A Tale of Two Tails: Commuting and the Fuel Price Response in Driving*

Kenneth Gillingham, *Yale University, NBER, and CESIfo*

Anders Munk-Nielsen, *University of Copenhagen*

September 28, 2018

Abstract

Pricing greenhouse gases is widely understood as the most efficient approach for mitigating climate change, yet distributional effects hamper political acceptance. These distributional effects are especially important in transport, the fastest growing sector for greenhouse gas emissions. Using rich data covering the entire population of vehicles and households in Denmark, this study uncovers an important feature of driving demand: two groups of much more responsive households in the lower and upper *tails* of the work distance distribution. We further estimate the causal effect of public transport–a critical determinant of the upper tail–and show how public transport access can both reconcile differences in fuel price elasticities between the United States and Europe, and considerably influence the distributional effects of fuel pricing.

Keywords: distributional effects, transportation, commuting, urban form, environmental taxes.

JEL classification codes: R4, R2, Q4, H2, L9

*Gillingham: Yale University, NBER, and CESIfo, 195 Prospect Street, New Haven, CT 06033 USA, kenneth.gillingham@yale.edu; Munk-Nielsen: University of Copenhagen, Department of Economics, Øster Farimagsgade 5, 26.3.01, DK-1353 Copenhagen K, Denmark, anders.munk-nielsen@econ.ku.dk.

The authors wish to thank Bertel Schjerning, John Rust, Matthew Kotchen, Ismir Mulalic, Fedor Iskhakov, and Søren Leth-Petersen, as well as numerous seminar participants, for helpful comments and feedback. The authors would also like to acknowledge funding from the IRUC research project, financed by the Danish Council for Strategic Research.

“Relocation is part of the solution to a serious problem. That growth and jobs are too unevenly distributed between country and city. It creates a risk of Denmark being ripped apart.” - Lars Løkke Rasmussen, Prime Minister of Denmark, 19 October 2016 (on the government’s plan to relocate public sector jobs)

1 Introduction

Greenhouse gas emissions from transportation amount to a quarter or more of all anthropogenic emissions in many countries around the world, and this share is rapidly growing. From 1990 to 2015, the share of greenhouse gas emissions from transport has increased from 15 percent to 26 percent in Europe and from 23 percent to 27 percent in the United States (EEA, 2017; EPA, 2017). With continued declines in the cost of renewables and continued decarbonization of the electricity sector, this trend is likely to continue and perhaps even accelerate. Thus, policymakers worldwide have become increasingly attentive to the transport sector. Economists are uniformly in support of pricing greenhouse gases as the first-best approach to mitigate greenhouse gas emissions. Yet, political acceptance of pricing policies is often hampered by concerns about the distributional effects of such policies, and this is especially true for pricing transportation fuel consumption (Borenstein, 2017). A common concern is that pricing policies will disproportionately affect a subset of households, such as less-wealthy households outside of urban areas.

This study uses millions of vehicle-level odometer readings matched to individual-level demographic information from the Danish registers to ask several questions. How do households across the population change their driving in response to fuel price changes? How is this response influenced by access to public transport? And what does the heterogeneity in response mean for the distributional effects of fuel price changes? We uncover a new finding to the literature: two groups of households who are much more responsive to changing fuel prices than most of the population. These households are in the *tails* of the work distance distribution; one group has the longest commutes and the other commutes very little. Our mean medium-run (one-year) elasticity estimate of -0.30 is considerably influenced by these two groups of tail households, each of which are much more elastic. These findings can be rationalized by switching costs incurred when substituting from driving to other modes of transport, such as public transport, and we estimate the causal effect of public transport on driving responsiveness. Danes have almost universal access to public transport and we posit that our results hold in similar settings around the world.

Our results have direct policy implications. By uncovering the economic mechanisms at work in the driving elasticity, we show that the two groups of tail households face substantially reduced impacts from fuel price increases. This is especially important because the roughly 15% of the population in the upper tail–households with the longest commutes–face a heavy burden from fuel price increases and thus may be more opposed to pricing greenhouse gases. This is a politically salient point in Denmark as the Danish government pays particularly close attention to rural voters who tend to have longer commutes. As was noted by Prime Minister Rasmussen, the Danish government is even going as far as moving 3,900 governmental jobs to outside the capital region at a cost of over 400 million DKK ($61 million at the July 9, 2017 exchange rate).[1] These issues are by no means isolated to Denmark, with the impact on groups of voters outside of cities playing a key role in policy debates about pricing greenhouse gases throughout the developed world.

Indeed, our findings also have implications for other countries where access to public transport is less universal. Our results suggest that access to public transport is a prerequisite for the existence of the upper tail. Without adequate access to public transport, households with long commutes are less able to substitute away from driving when fuel prices rise. Intuitively, previous work from countries with more limited access to public transport, such as the United States, show no evidence of the upper tail of responsiveness from households with longer commutes. Such previous work does however routinely find that households in cities (who would be expected to have shorter commutes) tend to be more responsive (e.g., Kayser, 2000; Gillingham, 2013, 2014; Gillingham, Jenn, and Azevedo, 2015), consistent with our lower tail, which we show is also primarily driven by households in the cities in Denmark. By identifying the upper tail in the work distance distribution in Denmark, we show that there is a group of more-responsive households in Europe that is not likely to exist in countries like the United States where public transport provision is lower. This directly impacts the effectiveness of fuel pricing in reducing driving and emissions and may crucially affect the political acceptability of pricing greenhouse gases.

This research contributes to several strands of literature. First, there is a growing literature on the distributional effects of policies to reduce greenhouse gas emissions from the transportation sector. Economists have worked on this issue for decades, primarily focusing on the vertical distributional effects (i.e., distributional effects over income) of gasoline taxes (Poterba, 1989, 1991). More recently, Borenstein and Davis (2016) estimate the distributional effects of U.S. Clean Energy Tax Credits (including subsidies for hybrid and electric vehicles) and find them to be quite regressive. Levinson (2016), Jacobsen (2013), and Davis and Knittel (2016) compare the vertical distributional impacts of fuel economy standards to gasoline taxes in the United States, generally finding that fuel economy standards are more

[1]See the Danish factsheet: `https://www.regeringen.dk/aktuelle-dagsordener/udflytning-af-statslige-arbejdspladser/se-status-paa-udflytning/`.

regressive than gasoline taxes.[2] There is much less work on geographic distributional consequences, despite the importance for policy. Bento et al. (2009) use survey data to estimate the efficiency and distributional consequences of gasoline taxes across states in the United States and show that households in more rural states face a much higher burden. Our study is the first to identify the two tails of more responsive drivers, link them to geographic location, shed light on the mechanisms creating them, and show how they influence the short-run distributional effects of policies affecting fuel prices.

This study also uses rich data to provide a new point estimate for the fuel price elasticity of driving, which is a dominant component in the modeling of gasoline or diesel demand. Understanding the responsiveness to fuel prices is a first-order question in economics. Not only is it valuable for anticipating responses to future swings in oil prices, it is also useful for measuring the macroeconomic effects of oil price fluctuations (e.g., Edelstein and Kilian, 2009) and providing insight into the role of speculators during oil price shocks (Hamilton, 2009). Not surprisingly, there is a vast literature aiming to estimate the price elasticity of gasoline demand (e.g., for some recent studies see Coglianese et al., 2016; Davis and Kilian, 2011; Hughes, Knittel, and Sperling, 2008; Li, Linn, and Muehlegger, 2014; Hymel and Small, 2015; Small and van Dender, 2007; Levin, Lewis, and Wolak, 2017). Most of these studies use aggregate data at the regional or national level.[3] More recently, several studies have estimated the elasticity of vehicle-miles-traveled with respect to the price of gasoline using disaggregated micro-level data, either from surveys or inspection odometer reading data (e.g., Linn, 2016; Bento et al., 2009; Gillingham, 2013, 2014; Munk-Nielsen, 2015; De Borger, Mulalic, and Rouwendal, 2016b; Knittel and Sandler, 2018). In a notable contrast, estimates for drivers in the United States tend to be in the range of -0.05 to -0.30, while similar benchmark estimates for European drivers tend to show a much more elastic response. For example, Frondel and Vance (2013) estimate a medium-run driving elasticity with respect to the gasoline price of -0.45 in Germany.[4] Similarly, in contemporaneous work, De Borger, Mulalic, and Rouwendal (2016a) focus on a subsample of two-vehicle households in Denmark and find the medium-run fuel price elasticity of driving to range between -0.32 and -0.45. We demonstrate how the removal of the upper tail–or the removal of adequate access to public transport–can reconcile these differing estimates between the United States and Europe.

[2]There is a substantial literature on the distributional effects of gasoline taxes and carbon taxes, including Hausman and Newey (1995), West (2004), West and Williams (2004), Bento et al. (2009), Sterner (2012), Williams et al. (2015).

[3]Review articles cover dozens of studies going back decades, most using aggregate data. For example, see Dahl and Sterner (1991), Espey (1998), Graham and Glaister (2004), and Brons et al. (2008).

[4]-0.45 is the fixed effects estimate, which we believe is better identified than other estimates in the paper, which are closer to -0.6.

Our results also contribute to a third vein of literature on the complex relationships between household location, public transport availability, gasoline prices, and consumer decisions about how much to drive. Since at least McFadden (1974), it has long been recognized that access to public transport is an important mediator of travel choices, with clear environmental implications (e.g., Glaeser and Kahn, 2010). Our work contributes to this literature by using an instrumental variables strategy to identify the causal effect of public transport on the responsiveness of driving to fuel price changes. Further, there is growing evidence that urban form and the spatial structure of commuting demand can affect travel mode choices (Bento et al., 2005; Grazi, van den Bergh, and van Ommeren, 2008; Brownstone and Golob, 2010). Our paper follows the literature estimating short-run and medium-run driving elasticities by holding the household location fixed, but shows how long-run equilibrium commute distances influence the short-run response to fuel prices.[5] These results are important for land use and transport policy, as they clarify the implications of policies that encourage shorter commute times (e.g., some smart growth policies) or greater access to public transport.

The remainder of this paper is organized as follows. The next section describes the rich Danish register data and provides descriptive evidence on the primary features of the data relevant to estimating the driving responsiveness. Section 3 describes our empirical strategy, while section 4 presents the results and a set of robustness checks. Section 5 discusses implications for policy, including an illustrative distributional effects analysis. Section 6 concludes.

2 Data

2.1 Data Sources

This study is based on rich data from the Danish registers on the population of both households and vehicles in Denmark from 1998 to 2011. There are three main sources. The first is the vehicle license plate register, which contains the vehicle identification number, gross vehicle weight rating (i.e., maximum operating weight including passengers and cargo), fuel type, date of registration, owner identification number, and whether the vehicle type is a personal car or a van.[6]

The second data source is the vehicle inspection database. Since July 1, 1998, all vehicles

[5] We also run a series of robustness checks where we exploit data on household residence and work locations and changes in these over time, finding that both our mean elasticity and tail results are quite robust.

[6] Company cars are not in our database and are not linked to a person but rather to the firm. However, individuals with access to a company car must pay a tax for this, and we observe that 3.7% of our households have at least one member paying this tax.

in Denmark are required to undertake a mandatory safety inspection at periodic intervals after the first registration of the vehicle. In Denmark, the first inspection is roughly four years out, and subsequent inspections are every other year.[7] Only a small number of used vehicles are imported into Denmark, in part because they pay a large vehicle registration fee and value-added tax that are assessed based on similar new vehicle prices. The fee and tax schedule are based on the value of the vehicle for all vehicles new to Denmark.[8] The inspection database contains odometer readings, which can be used to determine the kilometers driven between two inspections.

The third primary data source is the household register, which contains detailed demographic data at the calendar year-level. These data include the number of members of the household, ages and sex of these members, municipality of the household, income of the household members (including transfers), and a measure of work distance used to calculate the tax deduction for work travel.[9] On Danish tax forms, households must report their home address, workplace address, and number of days that work travel occurred (regardless of mode of transport). Multiplying the door-to-door work distance by the number of days yields a measure of work distance that accounts for how much total commuting is actually occurring. We then divide by 225 to give a measure of the effective average work distance in a workday. This measure is particularly useful because a household that lives very far away from work but rarely ever physically commutes (e.g., teleworks most of the time) would have a low (average) work distance by this measure. The tax authorities take this measure seriously, with random audits of employers, and checks on the home and work addresses, so it is considered a highly reliable measure. One downside of this measure is that the individual is only eligible for a deduction if the distance is greater than 12 kilometers but there is no minimum requirement on the number of days that travel to work occurred. The work distance measure will therefore be equal to zero if the individual lives closer than 12 kilometers from the work place or if the individual does not travel to their workplace. If an individual lives further than 12 kilometers and only occasionally commutes, it is possible for the measure to be between 0 and 12.

For 2000 to 2008, we also have data on the actual work distance for 79.6% of the households

[7] This is a very similar schedule to inspections in states in the United States, such as California. Details about the driving period lengths are in Appendix A.1.2. We control for the length of the period and perform robustness checks subsampling on the period length.

[8] After 2007, the vehicle registration fee assessed at the time of the transaction is also adjusted based on the fuel economy of the vehicle.

[9] For couples, there is a separate work distance variable for the male and female of the couple. In these cases, we use the maximum work distance of the two members but have also performed robustness checks to confirm that this does not appreciably change the primary results.

measured using a door-to-door shortest-path algorithm and provided by Statistics Denmark.[10] This second measure provides a useful check on our first measure, especially for households with commutes under 12 kilometers (See Appendix A.3.3). However, the actual work distance measure is a flawed measure for understanding the commute length of long work distance households, for some of these households may work from home or in nearby locations, and thus rarely travel the full commute. Fortunately, our robustness checks suggest similar results regardless of which measure we use.

In addition to the register data, we also bring in daily price data for 95 octane gasoline and diesel fuel from the Danish Oil Industry Association.[11] Similarly, we also bring in daily West Texas Intermediate crude oil price data for a robustness check.[12] Finally, we use data from Journey Planner on all bus and train stops in Denmark in 2013.[13] This is a single cross-section, but there were no substantial Denmark-wide changes in public transport over our time period; there were only minor extensions of certain lines and other tweaks to the system, as will be discussed further below.

We also have access to some additional car characteristics, including fuel economy in kilometers/liter and the manufacturer suggested retail price (MSRP). These data are from the Danish Automobile Dealer Association (DAF). However, these variables are not available for car vintages older than 1997. Finally, we bring in historical municipality-level population data from 1916 from Statistics Denmark.

2.2 Development of the Final Dataset

We combine the data from the various sources to create a final dataset where the unit of observation is a vehicle driving period between two inspections. So if a driver has a first inspection of her vehicle on June 1, 2004 and the next inspection on June 6, 2006, the driving period will be the 735 days between these two tests. We use the difference in odometer readings between these two inspections to calculate the total kilometers driven and the kilometers driven per day over the driving period. Similarly, we calculate the average gasoline, diesel, and oil price over the same driving period. If a car changes owners during a driving period, we include an observation for both households that have contributed to the driving and a variable for the fraction of the driving period the car is held by each owner.

To match our calendar year demographic data with driving periods, we construct a weighted average of the values of the demographic variables over the years covered by the

[10] Statistics Denmark has access to the actual addresses of individuals. This information, however, is anonymized in our dataset so we cannot perform any operations based on GIS information.

[11] See `www.eof.dk`, Accessed June 17, 2015.

[12] See `www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=RWTC&f=D`, Accessed June 15, 2015.

[13] See `www.journeyplanner.dk`, Accessed April 19, 2013.

driving period. For example, if a driving period covers half of 2001, all of 2002, and half of 2003, the values of the demographic variables would be given a weight of 0.25 for 2001, 0.5 for 2002, and 0.25 for 2003. The density of public transport stops is added to the data at the municipality level.

The final dataset after cleaning consists of 5,855,446 driving period observations covering nearly all driving periods by Danish drivers over the period from 1998 to 2011. Table 1 presents summary statistics for the final dataset. For our estimations, we demean all of the variables in order to facilitate interactions. Appendix A provides further details on the data sources and cleaning process.

2.3 Descriptive Evidence

There is considerable variation in fuel prices in Denmark over the period 1998 to 2011. Figure 1 shows average gasoline prices over time in our dataset. The x-axis denotes the time of the inspection at the beginning of the driving period. Figure 1 also plots the average daily vehicle-kilometers-traveled (VKT) over the driving period, illustrating a negative relationship between fuel prices and driving.[14]

The rich Danish register data allow us to explore the relationship between fuel prices and driving in greater detail. Figure 2 divides the sample into ten groups based on the percentiles of driving in each year. The vehicles in each group may change over time, as we recalculate the percentiles in each year. The figure illustrates that for most groups there appears to be very little change in driving over time, even as fuel prices change significantly. However, the 1 percent of drivers who drive the most show a noticeable decrease in VKT during driving periods that begin between 2003 and 2005, just as gasoline prices are rising. Even though this figure is based on VKT rather than work distance, we view this as initial suggestive evidence of the existence of an upper tail of more responsive drivers.

One would expect a high correlation between driving and work distance. Figure 3 uses a binned scatterplot to show the nonparametric relationship between per-vehicle driving and work distance. This plot shows that drivers exhibit remarkable heterogeneity over work distance. Longer commutes translate into more driving. There is a point mass at zero, which accounts for households with a work distance measure less than 12 kilometers. After this point mass, driving is a monotonically increasing function of work distance. The general concavity suggests that there is likely to be more non-commute driving for households with shorter work distances since households with much longer commutes drive only slightly more than households with much shorter commutes.

[14]See Appendix A for a figure showing the distribution of VKT.

We can further explore our data by examining the characteristics of households stratified by their work distance measure. Table 2 compares households with zero work distance to households with a work distance over 30 kilometers.[15] The drivers in the upper tail of work distance on average have a higher income, have more vehicles, have larger families, and drive more. They also tend to drive younger cars and are more likely to be driving a diesel car. These mean statistics mask heterogeneity in the upper tail: there are some very wealthy households in the upper tail that skew several of the averages upwards, while most in the upper tail are poorer rural households. It is also relevant that those in the upper tail also have only slightly less access to public transport than the rest of the population.

To visualize where the high-work distance households live, Panel (a) of Figure 4 shows a map of Denmark where each municipality is shaded according to the average work distance of the households living in that municipality (darker means shorter work distance). The figure shows that, conditional on owning a car, the high-work distance households tend to be in rural areas or on the outskirts of the major urban areas, while households that commute the least are in the urban areas. The regions of high-work distance municipalities also tend to be the municipalities with most driving (see Appendix).

An important way for a driver who commutes further to reduce driving is by switching to public transport. Panel (b) of Figure 4 illustrates the prevalence of public transport access throughout Denmark by showing each train or bus stop as a dot. There are bus or train stops nearly everywhere in Denmark. Moreover, there is on-call public transport available in rural municipalities where the stops are sparser ("telebusser"). This pervasiveness of public transport–which contrasts with other countries such as the United States–makes switching behavior possible for those with long commutes.

3 Empirical Approach

3.1 Empirical Specification

A first goal of this paper is to investigate the fuel price elasticity and explore how this elasticity varies with work distance. We follow a vast literature on estimating fuel price elasticities in using a linear log-log specification for driving and the fuel price. This specification not only provides for a ready interpretation of the coefficient of interest, but we find that it also fits the data well. A less common, but also intuitive, specification would use the price per kilometer of driving (i.e., the fuel price divided by the fuel economy) rather than the fuel price. We opt

[15] We choose 30 kilometers because it captures a small percentage–about 15%–of the population and because our results will show that this threshold is roughly where the upper tail elasticity becomes greater. We could choose another cut-off without changing the qualitative insights here.

not to use this specification because we do not observe fuel economy for a sizable portion of the sample.[16]

Consider the demand for driving for vehicle i in household h during a driving period t, which may cover several years y. Recall that a driving period is simply the period in between two odometer readings. We model the demand for driving as follows:

$$\log \mathrm{VKT}_{iht} = \gamma \log p_{ft} + \mathbf{x}_{iht}\boldsymbol{\beta} + \phi(f,t) + \mu_h + \varepsilon_{iht}. \tag{1}$$

VKT_{iht} is the average daily driving in kilometers, $\mathbf{x}_{iht}$ denotes a vector of controls, and p_{ft} is the average daily fuel price over the driving period t for the fuel type $f \in \{gasoline, diesel\}$ of vehicle i. The coefficient γ is our primary coefficient of interest–the fuel price elasticity. The controls in $\mathbf{x}_{iht}$ include variables for work distance, age of members of the household, gross income of members of the household, whether the household lives within one of the five major urban areas of Denmark, number of children, whether the vehicle is a company car, whether the household has at least one self-employed individual, and the density of bus or train stops in the municipality. The vector $\mathbf{x}_{iht}$ also includes variables for whether and by how much the driving period overlaps with other driving periods by the same household.[17]

We use $\phi(f,t)$ to denote fuel type-specific time controls. We consider several different types of time controls with increasing flexibility and our results are robust to which type we use. In the simplest case, we use a linear time trend in the mid-date of the driving period. A more flexible control uses year dummies for the mid-date of the driving period to allow a non-linear trend in average driving. We can then include a second set of year dummies interacted with a dummy for diesel fuel to allow the effect to be fuel-type specific. Next, we include year dummies equal to one if the driving period overlaps with the calendar year for even a single day. However, if a driving period only has a single day in a year, it would naturally be affected less by whatever unobserved time-factors affecting driving in that year. Thus, we examine a specification that includes year variables that are equal to the fraction of the driving period that takes place in the given year. These fuel-type specific year controls

[16]We also include household-vehicle fixed effects in a robustness check, and these fixed effects remove any vehicle-specific time-invariant factors, such as fuel economy. Our resulting elasticity is nearly identical. Further, we examine specifications where all covariates are included in logs as well as specifications where all covariates enter in levels, and again find similar results.

[17]Recall that if the car changes owner mid-way through the driving period, the driving period is included as an observation by both households and we add a control for the percent of the driving period each household owns the car. We also add controls for ownership of other vehicles that do not admit driving observations such as motorcycles, mopeds, trailers, etc.

are what we use in our preferred specification.[18]

In addition to the year controls, we also include controls for the length of the driving period as well as a dummy for whether the driving period is the first observation for the car (i.e., it is a new car), the car age, and the fraction of the driving period that it was owned by the current household. The controls for the length of the driving period may be particularly helpful for directly addressing any mean reversion due to shorter driving periods having greater variance in VKT. Finally, μ_h are household fixed effects.[19]

3.2 Identification

Of primary interest in this paper is discerning whether there is important heterogeneity in the relationship between driving and fuel price that influences the short-run distributional impacts of price changes. Our gasoline and diesel fuel price variables are time series variables, as there is negligible cross-sectional variation in fuel prices across Denmark. The primary source of the time series variation in these refined fuel price variables is variation in oil prices, as oil is the feedstock for gasoline and diesel production. Any remaining variation in the refined fuel prices may be due to Denmark-specific shocks to refining or fuel demand. The oil price is determined on the global market and Denmark is a small market, so it reasonably follows that Denmark-specific shocks are not likely to affect the global oil price. However, localized shocks may influence the non-oil price-related variation in the refined fuel prices. In addition, there may be correlated demand shocks across countries. For example, a common demand shock in Northern Europe due to a macroeconomic shock would be represented in the refined fuel price time-series variation.

These localized shocks and correlated demand shocks are likely to be a small part of the fuel price variation. Nevertheless, we consider each carefully. We address common regional demand shocks that may influence both driving and oil prices with our flexible time controls, and we perform a series of robustness checks with different time controls. We address the possibility of endogeneity due to localized shocks by performing a robustness check in which

[18] To illustrate these, if a driving period starts in the middle of 2001 and ends in the middle of 2003, then the controls will be 0.25 for 2001 and 2003 and 0.5 for 2002 and zero for all other years. These variables are interpreted exactly like year dummies but are more flexible in the sense that they accommodate the fact that our driving periods can have different lengths and may cover a given year to a smaller or greater extent. For clarity, note that if all driving periods started on Jan 1 and ended on Dec 31 the same year, there would be no difference between our time controls and using year dummies. Our effects are not fully flexible, however: the effect of covering 2003 by 10% is the same whether those 10% are in the start of the year or the end of the year. We have also estimated a specification with dummies at the year-and-month level, but if we furthermore allow each of these dummies to vary by fuel type, there is almost no variation left in the price variable.

[19] Note that only 15% of households in Denmark own more than one vehicle, so household fixed effects are nearly the same as household-vehicle fixed effects.

we instrument for the refined fuel price with the global oil price. Specifically, we use the WTI oil price index, which is based in the United States and captures variation in global oil prices that is quite removed from localized shocks in Denmark.

As fuel prices are an aggregate Denmark-wide variable, one might be concerned about other macroeconomic dynamics, such as fuel prices affecting the macroeconomy and thus indirectly affecting driving. Fortunately, we have access to panel data on household-level incomes, which are a far more precise variable to use to account for how macroeconomic shocks affect consumer decisions than commonly used variables such as GDP or unemployment. Importantly, our income variable is time-varying.[20]

Our specification includes household fixed effects to nonparametrically address time-invariant unobserved household heterogeneity. These household fixed effects are particularly important for identification because they allow us to focus on within-household variation (deviations from the mean) in driving over time. Any sorting into different locations based on time-invariant unobserved preferences will be captured by the fixed effects, as will any time-invariant unobserved household heterogeneity relating to the car choice decision. This means that the variation we use in the work distance variable is purged of longer-term issues of consumer preferences about where to live and where to work.

In using these household fixed effects, we are identifying our work distance-related coefficients largely from households that moved or had their workplace move. The identifying assumption for these coefficients to be quantifying a causal effect is that households move for a variety of reasons (e.g., for a better job, to be closer to family, to reduce their commute, to buy a house, etc.), but they do not move because of a change in unobserved preferences for driving, and similarly, workplaces change for exogenous reasons to the household (e.g., the company moves or the worker gets a new job). For this identification strategy to be problematic one must believe that households have a *time-varying* unobserved preference for driving that happens to be correlated with the choice to move or change jobs. While this seems unlikely, it may be possible. Thus, we run a series of robustness checks exploring different sources of variation using subsamples. For example, a firm relocation should be an exogenous shifter of work distance from the household perspective and we find that our basic results still hold on a subsample of households where only firms relocate.[21] We also find similar results when we explore a subsample of only households that choose to relocate. The fact that we get similar results when we utilize these very different sources of variation in the

[20] We also explore a robustness check where we include municipality-level unemployment, but find that adding this covariate makes no difference to the fuel price coefficient or the tail.

[21] We do not use this as our primary specification because of the unusual sample selection (people who have their firm relocate are not the same as the broader population) and because we lose power from using the much smaller sample. The details are in Appendix C.8.

work distance over time provides further evidence supporting the validity of our identifying assumption.

Another possible identification concern is that our variable for public transport access, the density of bus and train stops, is simultaneously determined with driving. For this to be an issue, the Danish government would have to set public transport access based at least in part on the expected responsiveness to fuel price in an area. If the Danish government set public transport access based on the *level* of driving, this would be fully addressed by the household fixed effects. However, it may be possible that the Danish government set public transport access based in part on responsiveness, perhaps due to a desire to alleviate congestion during times of lower fuel prices. To address any possible endogeneity concern relating to our variable for public transport access, we also examine a specification where we instrument for public transport access with municipality-level population data from 1916, which is the earliest year complete municipality-level population data are available.[22] This approach follows Duranton and Turner (2011), Mulalic, Pilegaard, and Rouwendal (2015), and other recent papers using historical data that determine the location of rail and road lines, but otherwise should not influence outcomes today.

A further possible concern is that urban amenities might change over time, and these changes could be correlated with both work distance and driving. While most urban amenities take decades to develop and can be safely considered constant in our empirical setting, it is possible that some amenities occasionally change. Public transport is an example. We researched this possibility carefully. In general, we find that there were few changes to urban amenities over our time period, and public transport also did not change much. Any Denmark-wide changes would be addressed directly with our time controls, so any concern would have to be a region-specific change that is correlated with work distance. The only meaningful change we uncovered was the addition of a line to the Copenhagen metro between 2002 and 2007 (Mulalic, Pilegaard, and Rouwendal, 2015). This would not be captured in our variable for public transport access. If households or firms happened to move in the 2007-2011 period because of this expansion of the metro line, this could bias our coefficient on the work distance. Thus, we perform a robustness check in which we remove the Greater Copenhagen area from our estimation sample, and again find similar results.

A next possible identification concern could be that we are not including controls for vehicle characteristics, such as fuel economy. Given our household fixed effects addressing time-invariant household preferences for vehicles and the fact that we are using time series variation in fuel prices, this is unlikely to be an issue. This is especially true because in

[22] See `http://www.dst.dk/Site/Dst/Udgivelser/GetPubFile.aspx?id=19910&sid=byersfolk1801`, Accessed June 12, 2017.

Denmark only a small percentage of households have more than one car (less than 15%), so household fixed effects are extremely similar to vehicle fixed effects. However, we also examine a robustness check with household-vehicle fixed effects to address any unobserved heterogeneity at the vehicle (rather than household) level, which again provides similar results only on a smaller subsample and with somewhat less power. Regarding multi-car households, we observe and control for the number of cars, motorcycles, mopeds, campers, vans, and trailers.

4 Results

4.1 The Mean Elasticity of Driving

Table 3 shows the results from estimating the linear fixed effects model in equation (1). A rich set of controls are included in the estimation, but for brevity, we report selected coefficients. Column (1) is the most parsimonious specification, which only controls for seasonality (% of the driving period covering each month), the driving period, and car characteristics. The coefficient on the log fuel price indicates a fuel price elasticity of driving of -0.87. When we add year controls and demographics in column (2), the elasticity changes to -0.30. This indicates the importance of controlling for individual-level demographics as well as using time controls. In columns (3) and (4), we add household fixed effects with and without time controls. Without the time controls, the elasticity is -0.51. Adding time controls brings the elasticity to -0.30, which is our preferred estimate and can be interpreted as a medium-run or one-year elasticity. It may not be surprising that the elasticity moves closer to zero when we nonparametrically control for general time trends in driving since larger economic trends could be correlated with both driving and fuel prices.[23] The ability to simultaneously control for household fixed effects and time controls is a unique advantage of our data, which combines full population data with a over a decade time horizon.

It is worth noting that we find the same fuel price elasticity in columns (2) and (4), which are identical except for the addition of household fixed effects. We take this as an indication that our rich set of controls are capturing the most important determinants of the fuel price elasticity. In particular, we expect that the variables for work distance, company cars, and personal income capture key components of driving demand.

As was also seen in our descriptive analysis, Table 3 shows that driving is increasing in work distance. Even the dummy for whether the work distance is non-zero (recall that it

[23] In Appendix Table 15 we show that the elasticity is robust to the exact functional form of the time controls. In fact, even in a specification with just a linear time trend in the starting year of the period, the elasticity is -0.31.

is censored at 12 km based on how the data are collected) has a positive and statistically significant coefficient. The results indicate that increasing the work distance by one additional km can be interpreted as increasing daily driving by approximately 0.6%. An increase of one standard deviation in work distance would correspond to a change in driving of over a third of the inter-quartile range–an economically significant effect.[24]

The statistically significant coefficients on income suggest that increasing income lowers driving demand for couples, while it increases driving demand for singles. This may be due to wealthier couples being able to afford to live in more geographically advantageous areas, while singles cannot. However, this effect is economically relatively small, which is important because it indicates that factors such as work distance are more economically significant than even income. This provides evidence supporting a focus on commuting in this paper.

The coefficient on the density of bus/train stops per km^2 is statistically significant and negative in columns (2) and (3), which might be expected: better access to public transport should reduce driving. However, because access to public transport is so universal in Denmark (recall Figure 4) and public transport access changes so rarely, there is relatively limited variation in this variable and no time-series variation. When households fixed effects are included, the remaining identifying variation stems from households that move between municipalities. Thus, it may not be surprising that in column (4), which includes the year-fuel controls, the effect becomes statistically indistinguishable from zero. As discussed above, a possible concern with this specification is that the public transport variable is endogenous. As we will see shortly, the estimated fuel price elasticity does not change in our IV specification, further supporting our identification approach here.

A natural question that arises when using the price elasticity for policy analysis is how well the functional form of the demand curve follows a constant elasticity assumption. To explore non-linearities further, Figure 5 plots the non-parametric relationship between log VKT and log fuel price, where both variables have been residualized in a first stage to account for partial correlations with remaining regressors.[25] A key finding is that over a broad range of fuel prices, the functional form in the log-log scale is approximately linear. This supports both the use of the log-log specification, as well as the use of the elasticity over a relatively broad range of fuel price changes. Only at the extremes of fuel price in our data, which are identified from fewer observations, do we observe a nonlinear relationship. Note that the

[24] The standard deviation of work distance is 19.7 and the inter-quartile range where it is observed is 19.3. The VKT distribution has standard deviation 40.2, a 25th percentile of 27.7 km, and a 75the percentile of 58.2.

[25] Note that this is different from a partially linear semi-parametric model. The literature includes a number of estimators for doing this (e.g., Robinson, 1988; Blundell, Horowitz, and Parey, 2012) but standard approaches do not permit fixed effects, which are key in our setting. This is why we focus on the non-parametric relationship between orthogonalized regressors instead.

extremes in fuel price are not the extremes in work distance or VKT, so these extremes are entirely unrelated to our tails. Having a slightly different relationship at very high and very low fuel prices accords with intuition and underscores that the estimates in this paper should be used with caution when the fuel price is much lower or higher than has been typically observed in our dataset.

4.2 Exploring the Tails

We next explore a similar specification to our primary model in equation (1) that includes interactions between the log of the fuel price and a subset of controls to better understand the determinants of the mean elasticity. We denote this subset with $\mathbf{x}^1_{iht}$. The linear model with interactions is given by

$$\log \mathrm{VKT}_{iht} = \gamma_0 \log p_{ft} + \boldsymbol{\gamma_1} \mathbf{x}^1_{iht} \times \log p_{ft} + \mathbf{x}_{iht}\boldsymbol{\beta} + \phi(f,t) + \mu_h + \epsilon_{iht}.$$

In $\mathbf{x}^1_{iht}$, we include variables of interest for our interactions, including work distance and public transport variables. A virtue of this approach is the simplicity of estimation using a standard fixed effects estimator. One feature of this approach is that the model places no restrictions on the values of γ_0 or γ_1, so it is possible to find positive values of the price elasticity of driving for certain groups of households.[26]

Table 4 shows the results from estimating the above equation. Column (1) repeats our primary results with no interactions. Column (2) adds work distance interactions. Column (3) further adds the public transport density interaction (estimated on the slightly smaller sample for which the data including the IV are complete). Column (4) is the same as our primary specification in column (1), only it instruments for public transport density with municipality population in 1916. The first-stage F-statistic is 49.5, demonstrating that we do not have a weak instrument concern. Column (5) instruments for public transport density, as well as the interaction of public transport density with the fuel price. Comparing across columns demonstrates the robustness of our results. Most notably, the elasticity at the mean, presented near the bottom of the table, does not substantially change across columns.

We first focus on the work distance interaction variables. In column (2) in Table 4, we include interactions with a quadratic in work distance. All three work distance coefficients are statistically significant. The interactions of the quadratic in work distance with price

[26] Blundell, Horowitz, and Parey (2012) formulate a nonparametric estimator that imposes negative elasticities, arguing that their findings of an upward sloping demand curve without this restriction must be due to a small sample size. Our sample size is very large and set of controls extensive, so we prefer to not to impose any non-negativity constraints on the elasticity. Of course, it is also theoretically possible that some people respond to rising fuel prices by increasing their driving (e.g., if driving is a complement to an activity that is strongly positively correlated with fuel prices).

tend to show an essentially decreasing linear slope, while the interaction of the dummy for a non-zero work distance implies an overall relationship that perhaps can be described as an inverted-U. To more clearly illustrate this relationship in a nonparametric fashion, we interacted the log fuel price with 19 dummies for work distance bins. Figure 6 illustrates the inverted-U shape. For the shortest work distances (< 12 km), the fuel price elasticity is relatively high in absolute value at nearly -0.40.[27] For slightly longer work distances (long enough where walking and biking are less viable options), the elasticity decreases in absolute value around -0.05 (a value comparable to some estimates in the United States). But then it increases in absolute value again, reaching roughly -0.6 for work distances over 70 km (note the data begins to become sparser by this point).

Figure 6 visually demonstrates two tails based on work distance. A first tail consists of households who own a vehicle but commute very little (either live very close to work or tend to work from home). A second tail consists of households with very long commutes. There is a clear economic intuition that can rationalize these two tails. For the first tail, nearly all of the driving is for non-work trips. Unlike trips between home and work, these trips are more likely to be diverse, which means that drivers may have more opportunities to substitute some of these trips from driving to biking, walking, or public transport (and these trips are likely to be city driving, which for non-electric hybrid cars is less efficient). For the second tail, smaller increases in fuel prices lead to much larger expenditures on driving, providing a strong incentive to consider substitutes. Both tails influence the mean driving elasticity, although the second tail would be expected to be more important for overall fuel consumption due to the greater expenditures on fuel by drivers in this upper tail. In Appendix D we develop a simple analytical model of travel decisions to build further intuition for the two tails.

Moving to public transport, Columns (3) and (5) in Table 4 both show that the coefficient on the public transport interaction with fuel price is statistically significant and negative, indicating that increased public transport access increases the responsiveness to fuel price changes. In column (5), we are instrumenting for public transport density in both the main effect and interaction with a valid and strong instrument, and thus we interpret this result as the causal effect of public transport density on the fuel price responsiveness. The coefficient indicates that an increase of one more bus or train stop per square kilometer will change the fuel price elasticity by -0.01 (recall the mean number of stops per square kilometer is just under 16, with a standard deviation of just over 18). Thus, subtracting one standard deviation from the mean can change the price elasticity by -0.18, which is an economically large difference in responsiveness when the mean elasticity is -0.30. This result is not only to

[27]This result is robust to using the shortest path measure of work distance, which reassures us that the coding at zero is not the reason for our finding of the lower tail. Furthermore, the shape is robust to using a log specification in work distance. The details are in Appendix C.10.

the best of our knowledge new to the literature, but it also underscores that public transport access is a key factor influencing the fuel price responsiveness–a result we will discuss further in section 5.

4.3 Geographic Heterogeneity in the Elasticity of Driving

The results thus far indicate two tails, the first of which involves households with long commutes (upper tail) and the second households who commute very little (lower tail). One might expect to see further evidence of the upper tail in a particularly high responsiveness to fuel price changes in the outskirts of cities, where households have the longest commutes. Similarly, a high responsiveness to fuel price changes in urban areas would build further evidence supporting our hypothesized mechanism for the lower tail. Figure 7 presents the results of a geographical analysis, illustrating the spatial location of the most responsive households. The shading in the figure indicates the predicted elasticity for each observation averaged over the municipalities (darker is more responsive). The three largest cities are labeled.

Two key findings emerge from Figure 7. First, some of the most responsive municipalities are in the largest cities. This aligns with the above evidence suggesting that there is a tail of more responsive drivers with shortest commutes. Second, many of the other most responsive municipalities are in the outskirts of cities. For example, the region just north of Copenhagen has some of the most elastic drivers. These areas tend to have wealthy, high-educated households who often drive for their commute to jobs in Copenhagen. Access to public transport is excellent (recall Figure 4). Similar findings emerge for other areas in the outskirts of urban areas, further building evidence in support of the existence of an upper tail.

4.4 Robustness Checks

We perform an extensive set of robustness checks to confirm our primary results. They are summarized briefly here and discussed in more detail in the Appendix. They broadly confirm our preferred point estimate of the fuel price elasticity, -0.30. Moreover, in our tests, we find that the result of the two tails generally continues to hold. Table 5 provides an overview of the different robustness checks, showing the highest and lowest elasticities that came out in each case; in many cases, the extreme elasticities are to be expected, so we discuss them in the text below, going through each case in turn.

Our first robustness check examines the time window of our sample. Rather than using driving periods that start between July 1998 and December 2007 (which run through 2011), we estimate the same model either starting the sample as late as 2001 or ending the sample as early as 2004. The results bound our preferred estimate in a relatively narrow window: -0.40 to -0.28. The differences may be due to a time-varying elasticity as much as to a lack of robustness. Our second check examines a subsample of the data either controlling for or restricting the sample to driving periods that are of a typical length, which in our setting is two years or four years, plus or minus three months. Our estimated elasticity is quite robust and demonstrates that timing of the inspections is not an identification concern. See Appendix C.1 for full results. The result of the two tails holds for these robustness checks.

The empirical design in this study models both gasoline and diesel car users. This essentially imposes the restriction that drivers of the two different types of cars respond similarly to the fuel price regardless of whether it is gasoline or diesel. The resulting mean elasticity is more useful from a policy perspective, but it masks differences in how diesel and gasoline vehicles are driven. We thus perform a third robustness check where we estimate the same model in equation (1) separately for diesels and gasoline vehicles. We also examine a specification with an interaction between the log fuel price and a diesel dummy. The interaction shows that the gasoline price elasticity is -0.26, while for the diesel segment it is -0.38. Estimating on separate samples yields corresponding elasticities of -0.27 and -0.55. These findings demonstrate that the elasticity is not primarily identified by the differential between gasoline and diesel fuel prices. They also highlight that the diesel segment is more price sensitive, which is consistent with the story behind the two tails since diesel drivers tend to have longer commutes. We perform a similar robustness check for a couples subsample and a singles subsample, finding elasticities of -0.32 and -0.20 respectively. In both cases, the upper tail persists, although the upper tail loses statistical significance in the subsample of singles and diesel cars, perhaps not surprisingly given how much of the data (including many of the people in the tails) are excluded.

The year controls employed in equation (1) are highly flexible, which is important for controlling for potentially correlated time-varying factors, but is also demanding on the data. We thus run robustness checks where we examine alternative time controls (see Appendix C.4). We find that the results are robust to removing our seasonality controls (the % of each month controls) and even removing the year controls for diesel vehicles. When we reduce the time controls to a single linear trend we find estimated elasticity of -0.31. The elasticity is also unchanged if we use year dummies based on the midpoint of the driving period. Both tails remain in all of these checks.

Next, we consider the possibility that fuel prices are endogenous. Denmark is a small country buying both gasoline and diesel on the larger European market, so it is not likely that Denmark-specific demand shocks lead to an endogeneity issue. However, it is possible that such an issue may occur. Thus, our robustness check instruments the fuel price using the WTI crude oil price, which is not only physically located in the United States, but is determined by global oil market movements (see Appendix C.6). It is hard to imagine a small localized demand shock in Denmark possibly affecting the WTI crude oil price. At the same time, the first stage regression indicates that it is a strong instrument, since oil is the primary feedstock for refined fuel. The 2SLS fuel price elasticity estimate is a statistically significant -0.36. This estimate is close to our preferred estimate of -0.30, and we view this as confirming our estimate. Given the standard errors, these two estimates are not statistically significantly different.

One possible concern with our study is that we use the work distance variable reported on tax returns and households with work distances less than 12 kilometers are not eligible for the deduction, so are coded to zero. While we are very confident that this coding is done correctly (due to the threat of random audits) our primary results do not have much variation between zero and 12 kilometers. Thus, we also estimate our model using the 79.6% subsample for which we have the actual distance between the home and workplace using a shortest-distance algorithm. The estimated elasticity is -0.37 and again is not statistically significantly different from our preferred estimate of -0.30, confirming our estimate on the larger sample using the better measure that captures intensity of commuting. Equally importantly, the result of the two tails again holds using this alternative work distance measure (see Appendix C.10).

We also perform a series of robustness checks examining the possibility of selection into different vehicles that may lead car characteristics to be endogenous (e.g., see Gillingham (2013) or Munk-Nielsen (2015)). We find our results quite robust to the exact choice of vehicle characteristics that are included. More importantly, we also run a specification with household-vehicle fixed effects, which would address any concern about the inclusion or exclusion of any particular vehicle characteristic, such as fuel economy. The estimated elasticity is -0.31 and the tail story continues to hold, confirming that our results are robust, albeit with some loss of statistical significance due to the smaller sample (see Appendix C.7).

Next, we consider the source of variation over time in the work distance measure. We explore different sources of variation, estimating the model on subsamples of only households that choose to relocate or only households where the firm of at least one spouse relocates at some point during the sample.[28] These are much smaller subsamples that are useful for understanding the variation driving our results. Using the 49,074 observation subsample where only the firm relocates, we find a mean elasticity of -0.36. Using the 3,244,793 observation subsample where a household moves, we find a mean elasticity of -0.41. In both cases, the

[28]A relocation is defined by Statistics Denmark based on a change in address for the firm or one of the firm's work locations, or if the work location of all workers in one location changed to a different location within the same firm.

two tails based on work distance are still present and the original mean elasticity is within the 95% confidence interval (due to the substantially higher standard error on the estimates). The details are in Appendix C.8.

Finally, we examine the possibility that gasoline prices affect driving both directly, through a consumer response to the price at the pump, and indirectly, through macroeconomic effects. Our empirical specification includes personal income at the household level as a covariate, which should address macroeconomic effects that may influence driving. We first added municipality-level unemployment as another covariate to the specification run in column 4 of Table 3. The coefficient on the log fuel price changes from -0.301 to -0.299 and is again statistically significant. The work distance interaction coefficients are again identical to a rounding error, indicating that the tail remains. We also include register data measure of unemployment as controls (see Appendix C.9), and the estimate only changes from -0.301 to -0.298 with the tail estimates similarly unchanged.

Second, we explored the relationship between unemployment rates and fuel prices by regressing the unemployment rate on fuel prices. When we look at the correlation without any other covariates, we see a very small and negative relationship. When we add further controls, especially time controls, the coefficient on fuel price becomes statistically insignificant. While not definitive, we view these findings as suggestive that fuel prices seem to have relatively little effect on the unemployment rate in the Danish economy. We further divide the sample into four bins based on the average work distance of the municipality and run the same regression of the unemployment rate on the fuel price on each subsample. We again find a small effect, and importantly, the effect is roughly the same over the four bins of work distance. This suggests that the effect of fuel prices on the unemployment rate does not vary based on commuting distance.

Third, we ran a type of placebo test by collecting data on the price of other fuels that might also influence the economy but that are not used for driving. We ran the same regression as in column 4 of Table 3 only we also include covariates for the log of the natural gas price and the log of three coal prices for coal that is occasionally imported into Northern Europe: from the United States, Australia, and South Africa.[29] For all four prices (regardless of whether we include them separately or together), we find that the coefficient values are very close to zero and statistically insignificant. This provides further evidence that the results we are finding are from fuel prices influencing directly influencing driving, rather than influencing the broader economy and indirectly influencing driving.

[29]The source of the natural gas prices is Statistics Denmark, while the source of the coal prices is the U.S. Energy Information Administration.

5 Implications for Policy

In this section, we discuss the implications for policy of our results. The finding of the two tails is particularly important in light of discussions in Europe about carbon taxation. An economy-wide carbon tax would raise the price of fossil fuels in all sectors, including transport. Perhaps the most challenging political obstacle to carbon pricing is the potential for perverse distributional consequences. Vertical distributional consequences (i.e., across income) are clearly important, but what we emphasize here is that distributional consequences across geography due to differences in work distance can also be very important. As has been well-documented in the news, in many countries around the world–including Denmark–there is political discontent in areas outside of the large cities. For example, most of the votes for Brexit in the United Kingdom were from areas outside of London and other large cities. There is similar discontent in Denmark, with the anti-immigration party receiving the most votes in areas outside of the largest cities. Politicians in Denmark have been acutely aware of this discontent and have gone to great lengths to provide support for these areas. For example, they have moved thousands of government jobs away from the capital region urban center, at a cost of millions of dollars. Thus, the differing distributional effects over geography–which are influenced greatly by the work distance–are directly relevant for the political acceptability of any policy that disproportionately impacts poorer rural regions more, including carbon pricing.

The tails have clear implications for the differing distributional effects. The upper tail is especially important for it tends to consist of drivers who live outside of the major urban areas and who drive relatively more (just under 15% of the population). These upper tail households face a high burden from fuel prices and a disproportionately higher burden when fuel prices rise. Thus, we perform a simple set of illustrative calculations to demonstrate the importance of our finding of the upper tail. The focus here is on short-run equilibria; the much longer-run relationships between location choice, labor markets, and driving responsiveness are outside the scope of this analysis.

For illustrative purposes, we focus our analysis on the short-run effects of an increase in fuel taxes that leads the fuel price to rise by 1 DKK/l for both gasoline and diesel.[30] The average gasoline price over the 1998 to 2008 period is 9.01 in 2005 DKK, so this represents a substantial price increase, but it is within the range of the variation in our data.[31] With a fuel price elasticity of driving of -0.30, the proposed increase of 1 DKK/l translates into

[30]This construction sidesteps issues of passthrough; this equal a 1 DKK/l tax if the passthrough rate was 100%.

[31]This maps to an increase in gasoline prices of $0.57 per gallon based on the June 18, 2015 exchange rate of 6.54 DKK per dollar.

a 3.3% short-run reduction in driving (it could be expected to be larger in the long-run). If this occurred due to a tax policy, then fuel tax revenue would increase by 13.2%.[32] We abstract from externalities in this analysis in order to focus on distributional effects.

5.1 Distributional Effects

We focus on the distributional impacts on consumers, calculating both the transfers from consumers to the government (i.e., government revenue, calculated as $(p_1 - p_0)VKT(p_1))$ and the direct loss in consumer surplus from reducing the amount driven ($\int_{p_0}^{p_1} VKT(p) - VKT(p_1)\mathrm{d}p$).[33] Ignoring externalities, these two quantities combined amount to the change in consumer surplus (ΔCS) prior to any redistribution of revenues. In Denmark, the revenues from fuel taxes are added to the general fund, for use on all government spending. For illustrative purposes, we assume that the funds are redistributed (or provided in equivalently-valued services) on an equal basis to households, an assumption that is reasonable for the Danish setting. Figure 8 provides a simple graphical presentation of the two areas being calculated (note in our calculations, we are not imposing linearity on the demand curve).

Table 6 shows the effects of the increase in fuel prices due to the fuel tax broken down by work distance group. A key finding emerges from this analysis: Households in the upper tail of the work distance distribution are more able to substitute away from driving to reduce their tax burden, as shown by the greater direct change in consumer surplus (ΔCS in column (1)). They drive much more (compare the the upper tail to the full sample in column (2)), but can substitute away from driving to a greater degree, consistent with a larger discrete shift to public transport. This comes about from a flatter demand curve, implying a larger direct loss in consumer surplus, but a greater ability to avoid the tax by changing behavior.

Of course, the revenue from the transfer is not lost, but can be used for a variety of purposes or redistributed. For illustrative purposes, we calculated the effects of a uniform redistribution. We find that those in the upper tail face a much larger burden from the fuel tax policy than average (-15 DKK per household per day), while those in the lowest tail receive a gain (5.82 DKK per household per day). Quantifying these effects is crucial for understanding the political acceptability of carbon taxes and more broadly the distributional effects after redistribution is a question that policymakers in Denmark care deeply about. Compensatory policies, such as a policy that redistributes more of the revenue to the upper

[32] At 9 DKK/l, the increase of 1 DKK/l is 11.1%, which at an elasticity of -0.30 translates to a change in driving of 3.33%. Over the sample period, taxes make up 64.87% of the gasoline price, corresponding to 5.84 DKK/l at 9 DKK/l. An increase in 1 DKK/l thus corresponds to an increase of 17.13% in taxes, giving a total relative change in taxes of $(1 + 0.1713) \times (1 - 0.0333) = 13.2\%$.

[33] This direct loss in consumer surplus is the Harberger triangle or deadweight loss assuming that Danes are price-takers for fuel and that any externalities have already been internalized from previous taxes.

tail, would be necessary if the policy goal is an an equal distributional effects by group.

In contrast to the above results, we can see in Figure 9 that this illustrative policy is actually quite progressive from a vertical equity standpoint, with a higher burden faced by wealthier households. This underscores the usefulness of focusing on more than just vertical distributional equity.

A natural question is how the Danes have come to accept taxes in excess of 50% of the price at the pump. We posit that public transport may play a key role. Our causal estimate of the effect of public transport on price sensitivity indicates that while the average elasticity is -0.30, if we reduce the public transport density by one standard deviation, the average elasticity becomes -0.13. Thus, public transport is instrumental in providing households with options for avoiding the large fuel taxes. A second implication though is that fuel taxes are especially distorting for households in the upper tail in leading to changes in behavior (which may be efficiency-improving to the extent that there are uninternalized externalities).

5.2 Public Transport and Reconciling Elasticities

Our results are also insightful for reconciling elasticities between Europe and the United States (recall the literature reviewed in section 1). First, note that in Figure 6 the average elasticity outside of the two tails is closer to zero than -0.3 and for much of the sample, closer to zero than -0.2. This squarely puts the elasticity in the range of elasticities commonly estimated in the United States. Our estimate of the causal impact of public transport provides an insight into why.

Recall that we also performed a simulation that reduces the density of public transport in Denmark by one standard deviation. This change brings the density to roughly in the range of levels common in the United States, and thus sheds light on the extent to which public transport can reconcile the elasticities. We found that the elasticity becomes -0.13. This finding helps to further bridge a gap in the literature by helping to reconcile the differences between elasticity estimates that are an important input to policy.

6 Conclusion

This paper uncovers two tails of more responsive drivers than most of the population in Denmark. The first tail is a small group of consumers living in the outskirts of cities with long commutes, but adequate access to public transport. The second tail is a group of drivers who commute very little and tend to live in cities. There is an economic intuition for each of these tails: households with long commutes can readily switch to public transport, while households who commute very little largely use their vehicles for a diverse set of non-work trips, many of which can easily be switched to other modes of transport. Households that are not in these two tails tend to be much more inelastic in their response to fuel price changes.

This finding is important for two reasons. First, it is particularly relevant for the political economy of carbon pricing. Those in the upper tail face the largest burden from a rise in fuel prices, but access to public transport allows them to more readily substitute away from driving, thereby reducing their burden. Our illustrative calculations quantify the importance of this effect, showing how ample access to public transport reduces the burden on the upper tail through access to an appealing substitute.

Second, the finding of the two tails helps to reconcile the results of studies in Europe and the United States that estimate fuel price elasticities. Most studies in the United States show fuel price elasticities in the range of -0.10 to -0.30, while those in Europe over the same time frame generally tend to be higher in absolute value. Our preferred mean price elasticity value is -0.30, but we estimate the causal impact of public transport to show that if we remove ample access to public transport, this elasticity changes to -0.13, which is much more in line with recent estimates from the United States for short-run fuel price elasticities of driving. One implication of these results is that if the United States improved its public transport opportunities, the upper tail of responsiveness could emerge there as well, potentially reducing some of the political challenges to fuel taxation and carbon pricing.

References

Bento, Antonio, Maureen Cropper, Ahmed Mushfiq Mobarak, and Katja Vinha. 2005. "The Effects of Urban Spatial Structure on Travel Demand in the United States." *Review of Economics and Statistics* 87(3):466–478.

Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *American Economic Review* 99(3):667–699.

Blundell, Richard, Joel Horowitz, and Matthias Parey. 2012. "Measuring the Price Responsiveness of Gasoline Demand: Economic Shape Restrictions and Nonparametric Demand Estimation." *Quantitative Economics* 3 (1):29–51.

Borenstein, Severin. 2017. "Creative Pie Slicing To Address Climate Policy Opposition." *Energy Institute at Haas Blog* .

Borenstein, Severin and Lucas Davis. 2016. "The Distributional Effects of US Clean Energy Tax Credits." In *Tax Policy and the Economy, Volume 30*. University of Chicago Press, 191–234.

Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld. 2008. "A Meta-analysis of the Price Elasticity of Gasoline Demand. A SUR Approach." *Energy Economics* 30 (5):2105–2122.

Brownstone, David and Thomas Golob. 2010. "The Impact of Residential Density on Vehicle Usage and Energy Consumption." *Journal of Urban Economics* 65 (1):91–98.

Coglianese, John, Lucas Davis, Lutz Kilian, and James Stock. 2016. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." *Journal of Applied Econometrics* forthcoming.

Dahl, Carol and Thomas Sterner. 1991. "Analysing Gasoline Demand Elasticities: A Survey." *Energy Economics* 13 (3):203–210.

Davis, Lucas and Lutz Kilian. 2011. "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 26 (7):1187–1214.

Davis, Lucas and Chris Knittel. 2016. "Are Fuel Economy Standards Regressive?" *NBER Working Paper 22925* .

De Borger, B., I. Mulalic, and J. Rouwendal. 2016a. "Substitution between cars within the household." *Transportation Research A* 85:135–156.

De Borger, Bruno, Ismir Mulalic, and Jan Rouwendal. 2016b. "Measuring the Rebound Effect with Micro Data: A First Difference Approach." *Journal of Environmental Economics and Management* 79:1–17.

Duranton, Gilles and Matthew Turner. 2011. "The Fundamental Law of Road Congestion: Evidence from the US." *American Economic Review* 101 (6):2616–52.

Edelstein, Paul and Lutz Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices." *Journal of Monetary Economics* 56 (6):766–779.

EEA. 2017. "Greenhouse Gas Emissions from Transport Indicator Assessment." *European Environment Agency* .

EPA. 2017. "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2015." *U.S. Environmental Protection Agency* .

Espey, Molly. 1998. "Gasoline Demand Revisited: An International Meta-analysis of Elasticities." *Energy Economics* 20 (3):273–295.

Frondel, Manuel and Colin Vance. 2013. "Re-Identifying the Rebound: What about Asymmetry?" *Energy Journal* 34 (4):43–54.

Gillingham, Kenneth. 2013. "Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices." *Yale University Working Paper* .

———. 2014. "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock." *Regional Science & Urban Economics* 47 (4):13–24.

Gillingham, Kenneth, Alan Jenn, and Ines Azevedo. 2015. "Heterogeneity in the Response to Gasoline Prices: Evidence from Pennsylvania and Implications for the Rebound Effect." *Energy Economics* 52 (S1):S41–S52.

Glaeser, Edward and Matthew Kahn. 2010. "The Greenness of Cities: Carbon Dioxide Emissions and Urban Development." *Journal of Urban Economics* 67 (3):404–418.

Graham, Daniel and Stephen Glaister. 2004. "Road Traffic Demand Elasticity Estimates: A Review." *Transport Reviews* 24 (3):261–274.

Grazi, Fabio, Jeroen van den Bergh, and Jos van Ommeren. 2008. "An Empirical Analysis of Urban Form, Transport, and Global Warming." *The Energy Journal* 29 (4):97–122.

Hamilton, James. 2009. "Understanding Crude Oil Prices." *Energy Journal* 30 (2):179–206.

Hausman, Jerry and Whitney Newey. 1995. "Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss." *Econometrica* 63 (6):1445–1476.

Hughes, Jonathan, Christopher Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29 (1):93–114.

Hymel, Kent M. and Kenneth A. Small. 2015. "The Rebound Effect for Automobile Travel: Asymmetric Response to Price Changes and Novel Features of the 2000s." *Energy Economics* forthcoming.

Jacobsen, Mark. 2013. "Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity." *American Economic Journal: Economic Policy* 5 (2):148–187.

Kayser, Hilke. 2000. "Gasoline Demand and Car Choice: Estimating Demand Using Household Information." *Energy Policy* 22 (3):331–348.

Knittel, Christopher and Ryan Sandler. 2018. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." *American Economic Review: Economic Policy* forthcoming.

Levin, Lewis, Matthew Lewis, and Frank Wolak. 2017. "High-Frequency Evidence on the Demand for Gasoline." *American Economic Journal: Economic Policy* 9 (3):314–347.

Levinson, Arik. 2016. "Energy Efficiency Standards Are More Regressive Than Energy Taxes?" *NBER Working Paper 22956* .

Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy* 6 (4):302–342.

Linn, Joshua. 2016. "The Rebound Effect for Passenger Vehicles." *Energy Journal* 37 (2):257–288.

McFadden, Daniel. 1974. "The Measurement of Urban Travel Demand." *Journal of Public Economics* 3 (4):97–122.

Mulalic, Ismir, Ninette Pilegaard, and Jan Rouwendal. 2015. "Does Improving Public Transport Decrease Car Ownership? Evidence from the Copenhagen Metropolitan Area." *Tinbergen Institute Discussion Paper 2015-139* .

Munk-Nielsen, Anders. 2015. "Diesel Cars and Environmental Policy." *Working Paper* .

Poterba, James. 1989. "Lifetime Incidence and the Distributional Burden of Excise Taxes." *American Economic Review* 79 (2):325–330.

———. 1991. "Is the Gasoline Tax Regressive?" In *Tax Policy and the Economy, Volume 5*. MIT Press, 145–164.

Robinson, Peter M. 1988. "Root-N-consistent semiparametric regression." *Econometrica* :931–954.

Small, Kenneth A. and Kurt van Dender. 2007. "Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect." *Energy Journal* 28 (1):25–51.

Sterner, Thomas. 2012. "Distributional Effects of Taxing Transport Fuel." *Energy Policy* 41:75–83.

West, Sarah. 2004. "Distributional Effects of Alternative Vehicle Pollution Control Technologies." *Journal of Public Economics* 88 (3-4):735–757.

West, Sarah and Roberton Williams. 2004. "Estimates from a Consumer Demand System: Implications for the Incidence of Environmental Taxes." *Journal of Environmental Economics and Management* 47 (3):535–558.

Williams, Roberton, Hal Gordon, Dallas Burtraw, Jared Carbone, and Richard Morgenstern. 2015. "The Initial Incidence of a Carbon Tax across Income Groups." *National Tax Journal* 68 (1):195–214.

Tables

Table 1: Summary Statistics

	Mean	Std dev
Panel A: Vehicle Characteristics		
Vehicle-km-traveled (km/day)	46.6	(40.2)
Vehicle-km-traveled (log)	3.027	(0.855)
Vehicle Weight (kilograms)	1,671	(331)
Car age at start of period (years)	6.97	(5.17)
$\mathbb{1}$(Diesel vehicle)	0.14	(0.35)
$\mathbb{1}$(Van)	0.08	(0.27)
Driving period length (years)	2.34	(0.89)
% of period owned by this owner	0.79	(0.30)
# additional cars owned	0.34	(0.60)
# vans owned	0.05	(0.24)
# motorcycles owned	0.05	(0.27)
# mopeds owned	0.03	(0.16)
Panel B: Household Characteristics		
Reported work distance (km)	12.2	(19.7)
$\mathbb{1}$(Work distance > 12km)	0.50	(0.50)
Reported work distance (log)	3.027	(0.855)
Actual work distance (km)^a	23.4	(35.6)
Gross income (DKK)	574,056	(627,921)
Gross income-couples (DKK)	646,638	(628,011)
Gross income-singles (DKK)	320,975	(558,098)
Gross income (log)	13.09	(0.579)
$\mathbb{1}$(Couple)	0.78	(0.42)
Age (oldest household member)	49.8	(14.4)
Number of children	0.76	(1.02)
$\mathbb{1}$(Urban municipality)	0.16	(0.36)
Bus/train stops per km^2	15.9	(18.4)
$\mathbb{1}$(Access to company car)	0.03	(0.18)
$\mathbb{1}$(Self-employed)	0.10	(0.30)
Observations	5,855,446	

An observation is a vehicle driving period between two odometer readings. All Danish kroner (DKK) are inflation-adjusted to 2005 DKK. We are not permitted to present the min and max due to Statistics Denmark rules. Average log variables are only averaged over observations with strictly positive values. ^{a}: The actual work distance is available for 79.6% of the sample.

Table 2: Means of Selected Variables Stratified by Work Distance

	WD=0	WD$\in (0; 30]$	WD>30
Daily Vehicle-kilometers-traveled	39.90	51.19	62.87
Gross income (DKK)	508,542	638,515	692,646
Gross income-couples (DKK)	589,690	690,251	739,434
Gross income-singles (DKK)	304,080	350,612	389,282
Couple dummy	0.72	0.85	0.8664
Reported work distance (km)	0.00	16.95	49.45
Work distance $>$ 12 km dummy	0.00	1.00	1.00
Work distance (door-to-door)[a]	12.66	24.38	49.44
Number of children	0.58	0.99	1.00
Urban dummy	0.19	0.13	0.10
Self employed dummy	0.12	0.07	0.05
Bus/train stops per km^2	17.45	14.14	13.29
Diesel dummy	0.11	0.16	0.25
Car age at start of period	7.33	6.68	6.22
# additional cars owned	0.26	0.39	0.51
Observations	3,253,413	1,762,952	839,081

WD denotes the work distance.

An observation is a vehicle driving period between two odometer readings.

[a]: The door-to-door work distance is available for 79.6% of the sample.

Table 3: Estimations of Driving Demand

	OLS		Household FE	
	(1)	(2)	(3)	(4)
$\log p^{\text{fuel}}$	-0.87***	-0.30***	-0.51***	-0.30***
	(0.028)	(0.016)	(0.012)	(0.016)
Work Distance (WD) controls				
WD		0.006***	0.006***	0.006***
		(0.0004)	(0.0001)	(0.0001)
WD squared		-0.00002***	-0.00002***	-0.00002***
		(0.000003)	(0.000001)	(0.000001)
WD non-zero		0.076***	0.010***	0.015***
		(0.0062)	(0.0016)	(0.0016)
Macro and Transport controls				
log gross income-couple		-0.03***	-0.03***	-0.03***
		(0.005)	(0.002)	(0.002)
log gross income-single		0.03***	0.03***	0.02***
		(0.002)	(0.003)	(0.003)
Bus/train stops per km^2		-0.002*	-0.0003**	0.00004
		(0.0006)	(0.0001)	(0.0001)
Year-fuel controls	No	Yes	No	Yes
Number of children	No	Yes	Yes	Yes
Self-employed	No	Yes	Yes	Yes
Access to company car	No	Yes	Yes	Yes
Household FE	No	No	Yes	Yes
R^2	0.20	0.34	0.68	0.68
N	5,855,446	5,855,446	5,855,446	5,855,446

Dependent variable is the log VKT. An observation is a driving period. All specifications include driving period controls, % of each month controls, a quadratic in age of both the male and female of the household, and some car characteristics (a quadratic in weight, diesel dummy, van dummy, a quadratic in vehicle age, and number of vehicles of each type owned by the household). The within R^2 is reported for the household fixed effects specifications. Robust standard errors clustered at the municipality level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 4: Effect of Work Distance and Public Transport on Responsiveness

	Household FE			FE-IV	
	(1)	(2)	(3)	(4)	(5)
$\log p^{\text{fuel}}$	-0.30***	-0.36***	-0.47***	-0.31***	-0.47***
	(0.016)	(0.02)	(0.01)	(0.01)	(0.01)
Work Distance (WD) interactions					
WD $\times \log p^{\text{fuel}}$		-0.01***	-0.01***		-0.01***
		(0.0006)	(0.0005)		(0.0004)
WD squared $\times \log p^{\text{fuel}}$		0.00001**	0.00001***		0.00001***
		(0.000005)	(0.000003)		(0.000003)
WD non-zero=1 $\times \log p^{\text{fuel}}$		0.37***	0.36***		0.32***
		(0.02)	(0.01)		(0.01)
Public transport interaction					
Bus/train stops per km^2 $\times \log p^{\text{fuel}}$			-0.004**		-0.011***
			(0.001)		(0.001)
Household FE	Yes	Yes	Yes	Yes	Yes
Mean elasticity	-0.30***	-0.29***	-0.31***	-0.30***	-0.32***
R^2	0.18	0.18	0.18	0.17	0.17
N	5,855,446	5,855,446	4,773,953	4,773,953	4,773,953

Dependent variable is log VKT. An observation is a driving period. All specifications include the main effects for each interaction. All specifications have a year-fuel controls, public transport density, driving period controls, % of each month controls, a quadratic in age of both the male and female of the household, and some car characteristics (a quadratic in weight, diesel dummy, van dummy, a quadratic in vehicle age, and number of vehicles of each type owned by the household). Column (4) instruments for public transport density with the municipality population in 1916. Column (5) does the same and also instruments for the public transport interaction with the municipality population in 1916 interacted with fuel prices. The mean elasticity takes the mean of the elasticity predicted over all observations. The within R^2 is reported for the household fixed effects specifications, but between is reported for the FE-IV specifications. Robust standard errors clustered at the municipality level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 5: Overview of Robustness Checks: Ranges of Elasticity Estimates

Name	Elasticity Range
Years in the sample	[-0.40;-0.26]
Length of driving periods	[-0.30;-0.27]
Fuel type	[-0.54;-0.26]
Singles or couples	[-0.32; -0.25]
Time controls	[-0.31;-0.29]
Instrumenting with oil price	-0.36
Using actual work distance	-0.44[a]
Household-vehicle fixed effects	-0.33
Alternative fuel prices	[-0.41; -0.39][b]

[a]: Result on a different subsample where the comparable elasticity estimate from our preferred model was -0.37.

[b]: Also a different subsample where the comparable elasticity is -0.39.

Table 6: Distributional Effects By Work Distance (per Household)

	Observations	(1) ΔCS	(2) Transfer	(3) DWL
Full Sample	4,773,953	-43.1	-42.4	-0.70
Lower Tail	2,666,643	-36.6	-35.9	-0.70
Upper Tail	683,347	-57.5	-56.2	-1.28

Notes: WD denotes the work distance. 'Lower Tail' refers to households with WD<12 km; 'Upper Tail' refers to households with WD>30km. 'ΔCS' denotes the uncompensated total change in consumer welfare. 'Transfers' refers to the tax revenues; the (negative) payment from the household to the government (given by $(p_1 - p_0)VKT(p_1)$). 'DWL' is the direct change in consumer surplus (the deadweight loss if Danes are price-takers and externalities have been internalized). All numbers indicate the mean within the relevant sub-population and are daily numbers.

Figures

Figure 1: Vehicle Kilometers Traveled (VKT) and Fuel Price

(a) Levels (b) Logarithms

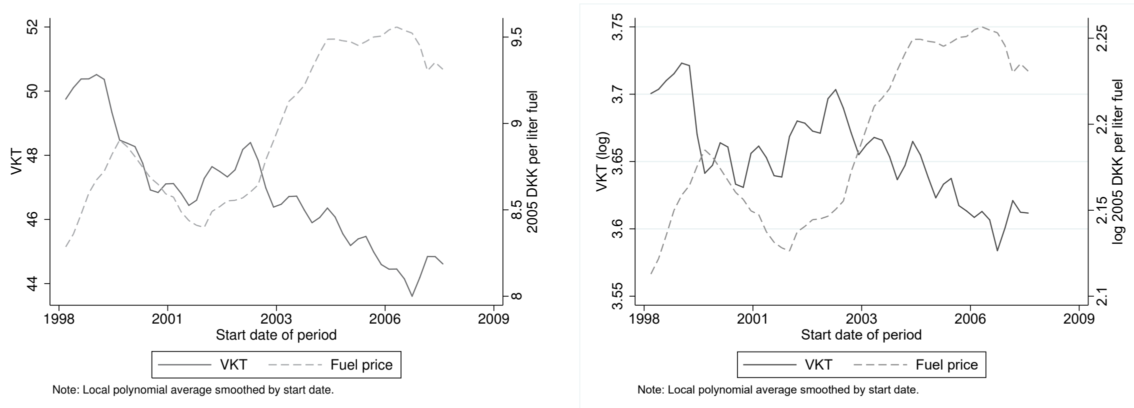

Figure 2: Percentiles of Driving Over Time

(a) VKT (b) log VKT

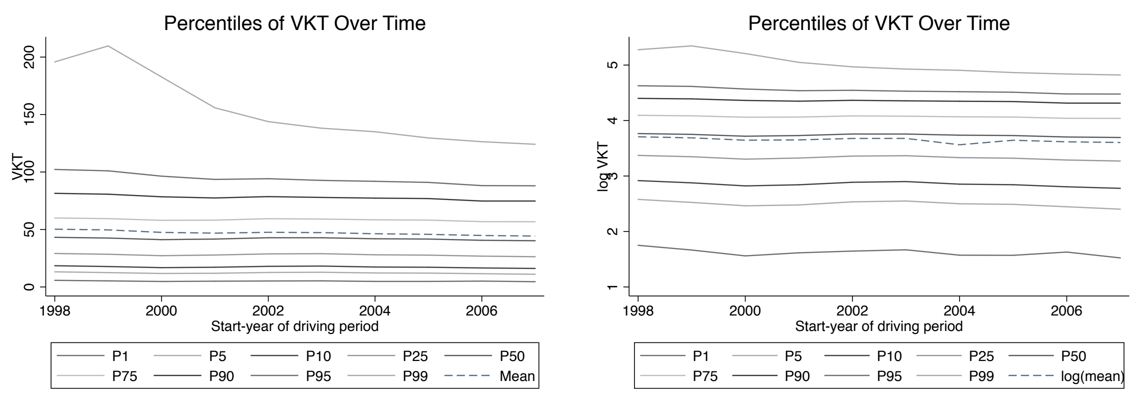

Figure 3: Nonparametric Relationship Between Driving and Work Distance

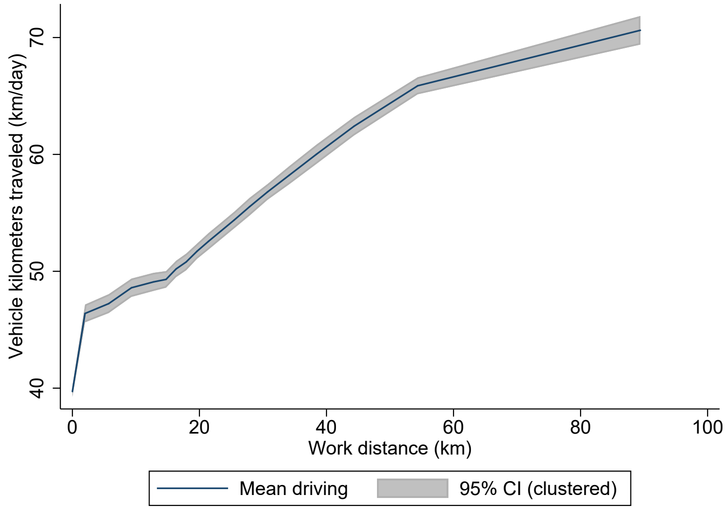

Notes: The line shows raw means within bins of work distance. We have split work distance according to 40 bins with equally many observations, except for the zero bin (which holds half of observations). To construct the confidence interval, we regressed driving on dummies for all bins, excluding a constant in the regression. We then used the standard errors (clustered at the municipality level) for these dummies to construct the 95% confidence intervals around the means, multiplying by ±1.96. There are bins for work distance between 0 and 12 due to households that work part time but have work distances above 12 km.

Figure 4: Average Work Distance by Municipality (Panel (a)) and Public Transport Stops (Panel (b))

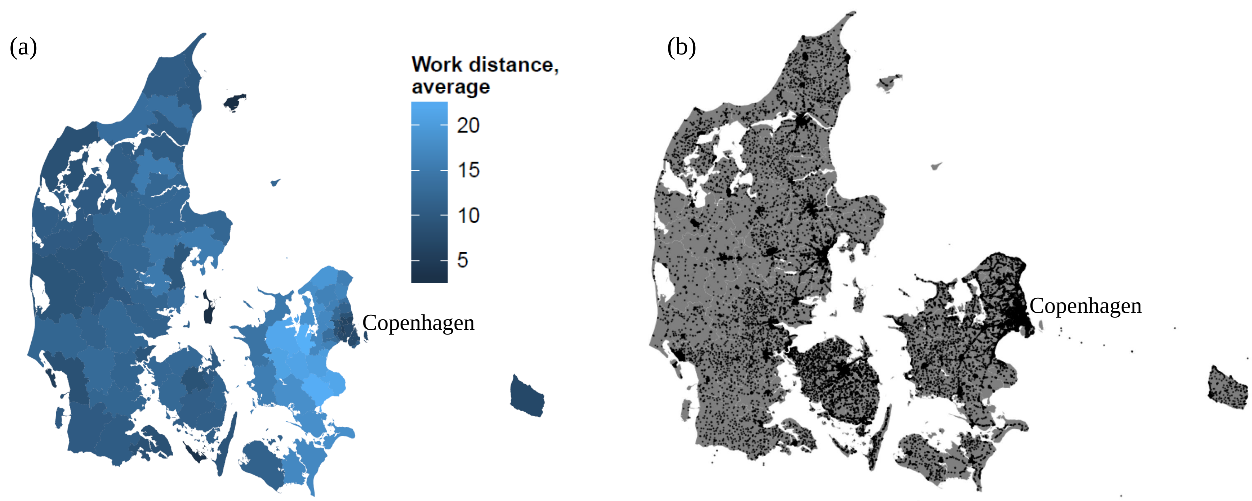

Figure 5: Nonparametric Demand Curve

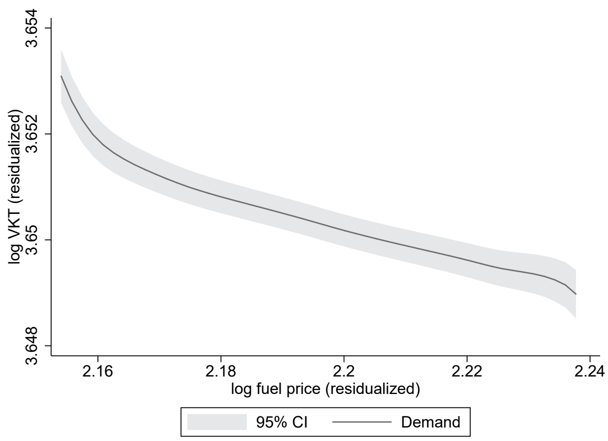

Note: This figure is constructed by first residualizing log VKT and log fuel price and then running a local polynomial regression with the default bandwidth in Stata. The first-stage residualization is done by regressing the variable on a full set of controls (excluding the fuel price) and household fixed effects and extracting the residual. We show the relationship for residualized log fuel prices between the 1st and 99th percentile of log fuel prices (2.15 and 2.24).

Figure 6: Relationship Between the Fuel Price Elasticity and Work Distance (in kilometers)

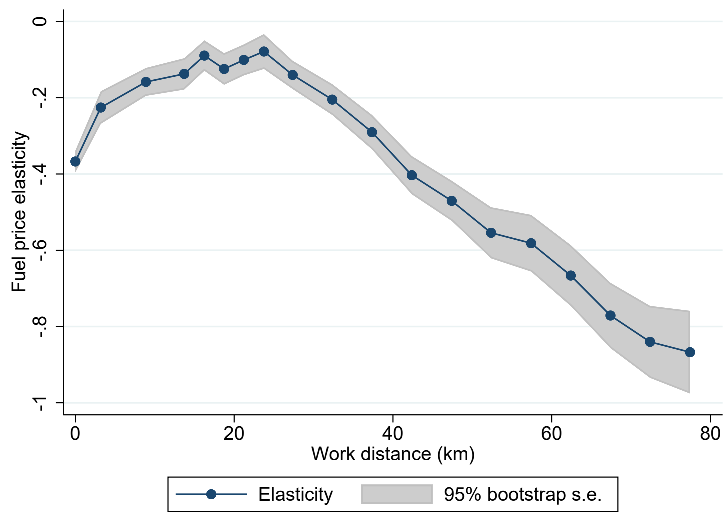

Note: We have split work distance according to bins with 5% of households each except for the zero bin (which holds half of all observations due to the censoring). Next, we regressed log VKT on all controls except WD but with a full set of dummies for bins of WD interacted with the fuel price. On the x-axis is shown the mean WD within bins; on the y-axis the coefficient estimate for that bin, i.e. the mean fuel price elasticity in that bin. It can be considered a semi-nonparametric model for the fuel price elasticity in the sense that as the sample size increases, one could add more dummies interacted with the fuel price.
Next, we ran 1000 bootstrap estimations and show ± 1.96 standard error around each coefficient estimate. To avoid the problem of having both a main effect (coefficient on log price) in addition to a coefficient on a bin-dummy times the fuel price, we instead opted to omit the log fuel price (main effect) and instead include a full set of all bin dummies interacted with the log fuel price. This does not change any point estimates but standard errors can now be directly used without taking into account the covariance with the main effect.

Figure 7: The Average Elasticity by Municipality

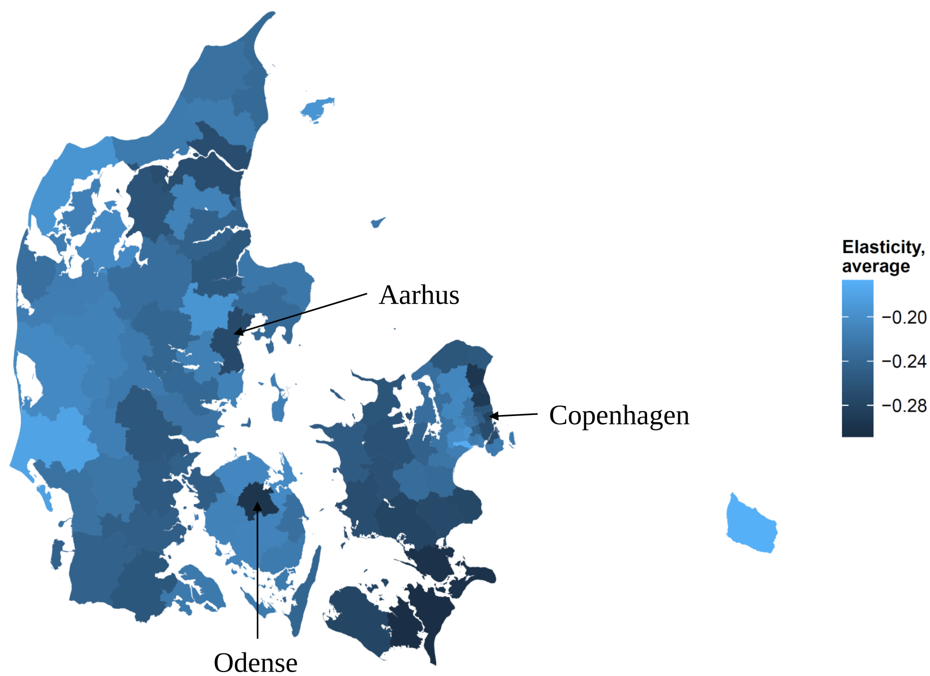

Figure 8: Welfare Effects of a Fuel Tax Increase

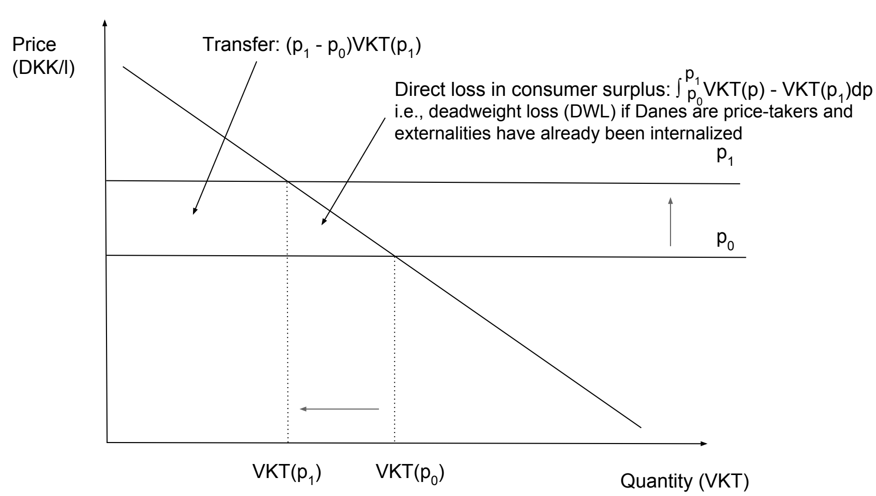

Note: The graph illustrates the welfare effects of a fuel tax that results in a price increase from p_0 to p_1.

Figure 9: The Illustrative Policy is Progressive

(a) Net ΔCS in levels (b) Net ΔCS relative to income

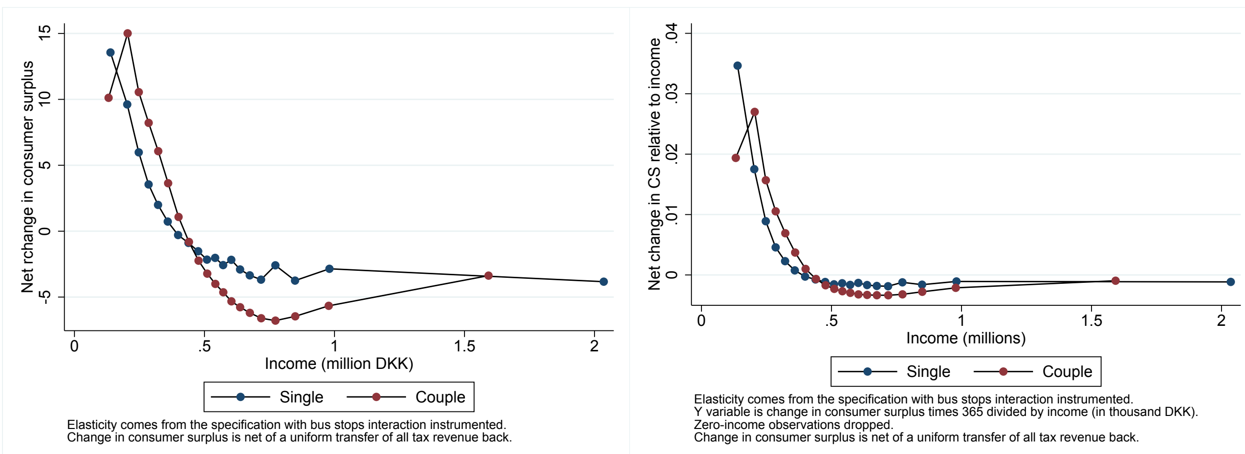

Note: Both panels show the change in consumer surplus from an increase in fuel prices of 1 DKK per liter by household income, split for couples and singles (singles obviously have lower total income). In panel (b), the x-axis, the change in consumer surplus is measured relative to household income.

ONLINE APPENDIX

A Data Appendix

Given the comprehensiveness and richness of the dataset used in this study, we include this online appendix to describe the data in more detail and elaborate on the sample selection.

A.1 Development of the Final Dataset

As a general rule, we focused our data cleaning efforts on avoiding dropping observations to maximize our dataset coverage. Our data cleaning proceeded as follows. First, we converted the dataset into one where the observation is a driving period and all negative driving observations or other observations with critical missing data are dropped. This dataset has 10,994,333 observations. We further restrict the data to have driving periods that began after July 1, 1998 (since tests were not mandatory prior to this date) and before January 1, 2008 (this ensures that our sample is not biased away from new vehicles). These restrictions leave us with 7,254,893 observations. After this, we delete observations where the length of the driving period, which we call *years to test*, is not either between 1 and 2.5 years or between 3.5 and 4.5 years. This is chosen to balance not getting too many observations with unexplained lengths of the driving periods while also accounting for early phase-in, which led to a number of 1-year periods in 1999 and 2000. This leaves us with 6,877,185 driving periods. We have missing demographic variables for 277,294 observations, which brings us to 6,599,891 observations. We drop 178 outlier observations with VKT greater than 10,000 km/day. Finally, because we use household fixed effects, we drop 744,267 households that are only observed once in the dataset. This brings us to our final sample size of 5,855,446.

To clarify how these observations are distributed over time, Table 10 shows a histogram of the start year of the driving period. The low number in the first year is due to the sample selection criterion keeping only periods starting after July 1, 1998.

The following sections provide more detail on the sources and cleaning of the data.

A.1.1 Car Ownership

The data we use on car ownership comes from the Danish Central Motor Register. This register contains the license plate, vehicle identification number (VIN), and personal identification number (i.e., CPR numbers, which allow us to merge these data in with other public registers). In the raw data, we observe some problematic observations. When we observe a car with a car ownership period for one owner that does not end and a car ownership

Figure 10: Observations by start year

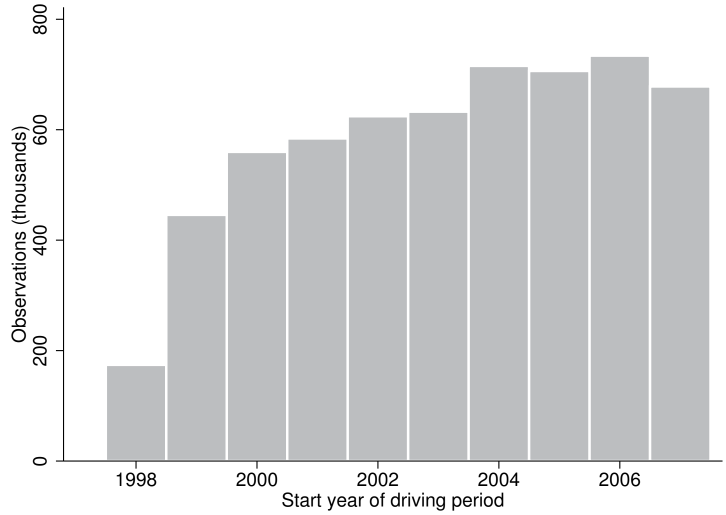

period for a different owner at a later date, we know that the transaction was not properly recorded. In this case, we assign the ending of the ownership period for the first owner at the date when the second owner is first observed with the vehicle. We also do the same for the reverse scenario. We also occasionally see problematic observations where there is an overlap of owners. In that case, we have no way of discerning which person truly owns the car and according to the data documentation such an observation should be impossible so we drop them from the dataset.

A.1.2 Driving Periods

The data on driving periods come from the Ministry of Transportation (MOT) tests that were introduced in 1997. These inspections are mandatory and must be performed at car ages 4, 6, 8, 10, 12, etc. This means that we have two types of driving periods; The *first* driving period is 4 years long (that is, it has *4 years to test*) and any *subsequent* driving period will be only 2 years long. The inspection date is set based on the date of the first registration of the car in Denmark. In practice, the years to test may deviate with plus or minus three months around these designated years. A person may choose to take the car in for inspection *earlier* than the set date if he or she wishes.

MOT tests were originally performed by public authorities directly but in more recent years, they have been performed by private companies approved by the MOT. The goal of the test is to verify that the car is in safe condition for driving on the roads. As a part of the test, the odometer of the car is recorded. A test may have four outcomes; 1) The car can be *approved*. 2) The car can be *conditionally approved*, meaning that certain repairs must be performed for the car to be in legal driving order but that no extra test will be required. 3) The car can be *approved after a re-inspection*, implying that repairs must be made and then the car must return for another test before 33 calendar days. Finally, 4) the car can be declared *not approved* in which case it will be illegal to drive the car and the police will withdraw the license plates. Some drivers may take their vehicle in for an inspection early prior to selling the car in order to give the buyer a signal that the car is in proper working order. Figure 11 shows the distribution of the driving period length. The vertical lines mark the sample selection described above.

Figure 11: Years to Test Distribution

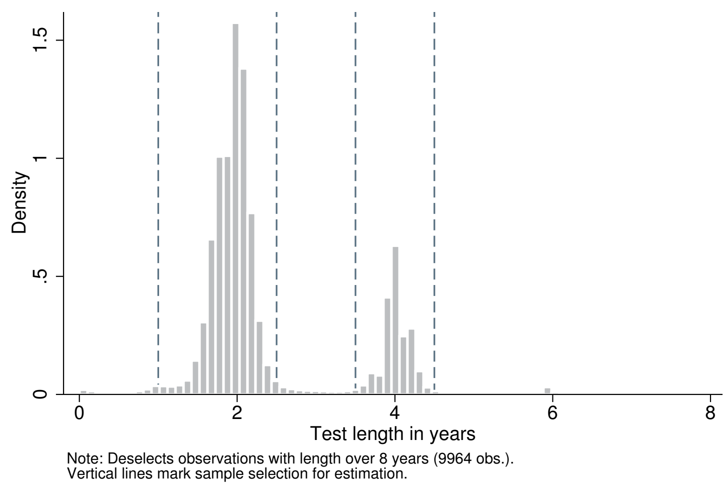

Note: Years to test is the time between two odometer readings. New cars come in for inspection around 4 years and used cars around 2 years.

The sample selection criteria mentioned above for the timing of the driving periods can also be seen in Figure 12. We have selected the sample for a period when the years to test is relatively constant, thus helping to alleviate any concerns of sample selection bias based on this variable.

A.2 Detailed Variable Description

Table 7 lists of all the variables used in this paper with details on each.

Table 7: Variables used in the paper

Variable	Description

VKT	Vehicle-kilometers-traveled in km per day. The variable is constructed by taking first-differences of the odometer readings from the dataset with vehicle inspections. For the first inspection we observe for a car, we assume that the odometer was zero at the time of the car's first registration in Denmark. This will be incorrect if the car was imported from abroad. However, then the car must have had a toll inspection, which we observe, so we can run a robustness check on this assumption. We find that this does not impact our results.
Couple	Dummy for there being two members of the household (married or co-habiting, of opposite genders and having at less than 15 years of age difference).
Real gross income	The sum of gross incomes for the member(s) of the household. The variable comes from the income tax registers. The variable includes all government transfers such as pension payments, unemployment benefits, etc.
Real gross income (couples)	As above but equal to zero for singles.
Real gross income (singles)	As above but equal to zero for couples.
WD	Work distance. The variable is based on the Danish deduction for work distance. Any working household having further than 12 km each way to work can deduct a fixed amount per km. Thus, the measure will be equal to zero if the individual lives closer than 12 km from his or her work. Between 12 and 25 km, there is a rate and above 25 km, the rate drops to half. The rate changes over the period. The total deduction is the daily rate times the number of days worked. The variable is self-reported but the tax authorities have access to both the home and work addresses for the individual. The deduction is the rate times the distance times the number of days worked. We do not observe the number of days worked so we assume 225 work days, which corresponds to the number of days in a typical Danish work year. For example, the official number of work days were 224 in 2007, 226 in 2008, 225 in 2009 and 228 in 2010. Most unions follows these, as do most public sector employees. Figure 16 shows the density of the work distance variable. Note that there is a positive mass on the interval $(0; 12)$ km even though the deduction is only given if the door-to-door work distance is above 12 km; this is due to the assumption about 225 work days per year. If an individual works part-time, say 110 days, but has a distance of 20 km to work, then the variable will be equal to 10. The positive mass will therefore contain many part time employees. For validity, we can compare it to the continuous WD measure, available for a subset of the period (see Appendix A.3.3).
WD non-zero	Dummy for the WD measure being observed. Thus, this is essentially a dummy for the individual living further than 12 km from the work place.
WD (door-to-door distance)	This is the door-to-door distance from home to work. The variable comes from the Danish Technical University's Department of Transportation. It is calculated using a shortest-path algorithm and the National Transport model with GIS data on households and their work places. The variable is only observed for households where the work place is observed and not for 1998 or 1999. In total, it is observed for 76.17% of our estimation sample (79.61% of the observations between 2000 and 2008). We use this measure to validate the tax-based WD variable.
# of children	The number of children aged less than 18 years living with the household.
Urban (dummy)	Dummy equal to one if the household lives in either Copenhagen, Frederiksberg, Aarhus, Aalborg og Odense municipalities, which constitute the major Danish urban areas.
Company car	Dummy equal to one if at least one member of the household has paid the tax penalty for having access to a company car. The use of company cars is restricted to avoid making it an alternative to buying your own car privately. The size of the tax depends on the value of the car. We collapse the variable to a dummy for having any car available to any of the members of the household. Individuals may have access to a company car and not pay this tax if the car is a "yellow license plate" car. These cars can have at most two seats and are typically vans used by craftsmen. The police enforce this very strictly and an individual caught using a company car privately and not paying the penalty is fined and some times forced to pay the registration tax.

Self employed	Dummy equal to one if the household has at least one self employed individual. This information comes from the tax registers.
# of periods observed	The number of driving periods observed for the household. Note that the other driving periods may be with different cars and that our sample selects only households with at least two driving periods.
Bus/Train stops per km^2	The number of public transport stops in the municipality in 2013 divided by the area of the municipality of residence at the start of the driving period in km^2. The data for this comes from the Travel Planner (`http://rejseplanen.dk`), which is a search engine for planning trips using public transportation. The data are only available for a cross-section in 2013. The highest number of stops is 79.9 stops per km^2 for Odense municipality and the lowest is Aaskov municipality with 0.3 stops per km^2.
Weight (ton)	The gross weight of the car in metric tonnes. This is the maximum allowed weight of the vehicle including cargo. The variable comes from the vehicle type approval documents.
Diesel	Dummy equal to one if the car uses diesel fuel. Note that the fuel price will then be based on the diesel price.
Van	Dummy equal to one if the vehicle is a van.
Percent owned of period	The fraction of the driving period where the car was owned by this household. That is, if the driving period starts on Jan 1st, 2001 and ends on Jan 1st 2003, but the car changed owner on Jan 1st 2002, this variable will be equal to 0.5 for both the observations of the two households driving the car.
Driving period length	The length of the driving period in years. For new cars, this will be 4 years and for older cars, it will be 2 years, both plus or minus 3 months and with some exceptions. Note that our sample selects on driving periods being either 1.0 to 2.5 years long or 3.5 to 4.5 years long.
Car age	Car age in years at the start of the driving period. Car age is defined as the time since the car's first registration in Denmark since we do not observe the actual production year of the vehicle. This will be very close to the number of years since the model year for most vehicles, but will be off for the small number of imported vehicles.
# cars / vans / motorcycles / mopeds / trailers owned	Continuous measure of the number of vehicles of the given type owned by the household. For example, if for a given household i and driving period t, the household owns another car for the entire duration of the period, then # of cars owned will be 2.0. If that other car is only purchased half-way through the driving period t, then it is equal to 1.5. That is, the variable is equal to the fraction of this driving period overlapping with the ownership of other vehicles.
First driving period	Dummy equal to one if it is the car's first driving period, i.e., the driving period's start date is equal to the first registration date of the car.
Fraction owned	For household i and driving period t, this is the percent of the driving period where houseohld i is the owner. That is, if the car changes owner midway through, there will be an observation in the dataset for each of the two households owning the car and they will both have this variable set to 0.5.
Years to test	The length of the driving period in years (continuous variable). Due to our sample selection, this will be in $[1.0; 2.5]$ or in $[3.5; 4.5]$.
% of each month	This is a set of variables for each month equal to the % of the driving period taking place in each of the 12 months. Thus, if a driving period is precisely 2 or 4 years long, these will all be equal to $\frac{1}{12}$. We omit April as the reference group in regressions since the fractions will always sum to 1.
Year controls	These are variables for each year, 1998, ..., 2011, each equal to the % of the driving period falling in that year. In the preferred specification, we exclude year 2003 as the reference year and include an additional full set of year controls interacted with the diesel dummy to allow a separate time trend for diesels.

A.3 Additional Descriptives

A.3.1 Driving and Demographics

Figure 13 shows the distribution of vehicles kilometers traveled (VKT). The figure is cut at 200 km/day for clarity. Note that there is still positive mass for very low VKT. This may be explained by vehicles such as vintage or specialty cars.

A.3.2 Additional Spatial Descriptives

Figure 14 shows the number of observations (i.e., driving periods) by municipality. The four major urban areas clearly stand out: Copenhagen (east), Odense (center, on the island of Fyn), Aarhus (midway up on the eastern side of Jutland) and Aalborg (Northern part of Jutland).

Figure 15 shows a map of Denmark where municipalities are colored by the average VKT of the households. We see that the households with high driving tend to be in the outskirts of the major urban areas with a few exceptions. Note that this figure plots observations in the estimation sample, so it should be interpreted recognizing that it conditions on households owning a car. Note that the car ownership rate is 40% in the five largest urban municipalities and 67% elsewhere in Denmark, so a map of the per capita driving would show even lower driving in the urban areas relative to rural areas.

A.3.3 Work Distance

In this subsection, we discuss the validity of the work distance variable. Table 8 shows summary statistics for work distances of males, females and singles. It shows both the measure based on the tax deduction for work distance as well as the "door-to-door work distance" variable, which measures the distance using GPS coordinates. The tax deduction is a deduction from taxable income and it is given as a fixed amount per kilometer per day but is equal to zero if the distance is shorter than 12 km. The number of days worked is not observed so we assume that all individuals work 225 days a year, which is very common in Denmark. Hence, if the individual actually worked fewer days, we will be undershooting the measure (which explains why the variable can take values below 12 km) and vice versa. The per km rate varies over time and there is a kink in the schedule at 50 km where it falls to half the rate.[34]

[34] In some years, a small number of *fringe municipalities* (Danish: *udkantskommuner*) also had the full rate after the 50 km threshold.

Figure 12: Years to Test by Start Date of the Driving Period

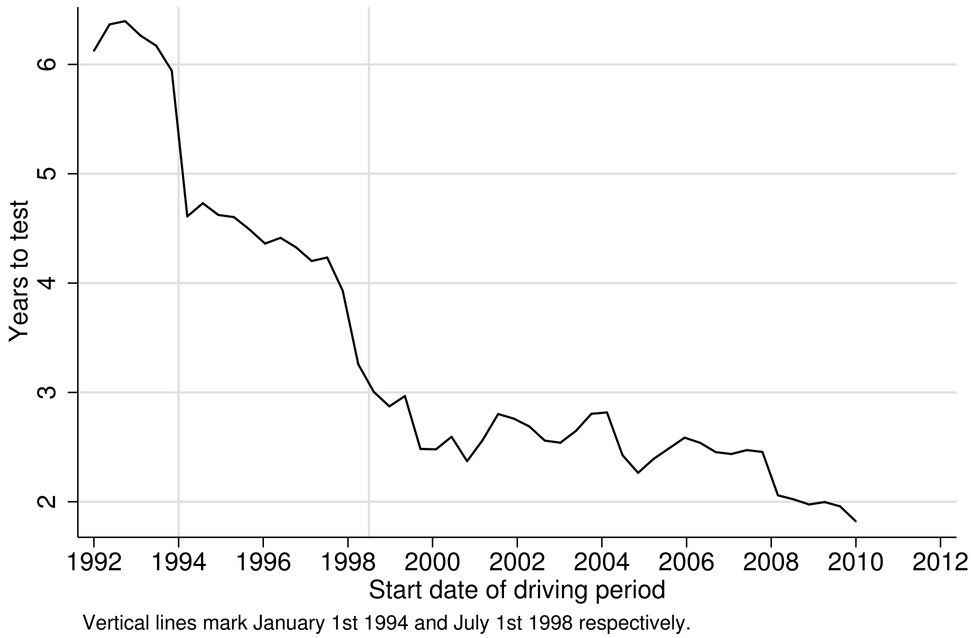

Note: Years to test is the difference between two odometer readings. Since there are more used than new cars, the average is closer to 2 than 4. We start our sampling period in mid 1998 because the average stabilizes here. Prior to that, cars were coming in that had never been to an inspection before and therefore had very long driving periods.

To explore the validity of the work distance variable, we exploit the aforementioned *door-to-door* work distance, which is based on the address of the home and work location. Thus, it directly captures the literal work distance. However, it is not available for the full sample and it is a massive over-estimate for households that work from home or work elsewhere than the primary office of their work. Thus, we see it as a useful robustness check and opt to use the tax return variable in our primary specification.

We compare the distribution of driving according to the two variables to validate the tax return measure. To make the comparison sensible, make the comparison for the subsample where both measures fall in the range [12 km ; 100 km]. The lower bound ensures that the tax-based measure is also observed, while the upper bound makes the graph easier to read. Figure 16 shows the comparison, demonstrating the comparability of the two work distance variables.

Figure 13: The Distribution of Vehicle Kilometers Traveled

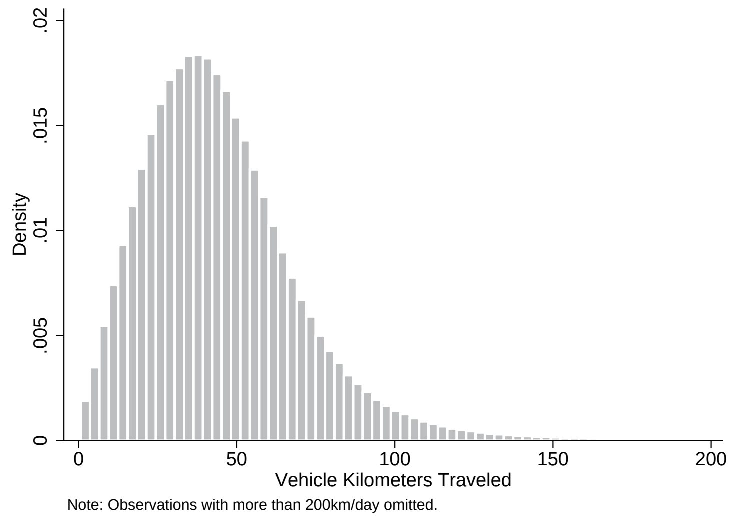

Figure 14: Observations in the Estimation Sample by Municipality

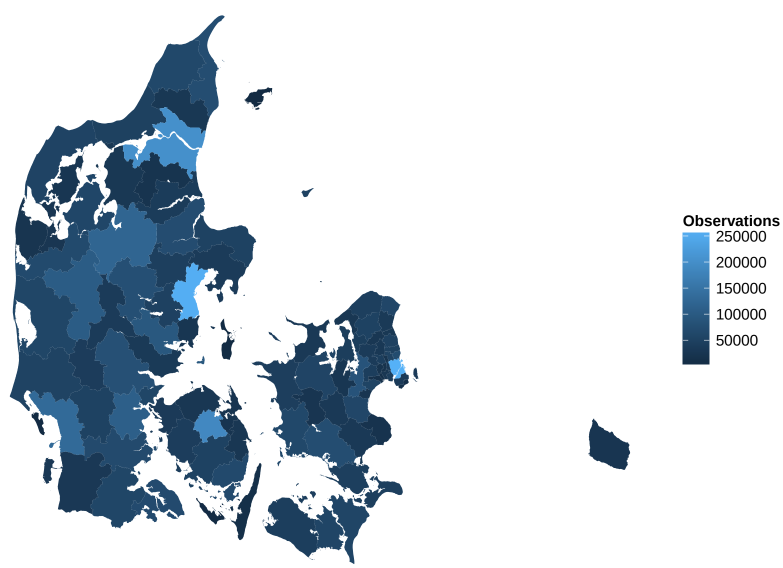

Figure 15: Average Driving by Municipality

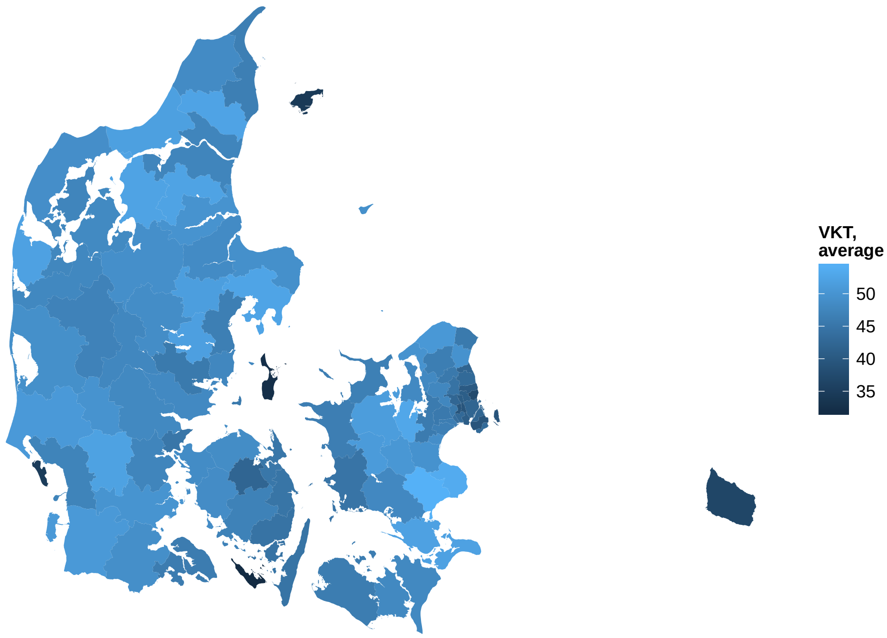

Table 8: Work Distance (WD) Variables

	N	mean	sd	p1	p10	p25	p50	p75	p90	p95	p99
WD, male	4550411	9.5932	18.51	0.0	0.0	0.0	0.0	15.5	32.2	44.7	80.7
WD, female	4550411	6.9385	13.77	0.0	0.0	0.0	0.0	10.7	24.8	33.7	58.3
WD, single	1305035	7.6966	16.87	0.0	0.0	0.0	0.0	9.0	27.5	39.3	75.1
WD non-zero, male	4550411	0.3493	0.48	0	0	0	0	1	1	1	1
WD non-zero, female	4550411	0.3137	0.46	0	0	0	0	1	1	1	1
WD non-zero, single	1305035	0.2917	0.45	0	0	0	0	1	1	1	1
WD, door-to-door, male	3343884	20.3157	34.36	0.0	0.0	2.7	9.8	23.7	46.5	71.8	196.3
WD, door-to-door, female	3094025	14.3657	22.45	0.0	0.6	2.8	8.1	18.1	32.1	45.0	99.3
WD, door-to-door, single	813453	18.6009	32.87	0.0	0.1	2.6	8.6	21.1	42.0	66.1	183.4

Note: WD refers to the work distance variable based on the travel tax deduction, which is censored at 12 km but contains information on the number of days commuted. "WD door-to-door" refers to the shortest path measure from home to work. The two measures should only be expected to be equal if the person has longer than 12 km to work and works precisely 225 days each year.

Figure 16: Comparing the Two Work Distance Measures

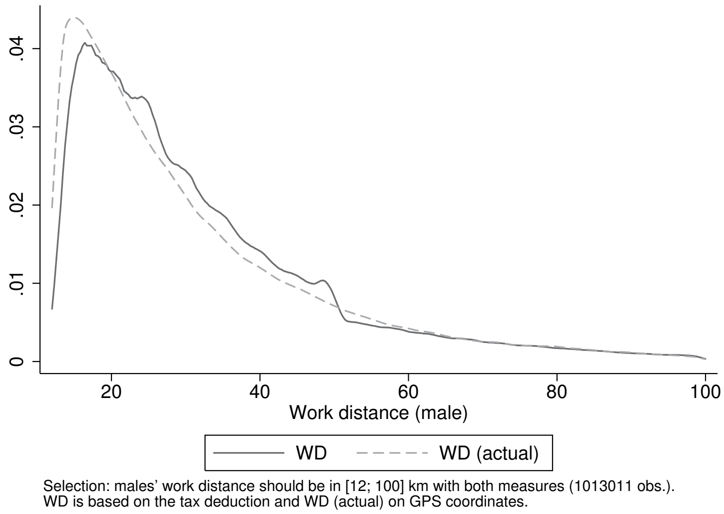

Note: Both curves are non-parametric kernel density estimates for the work distance for households where it is greater than 12 km for both measures. The tax-based measure features a few notable excess-mass points, which is most likely due to individuals rounding off if in doubt.

B Additional Regression Results

This appendix contains number of econometric results supplementing the primary results from section 4. To begin, Table 9 shows the coefficients pertaining to car characteristics and the driving period that were suppressed in the primary results table in our paper.

To further explore heterogeneity, Table 10 shows the coefficients for the demographic variables for the quantiles 1, 50 and 99 in the panel quantile regression estimates. They show that many of the coefficients do not vary over the conditional distribution of VKT. However, the fuel price elasticity, work distance, company car dummy, and transit stop density variables change.

Table 9: Main results — Car and Period Controls

	OLS		Household FE	
	(1)	(2)	(3)	(4)
	No demo	Base	FE	Main
$\log p^{\text{fuel}}$	-0.866***	-0.298***	-0.515***	-0.304***
	(0.00509)	(0.0143)	(0.00722)	(0.0154)
New car	-0.00350*	0.0128***	0.00838***	0.0394***
	(0.00148)	(0.00148)	(0.00160)	(0.00164)
Percent owned of period	-0.189***	-0.112***	-0.0537***	-0.0154***
	(0.000826)	(0.000862)	(0.00106)	(0.00110)
Driving period length	-0.0507***	-0.0541***	-0.0465***	-0.0242***
	(0.000634)	(0.000645)	(0.000681)	(0.000725)
Weight (ton)	0.00214***	0.00169***	0.00166***	0.00167***
	(0.00000523)	(0.00000506)	(0.00000799)	(0.00000798)
Weight squared	-0.000000471***	-0.000000369***	-0.000000354***	-0.000000354***
	(1.35e-09)	(1.30e-09)	(2.00e-09)	(2.00e-09)
Diesel	0.316***	0.311***	0.228***	0.259***
	(0.000918)	(0.00557)	(0.00139)	(0.00545)
Van	-0.236***	-0.199***	-0.204***	-0.205***
	(0.00117)	(0.00115)	(0.00171)	(0.00170)
Car age	-0.0302***	-0.0275***	-0.0284***	-0.0293***
	(0.0000932)	(0.0000911)	(0.000140)	(0.000141)
# cars owned	0.0482***	-0.0202***	-0.0581***	-0.0501***
	(0.000593)	(0.000759)	(0.00114)	(0.00109)
# vans owned	0.0111***	-0.0470***	-0.0711***	-0.0654***
	(0.00124)	(0.00122)	(0.00183)	(0.00179)
# motorcycles owned	0.0319***	-0.00420***	0.0102***	0.0118***
	(0.00101)	(0.000905)	(0.00178)	(0.00178)
# mopeds owned	0.136***	0.0415***	0.0232***	0.0204***
	(0.00138)	(0.00131)	(0.00218)	(0.00217)
# trailers owned	0.0123***	0.0258***	0.00334**	0.00595***
	(0.000519)	(0.000983)	(0.00106)	(0.00106)
Year controls	No	Yes	No	Yes
Household FE	No	No	Yes	Yes
R^2	0.20	0.34	0.68	0.68
N	5,855,446	5,855,446	5,855,446	5,855,446

Dependent variable is the log VKT. An observation is a driving period. All specifications have all of the other variables and controls in Table 9. Robust standard errors clustered at the household level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 10: Panel Quantile Regression for P01, P50 and P99: Demographics

	(1) Linear	(2) P01	(3) P50	(4) P99
$\log p^{\text{fuel}}$	-0.304***	-0.559***	-0.233***	-0.609***
	(0.0154)	(0.0663)	(0.00739)	(0.0592)
Work Distance (WD) controls				
WD, male	0.00242***	0.00137***	0.00260***	0.00347***
	(0.0000336)	(0.000106)	(0.0000118)	(0.0000947)
WD non-zero, male	0.0329***	0.0797***	0.0320***	-0.00688
	(0.00107)	(0.00427)	(0.000476)	(0.00382)
WD, female	0.00303***	0.00245***	0.00328***	0.00338***
	(0.0000443)	(0.000151)	(0.0000168)	(0.000135)
WD non-zero, female	0.0257***	0.0950***	0.0247***	-0.0327***
	(0.00111)	(0.00461)	(0.000513)	(0.00412)
WD, single	0.00419***	0.00360***	0.00448***	0.00562***
	(0.0000835)	(0.000216)	(0.0000241)	(0.000193)
WD non-zero, single	0.0724***	0.138***	0.0713***	-0.0174*
	(0.00243)	(0.00832)	(0.000927)	(0.00743)
Age controls				
Age, male	0.0212**	0.0224***	0.0213***	0.0199***
	(0.00813)	(0.00133)	(0.000148)	(0.00119)
Age, female	0.0468***	0.0534***	0.0469***	0.0403***
	(0.00813)	(0.00132)	(0.000148)	(0.00118)
Age, single	0.0598***	0.0631***	0.0604***	0.0549***
	(0.000939)	(0.000971)	(0.000108)	(0.000868)
Age squared, male	-0.0000930***	-0.000118***	-0.0000943***	-0.0000705***
	(0.0000112)	(0.0000128)	(0.00000143)	(0.0000115)
Age squared, female	-0.000195***	-0.000275***	-0.000197***	-0.000117***
	(0.0000115)	(0.0000134)	(0.00000149)	(0.0000120)
Age squared, single	-0.000206***	-0.000275***	-0.000213***	-0.000119***
	(0.00000767)	(0.00000933)	(0.00000104)	(0.00000834)
Other demographic controls				
log gross inc (couple)	-0.0242***	-0.0156***	-0.0176***	-0.0278***
	(0.00162)	(0.00304)	(0.000339)	(0.00272)
log gross inc (single)	0.0200***	0.0268***	0.0195***	0.0144***
	(0.00288)	(0.00226)	(0.000252)	(0.00202)
Urban (dummy)	-0.0249***	-0.0392***	-0.0254***	-0.0146**
	(0.00284)	(0.00519)	(0.000578)	(0.00463)
# of kids	-0.0168***	-0.0145***	-0.0169***	-0.0145***
	(0.000650)	(0.00146)	(0.000163)	(0.00130)
Company car	-0.0977***	-0.312***	-0.102***	0.0601***
	(0.00216)	(0.00695)	(0.000775)	(0.00621)
Self employed	0.000712	-0.0818***	0.00334***	0.0694***
	(0.00136)	(0.00426)	(0.000474)	(0.00380)
Bus/Train stops per km^2	0.0000419	-0.0000421	0.0000173	0.000300**
	(0.0000548)	(0.000103)	(0.0000114)	(0.0000916)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Linear Fixed Effects (FE)	Yes	No	No	No
Canay (2011) FE	No	Yes	Yes	Yes
N	5855446	5855446	5855446	5855446

Standard errors in parentheses. FE are at the household level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C Robustness Checks

C.1 Stratifying on Time

Tables 11 and 12 show the implications for the estimated fuel price elasticity of dropping certain years from the sample. These results demonstrate considerable robustness.

Table 11: Robustness: dropping earlier years

	(1) Full	(2) 1999-	(3) 2000-	(4) 2001-
$\log p^{\text{fuel}}$	-0.301***	-0.325***	-0.381***	-0.398***
	(0.0169)	(0.0186)	(0.0198)	(0.0219)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00830***	-0.00834***	-0.00855***
	(0.000386)	(0.000387)	(0.000392)	(0.000410)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.351***	0.345***	0.340***
	(0.0156)	(0.0155)	(0.0159)	(0.0161)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE		Yes	Yes	Yes
N	5855446	5681226	5235440	4675560
R^2	0.685	0.691	0.707	0.725

Note: In each column (2)–(4), data before year 97, 98, 99 are dropped respectively.
Standard errors clustered at the municipality level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Robustness: dropping later years

	(1) Full	(2) -2006	(3) -2005	(4) -2004
$\log p^{\text{fuel}}$	-0.301***	-0.255***	-0.307***	-0.278***
	(0.0169)	(0.0186)	(0.0191)	(0.0198)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00801***	-0.00818***	-0.00873***
	(0.000386)	(0.000403)	(0.000408)	(0.000506)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.335***	0.348***	0.347***
	(0.0156)	(0.0160)	(0.0169)	(0.0194)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE		Yes	Yes	Yes
N	5855446	5177147	4443035	3736630
R^2	0.685	0.697	0.712	0.727

Note: In each column (2)–(4), data after year 06, 05, 04 are dropped respectively.
Standard errors clustered at the municipality level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.2 Stratifying on Couples or Singles

Table 13 shows the results when estimating on the sample consisting exclusively of couples or singles, again demonstrating considerable robustness.

Table 13: Robustness: dropping couples or singles

	(1) Base	(2) Only couples	(3) Only singles
$\log p^{\text{fuel}}$	-0.301***	-0.317***	-0.192***
	(0.0169)	(0.0208)	(0.0312)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00910***	-0.000267
	(0.000386)	(0.000395)	(0.000776)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.330***	0.403***
	(0.0156)	(0.0165)	(0.0268)
Year controls	Yes	Yes	Yes
% of each month	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE		Yes	Yes
R^2	0.685	0.645	0.759
N	5855446	4550410	1305036

Note: columns (2) and (3) contain only couples or singles respectively.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.3 Stratifying on the Length of the Period

In our primary specification, we include as control variables both the length of the driving period as well as a dummy for whether it is the first driving period for the car. Since our outcome is the average daily driving, there should not be a mechanical relationship so this robustness check is just to confirm that there is not an issue. Nevertheless, we have included driving periods that are longer or shorter than expected and we now turn to examining robustness with respect to these. In table 14, we drop the driving periods that have years to test (length of the driving period) more than 3 months away from either 2 or 4 years. Recall that a normal test period will be 4 years for a new car and 2 years for a used car. However, during the phase-in of the inspections, cars were summoned for inspection for the first time and therefore did not necessarily drive the normal length early on. The results show that when we remove these driving periods with non-standard length we find a numerically lower elasticity of -0.275. In column (2), we include a dummy to control for the non-standard length, but this does not change the fuel price elasticity much at all (-0.304). We have also experimented with using the length of the driving period as an inverse probability weight as a robustness check. This results in a slightly higher mean elasticity, which is also what we find for newer cars — in that sense, it is consistent with the estimates applying more weight to the newer driving observations.

We have also experimented with a regression where we assign each observation a weight proportional to the length of the driving period. We found that this raises the elasticity a little, moving it towards the higher elasticity we find when we estimate on the subsample of households holding newer cars.

Table 14: Robustness: length of the driving period

	(1) Base	(2) Dummy	(3) Subsample
$\log p^{\text{fuel}}$	-0.301***	-0.296***	-0.268***
	(0.0169)	(0.0172)	(0.0187)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00807***	-0.00906***
	(0.000386)	(0.000386)	(0.000412)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.344***	0.300***
	(0.0156)	(0.0157)	(0.0151)
Non-standard test length		-0.00274**	
		(0.000834)	
Year controls	Yes	Yes	Yes
% of each month	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE	Yes	Yes	Yes
R^2	0.685	0.685	0.726
N	5855446	5855446	4535353

Note: Standard test length: years to test is ± 3 months from either 2 or 4 years.
Elsewhere, sample selection requires VKT in [1;2.5] or [3.5;4.5] years.
Standard errors clustered at the municipality level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.4 Year and Seasonality Controls

Table 15 shows the results when we change the way we control for time effects in decreasing complexity over the columns. The results show that even if we simplify down to a specification with only a linear time trend, our mean elasticity is nearly unchanged. However, if we remove time controls entirely, the elasticity changes substantially.

Table 15: Robustness: year controls

	(1) Baseline	(2) Common year FE	(3) No month	(4) Linear	(5) None
$\log p^{\text{fuel}}$	-0.301***	-0.296***	-0.290***	-0.307***	-0.508***
	(0.0169)	(0.0159)	(0.0157)	(0.0101)	(0.0117)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00810***	-0.00809***	-0.00832***	-0.00832***
	(0.000386)	(0.000380)	(0.000380)	(0.000384)	(0.000385)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.344***	0.345***	0.350***	0.344***
	(0.0156)	(0.0157)	(0.0157)	(0.0157)	(0.0157)
Linear time trend				-0.0409***	
				(0.000775)	
Year controls (gas)	Yes	Yes	Yes	No	No
Year controls (diesel)	Yes	No	No	No	No
% of each month	Yes	Yes	No	No	No
Car	Yes	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes
N	5855446	5855446	5855446	5855446	5855446
R^2	0.685	0.685	0.685	0.684	0.683

The columns sequentially reduce the flexibility of time controls.
In (2), there is only one set of year dummies (dropping the fuel-type specific ones);
in (3), month controls are dropped; in (4), the year dummies are swapped for a linear time trend;
in (5), all time controls are dropped. Standard errors clustered at the municipality level.
Standard errors clustered at the municipality level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In Table 16, we change the main specification to use the *number* of months covered by the driving period rather than the *fraction* of each month covered (as we use in the main specification). Our mean elasticity is almost unchanged (from –0.373 to –0.372).

Table 16: Robustness: month controls

	(1) Fraction	(2) Sum
$\log p^{\text{fuel}}$	-0.301***	-0.301***
	(0.0169)	(0.0169)
Feb	-0.155***	-0.00372***
	(0.0432)	(0.00108)
Mar	-0.0950*	-0.000295
	(0.0458)	(0.000746)
May	-0.0301	0.00135
	(0.0523)	(0.000842)
Jun	0.0513	0.00332***
	(0.0348)	(0.000871)
Jul	0.233***	0.00783***
	(0.0428)	(0.000963)
Aug	-0.0436	0.00105
	(0.0453)	(0.000984)
Sep	0.00683	0.00241**
	(0.0429)	(0.000804)
Oct	-0.0517	0.000977
	(0.0410)	(0.000881)
Nov	-0.144***	-0.00204*
	(0.0397)	(0.000867)
Dec	-0.163***	-0.00224*
	(0.0458)	(0.00102)
Apr		0.00187*
		(0.000829)
Year controls	Yes	Yes
Car	Yes	Yes
Period	Yes	Yes
Demographics	Yes	Yes
Household FE		Yes
N	5855446	5855446
R^2	0.685	0.685

(1): The share of the driving period falling in each month.

(2): The number of months covered by the driving period.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.5 Fuel Type

Table 17 explores heterogeneity in the fuel price elasticity by the fuel type of the car. Note that when we have household fixed effects, removing one or more rows will drop households entirely if they end up with one or zero remaining periods. Thus, we are removing some of the "switchers" who have responded on the extensive margin of choosing a different vehicle, which we do not model separately in this paper. We showed in a separate robustness check that this sample selection does not appreciably change the results, but it should be kept in mind in interpreting these results.

We see that allowing the elasticity to vary by fuel type results in a lower (in absolute value) mean estimate (–0.257), while the positive coefficient on the interaction of the diesel dummy and the log fuel price implies a higher elasticity for the diesel drivers (–0.392). Estimating only on the subsamples of each fuel type confirms these results, yielding a lower elasticity for gasoline drivers (–0.268) and a higher for diesel drivers (–0.541). Note that diesel cars generally cost more up-front but are cheaper to use due to a higher fuel efficiency and a lower price per litre of fuel (see e.g. Munk-Nielsen, 2015). Therefore, it is perhaps not surprising that the diesel sample appears to be more price responsive. Note also that the diesel sample is much smaller than the gasoline sample.

Table 17: Robustness: elasticity by fuel type

	(1) Base	(2) Interaction	(3) Gas only	(4) Diesel only
$\log p^{\text{fuel}}$	-0.301***	-0.260***	-0.268***	-0.552***
	(0.0169)	(0.0198)	(0.0186)	(0.0274)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00803***	-0.00383***	0.000709
	(0.000386)	(0.000386)	(0.000566)	(0.000422)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.345***	0.177***	0.110***
	(0.0156)	(0.0157)	(0.0199)	(0.0208)
Diesel=1 $\times \log p^{\text{fuel}}$		-0.120***		
		(0.0277)		
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
R^2	0.685	0.685	0.693	0.752
N	5855446	5855446	5018019	837427

In columns 3 and 4, we include only gas or diesel cars respectively.
Therefore, the diesel-specific year controls are dropped.
For column (4), note the substantial drop in the number of observations.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.6 Instrumental Variable Estimation

Here we present results from instrumenting for the fuel price. Our primary instrument is the WTI crude oil price in USD per barrel. The price is converted to DKK using the exchange rate from June 18, 2015 and then deflated using the Danish CPI. Figure 17 shows the oil price together with the Danish real fuel prices, illustrating the high correlation.

Figure 17: Danish Fuel Prices and the WTI Oil Price

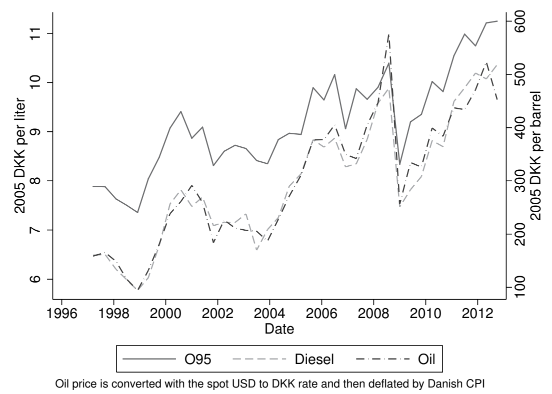

Table 18 shows the main two-stage least squares results, instrumenting log real fuel price with log real WTI oil price.

Table 18: Instrumental Variables Results

	No FE		FE	
	(1)	(2)	(3)	(4)
	OLS	2SLS	FE	2SLS FE
$\log p^{\text{fuel}}$	-0.298***	-0.484***	-0.301***	-0.364***
	(0.0177)	(0.0221)	(0.0169)	(0.0244)
WD $\times \log p^{\text{fuel}}$	-0.00548***	-0.000544	-0.00807***	-0.000336
	(0.000611)	(0.000691)	(0.000386)	(0.000471)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.404***	-0.205***	0.344***	0.121***
	(0.0235)	(0.0248)	(0.0156)	(0.0196)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	No	No		Yes
N	5855446	5855331	5855446	5855296

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 19 shows the first stage results. Note that the very high R^2 of 98% is partially due to the fact that overlapping periods are repeated. These results indicate that the log oil price is a very strong instrument. The F-statistic for both columns is well above 100.

Table 19: Instrumental Variables Results: First Stage

	(1)	(2)
	Simple	Full
	-0.973***	-0.855***
	(0.000299)	(0.000700)
$\log p^{\text{oil}}$	0.176***	0.177***
	(0.0000187)	(0.0000552)
$\mathbb{1}(Dieselcar) \times \log p^{\text{oil}}$	0.147***	0.124***
	(0.0000519)	(0.000129)
All controls	No	Yes
Household FE	No	No
N	5855331	5855331
R^2	0.972	0.982

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.7 Fuel Efficiency and Car Price

In this section, we argue why our estimate of the fuel price elasticity is not biased by our inclusion of most, but not all, vehicle characteristics. First, we show that adding the fuel economy as a control (in the subsample where the variable is observed) does not change the fuel price elasticity.

In Table 20, we show the results of our primary estimation only including fuel economy and car price (manufacturer's suggested retail price, MSRP). One major reason why these variables are not included in the main specifications is that they are only available for a subset of the period. The data source for these variables is the Danish Automobile Dealer Association (DAF). This dataset has been merged to the VINs used by the Motor Register.[35]

The results in Table 20 show how the sample where the characteristics are observed is different from the estimation sample used throughout this paper; switching to this subsample changes the fuel price elasticity from –0.30 to –0.59 (see column (2)). This can be at least partly explained by there being more households with newer cars in the subsample; from the interaction results, we saw that households who have newer cars tend to also be more price sensitive. Including the fuel efficiency variable in column (3) only very slightly changes the elasticity from -0.59 to -0.58. Further including the MSRP in column (4) leaves this almost entirely unchanged (-0.58). We take this as an indication that the included car characteristics are so highly correlated with these variables, that we have little to worry about by excluding them.

[35]The authors gratefully acknowledge Ismir Mulalic at DTU Transport for his assistance with this.

Table 20: Robustness: controlling for fuel efficiency and MSRP

	(1) Base	(2) Subsample	(3) Add km/l	(4) Add MSRP
$\log p^{\text{fuel}}$	-0.301***	-0.602***	-0.593***	-0.594***
	(0.0169)	(0.0220)	(0.0219)	(0.0222)
Fuel efficiency in km/l			-0.00246***	0.00167***
			(0.000508)	(0.000452)
Car price				0.000000477***
				(1.55e-08)
WD $\times \log p^{\text{fuel}}$	-0.00807***	-0.00504***	-0.00503***	-0.00516***
	(0.000386)	(0.000415)	(0.000415)	(0.000408)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.344***	0.437***	0.436***	0.430***
	(0.0156)	(0.0174)	(0.0174)	(0.0169)
Weight (ton)	0.00167***	0.00197***	0.00194***	0.00177***
	(0.0000267)	(0.0000282)	(0.0000274)	(0.0000305)
Weight squared	-0.000000354***	-0.000000385***	-0.000000381***	-0.000000366***
	(5.48e-09)	(5.62e-09)	(5.61e-09)	(6.21e-09)
Diesel	0.257***	0.217***	0.231***	0.198***
	(0.00646)	(0.00971)	(0.0109)	(0.0105)
Van	-0.204***	-0.230***	-0.232***	-0.137***
	(0.00304)	(0.00324)	(0.00339)	(0.00436)
Car age	-0.0293***	-0.0194***	-0.0199***	-0.0176***
	(0.000413)	(0.000298)	(0.000297)	(0.000306)
# cars owned	-0.0498***	-0.0258***	-0.0258***	-0.0270***
	(0.00146)	(0.00168)	(0.00168)	(0.00171)
# vans owned	-0.0653***	-0.0803***	-0.0802***	-0.0841***
	(0.00188)	(0.00241)	(0.00240)	(0.00247)
# motorcycles owned	0.0117***	0.00874***	0.00868***	0.00843***
	(0.00172)	(0.00213)	(0.00212)	(0.00211)
# mopeds owned	0.0210***	0.0167***	0.0167***	0.0168***
	(0.00234)	(0.00261)	(0.00261)	(0.00262)
# trailers owned	0.00592***	0.00870***	0.00866***	0.00876***
	(0.00126)	(0.00135)	(0.00134)	(0.00135)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE		Yes	Yes	Yes
R^2	0.685	0.779	0.779	0.780
N	5855446	3035301	3035301	3035301

(2), (3) and (4) restricts the sample to fuel efficiency and car MSRP being observed.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.8 Location Decisions and Work Distance

In this section, we address robustness with respect to household and firm location decisions. We have data on the home municipality of the household. This allows us to classify households as moving based on whether they ever change municipality. Table 21 shows the key specification estimated on the primary sample of 5.9m households and on the subsample of 3.2m households make a move across municipal borders in column (2). The mean elasticity is larger, changing from -0.30 to -0.36, but all coefficients relating to the work distance and its interaction with the fuel price are not statistically different. This indicates that our finding of a different elasticity for households in the tails of the work distance distribution is not driven by households relocating.

Next, we turn to the firm location decisions. In our data, we can identify the firm a worker is employed at as well as the individual sub-unit ("plant") within the firm. Our data contains information on whether the plant relocates in a given year. This information is created by Statistics Denmark based on the plant locations where individuals work.[36] As one might expect, relocations are not extremely common in our data, but common enough to leave us with 49,074 household-driving-periods to estimate our model on.

In column (3), we estimate on the subsample of 49,074 households where the work location of at least one spouse relocates. The 95% confidence intervals around these parameters all contain the original parameter estimates. In column (4), we estimate only with households where the firm did not relocate and the results are extremely close to the original results. Of course, the spouse where the firm moves may decide to look for a different job in response to advance information about the firm relocating. In columns (5) we only use the 26,803 observations where the firm relocates but where the household does not and again, the original parameter estimates are all contained in the 95% confidence interval. Finally, in column (6), we only estimate on the 3.2m households that moved but where neither spouse worked at a firm-plant that relocated. It almost does not change the results to exclude the households where the firm relocated.

[36] In their raw data, Statistics Denmark observes one address no longer being associated with a firm and a new one being, and they observe the addresses of many of the workers switching from the old address to the new. They require a minimum of the workers from one location appearing at the new in order for it to be classified as a relocation.

Table 21: Robustness: Stratifying on location choices

	(1) Baseline	(2) HH moves	(3) Firm moves	(4) Firm stays	(5) Firm, not HH	(6) HH, not firm
Estimation sample contains observations where:						
Household moves	Yes	Yes	No	Yes	No	Yes
Household never moves	Yes	No	Yes	Yes	Yes	No
Firm moves	Yes	Yes	Yes	No	Yes	No
Firms never moves	Yes	Yes	No	Yes	No	Yes
$\log p^{\text{fuel}}$	-0.356***	-0.418***	-0.374*	-0.357***	-0.457*	-0.416***
	(0.0200)	(0.0223)	(0.151)	(0.0201)	(0.208)	(0.0223)
WD	0.0270***	0.0233***	0.0426***	0.0269***	0.0509**	0.0231***
	(0.00135)	(0.00209)	(0.0110)	(0.00136)	(0.0155)	(0.00209)
WD non-zero	-0.807***	-0.769***	-0.643*	-0.808***	-0.907*	-0.767***
	(0.0376)	(0.0570)	(0.280)	(0.0374)	(0.435)	(0.0564)
WD $\times \log p^{\text{fuel}}$	-0.00980***	-0.00800***	-0.0172***	-0.00974***	-0.0204**	-0.00794***
	(0.000634)	(0.000964)	(0.00490)	(0.000641)	(0.00710)	(0.000965)
WD non-zero=1 $\times \log p^{\text{fuel}}$	0.374***	0.357***	0.310*	0.375***	0.421*	0.356***
	(0.0172)	(0.0260)	(0.126)	(0.0171)	(0.197)	(0.0257)
Mean elasticity	-0.302	-0.355	-0.414	-0.302	-0.493	-0.352
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
N	5855446	3244793	49074	5806372	26803	3168183

Standard errors clustered at the municipality-level in parentheses.

A household is defined as "moving" if it is observed in two different municipalities. Households are assigned to firms based on their registered primary employer. We match households to the firm they work and to the individual work location within the firm they work (which we refer to as a "plant"). Our data contains information on whether the particular plant relocated in a given year. If the plant relocates and the household still works with the firm in the year of the relocation, we classify that entire household observation as one where the firm moves.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.9 Macro effects: unemployment and alternative fuels

Table 22 includes controls for unemployment. The data is based on uninsurance benefits claims and we use both a dummy for any unemployment during the driving period as well as a measure of the fraction of the driving period covered by unemployment. Each of these controls are specific by gender and single households.

Table 22: Robustness: controlling for unemployment

	Only main effects		With interactions	
	(1)	(2)	(3)	(4)
	Baseline	Add UE	Baseline	Add UE
$\log p^{\text{fuel}}$	-0.301***	-0.298***	-0.301***	-0.299***
	(0.0162)	(0.0162)	(0.0169)	(0.0169)
Unemployment (%), male		0.0261***		0.0271***
		(0.00467)		(0.00469)
Unemployment (%), female		0.00523		0.00566
		(0.00292)		(0.00294)
Unemployment (%), single		-0.00291		-0.00217
		(0.00692)		(0.00698)
Unemployment (dummy), male		0.00570***		0.00575***
		(0.00122)		(0.00122)
Unemployment (dummy), female		-0.00130		-0.00136
		(0.000958)		(0.000958)
Unemployment (dummy), single		0.0141***		0.0139***
		(0.00219)		(0.00218)
WD $\times \log p^{\text{fuel}}$			-0.00807***	-0.00809***
			(0.000386)	(0.000386)
WD non-zero=1 $\times \log p^{\text{fuel}}$			0.344***	0.345***
			(0.0156)	(0.0157)
Constant	0.705***	0.688***	0.741***	0.723***
	(0.0749)	(0.0750)	(0.0771)	(0.0772)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
R^2	0.685	0.685	0.685	0.685
N	5855446	5855446	5855446	5855446

Note: Unemployment is measured in the percent of weeks of the middle year of the driving period.
Standard errors are clustered at the municipality-level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 23 shows the robustness checks where we control for alternative fuel prices. The "matched sample" results in column (2) show the results using exactly the same sample we are able to use for the alternative fuel price regressions. The results in the remaining columns show that the fuel price coefficient barely changes when including alternative fuel prices. And in general, the alternative fuel prices are not statistically significant.

Table 23: Robustness: controlling for alternative fuel prices

	Full sample (1) log_km	Matched sample (2) log_km	(3) log_km	(4) log_km	(5) log_km	(6) log_km
$\log p^{\text{fuel}}$	-0.301***	-0.391***	-0.394***	-0.385***	-0.391***	-0.403***
	(0.0162)	(0.0195)	(0.0260)	(0.0170)	(0.0172)	(0.0214)
log natural gas price			0.000730			
			(0.00808)			
log coal price (Australia)				-0.00255		
				(0.00397)		
log coal price (S. Africa)					-0.0000288	
					(0.00425)	
log coal price (Norway)						-0.00566*
						(0.00228)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	5855446	5190243	5190243	5190243	5190243	5190243

The natural gas price comes from Statistics Denmark. The coal price (Norway) is the US export price to Norway, available from the U.S. Energy Information Administration. The coal prices for Australia and South Africa come from the World Bank.
Standard errors clustered at the municipality level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.10 Alternative Specifications and the Tail

In this section, we address potential concerns relating to:

- The work distance variable,
- The linear functional form for the work distance.

First, one might be concerned that the lower tail we uncover is due to our work distance measure being censored at 12 km. To address this, we estimate our model on the subset where we have access to the door-to-door measure of work distance. As mentioned elsewhere, we find this measure inferior despite not suffering from censoring because our preferred tax-based measure also captures the number of days of commuting.

Figure 18: Price elasticity and door-to-door work distance

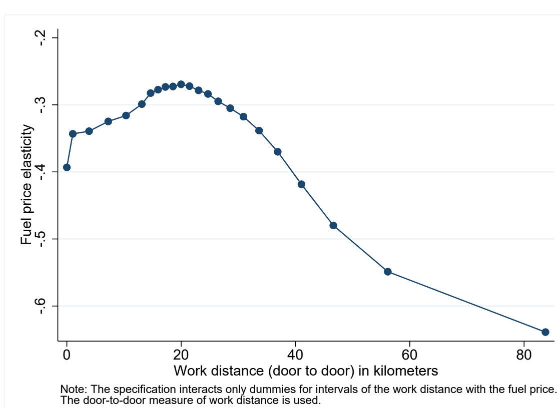

Second, we turn to the functional form for how the work distance enters. Omitting the other controls, our specification takes the form

$$\log \mathrm{VKT}_{it} = (\gamma_0 + \gamma_1 \mathrm{WD}_{it} + \gamma_2 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}}) \log p_{it}^{\mathrm{fuel}} + \delta_0 + \delta_1 \mathrm{WD}_{it} + \delta_2 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}} + \cdots$$

One may worry that our linear specification in work distance does not capture the true relationship in the elasticity, or find it unnatural to have log in price and quantity but not work distance. However, we find this to be the most natural specification because the work distance–in contrast to the fuel price and driving–takes the value of zero often. Nevertheless, to satisfy the curious reader, Table 19 shows that the finding of a U-shape in the elasticity over work distance also arises from such a specification. We are presenting the results from the following regression:

$$\log \mathrm{VKT}_{it} = (\gamma_0 + \gamma_1 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}} \log \mathrm{WD}_{it} + \gamma_2 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}}) \log p_{it}^{\mathrm{fuel}} + \delta_0 + \delta_1 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}} \log \mathrm{WD}_{it} + \delta_2 \mathbf{1}_{\{\mathrm{WD}_{it}>0\}} + \cdots$$

Figure 19 shows that the picture is qualitatively exactly the same as in the primary specification, shown in Figure 6, although the functional form of the relationship now naturally displays slight additional curvature due to logarithic form in work distance.

Figure 19: Price elasticity work distance measured in logs

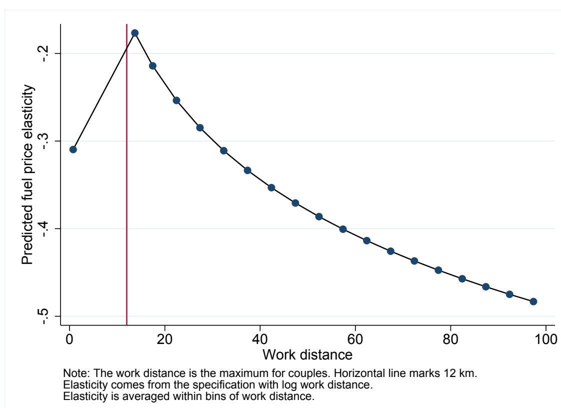

D A Simple Model of Travel Decisions

This section develops a simple model of the travel decision of a car-owning agent in order to build intuition for the economics underlying our empirical results. The focus of this model is on the economics of the short-run price responsiveness of driving and how it varies with the work distance of the household. For clarity of exposition, we abstract from other decisions that may influence driving in the long-run, such as where to live and what employment to accept. Our model is well-suited for a setting where the decision-maker has access to public transport. Such a setting is relevant to nearly all of Denmark, as well as much of Europe and many other areas in the world. For example, in 2014, 87% of Danes live within one kilometer of a public transport stop and nearly all the remainder are served by on-call buses ("telebusser").[37] We model a static setting for a given finite amount of time, such as one week.

Let the total number of kilometers traveled by the agent be given by T. The agent can travel by personal vehicle or by other modes of transport, including public transport, biking, or walking. Let the kilometers traveled by personal vehicle be denoted by v, so the remaining kilometers traveled is $T - v$. Consider two types of travel. The first type is repeated travel that occurs several times a week, such as for a commute to work. The second is discretionary, shopping, or leisure travel. Let $d^w \in [0, 1]$ be the fraction of commuting (work) trips driven. $d^w = 1$ if all of commuting is accomplished by driving and $d^w = 0$ if all of commuting is done by other modes of transport. Similarly, let $d^l \in [0, 1]$ be the fraction for non-commuting (leisure) trips.

Let w be the kilometers traveled for the commute and $g^w(d^w, w)$ be the additional utility from commuting to work by driving rather than other forms of transport. Similarly, let l be the kilometers traveled for non-work trips and $g^l(d^l, l)$ be the utility from driving for non-work trips. In the short-run, w is not a choice variable, so $g^w(d^w, w) \equiv g^w(d^w)$. In contrast, l is a choice variable. As driving is a more flexible form of transport, assume $\frac{\partial g^l(d^w)}{\partial d^w} > 0$ and $\frac{\partial g^l(d^l)}{\partial d^l} > 0$. However, there is an important difference between commuting trips and other trips that motivates our specification of these functions. While trips for shopping or leisure involve travel to a diverse set of locations, commute trips are very homogenous, from the same origin to the same destination and usually at the same time of day. Thus, for a given set of commute trips in a given time period, we expect the marginal utility from commuting by personal car to be constant, regardless of the amount of driving. This allows us to define

[37] See `http://passagerpulsen.taenk.dk/file/68/download?token=fy19yEeh`, Accessed June 16, 2015.

$g^w(d^w) \equiv \gamma^w d^w$, where γ^w is a constant.[38] In contrast, there is inherent heterogeneity in the ability to bike, walk, or take public transport for non-commute trips. For some shopping or leisure trips, public transport or biking are very attractive modes of travel; for others, they are highly unappealing due to the distance or destination.

Consider an agent who maximizes utility subject to a budget constraint:

$$\max_{d^w \in [0,1], d^l \in [0,1], l} u(x) + g^w(d^w) + g^l(d^l, l)$$
$$\text{s.t.} \quad y \geq p^v v + p^b(T - v) + x,$$

where x is the outside good (whose price is normalized to 1), y is total income, p^v is the price per kilometer of driving a personal vehicle (hence "v"), and p^b is the price per kilometer of the non-driving mode such as a bus (hence "b").

Inserting the assumed form of g^w, the Lagrangian for this problem can be written as

$$\max_{d^w \in [0,1], d^l \in [0,1], l} u(x) + \gamma^w d^w + g^l(d^l, l) + \lambda \left[y - (p^v - p^b)v - p^b T - x \right],$$

where λ is the shadow price or marginal utility of income.

Assuming standard regularity conditions and using $v = d^w w + d^l l$, the optimal d^l can be characterized by the following first-order condition:

$$\frac{\partial g(d^l, l)}{\partial d^l} = \lambda(p^v - p^b)l.$$

This condition indicates that the household will choose the fraction of non-commute driving, $d^l \in [0, 1]$, so that the marginal utility of an additional kilometer traveled by car is equal to the marginal cost (converted to be in terms of utility). Because shopping and leisure trips are heterogenous, the household will shift the least inconvenient trips to public transport, walking, or biking when fuel prices increase. Of course, corner solutions at 0 and 1 are possible if the marginal cost is sufficiently high or low. Otherwise, $\frac{\partial^2 g^l(d^l,l)}{\partial (d^l)^2} \neq 0$ and the monotonicity of $g^l(\cdot)$ assures an interior solution, as one would expect. Similarly, the first order condition for l is $\frac{\partial g^l(d^l,l)}{\partial l} = \lambda \left[(p^v - p^b)(d^l - d^w) - p^b \right]$, which indicates that the marginal benefits of additional non-commute driving are equal to the marginal costs.

The setting is different for commuting, as $\frac{\partial g^w(d^w)}{\partial d^w} = \gamma^w$. Given this, as long as we do not have exact indifference (i.e., $\gamma^w = \lambda(p^v - p^b)w$), a utility-maximizing household would never choose an interior solution. Instead, we obtain the following discrete solution for the choice

[38] This follows from the short-run nature of the decision. In the longer-run, the utility from the work commute may be increasing in w, as it may allow for the household to buy a better house or get a better job. But in the short-run, the house and job are fixed and thus the utility from commuting is fixed.

of mode for commute travel:

$$d^w = \begin{cases} 1 & \text{if } \gamma^w \geq \lambda(p^v - p^b)w \\ 0 & \text{else.} \end{cases} \tag{2}$$

If the marginal utility from driving is greater than marginal cost (converted to be in terms of utility), then $d^w = 1$ and all commute trips in this setting are done by driving. Otherwise, all commute trips are taken using other forms of transport, such as public transport, cycling, or walking.[39] We can think of γ^w intuitively as a type of *switching cost* that prevents a change in commute driving unless there is a sufficiently large change in the marginal cost.[40] It can be thought of as the marginal utility of driving instead of using other forms of transport, and it includes such factors as the effort in planning transport trips or the psychological cost of changing habits.

This framework has important implications for our empirical setting. We are not interested in d^w, d^l, l per se, but we are interested in the overall fuel price sensitivity of driving ($\frac{\partial v}{\partial p^v}$) and the heterogeneity in this sensitivity. Differentiating, we have $\frac{\partial v}{\partial p^v} = \frac{\partial d^w}{\partial p^v} w + \frac{\partial d^l}{\partial p^v} l + \frac{\partial l}{\partial p^v} d^l$. From the first order conditions and the implicit function theorem we also know that at the optimal values of l and d^l, $\frac{\partial d^l}{\partial p^v} = \frac{\lambda l}{\frac{\partial^2 g^l}{\partial (d^l)^2}}$ and $\frac{\partial l}{\partial p^v} = \frac{\lambda(d^l - d^w)}{\frac{\partial^2 g^l}{\partial (l)^2}}$.

For commute driving, the discontinuity in the optimal mode choice implies a discontinuity in the response so that the derivative is zero (almost) everywhere. We thus consider a change in gasoline prices leading to a change from p^v_0 to p^v_1. Consumers will switch from driving to other modes of transport at the threshold $p^v = p^b + \frac{\gamma^w}{\lambda w}$. So the change in commute driving with the given change in fuel prices is

$$\Delta d^w = \begin{cases} 1 & \text{if } d^w = 0 \text{ and } p^v_1 < p^b + \frac{\gamma^w}{\lambda w} < p^v_0, \\ -1 & \text{if } d^w = 1 \text{ and } p^v_1 > p^b + \frac{\gamma^w}{\lambda w} > p^v_0, \\ 0 & \text{otherwise.} \end{cases}$$

This expression highlights when substitution might occur with a fuel price rise. For example, in order for there to be a switch away from driving for commutes, the increase in the marginal cost of driving must be sufficient to overcome the marginal cost of the other option p^b plus the marginal utility of driving above other sources, scaled by the distance of the commute

[39] Note this corner solution holds *in any given empirical setting.* If the setting changes (e.g., there is very nice weather on a given day), we still have a corner solution in the new setting. The key point is that there is a switching cost influenced by the relative prices.

[40] Note that this is a static, rather than dynamic model, so our preferred interpretation of γ^w is as the threshold level of savings required for a substitution rather than as a classic switching cost in a dynamic model.

and put in monetary terms. Thus, for households with very long commutes (i.e., a large w), a fuel price change sufficiently large to induce a switch in modes would imply a more substantial decrease in driving, leading to our first hypothesis:[41]

Hypothesis 1. (Upper Tail) *With sufficient gasoline price variation, households with a longer work distances are more responsive to changes in gasoline prices than the average household.*

This first tail has a clear intuition in that households with long work distances spend proportionately more on fuel for commuting when fuel prices increase, so there would be a strong incentive to reduce driving. However, there must be other available modes of travel for a change in driving to be possible–such as public transport.

For drivers with the who commute the least, there is a very different decision problem. Most of these drivers live in the city or right in town (some may of course work from home in more rural areas, but this is likely to be a small percentage of drivers). For those who live in an urbanized area, walking and public transport would be particularly attractive modes. There is no discrete switching behavior, but rather a continuous substitution from driving due to a diverse array of non-work trips. With more attractive other options, the marginal trip for these drivers would be expected to be more likely to switch away from driving when fuel prices rise.

This can also be seen more formally using $\frac{\partial v}{\partial p^v} = \frac{\partial d^w}{\partial p^v} w + \frac{\partial d^l}{\partial p^v} l + \frac{\partial l}{\partial p^v} d^l$. For households with a negligible w, but a diverse array of attractive mode choices for a diverse array of non-work trips, we would expect the optimal l and d^l to be larger than similar households with a larger w (recall that these are households that own a car despite a short w). This would follow naturally from utility maximization under a budget constraint. Further, assume λ is constant or larger for those with very small w, which would follow if households with greater w are at least as wealthy as those with small w (e.g., households in the suburbs may be wealthier than those in the city). From the implicit function theorem, $\frac{\partial d^l}{\partial p^v} = \lambda l / \frac{\partial^2 g^l}{\partial (l)^2}$ and thus with concavity of $g^l(\cdot)$, $\frac{\partial d^l}{\partial p^v} l$ is more negative with a negligible w. Again using the implicit function theorem, the same is true for $\frac{\partial l}{\partial p^v} d^l$. Finally, we would expect households with very short work distances to have l and d^l sufficiently larger than those with average w (due to the diversity of trips and ability to use different modes of travel) that the more negative values of $\frac{\partial d^l}{\partial p^v} l + \frac{\partial l}{\partial p^v} d^l$ will outweigh $\frac{\partial d^w}{\partial p^v} w$ being close to zero, implying that $\frac{\partial v}{\partial p^v}$ becomes more negative.

Hypothesis 2. (Lower Tail) *Households that commute very little are more responsive to changes in gasoline prices than the average household.*

[41] This can be seen using $\frac{\partial v}{\partial p^v} = \frac{\partial d^w}{\partial p^v} w + \frac{\partial d^l}{\partial p^v} l + \frac{\partial l}{\partial p^v} d^l$. These drivers have large w and for a large-enough fuel price increase $\Delta d^w = -1$. With this major shift away from work driving, the signs of the changes in l and d^l are ambiguous. Relative to drivers with an average w, l and d^l are also ambiguous. Combining these observations with the large effect of Δd^w suggests that the other terms would be unlikely to fully offset Δd^w.

This second hypothesis stems naturally from our model, but is also consistent with previous empirical evidence from the United States indicating that households residing in cities are some of the most responsive to fuel price changes (Kayser, 2000; Gillingham, 2013, 2014; Gillingham, Jenn, and Azevedo, 2015). This formulation is also more general than households that reside in cities

To be clear, we do not claim that these two hypotheses hold under all conditions (e.g., without adequate public transport, the upper tail is much less likely). However, they hold under reasonable assumptions and provide empirically testable implications for the heterogeneity in the fuel price elasticity of driving in our empirical setting. We also do not claim that this model provides the *only* way that two tails might be generated.[42]

We intentionally developed a simple static model to build intuition for the shorter-run driving decisions. In a dynamic setting, we would expect to see similar substitution behavior, whereby households could "invest" in switching if the discounted savings from doing so outweigh the switching cost.[43] Because the savings equal w times the travel cost differential, households with longer w will switch for smaller changes in the fuel price. Such a mechanism could also be included in our model by allowing γ^w to be heterogeneous and increasing in w. For a longer-run analysis, the work distance could be endogenized, but this is outside the scope of our paper, which follows the literature in focusing on the short-run driving decision.

[42]For example, if one made the assumption that the elasticity increases with fuel demand as a share of income and also assumed that this share increases with distance from the city, one could generate the upper tail. Our data suggest that the share may not vary uniformly with distance from the city, so this would be a poor assumption.

[43]The intuition is similar to the intuition in an (S, s)-model of portfolio choice; for small changes in the fuel prices, most households will stick with their baseline mode choice and avoid paying the switching cost. For larger changes, however, they will be forced to re-optimize.