A TALE OF TWO TAILS:
COMMUTING AND THE FUEL PRICE RESPONSE IN DRIVING

Kenneth Gillingham
Anders Munk-Nielsen

A Tale of Two Tails: Commuting and the Fuel Price Response in Driving
Kenneth Gillingham and Anders Munk-Nielsen
NBER Working Paper No. 22937
December 2016
JEL No. Q4,Q5,R2,R4

ABSTRACT

The consumer price responsiveness of driving demand is central to the welfare consequences of fuel price changes. This study uses rich data covering the entire population of vehicles and consumers in Denmark to find a medium-run price elasticity of driving of -0.30. We uncover an important feature of driving demand: two small groups of much more responsive households that make up the lower and upper tails of the work distance distribution. The first group lives close to work in urban areas. The second group lives outside of major urban areas and has the longest commutes. Access to public transport appears to be the force behind the existence of the tails, enabling the switch away from driving. We find that a fuel price increase of 1 DKK/liter implies an average deadweight loss of 0.66 DKK/liter, but there is considerable heterogeneity and the tails bear a larger share of the loss.

Kenneth Gillingham
School of Forestry and Environmental Studies
Yale University
195 Prospect Street
New Haven, CT 06511
and NBER
kenneth.gillingham@yale.edu

Anders Munk-Nielsen
University of Copenhagen
Department of Economics
Oster Farimagsgade 5, 26.3.01
DK-1353 Copenhagen K, Denmark
superpronker@gmail.com

1 Introduction

Oil prices have historically been highly variable, with Brent crude spot prices ranging over the past decade from \$139/barrel at the peak in June 2008 to below \$50/barrel in January 2015. These large oil price gyrations lead to corresponding changes in refined fuel prices, influencing transport decisions, congestion, and environmental outcomes. Understanding the price elasticity of driving, which underpins the price elasticity of fuel consumption, is therefore of considerable policy interest. 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; Kilian and Murphy, 2014). Furthermore, it forms the basis for measuring the welfare consequences of changes in fuel prices.

This study estimates the price elasticity of driving and provides new insight into the underlying determinants of this elasticity. Using vehicle-level odometer readings matched to individual-level location and demographic information from the Danish registers, we uncover two small 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 very short commutes and the other has the longest commutes. Our mean medium-run (one-year) elasticity estimate of -0.30 is considerably influenced by these two groups of tail households, each of which have an elasticity estimate closer to -0.6. These findings can be rationalized with a model of switching costs incurred when switching from driving to other modes of transport, such as public transport. Danes have almost universal access to public transport and we posit that our results hold in similar settings around the world.

This research contributes to three strands of literature with major policy importance. First, it provides a new point estimate for the fuel price elasticity of driving, which is a dominant component in the modeling of gasoline or diesel demand. 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), largely using aggregate data at the regional or national level.[1] More recently, a handful of 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 (Linn, 2016; Bento et al., 2009; Knittel and Sandler, 2013; Gillingham, 2013, 2014; Munk-Nielsen, 2015; Levin, Lewis, and Wolak, 2014). Most of these recent medium-run elasticity

---

[1]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).

estimates are for drivers in the United States and are in the range of -0.10 to -0.35. In contrast, similar benchmark estimates for Europe tend to show a 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.[2] Similarly, in contemporaneous work, De Borger, Mulalic, and Rouwendal (2016a) focus on two-vehicle households in Denmark and find the medium-run fuel price elasticity of driving to range between -0.32 and -0.45. Our study not only helps to reconcile these differing estimates across countries, but it also sheds light on the mechanisms underlying the differences. In particular, by identifying the tails in the distribution of consumer response and the reason for these tails, we can posit that there are groups of more-responsive households in Europe that simply do not exist in the United States.

Identifying the composition of the tail households contributes to a second vein of literature on the complex relationships between urban form, gasoline prices, and consumer decisions about how much to drive. There is growing evidence that urban form and the spatial structure of labor force demand affect travel choices and commuting behavior (Bento et al., 2005; Grazi, van den Bergh, and van Ommeren, 2008; Brownstone and Golob, 2010). Since at least McFadden (1974), it has been long-recognized that access to public transport is an important mediator of travel choices, with clear environmental implications (e.g., Glaeser and Kahn, 2010). Denmark provides a very useful empirical setting for exploring these issues, for access to public transport is near-universal, yet there is considerable variation in commute distances and the degree of access to appealing substitutes to driving. Our findings are informative for the development of models of household location choice and access to public transport by revealing the detailed spatial relationship between location and driving.

The third strand of the literature to which we contribute is the analysis of environmental tax reforms on the light duty vehicle fleet. Several recent papers focus on vehicle registration tax reforms using discrete vehicle choice models (e.g. D'Haultfœuille, Givord, and Boutin, 2014; Adamou, Clerides, and Zachariadis, 2013; Huse and Lucinda, 2013). Without modeling the endogenous choice of driving, these papers can only calculate a rough estimate of the environmental implications of such policies. Other work incorporates the driving decision, for example in a discrete-continuous framework (Jacobsen, 2013; Gillingham, 2013; Munk-Nielsen, 2015; Grigolon, Reynaert, and Verboven, 2015), in order to evaluate environmental policies focused on vehicles. However, the computational complexity of such models prevents a sufficiently detailed modeling of the driving decision to fully capture the heterogeneity of the response. Our study provides a comprehensive picture of the consumer response on the

---

[2]-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.

intensive margin, which can help inform the choice of salient features to include in discrete-continuous models of vehicle choice and utilization designed to examine policies affecting both margins.

In this paper, we underpin our empirical analysis with a simple theoretical model that provides an economic explanation for the existence the tail households. A key feature in this model is the presence of switching costs incurred when changing transport modes. Consider households with very high work distances. When fuel prices increase, these households stand to gain more from switching to public transport and will therefore respond more strongly than households who do not commute as far. Households with very short commutes face a different decision problem, one in which nearly all driving demand is for non-work trips such as shopping or leisure travel. These households drive to a diverse set of destinations and for some of these destinations public transport or other mode choices are also attractive. This means that when fuel prices rise, there is greater ability to respond by switching from driving to other mode choices. We present empirical evidence consistent with these explanations using both a quantile regression framework and a standard linear framework with a rich set of interactions to explore the determinants of greater price responsiveness.

We illustrate what our results mean for policy through an illustrative counterfactual analysis of a price increase of 1 DKK/liter (l) for both gasoline and diesel fuel (i.e., just over $0.50/gallon). Decomposing the total response in driving, we find that the most-responsive 5% of drivers are responsible for 14.4% of the total reduction in driving. Moreover, we develop a new approach to calculate the deadweight loss from this increase in fuel prices based on our quantile regression results. We find a mean deadweight loss of 0.66 DDK/l, and show that this deadweight loss is highly heterogeneous. In fact, the deadweight loss for both highly responsive tails of households is more than four times greater than for the less responsive households in the middle of the work distance distribution, a result that is new to the literature.

The remainder of this paper is organized as follows. The next section lays out our simple theoretical model to provide a framework for the economics underlying our results. Section 3 describes the rich Danish register data and provides descriptive evidence on the primary features of the data relevant to estimating the driving responsiveness. Section 4 describes our empirical strategy, while section 5 presents the results and a set of robustness checks. Section 6 provides the illustrative policy counterfactual and section 7 concludes.

2 A Simple Model of Travel Decisions

This section develops an 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 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 (km) of a public transport stop and nearly all the remainder are served by on-call buses ("telebusser").[3] We model a static setting for a given finite amount of time, such as one week.

To simplify our setting, we hold the total number of km traveled by the agent fixed at T. The agent can travel by personal vehicle or by other modes of transport, including public transport, biking, or walking. Let the km traveled by personal vehicle be denoted by v, so the remaining km 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 decision of how much to drive for commuting trips. $d^w = 1$ 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 same decision for non-commuting (leisure) trips.

Let $g^w(d^w)$ be the additional utility from commuting to work by driving rather than other forms of transport. Similarly, let $g^l(d^l)$ be the utility from driving for non-work trips. 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 the 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 would expect the marginal utility from commuting by personal car to be constant, regardless of the amount of driving. This allows us to define $g^w(d^w) \equiv \gamma^w d^w$, where γ^w is a constant. In contrast, there is inherent heterogeneity in the ability to bike, walk, or take public transport for any specific 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. Thus, one would expect some curvature of $g^l(d^l)$, i.e., the marginal utility of driving will vary with the fraction of non-work

---

trips driven: $\frac{\partial^2 g^l(d^l)}{\partial (d^l)^2} \neq 0$ (and we might expect that g^l is concave, but it is not necessary to assume this).

Denote the km traveled for the commute by w and the km traveled for non-commute trips by l. Consider an agent who maximizes utility subject to a budget constraint:

$$\max_{d^w \in [0,1], d^l \in [0,1]} u(x) + g^w(d^w) + g^l(d^l)$$

$$\text{s.t.} \quad y \geq p^v v + p^b (T - v) + x,$$

where x is the outside good (whose price is normalized to 1), y is total income, p^v is the price per km of driving, and p^b is the price per km of the non-driving mode.

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

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

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

We can now solve for d^w and d^l. Assuming standard regularity conditions and using $v = d^w w + d^l l$, the optimal non-work travel decision can be characterized by the following first-order condition:

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

This condition is entirely standard; 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). In other words, since 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, $\frac{\partial^2 g^l(d^l)}{\partial (d^l)^2} \neq 0$ and the monotonicity of $g^l(d^l)$ assures an interior solution, as one would expect.

The setting is different for commuting, since $\frac{\partial g^w(d^w)}{\partial d^w} = \gamma^w$. 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 "bang-bang" solution for the choice of mode for commute travel:

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

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 are done by driving. Otherwise, all commute

trips are done by other forms of transport, such as public transport, cycling, or walking. 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.[4] 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 interested in the fuel price sensitivity of driving and the heterogeneity in this sensitivity. That is, we are interested in $\frac{\partial v}{\partial p^v}$ where $v = d^w w + d^l l$, holding w and l fixed. Consider the comparative statics with a change in gasoline prices at the optimal values of d^w and d^l. From the implicit function theorem we know that for non-work driving,

$$\frac{\partial d^l}{\partial p^v} = \frac{\lambda l}{\frac{\partial^2 g^l}{\partial (d^l)^2}} \tag{2}$$

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_0^v to p_1^v. 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 driving with the given change in gasoline prices is

$$\Delta d^w = \begin{cases} 1 & \text{if } p_1^v < p^b + \frac{\gamma^w}{\lambda w} < p_0^v, \\ -1 & \text{if } p_1^v > p^b + \frac{\gamma^w}{\lambda w} > p_0^v, \\ 0 & \text{otherwise.} \end{cases}$$

This expression highlights when switching might occur with a fuel prices 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 and put in monetary terms.

Now the response in total driving to the change in gasoline prices is given by:

$$\Delta v = \Delta d^w w + \Delta d^l l. \tag{3}$$

Thus, for households with very long commutes (i.e., a large w), a fuel price change sufficiently

---

[4]Note that this is a static, rather than dynamic model, so γ^w can be interpreted more as a threshold level of savings required than as a classic switching cost in a dynamic model. But the interpretation as a type of switching cost is still useful here.

large to induce a switch would imply a much greater decrease in driving. This can be restated as our first testable implication:

Proposition 1. *Households with a longer work distances are expected to be more responsive to changes in gasoline prices when there is sufficiently large gasoline price variation to induce transport mode switching for work trips.*

Equation (3) also shows that for the households with the shortest work distances, Δv becomes determined entirely by the change in non-work driving, Δd^l. According to equation (2), the price sensitivity for these households ultimately boils down to the curvature of the utility of driving for non-work travel ($g^l(\cdot)$). The underlying fundamentals determining the shape of $g^l(\cdot)$ are factors such as the availability of appealing substitutes to driving, the closeness of amenities, and the types of leisure activities that households with low work distances engage in. Our model imposes no a priori restriction on the curvature of $g^l(\cdot)$, so this is an empirical question. However, it is common to assume that non-work travel is more discretionary and thus may be more responsive to changes in gasoline price. Our second testable implication summarizes:

Proposition 2. *For households with very short work distances, the fuel price sensitivity of total driving approaches the fuel price sensitivity for non-work trips. To the extent that non-work trips are more discretionary, we would predict greater responsiveness to changing fuel prices for households with very short work distances.*

This simple model lays a theoretical foundation for analyzing the heterogeneity of fuel price elasticity of driving. It is intentionally a simple static model to build intuition for the short and medium-run driving decisions. We would expect to see the same switching behavior in a dynamic setting, whereby households could "invest" in switching if the discounted savings from doing so outweigh the switching cost.[5] Since the savings are just the work distance times the fuel price differential, this implies that households with longer work distances will switch for smaller changes in the fuel price. Such a mechanism can be explained in our model by allowing γ^w to be heterogeneous in the population and increasing in w. In the longer-run work distance could also be endogenized, but this is outside the scope of our paper.

---

[5]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.

3 Data

3.1 Data Sources

We use 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. Starting on July 1, 1998, all vehicles in Denmark have been required 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 km 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. This measure of work distance is the product of the reported work distance and the reported number of days that work travel occurred (regardless of mode of transport). Since the address of the work place is known to the tax authorities, this number is subject to auditing. The individual is only eligible for a deduction if the distance is greater than 12 km but there is no minimum requirement on the number of days worked. The work distance measure will therefore be equal to zero if the individual lives closer than 12 km from the work place or if the individual does not work. For 2000 through 2008, we have data on the actual work distance for 79.6% of the households measured using

---

[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).

[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.

[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.

a shortest-path algorithm and provided by Statistics Denmark.[9] We find that these two work distance measures are quite similar in a robustness check using the tax deduction measure (Appendix A.3.3). Further details on the dataset are in Appendix A.2.

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.[10] Similarly, we also bring in daily West Texas Intermediate crude oil price data for a robustness check.[11] Finally, we use data from Journey Planner on all bus and train stops in Denmark in 2013.[12]

We also have access to some additional car characteristics, including fuel economy in km/l and the manufacturer suggested retail price (MSRP). These data comes from a dataset from the Danish Automobile Dealer Association (DAF). However, these variables are not available for car vintages older than 1997 so we do not included them in the preferred specification.

3.2 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 km driven and the km 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 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 count of public transport stops is added to the dataset at the municipality level. For a detailed description of the variables used, see Appendix A.2.

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. Appendix A provides further details on the data cleaning process.

---

[9]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.

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

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

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

Table 1: Summary Statistics

	Mean	Std dev
Vehicle-km-traveled (km/day)	46.6	(40.2)
Gross income (DKK)	574,056	(627,921)
Gross income-couples (DKK)	646,638	(628,011)
Gross income-singles (DKK)	320,975	(558,098)
$\mathbb{1}$(Couple)	0.78	(0.42)
Age (oldest household member)	49.8	(14.4)
Reported work distance (km)	12.2	(19.7)
$\mathbb{1}$(Work distance > 12km)	0.50	(0.50)
Actual work distance (km)[a]	23.4	(35.6)
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)
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)
% of period owned by this owner	0.79	(0.30)
Driving period length (years)	2.34	(0.89)
# 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)
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.

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

We are not permitted to present the min and max due to Statistics Denmark rules.

3.3   Descriptive Evidence

There has been considerable variation in both gasoline and diesel prices in Denmark from 1998 to 2008. Figure 1 shows average gasoline and diesel 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.[13]

Figure 1: Vehicle Kilometers Traveled and Fuel Price

Note: Local polynomial average smoothed by start date.
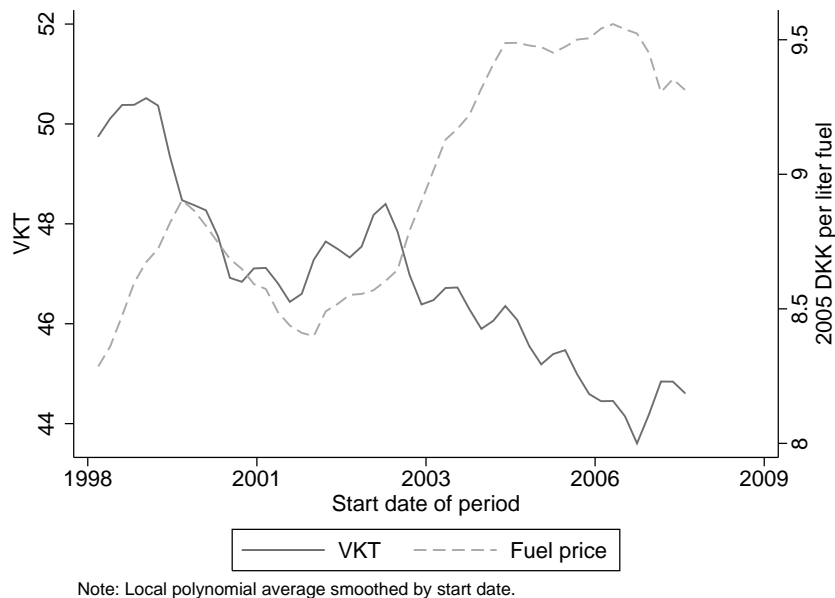
The rich Danish register data allow us to explore the relationship between fuel prices and driving in greater detail. Figure 2 divides the sample into ten groups based on the percentiles of driving in each year. The vehicles in each group may change over time, as we recalculate the percentiles in each year. The figure shows the interesting pattern 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. This provides the first evidence of the existence of the first tail of more responsive drivers, as suggested in Proposition 1.

Who are these drivers who drive the most? Table 2 shows the mean for selected demographics and other characteristics stratified by the amount driven. Not surprisingly, higher VKT drivers have a higher income, have more vehicles, have larger families, and live further

---

[13]See Appendix A for a figure of the unconditional distribution of VMT.

Figure 2: VKT Percentiles Over Time

Percentiles of VKT Over Time
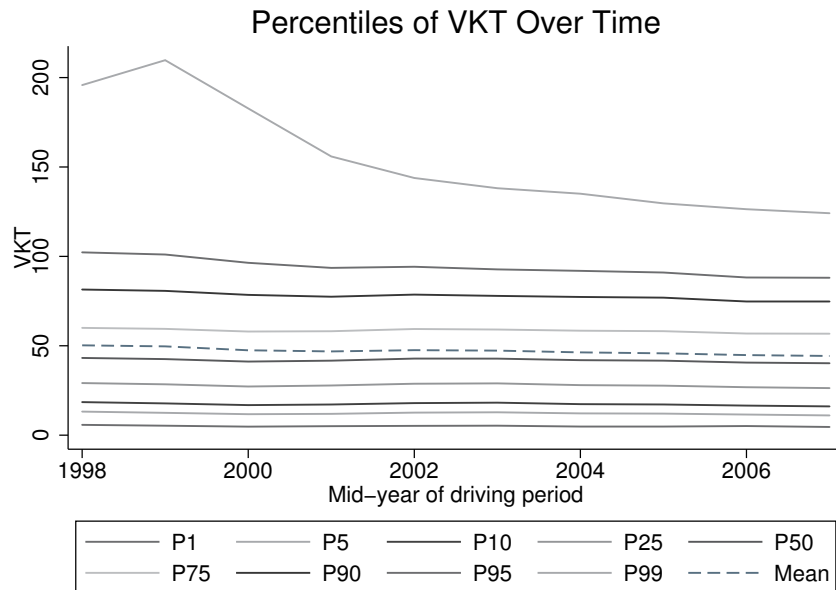
away from their workplaces. They also tend to drive heavier, younger, and diesel cars more frequently. Otherwise, these higher VKT drivers are similar to the average driver in most other characteristics, including access to public transport.

To visualize where the high-driving households live, Panel (a) of Figure 3 shows a map of Denmark where each municipality is shaded according to the average VKT of the households living in that municipality (darker means less driving). The figure shows that, conditional on owning a car, the high-VKT households tend to be in rural areas or on the outskirts of the major urban areas, while lowest-VKT households are in the urban areas.[14] The regions of high-VKT also tend to be the municipalities with longer work distances (see Figure 15 in the appendix).

An important way for a driver who lives further from work to be able to reduce driving is by switching to public transport. In many countries, such as the United States, access to public transport tends to be very limited. Panel (b) of Figure 3 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, as mentioned above. This pervasiveness of public transport makes switching behavior possible for those with long commutes.

---

[14]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.

Table 2: Means of Selected Variables Stratified by VKT

	VKT<100	VKT≥100
Gross income (DKK)	571,328	645,390
Gross income-couples (DKK)	644,262	705,238
Gross income-singles (DKK)	319,520	369,008
Couple dummy	0.78	0.82
Reported work distance (km)	11.6	26.9
Work distance > 12 km dummy	0.44	0.63
Actual work distance (km)[a]	22.7	39.3
Number of kids	0.75	1.00
Urban dummy	0.16	0.14
Self-employed dummy	0.10	0.16
Bus/train stops per km^2	15.9	14.2
Diesel dummy	0.13	0.50
Car age at start of period (years)	7.03	5.41
# additional cars owned	0.33	0.60
Observations	5,639,738	215,708

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

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

Figure 3: Average VKT by Municipality (Panel (a)) and Public Transport Stops (Panel (b))
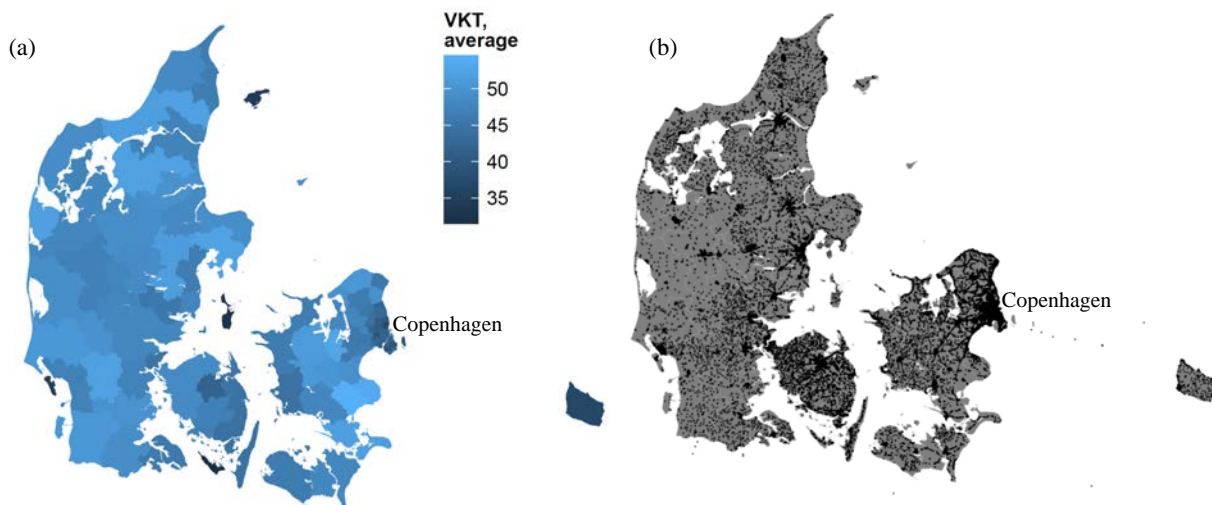
Figure 4 uses binned scatterplots to show the nonparametric relationship between per-vehicle driving and two key variables: household income (Panel (a)) and work distance (Panel (b)). Both exhibit remarkable heterogeneity. The graph of driving by income displays an inverted U-shape, although the amount of driving by the very wealthiest–who may live close to the center city and/or have multiple cars–is slightly lower than at the peak. The graph of driving by work distance demonstrates how longer commutes translate into more driving. There is a point mass at zero, which accounts for all households with a work distance less than 12 km. After this driving is a monotonically increasing concave function of work distance. This suggests that there is likely to be more non-work driving for households with shorter work distances since households with a much longer commute only drive slightly more than households with much shorter commutes. These descriptive results provide a useful context for our empirical results focusing on the change in driving with fuel price changes.

Figure 4: Nonparametric Relationship Between Driving and Work Distance and Income

(a) Household Income (DKK)   (b) Work Distance (km)
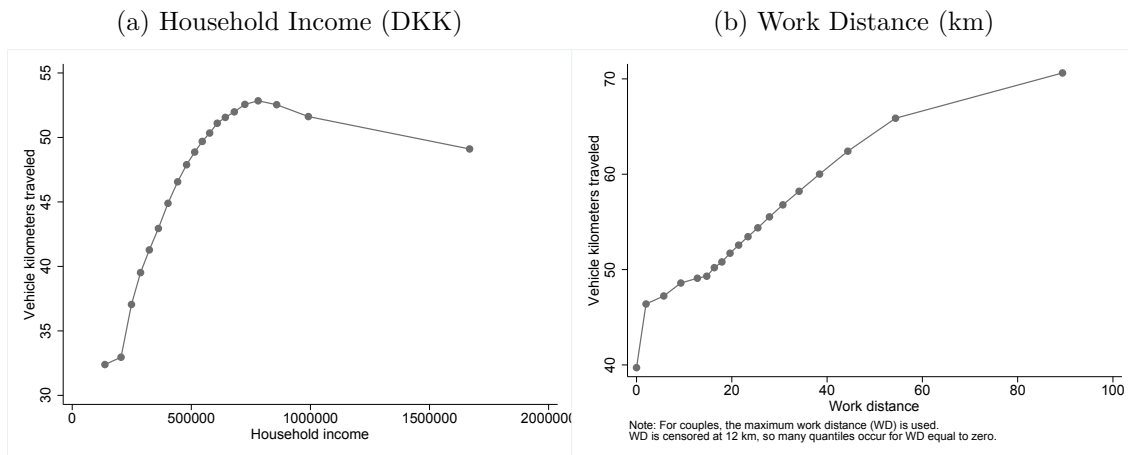
Notes: The dots are placed according to equi-distant percentiles of the conditioning variables and the corresponding y-value represents the average VKT within that percentile-group. The dots are connected by linear line segments for illustration purposes.

4 Empirical Approach

4.1 Empirical Specification

A primary goal of this paper is to investigate the fuel price elasticity and to explore the heterogeneity in this elasticity. 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.

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:

$$\log \text{VKT}_{iht} = \gamma \log p_{iht} + \mathbf{x}_{iht}\boldsymbol{\beta} + \sum_{y=1998}^{2011} \sum_{f=\text{gas,diesel}} \delta_{fy}\omega(i,t,y)\mathbb{1}(\text{gas}_i) + \eta_h + \varepsilon_{iht}. \quad (4)$$

VKT_{iht} is the average daily driving in kilometers and p_{iht} is the average daily fuel price over the driving period for vehicle i (gasoline or diesel price depending on the car type) and $\mathbf{x}_{iht}$ denotes a vector of controls. The coefficient γ is our primary coefficient of interest–the fuel price elasticity for vehicle i in driving period t. 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, the density of bus or train stops in the municipality, and vehicle characteristics. 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.[15]

The variable $\omega(i,t,y)$ denotes time controls, which vary by the vehicle fuel type $f \in \{\text{gas, diesel}\}$, year y, and driving period t in order to capture fuel type-specific factors that change over time. Our specification of these controls is motivated by the fact that a vehicle driving period is not exclusively in a single year, but generally covers two years and up to five years. This prevents us from using traditional year fixed effects. Instead, we allow $\omega(i,t,y)$ to denote the fraction of a driving period t that falls within the year $y \in \{1998, ..., 2011\}$. For example, if a driving period starts on July 1st 2001 and ends on June 30th 2003, $\omega(i,t,y)$ will be 0.25 for $y \in \{2001, 2003\}$ and 0.5 for $y = 2002$. The coefficients δ_{fy} will therefore act similarly to fuel type-specific year fixed effects, but since ω are continuous variables they afford the extra flexibility depending on the degree to which a driving period overlaps with year y. Since the weights sum to unity, we omit year 2003 as the reference year. In our robustness checks, we also examine alternative specifications for our time controls. Finally, η_h are household fixed effects, included to control for household time-invariant unobserved heterogeneity.

---

[15]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.

4.2 Identification

Of primary interest in this paper is the relationship between driving and fuel price–and how it varies across the population. 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 do not likely appreciably 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 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.

Another potential identification concern is the possibility of unobserved heterogeneity at the household level in vehicles and driving. We control for vehicle characteristics using our rich data, and more importantly, include household fixed effects to nonparametrically address time-invariant household unobserved heterogeneity. The household fixed effects are particularly important for identifying the coefficients on our work distance, urban area, and public transport variables since this allows us to focus on within-household variation (deviations from the household mean) in driving over time. Any sorting into different locations based on time-invariant unobserved preferences will be captured by the household fixed effects, and as such, we are identifying these coefficients largely from movers within our sample. The identifying assumption here is that people 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.[16]

---

[16]Because we control for commute distance, any change in unobserved preferences for driving would have to relate to non-work driving to be an issue. But non-work driving tends to be highly diverse, so we view it as highly unlikely that households will move for this reason. As a robustness check, we also estimate the model with municipality fixed effects and find extremely similar results.

5 Results

5.1 The Mean Elasticity of Driving

Table 3 shows the results from estimating the linear fixed effects model in equation (4). A very rich set of controls are included in the estimation, but for brevity, we only report selected coefficients (see the Appendix for the remainder). Column (1) is the most parsimonious specification, which only controls for car characteristics, seasonality (% of the driving period covering each month), and the driving period.[17] 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 drops 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 -0.52. Adding time controls reduces 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.[18] 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, the controls for work distance, company cars, and income appear to capture key components of driving demand, as evidenced by an R^2 of 0.34 in column (2), which is quite high for micro-data studies in a specification without household fixed effects.

The results in Table 3 also help fill out the story for how driving demand is determined. As was also seen in our descriptive analysis above, driving is increasing in work distance. This holds for males and females in a household with a couple, as well as for one-person households. Even the dummy for whether the work distance is non-zero (recall that it is censored at 12 km based on how the data are collected) has a positive and highly statistically significant coefficient. The results for both males and females indicate that increasing the work distance by one additional km can be interpreted as increasing daily driving by approximately 0.3%–an economically significant effect.

---

[17]For full list of variables, with details on each, see to Appendix A.2.

[18]In Appendix Table 16 we show that the elasticity is highly 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.

Table 3: Estimations of Driving Demand

	OLS		Household FE	
	(1)	(2)	(3)	(4)
$\log p^{\text{fuel}}$	-0.87***	-0.30***	-0.52***	-0.30***
	(0.01)	(0.01)	(0.01)	(0.02)
Work Distance (WD) controls				
WD, male		0.002***	0.003***	0.002***
		(0.00003)	(0.00003)	(0.00003)
WD non-zero, male		0.07***	0.03***	0.03***
		(0.001)	(0.001)	(0.001)
WD, female		0.003***	0.003***	0.003***
		(0.00004)	(0.00004)	(0.00004)
WD non-zero, female		0.058***	0.022***	0.026***
		(0.001)	(0.001)	(0.001)
WD, single		0.004***	0.004***	0.004***
		(0.0001)	(0.0001)	(0.0001)
WD non-zero, single		0.17***	0.07***	0.07***
		(0.002)	(0.002)	(0.002)
Other demographic controls				
log gross income (couple)		-0.03***	-0.03***	-0.02***
		(0.001)	(0.002)	(0.002)
log gross income (single)		0.03***	0.02***	0.02***
		(0.001)	(0.003)	(0.003)
$\mathbb{1}$(Urban area)		0.004***	-0.014***	-0.025***
		(0.001)	(0.003)	(0.003)
Bus/train stops per km^2		-0.001***	-0.0003***	0.00004
		(0.00002)	(0.0001)	(0.0001)
$\mathbb{1}$(Access to company car)		-0.19***	-0.10***	-0.098***
		(0.002)	(0.002)	(0.002)
Year controls	No	Yes	No	Yes
Age (quadratic)	No	Yes	Yes	Yes
Number of children	No	Yes	Yes	Yes
Self-employed	No	Yes	Yes	Yes
Household FE	No	No	Yes	Yes
R^2	0.20	0.34	0.18	0.18
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 car characteristics (a quadratic in weight, diesel dummy, van dummy, vehicle age, and number of vehicles of each type owned by the household), age of the male and female of the household, period controls, and % of each month control. The within R^2 is reported for the household fixed effects specifications. Robust standard errors clustered at the household level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Several other coefficients provide further new insight on driving demand. The highly 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. The urban area dummy is also highly statistically significant. The coefficient is positive in the specification without household fixed effects, column (2), which may seem surprising, since Figure 3 showed that major urban areas are associated with less driving. However, the coefficients should be interpreted as *conditional* on work distance and other covariates–so they suggest that holding work distance fixed, households in urban areas tend to drive more. When we add household fixed effects in column (4), the urban dummy coefficient changes sign to be negative. In this specification, the identifying variation comes from households that moved between urban and rural areas. It indicates that these households drive 2.5% less when they live in the urban areas.

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 3) and public transport access changes so rarely, there is limited variation in this variable and no time-series variation. Thus, it may not be surprising that the effect in column (3) is economically quite small and in column (4), the effect becomes statistically indistinguishable from zero.

The coefficient on the availability of a company car is negative and highly statistically significant in all specifications. In our sample, 3.8% of the households have access to a company car. The negative coefficient can easily be explained by the fact that the car doing the driving is not the company car, but is privately-owned. The coefficient implies that households with access to a company car drive nearly 10% less than other households, likely due to some switching of driving from the private car to the company car.[19]
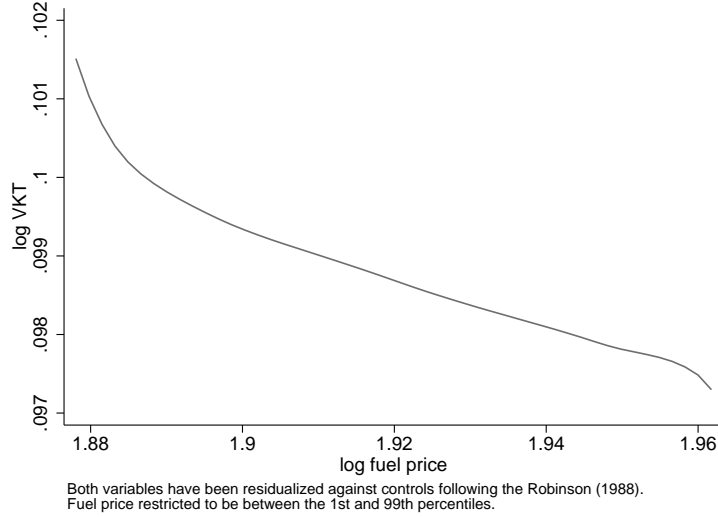
A natural question that arises when using the price elasticity for policy analysis is how well the functional form of the demand curve follows a constant elasticity assumption. Figure 5 plots a semi-parametric demand curve of log VKT versus log fuel price following the approach in Robinson (1988).[20] A key finding is that over a broad range of fuel prices, the functional form is approximately linear. This supports both the use of the log-log specification, as

---

[19] See Gutiérrez-i Puigarnau and Ommeren (2011) for a more complete treatment of the travel elasticity for households with company cars.

[20] Specifically, this is a double-residual approach, whereby first log VKT and log fuel price are *residualized* by regressing out the effect of all the remaining regressors from the primary specification, and then the residualized log VKT is regressed non-parametrically against the residualized log fuel price, using a local polynomial regression.

Figure 5: Semi-parametric Demand Curve

Both variables have been residualized against controls following the Robinson (1988).
Fuel price restricted to be between the 1st and 99th percentiles.

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, which would accord with intuition and underscore 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.

## 5.2 Heterogeneity in the Elasticity of Driving: Quantile Results

To better understand the mechanisms underpinning the fuel price elasticity, we leverage our rich dataset by using quantile estimation approaches. Specifically, we estimate a conditional quantile model at quantile τ:

$$\log \text{VKT}_{iht}(\tau) = \alpha(\tau) \log p_{iht} + \mathbf{x}_{iht}\boldsymbol{\beta}(\tau) + \sum_{y=1998}^{2011} \sum_{f=\text{gas,diesel}} \delta_{fy}(\tau)\omega(i,t,y)\mathbb{1}(\text{gas}_i) + \eta_h + e_{iht}(\tau).$$

This specification is particularly useful for examining the heterogeneity in the elasticity by estimating the quantiles of the conditional distribution of the coefficients. There are several approaches for estimating conditional quantile models with (quasi-)fixed or random effects. We estimate the parameters using the panel quantile estimator of Canay (2011) for

21

computational feasibility given the large size of our dataset.[21] The Canay (2011) estimator proceeds in three steps. First, we estimate a standard fixed effects estimator to obtain the within estimate $\hat{\eta}_h$. Second, we construct the regressand $\log \widetilde{\mathrm{VKT}}_{iht} := \log \mathrm{VKT}_{iht} - \hat{\eta}_h$. Third, we run the pooled quantile regression of $\log \widetilde{\mathrm{VKT}}_{iht}$ on all of our regressors (without the fixed effects). The first assumption underpinning the Canay (2011) quantile estimation approach is that the fixed effect is a pure location shift, i.e., it is not allowed to vary with the quantiles. Second, conditional on the observables, the fixed effect must be independent of the quantile error term. These assumptions also admit Chamberlain-style random effects, where the fixed effect is projected onto the time-average of the observables. However, the assumptions do not allow fully flexible, quantile-varying fixed effects, and consistency relies on both the number of cross-sectional and time-series observations growing without bound.

Figure 6: Elasticity by Conditional Quantile
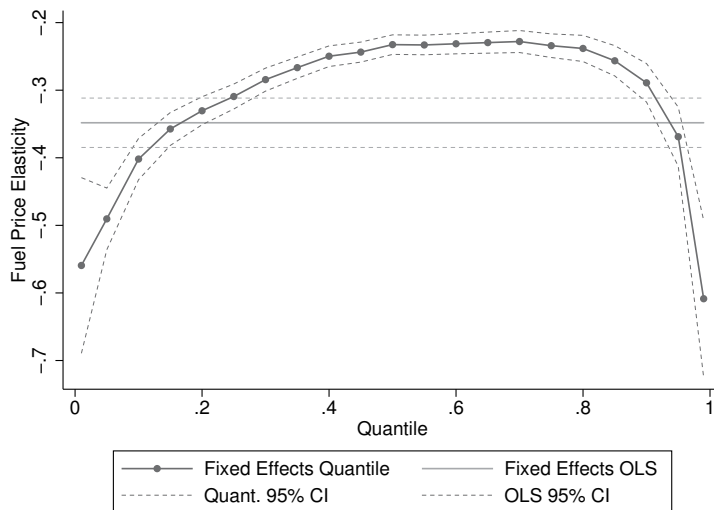
We estimate the panel quantile regression model for various quantiles. Figure 6 presents the results graphically (the full results are available in the Appendix). The figure shows a clear inverted U-shape in the fuel price elasticity. The lowest and the highest conditional quantiles of VKT have distinctly higher fuel price elasticities (in absolute value) than the

---

[21]In one alternative, Abrevaya and Dahl (2008) use a Chamberlain-style random effects estimator, projecting the η_hs on covariates from all periods or the time-averages. This essentially amounts to adding more regressors and running a pooled quantile regression. Koenker (2004) on the other hand takes a high-dimensional penalization approach, treating the η_hs as N additional parameters to be estimated, and penalizing the sum of absolute values of the fixed effects in the spirit of the LASSO estimator. Using clever computational tricks, he makes the approach computationally feasible for a small to medium size dataset. However, given the size of our dataset, this approach would be infeasible due to memory constraints, even run on a server with over 100GB of RAM.

middle region. These results provide strong evidence of two tails that increase the average responsiveness and considerably influence the price elasticity of driving.

The finding of two tails is consistent with the predictions of the theoretical model. The model suggested that households with long commutes but adequate access to public transportation stand to gain a great deal from switching their commute from driving to public transportation (Proposition 1). These may be the households in the upper tail, who are more responsive to changing fuel prices. Similarly, the model suggested that households with low driving demand would be likely to be taking a diverse set of trips with good substitutes available, and thus would also be more responsive (Proposition 2). The next sections provide further evidence on *who* is in these two tails.

5.3   Heterogeneity in the Elasticity of Driving: Interactions

To examine the characteristics of drivers in the two tails of greater responsiveness, we explore a similar specification to our primary model in equation (4) that includes interaction terms. We focus on interactions between the log of the fuel price and a subset of our controls. We denote this subset with $\mathbf{x}^1_{iht}$. This linear model with interactions is given by

$$\log \mathrm{VKT}_{iht} = \gamma_0 \log p_{it} + \boldsymbol{\gamma_1} \mathbf{x}^1_{iht} \times \log p_{it} + \mathbf{x}_{iht}\boldsymbol{\beta} + \sum_{y=1998}^{2011} \sum_{f=\mathrm{gas,diesel}} \delta_{fy}\omega(i,t,y)\mathbb{1}(\mathrm{gas_i}) + \eta_h + \epsilon_{iht}.$$

In $\mathbf{x}^1_{iht}$, we include a set of household demographic and car-related 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. 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.[22]

Table 4 shows the results from estimating the above equation. Column (1) shows the results without including interactions with car characteristics. Column (2) removes all interactions except car characteristics. Column (3) includes all interactions and is our preferred specification for interpretation. Comparing across columns demonstrates the robustness of

---

[22]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 negatively correlated with fuel prices).

Table 4: Estimations with Interactions to Demonstrate the Heterogeneous Elasticity

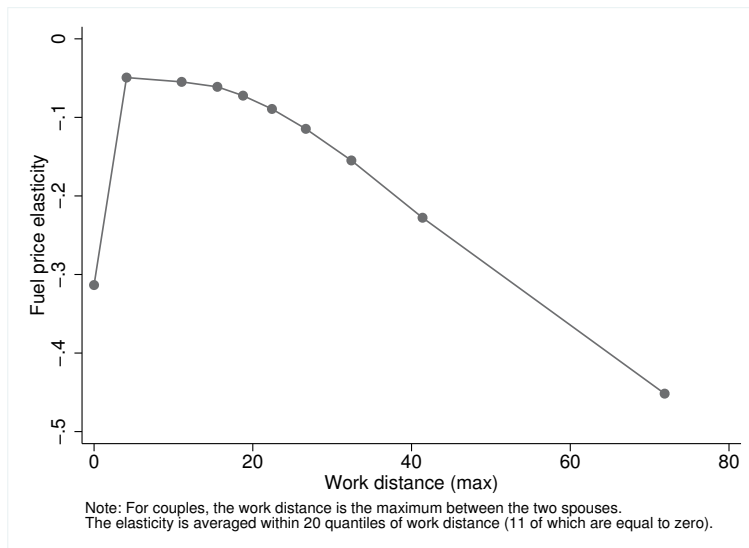|	(1)	(2)	(3)
$\log p^{\text{fuel}}$	-0.88***	-3.85***	-4.70***
	(0.24)	(0.09)	(0.28)
Work Distance (WD) interactions			
WD, male $\times \log p^{\text{fuel}}$	-0.01***		-0.01***
	(0.0003)		(0.0004)
WD, female $\times \log p^{\text{fuel}}$	0.003***		0.005***
	(0.0004)		(0.0004)
WD, single $\times \log p^{\text{fuel}}$	0.002**		0.003***
	(0.0007)		(0.0007)
WD squared, male $\times \log p^{\text{fuel}}$	-0.000006***		-0.000006***
	(0.0000002)		(0.0000002)
WD squared, female $\times \log p^{\text{fuel}}$	-0.000008***		-0.000008***
	(0.000001)		(0.000001)
WD squared, single $\times \log p^{\text{fuel}}$	-0.00001***		-0.00001***
	(0.0000006)		(0.0000006)
WD non-zero, male $\times \log p^{\text{fuel}}$	0.10***		0.12***
	(0.013)		(0.014)
WD non-zero, female $\times \log p^{\text{fuel}}$	0.19***		0.19***
	(0.013)		(0.013)
WD non-zero, single $\times \log p^{\text{fuel}}$	0.33***		0.33***
	(0.027)		(0.027)
Other demographic interactions			
log gross income (couple) $\times \log p^{\text{fuel}}$	-0.086***		-0.004
	(0.012)		(0.014)
log gross income (single) $\times \log p^{\text{fuel}}$	0.085***		0.040
	(0.022)		(0.027)
$\mathbb{1}$(Urban area) $\times \log p^{\text{fuel}}$	-0.014		-0.034
	(0.018)		(0.018)
Bus/train stops per km^2 $\times \log p^{\text{fuel}}$	-0.004***		-0.003***
	(0.0003)		(0.0003)
$\mathbb{1}$(Access to company car) $\times \log p^{\text{fuel}}$	-0.41***		-0.40***
	(0.023)		(0.024)
Car characteristics $\times \log p^{\text{fuel}}$	No	Yes	Yes
Household FE	Yes	Yes	Yes
Mean predicted elasticity	-0.25***	-0.29***	-0.24***
R^2	0.18	0.18	0.19
N	5,855,446	5,855,446	5,855,446

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 controls, car characteristics (a quadratic in weight, diesel dummy, van dummy, vehicle age, and number of vehicles of each type owned by the household), period controls, % of each month control, and interactions between the fuel price and a couple dummy, age, number of children, and the self-employed dummy. The mean predicted elasticity takes the mean predicted elasticity over all observations. The within R^2 is reported for the household fixed effects specifications. Robust standard errors clustered at the household level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

our results (e.g., the elasticity at the mean). The elasticity at the mean, presented near the bottom of the table, and the coefficients do not substantially change across columns.

We first focus on the work distance interaction variables. In column (3) in Table 4, we include interactions with a quadratic in work distance for males in a couple, females in a couple, and singles. We find coefficients that are highly statistically significant from zero for all of these work distance variables. There are also clear patterns that emerge. The coefficient on the work distance is negative for males, while it is positive for females and singles. As is shown in Appendix, it turns out that married males have by far the longest work distances. The smaller and positive coefficient for females in a couple, along with the finding that the female commute is much shorter, suggests that females either tend to have jobs that are more widely dispersed or that couples tend to locate closer to the female's workplace. Singles are more similar to females, perhaps because they can more easily locate close to their workplace.

The quadratic terms for the work distance variables are particularly useful for better understanding the how the responsiveness to fuel prices changes with work distance. Figure 7 illustrates this relationship. To develop this figure, we first calculated the individual-level predicted elasticities ($\hat{\gamma}_{it} = \hat{\gamma}_0 + \mathbf{x}_{it}^1 \hat{\gamma}$). For couples, we use the maximum work distance of the two members. Then we divided the work distance into 10 quantiles (since nearly half the sample has a work distance less than 12 km, we include all in this category in one quantile bin). Finally, we compute the average elasticity within each of these bins.

Figure 7: Fuel Price Elasticity and Work Distance (km)

Note: For couples, the work distance is the maximum between the two spouses.
The elasticity is averaged within 20 quantiles of work distance (11 of which are equal to zero).

We observe an inverted-U shape in Figure 7, just as we had earlier seen in Figure 6. For the shortest work distances (< 12 km), the fuel price elasticity is relatively high in absolute

value at -0.30. For slightly longer work distances, the elasticity decreases in absolute value to -0.05. But then it increases in absolute value again, reaching nearly -0.45 for work distances just over 70 km. These results help to lift the veil off the mechanisms leading to the two tails in Figure 6. Not only do we see the two tails here, but we show that they are closely related to the commute distance, further confirming our model in section 2. For drivers with very long commutes, smaller increases in fuel prices are required to decrease driving. For drivers with short commutes, most driving will be due to non-work trips, which are more diverse and thus it is likely that the drivers can substitute from driving to biking, walking, or using public transportation.

The remaining coefficients in Table 4 further clarify the heterogeneity in responsiveness. Column (1) indicates that high-income households are more responsive, but once we allow the elasticity to vary with the car characteristics as well in column (3), the interaction becomes statistically insignificant from zero. This result suggests that income effects are largely captured by the car choice. The coefficient on the interaction of the fuel price with the urban area dummy is statistically insignificant from zero in both columns (1) and (3). However, the coefficient on the interaction of the fuel price with the density of bus/train stops has a coefficient that is negative and highly statistically significant from zero. The coefficient implies that households living in a region with one standard deviation more stops available (18.4 stops/km^2) will have a 25.2% larger (in absolute value) elasticity at the mean. This result accords with intuition: when there is greater access to public transport, there is greater ability to substitute away from driving. The statistically significant coefficient on the bus/train stops variable and the insignificant coefficient on the urban area variable also suggest that previous results showing that drivers in urban areas are more responsive to fuel price increases (e.g., Gillingham (2014)) may be largely capturing a public transport effect. Our study the first we are aware of to disentangle the effect of public transport from being in an urban area. In addition, these results highlight the importance of public transport as a mechanism for fuel price responsiveness.

The coefficient on the interaction between the fuel price and access to a company car is negative and highly statistically significant from zero. This suggests that when fuel prices rise, households may switch over to their company car more. More broadly, this is consistent with multi-vehicle households switching to the least expensive vehicle option within the portfolio (De Borger, Mulalic, and Rouwendal, 2016a; Archsmith et al., 2016).[23]

---

[23]In the Appendix we provide additional evidence that multi-vehicle households are more responsive, as well as evidence that drivers of diesel vehicles are more responsive than drivers of gasoline vehicles.

5.4   Spatial Heterogeneity in the Elasticity of Driving

The results thus far indicate two tails, the first of which involves households with long commutes and the second households with short commutes. One might expect to see further evidence of the first tail in a particularly high responsiveness to fuel price changes in the outskirts of cities. Similarly, a high responsiveness to fuel price changes in cities themselves would build further evidence for the second tail. Figure 8 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.

Figure 8: The Average Elasticity by Municipality
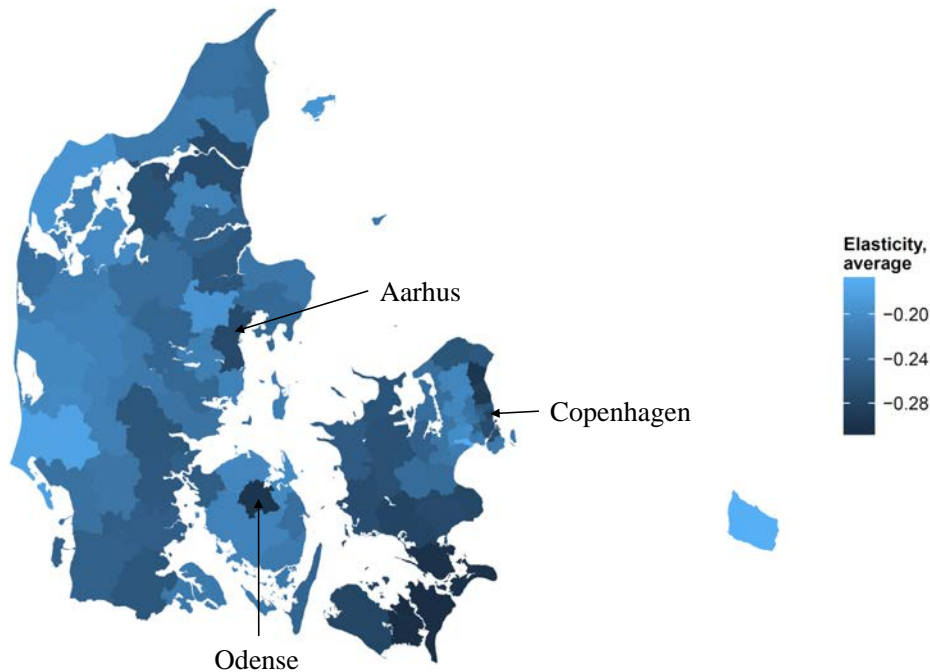
Two key findings emerge from Figure 8. First, some of the most responsive municipalities are in the largest cities. This aligns with Proposition 2 and the above evidence suggesting that there is a tail of more responsive drivers with short commutes. Second, many of the other most responsive municipalities 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 commute to jobs in Copenhagen. Access to public transport is excellent (recall Figure 3). Similar findings emerge for other areas in the outskirts of urban areas, further building evidence in support of Proposition 1.

5.5 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 have found 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 perfectly expectable, so we discuss them in the text below, going through each case in turn.

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

Name	Elasticity Range
Years in the sample	[-0.40;-0.28]
Length of driving periods	[-0.30;-0.28]
Fuel type	[-0.54;-0.26]
Singles or couples	[-0.32; -0.25]
Time controls	[-0.31;-0.30]
Instrumenting with oil price	-0.37
Household-vehicle fixed effects	-0.28

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 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 extremely robust to this robustness check and demonstrates that are results are not confounded by the timing of the inspections.

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 (4) separately for diesels and gasoline vehicles. We also examine a specification with an interaction between the log fuel price and a diesel dummy. The in-

teraction shows that the gasoline price elasticity is -0.26, while for the diesel segment it is -0.39. Estimating on separate samples yields corresponding elasticities of -0.27 and -0.54. 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 theoretical model 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 -0.25 respectively. These underscore the robustness of our primary estimated elasticity to the inclusion of either subsample.

The year controls employed in equation (4) 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. We find that the results are highly robust to removing our % of each month time 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.

Next, we consider carefully 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 a simultaneity 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. 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 (see the Appendix). The 2SLS fuel price elasticity estimate is highly statistically significant from zero at -0.37. This estimate is quite 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 highly statistically significantly different (e.g., the 99% confidence intervals overlap). For this reason and the computational complexity of many of our analyses, we chose not to use the 2SLS estimate as the preferred estimate.

Finally, we 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); Munk-Nielsen (2015)). We find our results quite robust to the exact choice of vehicle characteristics that are included. We also run a specification with household-vehicle fixed effects so that we were not making as much use of variation from households switching vehicles. The estimated elasticity with our preferred specification and household-vehicle fixed effects is -0.28. These results suggest that selection into vehicles is not a concern.

6 Illustrative Counterfactual Simulation

In this section, we analyze the implications of our empirical findings for the welfare effects of an illustrative increase in fuel prices by 1 DKK/l for both gasoline and diesel. The average gasoline price over the 1998 to 2008 period is 9.01 2005 DKK, so this represents a substantial price increase, but it is within the range of the variation in our data.[24] Such an increase may be due to exogenous swings in oil prices or a fuel tax policy. There is some evidence that consumers in the United States respond more to changes in gasoline taxes than to gasoline price swings (Li, Linn, and Muehlegger, 2014). To the extent that these differ in Denmark as well, then this counterfactual can best be thought of as an analysis of an exogenous fuel price increase rather than a tax policy.

With a fuel price elasticity of driving of -0.30, the proposed increase of 1 DKK/l translates into a 3.3% reduction in driving. If this occurred due to a tax policy, then fuel tax revenue would increase by 13.2%.[25] In terms of emissions, if households do not respond to the price on the extensive margin (i.e., by changing their cars) and if all vehicles respond in the same way, then our -0.30 estimate implies an elasticity of carbon dioxide or local air pollutant emissions with respect to the fuel price of -0.30. Munk-Nielsen (2015) estimates a discrete-continuous model of car choice and driving in Denmark and finds that when fuel prices increase by 1%, the fuel economy of newly purchased cars only increases by 0.1%. Hence, the -0.30 may be very close to the true medium-run change in emissions from this change in fuel prices. Fully analyzing the effect on other important vehicle externalities, such as congestion and accidents (Mayeres and Proost, 2013), is outside the scope of this paper, but our results can shed some light on how these external costs would change. For example, since the drivers in the two tails generally either live in the city or commute into the city, one might expect congestion to be alleviated from the price change.

The dramatic heterogeneity in the fuel price responsiveness maps into differing consequences of a change in fuel prices across quantiles of responsiveness. By computing the predicted response in VKT by quantile based on the quantile regression estimates, we find that the top 5th quantile accounts for 14.4% of the sum of the predicted responses in driving for the population and the bottom 5th quantile accounts for 4.8% of the total predicted response. This suggests that the tail of drivers with long work distances (the top 5th quantile)

---

[24]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.

[25]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 $(1 + 0.1713) \times (1 - 0.0333) = 13.23\%$.

are more important for the welfare and environmental implications of a fuel price change.[26]

6.1 Average Welfare Effects

When examining the welfare effects of a change in gasoline prices, there is both a transfer and a deadweight loss (DWL). We focus first on the calculating the classic DWL. We ignore external costs, follow the usual practice in assuming a constant supply curve over the relevant range of our price change (Hausman and Newey, 2016), and focus on the area under the demand curve (i.e., Harberger triangle in a linear model).

We can exponentiate both sides of equation (4) to rewrite demand as $\mathrm{VKT}(p) = \varepsilon p^\gamma \prod_{k=1}^{K} z_k^{\theta_k}$, where p is the fuel price, $z = (x, \omega, \eta)$, and $\theta = (\beta, \delta, 1)$. We then can calculate the DWL for a change in fuel price from p_0 to p_1 by

$$
\begin{aligned}
\mathrm{DWL}(p_0, p_1, \gamma) &= \int_{p_0}^{p_1} \mathrm{VKT}(p) - \mathrm{VKT}(p_1) \, \mathrm{d}p \\
&= \varepsilon \prod_{k=1}^{K} z_k^{\theta_k} \left[\frac{1}{1+\gamma} p^{1+\gamma} \right]_{p_0}^{p_1} - \varepsilon p_1^\gamma \prod_{k=1}^{K} z_k^{\theta_k} (p_1 - p_0) \\
&= \mathrm{VKT}(1) \left[\frac{1}{1+\gamma} (p_1^{1+\gamma} - p_0^{1+\gamma}) - p_1^\gamma (p_1 - p_0) \right].
\end{aligned}
$$

When we calculate this equation using our preferred specification (with a mean elasticity estimate of -0.30), we estimate a mean DWL of 0.59 DKK/l for a fuel price increase of 1 DKK/l. This suggests a substantial DWL from the increase in fuel prices, which might be expected given that fuel prices are very high in Denmark due to high fuel taxes. However, this calculation masks the heterogeneity that our results showed was so important. Recent work has shown that the DWL accounting for individual heterogeneity can be quite different than the average DWL (Hausman and Newey, 2016).[27] Using the same mean elasticity, we also calculate the transfer from consumers to producers (for a price increase) or the government (for a tax increase) based on p_1 times $\mathrm{VKT}(p_1)$, finding this transfer to be 37.3 DKK/l.

---

[26]Note this results does not stem from the log-log functional form, since the quantile regression allows the price parameter to vary freely over the conditional quantiles of VKT.

[27]Our panel dataset is much superior to the survey cross-section used in Hausman and Newey (2016), so many of the identification issues they address do not apply in our context.

6.2 Individual Welfare Effects

To account for individual heterogeneity in our DWL calculation, we develop a new procedure for obtaining a measure of the DWL for each observation based on the quantile regression results. We base our procedure on the argument in Koenker (2005, ch. 2.6) that suggests we can think of the quantile model as having each observation (i, t) randomly drawing a uniform quantile, u_{it}, and then being assigned parameters according to the quantile regression function, $\gamma_{it} = \gamma(u_{it})$ and $\boldsymbol{\beta}_{it} = \boldsymbol{\beta}(u_{it})$. In this case, we would write the model as $\log \text{VKT}_{it} = \gamma(u_{it}) \log p_{it} + \mathbf{x}_{it}\boldsymbol{\beta}(u_{it}) + \sum_y \delta_{f_i}\omega(i, t, y) + \eta_i$, with $u_{it} \sim \text{Unif}(0, 1)$. This formulation is advantageous because we do not observe u_{it}; if we did, we could simply plug in the quantile regression estimates into (6.1).

When u_{it} is not observed but has a known density, we can integrate it out in the spirit of Melly (2005) and Machado and Mata (2005). Thus, we replace the unobserved, latent deadweight loss with the *Integrated Deadweight Loss* (IDWL) given by

$$\text{IDWL}_i(p_0, p_1, 1) =$$
$$\int_0^1 \left(\prod_{k=1}^K z_{itk}^{\theta_k(u)}\right) \left[\frac{1}{1 + \gamma(u)} \left(p_0^{1+\gamma(u)} - p_1^{1+\gamma(u)}\right) - p_1^{\gamma(u)}(p_1 - p_0)\right] \mathrm{d}u. \tag{5}$$

This integral can be computed in a number of different ways; Machado and Mata (2005) use a simulation approach in a somewhat similar setting and Melly (2005) uses a grid. Given the computational requirements for estimating the model at even a single quantile (approximately 10 hours on a 64-core machine with 1 TB of RAM), we use a grid and let the computer time dictate the fineness of the mesh.[28] We use 21 grid points, so that $u_q \in \{0.01, 0.05, 0.10, 0.15, ..., 0.95, 0.99\}$.

Using this approach, we find an average DWL estimate of 0.56 DKK/l for the price increase of 1 DKK/l when we plug in the average characteristics over all observations. When we plug in the characteristics of each observation, the estimate is 0.66 DKK/l, which is our preferred estimate of the DWL. We can see the heterogeneity in the DWL through the standard deviation across observations, which is 0.38 DKK/l. Using the quantile estimation results, we again calculate the transfer from consumers (plugging in the characteristics of each observation), and find the mean transfer to be 43.7 DKK/l with a standard deviation of 25.9 DKK/l. This underscores the differing distributional effects across the population.
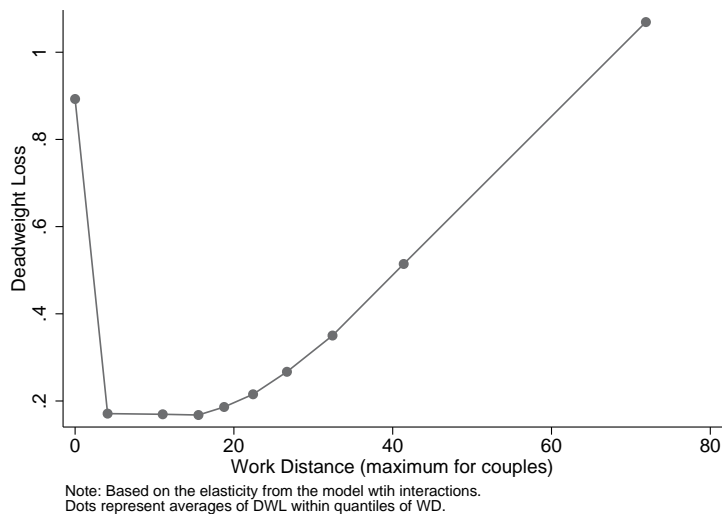
---

[28]Portnoy (1991) shows that with a finite sample, the estimated quantile regression function, $u \mapsto (\hat{\gamma}(u), \hat{\boldsymbol{\beta}}(u))$, will only change at a finite number of points on the interval $[0; 1]$ and that this number is $O(N \log N)$. Melly (2005) notes that for his estimator of the conditional distribution based on the quantile predictions, a mesh size on the order of $O(N^{-.5-\varepsilon})$ will ensure that the asymptotics still hold. For computational reasons, we are unable to scale up accordingly.

6.3 Heterogeneity in the Welfare Effects

We further explore the heterogeneity in the DWL by using the individual-level predicted elasticities from the model with interactions from Section 5.3. Given that we showed substantial heterogeneity in the elasticity based on the work distance, we are specifically interested in how the DWL varies with work distance. We divide the work distance (for couples we take the maximum) into 10 quantiles as before and calculate the DWL for each quantile using the predicted elasticity for each observation ($\hat{\gamma}_{it}$).

Figure 9 presents the results. As economic intuition would suggest, the most elastic households have the greatest DWL. These are households in the tails, with the highest and lowest work distances. For these households, the DWL is approximately 1 DKK/l, while at the middle part of the work distance distribution there is a much smaller DWL (around 0.2 DKK/l). These findings illustrate how heterogeneity in the elasticity by work distance carries over into substantial heterogeneity by work distance in the DWL–a novel finding in the literature. We see analogous heterogeneity across work distance in the transfer as well.

Figure 9: Deadweight Loss by Work Distance

Note: Based on the elasticity from the model wtih interactions.
Dots represent averages of DWL within quantiles of WD.

7 Conclusion

This paper estimates the medium-run (one-year) fuel price elasticity of driving demand for Denmark, with a preferred estimate of -0.30. We show that this elasticity is highly heterogenous, with two tails of much more responsive drivers than most of the population. 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 living in cities with short commutes. These two tails can be explained by a simple economic model in which households with long commutes can readily switch to public transport, while households in the city 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.

These results may help reconcile the results of studies in Europe and the United States. Most studies in the United States show fuel price elasticities in the range of -0.15 to -0.35, while those in Europe over the same time frame generally tend to be higher in absolute value. The first tail of drivers with long commutes relies on access to public transport for its existence. We posit that the difference in fuel price elasticities stems from the poor public transport options in the United States, which ensure that the first tail does not exist there. Our results indicating that public transport access has a very strong influence on the elasticity underscore this possibility. One implication of these results is that if the United States improved its public transport opportunities, the long-commute tail could emerge there as well.

The two tails also play an important role in the effects of gasoline price increases. We find that the high work distance tail households, defined as the households in the top 5% of the conditional driving distribution, account for 14.4% of the total response in driving. Other studies have found that the primary adjustment to fuel prices is in driving rather than the fuel economy of new cars (Bento et al., 2009; Munk-Nielsen, 2015), so this result implies that the tail households will bear the greater part of the aggregate reduction in driving and emissions in response to a fuel price increase. This has implications for other externalities from driving as well, as it suggests that fuel price increases may be particularly effective at reducing commuting congestion in Denmark.

We use our quantile regression estimates to estimate that the average DWL from a 1 DKK/l increase in the fuel price is 0.66 DKK/l. This seemingly very large DWL is most likely due to the high fuel prices in Denmark. The heterogeneity in responsiveness by work distance also carries over to DWL. We find that the two groups of tail households bear a four times larger DWL (and much larger transfers from consumers) than those with medium work distances, highlighting how the heterogeneity in the elasticity implies that fuel price increases affect the tails very differently than most of the population.

Finally, our findings have implications for the effect of fuel economy standards and other policies aimed at improving new vehicle fuel economy. Policies that improve new vehicle fuel economy lower the cost per kilometer of driving, leading to the well-known *rebound effect* of increased driving, reducing fuel savings and emissions reductions (Gillingham, Rapson, and

Wagner, 2015). If consumers respond the same way to changes in fuel prices as to changes in fuel economy, our results suggest a rebound effect of 30%, whereby 30% of the fuel savings from the fuel economy improvement are taken back by the induced driving. We are cautious in this interpretation because consumers may respond differently to fuel price changes as to fuel economy changes. We view this as an area worthy of continued study, perhaps building on the recent work by De Borger, Mulalic, and Rouwendal (2016b).

References

Abrevaya, Jason and Christian M Dahl. 2008. "The Effects of Birth Inputs on Birthweight: Evidence from Quantile Estimation on Panel Data." *Journal of Business & Economic Statistics* 26 (4):379–397.

Adamou, Adamos, Sofronis Clerides, and Theodoros Zachariadis. 2013. "Welfare Implications of Car Feebates: A Simulation Analysis." *The Economic Journal* 124:F420–F443.

Archsmith, James, Kenneth Gillingham, Christopher Knittel, and David Rapson. 2016. "The Household Vehicle Portfolio: Implications for Emissions Abatement Policies." *Working Paper* .

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.

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.

Canay, Ivan A. 2011. "A Simple Approach to Quantile Regression for Panel Data." *Econometrics Journal* 14 (3):368–386.

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.

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.

D'Haultfœuille, Xavier, Pauline Givord, and Xavier Boutin. 2014. "The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus." *The Economic Journal* 124 (578):F444–F480.

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

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, David Rapson, and Gernot Wagner. 2015. "The Rebound Effect and Energy Efficiency Policy." *Review of Environmental Economics & Policy* 10 (1):68–88.

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.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2015. "Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market." *Working Paper* .

Gutiérrez-i Puigarnau, Eva and Jos Van Ommeren. 2011. "Welfare Effects of Distortionary Fringe Benefits Taxation: The Case of Employer-Provided Cars." *International Economic Review* 52 (4):1105–1122.

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

Hausman, Jerry and Whitney Newey. 2016. "Individual Heterogeneity and Average Welfare." *Econometrica* 84 (3):1225–1248.

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.

Huse, Cristian and Claudio Lucinda. 2013. "The Market Impact and the Cost of Environmental Policy: Evidence from the Swedish Green Car Rebate." *The Economic Journal* 124:F393–F419.

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.

Kilian, Lutz and Daniel Murphy. 2014. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics* 29 (3):454–578.

Knittel, Christopher and Ryan Sandler. 2013. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." *National Bureau of Economic Research Working Paper No. 18849* .

Koenker, Roger. 2004. "Quantile Regression for Longitudinal Data." *Journal of Multivariate Analysis* 91 (1):74–89.

———. 2005. *Quantile Regression.* Cambridge University Press.

Levin, Lewis, Matthew Lewis, and Frank Wolak. 2014. "High-Frequency Evidence on the Demand for Gasoline." *mimeo, Stanford University* .

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.

Machado, José AF and José Mata. 2005. "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression." *Journal of Applied Econometrics* 20 (4):445–465.

Mayeres, Inge and Stef Proost. 2013. "The Taxation of Diesel Cars in Belgium - Revisited." *Energy Policy* 54 (0):33–41.

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

Melly, Blaise. 2005. "Decomposition of Differences in Distribution using Quantile Regression." *Labour Economics* 12 (4):577–590.

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

Portnoy, Stephen. 1991. "Asymptotic Behavior of the Number of Regression Quantile Breakpoints." *SIAM Journal on Scientific and Statistical Computing* 12 (4):867–883.

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.