QUANTIFYING THE PREMIUM EXTERNALITY OF THE UNINSURED

Stephen (Teng) Sun
Stanford University

Constantine Yannelis
Stanford University

Abstract
In insurance markets, the uninsured can generate a negative externality on the insured, leading insurance companies to charge higher premia. Using a novel panel data set and a staggered policy change that introduces exogenous variation in the rate of uninsured drivers at the county level in California, we find that uninsured drivers lead to higher insurance premia: a 1 percentage point increase in the rate of uninsured drivers raises premia by roughly 1%. We calculate the monetary fine on the uninsured that would fully internalize the externality and conclude that actual fines in most US states are inefficiently low. (JEL: G22, H23)

Keywords: Insurance, externality, uninsured, Pigouvian tax.

1. Introduction

In some insurance markets uninsured individuals drive up costs for the insured, creating an externality. The National Association of Insurance Commissioners estimated that Americans spent $186 billion on automobile insurance premia in 2009, and roughly 15% of American drivers lack automobile insurance. The externality imposed by the uninsured has been a large part of the impetus behind ten US states passing "No Pay, No Play" legislation in the past decades, which restricts the ability of uninsured drivers to sue for damages following a collision as well as the large fines for driving without insurance seen in many European countries.

The aim of this paper is to estimate the size of the externality caused by uninsured drivers in the automobile insurance market and discuss the optimal policy response.[1] In

The editor in charge of this paper was Claudio Michelacci.

Acknowledgments: We especially wish to thank Caroline Hoxby for guidance and helpful comments. We also thank Nick Bloom, Tim Bresnahan, Liran Einav, Han Hong, Xing Li, Florian Scheuer, Stephen Terry, and seminar participants at Stanford, École des Mines and the University of Illinois at Urbana-Champaign for helpful comments. We thank the California Department of Insurance and in particular Luciano Gobbo for providing us with data and regulatory information which was crucial to the development of this project. Constantine Yannelis gratefully acknowledges the financial support of the Alexander S. Onassis Foundation.

E-mail: doubles@stanford.edu (Sun); yannelis@stanford.edu (Yannelis)

1. There are numerous other externalities arising from driving, for example the vehicle size externality studied by Anderson (2008) and Anderson and Auffhammer (2014).

this market, when a collision occurs and an uninsured individual is at fault, the insured individual will typically be compensated by his own policy. When the uninsured driver has insufficient resources to cover the cost of the damage they can declare bankruptcy, passing the costs of the crash on to the insurance company and finally onto insured drivers via higher premia. Despite the theoretical and policy interests behind this externality, there is relatively little empirical evidence on this topic. We find credible empirical evidence that this externality is present, and that a 1 percentage point increase in the rate of uninsured drivers raises premia by roughly 1%.

Estimating the size of the effect of the uninsured on premia unavoidably poses a substantial empirical challenge. In theory, the most prominent concern is reverse causality due to the endogeneity of the rate of the uninsured with respect to insurance premia. In practice, measurement error in the standard proxy for the rate of uninsured motorists and the correlation of unobservable factors with both premia and the uninsured drivers' rate will also bias estimates of the externality. Although the literature on insurance is large, empirical research on the effect of the uninsured on premia has been lacking due to the aforementioned challenges. Our paper attempts to fill this gap for the case of the automobile insurance market. Exploiting a novel data set and a policy change varied at the county level in California, we quantify the extent of this negative externality. We find that the fines on the insured that would fully internalize the externality in most US states are inefficiently low.

We exploit variation in the rate of uninsured drivers resulting from an exogenous policy change to identify the effect of uninsured drivers on insurance premia. Between 1999 and 2007 the California Low Cost Automobile Insurance (CLCA) Program was introduced in the state of California and rolled out sequentially on a county-by-county basis. The program provided low cost insurance policies to drivers who met certain income, driver history and vehicle value criteria. The introduction of the CLCA program, together with a media campaign in areas in which the program was in effect, resulted in an approximate 1-2 percentage point decrease in the rate of uninsured drivers. The staggered rollout of the program makes it possible to achieve credible identification of the causal effect of uninsured drivers on premia.

We compiled novel panel data for 58 counties in California between 2003 to 2007. Our main data set consists of insurance premium quotes collected by the California Department of Insurance from licensed insurers based on several hypothetical risks including demographic and driving characteristics, policy limits, location, and coverage availability. Each observation is an offer price for one of two typical insurance plans, for hypothetical consumers with specific observable demographics from an insurance company. The main variation of interest is geographic variation – at the zip code level – in insurance premia as automobile insurance companies collect zip codes from clients and vary prices accordingly.

The use of policy-driven variation in the rate of uninsured motorists along with new administrative data on insurance premia leads us to conclude that uninsured drivers raise premia for other drivers, as predicted by theory. We find that a 1 percentage point increase in the share of drivers who are uninsured leads to a 1% increase in premia. To illustrate, this implies that consumers could save approximately $500 annually if the county with the highest uninsured drivers' rate, 29%, sees its uninsured drivers rate fall to that of the county with lowest uninsured drivers rate, 9%.

We discuss the optimal corrective Pigouvian tax on uninsured drivers using a simple model informed by our empirical estimates. Given that uninsured individuals increase premia paid by insured individuals, the government can levy a fine or tax on the uninsured to try to capture the effect of the externality. We find that the optimal tax is $2,240, which forces uninsured drivers to fully pay for the externality. Given that enforcement is stochastic, this is substantially higher than current fines in the US, although in line with fines in some European countries such as France. Such a high fine, if enforced rigorously, would effectively eliminate uninsured drivers as purchasing insurance on the private market would be cheaper than paying the fine.

We conduct a battery of robustness checks and examine alternative explanations. Our results are robust to concerns regarding weak instruments as well as to dropping any wave of the CLCA program and controlling for a county-specific time trend. We also vary the definitions of our instrumental variables and obtain consistent results. Utilizing the eligibility requirement of the CLCA program, we discuss and reject alternative explanations such as increased competition or unobserved selection on crash risk and moral hazard. Other phenomena, such as the introduction of the CLCA program inducing insurance companies to lower prices to compete with subsidized plans, or unobserved selection on accident risk could potentially explain our results. The structure of the CLCA program allows us to rule out such alternative explanations. We are able to test these alternative hypotheses by restricting our sample to individuals ineligible for the CLCA program, and reject these explanations for the observed effects following the introduction of the CLCA program.

The paper is organized as follows. Section 2 presents a concise model based on prior literature. Section 3 describes the data, which to our knowledge has not been used in the economics literature. Section 4 discusses our estimation strategy, explaining how we use a policy change to overcome the endogeneity problem. Section 5 presents our main empirical results, in which we find evidence of an externality. The section then discusses Pigouvian taxation. Section 6 presents robustness checks and rules out alternative explanations for our results. Section 7 concludes.

2. Theory

In this section we discuss the theory behind the externality caused by uninsured drivers on auto insurance premia, and we illustrate the endogenous relationship between premia and uninsured drivers which creates difficulties in estimation. There are three primary sources of bias complicating estimation of the externality. First, there is reverse causality which has been the focus of much of prior theory literature. This would tend to bias estimates upwards. Second, there is measurement error in standard measures of the rate of uninsured motorists. This biases estimates towards zero. Finally, there are unobservable factors associated with both the rate of uninsured motorists and premia, which can bias estimates in any direction.

In Section 5 we will use the model as a framework to discuss the optimal policy response to uninsured drivers. The basic intuition behind the theory underlying reverse causality is straightforward. Typically when a driver is found at fault in a crash, the at-fault driver's insurance covers the cost of damages. However, when an uninsured or underinsured driver causes a crash the driver may not have sufficient resources to cover damages and can declare bankruptcy. In this case the damaged party will be forced either to pay expenses out of pocket or collect payment from his own insurance plan. Thus in an area with a higher proportion of uninsured drivers, insurance companies will charge higher premia to obtain a given rate of return.

The model draws heavily from Smith and Wright (1992) and Keeton and Kwerel (1984). Define an individual i with wealth w_i and probability of being involved in a crash π_i. The individual purchases liability insurance from firm j with uninsured motorist coverage that costs p_{ij}. Liability insurance, which is the minimum insurance coverage required by law in most US states, pays for damage incurred by the holder of the policy to other individuals. In our discussion we restrict the externality to passing through the uninsured motorist coverage channel; this is not the only mechanism through which the externality is present. Auto insurance policies list uninsured motorist premia separately, but this does not capture the full extent of the uninsured driver externality due to several factors. First, while an entire package may be actuarially fair, the individual components may not. Second, uninsured drivers may still incur costs on insured drivers through property and collision coverage as the uninsured motorist coverage in California covers only bodily injury. See Smith and Wright (1992) for further discussion.

An individual i who purchases insurance has a payoff of $w_i - p_{ij}$ if he is not involved in a crash or if he is involved in a crash with another driver and found not to be at fault. For simplicity and without loss of generality,[2] assume that an individual has an equal probability of being found at fault or not at fault in a crash and a crash always involves two cars. If an individual is involved in a crash and is found at fault, the individual must either pay for the damage incurred to his vehicle or declare bankruptcy, hence the individual's payoff is $\max\{w_i - p_{ij} - L_i^s, 0\}$ where L_i^s is the stochastic cost of damage incurred by either party equally from the crash. In this case, the insurance company covers the losses L_i^s of the other driver who is not at fault.[3] This event occurs with probability $\pi_i/2$. Thus an insured driver has expected utility, given utility function $U(.)$ with standard properties:

$$V_{ins}(p_{ij}, w_i) = U(w_i - p_{ij})\left(1 - \pi_i + \frac{\pi_i}{2}\right) + \mathbb{E}\left[U(\max\{w_i - p_{ij} - L_i^s, 0\})\right]\frac{\pi_i}{2}$$

2. With the notable exception of moral hazard, which we discuss in Section 6. Allowing for any other arbitrary probability of being at fault will not change the basic intuition and prediction of our model.

3. Note that since he holds a liability only policy, which pays for the damage done to the other individual's car, the insured driver must still pay for the damage to his own vehicle, L_i^s.

Let λ be the fraction of uninsured motorists in a market, and note that λ is a function of premia, since when premia are high few drivers will purchase insurance. For an uninsured driver, if no crash occurs, or if a crash occurs with an insured driver and the uninsured driver is not found at fault, the driver obtains payoff w_i. The probability of not being involved in a crash is $1-\pi_i$ and the probability of being involved in a crash with an insured driver and not being found at fault is $(\pi_i/2)(1-\lambda)$. The expected utility for an uninsured driver if involved in a crash and found at fault is similar to that of a driver with liability insurance, with the exception of never having paid a premium to an insurance company, and that the driver must pay for the other driver's losses, rather than the insurance company paying: $\max\{w_i - 2L_i^s, 0\}$. Finally, if an uninsured driver is involved in a crash with another uninsured driver who is at fault, the driver receives a payoff $\min\{w_i - L_i^s + R_i, w_i\}$, which occurs with probability $\lambda\pi_i/2$. We let R_i refer to the amount the driver recovers from the uninsured individual who caused the crash, which is random. Assuming a continuous, increasing and concave utility function $U(.)$, the total expected utility $V_{unins}(w_i)$ for the uninsured driver becomes

$$V_{unins}(w_i) = \mathbb{E}\left[U(w_i)(1-\pi_i + \frac{\pi_i}{2}(1-\lambda))\right] + \mathbb{E}\left[U(\max\{w_i - 2L_i^s, 0\})\right]\frac{\pi_i}{2} + \mathbb{E}\left[U(\min\{w_i - L_i^s + R_i, w_i\})\right]\lambda\frac{\pi_i}{2}$$

A driver will choose to insure if $V_{ins}(p_{ij}, w_i) \geq V_{unins}(w_i)$. As we would expect, a driver is less likely to choose to insure when his premium is higher. Thus λ, the rate of uninsured drivers, is increasing in the premium p_{ij}. This property leads to simultaneity bias which, as we will see, presents significant empirical challenges to estimating the effect of uninsured drivers on insurance premia.

We assume a representative risk-neutral firm in a competitive insurance market and we compute the actuarially fair premium by equating revenues with expected indemnities, which amount to the expected liability loss from an insured driver as well as the expected loss from being involved in a crash with an uninsured driver who declares bankruptcy. We thus have

$$p_{ij} = \mathbb{E}\left[(\max\{L_i^s - R_i, 0\}\lambda + L_i^s)\frac{\pi_i}{2}\right].$$

Assuming that crash rates of the policy holder are a function of demographics X_i, we have $(\pi_i/2)(\mathbb{E}[L_i^s]) = X_i'\gamma$. Define

$$\beta_i = \mathbb{E}\left[\max\{\frac{L_i^s - R_i}{\mathbb{E}[L_i^s]}, 0\}X_i'\gamma\right] \geq 0$$

and we have the following equation for the premium that individual i pays to firm j

$$p_{ij} = \beta_i\lambda + X_i'\gamma.$$

The premia charged by the insurance company are weakly increasing in λ, the rate of uninsured drivers. Hence, ceteris paribus we would expect an area with a higher rate of uninsured drivers to have higher insurance premia. At the same time λ is increasing in p_{ij} as higher premia will cause fewer drivers to insure. This is our first major source of bias, as an area with high premia for reasons unrelated to the rate of uninsured drivers could also have a high rate of uninsured drivers. Another source of bias stems from the fact that λ is unobserved. Rather than observing the rate of uninsured motorists directly, researchers observe λ^*, an estimate of the rate of uninsured motorists coming from claims data. Since crashes are quite infrequent, this induces measurement error which will attenuate estimates of β. The final concern is that unobservable factors in X_i' may be correlated with both λ and p_{ij}. This could bias estimates in any direction, however in our case most of these factors would tend to bias estimates downwards. For example, it is likely that both the rate of uninsured motorists and the average cost of vehicles are correlated with local economic conditions. If local economic conditions deteriorate, fewer drivers will choose to insure, lowering insurance premia for all drivers. These endogeneity problems make it difficult to estimate the true effect of λ on p_{ij}, since λ will be significantly correlated with the error term in any regression. In the next section, we discuss how we overcome the endogeneity problem and estimate the true effect of uninsured drivers on insurance premia.

3. Institutional Background and Data

3.1. The CLCA Program

California mandates, as do other US states with the exception of New Hampshire, that drivers purchase basic liability automobile insurance. In California the basic liability insurance required by law consists of $15,000 of bodily injury insurance per individual, $30,000 of total bodily injury insurance per crash, and $5,000 of property damage insurance per crash. Despite the mandates, many drivers remain uninsured. For instance, in 1997, the Department of Insurance estimated 20.12 % of California drivers were uninsured. To reduce the share of drivers who are uninsured, California introduced the California Low Cost Automobile Insurance program (CLCA) in 1999, starting with two pilot counties: San Francisco and Los Angeles. CLCA offers basic liability insurance to eligible low-income individuals who live in California counties where the program is active. Rates under the CLCA program are set annually at the county level by the California Automobile Assigned Risk Plan (CAARP) commissioner. They are set well below rates for plans available in the market.[4] The rates set by CAARP are intended to cover the administrative costs of the program but not to allow insurance companies to make a profit. Premia are not subsidized by the government, and policyholders are assigned to insurance firms based on their share of the voluntary auto insurance market in each county. Eligibility for the program was

4. CLCA coverage is lower than the minimum required insurance coverage for holders of normal private automobile insurance plans.

determined by two main factors, income and a vehicle value threshold.[5] We do not observe income, as it is illegal in California for automobile insurers to price on income, however we do observe vehicle value.

The CLCA program was instituted in two pilot counties in 1999, and then expanded across the state in five different waves between April 2006 and December 2007. The introduction of the CLCA program was coupled with intense media campaigns in counties that were thought to be underserved or have a high proportion of uninsured drivers by the Department of Insurance. Advertisements were put out via print, radio, cable television, community organizations and government agencies. This media campaign about the legal requirement for carrying insurance would likely have had a second effect in decreasing the rate of uninsured drivers, as well as the primary effect of decreasing uninsured drivers via insurance plans under the CLCA program. Figure 1 illustrates the expansion of the CLCA program via waves between 1999 and 2007.

The CLCA program was introduced to different counties in California based on determination of need, which was interpreted as the number of uninsured drivers in a county between 1998 and 2007. For more details on the implementation of the CLCA program consult Schultz and Yarber (2006). The number of uninsured drivers depends largely on the size of counties rather than the rate of uninsured drivers. County borders are somewhat arbitrary, and the population size of California counties varies drastically while the rate of uninsured drivers, which is the driving force behind the externality, does not vary as much, ranging from 9% to 29%. Effectively, this means that CLCA program waves were assigned by the population of counties. Figure 2 illustrates the means of certain key variables of counties across county waves. There is a clear declining trend in population across the five waves, while other variables such as crash rates, rates of uninsured drivers, and premia are close to being identical. The exception to this rule is in the final wave, where the results are affected by several small counties in the Sierra Nevada mountains which have a very high measured crash rate:[6] Alpine, Placer, Nevada, El Dorado and Sierra.

3.2. Data

Our main dataset, which to our knowledge has not been used in the economics literature, comes from the California Department of Insurance. Following January 1, 1990, California law[7] required that the California Department of Insurance collect data on insurance rates in the state. The Department of Insurance ran the Automobile Premium Survey (APS) which collected data on automobile insurance premia from

5. See Appendix C.1 for a further discussion on eligibility and the CLCA program.

6. The sharp spike in crash rates likely represents that the crash rate is measured by the number of injury crashes over the total number of vehicles in a county. For the final wave this measure reports implausible crash rates several times higher than those of other counties. The Lake Tahoe region is a popular tourist destination, and it is likely that the high measured crash rates reflect tourists getting into crashes in counties with very low numbers of registered vehicles.

7. Specifically, the California Insurance Code Section 12959.

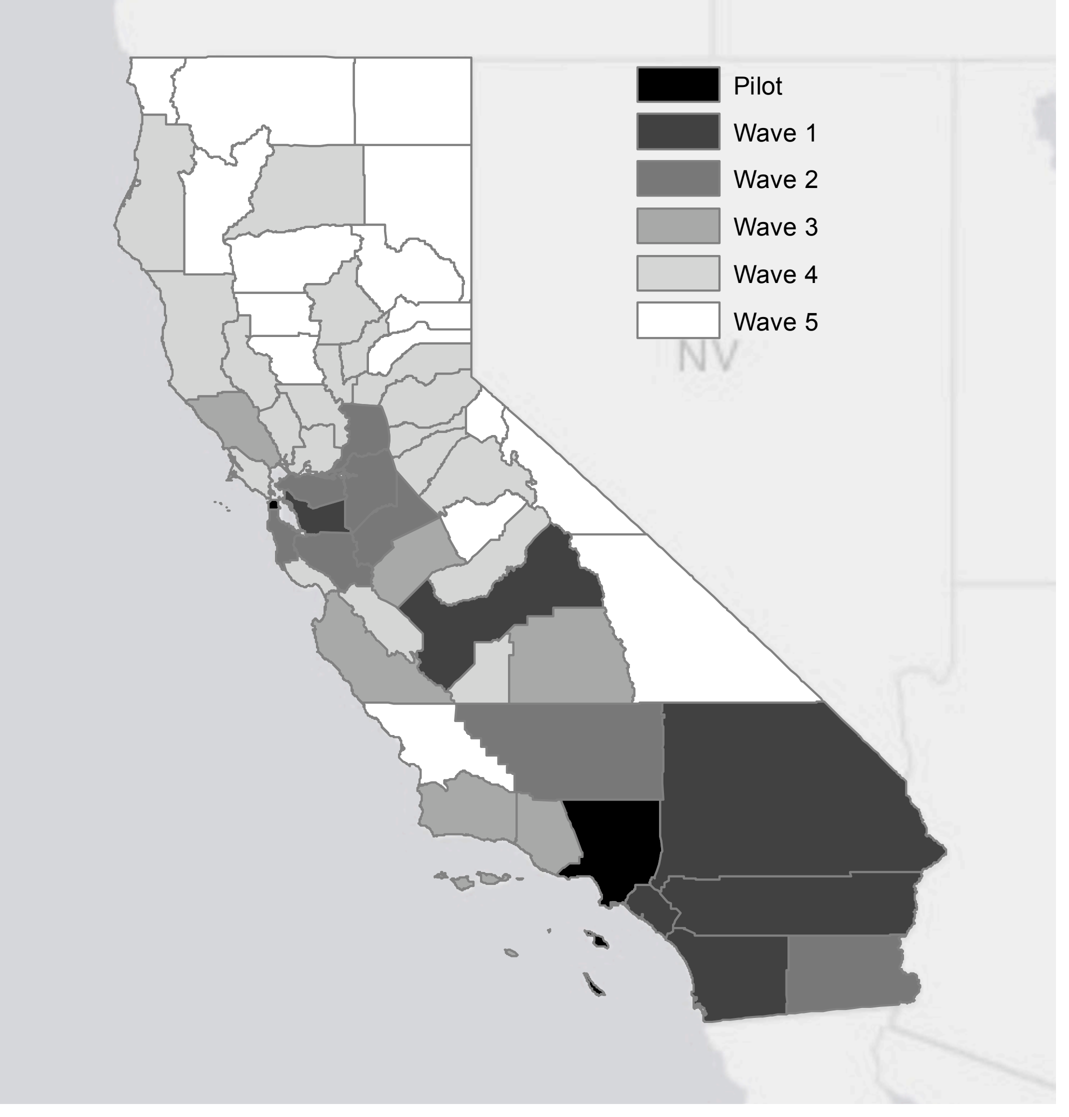

FIGURE 1. CLCA program waves.
Pilot Counties (1999)- Los Angeles and San Francisco.
Wave 1 (April 1, 2006)-Alameda, Fresno, Orange, Riverside, San Bernardino, San Diego.
Wave 2 (June 1, 2006)-Contra Costa, Imperial, Kern, Sacramento, San Joaquin, San Mateo, Santa Clara, Stanislaus.
Wave 3 (March 30, 2007)- Merced, Monterey, Santa Barbara, Sonoma, Tulare, Ventura.
Wave 4 (October 1, 2007)-Amador, Butte, Calaveras, El Dorado, Humboldt, Kings, Lake, Madera, Marin, Mendocino, Napa, Placer, San Benito, Santa Cruz, Shasta, Solano, Sutter, Tuolumne, Yolo, Yuba.
Wave 5 (December 10, 2007)-Alpine, Colusa, Del Norte, Glenn, Inyo, Lassen, Mariposa, Modoc, Mono, Nevada, Plumas, San Luis Obispo, Sierra, Siskiyou, Tehama, Trinity.

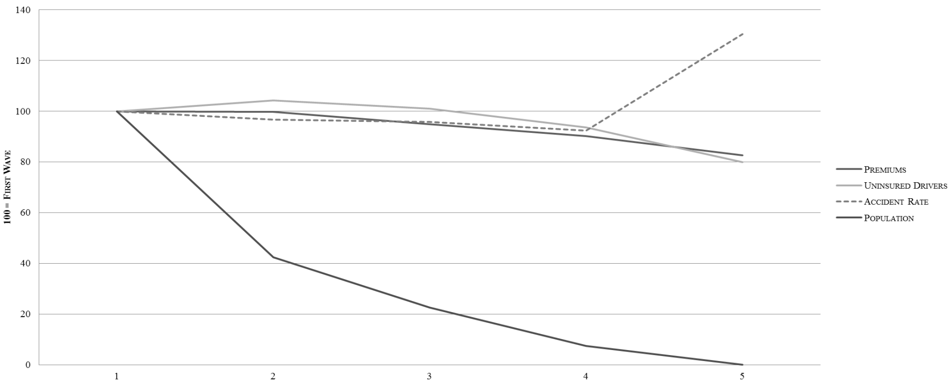

FIGURE 2. CLCA waves assigned by population. This figure plots average inflation adjusted premia, the rate of uninsured drivers, crash rates and population across each wave of the implementation of the CLCA program. The spike in crash rates in the fifth wave is driven by counties in the Lake Tahoe region. Including these counties does not change the results significantly. The CLCA program was effectively assigned by population size so we see a clear decreasing trend in population size across CLCA waves, while we do not see significant differences in other variables.

insurers licensed to provide automobile insurance in local California zip codes based on several hypothetical risks including demographic and driving characteristics, policy limits, location and coverage availability. Each observation represents an offer price for consumers with particular observable demographics from a firm operating in a particular zip code providing coverage beginning January 1st until December 31st of the year surveyed. Auto insurance pricing is heavily regulated in California, and insurers must charge customers prices based on formulae registered with the state Department of Insurance. The survey oversampled hypothetical drivers with speeding tickets and at fault crashes, leading to a higher average premium in comparison to the general populace. The premium survey data is available from 2003 to 2010 and our final sample is from 2003 to 2007, matching the uninsured drivers rate data, which is available from 2002 to 2007 due to significant delays in the release of the uninsured drivers' rate.[8]

The database consists of several million observations, the main variable of interest being the annual premium for an automobile insurance plan. The observations are indexed by zip codes, allowing the researcher to match the database to county-level data. The APS database also contains data on National Association of Insurance Commissioner (NAIC) codes of insurers as well as vehicle make and year, which we matched to vehicle value using pricing information.[9] The APS collected data on two types of plans from licensed insurers in zip codes, a basic plan and a standard coverage plan for different demographics. The basic plan represents a plan just above

8. The 2008 APS survey was not conducted for administrative reasons. The 2004 survey was not conducted statewide.

9. The website Auto Loan Daily was used as the source for vehicle values.

TABLE 1. Summary statistics.

	Mean of Analysis Variable					
Wave	Pilot	1	2	3	4	5
Premium	3167.96	2224.57	2225.00	2092.12	1992.88	1821.96
Uninsured Rate	20.66	20.79	22.13	22.09	23.53	18.62
Crash Rate	1.21	.87	.85	.82	.79	.96
At Fault	.47	.47	.47	.47	.47	.47
Standard Plan	.77	.75	.75	.75	.75	.75
Age	30.02	30.01	30.00	30.01	30.01	30.01
Mile Driven per Day	12.44	12.44	12.44	12.44	12.44	12.44
Speeding	.47	.47	.47	.47	.47	.47
Female	.47	.43	.42	.43	.42	.42
N	1,744,022	1,405,293	518,616	335,310	496,198	224,781

Notes: The table provides means of the analysis variables. The data source for uninsured motorists is the California Department of Insurance. The data source for the number of crashes is the California Highway Patrol. The data source for all other variables is the California Department of Insurance Automobile Premium Survey.

TABLE 2. Automobile insurance plan coverage.

	Basic Coverage	Standard Coverage	CLCA Plan
Bodily Injury	\$15,000/\$30,000	\$100,000/\$300,000	\$15,000/\$20,000
Property Damage	\$5,000	\$50,000	\$3,000
Medical Payments	\$2,000	\$5,000	–
Uninsured Motorist Bodily Injury	\$15,000/\$30,000	\$30,000/\$60,000	–
Comprehensive Deductible	–	\$250	–

Notes: Bodily Injury (BI) claims are the maximum that an insurance company will pay per person and the maximum an insurance company will pay for injuries from a specific crash. Uninsured Motorist Bodily Injury (UMBI) claims are the maximum that an insurance company will pay per person and the maximum an insurance company will pay for injuries from a specific crash where an uninsured motorist is at fault. California law mandates BI and property damage coverage according to the basic liability-only policy.

the minimum required threshold for coverage in California, while the standard plan was deemed by the Department of Insurance to be the most common automobile insurance plan in California. Table 2 summarizes the two types of private plans and the basic CLCA plan.

The APS survey data contains age, gender, the number of years an individual possessed a license, the number of miles an individual drives to work daily, the number of miles an individual drives in a year, the number of persons covered under a plan, the types of vehicles covered under the plan, the number of speeding tickets a hypothetical individual received in the three years prior to the survey date, and the number of at-fault automobile crashes in which an individual was involved in the three years prior to the survey. Since premia were unadjusted for inflation, we collected data on the Consumer Price Index from the Bureau of Labor Statistics using the BLS December CPI of each year in our adjustments.

The main APS survey data was matched to three other data sources which were obtained from the California Department of Insurance, the California Highway Patrol Integrated Traffic Records System, and the US Census Small Area Estimates Branch. Whether or not the CLCA program was in effect in various counties as well as premium rates was obtained from the California Department of Insurance 2011 Report to the Legislature. The matched data provided crucial time-varying geographic information that was not in the survey data. We matched data from our sample premium database to zip code level data from California using various sources. Zip code level data on uninsured bodily injury claims and all bodily injury claims was also obtained from the California Department of Insurance between 2002 and 2007. We used this data to construct a measure of uninsured drivers following Smith and Wright (1992) and Cohen and Dehejia (2004).[10] For each zip code, we use the average rate of uninsured motorists in zip codes within a 25 mile (40km) radius of the zip code area.[11]

To construct our measure of crash rates, county level data on injuries and fatalities resulting from automobile collisions was obtained from the California Highway Patrol. Since 2002, the California Statewide Integrated Traffic Records System has provided a database of information on monthly traffic collisions in California counties. The system provides data on all reported fatal and injury collisions occurring on public roads in California. The data is compiled from local police and sheriff jurisdictions and California Highway Patrol field offices. We used this data, and data on the total number of exposures and percentage of uninsured motorists from the Department of Insurance, to compute the injury collision and fatality collision rates in various California counties by taking the number of injury crashes over the number of registered vehicles.

4. Empirical Strategy

We exploit the staggered introduction of the CLCA program to generate variation in the rate of uninsured drivers, and the effect of the uninsured on premia is identified under two assumptions. The first assumption is that the instrumental variables are correlated with the rate of uninsured drivers, which is supported in results presented in Section 5.1. The second assumption is that the instrumental variables are orthogonal to unobserved determinants of insurance premia. Thus the identifying assumption for our empirical strategy is that, had it not been for the introduction of the CLCA program, there would have been no differential conditional on changes in the insurance premia across California counties in different waves over our sample period. It is important to note given that we control for year and zip-code fixed effects, any confounding factor must be a systematic time-varying zip-code-specific change that coincides with our observed trend in insurance premia.

10. See Appendix A for more on estimating the rate of uninsured drivers. Our measure is the number of Uninsured Motorist Bodily Injury claims divided by the number of Bodily Injury claims in a given zip code. This measure will be identical to the rate of uninsured motorists given two plausible assumptions, one, the probability of being involved in a crash is the same for both insured and uninsured motorists and two, in crashes between insured and uninsured motorists each party is equally likely to be found at fault.

11. According to the Bureau of Transportation Statistics (2006), this is roughly the number of kilometers that an average Californian drives per day. The main results are robust to varying the uninsured motorist zip code region. We use a standard equirectangular approximation to compute distance.

While our identifying assumption cannot be tested directly, Figure 3 provides further support that there was no significant pre-existing trend in the insurance premia across different CLCA program waves. Figure 3 shows wave-by-year fixed effects from regressing premia on controls for individual, geographic, temporal and vehicle controls. None of the fixed effects are significant at the 5% level, and there do not appear to be significant differences in the waves conditional on observables. The figure also provides graphical evidence for our hypothesis that the CLCA program reduced the rate of uninsured drivers, thereby reducing automobile insurance premia. In 2006, when the CLCA program begins, we see a sharp drop in premia for the first two waves, where the CLCA program took effect. We also plot the average annual rate of uninsured drivers by the number of years before and after the introduction of the CLCA program in Figure 4. This visual illustration makes the case that there is no clear declining pretrend in the rate of uninsured drivers before the implementation of the CLCA program. We conclude that examining the dynamic variation of both the insurance premia and the rate of uninsured drivers combined with the specific timing of CLCA waves provide strong support for our identifying assumption.

Our instruments from the CLCA program include the following: the average number of months during the year in which the CLCA program was active in a zip code cluster; the square of the previous term; and a dummy of whether the CLCA program was active for more than four months in a zip code cluster. We use the number of months during the year in which the CLCA program was active since typically the CLCA program was introduced in the middle of a year, and we avoid any arbitrary cutoffs associated with an indicator variable of whether or not the CLCA program was in effect. The results are robust if instead we use an indicator of whether or not the CLCA program was in effect for the entire year, or an indicator of whether or not the CLCA program was in effect for any part of the year.

The CLCA program being in effect is associated with a drop in the rate of uninsured drivers due to both the direct effect of uninsured drivers entering the program and through the media campaign associated with the introduction of the program. It is also highly plausible that the introduction of the CLCA program was exogenous to insurance premia in a county. California government documents regarding the introduction and expansion of the CLCA program do not make any mention of premia being used as a determinant of where the CLCA program was introduced, and from Figure 2 it appears that the California government simply rolled out the program in counties with a larger population first. We also find that population is not a significant determinant of premia when we control for population, and our results are robust to including population in the specification. Furthermore, the rate of uninsured drivers varies much more within zip code clusters in counties as opposed to across counties.

We also include as an instrument the number of months the CLCA program is in effect squared. If the average number of months that the CLCA program is in effect is a valid instrument, the square of the instrument will always mechanically be a valid instrument. There is also an intuitive reason to include the square of the program as an instrument– we expect the effect of the CLCA program to be greater in geographic areas where the program has been active for more time.

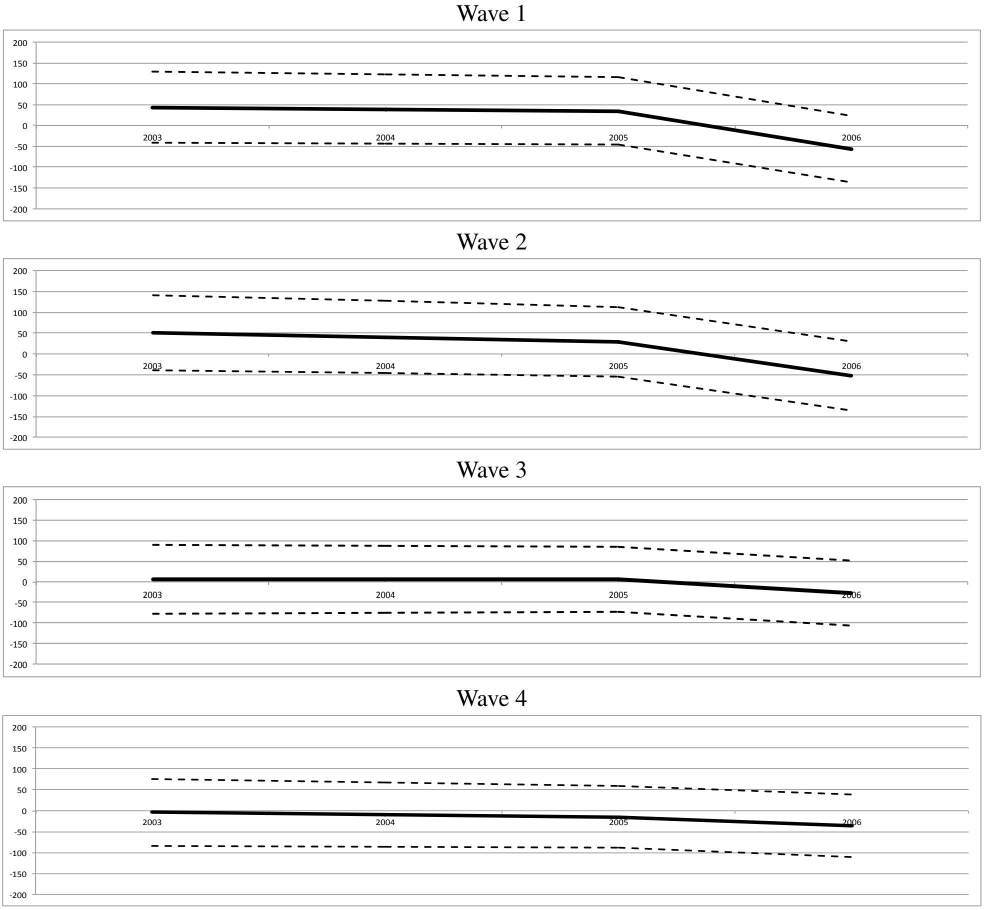

FIGURE 3. No significant pre-trend across waves. This figure plots the estimated difference (wave by year fixed effects, where the fifth wave is omitted to avoid multicollinearity) from a regression of premia on individual, geographic, temporal and vehicle controls. Confidence bands at the 95% level are included matching each line style. Note that in the first two waves, the CLCA program went into effect in 2006.

To implement the IV estimator, we first run the following regression (first stage):

$$\lambda_{gjit} = \alpha_g + \alpha_j + \alpha_t + \alpha_v + X_{it}^{'} b_1 + CLCA_{gt}^{'} b_2 + e_{gjit}, \tag{1}$$

where λ_{gjit} is the rate of uninsured drivers in geographic area g in which firm j offers an insurance premium to individual i at time t, $CLCA_{gt}^{'}$ is a vector consisting of our CLCA instruments, X_{it} is a vector of control variables and $\alpha_g, \alpha_j, \alpha_t$ and α_v are zip code, firm, year and vehicle fixed effects. We then estimate the second stage,

$$premium_{gjit} = \alpha_g + \alpha_j + \alpha_t + \alpha_v + X_{it}^{'}\gamma + \beta\hat{\lambda}_{gijt} + \varepsilon_{gijt}, \tag{2}$$

where $premium_{gjit}$ is the real (inflation-adjusted) premium offered in geographic area g by firm j to individual i at time t and $\hat{\lambda}_{gt}$ are predicted values of the rate of uninsured drivers from our first stage, equation (1). We use year fixed effects to control for any time-specific macro effects that shift the premium of automobile insurance in California. In our context, such macro effects could involve technological progress in automobiles that reduced loss in crashes or changes in the degree of competitiveness in automobile insurance markets that affect areas across California. We use zip code fixed effects to capture any unobserved zip code characteristics that are fixed over time, such as population characteristics, general weather conditions, traffic conditions, and any other bias associated with geographic characteristics. These zip code fixed effects are important for mitigating potential bias associated with the likely endogeneity of the rate of uninsured drivers. For example, the bias can arise from the fact that wealthier zip code areas have fewer uninsured drivers and tend to have higher insurance premia for reasons like price discrimination, which is difficult for the researcher to control directly. We also use company fixed effects to control for any time-invariant company-specific effects. For example, some firms may be more competitive and focus on thrift consumers while some firms charge higher premia for superior quality of service and brand capital. The vehicle fixed effects control for vehicle specific pricing factors, for example, more expensive vehicles may be more expensive to insure. We define the vehicle fixed effects by brand and model, and all results are robust to specifying the vehicle fixed effects by brand, model and year. Our coefficient of interest is β, which we interpret as the average effect of a 1 percentage point increase in the rate of uninsured drivers on the average premium. It is important to mention the caveat that our estimates are local. It is likely that there are nonconstant average effects in how uninsured drivers affect insurance premia. The average rate of uninsured drivers in California during our time period is 20.6%, with a standard deviation of 4%.

It is worth noting that while the different waves of the CLCA program were introduced by county, for our purpose, the minimal source of variation in the rate of uninsured drivers affected by this program is at what we call zip code cluster level, since a driver living in a border zip code could drive to other nearby zip codes of a different county. We define a zip code cluster as zip codes within a 25 mile (40 km) radius of the zip code under consideration. We choose 25 miles as this is the average number of miles a Californian drove daily in 2007. Our rates of uninsured drivers as well as the instruments are all averages of the raw variables within the zip code cluster. This feature also justifies why we are clustering at zip code cluster level instead of county level.

5. Main Results

5.1. Effects of CLCA on the Rate of Uninsured Drivers

Before showing our main results, we first examine how the introduction of the CLCA program affected the rate of uninsured drivers in California. We collapse the dependent

TABLE 3. Survey sample characteristics.

	No CLCA Mean	CLCA Mean	Difference	p-value	Observations
Female	.454	.427	.023	.260	4,724,220
	(.006)	(.020)			
Age	30.002	30.0276	-027	.946	4,724,220
	(.004)	(.006)			
Standard Plan	.763	.752	.011	.456	4,724,220
	(.003)	(.001)			
Crash Rate	.930	1.024	-.094	.077	4,724,220
	(.029)	(.095)			
Daily Miles Drive	12.446	12.437	.009	.202	4,724,220
	(.002)	(.004)			
At Fault Crash	.473	.473	-.001	.951	4,724,220
	(.003)	(.001)			
Speeding Ticket	.473	.472	-.001	.330	4,724,220
	(.004)	(.003)			

Notes: The first column presents the mean of the variable in the row before the CLCA program was active for at least four months. The second column presents the mean of the variable in the row after the CLCA program was active for at least four months. The third column is the difference. The fourth column is the p value from an F test for the hypothesis that these two means are the same from a regression with standard errors clustered at the county level. The final column is the number of observations.

TABLE 4. Effect of CLCA program on the rate of uninsured drivers.

	(1)	(2)	(3)	(4)
CLCA Program	-1.715***			-1.069***
	(0.290)			(0.387)
Months CLCA		-0.204***		0.164
		(0.0334)		(0.121)
Months CLCA^2			-0.0166***	-0.0148**
			(0.00231)	(0.00709)
Observations	995	995	995	995

Notes: * $p < .1$, ** $p < .05$, *** $p < .01$. The dependent variable in all columns is the rate of uninsured drivers in a zip-code cluster, measured by UMBI/BI. The independent variable in specification (1) is an indicator of whether or not the CLCA program was in effect in the zip code cluster for more than a third of the year. The independent variable in specification (2) is the average number of months the CLCA program is active in a 25 mile radius around the zip code where the premium quote is located. The independent variable in specification (3) is the average number of months the CLCA program is active in a 25 mile radius around the zip code where the premium quote is located squared. Column (4) shows the coefficients of the instruments in the first stage of the two stage least square estimation. Standard errors are in parentheses and are clustered at the zip code cluster level.

variable to be the average insurance premium in a zip-code, corresponding to the source of variation in the rate of uninsured drivers. In the first three columns of Table 4, we regress the rate of uninsured drivers on each of our three instrumental variables and find the rate of uninsured motorists decreased. We also run a first-stage regression for our collapsed sample in column (4) by including all our instrumental variables and controls. We find a F-statistic for the hypothesis that all instruments

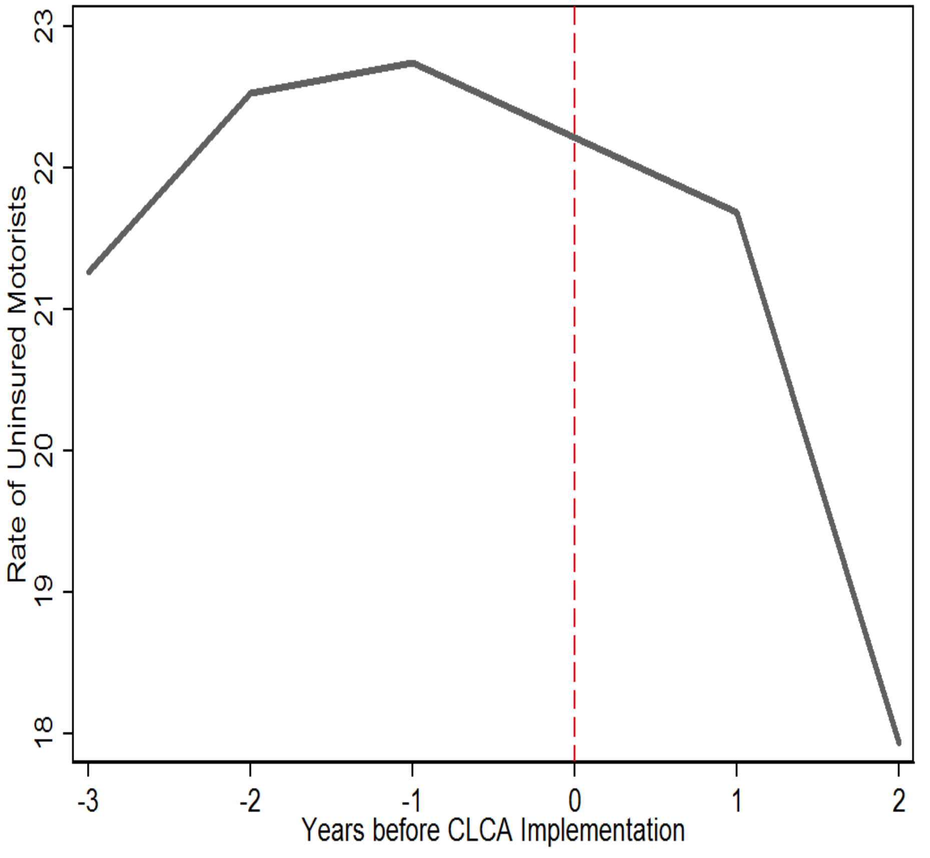

FIGURE 4. Rate of uninsured motorists. This figure plots the average rate of uninsured drivers leading up to and immediately after the implementation of the CLCA program. See figure 1 for the dates of implementation of the CLCA program in specific counties. The dashed line denotes the year in which the CLCA program was implemented in the county.

jointly have no effects to be 17.49, which is above the standard threshold for weak instruments. Overall, our results demonstrate that the CLCA program indeed reached the desired goal of reducing the number of uninsured drivers, which provides the variation essential to our empirical strategy. In terms of economic magnitude, the introduction of the CLCA program led to a roughly 1-2 percentage point drop in the rate of uninsured drivers.

5.2. Estimates of the Externality

Table 5 presents our main results of equation (2) with the level of aggregation collapsed to the zip code cluster by year level. We also present the first-stage results along with the F-statistic for this collapsed sample. The second column includes geographic level

TABLE 5. Main results for the collapsed sample.

	(1)	(2)	(3)	(4)	(5)
			Full Sample		
Uninsured Drivers	30.49***	28.96***	42.16***	17.99**	25.23***
	(8.136)	(7.983)	(11.59)	(7.875)	(8.031)
			First Stage		
Months CLCA in Effect	0.164	0.139	-0.0698	-0.0677	
	(0.121)	(0.117)	(0.101)	(0.0436)	
Months CLCA in Effect^2	-0.0148**	-0.0132*	-0.00482		-0.00559**
	(0.00709)	(0.00700)	(0.00675)		(0.00260)
CLCA Program in Effect	-1.069***	-1.062***		-0.733**	-0.759**
	(0.387)	(0.387)		(0.358)	(0.325)
F-Stat	12.14	12.30	13.33	15.95	17.49
Instrument Dropped			Indicator	Months^2	Months
Observations	995	995	995	995	995

Notes: * $p < .1$, ** $p < .05$, *** $p < .01$. We collapse the sample to the zip-code level, corresponding to the source of variation. The dependent variable in all columns is the average real premium quote in a zip-code. The rate of uninsured drivers is measured between 0 and 100 (in percentage points). In the IV estimates of column (1) and (2) the rate of uninsured drivers is instrumented using (i) the average number of months during which the CLCA program was in effect in a zip code cluster, (ii) the average number of months during which the CLCA program was in effect in a zip code cluster squared and (iii) an indicator of whether or not the program has been in effect for more than four months. Columns from (2) to (5) include controls for crash rate on top of the specification of column (1). Columns (3) to (5) drop one of the instruments as indicated in the table. The first stage of the instrumental variables regression is given below each estimate for the effect of uninsured motorists, along with an F test for the joint significance of the instruments. The rate of uninsured drivers is measured by UMBI/BI. The crash rate is measured by the number of injury exposures over the total number of registered vehicles in a county and is included in all specifications. All specifications include zip code cluster and year fixed effects. Standard errors are in parentheses and are clustered at the level of the zip code cluster.

controls such as the county crash rate and population on top of the specification in column (1). Columns (3) to (5) each drop one instrument. In all specifications, the coefficients are significant at the 1% confidence level. The results suggest that a one percentage point increase in the rate of uninsured motorists is associated with a 1% increase in insurance premia.

To illustrate the importance of using our IV approach we present OLS and fixed effects estimates of the externality in Table 6, using disaggregated observations. As well as serving as an additional robustness check for our main results, this allows us to include individual level controls.[12] The results do not change substantially when individual level controls are included, which suggests that our results are not driven by changes in survey sample composition. In the first two columns we are treating the rate of the uninsured as exogenous and do not control for zip code fixed effects

12. There are other useful aspects using disaggregated data. First, there may be entry and exit of insurers or changes in the types of vehicles in a geographic area. This could make aggregated estimates less precise. Second, there is substantial price dispersion so aggregating will introduce noise into our estimates. Moreover, the disaggregated data allows us to test explicitly in Table 3 that the composition of the survey did not change. This strengthens our identifying assumptions.

TABLE 6. Comparison of results for effects of the uninsured on premia.

	(1) OLS	(2) OLS	(3) FE	(4) FE	(5) IV	(6) IV
Uninsured Drivers	-11.19*	-13.01*	3.205**	3.169**	27.59***	27.58***
	(6.613)	(6.708)	(1.553)	(1.565)	(9.212)	(9.315)
Crash Rate	1337.6***	1352.8***	262.6***	257.9***	167.6*	162.1*
	(359.8)	(363.1)	(75.49)	(75.45)	(91.95)	(93.32)
At Fault Crash	770.5***	804.8***	739.4***	780.5***	739.4***	780.5***
	(13.18)	(13.79)	(13.64)	(14.43)	(13.61)	(14.40)
Standard Plan	1669.2***	1695.0***	1020.9***	1784.0***	1020.9***	1784.0***
	(24.69)	(25.63)	(14.72)	(29.94)	(14.68)	(29.86)
Age		-43.37***		-46.47***		-46.47***
		(0.928)		(0.967)		(0.965)
Daily Miles		38.59***		39.50***		39.50***
		(0.546)		(0.565)		(0.564)
Speeding Ticket		593.1***		556.8***		556.8***
		(9.984)		(10.59)		(10.57)
Female		-148.2***		-169.3***		-169.2***
		(8.947)		(4.046)		(4.059)
Observations	4,724,220	4,724,220	4,724,220	4,724,220	4,724,220	4,724,220

Notes:* $p < .1$, ** $p < .05$, *** $p < .01$. The dependent variable in all columns is the real premium quote offered by a firm. The rate of uninsured drivers is measured between 0 and 100. In the IV estimates the rate of uninsured drivers is instrumented using (i) the average number of months during which the CLCA program was in effect in a zip code cluster, (ii) the average number of months during which the CLCA program was in effect in a zip code cluster squared and (iii) an indicator of whether or not the program has been in effect for more than four months. The rate of uninsured drivers is measured by UMBI/BI. The crash rate is measured by the number of injury exposures over the total number of registered vehicles in a county. Columns (3) to (6) include zip code, year, firm and vehicle fixed effects. Standard errors are in parentheses and are clustered at the level of the zip code cluster.

in the OLS regression. In both specifications, the coefficient on the rate of uninsured drivers is negative and significant at 0.10 level, indicating the nonsensical result that more uninsured drivers reduce insurance premia. This is not surprising given that in these specifications our main source of variation is the cross sectional difference in the rate of uninsured drivers. Geographic factors such as wealth differences, leading to price discrimination, or low vehicle values leading to lower crash costs may result in a negative correlation between premia and the rate of uninsured drivers. These factors necessitate controlling for zip code and other fixed effects. Indeed, when we control for zip code and year fixed effects in Table 6, columns (3)-(4), the coefficient on the rate of the uninsured changes its sign and becomes positive and statistically significant.

The inclusion of zip code fixed effects corrects only part of the endogeneity problem that arises from cross-sectional differences across zip codes. The simultaneity bias illustrated in our simple model in Section 2 will lead the coefficient to be biased upwards in an OLS regression even after controlling for fixed effects. Another potential source of bias is unobserved time varying geographic variables that are determinants of premia and at the same time correlated with the rate of uninsured motorists. For example, deteriorating local economic conditions may both lower premia through inducing individuals to purchase lower cost cars, as well as induce some individuals to drive without insurance. At the same time, we face a third potential source of bias, measurement error in the rate of uninsured drivers. In the authors' view this is the most serious source of bias in the OLS and Fixed Effects estimates. We use a widely used measure for the rate of uninsured drivers, the uninsured motorist bodily injury claims over the insured motorist bodily injury claims. Since this measure is not a direct observation of the rate of uninsured motorists, but rather an estimate based on crash data, we expect this to be a noisy measure of the true rate of uninsured motorists. This measurement error effect could bias the coefficient towards zero.[13] We will return to the issue of measurement error shortly after our discussion on the main results using an instrumental variables approach.

The competing effects of simultaneity and omitted variables bias coupled with the measurement error make the OLS fixed effects estimates uninformative in regards to the true causal effect of the rate of uninsured drivers on insurance premia, other than providing us with evidence for the weak assertion that the effect is nonnegative. Fortunately, we can solve the above problems by instrumenting for the rate of uninsured drivers using the staggered introduction of the CLCA program that changes the rate of uninsured drivers. As reported in Table 6, columns (5)-(6), once instrumented for, the coefficient for the rate of uninsured drivers becomes higher in absolute value, with a positive value of \$27.5, or roughly 1% of the total value of a typical insurance contract in our data, showing a much larger effect of the uninsured on the insured than methods not controlling for the endogeneity and measurement error problem. Our empirical findings are consistent with theoretical predictions of Smith and Wright (1992) and Keeton and Kwerel (1984) in the auto insurance literature. The magnitude of our results does not change much when we add various demographic and driving record controls, providing an additional test that our instrument is uncorrelated with these controls. The (untabulated) R^2 is quite high when we include all controls, at .722, suggesting that our controls explain a great deal of the variation in automobile insurance premia. This is not surprising given that we control for most factors on which firms are legally allowed to price in California, and that we include zip code fixed effects.

Insurance premia are also increasing with the crash rate in a county, which is again consistent with Smith and Wright (1992). If we drop the crash rate from the specification, the coefficient on the rate of uninsured drivers does not change substantially, which suggests that moral hazard does not play a large part in explaining

13. If instead of observing a variable x_i, we observe a noisy measure $x_i^* = x_i + \eta_i$ where $\eta_i \perp x_i$, $E[\eta_i | x_i] = \mathbf{0}$ and $\mathrm{Var}[\eta_i | x_i] = \sigma_\eta^2$ and $\mathrm{Var}[x_i] = \sigma_x^2$ the coefficients $\hat{\beta}$ of the regression $y_i = x_i^* \beta + \varepsilon_i$, under standard assumptions, will be consistent for $\sigma_x^2/(\sigma_\eta^2 + \sigma_x^2)\beta$. When we follow Cohen and Dehejia (2004) and estimate our main specification in logs, which is more robust to measurement error, we find that the difference between the fixed effects and instrumental variables estimates is smaller supporting our hypothesis that measurement error accounts for much of the bias.

our results. The sign and magnitude of other coefficients in the results presented in Table 6 are also consistent with riskier drivers paying higher premia. Premia are also lower for women and middle aged drivers, which is likely to reflect lower crash rates for women and higher crash rates for inexperienced drivers. The latter point is supported by adding in the number of years licensed to the specifications as controls. Insurance premia are also increasing in the number of miles an individual drives to work daily as well as in speeding tickets and at-fault crashes, both of which are likely to be correlated with an increased risk of being involved in a crash. While our main variable of interest is the rate of uninsured drivers, the other coefficients in the regression support basic theoretical underpinnings of Smith and Wright (1992), Keeton and Kwerel (1984) and Arrow (1963), namely that premia will be increasing in crash rates and the inherent riskiness of a driver.

To examine the seemingly large discrepancy between the fixed effects OLS estimates and consistent IV estimates, we make use of results from Griliches and Hausman (1986) which explain how measurement error could bias the estimation in panel data compared with cross-sectional results.[14] Griliches and Hausman (1986) demonstrate that in certain situations measurement error could lead to varying degrees of bias depending on whether one does within-panel fixed effects estimation, first-difference or higher-order-difference OLS estimation. This will be true in our context given a high degree of positive serial correlation among the true uninsured drivers' rate and a serially uncorrelated measurement error term. They show that the larger the gap taken for differences, the larger the estimate and the smaller the bias; and the first differenced OLS estimate will be smaller than the within-panel fixed effects estimate. To see if this is indeed the case, we collapse our observations to the zip-code by year level and present our results using within-panel fixed effects, first-difference OLS, second-difference OLS and third-difference OLS in Appendix Table 1. Consistent with the predictions regarding measurement error from Griliches and Hausman (1986), we find the first-difference OLS estimate is the lowest, with the within-panel estimate being slightly larger, which is itself smaller than the second-difference OLS and with the third-difference OLS being the largest estimate. Griliches and Hausman (1986) provide a bias-correction formula to estimate the true parameter of interest based on the within-panel estimate, the first-differenced OLS estimate and the sample variance of the uninsured rate. The true estimate will be

$$\beta = \left[\frac{2\beta_w}{\mathrm{Var}(\lambda_{gt}-\lambda_{gt-1})} - \frac{(T-1)}{T}\frac{\beta_d}{\mathrm{Var}(\lambda_{gt}-\lambda_g)}\right] \times \left[\frac{2}{\mathrm{Var}(\lambda_{gt}-\lambda_{gt-1})} - \frac{(T-1)}{T}\frac{1}{\mathrm{Var}(\lambda_{gt}-\lambda_g)}\right]^{-1}$$

where $\beta_w = 4.2$ and $\beta_d = 0.8$ are the estimates from within-panel fixed effects and first-differenced OLS, respectively; $\mathrm{Var}(\lambda_{gt}-\lambda_{gt-1}) = 12.9$ is the variance of the

14. We thank an anonymous referee for this suggestion. See Chan and Stevens (2004) for another application of the Griliches and Hausman approach. Also see Baltagi (2008) pp.187-190 for more discussion of measurement error in panel data models.

first-differenced uninsured driver's rate while $\mathrm{Var}(\lambda_{gt} - \lambda_g) = 6.4$ is the variance of the uninsured driver's rate minus the cross-period average. $T = 5$ is the time period of panel in our sample. Using these values, we find the bias-corrected parameter value is 17.51, which is very close to our preferred IV estimate and within standard confidence intervals. The above calculation illustrates that measurement error seems to be present in our sample. As was mentioned earlier, the fact that measurement error is present in standard measures of the rate of uninsured motorists is not surprising. The metric of the rate of uninsured motorists, UMBI/BI, is built from claims data from crashes. Crashes are infrequent and occur stochastically, so it is natural that the standard measure of uninsured motorists is measured with error. Overall, our analyses show that measurement error can account for the large discrepancy between the OLS and IV estimates.

The magnitude of the IV estimates are consistent with a stylized model of automobile insurance pricing in which the externality is fully passed onto consumers. An insurance company will be required to pay for damages in two scenarios, one, if the driver is involved in a crash and found to be at fault, and two, if the driver is involved in a crash with an uninsured driver. Recall that λ is the rate of uninsured drivers and assume that (1) the probability $\theta = \pi_i/2$ of being at fault in a crash for insured or uninsured drivers is the same and (2) insured and uninsured drivers in expectation cause the same amount of loss, L. Under competitive pricing, the insurance premium is $P = L(\theta + \lambda\theta)$. Thus, for a 1 percentage point decline in the rate of uninsured drivers, the percentage decrease in premium will be $dP/P = 0.01L\lambda\theta/(L(\theta + \lambda\theta)) = 0.01/(1 + \lambda)$. In California the rate of uninsured drivers is roughly 20%, so a 1 percentage point drop in the rate of uninsured drivers will reduce the premium by 0.83%. Given that the average premium in our data is roughly \$2,356,[15] and we estimate that a 1% decrease in the rate of uninsured drivers reduces premia by \$27 in column (2) of Table 6, the aforementioned logic is very much in line with our results.

While the instrumental variables results above provide evidence for the presence of an externality, there are two potentially serious confounds. First, the CLCA plan is an insurance plan itself, and the presence of the plan in the market may directly cause insurers to lower premia. Second, riskier drivers may sort into the CLCA plan, lowering premia through changing the risk composition of drivers. Section 6 shows that the potential effects of both concerns are at most very small quantitatively.

5.3. *Pigouvian Taxation*

When aggregated over all insured drivers in California the social costs of the externality[16] are substantial. Based on our main specification, and uninsured motorists

15. The average premium in our data is larger than the typical premium paid in California since the survey data oversamples drivers with at fault crashes and speeding tickets.

16. See Parry et al. (2007) for a survey of externalities associated with automobile use and Edlin and Karaca-Mandic (2006) for a discussion of the general externality caused from miles driven.

rates in California in 2007 as well as rates of uninsured motorist coverage,[17] the total cost of the externality to California is about $6 billion. If the magnitude of the effect in other US states is similar in size to California on a per-person basis, the size of the externality would be quite large, which we calculated to be $27 billion nation-wide using NAIC estimates of average premia. If the magnitude of the effect is similar in the United Kingdom, we would estimate the size of the externality to be roughly £1.6 billion. This is substantially smaller than in the United States, given that the rate of uninsured motorists in the United Kingdom is only 3.5%. The Motor Insurers' Bureau levies a £33 surcharge on automobile insurance premia to fund damage arising from uninsured motorists. We note that this is quite close to our estimates in California– we would predict that uninsured motorists would raise premia by $90 (£50) if the rate of uninsured motorists is 3.5%.

The presence of externalities can be corrected by pricing the damage caused by uninsured drivers to other drivers. One way to accomplish this task is by levying a Pigouvian tax, or equivalent fine on uninsured drivers. Individuals would then only fail to purchase insurance if their private benefit exceeds the external social cost of being uninsured. This is in effect the system already in place in most of the United States directly or indirectly, as well as many other countries. While ostensibly it is illegal for motorists to drive without insurance in most US states, the current system closely mimics a Pigouvian tax. In most US states drivers who are caught without insurance are forced to pay a citation, which is essentially equivalent to a stochastic Pigouvian tax on driving uninsured. In theory authorities could set fines large enough so that very few drivers drive without insurance,[18] but intuitively the welfare effects of forcing uninsured motorists to buy insurance without a subsidy are ambiguous. The fine would disproportionately affect low income households, where most uninsured drivers tend to be located.

There is a long tradition since Pigou (1920) of economists advocating corrective taxes on externalities.[19] Despite the optimality of the Pigouvian approach, determining what corrective taxes should be levied is often difficult in practice. Typically, the most daunting challenge is measuring the size of the externality, which we did in the previous section of this paper. To accomplish our objective, we can levy a Pigouvian tax on uninsured drivers in a fashion similar to how most US states currently fine uninsured motorists. Authorities force uninsured drivers to pay a tax τ if they are uninsured and redistribute a subsidy s to all drivers. Given the framework outlined in the theory

17. In 2007, Department of Insurance data indicate that 17.83% of motorists were uninsured, and there were 19,280,329 vehicles with uninsured motorist coverage in the state of California.

18. This is the case in some European countries. For example, in France in 2012 the fine for driving without insurance is €3,750 with a three-year license suspension. Given these high fines, the rate of uninsured motorists in France is low, at .1% of registered vehicles compared to 14% in the US. Many European countries also have rates of uninsured motorists substantially lower than the US, as well as higher penalties for driving without insurance.

19. For the sake of brevity, we do not offer a full treatment of Pigouvian taxation. See Sandmo (1978) for a classic treatment of the problem or Mankiw (2009) for a more recent discussion of Pigouvian taxes.

section and under some weak assumptions, we can compute the optimal fine which only depends on observables. Implicitly, the probability of being caught uninsured must be factored into the tax, as currently drivers will only pay the tax if they are stopped by law enforcement officials. The tax will reduce the size of the externality by discouraging uninsured driving, while at the same time directly lowering premia by subsidizing insured drivers. Essentially the government can use a tax to correct the externality, fining uninsured drivers and redistributing the proceeds to all drivers.

Given three possible states, no crash, a crash with an insured driver, and a crash with an uninsured driver, consumers choose optimal amounts of insurance to purchase along the lines presented in Section 2. After consumers have made optimal insurance choices, the government solves for a representative consumer with insurance choice determined by consumers' optimization, $\max_{\tau} V(s,\tau)$ for given tax τ and subsidy s, subject to the government budget being balanced, $s = \lambda(\tau)\tau$. Solving the government's problem and applying the envelope theorem, after some algebra we obtain that the optimal corrective tax depends only on β, our estimated effect of the change in dollars in insurance premium with respect to a one percentage point change in the rate of uninsured drivers, and $\lambda(\tau)$ given by

$$\tau^* = \beta(1 - \lambda(\tau))$$

See Appendix B for a detailed derivation of the formula, which follows Chetty (2006). The optimal tax formula is simple and intuitive, depending on β, the amount of premia increase from uninsured drivers and $\lambda(\tau)$, the rate of uninsured drivers. The result indicates that uninsured individuals should fully bear the cost of the externality, which is similar to the Pigouvian tax in Edlin and Karaca-Mandic (2006). The fine is unambiguously increasing in β , which is the externality that the Pigouvian tax is designed to correct. A larger effect stemming from this externality would mean a larger corrective fine. As we would expect, the fine is zero if there is no externality. The optimal tax is always positive and thus will be a fine on the uninsured and a subsidy for the insured.

The results indicate that redistributive fines for driving without insurance should be \$2,240. This value is substantially higher than current fines in California, where individuals pay between \$100–\$200 for the first offense and \$500 for the second. This difference becomes even clearer when we note that enforcement is stochastic.[20] It is thus quite possible that, if relatively few drivers are caught driving without insurance, current fines are substantially below the optimum. It is difficult to determine the expected fine that California residents would pay for driving uninsured, as statewide data does not exist on tickets for driving uninsured. If the optimal fine of \$2,240 were enforced rigorously, this would effectively eliminate the uninsured driver problem as it would be cheaper for nearly all individuals to purchase a basic insurance plan rather than pay a heavy fine. While in theory a Pigouvian tax would be welfare improving, in

20. See Polinsky and Shavell (1979) for a discussion of the tradeoff between the probability and magnitude of fines.

practice there are several concerns. First, the uninsured may be liquidity constrained and unable to pay. Second, uninsured drivers are typically poor and may have a very high marginal utility of consumption. Third, enforcement and administrative costs may be costly and outweigh the benefits.

6. Robustness

Table 7 and Table 8 presents several robustness checks using both the disaggregated sample and the collapsed sample which show that our basic result holds controlling for several potential confounds. In all cases the coefficients remain significant at the .1 level and are similar in magnitude to the main results. In the main dataset, we restrict the sample to only observations where there is one driver on the insurance plan. Our main results are robust to including multi-driver policies as well. Concerns with the data and our measure of uninsured drivers are addressed in Appendix A.

Increased Competition: A potential concern to our empirical strategy and results is that the CLCA program, being an insurance plan itself, affects the insurance premium in the commercial market through an increased competition channel. As well as lowering the rate of uninsured drivers, introducing the CLCA program also offered another low-cost plan to consumers which may have led insurance providers to react by lowering premia. Thus it is possible that our results are partially or entirely driven by increased competition rather than the effect of the CLCA program on uninsured drivers. While we have no data on income to determine eligibility for the CLCA program,[21] we exploit another part of the eligibility requirement of the CLCA program to produce results that are free from this potential confound. In years prior to 2005, only vehicles worth less than $12,000 could be insured under the CLCA program, and this cap was raised to $20,000 in 2006 and following years. We can thus restrict our sample to only those surveyed insurance plans covering vehicles of higher value as to be ineligible for the CLCA program.[22]

Our findings restricting vehicles to be above certain threshold values and ineligible for the CLCA program are reported in Table 7. In column (1) and (2), we restrict the sample to vehicles above their survey year's maximum allowed car value for the CLCA program, while we restrict to vehicles above $20,000, the maximum allowed car value throughout the years in the CLCA program in column (3) and (4). To the extent that one might be concerned about potential spillover from a lower car-value plan to higher car-value plan or a coarse pricing strategy by insurance companies, we

21. It is illegal for insurers in California to price on factors such as income or race.

22. If there are spillovers across different products due to insurers solving a multi-product problem; or the CLCA program introduces competition in the liability and property damage markets and affect premia for all types of vehicles, then we can not completely rule out the effects from increased competition. With this caveat in mind, we note that the CLCA Program is a relatively small program with the average effect of decreasing the rate of uninsured drivers by 1%. Compared with 80% of drivers who insure in California, this program only targets the very low end segment of the insurance market and likely will have limited impact on such a general equilibrium effect due to strategic pricing in all types of auto insurance products.

TABLE 7. Results using various vehicle value thresholds above eligibility.

	(1) Ineligible	(2) Ineligible	(3) \$20,000	(4) \$20,000	(5) \$25,000	(6) \$25,000
			Disaggregated Sample			
Uninsured Drivers	27.73***	27.81***	26.46***	26.43***	33.84***	33.82***
	(9.232)	(9.223)	(7.151)	(7.136)	(11.82)	(11.79)
Observations	3,802,252	3,802,252	3,230,538	3,230,538	1,699,610	1,699,610
			Collapsed Sample			
Uninsured Drivers	29.20***	22.28***	25.02***	21.16***	24.63***	20.77***
	(3.760)	(3.308)	(5.921)	(4.709)	(5.850)	(4.652)
Observations	995	995	995	995	995	995
Controls		✓		✓		✓
Zip Code Fixed Effects	✓	✓	✓	✓	✓	✓
Year Fixed Effects	✓	✓	✓	✓	✓	✓

Notes:* $p < .1$, ** $p < .05$, *** $p < .01$. The dependent variable in the disaggregated sample panel is the real premium quote offered by a firm in a zip code and in the collapsed sample panel is the average premium quote in a zip code. The rate of uninsured drivers is measured between 0 and 100. Columns (1) to (2) restrict the sample to vehicles ineligible for the CLCA program in the current year, while columns (3) to (4) restrict the sample to vehicles ineligible for the CLCA program during the entire sample period. We restrict the sample to vehicles with values above \$25,000 in columns (5) to (6). In all estimates the rate of uninsured drivers is instrumented using (i) the average number of months during which the CLCA program was in effect in a zip code cluster, (ii) the average number of months during which the CLCA program was in effect in a zip code cluster squared and (iii) an indicator of whether or not the program has been in effect for more than four months. The rate of uninsured drivers is measured by UMBI/BI. The crash rate is measured by the number of injury exposures over the total number of registered vehicles in a county. The disaggregated sample also includes firm and vehicle fixed effects. Standard errors are in parentheses and are clustered at the level of the zip code cluster.

conduct a "stress test" by restricting the sample to vehicles above \$25,000 in column (5) and (6). If increased competition due to a new plan being offered could explain the bulk of our findings, we would expect the coefficient on the rate of uninsured drivers to drop substantially. Results from our three different sub-samples show this is not the case: while the point estimates of these regressions are slightly lower than that of the regression in Table 6, they are statistically indistinguishable from our main result, given the magnitude of standard errors. Our results are again statistically significant at 1% level. This demonstrates that increased competition cannot explain our findings, and that the effect of the CLCA program on premia comes almost entirely from decreasing the rate of uninsured drivers.

Weak Instruments. A typical concern regarding the use of the instrumental variables approach is the strength of the instruments. While the F-statistic of 17.49 in our first-stage regression exceeds the standard rule-of-thumb threshold of 10 for weak instruments (Staiger and Stock (1997)), one might still be concerned about whether instruments with this degree of power were able to produce stable estimates. As in Stock et al. (2002), we estimate the effect using limited information maximum likelihood (LIML) methods. The implementation follows the approach in Moreira and Poi (2003). The corrected confidence intervals are constructed from tests of coefficients based on the conditional distributions of nonpivotal statistics. We find in column (1) of Table 8 that our results from LIML are extremely close to the IV results, which indicates that weak instruments should not be a serious concern in our context.

County Waves. As we show in Figure 2, counties in the last wave of the CLCA program have higher crash rates and are smaller and more sparsely populated. These counties also tend to have slightly lower premia than other counties. To guard against the potential confound that the results are driven by some particular counties that have different characteristics than others, we exclude the counties in the final wave. The results are presented in column (3) of Table 8. In this specification the coefficient on the rate of uninsured drivers turns out to be close to our main results. We conclude that our results are not driven by the counties in the final wave being different from other counties. Our results are robust to excluding counties in any wave, for example, Los Angeles and San Francisco in the pilot wave, and we conclude that our results are not driven by any single wave.

Unobserved Selection. A potential concern is that unobserved selection on crash risk could play a major role in determining premia. For example, drivers switching to the CLCA program could be unobservably riskier than those remaining in traditional plans. This effect could lead insurance premia to fall for those remaining in traditional insurance plans. We view unobserved selection as unlikely given the regulation of automobile insurance pricing in California. First, following Proposition 103, automobile insurers are only allowed to price on certain factors, the vast majority of which are in our dataset. It is not clear why unobservably risky individuals would prefer the CLCA plan to traditional insurance plans with higher coverage limits. At the same time, we refer to our results in Table 7 where we restrict to several subsamples with car-values above the eligible value for the CLCA program. Drivers purchasing these types of insurance plans would not be affected by any unobserved selection into the CLCA program as they are ineligible for the program. We see nearly identical effects for individuals who were ineligible to enter the CLCA program due to high vehicle values, and for this group unobserved selection into the CLCA program cannot explain the price effects. Thus we conclude that unobserved selection is not a major driving force of our results.

County Specific Time Trend. It is possible that the introduction of the CLCA program at the county level was simply correlated with other factors that reduced premia. We test for this possibility by including a county-specific linear time trend in our main specification and the result is in column (4) of Table 8. The source of variation is at the zip code cluster level, given that a zip code could border other zip codes within its 25 miles radius in a different county that had the CLCA program introduced in a different wave. This feature will make sure that including a county-specific linear time trend will not fully absorb our source of variation for identification. Indeed we still obtain a significant effect of the effect of uninsured drivers on premia at the 1% level. The point estimate changes from 29 to 34, but given the relatively large standard error, we can not reject that it is different from the results in our main regression in Table 6 at the 5% level of significance.

Choice of Instruments. Our instruments in the main regression include an indicator of whether the CLCA program was active for over 4 months in a zip code cluster. We vary the definition of this instrument and our results remain robust. Column (5) of Table 8 reports results when we replace it with an indicator of whether or not the CLCA program was active at all during the year. We use an indicator of whether or not the CLCA program was active for over 6 months during the year instead in column (6) and dropped this indicator in column (7). In either of the above cases, altering the definition of our instrument creates almost no change to our main result. Our results are also robust to dropping any individual instrument. We conclude that our results are robust to changes in instrument specification.

Omitted Variables. Another potential concern is that coefficients in our specifications are subject to omitted variables bias. We do not think that this is a significant source of bias given the richness of our data and the regulatory framework in California. Automobile insurance is highly regulated in California, and we have all factors on which insurers are required to price, as well as, in the authors' view, the more important optional pricing factors. Proposition 103, passed in 1988, modified the California Insurance Code to mandate that automobile insurers in California were required to price on driving record, miles driven annually, and the number of years licensed. In addition, insurers were also allowed to price on secondary factors permitted by the insurance commissioner. For the period in which the authors have data (2003-2007), insurance companies were permitted to price on location (zip code), vehicle type and performance, number of vehicles owned by the household, the use of vehicles, gender, marital status, age, demographic characteristics of secondary drivers, persistency, the academic standing of any student in the household, completion of a driver training course, smoking, bundling of products with the same company and claims frequency and severity. Automobile insurers were not allowed to price on any other characteristics, and firms were required to report rate changes in their pricing formulae to the Department of Insurance. The mandatory pricing factors were also required to have a larger weight in the pricing formula than the optional pricing factors.[23] Given that our data includes information on all mandatory pricing factors, as well as the major optional pricing factors for automobile insurance pricing in California, we think it is unlikely that our results are significantly biased by omitted variables.

Underlying Change in Variables that Drive CLCA Wave Timing. In Section 4, we show that the specific timing of waves in the CLCA program is driven by the level of county population size. While county population should be a stationary variable that

23. As stipulated under California Insurance Code Section 1861.02(a).

we do not expect to change rapidly, we guard against the effect of changing population on the demand for automobile insurance in each county by including county population interacted with year dummies in column (8) in Table 8. We find the effect of uninsured drivers on premia increases slightly, but still can not reject that it is identical to what we find in our main results at a 5% level of significance.

Another potential concern is that our results could be driven by compositional changes in the survey data. Our premium data comes from an administrative survey, which uses a host of hypothetical risk profiles of drivers. A priori, there is no reason to believe that the government surveyed insurance premia for different groups of drivers after the CLCA program took effect. In Table 3, we demonstrate this is indeed the case. Since the insurance companies set prices based on several individual-specific characteristics, we directly examine the characteristics of drivers surveyed before and after the CLCA program to make sure that we compare prices for the same group of people. We compare the mean of major risk factors used in the main analysis for insurance pricing in the period before and after the CLCA program has been active for at least four months. These factors include sex, age, plan type, crash rate, daily miles driven, whether the driver has incurred an at-fault crash as well as whether the driver has a recent history of speeding tickets. Our F-test can not reject at 5% level the hypothesis that these characteristics ever changed after the CLCA program took effect. We reject at the 10% level that the crash rate is the same, which is consistent with moral hazard, insured drivers being less cautious and being involved in more crashes. We discuss this issue, which will not bias our results as we control for crash rates, further below. Another potential concern regards the CLCA program attracting some particular group of drivers whose behaviors could affect the insurance premium independent of the uninsured drivers' externality effect. This concern is also dealt with in our robustness check for unobserved selection, as we restrict the sample only to individuals who would have been ineligible for the CLCA program.

Moral Hazard. One potential concern is that our results may slightly underestimate the effect of uninsured drivers, as the CLCA program also introduced moral hazard. That is, increased insurance coverage could increase the risk of a crash,[24] and therefore lead to an increase in insurance premia. By covering previously uninsured individuals, the program may have given some drivers an incentive to drive in a less safe manner. Empirical evidence on this topic has been inconclusive so far. Chiaporri and Salanié (2000) find no evidence of asymmetric information in the automobile insurance market using the positive correlation test. Dionne et al. (2013) find evidence of moral hazard among people with less driving experience in France from 1995-97. They find no information problem for people with more than 15 years of experience. Abbring et al. (2003a) also find no evidence of moral hazard using dynamic insurance data and a test similar to Abbring et al. (2003b). However Cohen (2005) notes that the results of

24. See Shavell (1979) or Arrow (1971) for early discussions of this effect.

TABLE 8. Robustness checks for effects of the uninsured on premia.

	(1) Weak Instrument	(2) No Final Wave	(3) County Year Trend	(4) Indicator Ever Active	(5) Indicator Active 6M	(6) No Indic. Inst.	(7) Population X Year
				Disaggregated Sample			
Uninsured Drivers	29.77*** (5.808)	21.91*** (8.009)	19.51** (9.046)	27.78*** (9.224)	39.03*** (12.11)	38.07*** (12.29)	20.11** (9.179)
Observations	4,724,220	4,228,022	4,724,220	4,724,220	4,724,220	4,724,220	4,724,220
				Collapsed Sample			
Uninsured Drivers	30.20** (12.57)	24.16*** (6.832)	27.62*** (9.371)	40.53*** (11.64)	46.31*** (11.67)	42.52*** (11.88)	23.34*** (8.900)
Observations	995	936	995	995	995	995	995
Controls	✓	✓	✓	✓	✓	✓	✓
Zipcode Effects	✓	✓	✓	✓	✓	✓	✓
Year Effects	✓	✓	✓	✓	✓	✓	✓

Notes:* $p < .1$, ** $p < .05$, *** $p < .01$. The dependent variable in the disaggregated sample panel is the real premium quote offered by a firm in a zip code and in the collapsed sample panel is the average premium quote in a zip code. We use Limited Information Maximum Likelihood to estimate results in column (1). We drop the final wave counties in column (3). In column (4), county-specific year trends are used as extra controls. The rate of uninsured drivers, which is measured between 0 and 100, is instrumented using (i) the average number of months during which the CLCA program was in effect in a zip code cluster, (ii) the average number of months during which the CLCA program was in effect in a zip code cluster squared and (iii) an indicator of whether or not the program has been in effect for more than four months in all columns except 5 to 7. We replace instrument (iii) with an indicator of whether or not the program was active at all in column (5) and use the indicator of whether or not the program was active for over six months instead in column (6). We dropped instrument (iii) in column (7). We include controls for the county population interacted with year fixed effects in column (8). The rate of uninsured drivers is measured by UMBI/BI. The crash rate is measured by the number of injury exposures over the total number of registered vehicles in a county. Each specification includes the crash rate, driving history variables, age, plan type and gender. All specifications include zip code and year fixed effects. The disaggregated sample also includes firm and vehicle fixed effects. Standard errors are in parentheses and are clustered at the level of the zip code cluster.

Chiaporri and Salanié (2000)[25] are also consistent with asymmetric information and learning. Cohen and Dehejia (2004) estimate the effect of automobile insurance on traffic fatalities and find significant effects of moral hazard in the automobile insurance market. In our sample, Table 3 indicates that there is a small but marginally significant (at the 10% level) increase in crash rates following the roll out of the CLCA program. This is consistent with the moral hazard hypothesis. However, we already control for moral hazard effects of the CLCA program by including the local crash rate in our main regression. Thus our estimates may overcontrol for moral hazard as any increase in

25. Chiaporri and Salanié (2000) use a French dataset which focuses on younger drivers. Cohen (2005) finds a positive correlation between a lower deductible and crashes for drivers with three or more years of experience, but not for drivers with less than three years of experience. Dionne et al. (2013) also find no information problem for experienced drivers using a French dataset. This relationship is consistent with classic adverse selection theory, however the study cannot disentangle adverse selection and moral hazard. Cohen and Einav (2003) find that seat belt use does not increase reckless driving, providing further evidence against another type of moral hazard in automobile crashes. See Cohen and Siegelman (2011) for a review of the empirical literature on adverse selection in insurance markets and Einav et al. (2010) for a treatment of welfare effects.

traffic crashes due to moral hazard will be reflected in the coefficient on the crash rate. As a test for whether or not moral hazard significantly affects automobile insurance premia, we drop the county level crash rate from the specification. When we run this alternative specification, the coefficient on the rate of uninsured drivers changes only slightly, and the difference is not significantly different from zero. We take this as evidence that moral hazard does not play a significant part in our results. Furthermore, we regressed the crash rate on the uninsured drivers' rate, instrumented by our policy change. We find a zero effect in the point estimate, which is also not statistically significant. This finding suggests that the crash rate is not changed by newly insured drivers due to introduction of the new insurance program and helps us guard against confounds from the moral hazard channel.

7. Concluding Remarks

This paper makes two contributions. First, we empirically gauge the magnitude of the negative externality generated by uninsured parties in insurance markets, and second, we discuss the optimal corrective Pigouvian tax for this externality based on our empirical analysis. This paper uses a novel panel data set on auto insurance premia in California to quantify the negative externality generated by uninsured drivers on the insured. We overcome the endogeneity challenge inherent in the relationship between insurance premia and the rate of the uninsured, utilizing exogenous variations from the staggered introduction of a policy that lowers the rate of uninsured drivers.

Our data set and empirical strategy enable us to directly estimate the effect of uninsured on premia. Consistent with predictions of the theory, our study suggests that a higher rate of uninsured drivers has a significant effect on the auto insurance premium. We estimate that a 1-percentage-point increase in the rate of uninsured drivers leads to a roughly 1% increase in the total value of the insurance contracts in our data. These estimates imply that each driver could save almost $500 if every motorist became insured in the state of California, which would reduce automobile insurance costs by roughly a third. This study also develops a new formula for computing the optimal corrective tax or fine on uninsured individuals. This formula is parsimonious, relying only on the size of the externality and the rate of uninsured drivers. We compute that the optimal fine should be $2,240, which is substantially higher than current fines in most US states, although similar to fines in some European countries like France.

A fruitful avenue for further research would be to estimate the effect, if any, of the uninsured on health insurance premia. Theory work has noted that there may be a similar effect in the health insurance market resulting from the regulatory requirement that hospitals cross-subsidize the uninsured, and this effect has of late become an important policy issue in the United States due to the passage of the Patient Protection and Affordable Care Act in 2010.[26] The direct effect of the uninsured not paying

26. See Gruber (2008) for a survey of the literature on the uninsured in the health care market. There is a new and growing empirical literature on externalities in insurance markets. Mahoney (2015) and Cabral

medical bills is similar to the effect of uninsured motorists not paying for collision damages after crashes in which they are at fault, but there are a host of significant moral hazard risks associated with medical care as well as externalities from communicable diseases and other effects which could make the effect of the uninsured on premia substantially different. While our quantitative results concern only the automobile insurance market, estimating the effect, if any exists, of the uninsured on health insurance premia would serve both to test the predictions of economic theory and better inform the policy debate about health care.

References

Abbring, Jaap, Pierre-André Chiappori, and Jean Pinquet (2003a). "Moral Hazard and Dynamic Insurance Data." *Journal of the European Economic Association*, 1(4), 767–820.

Abbring, Jaap, James Heckman, Pierre-André Chiappori, and Jean Pinquet (2003b). "Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish." *Journal of the European Economic Association*, 1(2-3), 512–521.

Anderson, Michael (2008). "Safety for Whom? The Effects of Light Trucks on Traffic Fatalities." *Journal of Health Economics*, 27(4), 973–89.

Anderson, Michael and Maximillian Auffhammer (2014). "Pounds that Kill: The External Costs of Vehicle Weight." *The Review of Economic Studies*.

Arrow, Kenneth (1963). "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review*, 53(5), 941–973.

Arrow, Kenneth (1971). "Insurance, Risk and Resource Allocation." *Essays in the Theory of Risk Bearing*, Chicago: Markham Publishing Company.

Baltagi, Badi (2008). *Econometric Analysis of Panel Data*. John Wiley & Sons.

Bureau of Transportation Statistics (2006). "Highway Vehicle Miles Traveled." *Federal Highway Administration, Washington, DC*.

Cabral, Marika and Neale Mahoney (2014). "Externalities and taxation of supplemental insurance: A study of Medicare and Medigap." Tech. rep., National Bureau of Economic Research.

Chan, Sewin and Ann Huff Stevens (2004). "Do changes in pension incentives affect retirement? A longitudinal study of subjective retirement expectations." *Journal of Public Economics*, 88(7), 1307–1333.

Chetty, Raj (2006). "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics*, 90(10), 2351–2356.

Chiaporri, Pierre-Andre and Berndard Salanié (2000). "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1), 56–78.

Cohen, Alma (2005). "Asymmetric Information and Learning: Evidence from the Automobile Insurance Market." *Review of Economics and Statistics*, 87(2), 197–207.

Cohen, Alma and Rajeev Dehejia (2004). "The Effect of Automobile Insurance and Accident Liability Laws on Traffic Fatalities." *Journal of Law and Economics*, 47(2), 357–93.

Cohen, Alma and Liran Einav (2003). "The Effect of Mandatory Seat Beat Laws of Driving Behavior and Traffic Fatalities." *Review of Economics and Statistics*, 87(4), 823–843.

Cohen, Alma and Peter Siegelman (2011). "Testing for Adverse Selection in Insurance Markets." *The Journal of Risk and Insurance*, 77(1), 39–84.

and Mahoney (2014) study related externalities in the health insurance market, while Kolstad and Kowalski (2014) and Hackmann et al. (2014) also study the effects of mandates which decreased the rate of uninsured.

Dionne, Georges, Pierre-Carl Michaud, and Maki Dahchour (2013). "Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France." *Journal of the European Economic Association*, 11(14), 897–917.

Edlin, Aaron S. and Pinar Karaca-Mandic (2006). "The Accident Externality from Driving." *Journal of Political Economy*, 114(5), 931–955.

Einav, Liran, Amy Finkelstein, and Mark R Cullen (2010). "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 124(3), 877–921.

Griliches, Zvi and Jerry A Hausman (1986). "Errors in variables in panel data." *Journal of Econometrics*, 31(1), 93–118.

Gruber, Jonathan (2008). "Covering the Uninsured in the United States." *Journal of Economic Literature*, 46(3), 571–606.

Hackmann, Martin, Jonathan Kolstad, and Amanda Kowalski (2014). "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *NBER Working Paper 19149*.

Keeton, William R. and Evan Kwerel (1984). "Externalities in Automobile Insurance and the Uninsured Driver Problem." *Journal of Law and Economics*, 27, 149–79.

Kolstad, Jonathan and Amanda Kowalski (2014). "Mandate-Based Health Reform and the Labor Market: Evidence from the Massachusetts Reform." *NBER Working Paper 17933*.

Mahoney, Neale (2015). "Bankruptcy as Implicit Health Insurance." *American Economic Review*, 105(2).

Mankiw, N Gregory (2009). "Smart Taxes: An Open Invitation to Join the Pigou Club." *Eastern Economic Journal*, 35, 14–23.

Moreira, Marcelo J. and Brain P. Poi (2003). "Implementing tests with correct size in the simultaneous equations model." *Stata Journal*, 3(1), 57–70.

Parry, Ian, Margaret Walls, and Winston Harrington (2007). "Automobile Externalities and Policies." *Journal of Economic Literature*, 45(2), 373–399.

Pigou, A (1920). "The Economics of Welfare." London: MacMillan.

Polinsky, Mitchell and Steven Shavell (1979). "The Optimal Tradeoff between the Probability and Magnitude of Fines." *American Economic Review*, 69(5), 880–891.

Sandmo, A (1978). "Optimal Taxation in the Presence of Externalities." *The Swedish Journal of Economics*, 77(1), 86–98.

Schultz, Chris and Linda Yarber (2006). "Report to the Legislature & Consumer Education and Outreach Plan." *California Department of Insurance*.

Shavell, Steven (1979). "On Moral Hazard and Insurance." *The Quarterly Journal of Economics*, 93(4), 541–562.

Smith, Eric F. and Randall Wright (1992). "Why is Automobile Insurance in Philadelphia So Damn Expensive." *American Economic Review*, 82(4), 756–72.

Staiger, Douglas and James Stock (1997). "Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3), 557–586.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo (2002). "Testing for Weak Instruments in Linear IV Regression." *Journal of Business and Economic Statistics*, 20(4), 518–529.