I am unable to claim if the price effect is due to eroding market power or a result of heterogeneous station characteristics.

FIGURE 2.9. Effect of Station Entry Within 1 Mile on Incumbent Pricing

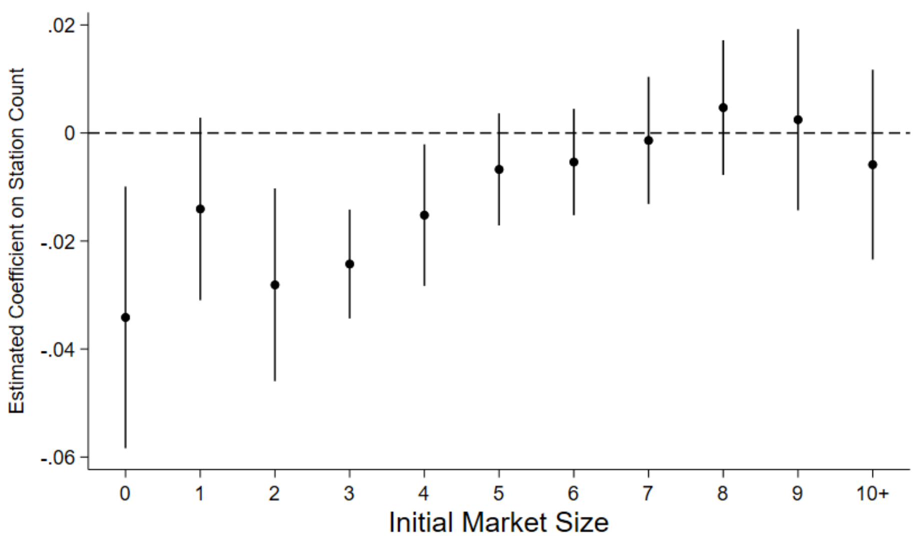

Notes: Coefficient estimates for the effect of station entry within 1 mile of an incumbent station on incumbent pricing of regular unleaded gasoline are shown. Coefficients are estimate by running separate regressions by initial market size. Regressions include day-of-sample, station fixed effects, and city specific time trends. Standard errors are clustered at the city level. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018.

2.5.2. Mystery Surcharge. California has the highest gasoline prices in the country and research has recently focused on trying to explain the divergence in prices between the state and the rest of the country. The price of crude oil, traded on a global market, is common across the country after accounting for transportation costs, as is the federal gas tax, and thus do not contribute to the difference in prices. Much of the difference is explained by increased regulation in California. California's gasoline excise tax, sales tax, environmental fees from the state's cap-and-trade program, low carbon fuel standards, and underground storage tank fees combine to be 2-3 times greater than the average state-level taxes in the rest of the country. Lastly, California has strict emission standards on gasoline sold within the state, resulting in a blend (CARBOB) that is uniquely used in California. The blend, almost uniquely supplied by refiners located in California, results in refinery level prices higher than in the rest of the country. However, even after decomposing the components of the pump-level gasoline price and accounting for the observable differences outlined above, there remains a portion of the price difference that is unexplained.

Recent work by Severin Borenstein has noted that prior to the Torrance refinery fire on February 18, 2015, the difference in pump-level prices between California and the rest of the country was nearly exclusively explained by these components. After the fire, the ensuing supply shock resulted in prices increasing in California. However, prices never converged back in line with national prices once the refinery resumed activity. An unexplained price difference averaging 40 cents per gallon emerged, a phenomenon referred to as the "Mystery Gasoline Surcharge". This has resulted in $48 billion dollars of additional expenditures on gasoline for Californians as of 2023. The mystery surcharge is concentrated in the supply chain after the gasoline has left the refinery or wholesale rack, which leaves the retailing industry, contracting between refiners and retailers, and transportation costs between wholesalers and the pump as potential sources of the difference.

If station markups increase, this can potentially impact the entry effect of a new station. Using the timing of the refinery fire as an exogenous supply shock, we can modify the difference-in-differences equation above to test for differential entry effects before and after the refinery fire by including an an additional term interacting the post-entry ($Entry_{st}$) indicator with a indicator ($Fire_s$) for whether the entry occurred in the period after the refinery fire. The time and station fixed effects absorb the other standard terms required in a difference-in-differences setting.

$$P_{st} = \alpha + \beta_1 Entry_{st} + \beta_2 Entry_{st} * Fire_s + \sigma_s + \delta_t + \rho_c(t) + \varepsilon_{st} \tag{2.5}$$

The outcome variable, P_{st} is price at the station-date level. The coefficient, β_1 reflects the reduced form effect of the entry of a station within a 1-mi. radius of station s. The coefficient β_2 represents the incremental change to the entry effect when an entry occur in the post-fire period, with linear combination of $\beta_1 + \beta_2$ capturing the full effect of entry for post-fire entry events. As before, it is important to include station fixed effects and linear time trends to account for the endogeneity of the location of entry.

Results are presented in Table 2.6 below. In Column 1, I present the baseline estimation of the entry effect using the difference-in-differences analog to the event study results shown above for entry. Entry within a 1 mi. radius is associated with a 2.0 cent decrease in incumbent pricing on average over the full sample period. Column 2 reports the estimated coefficients from estimation of equation (2.5). We fail to reject similar magnitudes of entry for events which occurred before and after the refinery fire, shown by the coefficient on Entry x Post being small and statistically insignificant.

Table 2.6. Effect of Station Entry Post Torrance Refinery Fire

	(1)	(2)
Entry	-0.020***	-0.024***
	(0.004)	(0.009)
Entry x Post		0.005
		(0.010)
Station FE	Yes	Yes
Day FE	Yes	Yes
Linear Time Trend	Yes	Yes
R Sq.	0.958	0.958
Obs	14,745,166	14,745,166

Notes: OLS estimates for the effect of entry within 1-mile on incumbent firm pricing for regular unleaded gasoline are shown. Column 2 interacts an indicator for whether the entry occurred post Torrance refinery fire with the post-treatment variable. Model standard errors are reported in parentheses and clustered at the city level. Data are sourced from the Oil Price Information Service (OPIS) for all stations in California for 2014-2018. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

2.6. Conclusion

Using daily price data and the timing of the entry of new gas stations and exit of existing gas stations, I estimate the effect of market size changes on the pricing for incumbent stations. The use of high-frequency data and the ability to restrict identification to within-station variation allows the approach to account for the endogenous entry and exit decisions of firms. I find that an increase in market size from entry is associated with statistically significant 2 cent decrease in the price at incumbent stations. This is compared to a precise null effect for the exit of a nearby firm. Both results are robust to various specifications, new estimators that correct for heterogeneous treatment effect in the two-way fixed effects specification, and across the various blends of gasoline sold. I also show that the results are strongest for the closest entries and dissipate as the market definition broadens. Exit is a precise null effect at all market definitions. These results are in line with and of similar magnitude to recent studies in other countries (Davis et al., 2023; Fischer et al., 2023).

I also contribute to the conversation about market power in the California gasoline market. California has gas prices that have diverged from the rest of the country for unexplained reasons after the Torrance refinery fire in February 2015. I provide some of the first evidence in the literature discussing changes in market conditions after this exogenous shock by showing that the magnitude of entry effects are unchanged despite the elevated prices. If the unexplained markup were concentrated at the retail pump, then we may expect larger entry effect in the post-period.

Figure 2.10. Number of Entrants per Incumbent Station by Initial Market Size

		Number of Entrants						
		0	1	2	3	4	5	Total
Initial Market Size	0	783	43	3	1	0	0	830
	1	781	81	12	3	0	0	877
	2	968	124	9	3	0	0	1,104
	3	1,084	163	25	2	0	0	1,274
	4	919	158	19	1	1	0	1,098
	5	970	153	15	2	6	1	1,147
	6	765	120	9	2	1	0	897
	7	636	165	14	2	0	0	817
	8	438	101	13	7	2	0	561
	9	289	54	12	0	0	0	355
	10	207	37	11	0	0	0	255
	11	114	25	2	0	0	0	141
	12	57	16	1	1	0	0	75
	13	48	13	0	0	0	0	61
	14	16	6	1	1	0	0	24
	15	9	3	0	0	0	0	12
	16	3	1	0	0	0	0	4
	17	3	0	0	0	0	0	3
	19	1	0	0	0	0	0	1
	Total	8,091	1,263	146	25	10	1	9,536

Notes: The number of entrants experienced for each station is shown by the initial market size for the incumbent station.

CHAPTER 3

The Interplay Between Health Insurance and Auto Insurance: Lessons from the Affordable Care Act (With Andrew Friedson)

3.1. Introduction

In the event of an auto accident resulting in injury, the responsible party for the medical expenses incurred depends on the determination of liability, with the at-fault party assuming responsibility for both parties' injuries. However, this determination is often not made immediately, leaving each party responsible for their own medical costs. The decision to seek medical care during this period is influenced by several factors including the severity of injuries, the likelihood of being held liable, and the out-of-pocket costs associated with care. The presence and level of health insurance coverage directly affects out-of-pocket costs and has been found to be associated with healthcare utilization and overall expenditures as demonstrated in studies by Manning et al. (1987), Card et al. (2008), Finkelstein et al. (2012), and Baicker et al. (2013). As a result, individuals with health insurance are more likely to seek care and accumulate medical expenses prior to liability determination, all else being equal.

Once liability is determined, all accumulated medical bills become the responsibility of the liable party, and their auto insurance becomes "primary," meaning that it supersedes any health insurer for payment of medical bills, including bills that have already been paid through a process known as subrogation.[1] Thus it is reasonable to posit that, all else being equal, auto accidents where the parties have health insurance will have larger outlays on average for the auto insurer than auto accidents where the parties do not have health insurance.[2]

[1] The health insurance company which initially paid for medical care is entitled to compensation from the auto insurance settlement. For more on subrogation, see: https://www.phiagroup.com/Media/Guide-to-Understanding-Subrogation

[2] In reality, states have varying tort law approaches to negligence in auto accidents. We present the case that is true for the majority of states in which the party deemed primarily at fault pays for the damages in the accident. 12 states operate under pure comparative negligence laws that attempt to align damage recovery with the determined percent negligence of each driver, while 4 states operate under pure contributory negligence structures which restrict all damage recovery for any driver found 1% at-fault or greater.

This creates a potential spillover effect from changes in health insurance coverage in the driving population on to auto insurance companies. As health insurance coverage expands, so might the volume of medical expenses paid by auto insurers, leading to higher premiums for auto liability insurance (since premiums are set as a function of prior outlays). In this study, we test for such a spillover in the context of the largest expansion of health insurance coverage in the U.S. in a generation, the Patient Protection and Affordable Care Act (ACA). The increase in health insurance coverage in the United States due to the ACA is well documented (Courtemanche et al., 2019, 2017; Frean et al., 2017; Kaestner et al., 2017; Simon et al., 2017; Sommers et al., 2015; Wherry and Miller, 2016). The ACA has also been linked to increases in utilization of medical care, including both primary care (Wherry and Miller, 2016) and emergency care (Akosa Antwi et al., 2015; Courtemanche et al., 2019; Nikpay et al., 2017).

In this study, we use geographic variation in the ACA's impact on health insurance coverage growth to show that locations that experienced larger expansions in health insurance coverage also experienced the largest increases in the frequency of auto insurance claims. These findings are most plausible (i.e., best satisfy the identifying assumptions of our model) in states with the largest health insurance growth. We further find that these increases in claim frequency are attributable to the ACA's private market reform portion of the legislation as opposed to the expansion of Medicaid eligibility. The private market effects are largely offset by opposite signed effects from Medicaid expansion in the states that expanded their programs as a part of the ACA. Our results imply that the non-Medicaid expansion portion of the ACA caused an additional outlay of $21.67 per earned exposure across coverages, with this increased expense eventually passing through to auto insurance customers as insurers increase future rates. Lastly, we show that this effect is isolated to the extensive margin, with no discernible effect on the average severity of claims due to the ACA.

3.2. Background

3.2.1. Auto Insurance. Automobile insurance is the largest property and casualty line of insurance in the United States, accounting for over $240 billion in direct written premiums in 2018.[3] [4] Proof of automobile liability insurance is required in 49 states and the District of Columbia for a

[3] http://www.ambest.com/review/displaychart.aspx?Record_Code=274410

[4] Direct written premiums are total written premiums net of premiums ceded for reinsurance.

vehicle to be legally operated.[5] Oversight of the automobile insurance industry is a power left by the federal government to the states, resulting in heterogeneity across states in regulation of the industry.

Two general legal approaches to auto insurance have emerged: the tort law system and the no-fault system. The majority of states operate under a tort law approach where the party deemed at-fault for an accident is held legally liable for damages caused to the injured party.[6] The alternative legal structure is a no-fault approach, where parties involved in an accident recover their damages from their insurance company on a first-party basis, with the goal of reducing litigation and transaction costs without the need to prove fault. In practice, no-fault states operate under a partial no-fault structure where the initial damages are treated under the no-fault structure up to a threshold, after which the driver may seek further recovery for damages from the at-fault party via tort law.[7]

An auto insurance plan is actually a bundle of several different insurance products (coverages). Two main components comprise the third-party liability portion of the policy; Bodily Injury (BI), which covers costs associated with injuries to others for which you are legally liable, and Property Damage (PD), which covers costs associated with damage to the physical property of others you have caused. There are also several first-party insurance products that may be present, but which are not legally required; Collision, which covers damage to your vehicle you caused during the act of operating the vehicle, and Comprehensive, which covers acts of nature, theft, fire, and other causes of damage outside of your immediate control.[8] [9] Finally, Personal Injury Protection (PIP) is first-party medical coverage offered in partial no-fault states to cover injuries to the insured and to others in their vehicle during an accident.

[5] In 30 states, financial responsibility laws allow for bonds or deposits with the state up to the amount of the minimum insurance requirement for that state may be used in lieu of insurance. New Hampshire is the exception state which does not have a compulsory insurance law.

[6] Heterogeneity also exists amongst states with regards to how fault is determined for an accident using shared and comparative negligence in an accident, the study of which is beyond the scope of this paper. Our empirical models control for time invariant heterogeneity at the state level, accounting for any differences along this dimension as well as other persistent regulatory differences at the state level.

[7] The distinction between the two approaches to auto insurance is important since in terms of how injuries are treated, first-party medical coverage under Personal Injury Protection is less likely to be impacted by insurance expansion due to the ACA since there is much less ambiguity at the time of the accident as to whether or not injuries will be covered by the auto insurance company. In the analyses that follow we first examine Bodily Injury coverage as this is the prevailing insurance coverage for injuries nationwide, but then perform analyses on other coverages as well.

[8] These two forms of insurance are commonly grouped together as "Collision and Comprehensive."

[9] As with most insurance products, there are subtleties to the definitions of what is covered under each coverage which are beyond the scope of this paper.

In terms of the interaction between health insurance and auto insurance, automobile liability insurance is typically primary to collectible health insurance, meaning health insurance will cover expenses only once the at-fault party's auto liability insurance has been exhausted. If a party to an accident seeks medical treatment initially using health insurance coverage when auto insurance is available, the health insurer is entitled to seek indemnification from the Bodily Injury settlement of the automobile accident once payment is disbursed in a process known as subrogation.

As with other forms of insurance, auto liability premiums are ideally set to approximately cover expected insurance outlays as well as necessary overhead expenses. Private passenger automobile insurers are required to file rating plans for their policies for scrutiny and approval in the majority of states. This introduces a lag between when circumstances change in terms of the underlying risk of the insured population and when insurance rates can be adjusted to reflect said changes. This lag is extended further by the fact that insurance policies can take up to 12 months to reach their renewal date, at which time insurers are permitted to process a rate increase.

Due to these lags, auto insurance companies have a financial interest in accurately forecasting future risk, as any adjustments to overall premiums need to begin the approval process 18-24 months in advance.[10] These lags also imply that unforeseen shocks to outlays can powerfully impact firm balance sheets: large unforeseen drops in outlays can lead to windfalls for firms, whereas large unforeseen increases in outlays can put insurers in a situation where they have collected insufficient premiums to cover incoming claims.[11]

3.2.2. The Affordable Care Act and Health Insurance. Passed in 2010, the Patient Protection and Affordable Care Act was aimed at expanding enrollment in health insurance. The law included a variety of industry reforms, government subsidies, and mandates which came into effect in waves over the years following its passage.[12] Some of the largest changes took effect in 2014 when the ACA launched private health insurance "Marketplaces" or "Exchanges" where individuals and small groups could purchase subsidized plans, introduced the "individual mandate"

[10]This refers to the base premium charged to the average customer. This is not the case for rate changes due to individual level experience rating (i.e. individual rates increasing immediately following an accident or traffic citation) which are implemented far more quickly - that said, the details of how experience rating is to be done in a given cycle are filed as part of the rate plan.

[11]A recent example being the large reduction in vehicle miles travelled due to the Covid-19 pandemic where auto insurance companies issued refunds of up to 15% of premiums to customers. https://www.aarp.org/auto/car-maintenance-safety/info-2020/coronavirus-car-insurance-premium-refund.html

[12]For a detailed catalog of the components of the ACA, see Kaiser Family Foundation (2013).

tax penalty for not having health insurance, and also expanded Medicaid eligibility to 138 percent of the federal poverty line (Gruber, 2011). However, in the 2012 U.S. Supreme Court case *National Federation of Independent Business v. Sebelius*, the court ruled that states were allowed to opt out of the Medicaid expansion portion of the ACA. At the beginning of 2014, 25 states opted to expand Medicaid under the ACA, with 7 additional states reversing course and choosing to expand Medicaid by the end of our sample in 2016.

An expansive literature has documented the impact of the ACA on insurance coverage in the U.S. population, finding that both the state specific Medicaid expansions and the implementation of the 2014 package of reforms increased insurance coverage (Courtemanche et al., 2019, 2017; Frean et al., 2017; Kaestner et al., 2017; Simon et al., 2017; Sommers et al., 2015; Wherry and Miller, 2016). The majority of the newly insured had existing lapses in health insurance of more than 3 years (Decker and Lipton, 2017). The change in the national insurance level can be observed plainly in summary data: in 2013, according to the American Community Survey (ACS), 44.4 million (or 16.8 percent of) Americans were uninsured, as of 2015 that number had dropped to 29.1 million (or 10.9 percent of) Americans

Health insurance coverage typically lowers the out-of-pocket expense of purchasing medical care, meaning that under most circumstances, the presence of health insurance is expected to increase consumption of medical care and total expenditures on medical care (by the insured and insurer combined). This prediction has been repeatedly verified by both randomized studies (for example, Manning et al. (1987); Finkelstein et al. (2012); Baicker et al. (2013)) as well as by quasi-experimental studies (for example, Currie and Gruber (1996); Anderson et al. (2012); Miller (2012b)).[13]

Studies on the ACA have documented increases in utilization of medical care due to both the Medicaid expansions and the other 2014 reforms (see Mazurenko et al. (2018) for a review of this literature).[14] In terms of care that would be possibly consumed post auto accident: Miller and Wherry (2017) and Wherry and Miller (2016) find evidence of increased utilization of both primary and specialist care due to the ACA. Nikpay et al. (2017) as well as Gotanda et al. (2020) find

[13] There are instances where health insurance decreases certain types of medical care utilization, for example emergency care that could be avoided by proper preventative care or treated in a different and more appropriate venue (Miller, 2012a).

[14] Less pertinent to this study, Jhamb and Colman (2015) find increases in utilization due to the ACA's 2010 dependent coverage mandate.

evidence of increased emergency department utilization due to the ACA, and Courtemanche et al. (2019) as well as Courtemanche et al. (2019) find evidence that the ACA led to increased utilization of ambulance services.

3.2.3. Potential Mechanisms. There are several avenues through which an increase in the prevalence of health insurance can impact automobile insurance claims, leading to a theoretical ambiguity as to the magnitude and direction of the relationship.

One possible avenue is that at the time of an accident, fault for any damage may not be clear and would be determined after the date of the incident.[15] Each party's decision of whether to seek medical attention (as well as how much care to consume) could in many cases be made *before* liability is established. The presence of health insurance has been shown empirically to be important to care decisions post-accident: Doyle (2005) found that those without health insurance who suffered an auto accident consumed 20 percent less care than those with health insurance. In these cases, each party would bear the up-front cost of any care, which would be mitigated by the presence and quality of any health insurance. As health insurance typically lowers out of pocket costs, we would expect individuals with health insurance to be more likely to seek care, all else equal. If that individual is later determined to be not-at-fault, then this health expense would become the responsibility of the at-fault party's auto liability insurance.

A second possible connection between health insurance and auto insurance is that health insurance may create *ex post* moral hazard with regards to driving behavior. In this scenario, the newly insured could experience a Peltzman effect (Peltzman (1975)), where additional risk protection (in this case, health insurance mitigating the potential liability cost of poor driving) could cause individuals to drive more recklessly, increasing the likelihood of accidents and subsequent claims. We find this possibility to be less likely given the evidence from Courtemanche et al. (2019) showing that the ACA did not alter risky behaviors observed in individuals involved in motor vehicle fatalities.[16]

A third possible link between the two types of insurance is that health insurers have greater bargaining power than auto insurers *vis-a-vis* medical providers and are able to extract lower prices

[15] A determination of liability may require evidence from a police accident report, which, for example, in Colorado can take up to 90 days. https://dmv.colorado.gov/report-accident

[16] These behaviors were drug, alcohol, and seat belt use (all as recorded by responders to the accident). The authors also found that the Medicaid portion of the ACA expansion was associated with an increase in the number of fatal accidents, this finding did not carry over to other components of the ACA.

for medical care. However, auto liability insurance payments are typically made in lump sums directly to the insured based on characteristics of the injury, so any savings along this dimension are likely captured by the auto insured and not the auto insurer.

Another possibility is that in the absence of health insurance, auto insurance can be used to cover health expenditures via fraudulent or padded claims. In this case, when health insurance coverage is expanded, the number of auto insurance claims and severity of claims would decrease as the incentive to commit auto insurance fraud is greatly diminished.

Finally, it is also possible that when health insurance is expanded, individuals on average have a larger pool of health stock (for example due to better case management of chronic conditions) and in some cases may be less likely to need care due to an auto accident.

3.2.4. Existing Literature. To our knowledge, two prior studies have causally studied the relationship between the ACA and automobile insurance. The first, Kadiyala and Heaton (2017), examine the 2010 dependent coverage mandate (DCM), which increased the age that children may remain on their parents' health insurance policy from 18 to 26. The authors found a reduction in total bodily injury claim counts and claim severity for those aged just below the new cutoff of 26 when compared to the control group just older than the new age cutoff. They did not find any discernible effect of the DCM on first party medical coverage PIP claims.

There are reasons to believe the impact due to the DCM on auto insurance may be different than the impact due to the 2014 ACA provisions. The DCM specifically targeted individuals 18-26 who may have different driving, accident, and health care use patterns than the relatively older population impacted by the 2014 components of the ACA. Further, many of the 2014 provisions were aimed at allowing uninsured individuals to acquire insurance when none was otherwise available, in many cases, as part of Medicaid. On the other hand, the DCM allowed young adults to remain on existing private plans. This means that the DCM carried a (by definition) larger proportion of newly insured individuals with private insurance, which differs from Medicaid in terms of generosity of coverage. Additionally, much of the increase in dependent coverage due to the DCM was due to plan upgrading rather than individuals for whom the counterfactual was to be uninsured: for example, Fone et al. (2020) estimate that 1.24 million out of 2.83 million (or 43% of) young adults who gained dependent coverage post DCM transitioned from a different form of health insurance coverage.

The second study on the relationship between the ACA and automobile insurance, Heaton and Flint (2021), explores the effect of the Medicaid expansion portion of the ACA on insurer profitability across the spectrum of property and casualty lines of insurance. The authors find a 7% decrease to calendar-year loss ratios (as a loss ratio is the ratio of losses to premiums, this represents an increase in profitability), and a similar decrease using the alternative outcome of aggregate loss dollars paid for the aggregate sum of automobile liability coverages. The authors are limited by the level of data aggregation and outcome measures used to detect if the result is driven by changes in the value of claims or changes in the number of claims, or which specific coverages are impacted.

Our study contributes to the prior literature along two dimensions. First, we utilize more granular data from the National Association of Insurance Commissioners which allow us to observe both average severity of claim (severity) and claim frequency per exposure (frequency) at the state-year level for each individual coverage. This allows us to examine not only the net effect of health insurance expansion on insurance company risk, but also the component parts of severity and frequency of claims for each unique coverage. Second, we utilize an estimation strategy adapted from Courtemanche et al. (2017) which decomposes the Medicaid component of the 2014 expansions from the other provisions, allowing us to present the first estimates in the literature for each component's impact separately while jointly estimating the impact of the full legislation. This decomposition is useful given that the Medicaid expansion component of the ACA and the other components of the ACA impacted different demographics. For example, the Medicaid expansion portion changed health insurance enrollment for an, on average, much younger population than the other components of the reform (Courtemanche et al. (2019)). As these age groups (and other demographics groups impacted differentially by the law may) have different driving and auto insurance claiming patterns, it is possible that the impact of the Medicaid component on auto liability insurance is quite different from the impact of the remainder of the ACA.

3.3. Data

Our primary data source is the National Association of Insurance Commissioners' (NAIC) annual *Auto Insurance Database Report* which contains claims information, policy exposure measures, and premium amounts. The NAIC collects data from state insurance commissioner offices and private vendor companies, aggregated to the state-year-coverage level. Our dataset covers the years 2010-2016 for all states, excluding Texas.[17] We follow Heaton and Flint (2021) in restricting our sample to 2010 to avoid contamination from the financial effects that the Great Recession had on the insurance industry and auto insurance specifically. Each annual publication reports three consecutive years of data on a three-year lag: the two prior years' data are published alongside the latest year's first reporting. Revisions may be made to annual numbers in subsequent publications, as such, we use the most recently published data for each year's observations.

For the analyses that follow, we construct two statistics that are commonly used in the insurance industry to describe claiming patterns: claim frequency and average claim severity. Claim frequencies are calculated as the number of incurred claims per 100 earned exposures during an accident year, where an earned exposure is measured as one vehicle covered for one full year. Claim frequency measures the probability that an insured vehicle will have a valid claim, not that the vehicle will be involved in an accident. Average claim severity is calculated as the total amount paid out under a given coverage divided by the total number of incurred claims for that same coverage during the accident year. Importantly, claim severity represents the expected payout for the insurer conditional on the accident being reported to the insurer and there being valid coverage provided under the insurance policy, subject to any deductibles or policy limits.

Our two outcome measures of interest will not only be sensitive to changes in the proclivity to seek treatment post-ACA, but also to changes in driving patterns, changes to customer claiming patterns, and changes in how auto insurers administer their plans. The nature of these statistics limits the generalization of our results to all automobile accidents, and instead frames the interpretation of the estimates as the reduced form impact of the ACA on auto insurance claims.

In the data, incurred claim counts and payouts are reported based on the timing of the accident, not the payout. As an example, if an accident occurred on 11/5/2015 and is reported to the insurer the same day, a payout made on 1/15/2016 will still be accounted as part of accident year 2015. Losses for liability coverages have been "developed forward" according to actuarial convention, meaning that they include expected payouts for contingencies that occurred in that year but have not yet been paid.[18]

[17] Earned exposure and incurred claims data are not available for Texas through the NAIC.

[18] The method for "loss development" is done to account for the long-tail nature a coverage where payments can be made substantially later than the occurrence due to litigation and prolonged medical treatment. Loss data in the *Auto Insurance Database Report* is reported on a three-year lag, which allows for seasoning of the data, however

We present estimates for the impact to the frequency and severity of the major auto insurance coverages: bodily injury (BI) liability, property damage (PD) liability, and collision (COLL). Bodily injury is the main coverage for medical damages suffered during accidents in tort liability states, and a required coverage. In no-fault states, PIP provides for medical damages on a first-party basis up to a determined limit, after which bodily injury liability coverage begins to function as in tort states for the excess amount. Importantly, no states switched between a no-fault and a tort liability structure during the period of study, which would structurally change the interpretation of the main outcome variables. Additionally, we are unaware of any systematic change to the operation of individual coverages or insurer plan administration correlated with the timing of the ACA.

Figure 3.1. Map of States by 2013 Health Uninsured Rate

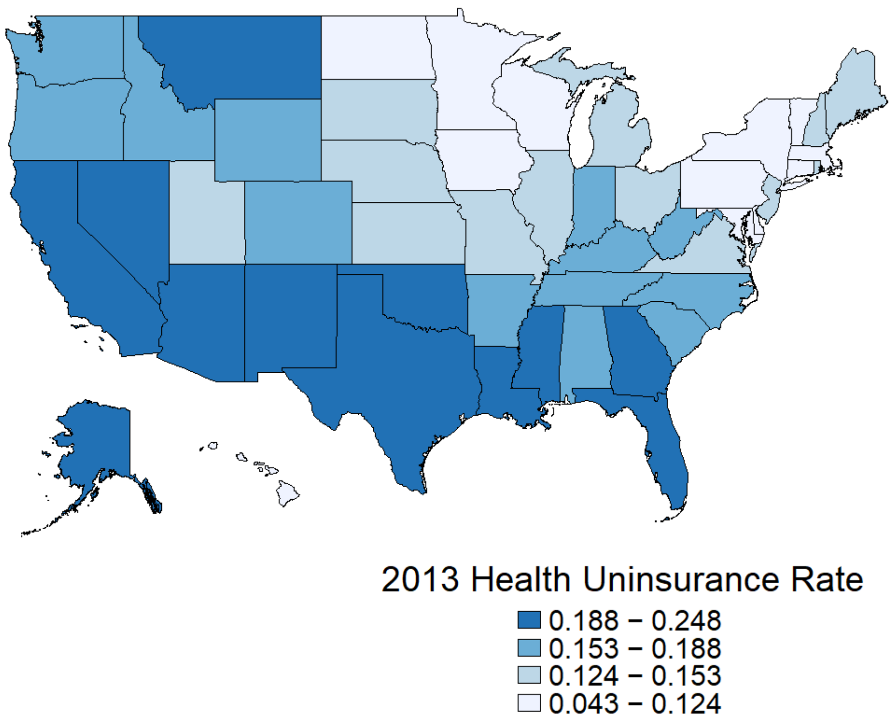

Notes: Health uninsured rates for 2013 are shown at the state level. Data are sourced from the U.S. Census Bureau's Small Area Health Insurance Estimates (SAHIE).

To measure the potential impact of the ACA on an individual state, we use the estimated health uninsured rate for the two years preceding ACA implementation, 2012 and 2013 separately,

bodily injury liability coverage is developed (projected) for an additional 2 years. Physical damage coverages, by contrast, are not developed further due to the short tail nature of these coverages characterized by quick settlement times. Methodologically, loss development is a proprietary actuarial projection based on historical patterns.

from the U.S. Census Bureau's Small Area Health Insurance Estimates (SAHIE). The SAHIE provides annual estimates of the proportion of the population insured and uninsured at various levels of geographic and demographic disaggregation. Figure 3.1 displays a map of the states by health uninsured rate in 2013. Dates for Medicaid expansions are collected from the Kaiser Family Foundation and used to create $Medicaid_s$, a binary indicator variable for whether or not a state expanded Medicaid during our sample and $Post_{st}$ which takes fractional value for the portion of year t state s expanded Medicaid, a value of one for all subsequent years, and a value of zero otherwise.[19] Figure 3.2 highlights the 32 states which adopted and enacted Medicaid expansions as part of the ACA as of the end of the sample period in 2016 as well as the timing of the various expansions.

Figure 3.2. Medicaid Expansion Dates by State

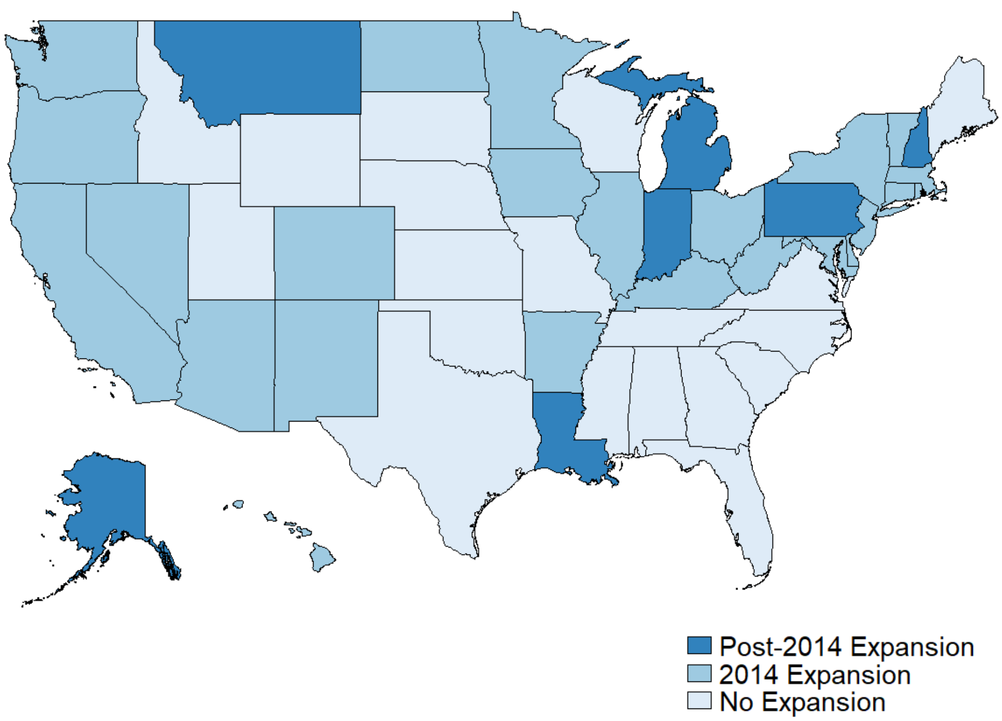

Notes: Medicaid expansion dates are shown as of the end of our sample period at the end of 2016.

We include a matrix of state-year level controls of related economic and traffic variables which are plausibly correlated with both the proportion of the population with health insurance and auto insurance claim experience. We collect vehicle miles travelled per registered vehicle and

[19] Expansions which go into effect on Jan 1st are given a value of 1 for that year.

miles of highway per capita from the Federal Highway Administration's *Highway Statistics* annual publication, state average personal disposable income from the U.S. Bureau of Economic Analysis, and unemployment rates from the U.S. Bureau of Labor Statistics. Lastly, we construct racial composition variables and total population counts from the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) project.

Table 3.1 reports summary statistics for variables used in the data set. The panel is balanced with 350 total observations for 49 states and the District of Columbia over the 7-year sample period.

3.4. Empirical Strategy

To estimate the impact of the ACA on the frequency and severity of auto insurance claims we adapt a version of the triple differences model from Courtemanche et al. (2017):

$$\begin{aligned}(3.1)\quad Y_{st} = \beta_o + \beta_1[Post_t x Uninsured2013_s] + \beta_2[Post_{st} x Uninsured2013_s x Medicaid_s] + \\ \beta_3[Post_{st} x Medicaid_s] + \beta_4\mathbf{X_{st}} + \theta_s + \delta_t + \epsilon_{st}\end{aligned}$$

where the dependent variables are frequency and severity for various automobile insurance coverages in state s in year t. $Uninsured2013_s$ is a binary variable that takes on the value one if state s is above the p percentile of states in terms of percent of population uninsured for 2013 and zero otherwise, where p is either the 25th, 50th, or 75th percentile. $Uninsured2013_s$ denotes a state having more potential room for health insurance enrollment to grow due to the enactment of the ACA starting in 2014. Courtemanche et al. (2017) show that the level of uninsured in a state prior to the ACA strongly predicts actual growth in health insurance enrollment as a result of the ACA. $Post_t$ is an indicator variable taking a value of 1 for all years post-2014 ACA expansion, and a value of zero otherwise. $Post_{st}$ is a state-specific variable denoting whether or not state s has expanded Medicaid as of year t, and takes on fractional values for mid-year expansions as described above. $Medicaid_s$ is an indicator for states which expanded their Medicaid programs under the ACA during our sample period. Its inclusion in the interaction term $Post_{st} x Uninsured2013_s x Medicaid_s$ allows for the estimation of the impact of Medicaid expansion separate from the private market reforms.

Table 3.1. Summary Statistics

	No Medicaid	Medicaid Expansion	Total
Panel A: Frequencies			
Bodily Injury	1.126	0.947	1.006
	(0.362)	(0.434)	(0.419)
Propoerty Damage	3.790	3.961	3.905
	(0.601)	(1.029)	(0.914)
Personal Injury Protection	2.497	1.530	1.738
	(0.837)	(0.548)	(0.734)
Collision	5.319	6.183	5.902
	(0.484)	(0.957)	(0.926)
Panel B: Severities			
Bodily Injury	15,551	20,457	18,857
	(3,797)	(10,179)	(8,927)
Propoerty Damage	3,221	3,224	3,223
	(373)	(436)	(416)
Personal Injury Protection	6,223	12,845	11,429
	(1,760)	(18,958)	(17,034)
Collision	3,669	3,652	3,658
	(395)	(436)	(423)
Panel C: State Characteristics			
Count	19	32	
Health Uninsured Rate-2013	0.182	0.150	0.160
	(0.042)	(0.040)	(0.043)
Health Uninsured Rate-2012	0.184	0.152	0.162
	(0.041)	(0.042)	(0.044)
Personal Disposable Income	36,778	41,650	40,061
	(3,858)	(5,921)	(5,800)
Unemployment Rate	6.975	7.446	7.292
	(2.106)	(2.155)	(2.148)
Population	9,091,444	14,745,693	12,901,434
	(6,002,494)	(12,632,107)	(11,228,906)
Percent White-Pop	0.772	0.792	0.786
	(0.091)	(0.089)	(0.090)
Percent Black-Pop	0.184	0.118	0.139
	(0.094)	(0.072)	(0.085)
Percent Other-Pop	0.043	0.090	0.075
	(0.023)	(0.076)	(0.067)
VMT per Registered Vehicle	0.013	0.012	0.012
	(0.002)	(0.001)	(0.002)
Miles of Roadway	103,752	108,453	106,919
	(26,076)	(52,756)	(45,803)

Notes: Data are from the NAIC's annual *Auto Insurance Database Report* for years 2010-2016. Column 1 and 2 report mean and standard deviation weighted by state population for states that expanded Medicaid during the sample period and states which did not expand Medicaid, separately. Column 3 presents the population weighted totals across all observations.

We include a vector of state-year control variables, $\mathbf{X_{st}}$, to account for economic and demographic variables which may be correlated with both the outcomes of interest and the treatment variables.[20] θ_s is a state fixed effect which accounts for time invariant heterogeneity at the state level and subsumes the time invariant state-level terms in the traditional triple differences specification, specifically, $Uninsured2013_s$ and $Medicaid_s$. δ_t is a year fixed effect accounting for common annual shocks and subsumes $Post_t$. Standard errors are clustered at the state level and regressions are weighted by state population.[21]

The above model has two coefficients of interest, β_1 which captures the effect of being in the top percentile of growth due to the 2014 ACA health insurance market reforms, and β_2 which captures the additional effect of a state expanding its Medicaid program under the ACA. This makes $\hat{\beta}_1$ the estimated effect of the non-Medicaid component of the ACA in a top percentile uninsurance state, and makes $(\hat{\beta}_1 + \hat{\beta}_2)$ the estimated effect of being in a top percentile uninsurance state that also expanded Medicaid.

Causal identification of the estimated coefficients in a triple differences empirical setting relies on two main assumptions. First, outcome variables across treated and control units must follow common trends prior to the policy change, the assumption being that these trends would have continued into the post period absent treatment. The pre-trends portion of the assumption is testable using event study analyses. Second, no other change in circumstances that could explain the outcomes can have occurred concurrently with the policy changes of interest. Under the model specified in equation 3.1, such events would have to be concurrent with ACA and the state-specific Medicaid expansions across a variety of states (each with its own timing), and would have to be correlated with the 2013 rate of health uninsurance to result in biased estimates. We are unaware of any such events.

Due to auto insurance being regulated and administrated at the state level, there are state-specific reporting and data idiosyncrasies.[22] Our research design accounts for these differences as

[20]Specifically, we control for vehicle miles traveled per registered vehicle, miles of roadway, personal disposable income, unemployment rate, and population race demographics.

[21]Alternative weighting schemes are explored in Appendix Table 3.5. We alternatively weight by property damage earned exposures and driving age population 15-75. Property Damage is a required coverage in all states, except for NH, making property damage exposures the best reflection of the number of vehicle-years insured for a given state in a given year. Both weighting schemes yield quantitatively similar results.

[22]Examples include variation in reporting incurred claim counts on a per-accident or per-claimant basis, differences in at-fault thresholds, or required insurance policy limits.

our empirical approach relies solely on within-state variation, and thus idiosyncrasies across states, which do not vary with time, will not bias our results. Additionally, only systematic changes correlated with the timing of the introduction of provisions of the ACA and pre-ACA health uninsured rates would introduce bias into the analysis. We are also unaware of any such systematic changes during our sample period.

3.5. Results

3.5.1. Descriptive Data Patterns. First, we present simple graphical evidence in Figure 3.3 and Figure 3.4 which plot annual BI claim frequency and severity for four groupings of states: combinations of states that were above and below the median 2013 health uninsurance rate, and those that did and did not expand Medicaid during our sample period. For BI frequency (Figure 3), all four groups follow roughly parallel trends in claims prior to 2014. States that had a high level of unisurance (and consequently high health insurance enrollment potential), but did not expand Medicaid, had a noticeable increase in BI claim frequency post ACA implementation.

This data pattern is consistent with the 2014 reforms increasing BI claim frequency in states with large health insurance growth.Our causal estimates focus on the within-group change over time as opposed to differences in levels as factors such as differences in the way insurance coverages are administered across states, inherent road safety, and climate, generate level differences. The pattern observed in Figure 3.3 is also potentially consistent with findings from Heaton and Flint (2021) who find that the Medicaid expansion component of the ACA decreased loss ratios, of which claims frequency is a main component.[23] The lack of growth in BI frequency in Medicaid expansion states, particularly those with high pre-ACA health uninsurance rates, could be due to the effects found by Heaton and Flint (2021) offsetting the growth in BI claim frequency from the non-Medicaid expansion portion of the 2014 ACA reforms. Different effects for the Medicaid and non-Medicaid portions of the ACA are plausible as the reforms impacted groups with different demographic makeups (Courtemanche et al., 2019).

For BI severity, shown in Figure 3.4, the descriptive results fit a positive linear time trend across all groups with no visible break in the trend at the time of ACA implementation.

[23]Loss ratios are calculated as (incurred losses / earned premiums), where losses are a direct function of both the frequency and severity of accidents. All else being equal, lower frequency of accidents would lead to lower loss ratios.

Figure 3.3. BI Frequency by Health Uninsured Rate and Medicaid Expansion Status

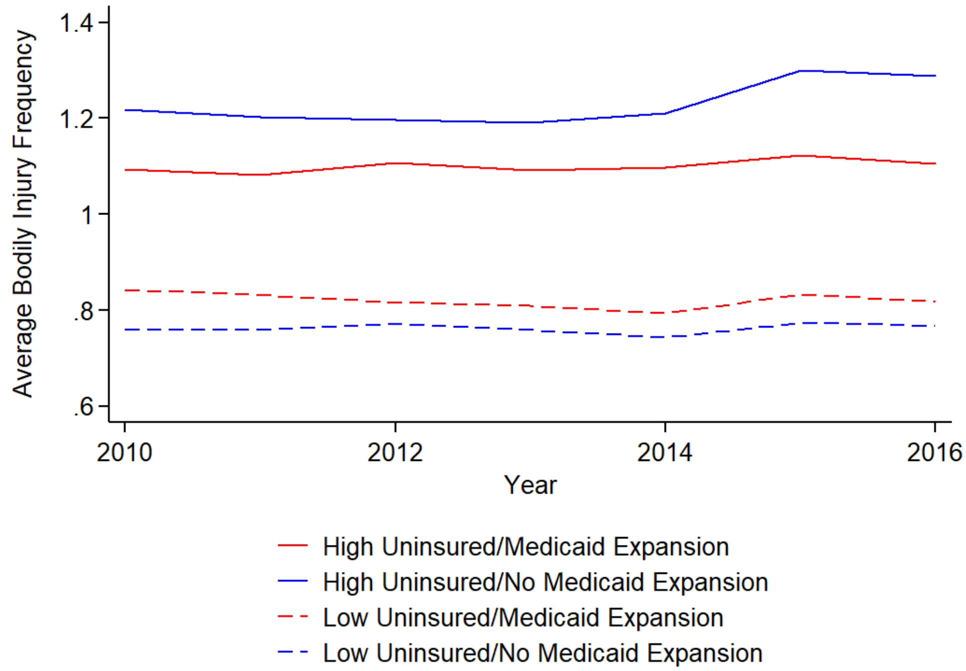

Notes: Average bodily injury frequency rates for each group are shown for the sample period, 2010-2016. High uninsured states are those having a health uninjured rate greater than the national population weighted average. Medicaid expansion states are those which expanded their medicaid programs during the sample period from 2014 to 2016. Frequencies are weighted by state population.

FIGURE 3.4. BI Severity by Health Uninsured Rate and Medicaid Expansion Status

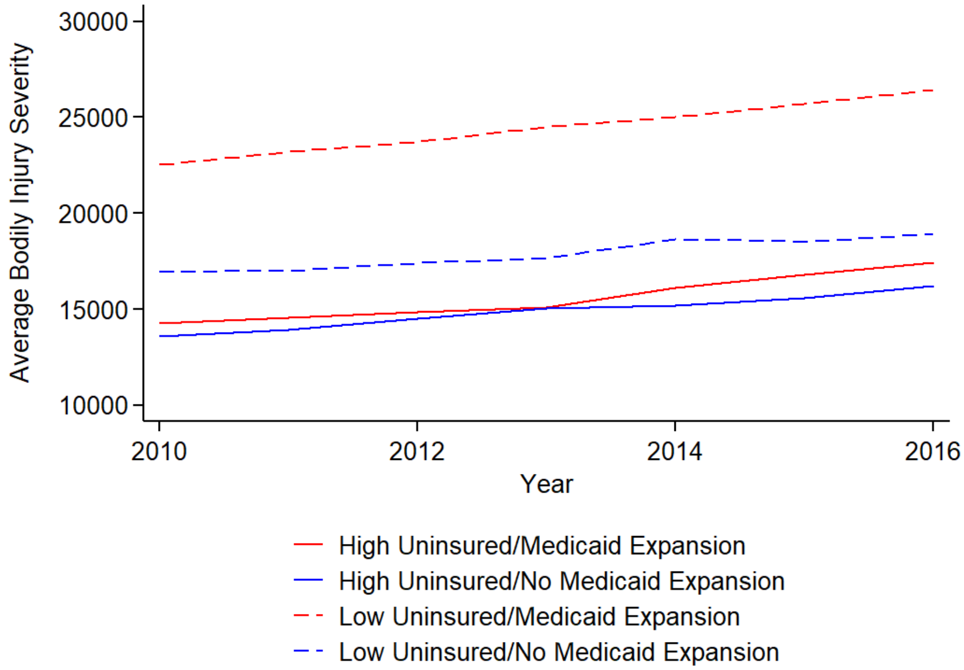

Notes: Average bodily injury severity for each group is shown for the sample period, 2010-2016. High uninsured states are those having a health uninjured rate greater than the national population weighted average. Medicaid expansion states are those which expanded their medicaid programs during the sample period from 2014 to 2016. Severities are weighted by state population.

3.5.2. Triple Differences Model. We next move to results from estimation of the triple differences model shown in equation (3.1). We report results for BI claim frequency and severity in Table 3.2. The first two columns report results for claim frequency with and without the inclusion of control variables and Columns 3-4 report the results for claim severity again with and without controls.

The top panel reports $\hat{\beta}_1$ and $\hat{\beta}_2$ from equation (3.1), the estimated coefficients on the variables $Post_t xUninsured2013_s$ and $Post_{st} xUninsured2013_s xMedicaid_s$. The results are reported three times, each for a different percentile cutoff used to designate the binary variable, $Uninsured2013_s$. From top to bottom the table reports results for defining $Uninsured2013_s$ as states above the 25th, 50th, and 75th percentiles of uninsurance in 2013.

Estimates of $\hat{\beta}_1$, the effect of being in the high uninsurance group relative to the low uninsurance group, are all positive and all statistically significant at conventional levels when $Uninsured2013_s$

Table 3.2. Effect of ACA on Bodily Injury Claim Frequency and Severity by Cutoff

	BI Freq.		BI Sev.	
	(1)	(2)	(3)	(4)
25%				
Post x High Uninsured	0.035	0.032	1,221.9***	1,280.8***
	(0.026)	(0.020)	(253.4)	(248.6)
Post x High Uninsured x Medicaid Expansion	-0.020	-0.020	-649.5	-873.2*
	(0.035)	(0.037)	(498.4)	(485.1)
50%				
Post x High Uninsured	0.064**	0.047*	-79.2	-47.4
	(0.030)	(0.027)	(526.5)	(655.7)
Post x High Uninsured x Medicaid Expansion	-0.039	-0.030	-78.0	-339.0
	(0.034)	(0.034)	(706.0)	(725.2)
75%				
Post x High Uninsured	0.090**	0.080**	370.6	559.1
	(0.040)	(0.035)	(451.0)	(536.9)
Post x High Uninsured x Medicaid Expansion	-0.050	-0.049	-78.3	-452.0
	(0.044)	(0.048)	(599.0)	(765.4)
Economic Controls	No	Yes	No	Yes
Driving Controls	No	Yes	No	Yes
State Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	350	350	350	350
2013 Health Uninsured Rate	0.160	0.160	0.160	0.160
Pre-ACA Mean BI Frequency or Severity	0.997	0.997	18,043	18,043

Notes: OLS estimates for the effect of the private and Medicaid expansion portions of the ACA on total bodily injury claim frequency (columns 1-2) and average claim severity (columns 3-4) are shown. Data are from the NAIC's annual *Auto Insurance Database* report for years 2010-2016. Columns 1 and 3 report estimates from the D-D-D regression in equation (3.1) without economic and driving condition control variables. Columns 2-4 report results using the same model, but with the addition of the control variables. Results are grouped by the cutoff used to designate the high uninsurance level, varying from 25% to 75%. Model standard errors are reported in parentheses and clustered at the state level, with all regressions weighted by state population. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

is defined as at or above the 50th or 75th percentile. The estimates are uniformly increasing in magnitude as the percentile used to define $Uninsured2013_s$ increases. For $\hat{\beta}_2$, we find a decrease in BI claim frequency, and a similar pattern of increasing magnitude as the percentile used to define $Uninsured2013_s$ increases. Taken together, these results imply that insurance enrollment from the non-Medicaid portion of the ACA increases BI claim frequency, whereas the Medicaid expansion decreases BI claim frequency (which is consistent with Heaton and Flint (2021)). In terms of magnitude, being above the median 2013 uninsurance rate yields a 0.047 point increase which is a 4.7 percent increase over the pre-ACA BI frequency of 0.997. For Medicaid expanding states, slightly over half of that increase is eroded by the decrease in claim frequency linked to the Medicaid component.

The next two columns repeat the analysis with BI severity as the outcome variable. Column 4 reports the estimated effect for both non-expansion and Medicaid expansion states. The implied effects are both statistically indistinguishable from zero and economically small in magnitude when compared to the pre-period average severity level of $18,043. These results are in line the ACA not impacting claim severity, which when taken together with the previous results on frequency, is consistent with an increase of care use along the extensive margin (frequency any care is used) following accidents. While we are limited in our ability to make an assertion about the intensive margin impacts due to the aggregation of the data, the null result on BI severity provides evidence that the newly filed claims are not just concentrated at the lower end of the severity distribution as would be expected if only minimal damage claims were avoided in the pre-ACA period due to lack of health insurance.

The states above the median 2013 uninsurance rate had an average of 1.17 million incurred BI claims yearly, or approximately 45 thousand claims per state. Our estimate implies that the above median group had an additional 55 thousand claims (or a little over 2,100 additional claims per state) due to the non-Medicaid portion of the ACA. Given that the average severity of a BI claim in these states in 2013 was $14,967, this implies an increase of approximately $31 million in paid out BI claims per state (or $820 million for the entire above median group).

3.5.3. Event Studies. We next use an event-study design where the key treatment variables, $Uninsured2013_s$ and $Uninsured2013_s xMedicaid_s$, and the $Medicaid_s$ indicator variable are interacted with a panel of event time dummy variables. The outcome of interest remains BI claim frequency. Figure 3.5 reports the coefficient estimates and 95 percent confidence intervals for the effect of the non-Medicaid ACA reforms when $Uninsured2013_s$ is defined as at or above the 25th, 50th or 75th percentile. Event time "zero" is set to 2014, the year the market reforms took place and before the bulk of the state Medicaid expansions. Figure 3.6 repeats this for the additional effect of the Medicaid expansions. For the Medicaid portion, event time "zero" is the set as the year of a state's Medicaid expansion. Both figures show clear changes in trajectory post-ACA that correspond with the respective findings in Table 3.2.

FIGURE 3.5. Bodily Injury Frequency: Event Study By High Uninsurance Cutoff for Non-Expansion States

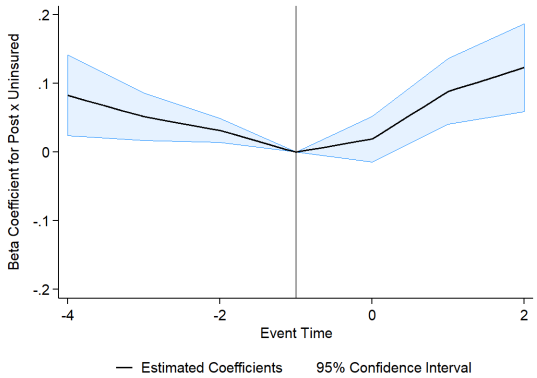

(a) 25%

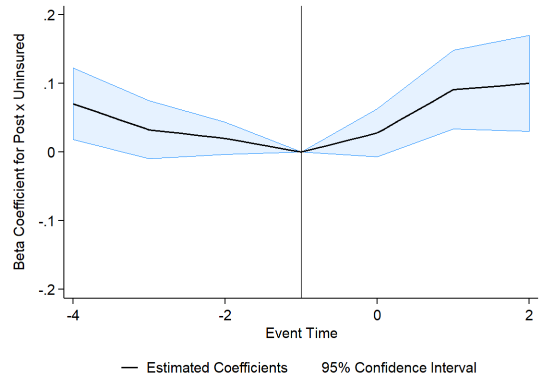

(b) 50%

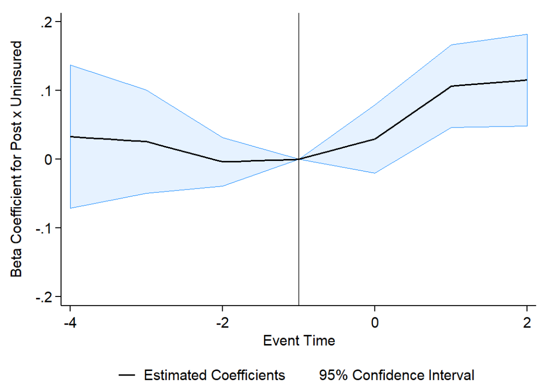

(c) 75%

FIGURE 3.6. Bodily Injury Frequency: Event Study By High Uninsurance Cutoff for Medicaid Expansion States

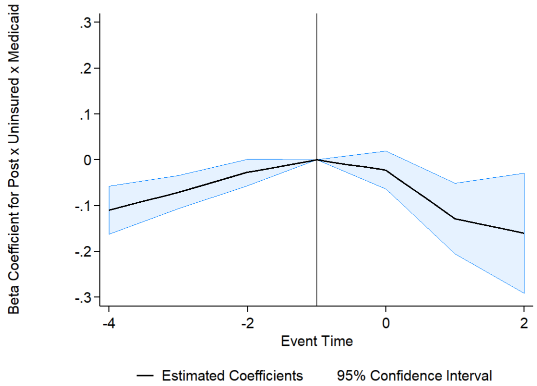

(a) 25%

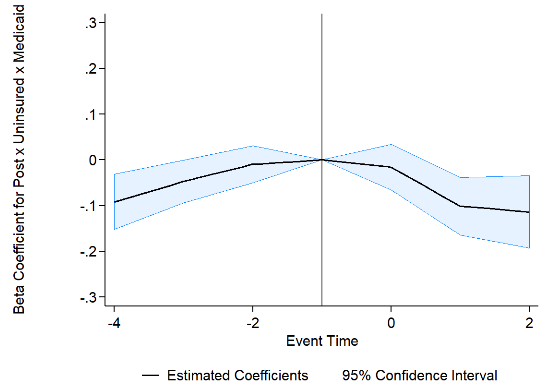

(b) 50%

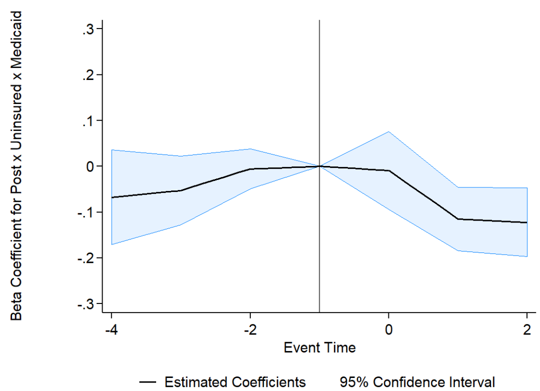

(c) 75%

As discussed earlier, in order for the estimated coefficients from the triple differences model in equation (3.1) to be interpreted as causal, parallel pre-trends are required. We can test this assumption using the event study estimates for coefficients before implementation. Coefficient estimates for the non-Medicaid portion of the ACA do not exhibit a pre-existing positive trend, and perhaps demonstrate a slight downward trend that reverses at the time of policy implementation. The higher the percentile cutoff used to define $Uninsured2013_s$, the less noticeable the pre-trend and the closer to 0 the pre-reform estimates. Estimates for the additional effect of the Medicaid expansion reported in Figure 3.6 show a similar pattern in the opposite direction.

3.6. Robustness and Extensions

In Table 3.3, we explore sensitivity of our results to several empirical decisions. The first row of Table 3.3 replicates the implied effects from estimation of equation (3.1), our baseline model, for reference. The second row repeats this estimation using a logged dependent variable. This results in the implied effects being interpreted as a semi-elasticity.

If there was anticipation of the upcoming ACA changes on the demand side, we could expect selection into or out of health insurance in the year prior, making the the 2013 health uninsurance rate endogenous to the timing of the ACA expansion and a poor measure of the policy exposure. In row 3, we replace the 2013 rate of uninsurance with the 2012 rate of uninsurance in equation (3.1), yielding similar effects to our main results.

In the next two rows we alternately drop states that expanded their Medicaid programs prior to January 1, 2014 (CA, CO, CT, DC, MN, MO, NJ, and WA) and those which expanded after January 1, 2014, but still within our sample time frame (AK, IN, LA, MI, MT, NH, and PA). The results for the non-Medicaid portion of the ACA are robust to the exclusion of these states, while our estimate of the Medicaid portion of the ACA is similar in magnitude and retains statistical significance when dropping late expansion states from the analysis.

A recent literature has emerged which notes that difference-in-differences estimators, and their corresponding event studies, can yield biased estimates when there is variation in treatment timing and heterogeneous policy effects.[24] A major concern being that previously treated groups are used as counterfactual control groups for later treated observations, an inappropriate comparison. In the next row, we report results from a model where we eliminate both early and late expanding states from the analysis, removing any differential treatment timing and yielding comparable results.

Finally, in the bottom two rows we test the sensitivity of the results to sample changes related to the auto insurance market. First, we drop no-fault states from the analysis given that bodily injury liability coverage operates only in excess of the underlying personal injury protection coverage. Here, we estimate qualitatively similar coefficients to the main analysis but lose statistical power

[24]See for example: Callaway and Sant'Anna (2021), Goodman-Bacon (2021), and Sun and Abraham (2021)

Table 3.3. Effect of ACA on Bodily Injury Claim Frequency–By cutoff

	ACA w/o Medicaid			Medicaid Expansion		
	25%	50%	75%	25%	50%	75%
Base Results	0.032	0.047*	0.080**	-0.020	-0.030	-0.049
	(0.020)	(0.027)	(0.035)	(0.037)	(0.034)	(0.048)
Log BI Frequency	0.008	0.021	0.047*	0.008	-0.008	-0.034
	(0.014)	(0.023)	(0.025)	(0.028)	(0.028)	(0.036)
2012 Uninsured Rate	0.032	0.053*	0.094***	-0.020	-0.036	-0.059
	(0.020)	(0.028)	(0.033)	(0.037)	(0.035)	(0.045)
Drop Early Expansion States	0.022	0.040	0.067**	0.021	-0.029	-0.010
	(0.018)	(0.029)	(0.030)	(0.043)	(0.036)	(0.049)
Drop Late Expansion States	0.030	0.042	0.079**	-0.054	-0.054	-0.086**
	(0.019)	(0.027)	(0.034)	(0.040)	(0.035)	(0.041)
Drop Late & Early Expansion States	0.019	0.033	0.066**	-0.016	-0.058	-0.050
	(0.019)	(0.028)	(0.030)	(0.058)	(0.041)	(0.040)
Drop No-Fault States	0.022	0.030	0.062	0.001	-0.011	-0.018
	(0.023)	(0.032)	(0.050)	(0.051)	(0.049)	(0.069)
Voluntary Market	0.033	0.050*	0.080**	-0.027	-0.036	-0.052
	(0.020)	(0.028)	(0.035)	(0.038)	(0.034)	(0.048)

Notes: This table reports coefficient estimates for variations on the main regression shown in equation (3.1). The estimates for the effect of the private and Medicaid expansion portions of the ACA on bodily injury frequency are shown separately. Columns represent the level of health uninsurance used to designate high uninsurance in the regression model, varying from 25% to 75%. The base results are the main estimates shown in Table 3.2. Specification 1 measures the outcome variable of bodily injury frequency in logs rather than levels. Specification 2 uses the 2012 health uninsurance rate to designate the high uninsurance groupings. Specification 3-5 alternatively drop states which expanded Medicaid programs before and/or after the 2014 ACA to account for variable timing. Specification 6 removes states which operate under a no-fault insurance structure while specification 7 removes claims from policies insured through a state's residual risk pool from consideration in the outcome variable, bodily injury frequency. Data are from the NAIC's annual *Auto Insurance Database* report for years 2010-2016. All models include a panel of driving and demographic controls and weighted by state population. Model standard errors are reported in parentheses and clustered at the state level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

due to the reduction in sample size. Lastly, our main results used frequency data for the entire auto insurance business in a state. As a way to provide liability insurance to high-risk drivers unable to find insurance in the insurance market, called the voluntary market, some states have established residual risk pools that guarantee coverage to any driver. We would expect different results from the original sample if auto insurance companies anticipated cost changes as a result of the ACA and rejected coverage for a different population. In the final row of Table 3.3, we find qualitatively similar results when restricting the BI frequency variable to those who find coverage through the voluntary auto insurance market, implying no discernible change in the rejection process of auto insurance companies correlated with the health insurance changes.

We next extend our analysis to the other main auto insurance coverages. In Table 3.4, we present regression estimates of our triple differences model with frequency and severity outcomes for property and collision coverages. For comparison purposes, we replicate the results from column 2 and 4 of Table 3.2 for BI coverage in the first column of each panel of Table 3.4.

Both PD and COLL coverages complement bodily injury as, most often, bodily injury is associated with an accompanying claim for damage to property.[25] The results for PD and COLL move in tandem with the results for BI: the direction of the estimates are always the same, they are statistically significant for similar percentile cutoffs, and both sets of estimates increase in magnitude as the percentile cutoff increases. While these estimates follow similar patterns, the aggregated nature of the data does not allow us to explicitly test if one coverage type is growing as an "add-on" to claims in other coverages.

The estimates for the effect of the Medicaid portion of the ACA on PD and COLL claim frequency are negative. This pattern is in-line with the findings of Heaton and Flint (2021) who find decreases in loss ratios due to the Medicaid expansion. Turning to the event time analyses for these coverages, Appendix Figures 3.7 and 3.8 plot estimates for PD while Appendix Figures 3.9 and 3.10 plot estimates for COLL frequency. These figures are slightly noisier than the figures for BI, but are still suggestive of a causal relationship for the analyses when using higher percentile cutoffs.

States above the median 2013 uninsurance rate had an average of 4.04 million incurred PD claims and an almost identical average of 4.04 million incurred COLL claims yearly, or approximately 155 thousand claims per state per coverage. Our estimate implies that the above median group had an additional 141 thousand PD claims (or approximately 5,423 additional claims per state) and an additional 173 thousand COLL claims (or approximately 6,654 additional claims per state) due to the non-Medicaid portion of the ACA. Given that the average severity of a PD claim in these states in 2013 was \$3,224, this implies an increase of approximately \$17.4 million in paid out PD claims per state (or approximately \$455 million for the entire above median group). Additionally, given that the average severity of a COLL claim in these states in 2013 was \$3,823, this implies an increase of approximately \$25 million in paid out COLL claims per state (or approximately \$661 million for the entire above median group).

[25]An example of a car crash with BI but not a corresponding claim for either PD or COLL would be if an insured driver hits a pedestrian.

Table 3.4. Effect of ACA on Claim Frequency and Severity by Cutoff and Coverage

	Frequency			Severity		
	(BI)	(PD)	(COLL)	(BI)	(PD)	(COLL)
25%						
Post x High Uninsured	0.032	0.080	0.077	1,280.8***	51.5	144.8***
	(0.020)	(0.051)	(0.061)	(248.6)	(35.4)	(40.9)
Post x High Uninsured x Medicaid Expansion	-0.020	-0.098	-0.157	-873.2*	51.2	-42.0
	(0.037)	(0.079)	(0.107)	(485.1)	(47.1)	(53.6)
50%						
Post x High Uninsured	0.047*	0.137**	0.247**	-47.4	48.0	87.9
	(0.027)	(0.066)	(0.102)	(655.7)	(74.1)	(77.6)
Post x High Uninsured x Medicaid Expansion	-0.030	-0.076	-0.150	-339.0	22.6	-40.1
	(0.034)	(0.077)	(0.114)	(725.2)	(78.0)	(87.3)
75%						
Post x High Uninsured	0.080**	0.216***	0.271***	559.1	45.7	175.7***
	(0.035)	(0.066)	(0.075)	(536.9)	(69.6)	(51.3)
Post x High Uninsured x Medicaid Expansion	-0.049	-0.151*	-0.088	-452.0	-5.8	-98.2
	(0.048)	(0.083)	(0.116)	(765.4)	(76.2)	(68.0)
Observations	350	350	350	350	350	350
Dependent variable pre-ACA mean	0.997	3.873	5.763	18,043	2,998	3,464

Notes: OLS estimates from equation (3.1) for the effect of the private and Medicaid expansion portions of the ACA are shown separately for bodily injury (BI), property damage (PD), and collision (COLL) frequencies and average severities. Results are grouped by the cutoff used to designate the high uninsurance level, varying from 25% to 75%. Standard errors, reported in parentheses, are clustered at the state level, all regressions are weighted by state population and include a panel of state-year control variables. Data are from the NAIC's annual Auto Insurance Database Report for year 2010-2016. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

3.7. Conclusion

In the decade following the initial implementation of the Affordable Care Act, a sizable literature has developed studying the direct health impacts from various portions of the legislation as well as unintended consequences caused by the law. Given the structure of automobile insurance and its coverage of physical injuries sustained during car accidents, auto insurance claims were plausibly impacted by changes to health behaviors due to the ACA such as the increase in the propensity to seek treatment.

Our findings suggest that i) such a linkage exists along the extensive margin for filing auto insurance claims but not the intensive margin of claim severity, and ii) a positive relationship manifests for the private market portions of the ACA changes, but not Medicaid insurance expansions. This is consistent with a story where individuals are more likely to seek medical care and then initiate a claim, but do not select into doing so based on claim severity.

The link between auto insurance and health insurance potentially generated sizable increases in auto insurance premiums, particularly in states which did not expand Medicaid. Based on our estimates, the non-Medicaid portion of the ACA was responsible for a $820 million increase in auto insurer outlays for bodily injury coverage claims. These states had roughly 100 million earned exposures (1 car insured for one year) in 2013 with an average bodily injury premium of $208, which yields an average increase of $8.19 per exposure or a 3.9 percent increase in the bodily injury portion of auto insurance premiums. If we consider property damage and collision coverages, then there was a total additional increase of $13.48 per exposure, or a 6.3 percent increase.

Overall, this represents a sizable monetary link between the two forms of insurance with reform to the health insurance sector generating spillovers into the auto insurance market. This is a secondary effect of health insurance reform that carries a non-trivial cost borne in another sector.

3.8. Appendix

FIGURE 3.7. Property Damage Frequency: Event Study By High Uninsurance Cut-off for Non-Expansion States

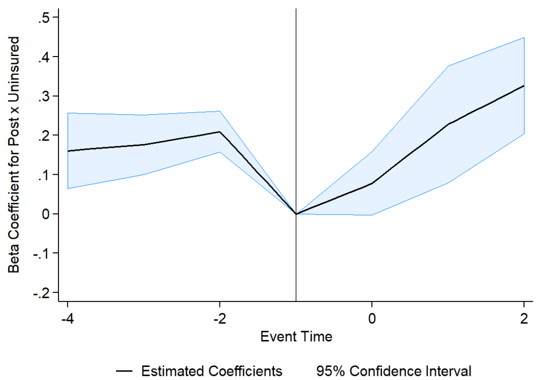

(a) 25%

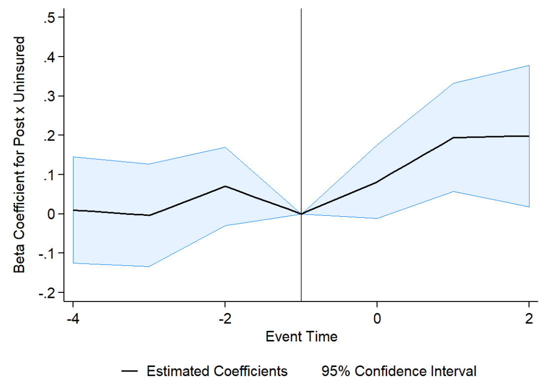

(b) 50%

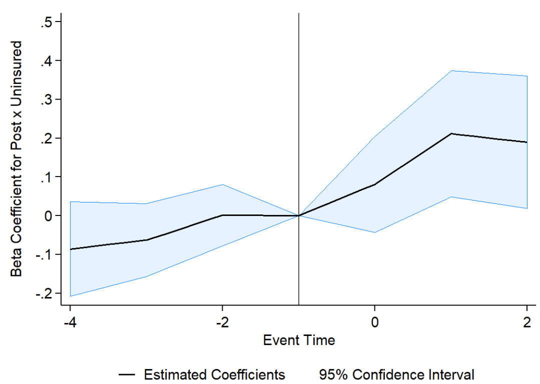

(c) 75%

FIGURE 3.8. Property Damage Frequency: Event Study By High Uninsurance Cut-off for Medicaid Expansion States

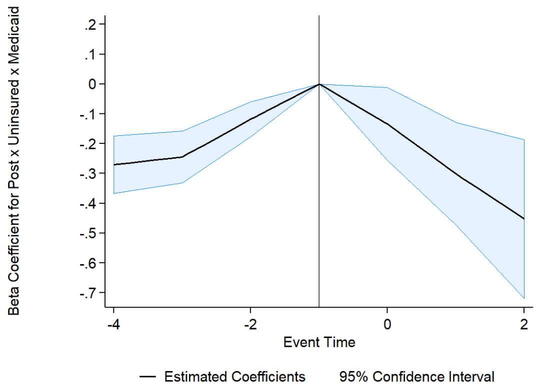

(a) 25%

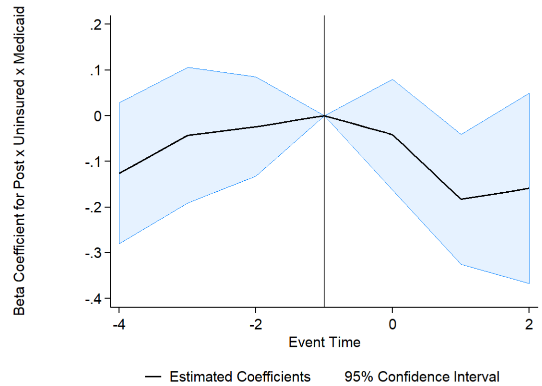

(b) 50%

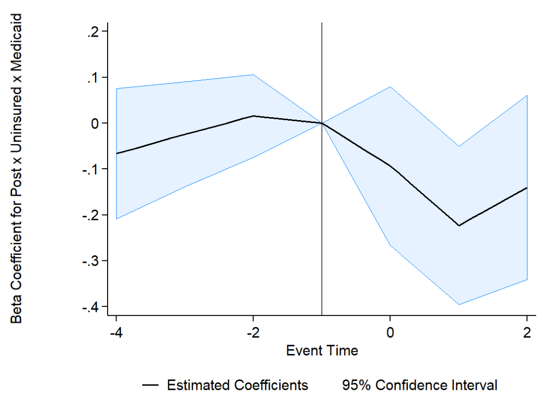

(c) 75%

Figure 3.9. Collision Frequency: Event Study By High Uninsurance Cutoff for Non-Expansion States

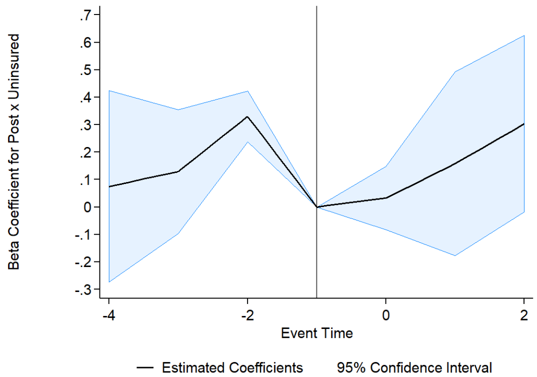

(a) 25%

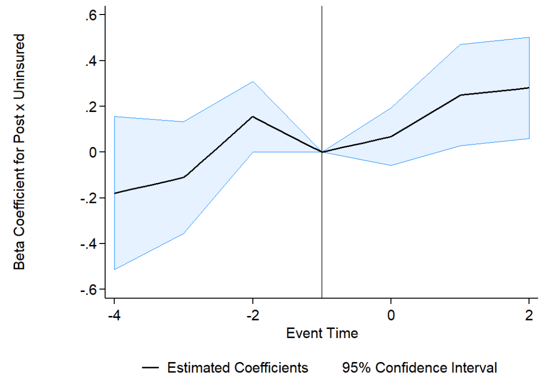

(b) 50%

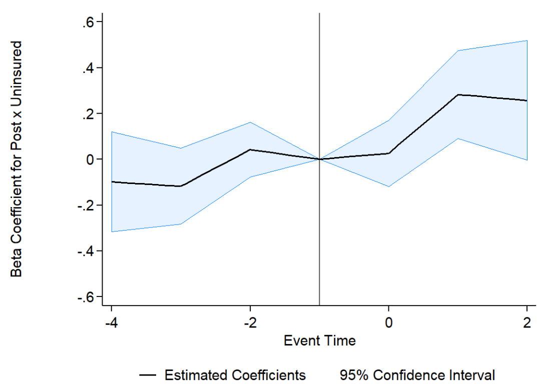

(c) 75%

FIGURE 3.10. Collision Frequency: Event Study By High Uninsurance Cutoff for Medicaid Expansion States

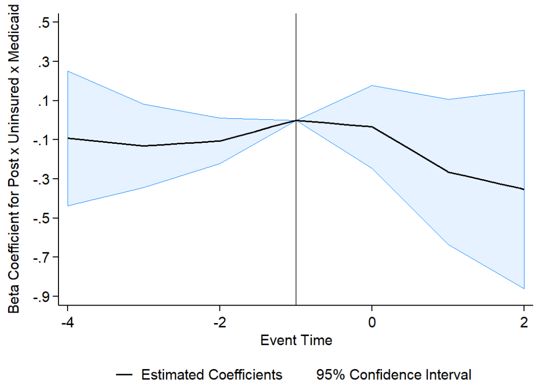

(a) 25%

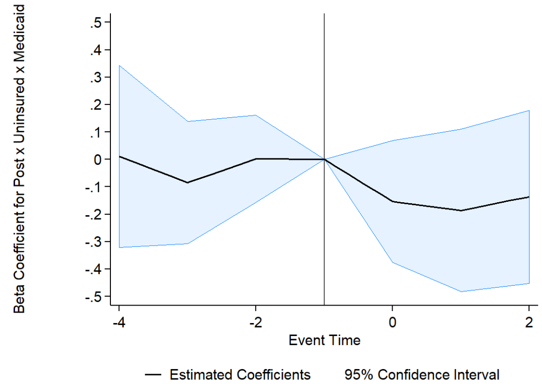

(b) 50%

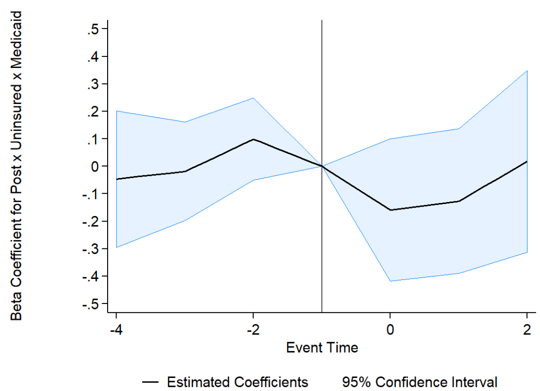

(c) 75%

Table 3.5. Bodily Injury Claim Frequency Results by Weights

	Frequency			Severity		
	(Pop)	(PD-EE)	(Drv Pop)	(Pop)	(PD-EE)	(Drv Pop)
25%						
Post x High Uninsured	0.032	0.030	0.033	1,280.768***	1,273.2***	1,271.3***
	(0.020)	(0.019)	(0.020)	(248.636)	(248.2)	(248.5)
Post x High Uninsured x Medicaid Expansion	-0.020	-0.015	-0.022	-873.159*	-940.9*	-851.1*
	(0.037)	(0.036)	(0.038)	‘(485.129)	(484.8)	(484.3)
50%						
Post x High Uninsured	0.047*	0.043	0.048*	-47.396	-154.0	-44.8
	(0.027)	(0.027)	(0.028)	(655.721)	(647.5)	(653.5)
Post x High Uninsured x Medicaid Expansion	-0.030	-0.026	-0.031	-339.004	-178.2	-336.7
	(0.034)	(0.034)	(0.034)	‘(725.180)	(732.1)	(722.5)
75%						
Post x High Uninsured	0.080**	0.076**	0.081**	559.056	532.5	561.6
	(0.035)	(0.034)	(0.035)	(536.930)	(537.9)	(537.6)
Post x High Uninsured x Medicaid Expansion	-0.049	-0.045	-0.050	-452.032	-348.8	-446.8
	(0.048)	(0.047)	(0.048)	‘(765.447)	(778.2)	(765.6)

Notes: OLS estimates for the effect of the private and Medicaid expansion portions of the ACA on bodily injury claim frequency and average claim severity are shown using various regression weighting techniques. Columns 1 and 4 show the main results from Table 3.2 and are weighted by state population. Columns 2 and 5 use Property Damage Earned Exposures for weights as this is the best representation for total vehicles covered under insurance policies throughout the year. Columns 3 and 6 weight by driving age population in a state (15-79). Data are from the NAIC's annual *Auto Insurance Database* report for years 2010-2016. Regressions include state-year level control variables. Model standard errors are reported in parentheses, clustered at the state level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Bibliography

Akosa Antwi, Y., A. S. Moriya, K. Simon, and B. D. Sommers (2015, jun). Changes in emergency department use among young adults after the patient protection and affordable care act's dependent coverage provision. *Annals of Emergency Medicine 65*(6), 664–672.e2.

Anderson, M., C. Dobkin, and T. Gross (2012, feb). The Effect of Health Insurance Coverage on the Use of Medical Services. *American Economic Journal: Economic Policy 4*(1), 1–27.

Arcidiacono, P., P. B. Ellickson, C. F. Mela, and J. D. Singleton (2020). The Competitive Effects of Entry. *American Economic Journal. Applied Economics 12*(3), 175–206.

Baicker, K., S. L. Taubman, H. L. Allen, M. Bernstein, J. H. Gruber, J. P. Newhouse, E. C. Schneider, B. J. Wright, A. M. Zaslavsky, and A. N. Finkelstein (2013, may). The Oregon Experiment — Effects of Medicaid on Clinical Outcomes. *New England Journal of Medicine 368*(18), 1713–1722.

Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. Shapiro (2016). Adapting to climate change: the remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy 124*.

Barron, J. M., B. A. Taylor, and J. R. Umbeck (2004, nov). Number of sellers, average prices, and price dispersion. *International Journal of Industrial Organization 22*(8-9), 1041–1066.

Bernardo, V. (2018, feb). The effect of entry restrictions on price: evidence from the retail gasoline market. *Journal of Regulatory Economics 53*(1), 75–99.

Bikales, J. (2020). California protects homeowners from having fire insurance dropped- again.

Bilal, A. and E. Rossi-Hansberg (2023, June). Anticipating climate change across the united states. Working Paper 31323, National Bureau of Economic Research.

Biswas, S., M. Hossain, and D. Zink (2023). California Wildfires, Property Damage, and Mortgage Repayment. *Working Paper*.

Boomhower, J., M. Fowlie, J. Gellman, and A. J. Plantinga (2023). How Are Insurance Markets Adapting to Climate Change? Risk Selection and Regulation in the Market for Homeowners

Insurance. *NBER Working paper*.

Borenstein, S. (2023). What's the Matter with California's Gasoline Prices?

Borenstein, S., J. Bushnell, and M. Lewis (2004). Market Power in California's Gasoline Market. *CSEM Working Paper 132*.

Born, P. and B. Klimaszewski-Blettner (2013). Should I stay or should I go? The impact of natural disasters and regulation on US property insurers' supply decisions. *Journal of Risk and Insurance 80*(1), 1–36.

Born, P. and K. Viscusi (2006). The catastrophic effects of natural disasters on insurance markets. *Journal of Risk and Uncertainty 33*, 55–72.

Botzen, W. J., O. Deschenes, and M. Sanders (2019). The economic impacts of natural disasters: A review of models and empirical studies. *Review of Environmental Economics and Policy 13*(2), 167–188.

Boyer, M. M., P. D. Donder, C. Fluet, M.-L. Leroux, and P.-C. Michaud (2020). Long-term care insurance: Information frictions and selection. *American Economic Journal: Economic Policy 12*(3), pp. 134–169.

Bradt, J. T., C. Kousky, and O. E. Wing (2021). Voluntary purchases and adverse selection in the market for flood insurance. *Journal of Environmental Economics and Management 110*(0), 102515.

Bresnahan, T. F. and P. C. Reiss (1991). Entry and Competition in Concentrated Markets. *Journal of Political Economy 99*(5), 977–1009.

Cabral, M. and M. Cullen (2019). Estimating the value of public insurance using complementary private insurance. *American Economic Journal: Economic Policy* (3), 88–129.

California Ballot Propositions and Initiatives (1988). Proposition 103: Insurance Rates, Regulation, Commissioner.

California Code of Regulations (2023). Title 10 § 2644.5.

California Code of Regulations (2024). Title 10 - investment, chapter 5 - insurance comissioner, subchapter 4.8 - review of rates, article 4 - determination of reasonable rates, section 2644.5 - catastrophe adjustment.

California Department of Insurance (2021). New data shows insurance companies non-renewed fewer homeowners in 2020 while fair plan 'insurer of last resort' policies increased. *2021 Press*

Releases.

Callaway, B. and P. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Card, D., C. Dobkin, and N. Maestas (2008, dec). The impact of nearly universal insurance coverage on health care utilization: Evidence from medicare. *American Economic Review 98*(5), 2242–2258.

Carranza, J. E., R. Clark, and J. F. Houde (2015). Price controls and market structure: Evidence from gasoline retail markets. *Journal of Industrial Economics 63*(1), 152–198.

Castro-Vincenzi, J. (2022). Climate Hazards and Resilience in the Global Car Industry. *Working paper*.

Courtemanche, C., A. I. Friedson, and D. I. Rees (2019, jun). Association of Ambulance Use in New York City with the Implementation of the Patient Protection and Affordable Care Act. *JAMA Network Open 2*(6), e196419–e196419.

Courtemanche, C., J. Marton, B. Ukert, A. Yelowitz, and D. Zapata (2017, jan). Early Impacts of the Affordable Care Act on Health Insurance Coverage in Medicaid Expansion and Non-Expansion States. *Journal of Policy Analysis and Management 36*(1), 178–210.

Courtemanche, C., J. Marton, B. Ukert, A. Yelowitz, D. Zapata, and I. Fazlul (2019, feb). The three-year impact of the Affordable Care Act on disparities in insurance coverage. *Health Services Research 54*(Suppl 1), 307.

Currie, J. and J. Gruber (1996, may). Health Insurance Eligibility, Utilization of Medical Care, and Child Health. *The Quarterly Journal of Economics 111*(2), 431–466.

Davis, L. W., S. McRae, and E. Seira (2023). The Competitive Effects of Entry in the Deregulated Mexican Gasoline Market. *Working Paper*.

de Chaisemartin, C. and X. D'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Decker, S. L. and B. J. Lipton (2017, aug). Most newly insured people in 2014 were long-term uninsured. *Health Affairs 36*(1), 16–20.

Diaz, D. and F. Moore (2017). Quantifying the economic risks of climate change. *Nature Climate Change 7*(11), 774–782.

Doyle, J. J. (2005, may). Health Insurance, Treatment and Outcomes: Using Auto Accidents as Health Shocks. *The Review of Economics and Statistics 87*(2), 256–270.

Dwyer, J. (1978). Fair plans: History, holtzman and the arson-for-profit hazard. *Fordham Urban Law Journal 7*(3), 617–648.

Einav, L. and A. Finkelstein (2023). Empirical analyses of selection and welfare in insurnace markets: a self-indulgent survey. *The Geneva Risk and Insurance Review 48*(4), 167–191.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *The Quarterly Journal of Economics 125*(3), 877–921.

Finkelstein, A. and J. Poterba (2014). Testing for Asymmetric Information Using "Unused Observables" in Insurance Markets: Evidence from the U.K. Annuity Market. *The Journal of Risk and Insurance*.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, and K. Baicker (2012, aug). The oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics 127*(3), 1057–1106.

Fischer, K., S. Martin, and P. Schmidt-Dengler (2023). The Heterogeneous Effects of Entry on Prices.

Fone, Z. S., A. I. Friedson, B. J. Lipton, and J. J. Sabia (2020). The Dependent Coverage Mandate Took a Bite out of Crime. *IZA Discussion Paper No. 12968*.

Frean, M., J. Gruber, and B. D. Sommers (2017, may). Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act. *Journal of Health Economics 53*, 72–86.

Gallagher, J. (2014). Learning about an infrequent event: Evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics 6*(3), 206–233.

González, X. and M. J. Moral (2023, sep). Competition and Competitors: Evidence from the Retail Fuel Market. *The Energy Journal 44*(01).

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

Gotanda, H., G. Kominski, and Y. Tsugawa (2020, mar). Association Between the ACA Medicaid Expansions and Primary Care and Emergency Department Use During the First 3 Years. *Journal of General Internal Medicine 35*(3), 711–718.

Gruber, J. (2011, mar). The impacts of the affordable care act: How reasonable are the projections? *National Tax Journal 64*(3), 893–908.

Gu, G. W. and G. Hale (2022). Climate Risks and FDI . *NBER Working paper*.

Hastings, J. S. (2004, mar). Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in Southern California. *American Economic Review 94*(1), 317–328.

Haucap, J., U. Heimeshoff, and M. Siekmann (2017). Fuel prices and station heterogeneity on retail gasoline markets.

Heaton, P. and C. Flint (2021, mar). Medicaid coverage expansions and liability insurance. *Journal of Risk and Insurance 88*(1), 29–51.

Issler, P., R. Stanton, C. Vergara-Alert, and N. Wallace (2024). Housing and Mortgage Markets with Climate Risk: Evidence from Wildfires in California. *SSRN Working Paper*.

Jhamb, J. and G. Colman (2015). The Patient Protection and Affordable Care Act and the Utilization of Health Care Services among Young Adults. *International Journal of Health and Economic Development 1*(1), 8–25.

Kadiyala, S. and P. Heaton (2017, dec). *The Effect of Health Insurance Coverage Expansions on Auto Liability Claims and Costs*. RAND Corporation.

Kaestner, R., B. Garrett, J. Chen, A. Gangopadhyaya, and C. Fleming (2017, jun). Effects of ACA Medicaid Expansions on Health Insurance Coverage and Labor Supply. *Journal of Policy Analysis and Management 36*(3), 608–642.

Kahn, M. E. (2021). Adapting to Climate Change: Markets and the Management of an Uncertain Future.

Knowles, S. G. and H. C. Kunreuther (2014). Troubled waters: The national flood insurance program in historical perspective. *Journal of Policy History 26*(3), 327–353.

Kousky, C. (2011). Managing Natural Catastrophe Risk: State Insurance Programs in the United States. *Review of Environmental Economics and Policy 105*(3).

Kousky, C. (2019). The role of natural disaster insurance in recovery and risk reduction. *Annual Review of Resource Economics 11*, 399–418.

Kousky, C. (2022). Understanding Disaster Insurance: New Tools for a More Resilient Future.

Kunreuther, H. (1996). Mitigating disaster losses through insurance. *Journal of Risk and Uncertainty 2*(12), 171–187.

Kunreuther, H. (2001). Mitigation and financial risk management for natural hazards. *The Geneva Papers on Risk and Insurance, Issues and Practice 2*(26), 277–296.

Lewis, M. S. (2015, sep). Odd Prices at Retail Gasoline Stations: Focal Point Pricing and Tacit Collusion. *Journal of Economics and Management Strategy 24*(3), 664–685.

Liao, Y., M. Walls, M. Wibbenmeyer, and S. Pesek (2022). Insurance Availability and Affordability under Increasing Wildfire Risk in California. *Resources for the Future*.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, A. Leibowitz, and M. S. Marquis (1987, jun). Health insurance and the demand for medical care: evidence from a randomized experiment. *American Economic Review 77*(3), 251–277.

Marcoux, K. and K. R. H. Wagner (2024). Fifty Years of U.S. Natural Disaster Insurance Policy. *Handbook of Insurance, forthcoming*.

Mazurenko, O., C. P. Balio, R. Agarwal, A. E. Carroll, and N. Menachemi (2018, jun). The effects of medicaid expansion under the ACA: A systematic review. *Health Affairs 37*(6), 944–950.

Miller, S. (2012a, dec). The effect of insurance on emergency room visits: An analysis of the 2006 Massachusetts health reform. *Journal of Public Economics 96*(11-12), 893–908.

Miller, S. (2012b, dec). The effect of the Massachusetts reform on health care utilization. *Inquiry (United States) 49*(4), 317–326.

Miller, S. and L. R. Wherry (2017, mar). Health and Access to Care during the First 2 Years of the ACA Medicaid Expansions. *New England Journal of Medicine 376*(10), 947–956.

Mulder, P. (2022). Mismeasuring Risk: The Welfare Effects of Flood Risk Information. *Working paper*.

Nikpay, S., S. Freedman, H. Levy, and T. Buchmueller (2017, aug). Effect of the Affordable Care Act Medicaid Expansion on Emergency Department Visits: Evidence From State-Level Emergency Department Databases. *Annals of Emergency Medicine 70*(2), 215–225.e6.

Nyce, C., R. Dumm, S. Sirmans, and G. Smersh (2015). The capitalization of insurance premiums in house prices. *Journal of Risk and Insurance 82*(4), 891–919.

Oh, S., I. Sen, and A.-M. Tenekedjieva (2023). Pricing of Climate Risk Insurance: Regulation and Cross-Subsidies. *SSRN Working paper*.

Peltzman, S. (1975, oct). The Effects of Automobile Safety Regulation. *Journal of Political Economy 83*(4), 677–725.

Prankatz, N. M. and C. Schiller (2021). Climate Change and Adaptation in Global Supply-Chain Network. *SSRN Working paper*.

Sastry, P. (2021). Who Bears Flood Risk? Evidence from Mortgage Markets in Florida. *Working Paper*.

Sastry, P., I. Sen, and A.-M. Tenekedjieva (2023). When Insurers Exit: Climate Losses, Fragile Insurers, and Mortgage Markets. *Working Paper*.

Simon, K., A. Soni, and J. Cawley (2017, mar). The Impact of Health Insurance on Preventive Care and Health Behaviors: Evidence from the First Two Years of the ACA Medicaid Expansions. *Journal of Policy Analysis and Management 36*(2), 390–417.

Sommers, B. D., M. Z. Gunja, K. Finegold, and T. Musco (2015, jul). Changes in self-reported insurance coverage, access to care, and health under the Affordable Care Act. *JAMA - Journal of the American Medical Association 314*(4), 366–374.

Spinnewijn, J. (2017, February). Heterogeneity, demand for insurance, and adverse selection. *American Economic Journal: Economic Policy 9*(1), 308–43.

State of California Department of Insurance (2023). Invitation to Workshop examining catastrophe modeling and insurance.

Sun, l. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Taylor, C. T., N. M. Kreisle, and P. R. Zimmerman (2010, jun). Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in southern california: Comment.

Wagner, K. R. H. (2022). Adaptation and Adverse Selection in Markets for Natural Disaster Insurance. *American Economic Journal: Economic Policy 14*(3), 380–421.

Watkins, N. and R. Lee (2022). Use of Catastrophe Models in California Homeowners Ratemaking Formula.

Weill, J. A. (2023). Flood Risk Mapping and the Distributional Impacts of Climate Information. *FEDS Working Paper*.

Wherry, L. R. and S. Miller (2016, jun). Early coverage, access, utilization, and health effects associated with the affordable care act medicaid expansions: A quasi-experimental study. *Annals of Internal Medicine 164*(12), 795–803.

Wooldridge, J. M. (2021, aug). Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators. *SSRN Electronic Journal*.

Xudong, A., S. Gabriel, and N. Tzur-Ilan (2024). Extreme Wildfires, Distant Air Pollution, and Household Financial Health. *SSRN Working Paper*.