

The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation[†]

By CHRISTOPHER R. KNITTEL AND RYAN SANDLER*

When consumers or firms don't face the true social cost of their actions, market outcomes are inefficient. In the case of negative externalities, Pigouvian taxes are one way to correct this market failure, but it may be infeasible to tax the externality directly. The alternative, taxing a related product, will be second-best. In this paper, we show that in the presence of heterogeneous externalities and elasticities, this type of indirect tax performs poorly. In our empirical application, gasoline taxes to address pollution externalities, less than a third of the deadweight loss of the externality is addressed by second-best optimal taxes. (JEL D62, H21, H23, H71, H76, Q53, R48)

A basic tenet of economics posits that when consumers or firms do not face the true social cost of their actions, market outcomes are inefficient. In the case of externalities, Pigouvian taxes can correct this market failure, and the optimal tax leads agents to internalize the true cost of their actions. Although technology is increasingly allowing policymakers to implement Pigouvian taxes that precisely match externalities, in practice directly taxing the externality is often either technically or politically infeasible. In such cases, policymakers might tax a product correlated with the externality. This introduces an additional complication: the level of externality generated can vary across agents. For example, instead of taxing vehicle emissions, policymakers tax gasoline even though emissions per gallon of gasoline consumed varies across vehicles. Similarly, a uniform alcohol tax may be imposed to reduce the negative externalities associated with use, even though externalities likely vary by the type of alcohol or who is consuming it. We refer to uniform taxes intended to address a heterogeneous externality as *second-best optimal*

*Knittel: William Barton Rogers Professor of Energy Economics, Sloan School of Management, MIT and NBER, 100 Main St., Cambridge, MA 02142 (email: knittel@mit.edu); Sandler: Office of Research, Bureau of Financial Protection, 1700 G Street NW, Washington, DC 20001 (email: ryan.sandler@cfpb.gov). This paper has benefited from conference discussion by Jim Sallee and conversations with Severin Borenstein, Joseph Doyle, Michael Greenstone, Michael Grubb, Jonathan Hughes, Dave Rapson, Nicholas Sanders, and Catherine Wolfram. The paper has also benefited from participants at the NBER Energy and Environmental Economics Spring Meeting and seminar participants at Northeastern University, University of Chicago, MIT, and Yale University. We gratefully acknowledge financial support from the University of California Center for Energy and Environmental Economics. The research was also supported by a grant from the Sustainable Transportation Center at the University of California, Davis, which receives funding from the US Department of Transportation and Caltrans, the California Department of Transportation, through the University Transportation Centers program. The views expressed in this article are those of the authors and do not necessarily reflect those of the Bureau of Financial Protection or the United States.

[†]Go to <https://doi.org/10.1257/pol.20160508> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

(SBO) taxes. When the level of externalities produced differs across consumers, a uniform tax will be second best and deadweight loss will remain. Moreover, if price responsiveness and externalities are correlated, the SBO tax will differ from the more easily calculated average externality.

In this paper, we study the efficiency and equity implications of using a uniform SBO tax in place of the optimal Pigouvian tax. As an empirical example, we focus on the local pollution externalities of the personal transportation market in California between 1998 and 2008. Policymakers are often concerned about four externalities in the transportation sector: local pollution from tailpipe emissions, known as criteria pollutants;¹ climate change externalities resulting from carbon dioxide associated with the engine's combustion process; road congestion; and externalities associated with accidents. For all but the climate change externality, a gasoline tax is an imperfect instrument. While fuel consumption is positively correlated with criteria pollutant emissions, congestion, and accident externalities, it is not perfectly correlated.²

In any market with heterogeneous externalities, the relationship between the SBO tax and the first-best Pigouvian tax depends on three empirical relationships: the distribution of externalities across individuals; the extent to which prices affect the implicit demand for the externality; and the correlation between individual-specific demand responses and externality levels. If individual demand responses do not differ, the SBO gasoline tax will simply be the average per-unit externality. However, if price responsiveness and externalities are correlated, Diamond (1973) shows that the SBO tax will be a weighted average of individual per-unit externalities, where the weights are the price derivatives of the individual-specific gasoline demand curves. In our empirical work, we allow the elasticity of vehicle miles traveled (VMT) with respect to gasoline prices—an elasticity that we hereafter call the VMT elasticity—to vary depending on a vehicle's emissions per mile traveled, which we observe in our data.

An important empirical result of this paper is that we find that vehicle-level emissions are correlated with vehicle-specific VMT elasticities; dirtier vehicles are more price responsive.³ Using detailed vehicle-specific data on miles driven, we show a positive correlation between criteria pollutant emissions and the VMT elasticity (in absolute value) holds for all three pollutants for which we have data: carbon monoxide (CO), hydrocarbons (HCs), and nitrogen oxides (NO_x). VMT elasticities are also positively correlated with greenhouse gas emissions and vehicle weight.

¹ Criteria air pollutants are the only air pollutants for which the administrator of the US Environmental Protection Agency has established national air quality standards defining allowable ambient air concentrations. Congress has focused regulatory attention on these pollutants (i.e., carbon monoxide, lead, nitrogen dioxide, ozone, particulate matter, and sulfur dioxide) because they endanger public health and are widespread throughout the United States.

² In contrast, burning a gallon of gasoline leads to the same amount of carbon dioxide emissions regardless of the vehicle, so a gasoline tax is the optimal instrument for climate change externalities.

³ While we tend to discuss the responsiveness of individuals to changes in prices because drivers can shift miles from one car to another, the more relevant response is the response of the number of miles driven by a particular vehicle. Therefore, throughout, we focus on the response of miles driven by a particular vehicle, not by a particular driver.

We find the average VMT elasticity is -0.13 , but differences between cleaner and dirtier vehicles are substantial. When we allow VMT elasticity to vary by within-year quartiles of NO_x emissions, the elasticity for vehicles in the highest (i.e., dirtiest) quartile is -0.28 . The VMT elasticity then falls monotonically with NO_x quartiles. The VMT elasticity is -0.15 in the third quartile, -0.05 in the second quartile, and 0.04 in the first quartile. Similar correlations between emissions and VMT elasticities hold for CO and HCs.

These correlations drive a wedge between the SBO gasoline tax associated with emissions and what we call the “naive” tax, which we define as the tax based only on the *unweighted* average externality across vehicles. We show the SBO gasoline tax is larger, on the order of 50 percent, than the naive gasoline tax in each of the years of our sample. However, we also show that for local pollutants, a uniform gasoline tax performs poorly in eliminating DWL. Across our sample, we estimate the SBO gasoline tax eliminates only 30 percent of DWL associated with the pollutants studied. During the second half of our sample, the SBO gasoline tax eliminates only 25 percent of the DWL.

To determine whether our DWL results extend to externalities beyond vehicle emissions, we next investigate which features of the personal transportation market lead to the failure of the SBO gasoline tax to eliminate a substantial portion of the DWL from emissions of criteria pollutants. First, we show that the failure of the SBO gasoline tax is not simply due to variation in the harmfulness of pollution across counties—a county-specific SBO tax would only remove slightly more DWL than a statewide tax. Instead, the failure of the SBO gasoline tax stems from two related factors: the heterogeneity in local pollution externalities across vehicle vintages and the overall right skew of the distribution of pollution externalities. If these features were not present, a uniform tax would address more than 80 percent of the DWL of local pollution externalities. These features are likely present in other vehicle externalities such as accidents and congestion, suggesting that policymakers will need to implement direct Pigouvian taxes or more complete policies to address the externalities of vehicle usage.

The equity and distributional consequences of SBO taxes are also important—in addition to normative considerations, political support for a tax scheme hinges on equity. Because gasoline demand is income inelastic, gasoline taxes are generally regressive, although there is debate in the literature as to just how regressive such taxes are (Poterba 1991, Chernick and Reschovski 1997). Gasoline expenditures and gasoline taxes paid are relatively constant across the income distribution and necessarily make up a larger share of income at the lower end of the distribution. One might expect that an emissions tax might be more regressive still, especially if poorer households are more likely to own higher polluting vehicles. Our results show the opposite, however. Our empirical model predicts that the average household in *every* income decile would pay a higher percentage of income under the SBO gasoline tax than under an optimal emissions tax. This is more pronounced for households at the bottom of the income distribution. Because dirtier vehicles are more price responsive and have lower VMT *ex ante*, the SBO gasoline tax revenues are higher than under the optimal emissions tax. Furthermore, although on average lower income households are more likely to drive higher polluting vehicles, the

correlation is weak. Vehicles that would have a higher tax burden under an optimal emissions tax are owned by households throughout the income distribution and make up a small minority of vehicles in every income bracket.

Our results have implications for regulating externalities beyond local pollutants emitted by personal automobiles. Indeed, our focus on the local pollution externalities is driven partly by data availability. Increasingly stringent emissions standards for new vehicles have gradually made local pollution less important compared to other externalities of personal automobiles. Nonetheless, both our result that the SBO tax and the naive uniform tax diverge substantially and our result that the SBO tax performs poorly likely transfer to other transportation externalities, and indeed other externalities entirely. For example, congestion externalities also exhibit substantial heterogeneity, with miles traveled during peak hours having much larger external costs, while being less price responsive as well. Our results suggest that the marginal congestion externality is probably much less than the average, and that a gasoline tax will address little of the overall congestion externality. In addition, although the significance of local pollution externalities has declined in the United States, local air quality continues to be an issue in other countries, and our specific results may be useful in those contexts.

We are not the first to analyze the optimal level of gasoline taxes. Parry and Small (2005) calculate the optimal gasoline tax for the United States and United Kingdom accounting for local and global pollution, accidents, congestion, and the inefficiencies associated with income taxes. Our analysis differs in four key respects. First, in general Parry and Small (2005) implicitly assume vehicle externalities are uncorrelated with the sensitivity of each vehicle's demand for gasoline and gasoline prices, whereas we allow for, and find, such correlation.⁴ Second, we account for the possibility that marginal damage of vehicle emissions may vary geographically. Third, Parry and Small (2005) do not estimate the DWL that remains from instituting a gasoline tax, as opposed to the first-best set of optimal emissions taxes, which is one of the main focuses of our paper. Fourth, our focus is more narrow in some respects, as our empirical work focuses on externalities associated with local and global pollution, although our results bear on the external costs related to accidents and congestion.

The closest paper to ours, in terms of our DWL results, is Fullerton and West (2010), who also investigate the amount of DWL eliminated by a uniform gasoline tax. They do so by calibrating a numerical model with approximate miles and emissions obtained by matching inspection data from a small pilot study in California to quarterly gasoline expenditures in the Consumer Expenditure Survey. In contrast, our estimates are based on actual emissions, miles traveled, and gasoline prices for

⁴Parry and Small (2005) do apply an *ad hoc* correction to the congestion component of their optimal tax to account for peak period driving being less price sensitive than nonpeak driving. Although it is surely true that peak period driving is less price sensitive, Parry and Small do not estimate the elasticities, nor the correlation between price responsiveness and congestion externalities. Parry and Small's correction is significantly smaller than the wedge we find between the SBO and naive pollution taxes.

the universe of California vehicles. We find a uniform tax removes much less of the DWL of pollution compared to their calculations.⁵

The paper proceeds as follows. Section I draws on Diamond (1973) to derive the SBO gasoline tax and the amount of remaining DWL. Section II discusses the empirical setting and data. Section III provides descriptive support for the empirical results through graphical analysis. Section IV presents the main empirical model and results on miles driven. Section V estimates empirically the optimal uniform tax and welfare effects, and Section VI presents results on the incidence of gasoline and emissions taxes. Section VII concludes the paper.

I. Optimal Uniform Taxes

In this section, we derive the second-best optimal uniform tax to internalize an externality in the presence of heterogeneity in the externality. We closely follow the model of Diamond (1973) in deriving the optimal tax. We then add more structure to the problem to analytically solve for the amount of remaining DWL.

Consumer h derives utility indirectly from her consumption of a good, α_h , and a numeraire, μ_h , but is also affected by the consumption of others, α_{-h} (the externality). Assuming quasi-linear preferences, consumer h 's utility can be written as

$$(1) \quad U^h(\alpha_1, \alpha_2, \dots, \alpha_h, \dots, \alpha_n) + \mu_h.$$

We assume utility is monotone in own consumption, i.e., $\partial U^h / \partial \alpha_h \geq 0$. This yields demand curves α_h^* , given by

$$(2) \quad \alpha_h^* = \alpha_h(p_g + \tau),$$

where p_g denotes the per-unit price of the good and τ a per-unit tax on purchases.

These assumptions, along with assuming an interior solution for each consumer, lead to the following result.

PROPOSITION 1: *The second-best optimal uniform per-unit tax, τ^* , is (from Diamond 1973)*

$$(3) \quad \tau^* = \frac{-\sum_i \sum_{h \neq i} \frac{\partial U^h}{\partial \alpha_i} \alpha'_i}{\sum_i \alpha'_i},$$

where $\sum_{h \neq i} \partial U^h / \partial \alpha_i$ is the external costs associated with one unit of consumption by individual i and α'_i is the derivative of consumer i 's demand for the good with respect to the price.

⁵ In addition, there is a broad literature aimed at estimating how vehicle owners' driving and scrappage decisions respond to gasoline prices and vehicles' policies, typically using either aggregate data or NHTS survey data. See, for instance, Li, Timmins, and von Haefen (2009); Gillingham (2010); and Jacobsen and van Benthem (2013).

PROOF:

See online Appendix A.

The SBO tax becomes a weighted average of vehicles' externalities where the weights are the derivative of the demand with respect to the tax. If there is a positive correlation between price responsiveness and externalities (in the vehicle context, dirtier cars are more responsive), this correlation will increase the SBO tax.⁶ Note that if consumption of α creates multiple externalities, the SBO tax addressing all externalities will be the sum of the SBO taxes that would address each externality independently. To see this, notice that the externalities appear in equation (1) only through $\partial U^h / \partial \alpha_i$. So long as $\partial U^h / \partial \alpha_i$ is additively separable (i.e., the externalities do not interact), the terms can be rearranged to express the total SBO tax as the sum of the SBO taxes for the component externalities.

As Diamond explicitly discusses, there is no requirement that all of the α_i terms must be negative.⁷ In the context of personal automobiles, this could occur if households hold multiple vehicles, as they may shift miles from their least fuel efficient, most polluting vehicle to a cleaner, more efficient vehicle.⁸

The presence of heterogeneity in the externality also implies that a uniform tax will not achieve the first-best outcome. A uniform tax will under-tax high externality agents and overtax low externality agents. We extend Diamond (1973) to solve for the amount of DWL remaining in the presence of an SBO tax applied to a market with heterogeneous externalities. We make assumptions on the distribution of the externality in order to obtain a closed-form solution with intuitive content. When we turn to our empirical application, we relax these assumptions entirely and leverage the empirical distribution.

We start with the case where demand elasticities and emissions are uncorrelated and then relax this assumption.

PROPOSITION 2: *Suppose consumers are homogeneous in their demand for the good, but individuals' per-unit externalities differ. In particular, let α' denote the derivative of demand with respect to price.*

If the distribution of the per-unit externality, E , is log normal, with probability density function

$$(4) \quad \varphi(E_i) = \frac{1}{E_i \sqrt{2\sigma_E^2}} \exp\left(-\frac{(E_i - \mu_E)^2}{2\sigma_E^2}\right),$$

⁶ As an intuitive example, imagine the case where there are only two vehicle types. The first emits little pollution, while the second is dirtier. Also imagine the clean vehicles are completely price insensitive, while the dirty vehicles are price sensitive. The naive tax would be calculated based on the average emissions of the two vehicle types. However, the marginal emission is the emissions rate of the dirty vehicles; the clean cars are driven regardless of the tax level. In this case, we can achieve first best by setting the tax rate at the externality rate of the dirty vehicle. There is no distortion to owners of clean vehicles since their demand is completely inelastic, so we can completely internalize the externality to those driving the dirty vehicles.

⁷ However, second-best optimal tax loses the interpretation as a weighted average if some α_i terms are positive.

⁸ Also note that the elasticity of gasoline consumption with respect to price accounts for households selling or scrapping their vehicles and buying different ones. That is, if the gasoline tax increases the scrappage rate of some vehicles, then the relevant derivative of the externality with respect to price is the expected change in gasoline consumption, not the change in gasoline demand, conditional on survival.

the DWL absent any market intervention will be given as:

$$D = \frac{1}{2\alpha'} e^{2\mu_E + 2\sigma_E^2}.$$

PROOF:

See online Appendix A.

This leads to the following calculation of remaining DWL under the SBO tax.

PROPOSITION 3: *Under the assumptions in Proposition 2, the ratio of remaining DWL after the tax is imposed to the DWL absent the tax is*

$$(5) \quad R = \frac{D - \frac{e^{2\mu_E + \sigma_E^2}}{2\alpha'}}{D} = 1 - \frac{e^{2\mu_E + \sigma_E^2}}{e^{2\mu_E + 2\sigma_E^2}} = 1 - e^{-\sigma_E^2}.$$

PROOF:

See online Appendix A.

With externalities uncorrelated with the demand for the good, the remaining DWL from a uniform tax depends only on the shape parameter of the externality distribution. The larger σ_E^2 is, the wider and more skewed will the distribution of the externality be, causing the uniform tax to “overshoot” the optimal reduction in consumption for more individuals.

If demand is not homogeneous and in fact is correlated with per-unit externalities, the calculation changes. Let α'_h denote the derivative of the demand associated with individual i with respect to the price of the good. For ease, define $B_i = 1/\alpha'_h$, and assume that B_h is distributed lognormal with parameters μ_B and σ_B^2 . Define ρ as the dependence parameter of the bivariate lognormal distribution (the correlation coefficient of $\ln E$ and $\ln B$). We then have the following result.

PROPOSITION 4: *When B_h and E_h are distributed lognormal with dependence parameter ρ , the SBO tax is*

$$\tau^* = e^{\mu_E + \frac{\sigma_E^2}{2} + \rho\sigma_E\sigma_B}.$$

PROOF:

See online Appendix A.

As we would expect, the SBO tax does not depend on the scale of the elasticity distribution, only on the extent to which externalities are correlated with elasticities. We can then calculate the amount of remaining DWL under both the naive and SBO tax.

PROPOSITION 5: *When B_i and E_i are distributed lognormal with dependence parameter ρ , the ratios of the remaining DWL after the SBO tax to the original DWL will be*

$$(6) \quad R(\tau^*) = 1 - e^{-\sigma_E^2},$$

and the ratios of the remaining DWL after the naive uniform tax to the original DWL will be

$$(7) \quad R(\tau_{naive}) = 1 - e^{-\sigma_E^2} (2e^{-\rho\sigma_E\sigma_B} - e^{-2\rho\sigma_E\sigma_B}).$$

PROOF:

See online Appendix A.

As we would expect, the optimal tax correctly accounts for the correlation between the externality and demand responses, and thus the remaining DWL depends only on the variance and skewness of the externality distribution. However, in the presence of correlation, the naive tax reduces less of the DWL from the externality, reducing it by a proportion related to the degree of correlation and the spread of the two distributions. The term in parentheses in equation (7) is strictly less than one and strictly greater than zero if $\rho > 0$, but may be negative if $\rho < 0$ and the shape parameters are sufficiently large.

II. Empirical Setting

A. Data

Our empirical setting is the California personal transportation market. We bring together a number of large datasets. Our analysis is primarily based upon the universe of emissions inspections from 1996 to 2010 from California's vehicle emissions testing program, the Smog Check Program, which is administered by the California Bureau of Automotive Repair (BAR). A vehicle appears in the data for a number of reasons. First, vehicles more than four years old must pass a Smog Check within 90 days of any change in ownership. Second, in parts of the state (details below), an emissions inspection is required every other year as a prerequisite for renewing the registration on a vehicle that is six years or older. Third, a test is required if a vehicle moves to California from out-of-state. Vehicles that fail an inspection must be repaired and receive another inspection before they can be registered and driven in the state. There is also a group of exempt vehicles. These are: vehicles of 1975 model year or older, hybrid and electric vehicles, motorcycles, diesel-powered vehicles, and large natural gas powered trucks.⁹

These data report the location of the test, the unique vehicle identification number (VIN), odometer reading, the reason for the test, and test results. We decode the VIN to obtain the vehicles' make, model, engine, and transmission. Using this information, we match the vehicles to EPA data on fuel economy. Because the VIN decoding is only feasible for vehicles made after 1981, our data are restricted to these models. We also restrict our sample to 1998 and beyond, given large changes that occurred in the Smog Check Program in 1997. This yields roughly 76 million observations.

⁹In 2010, California began requiring Smog Checks for diesel-powered vehicles of model years 1997 and newer. We have fewer than 500 diesel-powered vehicles in our data, however.

The Smog Check data report measurements for NO_x and HCs in terms of parts per million and CO levels as a percentage of the exhaust, taken under two engine speeds.¹⁰ As we are interested in the quantity of emissions, the more relevant metric is a vehicle's emissions per mile. We convert the Smog Check emissions readings into emissions per mile using conversion equations developed by Sierra Research for the California Air Resources Board in Morrow and Runkle (2005). The conversion equations are functions of both measurements of all three pollutants, vehicle weight, model year, and truck status. For more details on cleaning the Smog Check data, including the conversion equations for the three criteria pollutants, see online Appendix B.

As part of our simulation exercise, we also use data obtained from CARFAX Inc. to estimate scrappage decisions. These data contain the date and location of the last record of the vehicle reported to CARFAX for 32 million vehicles in the Smog Check data. This includes registrations, emissions inspections, repairs, import/export records, and accidents. Because the CARFAX data include import/export records, we are able to correctly classify the outcomes of vehicles which are exported to Mexico as censored, rather than scrapped, thus avoiding the issues identified in Davis and Kahn (2010).

For a subset of our Smog Check data, we are able to match vehicles to households using confidential data from the California Department of Motor Vehicles (DMV). These data track the registered address of every vehicle in the state, with one address given for each year. We use the registration information to attach demographic information on income from US census data. Online Appendix C discusses the process of cleaning the registration data. The DMV data are only available for the years 2000 to 2008.

For a portion of our analysis, we use data from the 2009 National Highway Transportation Survey, which contains information on household vehicles, annual VMT, and household income for a sample of households.

Finally, we use gasoline prices from EIA's weekly California average price series to construct average prices between inspections.

Table 1 reports means and standard deviations of the main variables used in our analysis, for all observations and broken down by vehicle age and by year of Smog Check. The average fuel economy of vehicles in our sample is 23.5 MPG, with fuel economy falling over the period of the sample. The change in the average dollars per mile has been dramatic, almost tripling between 1998 and 2008. The dramatic decrease in vehicle emissions is also clear in the data, with average per-mile emissions of HCs, CO, and NO_x falling considerably from 1998 to 2008. The tightening of standards has also meant that more vehicles fail Smog Checks late in the sample, although some of this is driven by the aging of the vehicle fleet.¹¹

¹⁰ HCs, as measured by the Smog Check Program, are similar to, but slightly different from, Non-Methane Organic Gases (NMOG), the pollutant measured by the EPA and California Air Resources Board (CARB)'s testing standards for new vehicles. HCs include emissions of methane, but do not include oxygenated compounds such as aldehydes. For a primer on the differences from CARB, see <https://www.arb.ca.gov/ei/speciate/orgtermnsmhctctogsummary.pdf>. Regardless of the metric, emissions of this type from vehicle sources largely consist of uncombusted gasoline fumes.

¹¹ Although the failure rate increases with time, the average emissions rate for vehicles of a given model year stays relatively constant over our sample. That is, emissions do not, on average, rise with vehicle age, holding

TABLE 1—SUMMARY STATISTICS

	All	Vehicle age			Year	
		4–9	10–15	16–28	1998	2008
Weighted fuel economy	23.53 (5.320)	23.30 (5.235)	23.70 (5.331)	23.80 (5.507)	24.27 (5.478)	23.07 (5.169)
Average \$/mile	0.0973 (0.0416)	0.0928 (0.0403)	0.0977 (0.0417)	0.109 (0.0426)	0.0581 (0.0134)	0.143 (0.0361)
Odometer (00000s)	1.214 (0.605)	0.932 (0.454)	1.376 (0.567)	1.626 (0.688)	1.023 (0.528)	1.323 (0.622)
Average VMT/day	24.50 (95.96)	30.23 (43.29)	23.09 (101.2)	15.64 (148.9)	29.48 (65.64)	22.62 (28.24)
Grams/mile HC	0.762 (1.177)	0.219 (0.270)	0.739 (1.019)	2.017 (1.670)	1.403 (1.506)	0.542 (1.022)
Grams/mile CO	5.525 (12.91)	0.510 (1.646)	4.814 (10.90)	18.15 (20.35)	12.44 (18.97)	3.488 (10.49)
Grams/mile NO _x	0.664 (0.638)	0.317 (0.303)	0.731 (0.599)	1.297 (0.728)	1.042 (0.904)	0.516 (0.547)
Failed Smog Check	0.0868 (0.282)	0.0435 (0.204)	0.106 (0.307)	0.165 (0.371)	0.0515 (0.221)	0.0992 (0.299)
Average HH income	48066.3 (17,031.0)	49,955.1 (17,685.0)	47,117.3 (16,556.3)	44,970.8 (15,555.7)	49,768.5 (17,952.5)	47,778.3 (16,791.9)
Truck	0.385 (0.487)	0.406 (0.491)	0.368 (0.482)	0.367 (0.482)	0.322 (0.467)	0.426 (0.494)
Vehicle age	10.68 (4.587)	6.694 (1.615)	12.14 (1.686)	18.54 (2.478)	9.244 (3.552)	11.77 (4.854)
Observations	76,510,820	34,713,936	29,775,806	12,008,157	4,172,978	5,849,644

Notes: Statistics are means with standard deviations presented below in parentheses. Weighted fuel economy is from EPA. Dollars per mile is the average gasoline price from EIA in between the vehicle's current and previous Smog Checks, divided by the vehicle's fuel economy. Average household income is taken from the 2000 census ZCTA where the smog check occurred. The dataset used for this table contains one observation per vehicle per year in which a smog check occurred.

B. Automobiles, Criteria Pollutants, and Health

The vehicle inspection data report emissions of three criteria pollutants: NO_x, HCs, and CO. All three of these directly result from the combustion process within either gasoline or diesel engines. Both NO_x and HCs are precursors to ground-level ozone, but all three have been shown to have negative health effects on their own.¹²

While numerous studies have found links between exposure to the ozone or the three criteria pollutants and health outcomes, the mechanisms are still uncertain. These pollutants, as well as the ozone, may directly impact vital organs or indirectly cause trauma. For example, CO can bind to hemoglobin, thereby decreasing the amount of oxygen in the bloodstream. High levels of CO have also been linked to heart and respiratory problems. NO_x reacts with other compounds to create nitrate

vintage constant. See online Appendix Figure A.1. This is likely at least in part due to the Smog Check Program itself, both directly (by fixing cars with high emissions) and indirectly (by forcing high emissions cars off the road).

¹²CO has also been shown to speed up the smog-formation process. For early work on this, see Westberg, Cohen, and Wilson (1971).

aerosols, which are fine-particle particulate matter (PM). PM has been shown to irritate lung tissue, lower lung capacity, and hinder long-term lung development. Extremely small PM can be absorbed through the lung tissue and cause damage on the cellular level. On their own, HCs can interfere with oxygen intake and irritate lungs. Ground-level ozone is a known lung irritant, has been associated with lowered lung capacity, and can exacerbate existing heart problems and lung ailments such as asthma or allergies.

III. Preliminary Evidence

One of the main driving forces behind our empirical results is how gasoline demand elasticities for different vehicles vary systematically with emissions levels. In this section, we present evidence that significant variation exists in terms of vehicle externalities within a year, across years, and even within the same vehicle type (make, model, model year, etc.) within a year. Further, simple statistics, such as the average miles traveled by vehicle type, suggest that elasticities may be correlated with externalities.¹³

Figure 1 plots the distributions of NO_x, HCs, and CO emissions in 1998, 2004, and 2010. The distribution of criteria pollutant emissions tends to be right skewed in any given year, with a standard deviation equal to roughly one to three times the mean, depending on the pollutant. The skewness implies that some vehicles on the road are quite dirty relative to the mean vehicle. Over time, the distribution has shifted to the left, as vehicles have gotten cleaner, but the range remains. A major reason for the leftward shift is the progressive tightening in the limits on emissions per mile imposed by the US EPA, such that more recent model years have substantially lower emissions per mile than vehicles produced in the 1980s and early 1990s. As these newer, cleaner vehicles have comprised a greater share of the fleet, average emissions have fallen.

The variation in emissions is not only driven by the fact that different types of vehicles are on the road in a given year, but also variation within the *same* vehicle type, defined as a make, model, model year, engine, number of doors, and drivetrain combination. To see this, Figure 2 plots the distributions of emissions for the most popular vehicle/year in our sample, the 2001 four-door Toyota Corolla in 2009. The vertical line is at the mean of the distribution. Here, again, we see that even within the same vehicle type in the same year, the distribution is wide and right skewed. The distribution of HCs is less skewed, but the standard deviation is 25 percent of the mean. CO is also less skewed and has a standard deviation that is 36 percent of the mean. Across all years and vehicles, the mean emissions rate of a given vehicle in a given year, on average, is roughly four times the standard deviation for all three pollutants (online Appendix Table A.1).

To understand how the distribution within a given vehicle type changes over time, Figure 3 plots the distribution of the 1995 3.8L, front-wheel drive, Ford Windstar in

¹³ We are not the first to document the large variation across vehicles in emissions. See, for example, Kahn (1996). Instead, our contribution is in finding a link between elasticities and emissions.

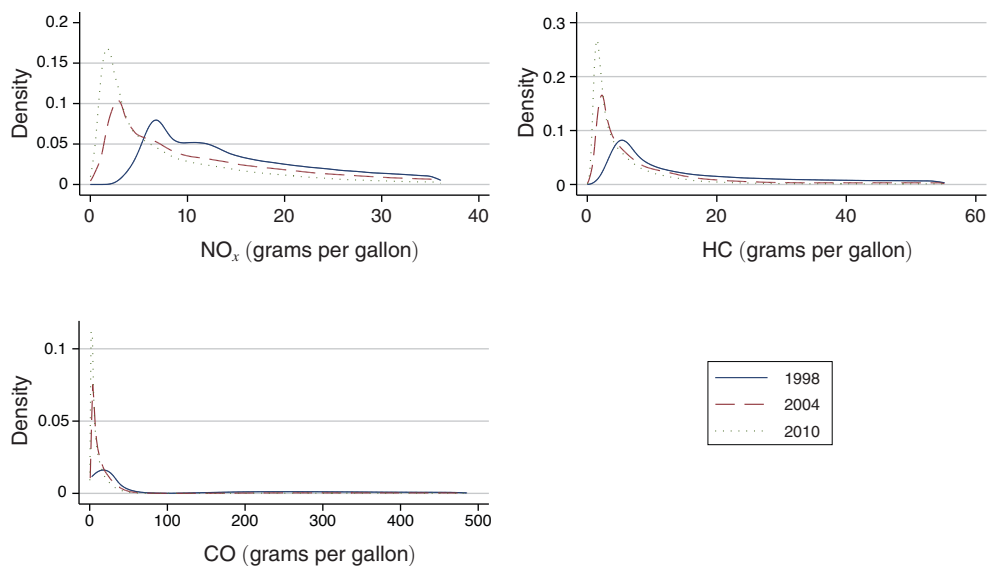


FIGURE 1. DISTRIBUTION OF THREE CRITERIA POLLUTANT EMISSIONS ACROSS ALL VEHICLES IN 1998, 2004, AND 2010
(observations above the ninetieth percentile are omitted)

1999, 2001, 2004, and 2007.¹⁴ These figures suggest that over time the distributions shift to the right, become more symmetric, and the standard deviation grows considerably, relative to the mean. Across all vehicles, the ratio of the mean emission rate of NO_x and the standard deviation of NO_x has increased from 3.16 in 1998 to 4.53 in 2010. For HCs, this increased from 3.59 to 5.51, and for CO, the ratio increased from 3.95 to 5.72.

These distributions demonstrate significant variation in emissions across vehicles and within vehicle type, and thus significant scope for meaningful emissions-correlated variation in elasticities along those lines.¹⁵ We next present suggestive evidence that VMT elasticities may be correlated with emissions. For each criteria pollutant within each calendar year, we rank vehicles by their observed emissions per mile and divide them into quartiles. We do the same for fuel economy. Next, for each quartile year, we compute the median annual VMT and plot how this has changed over our sample, normalized by the 1998 level for each quartile. Figure 4 foreshadows our results on VMT elasticities and externalities. For each pollutant, we see that the dirtiest quartile saw the largest decreases in miles driven during the run-up in gasoline prices from 1998 to 2008 when prices increased from

¹⁴ We chose this vehicle because the 1995 3.8L, front-wheel drive Ford Windstar in 1999 is the second-most popular entry in our data, and it is old enough that we can track it over four two-year periods.

¹⁵ Because of the way we handle multiple tests of a given vehicle within a year, our distributions likely understate the degree of on-road heterogeneity. In order for a vehicle to be registered, the vehicle must pass a Smog Check. In our data, we see multiple tests of the same vehicle over a short time frame. We use the final test, which will necessarily have been passed, for our calculations. Furthermore, our calculations omit unregistered vehicles, many of which are likely to have high emissions.




FIGURE 2. DISTRIBUTION OF THREE CRITERIA POLLUTANT EMISSIONS OF A 2001 FOUR-DOOR, 1.8L, TOYOTA COROLLA IN 2009 (observations above the ninetieth percentile are omitted)

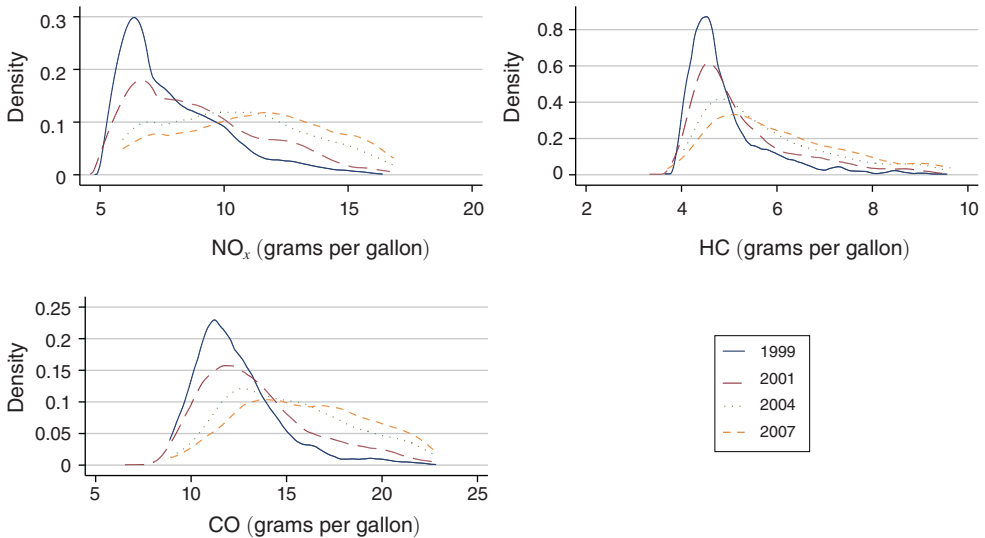


FIGURE 3. DISTRIBUTION OF THREE CRITERIA POLLUTANT EMISSIONS OF A 1995 3.8L, FWD, FORD WINDSTAR IN 1999, 2001, 2005, AND 2009 (observations above the ninetieth percentile are omitted)

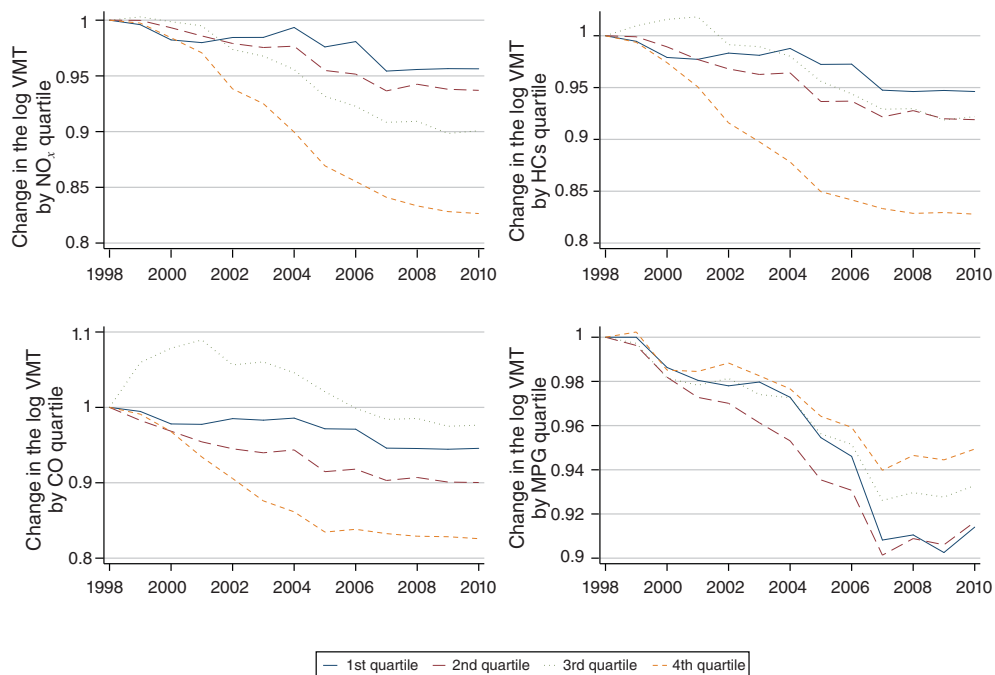


FIGURE 4. CHANGE IN THE LOG OF VMT OVER SAMPLE BY POLLUTANT QUARTILE

roughly \$1.35 to \$3.20.¹⁶ The ordering of the relative decreases suggests that dirtier vehicles may have been more responsive over this period.

IV. Vehicle Miles Traveled Decisions

We now estimate how changes in gasoline prices affect decisions about vehicle miles traveled (VMT) and how this elasticity varies with vehicle characteristics. Our empirical approach mirrors Figure 4. For each vehicle receiving a biennial Smog Check, we calculate average daily miles driven and the average gasoline price during the roughly two years between Smog Checks. Obviously vehicle owners with more fuel-efficient vehicles will respond less to changes in the per-gallon gasoline price, and to abstract from this, we specify the elasticity with respect to the price in dollars per mile (DPM) by dividing the average per-gallon price by fuel economy in gallons per mile. Thus, the price faced by each vehicle's owner will vary both with the exact period in between Smog Checks and with the specific vehicles' fuel economy. We then allow the elasticity to vary based on the emissions of the vehicle. We begin by estimating

$$(8) \quad \ln(VMT_{ijgt}) = \beta \ln(DPM_{ijgt}) + \gamma X_{it} + \mu_t + \mu_j + \mu_g + \mu_v + \varepsilon_{igt},$$

¹⁶ The levels also differ. Online Appendix Figure A.2 plots the median of daily miles traveled across our sample split up by the emissions quartile of the vehicle.

where i indexes vehicles, j vehicle types, g geographic locations, t time, and v vehicle age, or vintage; DPM_{ijgt} is the average gasoline price per mile faced by vehicle i between time t and the date of the previous Smog Check, X_{it} is a vector of time-varying covariates, and ε_{igt} is a residual.¹⁷ The covariate vector X_{it} includes an indicator for whether the vehicle is a truck, a quadratic time trend in days, and a sixth order polynomial in odometer reading, using the odometer reading from the previous Smog Check (i.e., the odometer reading before the choice of VMT). Our baseline specification assumes that gasoline prices are exogenous to individual driving decisions. Such an assumption is common in the literature, as gasoline prices are largely driven by movements in the world price of crude oil, which saw dramatic changes during the 2000s for reasons unrelated to driving choices in California.¹⁸ However, we have also estimated our main analyses instrumenting for DPM with the Brent Crude oil price, and we obtain very similar results. Online Appendix D describes a variety of robustness checks, including using a complete set of month-by-year fixed effects, thus relying on cross-sectional variation in gas prices, as opposed to time-series variation. All of our robustness checks yield results that are qualitatively similar to our baseline specification.

We begin by including demographic characteristics by the zip code of Smog Checks and year and vintage fixed effects. We then progressively include finer vehicle-type fixed effects by including make, then make/model/model year/engine, and finally individual-vehicle fixed effects. We also differentiate the influence of gasoline prices by vehicle attributes related to the magnitude of their negative externalities—criteria pollutants, CO₂ emissions, and weight.

We allow the VMT elasticity to vary with the magnitude of their externalities in two ways. For both approaches, we begin by ranking vehicles within each calendar year by their emissions per mile of NO_x, HCs, CO, fuel economy, or vehicle weight in pounds. In one set of specifications, we split vehicles up by the quartile of these variables and allow each quartile to have a separate β .¹⁹ In another set, we include a linear interaction of centiles of these variables and the log of gasoline prices in dollars per mile.

Table 2 shows our results, focusing on NO_x. The changes from models 1 to 4 illustrate the importance of controlling for vehicle-type fixed effects. Initially, the average elasticity falls from -0.269 to -0.123 when including make fixed effects, but then rises when including finer vehicle-type fixed effects. Model 4 includes

¹⁷ The fuel economy in gallons per mile used to calculate our DPM variable uses the standard assumption that 45 percent of a vehicle's miles driven are in the city and 55 percent are on the highway. This is the standard approach used by the EPA for combined fuel economy ratings.

¹⁸ See, for example, Busse, Knittel, and Zettelmeyer (2013).

¹⁹ We note that in addition to accounting for relative emissions, specifying emissions levels as quartiles reduces the effect of random noise from the Smog Check data. The Smog Check emissions testing itself is an estimate of true emissions, and measured emissions can vary even without changes in actual emissions. We can test the extent of this by comparing vehicles that pass an inspection for a change of ownership shortly before or after a biennial inspection. Although the changes between closely spaced inspections like this are mean zero, the average absolute change in the emissions reading is about 10-15 percent of a standard deviation. Specifying emissions as quartiles will mute this variance, since it is relatively unlikely that an unusually high or low emissions reading will move a vehicle to a different quartile than the true emissions level would indicate.

individual-vehicle fixed effects yielding an average elasticity of -0.134 .²⁰ In models 5 and 6, we examine heterogeneity with vehicle fixed effects. Model 5 includes interactions with quartiles of NO_x . The VMT elasticity for the cleanest vehicles, quartile one, is positive at 0.043, while the VMT elasticity for the dirtiest vehicles is twice the average elasticity at -0.280 . To put these numbers in context, the average per-mile NO_x emissions of a quartile one vehicle is 0.163 grams, while the average per-mile NO_x emissions of a quartile four vehicle is 1.68 grams. As a further point of comparison, the Tier 2 emissions standards established by the EPA in 2000 called for fleet average NO_x emissions of 0.3 grams per mile for passenger cars of model years 2004–2006 and 0.07 grams per mile starting with model year 2007.²¹ Model 6 assumes the relationship is linear in centiles of NO_x and finds that each percentile increase in the per-mile NO_x emission rate is associated with an elasticity 0.001 larger in absolute value from a base of essentially zero.

We find similar patterns across the other externalities. The range of the estimated VMT elasticities is somewhat larger when using quartiles of HCs and CO emissions compared to NO_x , with the dirtiest quartiles around -0.30 and the cleanest around 0.05. For CO_2 , the cleanest vehicles are those with the highest fuel economy, and here, we see the least fuel-efficient vehicles having a VMT elasticity of -0.183 , compared to -0.108 for vehicles with fuel economy in the highest quartile. We observe some heterogeneity in the VMT elasticity across vehicle weights as well, although it is smaller than the other externalities. For the full set of results, see online Appendix Table A.2.

The SBO gasoline tax is not (necessarily) affected by the mechanism behind the heterogeneity in elasticities we observe, and a full exploration of the mechanism is outside the scope of the present paper. In a related working paper (Knittel and Sandler 2013), we investigate several potential mechanisms, including multiple vehicle households switching VMT to a cleaner vehicle, older vehicles (which necessarily have higher emissions) being more responsive, and low-income consumers being more responsive and owning dirtier vehicles. In Knittel and Sandler (2013), we find that each of these explanations account for some of the heterogeneity in elasticities, but that a portion remains unexplained.

V. Efficiency of the Second-Best Optimal Gasoline Tax

In this section, we use our empirical results on driving responsiveness to evaluate the efficiency of SBO taxes in the presence of heterogeneity. We first calculate the SBO gasoline tax to address the externalities from emissions of NO_x , HCs, and CO and compare this to the naive tax equal to the average externality. We then compare the remaining DWL left over from these second-best taxes to the optimal outcome obtained by a Pigouvian tax on emissions. Although we focus on the local pollution

²⁰ Our average elasticity is larger than that found in Hughes, Knittel, and Sperling (2008), reflecting the longer run nature of our elasticity.

²¹ The Tier 2 standards also required light duty trucks (including SUVs) to have a fleet average of 0.2 grams per mile of NO_x emissions for model years 2004–2008 and meet the same 0.07 grams per mile standard starting with model year 2009. Tier 2 standards have since been superseded by tighter Tier 3 standards, which began taking effect with model year 2017.

TABLE 2—VEHICLE MILES TRAVELED, DOLLARS PER MILE, AND NITROGEN OXIDES (*quartiles by year*)

	Model 1 (1)	Model 2 (2)	Model 3 (3)	Model 4 (4)	Model 5 (5)	Model 6 (6)
ln(DPM)	-0.269 (0.044)	-0.123 (0.038)	-0.183 (0.027)	-0.134 (0.022)		-0.038 (0.028)
ln(DPM) × NO _x Q1					0.043 (0.021)	
ln(DPM) × NO _x Q2					-0.054 (0.022)	
ln(DPM) × NO _x Q3					-0.152 (0.025)	
ln(DPM) × NO _x Q4					-0.280 (0.028)	
ln(DPM) × NO _x Centile						-0.001 (0.000)
NO _x Q2					0.216 (0.663)	
NO _x Q3					-1.742 (0.881)	
NO _x Q4					-2.417 (1.003)	
NO _x Centile						-0.001 (0.001)
Truck	0.054 (0.033)	0.057 (0.045)	0.005 (0.055)			
Time trend	-0.244 (0.037)	-0.314 (0.024)	-0.278 (0.015)	-0.035 (0.028)	-0.057 (0.032)	-0.062 (0.025)
Time trend squared	0.002 (0.000)	0.002 (0.000)	0.002 (0.000)	0.000 (0.000)	0.000 (0.000)	0.001 (0.000)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Vintage fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Sixth-order polynomial in lagged odometer	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Make fixed effects	No	Yes	No	No	No	No
Vin prefix fixed effects	No	No	Yes	No	No	No
Vehicle fixed effects	No	No	No	Yes	Yes	Yes
Observations	36,387,455	36,387,455	36,387,455	36,387,455	29,779,909	29,779,909
R ²	0.210	0.218	0.143	0.121	0.117	0.118

Notes: Each observation is a vehicle's Smog Check inspection. The dependent variable is the log of the average daily vehicle miles traveled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy. Quartiles and centiles of NO_x are based on rankings of emissions per mile within the calendar year in which the Smog Check occurs. Standard errors clustered by vehicle make are reported in parentheses.

externalities of driving, our results will extend to any externality with a similar pattern of heterogeneity, and we discuss the implications of our results to SBO taxes for other driving externalities such as accidents and congestion.

A. Tax Calculations

We calculate the naive tax per gallon of gasoline as the simple average of the externality per gallon caused by all vehicles on the road in California in a particular

year. We value the externalities imposed by NO_x and HCs using the marginal damages calculated by Muller and Mendelsohn (2009), based on the county in which each vehicle has its Smog Check.²² For CO, we use the median marginal damage estimate from Matthews and Lave (2000).

Let the marginal damage per gram of pollutant p in county c be θ_c^p , with emissions rates in grams per mile by vehicle i of ϵ_i^p . Then the externality per mile of vehicle i , E_i is

$$(9) \quad E_i = \theta_c^{HC} \cdot \epsilon_i^{HC} + \theta_c^{\text{NO}_x} \cdot \epsilon_i^{\text{NO}_x} + \theta_c^{\text{CO}} \cdot \epsilon_i^{\text{CO}}.$$

The naive tax in year y will then be

$$(10) \quad \tau_{naive}(y) = \frac{1}{N^y} \sum_{i=1}^{N^y} \frac{E_i}{MPG_i},$$

where N^y denotes the number of vehicles on the road in year y and MPG_i denotes the fuel economy rating of vehicle i . In practice, since the stock of vehicles represented in the Smog Check data in any given year will be less than the total stock of vehicles in the vehicles fleet, we weight each Smog Check observation by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

Following Proposition 1, we calculate the SBO gasoline tax, taking into account the heterogeneity in both levels of the externality and the responsiveness to gasoline prices. We estimate a regression similar to equation (8), but allow the elasticity of VMT with respect to DPM to vary over quartiles of all three criteria pollutants, fuel economy, vehicle weight, and three groups of vehicle age. For more details, see online Appendix E. Let the group-specific elasticity for vehicle i be β_i^q , where q indexes cells by HC emissions, NO_x emissions, CO emissions, MPG, weight, and age, with the externalities again in quartiles by year. Further, let the average price per gallon and the quantity of gasoline consumed per year in gallons in year y be P_y^g and Q_y^g , respectively.²³ Then the optimal tax in year y will be

$$(11) \quad \tau^*(y) = \frac{-\sum_h \sum_{i \neq h} \frac{\partial U^h}{\partial \alpha_i} \alpha_i'}{\sum_h \alpha_h'},$$

²²Note that while we use Muller and Mendelsohn's county-specific damages for completeness, county-level variation in damages is of relatively little importance. Most of the variation in vehicle emissions is between vehicles within counties and even within zip codes. Online Appendix Table A.7 demonstrates this, presenting the total within-zip code and between-zip code standard deviations of the Smog Check emissions readings for each of the criteria pollutants we study. For all three pollutants, the average within-zip code standard deviation is at least 83 percent of the total standard deviation. The relative unimportance of county-level variation is also demonstrated later in Table 4, where we find that a county-specific SBO tax is a little better than a statewide tax. Note, also, that the values used in this paper differ from those used in the published version of Muller and Mendelsohn (2009). The published values were calculated using incorrect baseline mortality numbers that were too low for older age groups. Using corrected mortality data increases the marginal damages substantially. We are grateful to Nicholas Muller for providing updated values and to Joel Wiles for bringing this to our attention.

²³Again, we also weight vehicles based on the number of vehicles of that age and class that appear in the fleet as a whole; see online Appendix E. We also account for vehicle owners' decisions to scrap their vehicles to the extent these are affected by gasoline prices. Online Appendix G discusses the details and results of this exercise. To summarize, we allow gasoline price to affect scrappage decisions and allow this to vary over emissions profiles and

with

$$(12) \quad \alpha_i^t = -\beta_i^q \cdot \frac{Q_i^y}{P_i^y},$$

and

$$(13) \quad \frac{\partial U^h}{\partial \alpha_i} = \frac{E_i}{MPG_i}.$$

Table 3 shows the naive and SBO taxes for each year from 1998 to 2008. The naive tax would be 61.5 cents per gallon of gasoline consumed in 1998, while the SBO tax is 91 cents, 48 percent higher. The ratio of the naive and SBO gasoline tax increases even as the levels of the externalities decline over time. From 2000 on, the SBO gasoline tax is at least 50 percent larger than the naive tax in each year.

B. Welfare with Uniform Taxes

We have shown that because of the correlation between elasticities and externality rates, the SBO gasoline tax is much higher than the naive tax calculated as the average of per-gallon externalities. However, even the optimal uniform tax is still a second-best policy. Because of the heterogeneity in externality levels, the highest externality individuals will be taxed by less than their external costs to society, leaving the remaining DWL. Those with lower externalities than the weighted average will be taxed too much, overshooting the optimal quantity of consumption and creating more DWL. We now calculate how far both the naive and SBO gasoline taxes are from the optimal Pigouvian tax for the local pollutant emissions from driving. We also examine how much of any shortfall is due to unique features of the local pollution externality, rather than general factors that would apply to other externalities.

Simulation Results.—We begin by approximating the ratios of DWL with and without the taxes using our data to simulate the change in miles driven and thus in gasoline consumption from a tax. Let $miles_i^y$ be the actual average miles per day traveled by vehicle i between its last Smog Check and the current one, observed in year y , and let $\widehat{miles}_i^y(\tau)$ be the miles per day that a vehicle would travel if the average price of gasoline were raised by a tax of $\$ \tau$ per gallon that is fully passed through to consumers. In principle, a higher gasoline tax might induce some consumers to drive less aggressively to save fuel, which might also have effects on emissions of local pollutants.²⁴ However, we assume that vehicle fuel economy and

vintages. We find that the main source of heterogeneity occurs across vintages; specifically, increases in gasoline prices decrease the hazard rate of very old vehicles, but increase the hazard rate of middle-aged vehicles. Because emissions of criteria pollutants are positively correlated with age, this has the effect of increasing criteria pollutants, although the aggregate effect is small.

²⁴The direction of any such effect is hard to predict. Lower fuel consumption overall would likely reduce emissions per mile with all else equal. However, changing driving patterns would also influence the operating

TABLE 3—AVERAGE AND MARGINAL POLLUTION EXTERNALITY

	Average externality (¢/gallon)	Marginal externality (¢/gallon)
1998	61.48	91.27
1999	54.78	81.62
2000	48.55	74.31
2001	40.96	64.29
2002	34.18	54.09
2003	28.77	46.89
2004	24.31	39.26
2005	21.25	33.95
2006	18.61	29.52
2007	16.23	25.81
2008	14.36	22.84

Notes: Average externality is the simple average of damages from emissions of criteria pollutants produced by each car in each year, divided by fuel usage. We refer to a tax on the average externality as the “naive tax.” The marginal externality is computed as the weighted average of externality per gallon, using the negative slope of the vehicle’s demand curve as the weight. A tax on the marginal externality is the SBO gasoline tax. Both calculations also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole. The dollar figures are inflation adjusted to year 2008.

per mile emissions are constant despite the hypothetical tax. We approximate DWL as a triangle, such that the ratio of interest is

$$r(\tau) = \frac{\sum_i \frac{1}{2} \cdot \left| \frac{\text{miles}_i^y - \widehat{\text{miles}}_i^y(\tau)}{\text{MPG}_i} \right| \cdot \left| \frac{E_i}{\text{MPG}_i} - \tau \right|}{\sum_i \frac{1}{2} \cdot \left| \frac{\text{miles}_i^y - \widehat{\text{miles}}_i^y\left(\frac{E_i}{\text{MPG}_i}\right)}{\text{MPG}_i} \right| \cdot \frac{E_i}{\text{MPG}_i}}.$$

The fully optimal tax would have a ratio of zero, while a tax that actually increased the DWL from gasoline consumption would be greater than one. We note that, given other distortions in the market such as other taxes and externalities, the initial DWL might be more correctly approximated with a trapezoid, rather than a triangle. That is, some DWL from other inefficiencies would exist even with the first-best Pigouvian tax for local pollution. We return to the issue of other externalities in the next subsection. Note also that we have not fully accounted for the extensive margin in our simulation—that is, on the number and type of vehicles on the road. We account for scrappage and have a rough adjustment for purchases of new cars, but we cannot account for changes in technology or substitution patterns between old and new cars that might be influenced by a permanent tax on gasoline.²⁵

Table 4 shows the ratios of DWL with various SBO taxes to the DWL with no additional tax. The first two columns show ratios for a statewide tax based on the

temperature of vehicles’ engines with additional affects on emissions, particularly of NO_x. NO_x emissions are primarily limited by the catalytic converter, which operates more efficiently when warm.

²⁵For instance, Archsmith et al. (2017) shows that gasoline prices influence substitution patterns when households replace one vehicle in their portfolio, and leverage this to identify substitution preferences.

TABLE 4—RATIOS OF DWL WITH TAX TO DWL WITH NO TAX

	Baseline		No between-county variation		No between-vintage variation		Total DWL \$
	Naive	SBO	Naive	SBO	Naive	SBO	
1998	0.616	0.568	0.573	0.523	0.348	0.341	196,466,743.8
1999	0.636	0.577	0.592	0.529	0.330	0.325	158,104,022.2
2000	0.635	0.583	0.587	0.532	0.320	0.317	131,221,905.8
2001	0.690	0.627	0.649	0.582	0.348	0.345	100,426,397.4
2002	0.700	0.675	0.652	0.625	0.348	0.346	76,704,234.1
2003	0.716	0.699	0.661	0.643	0.316	0.314	58,869,859.8
2004	0.746	0.740	0.699	0.693	0.313	0.312	42,633,364.9
2005	0.766	0.762	0.723	0.718	0.319	0.318	27,431,776.5
2006	0.801	0.796	0.762	0.757	0.338	0.337	20,756,465.9
2007	0.817	0.817	0.780	0.780	0.328	0.327	15,589,665.8
2008	0.838	0.836	0.805	0.802	0.331	0.331	12,340,287.7
Average	0.724	0.698	0.680	0.653	0.331	0.329	76,413,156.7

Notes: DWL with no tax calculated is based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax is computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

average and marginal externalities (i.e., the naive and SBO taxes), respectively, of all vehicles in California in each year. DWL with the naive tax averages 72.4 percent of the DWL with no additional tax over the sample period and rises over time as the fleet becomes cleaner. The SBO gasoline tax is a little better, averaging 69.8 percent of DWL with no tax during our sample period.²⁶

Our simulation shows that the SBO gasoline tax is a remarkably poor instrument to address the local pollution externalities of personal automobiles. We next explore whether this result is likely to be specific to the context of local pollution externalities, or apply more broadly to other heterogeneous externalities. The purpose of this analysis is to explore the nature of the failure of the uniform gasoline taxes.

We first examine the role of between-county variation in local pollution externalities. The marginal damages from Muller and Mendelsohn (2009) are designed to vary at the county level, and within California, they vary substantially across counties, due to both baseline emissions levels and the extent to which population is exposed to harmful emissions. The third and fourth columns of Table 4 show the remaining DWL from a set of naive and SBO gasoline taxes computed at the county level. In essence, we eliminate the between-county variation in the local pollution externality by simulating separate taxes in each county equal to the average or marginal externality in each county and year. We find that the county-by-county variation in emissions and elasticities does not explain the failure of a single, uniform tax to remove a substantial amount of deadweight loss. The average ratio over our sample is 0.68 for the naive tax and 0.653 for the optimal uniform tax, just over 4 percentage points less than either the naive or SBO statewide tax.

²⁶ We can also calculate the ratio of remaining DWL to original DWL by calibrating equations (6) and (7) with the moments in our data. The average values in our sample for the lognormal shape parameters σ_E^2 and σ_B^2 are 1.47 and 1.51, respectively. The average value of ρ , the correlation coefficient for the logs of externality and inverse elasticity, is 0.28.

Another important feature of local pollution externalities that makes a uniform tax fall short is the variation across vintages. Older vehicles are on average much more polluting than newer ones, although as we showed in Section III, there is substantial variation even within vehicle type. To test the importance of between-vintage variation, the fifth and sixth columns of Table 4 show the proportion of remaining DWL after a set of naive or SBO gasoline taxes based on the local pollution externalities by vintage and year, eliminating between-vintage variation as a factor. Here we see a substantial improvement: 0.331 for the naive tax and 0.329 for an SBO gasoline tax.

The other major factor driving the failure of a uniform SBO tax to address local pollution externalities is the skewness in the externality distribution. The roughly 50 percent increase in the tax level from an SBO gasoline tax correctly abates more emissions from the dirtiest vehicles, but also overtaxes the cleanest vehicles by a larger amount. The welfare benefits of the SBO gasoline tax are around 10 percent higher than those from a naive tax, but still fall far short of the benefits from a true Pigouvian tax linked to actual vehicle emissions. The number of vehicles for which the uniform tax overshoots is remarkable. Specifically, the distribution of emissions is so strongly right skewed that the naive uniform tax overshoots for more than 72 percent of vehicle years and the SBO gasoline tax for even more.

The variance and skewness in the distribution of externality per gallon causes a uniform tax to be less efficient than might otherwise be expected. Figure 5 shows this clearly, plotting the kernel density of the externality per gallon in 1998 and 2008, with vertical lines indicating the naive tax and the optimal tax, respectively. The long right tail of the distribution requires that either tax greatly exceed the median externality. As striking as these results are, they are likely to be more skewed still in other states besides California. California has more stringent new-vehicle emissions standards and a more widespread emissions inspection program than most other US states and indeed more so than EU countries. These regulations will tend to reduce the number of vehicles in the right tail of emissions. For a more in-depth discussion of how our results would differ in other states, see in online Appendix F.

To test the importance of the shape of the externality distribution to the performance of an SBO tax, we next examine how the SBO gasoline tax would compare to the optimal Pigouvian emissions tax if the distribution became less skewed. To do this, we focus on the California personal vehicle fleet in 2008 and progressively remove more of the right tail of the externality distribution. These results are reported in Table 5. The first column reports the ratios of DWL with the SBO gasoline tax to DWL with no tax for 2008, with both tax and DWL recalculated after dropping vehicles whose externality is above the ninety-fifth, ninetieth, seventy-fifth, fiftieth, or twenty-fifth percentiles. The second column reports the SBO tax in dollars, while the last two columns give the standard deviation and skewness of the resulting externality distribution. With the top 5 or 10 percent of the externality distribution removed, the distribution remains highly right skewed, and the SBO tax still removes less than half the DWL of local pollution externalities. Cutting our sample at the seventy-fifth percentile makes the externality distribution much less skewed, and the remaining DWL after an SBO tax falls by more than half compared to our full sample. When vehicles above the twenty-fifth percentile are removed, the




FIGURE 5. DISTRIBUTION OF EXTERNALITY PER GALLON—VERTICAL LINES INDICATE NAIVE AND MARGINAL UNIFORM TAX

remaining externality distribution becomes roughly symmetrical with a skew statistic near zero, and the (now very small) SBO tax removes more than 84 percent of the DWL.

C. Treatment of Other Externalities

Our results in the previous subsection show that the heterogeneity in the local pollution externalities of personal automobiles make a uniform tax a remarkably poor tool to address those externalities. We now consider the implications of this result for other externalities, including other externalities from personal transportation. To the extent that a uniform tax can address other externalities, we are underestimating the effectiveness of a uniform tax. However, if the features of local pollution that lead an SBO tax to fail are present in other externalities, our results apply more broadly and point to the need for other policies to control externalities.

Indeed, many of the other externalities of personal transportation are similar to criteria pollution emissions in the sense that they also vary across vehicles. Congestion and accident externalities depend on when and where vehicles are driven. Accidents and infrastructure depreciation depend to some degree on vehicle weight.²⁷ We lack vehicle-specific measures of these other externalities to measure how imposing an SBO gasoline tax would address the DWL from these externalities. However, we can use our results of the previous subsection to assess qualitatively the effectiveness of a uniform gasoline tax. This is important in part because, to the extent that a uniform gasoline tax can address other externalities, it is possible that current policy already does so—the combined state and federal gasoline tax in California was \$0.47 per gallon during our sample period.

²⁷ For estimates on the degree of this heterogeneity, see Anderson and Auffhammer (2011) and Jacobsen (2013).

TABLE 5—RATIOS OF DWL WITH TAX TO DWL WITH NO TAX IN 2008,
ELIMINATING THE RIGHT TAIL OF THE EXTERNALITY DISTRIBUTION

Exclude above:	DWL ratio	SBO tax \$	Externality moments	
			SD	Skewness
Ninety-fifth percentile	0.745	0.12	0.144	3.171
Ninetieth percentile	0.628	0.07	0.066	2.973
Seventy-fifth percentile	0.373	0.04	0.020	1.128
Fiftieth percentile	0.249	0.02	0.008	0.520
Twenty-fifth percentile	0.157	0.01	0.004	0.158

Notes: DWL with no tax is calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax is computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

The harms of local pollution can vary by the geographic area where the pollutants are emitted. Externalities such as accidents and congestion are unlikely to have the same type of geographic heterogeneity—geographic variation for these externalities is more likely to be between particular roads and intersections rather than across counties. However, in the previous subsection, we showed that geographic variation is a minor factor in the failure of a uniform gasoline tax to address local pollutant externalities. A more important factor was variation across vintages of vehicles. We expect that a similar type of variation will exist for accident externalities, and to an extent with congestion externalities, through variation across hours of the day. More generally, variation across classes of individuals is a common feature in heterogeneous externalities. Perhaps the most important factor leading to the failure of the SBO gasoline tax for local pollution is the strong skew to the externality distribution. The harms of accident and congestion externalities are highly skewed as well—large trucks disproportionately contribute to the former, while peak period driving disproportionately contributes to the latter.

These factors indicate that the SBO gasoline tax that addresses accidents and congestion externalities as well as local pollution externalities will fail to remove a significant portion of the DWL. Therefore, the actual amount of DWL will be the sum of the DWL that we measure plus the DWL loss arising from the externalities that we cannot measure.²⁸

One externality that does not vary across vehicles is the social cost of CO₂ emissions due to their contribution to climate change. Because CO₂ emissions are, to a first-order approximation, directly proportional to gasoline consumption, in this case a per-gallon gasoline tax is the optimal policy instrument. The larger the climate change externality, the greater the *share* of DWL eliminated by the SBO gasoline tax will be. To get a sense of how climate change externalities affect our calculations, we repeat the analysis for a range of social costs of carbon (SCC).

²⁸ In fact, Parry and Small (2005) find that the contribution of these other externalities to the second-best optimal gas tax may, in fact, be larger than the contribution arising from local pollutants. This would suggest that the degree for which we understate the remaining deadweight loss might be large.




FIGURE 6. REMAINING DEADWEIGHT LOSS UNDER ALTERNATIVES GASOLINE-SPECIFIC EXTERNALITIES

We calculate the remaining DWL after imposing the SBO gasoline tax based on local pollution externalities, varying the SCC from 0 cents per gallon to \$1.00 per gallon.²⁹ While our discussion focuses on the externalities associated with CO₂, we stress that these calculations are relevant for *any* externalities for which a per-gallon tax is the first-best instrument. They also represent the lower bound on the remaining DWL when we consider any other externality for which a per-gallon tax is a second-best instrument.

Figure 6 summarizes the results across all years in our sample. The values for an extra per-gallon externality of zero roughly correspond to the ratios reported in Table 4.³⁰ Not until the extra per-gallon externality exceeds \$0.20 per gallon does a uniform gasoline tax eliminate the majority of DWL associated with both the criteria pollutants and a per-gallon externality. Even if the per-gallon externality is \$1.00, nearly 20 percent of combined DWL remains under both the naive and SBO gasoline taxes.

VI. Incidence of Gasoline and Pigouvian Taxes

Our results in Section V demonstrate that a uniform gasoline tax is an ineffective policy tool on efficiency grounds. In this section, we consider the implications of the SBO gasoline tax and the first-best optimal Pigouvian emissions tax for equity.

²⁹For comparison, Greenstone, Kopits, and Wolverton (2011) estimate the SCC for a variety of assumptions about the discount rate, relationship between emissions and temperatures, and models of economic activity. For each of their sets of assumptions, they compute the *global* SCC; focusing only on the US impacts would reduce the number considerably. For 2010, using a 3 percent discount rate, they find an average SCC of \$21.40 per ton of CO₂ or roughly 23.5 cents per gallon of gasoline, with a ninety-fifth percentile of \$64.90 (71 cents per gallon). These calculations assume that the life-cycle emissions of gasoline are 22 pounds per gallon. Using a 2.5 percent discount rate, the average SCC is \$35.10 (38.6 cents per gallon).

³⁰The figure plots the weighted averages across the years, while the last row in Table 4 is a simple average of the annual weighted averages, hence a slight difference.

Gasoline taxes are generally thought to be regressive, placing a greater burden on the poorest households. Fuel consumption is relatively constant across the income distribution, and thus the argument goes that a tax on gasoline expenditures will represent a larger fraction of the income of low-income households.³¹ However, it is possible that a uniform gasoline tax is less regressive than an emissions tax, particularly if poorer households tend to own dirtier vehicles. We begin by describing our methodology for assigning household income to the vehicles in our Smog Check data and then present our results on the regressivity of the SBO gasoline tax and the optimal Pigouvian emissions tax.

We note that there is an extensive literature analyzing the distributional impacts of gasoline taxes and carbon taxes more generally. Poterba (1991) argues that annual expenditure is a better proxy for permanent income than annual income and shows that by this metric a gasoline tax is less regressive than it might otherwise seem. However, Chernick and Reschovski (1997) use long-term income directly, and find gasoline taxes are still borne predominately by the poor. More recently, Teixidó and Verde (2017) find that gasoline taxes are regressive under analysis similar to Poterba (1991) if wealth is considered as well. Recent works focus on various means of “recycling” revenues from a gasoline tax to compensate for the burden of the tax, either by cutting more distortionary taxes or providing subsidies or lump-sum rebates to some households. Williams et al. (2015) argue that the uses of tax revenue may be more important than the tax itself to determine the regressivity of a carbon tax. While important, these issues are outside the scope of our analysis here, which focuses on the *relative* distributional effects of a gasoline tax and an emission tax.

A. Assigning Income to Vehicles

For the analyses in this section, we link the Smog Check data to DMV registration information. We geocode the addresses from the DMV data and match them to census block groups (CBGs) then link these data to CBG demographics from the 2000 Decennial Census. The DMV data allow us to match an address to each vehicle by calendar year for the period of 2000–2008. We then predict the annual average tax paid by the owners of each vehicle in the Smog Check data, using our estimates of optimal taxes and counterfactual VMT from Section V. Note that the because our predicted counterfactual VMT is vehicle specific, we implicitly allow for lower income households to be more responsive to gas prices to the extent that lower income households tend to own older and dirtier cars. However, our analysis only considers partial equilibrium effects, as we do not account for interactions with other markets, and in particular the use of the tax revenues. This should not materially affect the comparison between taxes, but our estimates of the absolute regressiveness of each tax may be biased.

³¹ In principle, a gasoline tax could be more progressive than this argument implies, if the poorest households do not own cars. For instance, del Granado, Coady, and Gillingham (2012) show that fuel subsidies in developing countries primarily accrue to higher income individuals. In practice, the 2009 NHTS indicates that over 95 percent of US households own at least one car; the fraction of California households with a car is slightly higher.

Panel A. SBO tax

Panel B. Pigouvian emissions tax




FIGURE 7. SIMULATED ANNUAL TAX PAID AS A PROPORTION OF CBG MEDIAN INCOME

All of the analyses in this section consider the average of predicted taxes over the whole period from 2000–2008. Although the levels of tax and externality are higher earlier in the period than later in the period, the distributional patterns are almost identical regardless of the time period used.

We note that annual income is not an ideal measure of well-being for the purpose of analyzing distributional impacts of a gasoline tax, as annual income may not reflect a household's long-term ability to bear a tax burden. Permanent or lifetime income is more appropriate, with annual consumption being the most commonly used proxy. We use annual income for practical reasons—we do not have consumption at the household or even the CBG level, and moreover the differences between annual expenditures and annual income will be muted in aggregate CBG data that we use.³²

Figure 7 shows a map of California CBGs, shaded to show the average predicted annual tax paid as a fraction of year 2000 CBG median income. Panel A shows the average tax burden of the SBO tax, while panel B shows the average tax burden of an optimal Pigouvian emissions tax. Two patterns are evident from the maps. First, there is substantial geographic dispersion across the state of California, with vehicle owners in the urban cores of Los Angeles and San Diego paying a much higher proportion of CBG median income in taxes. Second, although the geographic dispersion differs slightly between the SBO gasoline tax and the optimal emissions tax, the *levels* differ substantially—the annual burden of the optimal tax is much lower than that of the SBO tax. We discuss the intuition behind this surprising result below.

³² Poterba (1991) notes that the wedge between annual income and annual expenditure stems from a relatively small fraction of households whose expenditure ranking deviates substantially from their income ranking.

The map analysis in Figure 7 is instructive, but is not enough to show the extent to which the SBO gasoline tax or an optimal Pigouvian emissions tax would be progressive or regressive. To estimate the progressivity of these taxes, median CBG income is insufficient. Borenstein (2012) shows that CBG median income masks substantial within-CBG variation in household income, which causes it to do a poor job of capturing effects on the top and bottom of the income distribution. We implement Borenstein's suggested correction for CBG income, which uses the full distribution of household incomes in each CBG from the census data combined with a separate dataset that contains both annual VMT and household income. In brief, Borenstein's method requires the correlation between annual VMT and income within a CBG. In short, using the 2009 NHTS, we calculate the correlation between VMT and income. This allows us to assign vehicles in the Smog Check data to income brackets based on each vehicle's annual VMT and the proportion of households in each income bracket within the CBG the vehicle is registered in. For more details, see online Appendix H and Borenstein (2012).

Using the Borenstein (2012) correction, we assign vehicles to 1 of 10 income brackets, which aggregate the 16 income categories contained in the census data into groups roughly approximating deciles of the California household income distribution.³³ For purposes of calculating the tax burden and progressivity, we use the midpoint of each income bracket. Because our data are at the vehicle level, not the household level, we account for multiple vehicle households by dividing the estimated household income by the average number of vehicles per household for that income bracket, taken from the 2009 NHTS.

B. Regressivity Results

Figure 8 plots for each decile of household income the average tax burden as a percentage of estimated income for the naive uniform tax, the SBO tax, and the optimal Pigouvian emissions tax.³⁴ The figure also plots the average of pretax externality in dollars per year using the right axis. Aside from the tenth, highest, income decile, which has a much higher average annual VMT than the ninth decile, the average pretax externality is declining with income. In other words, we find that poorer households have dirtier cars, and pollute more in total even though their annual VMT is lower than richer households. As such, we expect an emissions tax to be regressive to some extent. Indeed, all three taxes are regressive, with the lowest income brackets predicted to pay the highest percentage of household income toward the tax. However, the curve is most steeply sloped for the SBO gasoline tax, indicating that it is the most regressive of the three. The optimal emissions tax imposes a lower average tax burden than the SBO gasoline tax in every income decile.

On its face, it seems surprising that an emissions tax would result in a lower tax burden for all parts of the income distribution. The explanation for this result

³³ Specifically, the break points for the groupings are at the 10.41, 19.87, 29.02, 41.68, 49.31, 59.30, 72.10, 81.29, and 93.5 percentiles of all households in California.

³⁴ Note that this is more precisely the average burden by decile for car owners. However, given that more than 95 percent of California households own cars, including carless households would not significantly change our results.

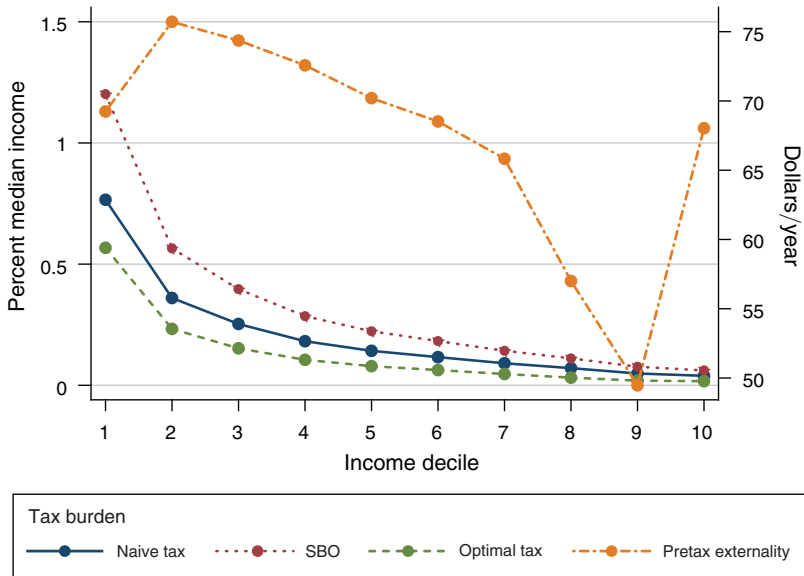


FIGURE 8. ANNUAL TAX BURDEN AS A PERCENT OF INCOME AND ANNUAL PRETAX EXTERNALITY BY HOUSEHOLD INCOME DECILES

is that under an emissions tax, vehicles with the highest per-gallon tax rate have the fewest posttax VMT, and vice versa. Even if all vehicles were equally price responsive, an emissions tax would raise less revenue on average than a uniform tax simply because the highest polluting vehicles pay the highest tax rate and thus reduce VMT and gasoline consumption the most. Of course, the core result of this paper is that vehicles are *not* equally responsive to gasoline prices, with the dirtiest vehicles having the greatest VMT elasticity. This further reduces the average burden of an emissions tax, while the variation in elasticities pushes the level and average burden of the SBO gasoline tax higher. Moreover, in practice, older and dirtier vehicles have lower VMT pretax. Thus, the SBO gasoline tax would only impose a lower tax burden on an income bracket if high-polluting, low VMT vehicles were concentrated in one decile, which is not the case. Indeed, the SBO tax is higher than optimal on a per-mile basis for more than 80 percent of vehicles. Although households in the lower income brackets are more likely to have higher polluting vehicles, more than 75 percent of vehicles in every income group have emissions below the marginal externality that determines the SBO gasoline tax. As a result, switching from the SBO gasoline tax to an emissions tax lowers the tax burden for the vast majority of vehicles in every income bracket.

It is also important to note that while Figure 8 illustrates that gasoline taxes are regressive, the figure hides a tremendous amount of variation within income deciles. We find that the variance falls as income rises. Figure 9 shows this clearly, plotting several percentiles of SBO gasoline tax expenditures as a share of income within each income bracket. For instance, the interquartile range of the share of SBO gasoline tax as a fraction of income for households in the lowest income decile is




FIGURE 9. DISTRIBUTION OF TAX BURDEN BY HOUSEHOLD INCOME DECILES

between 0.5 percent of household income to over 1.5 percent of household income. In contrast, the interquartile range for the higher deciles is extremely small.

VII. Conclusions

In this paper, we present three empirical results, all stemming from the stylized fact that vehicle emissions are heterogeneous and highly right skewed. First, the sensitivity of a given vehicle’s miles traveled to gasoline prices is correlated with the vehicle’s emissions. Dirtier vehicles are more price responsive. This increases the size of the second-best optimal uniform gasoline tax by as much as 50 percent.

Second, uniform indirect taxes are an inefficient policy tool to reduce vehicle emissions. In this paper, we demonstrate this deficiency through our empirical example of motor vehicle emissions of criteria pollutants. The optimal policy would

differentially tax vehicles based on their emissions, not on consumption of gasoline. While gasoline consumption and emissions are positively correlated, we show that gasoline taxes are a poor substitute for a true Pigouvian emissions tax. The remaining DWL under the second-best optimal gasoline tax exceeds 75 percent in the second half of our sample. Moreover, the reasons for this shortfall, primarily skewness of the externality distribution and marked differences by vehicle vintages, likely apply to other externalities. Although it comes as no surprise that an indirect tax fails to achieve the optimal result, the magnitude of that failure is striking.

Finally, we find that gasoline taxes are not only regressive, but are more regressive than a Pigouvian tax on emissions. Because the distribution of emissions is so strongly right skewed, with a small number of very high polluting vehicles contributing the bulk of total emissions, a uniform gasoline tax will tend to overtax relative to the social optimum, leaving the vast majority of vehicle owners paying more, and with the poorest households paying substantially more as a fraction of their income.

REFERENCES

- Anderson, Michael, and Maximilian Auffhammer.** 2011. "Pounds that Kill: The External Costs of Vehicle Weight." National Bureau of Economic Research (NBER) Working Paper 17170.
- Archsmith, James, Kenneth Gillingham, Christopher R. Knittel, and David S. Rapson.** 2017. "Attribute Substitution in Household Vehicle Portfolios." National Bureau of Economic Research (NBER) Working Paper 23856.
- Borenstein, Severin.** 2012. "The Redistributive Impact of Nonlinear Electricity Pricing." *American Economic Journal: Economic Policy* 4 (3): 56–90.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer.** 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1): 220–56.
- Chernick, Howard, and Andrew Reschovski.** 1997. "Who Pays the Gasoline Tax?" *Nation Tax Journal* 50 (2): 233–59.
- Davis, Lucas W., and Matthew E. Kahn.** 2010. "International Trade in Used Vehicles: The Environmental Consequences of NAFTA." *American Economic Journal: Economic Policy* 2 (4): 58–82.
- del Granado, Francisco Javier Arze, David Coady, and Robert Gillingham.** 2012. "The Unequal Benefits of Fuel Subsidies: A Review of Evidence for Developing Countries." *World Development* 40 (11): 2234–48.
- Diamond, Peter A.** 1973. "Consumption Externalities and Imperfect Corrective Pricing." *Bell Journal of Economics and Management Science* 4 (2): 526–38.
- Fullerton, Don, and Sarah E. West.** 2010. "Tax and Subsidy Combinations for the Control of Car Pollution." *B.E. Journal of Economic Analysis and Policy* 10 (1): Article 3.
- Gillingham, Kenneth.** 2010. "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock in California." http://www.iaee.org/en/students/best_papers/Gillingham.pdf.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton.** 2011. "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation." National Bureau of Economic Research (NBER) Working Paper 16913.
- Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling.** 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29 (1): 113–34.
- Jacobsen, Mark R.** 2013. "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior." *American Economic Journal: Applied Economics* 5 (3): 1–26.
- Jacobsen, Mark R., and Arthur A. van Benthem.** 2013. "Vehicle Scrapage and Gasoline Policy." National Bureau of Economic Research (NBER) Working Paper 19055.
- Kahn, Matthew E.** 1996. "New Evidence on Trends in Vehicle Emissions." *RAND Journal of Economics* 27 (1): 183–96.
- Knittel, Christopher R., and Ryan Sandler.** 2013. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." National Bureau of Economic Research (NBER) Working Paper 18849.

- Knittel, Christopher R., and Ryan Sandler.** 2018. "The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation: Dataset." *American Economic Journal: Economic Policy*. <https://doi.org/10.1257/pol.20160508>.
- Li, Shanjun, Christopher Timmins, and Roger H. von Haefen.** 2009. "How Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy* 1 (2): 113–37.
- Matthews, H. Scott, and Lester B. Lave.** 2000. "Applications of Environmental Valuation for Determining Externality Costs." *Environmental Science and Technology* 34 (8): 1390–95.
- Morrow, Silvia, and Kathy Runkle.** 2005. April 2004. *Evaluation of the California Enhanced Vehicle Inspection and Maintenance (Smog Check) Program. Final Report (Revised)*. California Air Resources Board. Sacramento, September.
- Muller, Nicholas Z., and Robert Mendelsohn.** 2009. "Efficient Pollution Regulation: Getting the Prices Right." *American Economic Review* 99 (5): 1714–39.
- Parry, Ian W. H., and Kenneth A. Small.** 2005. "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review* 95 (4): 1276–89.
- Poterba, James M.** 1991. "Is the Gasoline Tax Regressive?" *Tax Policy and the Economy* 5: 145–64.
- Teixidó, Jordi J., and Stefano F. Verde.** 2017. "Is the Gasoline Tax Regressive in the Twenty-First Century? Taking Wealth into Account." *Ecological Economics* 138: 109–25.
- Westberg, Karl, Norman Cohen, and K. W. Wilson.** 1971. "Carbon Monoxide: Its Role in Photochemical Smog Formation." *Science* 171 (3975): 1013–15.
- Williams, Robertson C. III, Hal Gordon, Dallas Burtraw, Jared C. Carbone, and Richard D. Morgenstern.** 2015. "The Initial Incidence of a Carbon Tax Across Income Groups." *National Tax Journal* 68 (1): 195–214.

This article has been cited by:

1. Nicholas Ryan, Anant Sudarshan. 2022. Rationing the Commons. *Journal of Political Economy* **130**:1, 210-257. [[Crossref](#)]
2. Javier D. Donna. 2021. Measuring long-run gasoline price elasticities in urban travel demand. *The RAND Journal of Economics* **52**:4, 945-994. [[Crossref](#)]
3. Cody Nehiba. 2021. Correcting heterogeneous externalities: Evidence from local fuel taxes. *Journal of the Association of Environmental and Resource Economists* . [[Crossref](#)]
4. Cody Nehiba. 2021. The Time-of-Day Travel Demand Elasticity Paradox. *SSRN Electronic Journal* **7** . [[Crossref](#)]
5. Bruno De Borger, Stef Proost. Pricing Principles in the Transport Sector 95-101. [[Crossref](#)]
6. Felix Creutzig, Aneeqe Javaid, Nicolas Koch, Brigitte Knopf, Giulio Mattioli, Ottmar Edenhofer. 2020. Adjust urban and rural road pricing for fair mobility. *Nature Climate Change* **10**:7, 591-594. [[Crossref](#)]
7. Christian Haas, Karol Kempa. 2020. Low-Carbon Investment and Credit Rationing. *SSRN Electronic Journal* . [[Crossref](#)]
8. Raoul Minetti, Tao Peng, Tao Jiang. 2019. Keeping up with the Zhangs and house price dynamics in China. *Journal of Economic Dynamics and Control* **109**, 103778. [[Crossref](#)]
9. Morteza Taiebat, Samuel Stolper, Ming Xu. 2019. Forecasting the Impact of Connected and Automated Vehicles on Energy Use: A Microeconomic Study of Induced Travel and Energy Rebound. *Applied Energy* **247**, 297-308. [[Crossref](#)]
10. Fredrik Carlsson, Christina Annette Gravert, Verena Kurz, Olof Johansson-Stenman. 2019. Nudging as an Environmental Policy Instrument. *SSRN Electronic Journal* **95** . [[Crossref](#)]
11. Javier D. Donna. 2018. Measuring Long-Run Price Elasticities in Urban Travel Demand. *SSRN Electronic Journal* . [[Crossref](#)]