Estimating the Fuel-tax Elasticity of Vehicle Miles Travelled from Aggregate Data

WERNER ANTWEILER
Sauder School of Business
University of British Columbia

SUMEET GULATI
Food and Land Systems
University of British Columbia

March 19, 2018

Abstract

This paper provides a novel and intuitive way to identify the fuel-tax elasticity of Vehicle Miles Travelled (VMT) using aggregate data. We approximate the VMT elasticity by subtracting the elasticity of fuel efficiency (a vehicle investment elasticity) from the overall elasticity of fuel consumption in a region. Our data from Canada allows significant variation in fuel taxes across regions and over time, giving us reliable estimates of the fuel-tax elasticity of VMT. Our estimates indicate that the long run impact from investment in fuel-efficient vehicles is slightly larger than the short run impact of reduced consumption, and our main specification yields an estimate of the fuel tax elasticity of VMT at approximately –0.65. As a part of this analysis we also find unique evidence of economically small but empirically significant "carbon leakage" from cross-border travel. Our methodology and results inform future work estimating the impact of fuel tax policies on distance-related negative externalities.

Keywords: carbon tax; fuel tax; carbon leakage; automobiles.

JEL Codes: Q53.

Antweiler: Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, V6T 1Z2, Canada. Phone: 604-822-8484. E-mail: werner.antweiler@ubc.ca. Gulati: Food and Resource Economics, The Faculty of Land and Food Systems, University of British Columbia, 2357 Main Mall, Vancouver, BC, V6T 1Z4, Canada. Phone: 604-822-2144. E-mail: sumeet.gulati@ubc.ca.
This paper was previously circulated as "Frugal Cars or Frugal Drivers? — How Carbon and Fuel Taxes Influence the Choice and Use of Cars."

1 Introduction

Gasoline taxes are used worldwide as a second-best instrument to address numerous unpriced externalities from driving,[1] and raise government revenue. Among the most important parameters determining its value is the fuel-tax elasticity of Vehicle Miles Travelled (VMT). Parry and Small (2005) demonstrate that externalities from congestion, accidents, and distance-related pollutants determine approximately 68% of the optimal tax in the United States. In their estimation, the distance-related response of the fuel tax—captured by the fuel tax elasticity of vehicle miles travelled (VMT)—is a critical component of the optimal gasoline tax.

Estimating a VMT elasticity is difficult in practice. Researchers require data on individual driving behaviour and sufficient variation in taxes over time (Graham and Glaister, 2002). Data on individual driving behaviour is hard to access due to privacy reasons, and where available, it suffers from limited geographical scope, a lack of longitudinal variation, or both. Data on fuel consumption by jurisdiction is easily available. Researchers commonly estimate the fuel price elasticity, but this elasticity reflects the joint effect of reduced driving and investment responses. In response to a fuel tax, motorists drive less and invest in fuel-efficient vehicles. The associated reduction in congestion, accidents, and other distance related pollutants such as noise, depends on the size and relative distribution of the motorists' response. If driving volume is price-inelastic, but motorists are relatively flexible in their choice of vehicles, then congestion, accidents, and the depreciation of infrastructure may remain unaffected by a fuel-tax increase. With data on fuel consumption by jurisdiction, separating this response into driving and investment components, remains a challenge.

In this paper, we present a method to infer the fuel tax elasticity of VMT from aggregate data. We use fuel demand data along with model-level new vehicle sales in provinces in Canada in 2001–2014. We are fortunate to have significant variation in fuel taxes across space and time. We first estimate the (long-run) effect of fuel taxes on fuel demand using an improved version of the estimation strategy in Li et al. (2014). This is the overall effect. Second, we estimate the effect of fuel taxes on the fuel efficiency of new vehicle purchases using two alternative strategies. This is the investment effect. Third, we infer the fuel tax elasticity of VMT as a residual component, subtracting the investment effect from the overall effect describing the assumptions needed for this approach to work.

Our paper makes several contributions. By using Canadian data on fuel consumption and new vehicle purchases in conjunction with Canada's uniquely large variation in fuel taxes, we are able to conduct a decomposition that identifies the fuel tax elasticity of VMT from aggregate data convincingly. Our approach and results provide the empirical foundation for future work on designing optimal fuel tax policies addressing the multitude of negative externalities as described by Parry and Small (2005).

[1]This includes externalities from emissions, depreciation of transportation infrastructure, road congestion, the risk of vehicle accidents, and others.

Empirically, our results provide estimates of the fuel tax elasticity of VMT at –0.65, which is derived as the difference between the overall tax elasticity of fuel demand of –1.44 and the tax elasticity of improved fuel efficiency from new vehicles of –0.79. Roughly interpreted, the behavioural component accounts for about 45% of the overall effect of fuel taxes, and the investment component accounts for about 55% of the overall effect.

We are also the first to empirically identify the effect of cross-border fuel tax differences on cross-border travel and domestic fuel demand. By employing an empirical strategy that follows Chandra et al. (2014), our paper finds a small amount of "carbon leakage". We also provide theoretical as well as statistical underpinnings for the price disaggregation strategy in Li et al. (2014), thus re-emphasizing the importance of the results established in their work. In addition to the work on aggregate fuel consumption, we also build on prior work by Klier and Linn (2011) on the fuel efficiency of vehicles.

In what follows we first present a review of the relevant literature in section 2. Then we provide a decomposition framework for estimating the fuel tax elasticity of VMT in section 3. We continue with a preview of our data in section 4. In section 5 we introduce a theoretical framework for the differential effect of permanent and temporary fuel price changes in support of the Li et al. (2014) approach that we rely upon. Building on this framework we present our empirical findings in section 6. We first present results for the overall effect of fuel taxes on fuel demand, and then the effect of fuel taxes on fuel efficiency of new vehicles. We conclude in section 7. At the end of the paper we include two appendices providing further insights into the statistical properties of our fuel-price/fuel-tax decomposition, and the theoretical derivation of the fuel-tax and fuel-price elasticities.

2 Related Literature

Our research is related to many subsets of literature. First it is related to those recognizing the second best nature of fuel and gasoline taxes. This includes, Parry and Small (2005), Fullerton and West (2002), and Feng et al. (2005), among others. Parry and Small (2005) derives the optimal fuel tax and then applies this to evaluate taxes in Britain and the USA. Fullerton and West (2002), and Feng et al. (2005) discuss second-best gasoline taxes to correct emissions from automobiles when individual emissions vary by vehicle type and fuel use.

Our research, contributes to the literature analyzing short-term and long-term gasoline demand. A meta-study Graham and Glaister (2002) summarizes the results from numerous earlier studies on the price and income elasticity of gasoline demand.[2] They find that there is a relatively wide range of estimates across OECD countries, with short-run price elasticities between –0.2 and –0.5, and long-run price elasticities between –0.75 and –1.35. Overall, they conclude "the overwhelming evidence from our survey suggests that long-run price elasticities will typically

[2] For an earlier survey of price elasticities see Dahl and Sterner (1991).

tend to fall in the –0.6 to –0.8 range." A more recent meta analysis, Brons et al. (2008) confirms that the demand for gasoline is price inelastic in the short-run and long-run. Several of these studies use estimates of gasoline price elasticity to evaluate the effectiveness of gasoline taxes as a means to reduce emissions. Sipes and Mendelsohn (2001) find that gasoline consumption is price inelastic in the short and long run. They conclude that using gasoline taxes to reduce emissions will likely require high tax levels. They also find that the income elasticity of gasoline is very low, which means that the incidence of a gasoline tax falls heavily on the poor. With increasing interest on carbon taxes, Davis and Kilian (2011) estimate gasoline price elasticities to evaluate the impact of a gasoline tax on carbon emissions. They recognize the endogeneity of prices, and to obtain consistent estimates include gasoline taxes as an instrument for gasoline prices. Nicol (2003) uses Canadian data. Using a demand system approach, they estimate gasoline own-price and income elasticities for five different regions in Canada as well as six different household groups.

We also contribute to a growing body of literature evaluating the impact of gasoline prices and taxes on the fuel economy of purchased vehicles. Consider some examples. Li et al. (2009) studies the impact of gasoline prices on fleet fuel economy. Busse et al. (2013) uses individual vehicle purchase data to evaluate consumer myopia in response to changes in gasoline prices. Barla et al. (2009) evaluates the distance traveled and fuel efficiency of Canadian vehicles in response to a change in gasoline prices. Barla et al. (2016) evaluates how gasoline prices affect the fuel economy of new vehicles in Québec, while Dorval and Barla (2017) show that the optimal gasoline tax in Québec should be differentiated regionally and is overall too low when considering the effect of a broad range of negative externalities.

We synthesize how the same changes in fuel taxes impacts gasoline consumption and new vehicle fuel economy. This is unique in the literature.[3] Our synthesis will allow policy makers to understand the contribution from short-run and long-run effects of changing fuel taxes to the overall reduction. It will help them adjust response constraints accordingly. Our second contribution is to clarify the framework for why responses differ depending on whether the changes are expected to be permanent—through changes in fuel taxes, or transitory—through a variation in non-tax fuel prices. We present a theoretical framework to explain why a consumption response to a fuel tax change should be larger than one to a non-tax price change. Then we illustrate this difference through our empirical analysis.

[3] See Basso and Oum (2007) for a survey the literature on gasoline and automobile demand. They compare estimates of changes in fuel consumption, mileage, and fuel efficiency due to changes in price and taxes from the different studies.

3 A Framework for Estimating the Fuel Tax Elasticity of Vehicle Miles Travelled

Our main aim is to pinpoint the fuel price (tax) elasticity of VMT. We do this by decomposing the overall effect of fuel taxes on gasoline demand into separate investment-related and behavioural components. By identifying the total effect and subtracting the investment effect, we can infer the behavioural effect that is our proxy for the VMT price elasticity.

Fuel demand—and in particular gasoline demand for automobiles—adjusts at different margins: the intensive margin of mileage driven; the extensive margin of replacing a fuel-inefficient vehicle with a more fuel-efficient vehicle; and the extensive margin of replacing a motor vehicle with alternative forms of transport. Fuel taxes have both short-term and long-term effects. The short-term effect is on gasoline consumption of the existing fleet. The long-term effect is on the, size, composition, and fuel efficiency of the fleet.[4]

Total fuel consumption F_t in year t can be decomposed into contributing elements. Let V_t denote the number of vehicles and M_t the total miles driven in year t. Then:

$$F_t \equiv V_t \cdot \left(\frac{M_t}{V_t}\right) \cdot \left(\frac{F_t}{M_t}\right) \tag{1}$$

Total demand for gasoline is the product of the number of vehicles, V_t, the kilometres driven per vehicle, M/V, and the fuel consumption per kilometre, F/M. Altering fuel consumption F/M requires long-term investments by motorists—by adopting fuel-efficient vehicles. Miles per vehicle M/V can change in the short term as motorists can decide to drive more or less with the vehicles they currently own.[5]

We can also introduce population, N_t, into our decomposition. Then:

$$F_t \equiv V_t \cdot \left(\frac{M_t}{N_t}\right) \cdot \left(\frac{F_t}{M_t}\right) / \left(\frac{V_t}{N_t}\right) \equiv F_t^h + F_t^f. \tag{2}$$

The miles-per-vehicle term M/V can be split into a miles-per-person term M/N and a vehicles-per-person term (V/N). Increasing urban density could lead to households scrapping second or third vehicles, or people moving closer to where they work. Thus V/N captures the composition effect of transportation mode choice. We can also decompose total fuel demand by British Columbians into fuel purchased within British Columbia, F_t^h, and abroad, F_t^f—superscripts denote

[4] In the long term, higher carbon and fuel taxes also affect where people live. However, comprehensive data on where people live and work and commute is outside the scope of this paper. We focus on consumption and investment decisions alone.

[5] According to estimates within Gillingham (2014), driving might be relatively inelastic in the medium to short-term. Increases in fuel prices and, importantly, expectations of future prices, will probably trickle slowly into (a) purchases of more fuel-efficient vehicles, or (b) adoption of alternative transportation.

'home' and 'foreign.' As many Canadians live close to the border, total fuel demand may be spread across the border especially when price differences are large, or other factors raise the volume of cross-border traffic.

Accordingly, in our empirical application we control for vehicles per capita, and also account for fuel consumption across the border.

Similarly, we can also consider the impact of new versus old vehicles in the fleet. We use the superscripts o and n to signify old (continuing) and new (recently purchased) vehicles. Then:

$$F_t \equiv F_t^o + F_t^n \equiv V_t^o \cdot \left(\frac{M_t^o}{V_t^o}\right) \cdot \left(\frac{F_t^o}{M_t^o}\right) + V_t^n \cdot \left(\frac{M_t^n}{V_t^n}\right) \cdot \left(\frac{F_t^n}{M_t^n}\right), \tag{3}$$

According to equation (3), total fuel consumption can be decomposed into fuel consumption by old and new vehicles in the fleet. Further, if miles per vehicle are similar across old and new vehicles—i.e., $M_t^o/V_t^o = M_t^n/V_t^n = M_t/V_t$—then we can rewrite the above equation as:

$$F_t \equiv \left(\frac{M_t}{V_t}\right) \left[V_t^o \cdot \left(\frac{F_t^o}{M_t^o}\right) + V_t^n \cdot \left(\frac{F_t^n}{M_t^n}\right)\right]. \tag{4}$$

Equation (4) corresponds to our estimates that follow. We have data on new vehicles sold by province and their fuel consumption per kilometre/mile driven from 2001-2014. This data helps us track the impact of changes in the second term of the square bracket on overall fuel consumption.

We rewrite equation (4) as functions of fuel taxes (τ), and introduce $m \equiv M/V$ as the vehicle mileage and $f \equiv F/M$ as the fuel consumption. Both $m(\tau)$ and $f(\tau)$ are functions of the tax rate with $m'(\tau) < 0$ and $f'(\tau) < 0$. Also denote the level of fuel consumption for old cars as $\bar{f} = F_t^o/M_t^o$, and define $\phi_t \equiv V_t^n/V_t$ as the share of new vehicles in a given period. Then:

$$F_t(\tau) = m(\tau)\left[(1-\phi_t)\bar{f} + \phi_t f(\tau)\right] V_t \tag{5}$$

Differentiating (5) with respect to τ and transforming the expressions into elasticities yields

$$\left[\frac{\mathrm{d}F_t}{\mathrm{d}\tau}\frac{\tau}{F_t}\right] = \left[\frac{\mathrm{d}m}{\mathrm{d}\tau}\frac{\tau}{m}\right] + \left[\frac{\mathrm{d}f}{\mathrm{d}\tau}\frac{\tau}{f}\right]\frac{\phi_t f}{(1-\phi_t)\bar{f} + \phi_t f} \tag{6}$$

Important assumptions. *First*, $\bar{f}$ is unchanged from taxes. This implies that the only change in fuel efficiency in the fleet comes from the purchase of new vehicles (for which we have data). This assumption is likely violated because there is an active used car market through which cars can be imported when taxes are imposed. Further, there are active retirements in the old vehicle fleet, which are likely determined by fuel taxes. The *second* assumption we make is that the proportion of new vehicles in the year ϕ_t are also unaffected by taxes. In other words, the decision to buy a new car is not accelerated by the level of gasoline taxes. Note that in reality both assumptions, tend to offset each other.

We introduce elasticities with respect to changes in τ into equation (6): the tax elasticity of overall fuel consumption, η_F; the tax elasticity η_f of fuel consumption per mile for the new vehicle fleet (for a particular year); and the tax elasticity η_m for VMTs. The vehicle fleet turns over slowly, and thus ϕ_t is small. However, the full effect of the policy change is realized only after the vehicle fleet turns over fully and $\phi_t \to 1$. Recall that by assumption there is no change in fuel efficiency in the old vehicle fleet. The short-term elasticity above turns into the long-term elasticity once fleet turnover is complete, and it follows that

$$\underbrace{\eta_F}_{(-)} = \underbrace{\eta_m}_{(-)} + \underbrace{\eta_f}_{(-)} \tag{7}$$

Equation (7) illustrates our method to estimate the fuel tax elasticity of VMT. We estimate η_F—the overall effect, and we estimate η_f, the investment effect separately for Canada. Given our assumptions, the fuel tax elasticity of VMT η_m is then just the residual calculated as $\eta_F - \eta_f$. In other words, subtracting the investment effect from the overall effect gives us the behavioural effect from the increase of a fuel tax.

4 Data Overview

4.1 Sources

We have two components for our analysis. An estimation of gasoline demand, and an estimation of new vehicle sales. For our analysis of gasoline demand we combine data from a number of sources. From Statistics Canada we obtain data on monthly gasoline demand by province (Table 405-0003); annual vehicle registration by province and type (Table 405-0004); annual disposable income by province (Table 384-5000); monthly population estimates by province (Table 051-0005); and population of metro areas on July, 2011 (Table 051-0056). From the Bank of Canada we obtain average monthly exchange rates, and from Natural Resources Canada we get average weekly fuel prices and taxes by metro area. From this data we derive monthly provincial fuel prices and taxes using population-weighted weekly averages by metro areas within each province. From the US Department of Transportation's Bureau of Transportation Statistics we obtain border crossings from Canada into the United States (monthly, by border crossing and US state). We also use corresponding data from Statistics Canada (Table 427-0002) that records entries into Canada from the US. From the US Energy Information Administration we obtain daily West Texas Intermediate crude oil prices in US dollar, and construct a geometric monthly average for use in our analysis. We combine all these data into a provincial panel for 2001–2014.

For our analysis of vehicle purchases we combine data from two major sources. Yearly new vehicle sales counts by model and province were generously provided by *Desrosiers Automotive Consultants Inc.*, also a panel from 2001-2014. We match the vehicle sales counts with vehicle-specific fuel economy data from the US Environmental Protection Agency (EPA).[6] The sales data from *Desrosiers* does not differentiate trim, engine, transmission, or even fuel used in their sales. Our sales counts are a sum of all sales for that model in that year. The data provided by the EPA lists fuel economy by model trim, and transmission. They also list fuel economy separately for highway, city and combined (55% city, and 45% highway) driving. We use the combined fuel economy, and match the median fuel economy rating (within automatic transmissions among the listed trims) for each model.

4.2 Background and Preview

Changes in the monthly consumption of gasoline in a province reflects short-run and long-run changes by consumers. In the short run, drivers can reduce wasteful driving, reduce, cancel, or combine trips, increase their use of public and alternative transportation (walking, biking etc.). In the long run consumers can buy more fuel efficient vehicles, or relocate to reduce their gasoline demand by moving closer to work or closer to public transit. Monthly sales of gasoline (as illustrated above) reflect a combination of these short and long run changes.

In Figure 1 we graph per capita fuel consumption for the four large provinces in Canada. Alberta has the highest per capita fuel consumption of the four, followed by Ontario. Quebéc, and British Columbia alternatively have the lowest consumption per capita. British Columbia's consumption displays a marked downward trend in fuel consumption per capita post 2008—the year the carbon tax was introduced. For the other provinces there is no discernible trend.

To illustrate how investment in fuel efficiency has changed over time, in Figure 2 we graph the sales-weighted average fuel consumption per 100KM for the new vehicles purchased every year in the four large provinces. British Columbia has a higher average sales weighted fuel consumption than QuÃ©bec and Ontario, but lower than Alberta. All four provinces have a seemingly common downward trend, especially around 2011.

We explain changes in monthly consumption, or changes in fleet fuel economy measures by exploiting variation in fuel prices and fuel taxes. In Figures 3 and 4 we illustrate relevant price and tax movements. Figure 3 shows the population weighted nominal price for gasoline in Canada for the years 2001-15. It highlights several important features. There is a seasonal component. Prices tend to be higher during the summer "driving season". The major recession of 2008 led to a significant drop in crude oil prices and recovery in 2009 and 2010. During 2011-2013 prices remained relatively flat, and the end of 2014 witnessed the collapse of oil prices that continued throughout 2015. Because inflation was low, real and nominal prices are very similar. Figure 4, illustrates all fuel (including carbon taxes) for the four large provinces (Ontario [ON], Quebec [QC], Alberta [AB], and British Columbia [BC]). Taxes also include Canada's value-added federal tax, the General Sales Tax [GST], as well as provincial counterparts where applicable. In particu-

[6] Freely available online at http://www.fueleconomy.gov/feg/download.shtml.

Figure 1: Gasoline Demand Across Provinces (L per Capita), 2001–2015

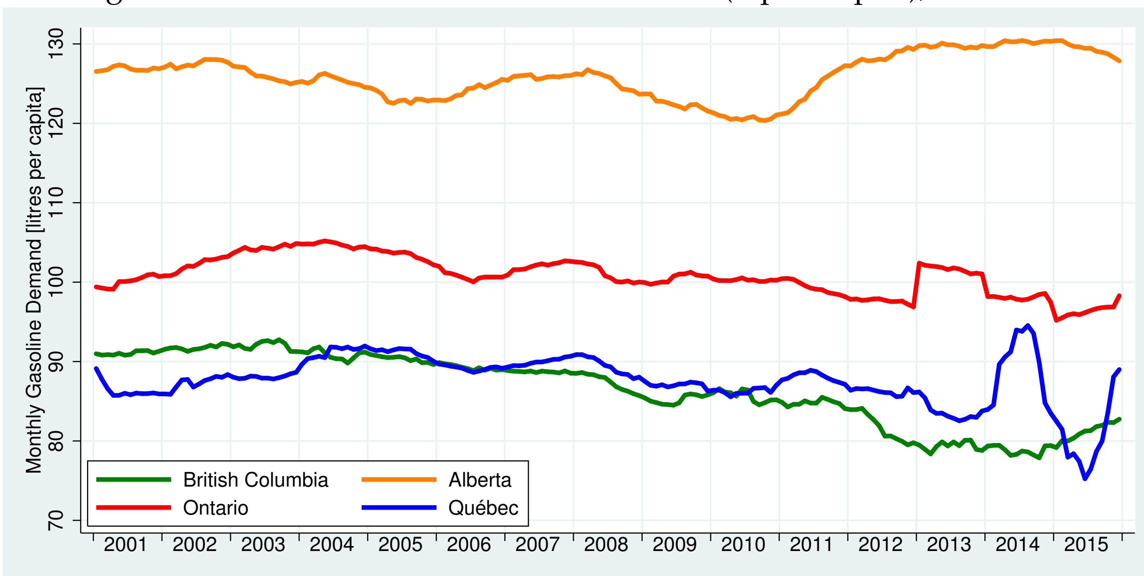

Source: Statistics Canada, CANSIM Table 405-0003. Gasoline demand is a 12-month moving average in order to suppress seasonality. Reporting anomalies in Ontario and Québec since 2013 distort the moving averages.

Figure 2: Fuel Consumption Across Provinces (L/100km), 2001–2014

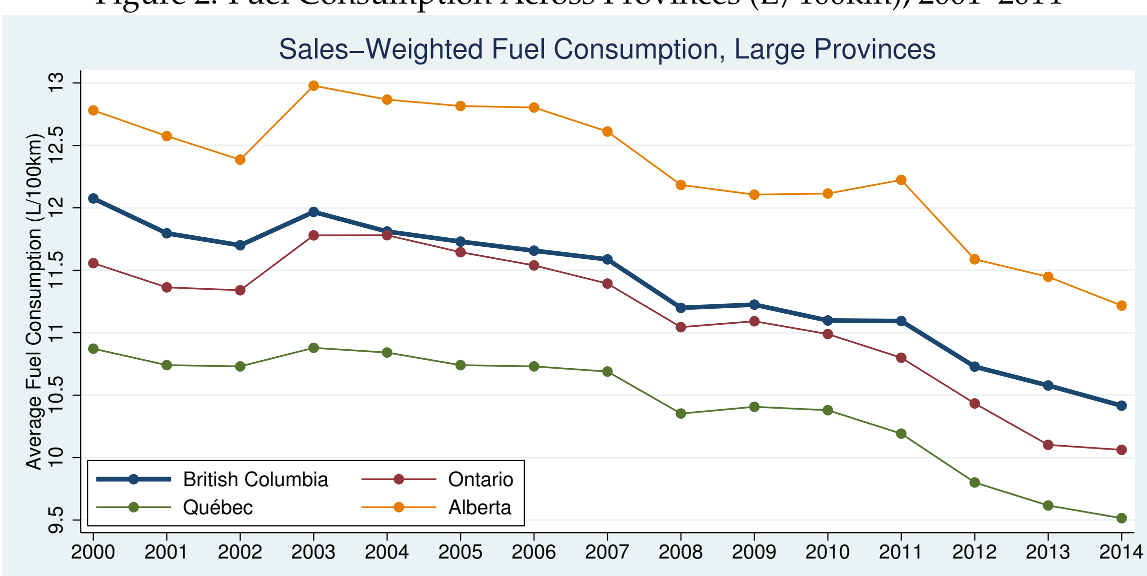

Source: Desrosiers Automotive Consulting Inc. and the Environmental Protection Agency.

lar, Ontario adopted the Harmonized Sales Tax [HST] in July 2010. The tax reform combined the 5% federal GST with the 8% provincial sales tax, and the combined 13% HST was applied to fuel. This amounted to a sudden 8% increase in the gasoline price, which is about the same magnitude in cents per litre as the carbon tax in BC. The jump is clearly visible in figure 4 (red line). BC's carbon tax took off in 2008 with several increments and the green line tracks these increases through 2012. The province of Québec has even higher fuel taxes than BC as the province increased the excise tax over the last few years, primarily to fix a widening budget deficit. Québec has excise taxes as high as those in BC, but also applies the provincial sales tax on top (14.975% compared to 5% in BC). The Greater Montreal region also imposes an additional 3-cent public transportation tax. Taken together, the fuel tax level in Québec (blue line) exceeded that in BC slightly during the most recent years.

5 The Differential Effect of Permanent and Temporary Fuel Price Changes

A crucial element in our analysis concerns the proper identification of the effect of policy on VMT. Fuel prices change rapidly because of volatile prices for crude oil. However, fuel taxes have a much more permanent effect because they change slowly and are fully taking into account by motorists when planning vehicle purchases. The differential effect of permanent and temporary fuel price changes has been established convincingly in Li et al. (2014). We provide a theoretical underpinning for this result, which is crucial for identifying the VMT price elasticity that is relevant for policy analysis. Specifically, our methodology derives the ratio of the fuel tax elasticity relative to the wholesale price elasticity. This demonstrates why the ratio can be quite large, which is what our estimates support.

The literature on gasoline demand traditionally distinguishes between short-term and long-term elasticities, but not between permanent and temporary (volatile) components of fuel prices. We expect that changes in fuel prices due to a permanent change in taxes provide greater incentives for purchasing fuel-efficient vehicles, and thus a larger reduction in fuel consumption than transitory fuel price changes. In this section we illustrate this mechanism through a motorists' investment decision problem under uncertainty. We employ a certainty-equivalent model to avoid algebraic complications, but preserve the insights necessary.

Owning a vehicle is repeated investment problem. The vehicle owners objective is to choose the time of replacement T to minimize the present value of the cost of vehicle ownership subject $V(T; a)$ subject to achieving a required amenity level a. The value function includes the cost of purchasing the vehicle: $K(f, a)$—strictly increasing with amenity a and strictly decreasing in fuel consumption in litres per kilometre f, the cost of fuel consumption: $fmp(t)$—where m is the kilometrage per period (assumed fixed), and $p(t)$ is the price of fuel per litre, a schedule of maintenance costs $c(t)$, a discount rate ρ, and a term to capture the recursivity

Figure 3: Regular Gasoline Retail Price, Population Weighted Canadian Average, 2001–2015

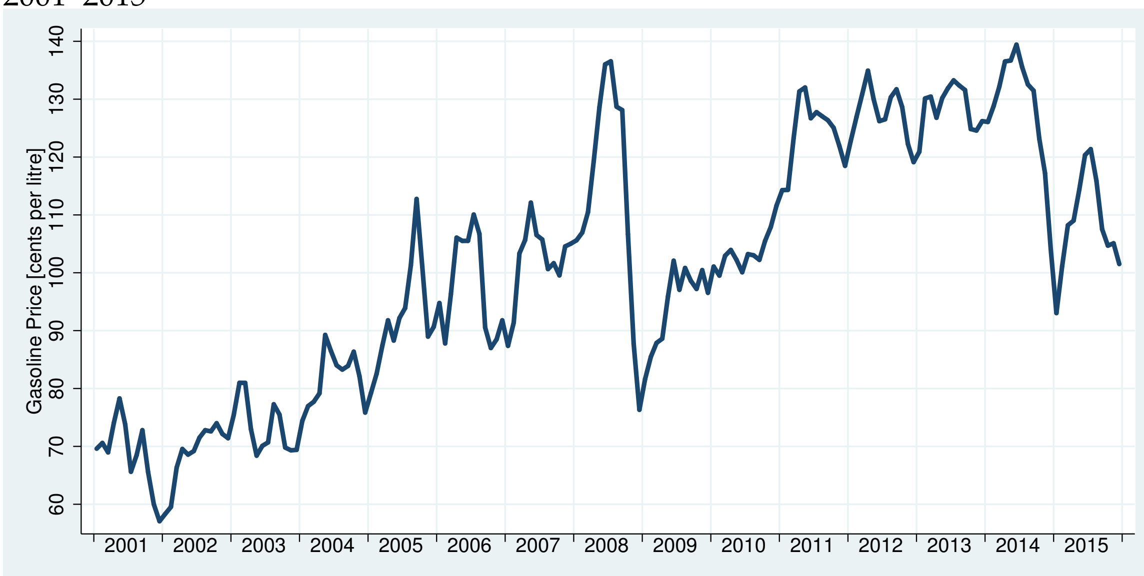

Source: Natural Resources Canada.

Figure 4: Fuel Taxes, Large Canadian Provinces, 2001–2015

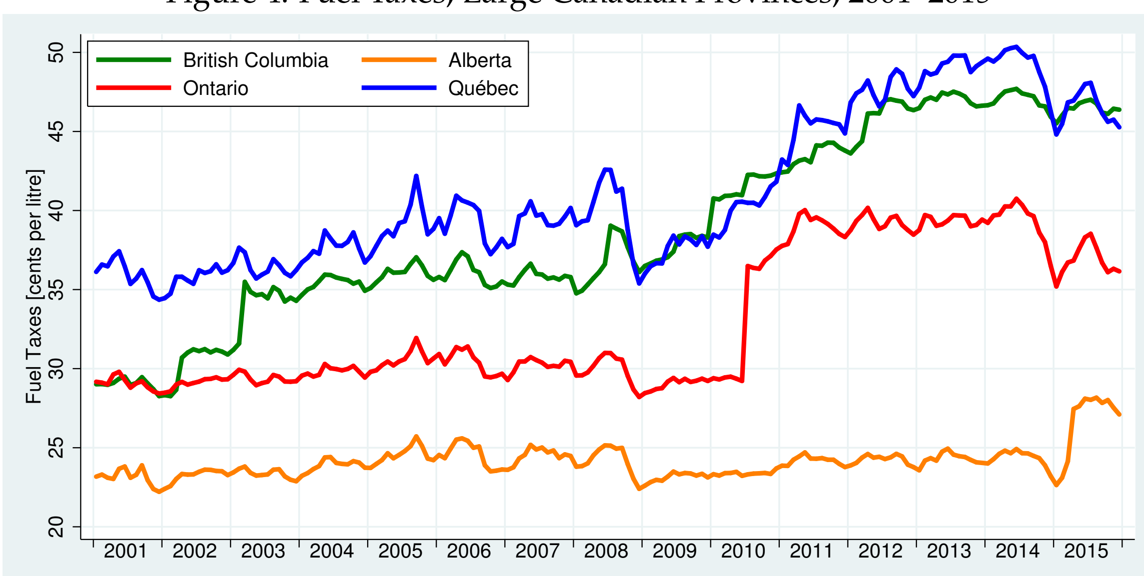

Source: Natural Resources Canada.

$V(T;a)\exp(-\rho T)$. Formally,

$$V(T;a) = K(f,a) + \int_0^T [fmp(t) + c(t)] \exp(-\rho t) dt + V(T;a) \exp(-\rho T). \tag{8}$$

We assume that the scrap value of the vehicle is zero. We also assume that at the end of the vehicle's life it is replaced with a vehicle of equal value (stationarity assumption). Finally, we assume that maintenance and operating costs are assumed to evolve at a constant growth rate γ so that $c(t) = c_0 \exp(\gamma t)$.

To illustrate the differential impact of permanent and transitory fuel price changes we decompose the price of fuel $p(t)$ into a permanent price $\bar{p}$, which in turn is composed of a tax element and a long-term gasoline cost element, and volatile temporary price $\hat{p}$ that may be larger or smaller than the long-term price $\bar{p}$. We model transitory price changes by assuming that a price shock $\hat{p}$ decays over time at rate $\lambda > 0$ and reverts to the fundamental price $\bar{p}$. Formally,

$$p(t) = \hat{p}e^{-\lambda t} + \bar{p}[1 - e^{-\lambda t}], \tag{9}$$

with initial condition $p(0) = \hat{p}$ and $\lim_{t\to\infty} p(t) = \bar{p}$. The half-life of the adjustment speed is thus $\ln(2)/\lambda$. For example, if $\lambda = 0.347$, a price-shock deviation decays by half in two years. Mean reversion in commodity prices has a strong empirical foundation (Schwartz, 1997; Pindyck, 1999); it also an important feature in recent work by Allcott and Wozny (2014) concerning gasoline prices and vehicle purchases.

Model (8) and (9) can be solved for T. Intuitively, as fuel and maintenance cost rise they trigger vehicle replacement at time

$$T^* = \text{argmin}_{T>0}(V(T;a)). \tag{10}$$

The resulting equation can be explored through implicit differentiation to determine the two fuel consumption elasticities $\hat{\phi}$ and $\bar{\phi}$ with respect to changes in the temporary price $\hat{p}$ and the permanent price $\bar{p}$:

$$\hat{\phi} \equiv \frac{\partial f}{\partial \hat{p}} \frac{\hat{p}}{f} \quad \text{and} \quad \bar{\phi} \equiv \frac{\partial f}{\partial \bar{p}} \frac{\bar{p}}{f} \tag{11}$$

In our mathematical appendix we show the details of how we derive the result that the ratio of the two elasticities, evaluated when $\hat{p} = \bar{p}$ is

$$\omega \equiv \left.\frac{\bar{\phi}}{\hat{\phi}}\right|_{\hat{p}=\bar{p}} = \left[1 + \frac{\lambda}{\rho}\right] \left[\frac{\exp(\rho T) - 1}{\exp((\rho+\lambda)T) - 1}\right] - 1 > 0 \tag{12}$$

The elasticity ratio depends on three variables: the discount rate ρ, the price-shock decay rate λ, and the vehicle lifespan T. The ratio ω does not depend on how much a vehicle is driven. Consider a numerical example with $T = 12$, $\rho = 0.05$ and $\lambda = 0.391$ (a shock half-life of 21 months), then $\omega = 3$. The fuel-efficiency response of permanent price changes would be three times larger than the corresponding response to a temporary price shock. Empirically, the relationship between λ and

ω is nearly linear for a given ρ and T. Our model links empirical estimates of the decay rate λ of price shocks directly to a predicted elasticity ratio. Conversely, empirical estimates of observed elasticity ratios imply a decay rate for price shocks. As it turns out, our empirical results are consistent with plausible values for T, ρ, and λ. We discuss further details in our Mathematical Appendix, also with respect to the magnitudes of $\hat{\phi}$ and $\bar{\phi}$, and the plausible magnitude of the decay rate λ. With a higher decay rate, this ratio is higher. In other words, if the transitory price change is expected to dissipate faster, fuel consumption will respond more to permanent price changes than to transitory ones. This ratio will be illustrated in our estimates in the following section.

6 Empirics

We have two distinct empirical strategies. Our first empirical strategy estimates the impact of gasoline prices and taxes on fuel demand—the overall effect, or our estimate of η_F (see Section 3). This strategy uses insights from Li et al. (2014) and Chandra et al. (2014). We modify the empirics from Li et al. (2014) to account for 'carbon leakage' from cross-border vehicle trips associated with fuel taxes.[7] Based on Chandra et al. (2014) we recognize that vehicular trips from Canadian provinces to bordering US states are determined by deviations in the nominal CAD-USD exchange rate from purchasing power parity. We also recognize that these trips are influenced by a variation in gasoline prices and taxes across provinces and time. These are some of the same covariates determining gasoline consumption in Li et al. (2014). For this reason we cannot just include cross-border trips in the regression equation proposed by Li et al. (2014). To account for endogeneity we instead use the variation in the Canadian Dollar-US Dollar exchange rate from purchasing power parity (among others) as an instrument for cross-border trips. This instrument is highly correlated with cross-border trips, and readily satisfies the exclusion restriction.

Our second empirical strategy estimates the impact of gasoline taxes on the fuel economy of new vehicles sold in Canadian provinces—the investment effect, or our estimate for η_f (see Section 3). We use two approaches. First, we analyze the impact of gasoline prices and taxes on the sales-weighted fuel efficiency of the average vehicle sold in each province. This provides us with a fleet-wide elasticity of taxes on fuel economy. Second we exploit the richness of our vehicle-level data to estimate how fuel and carbon taxes influence the market share of models according to their fuel efficiency, based on Klier and Linn (2010, 2011).

[7] Banfi et al. (2005) study cross border travel into Switzerland due to a differential in tax rates. They find that cross-border travel varies significantly as tax rates change.

6.1 Prices, Taxes, Gasoline Demand, and Carbon Leakage

In this subsection we estimate the impact of gasoline prices and taxes on fuel demand—the overall effect. Our aim is to derive an estimate of η_F (see Section 3), the elasticity of gasoline demand from gasoline taxes.

As described in Section 3, Canada's domestic demand for gasoline is influenced by the proximity to the United States. Many Canadians cross the border to the United States for business or pleasure, and a significant difference in gasoline taxes makes it worthwhile to refuel in the U.S. before returning to Canada. In our case, 'carbon leakage' measures gasoline purchases shifted from Canada to neighbouring US states due to higher fuel and carbon taxes.

Figure 5: Cross-Border Trips, Washington State and B.C., 2001–2015

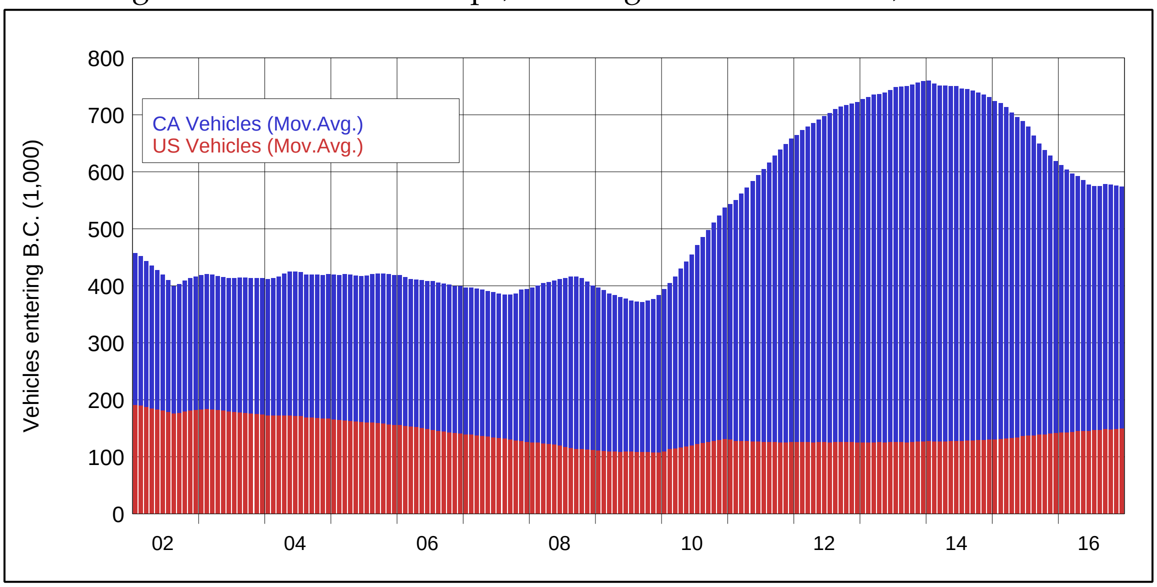

Source: Statistics Canada CANSIM Table 427-0002, and United States Bureau of Transportation Statistics. CA vehicles (top, blue) are returning to Canada, while US vehicles (bottom, red) are visiting Canada. Monthly data are shown as a 12-month moving average

In Figure 5 we illustrate cross-border travel from Washington State to British Columbia in recent years. The upper (blue) curve are actual crossings, and the lower (red) curve is US traffic. We use both data from Statistics Canada about vehicles entering (or returning to) Canada as well as data from the US Bureau of Transportation Statistics that shows the number of passenger vehicles entering Washington State (denoted by symbol b_{it} in our following analysis). Figure 5 shows that cross-border trips increased from an average of about 400,000 per month before 2009 to an average of over 750,000 by 2013. Peak months in 2013 and 2014 recorded more than 900,000 border crossings. The volume of border traffic has been receding in 2014 and 2015 as the Canadian Dollar has dropped below purchasing power parity.

To estimate gasoline demand in Canada, we build on Li et al. (2014) where a cross-country panel of d_{it} (per-capita demand for gasoline in Province i and month

t) is estimated by disaggregating the gross price of gasoline (p_{it}) in province i and month t into its net price (p^0_{it}) and the tax component ($p_{it} = p^0_{it} + \tau_{it}$). Additional regressors include per-capita disposable income y_{it} and province (μ_i) and time (ν_t) fixed effects. In some specifications we have also used more exhaustive province-year fixed effects.

Our estimating equation (13) differentiates two price components: (i) the volatile wholesale price net of taxes p^0_{it}; and (ii) the relatively stable tax component τ_{it}. Formally, p_{it} is decomposed into the product $p^0_{it}(1 + \tau_{it}/p^0_{it})$ and each term is estimated separately. The assumption underlying Li et al. (2014) is that consumers have a greater reaction to the price component perceived as permanent and long-term (the taxes) than the price component perceived as volatile and short-term (the wholesale price). This also follows directly from our theory presented earlier.

This decomposition, as proposed in Li et al. (2014), is desirable from an econometric point of view. The noise in the gasoline price data, attributable to the variations in the price of crude oil, is absorbed in our short-term (uncertainty driven) wholesale price variable p^0 rather than the long-term (certainty-based) tax variable $(1 + \tau/p^0)$. In Appendix A we demonstrate this formally. We show how the variance of the the tax mark-up $(1+\tau/p^0)$ is a fraction of the variance of p^0.[8] The decomposition into transitory and permanent effects is also rooted in the literature on investment under uncertainty (Dixit and Pindyck, 1994). New vehicle purchases are made in an environment where gasoline prices fluctuate significantly. Buyers of new cars observe separate signals from the tax component, which is certain, and the wholesale gas price, which is volatile. The different magnitudes of the elasticities $|\beta| < |\gamma|$ are, in part, attributable to the effect of making buying decisions under uncertainty or certainty. Similarly, decisions to switch transportation modes, move closer to work, or make other permanent adjustments to driving behaviour, are all made more readily when prices are more predictable.

To capture carbon/gasoline leakage, we first adopt a naïve approach and modify Li et al. (2014) to include cross border travel, resulting in

$$\ln(d_{it}) = \alpha + \beta \ln(p^0_{it}) + \gamma \ln\left(1 + \frac{\tau_{it}}{p^0_{it}}\right) + \delta \ln(b_{it}) + \theta \ln(y_{it}) + \mu_i + \nu_t + \epsilon_{it}, \qquad (13)$$

where b_{it} are counts for all cars crossing over from Canada to the US from Province i in month t. This measure includes Canadian vehicles visiting the United States as well as US vehicles visiting Canada. As figure 5 shows for British Columbia, Canadian vehicles dominate the picture, and the effect of the strong Canadian Dollar during 2011-2014 is clearly visible.

[8] Note that the tax variable can be decomposed further into $\tau = \tau^f(1 + \tau^s) + p^0\tau^s$, where τ^f is a fuel tax that is fixed by volume (per litre) and τ^s is a value-added tax rate (GST, HST, or QST). Therefore, the numerator in τ/p^0 also varies with p^0 due to the presence of a sales tax component, and this in turn offsets some of the variation introduced by having p^0 in the denominator. Alternative forms of decomposition are econometrically less desirable, including data smoothing. Smoothing the time series τ_t would not improve our estimation strategy because it would eliminate the variation in the numerator of τ/p^0 that offsets the variation in the denominator. Smoothing would also obscure the immediacy of fuel and sales tax changes, thus biasing our estimates of the tax effect downwards significantly.

We first present the results of our naïve approach in Table 1 where we do not account for endogeneity of cross-border travel. Columns (A) and (B) include data from seven Canadian provinces, excluding Prince Edward Island, Nova Scotia, and Newfoundland & Labrador—as they do not have all-season border traffic to the United States. The specifications in the first columns differ due to different fixed effects included. In the first column we include province, month, and year fixed effects. In the second column we exclude the year fixed effects. We consider the set of fixed effects in column (A) as most extensive and consider it our preferred specification. In column (A) we observe that gasoline prices have a negative effect on gasoline demand. The impact of gasoline taxes is also negative, and also significantly larger than the effect of gasoline prices. This is consistent with the result in Li et al. (2014). We also observe that as more cars travel south to the US from BC, fuel demand falls.

Table 1: Estimation Results: Li/Lin/Muehlegger Model

	(A)		(B)	
Log net price	-0.301^c	(7.66)	-0.252^c	(8.77)
Log all taxes	-0.952^c	(8.48)	-0.616^c	(6.84)
Log per-capita income	0.419^c	(6.14)	0.307^c	(7.97)
Log border crossings	-0.045^c	(3.75)	-0.033^b	(2.85)
Regression fit (R^2)	0.810		0.800	
Observations	1,259		1,259	
Province Fixed Effects	yes		yes	
Month Fixed Effects	yes		yes	
Year Fixed Effects	yes		no	

Note: The dependent variable in the models above is the log of per-capita gasoline demand in each region and each month. Statistical significance at the 95%, 99%, and 99.9% confidence levels is indicated by the superscripts ^{a}, ^{b}, and ^{c}, respectively. Absolute values of t-statistics are reported in parentheses. All models include a constant, which is not reported.

The main problem problem with the specification presented in Table 1 is endogeneity. In the presence of carbon leakage, gasoline prices and taxes also influence cross-border travel. A lower price for gasoline makes it cheaper to travel to the border, but a higher price differential between the domestic and foreign gasoline price increases the propensity to cross the border for refuelling. Higher domestic taxes increase this price differential and raise the number of border crossings—especially in jurisdictions close to the border. The remaining covariates from equation (13)—per-capita disposable income, province fixed effects (capturing the distance to the border), and time fixed effects (capturing seasonality)—also influence cross-border travel. This endogeneity impacts the coefficients for cross-border travel, and the coefficient on our main variable of interest: gasoline taxes.

Theoretical and empirical work in Chandra et al. (2014) explains cross-border travel. Their research identifies deviations of nominal exchange rates from purchasing power parity as the key driver for increased cross-border traffic.[9] Based on their study, we instrument for cross-border travel based on measures that capture the deviations of the exchange rate from the PPP rate. We follow Chandra et al. (2014) by allowing for a non-linear response of cross-border travel to the arbitrage opportunities that emerge. Cross-border shopping becomes more attractive when the deviations are sufficiently large to compensate for travel time and cost. Moreover, we allow this response to be asymmetric because of the different magnitudes of cross-border travel. Canadians tend to take advantage of cross-border shopping opportunities in the United States in larger numbers than Americans tend to take advantage of cross-border shopping opportunities in Canada.

One of our variables included in our regressions is the net (wholesale) gasoline prices. Gasoline prices in Canada are determined by the international price of crude oil, and the exchange rate. Further complicating things, the exchange rate is strongly influenced by the price of crude oil (Ferraro et al., 2015). This creates severe issues of collinearity for both our first stage and second stage regressions. Thus we jointly instrument both cross-border crossings and the net price of gasoline in equation (13) using four variables:

- the deviation of the nominal exchange rate from PPP, captured by $\ln(x_t/\bar{x})$, where x_t is the nominal exchange rate (in CAD/USD) in month t and $\bar{x}$ is the (average) PPP rate of 1.20 CAD/USD;
- the non-linear effect of a strong Canadian Dollar, captured by the squared positive deviation $\max(0, \ln(x_t/\bar{x}))^2$;
- the non-linear effect of a weak Canadian Dollar, captured by the squared negative deviation $\max(0, \ln(\bar{x}/x_t))^2$; and
- the log price of crude oil (expressed in USD) using the West Texas Intermediate (WTI) benchmark that is relevant for North America.

The two squared exchange rate terms can be expected to have positive signs because of the non-linear asymmetric effect explained above. The linear exchange rate term can have either sign and simply ameliorates the functional fit of the non-linear asymmetry. The fourth instrument is not particularly relevant for cross-border travel; it captures the world price of crude oil.[10]

Our instrumental variable (IV) results are presented in table 2, where we report the two first-stage regressions along with the second-stage regression. This first stage for cross-border crossing modifies Chandra et al. (2014) to include gasoline prices and taxes. It is also our first credible test for carbon leakage. If cross-border travel is influenced by fuel and carbon taxes, fuel demand and carbon emissions

[9] During our sample period, the US Dollar to Canadian Dollar exchange rate appreciated from approximately 0.8 USD/CAD in early 2009 to nominal parity (1.0) in 2011-12. It then fell to 0.9 in 2014 and 0.8 by early 2015 along with falling oil and commodity prices. In other words, 2010–2014 is characterized by an overvaluation of the Canadian Dollar relative to the level of purchasing power parity (PPP), which according to the OECD is approximately 0.833 USD/CAD.

[10] Our instruments satisfy conditional exogeneity. The exchange rate does not influence gasoline demand directly, only through cross-border travel. The exchange rate influences gasoline prices (as crude oil and gasoline are an internationally traded good), but prices are included in equation 13, thus assuring conditional exogeneity.

are exported, which is the definition of carbon leakage.

There are two first stage regressions listed. Regressions in Columns (A) and (B) include seven Canadian provinces. Consider first the regression panel where the dependent variable is: log Border Crossings. During the time period of investigation, fuel taxes in the US remained flat and significantly lower than in Canada, making Canadian taxes a good proxy for the cross-border price gap. In both specifications, the effect of taxes is strongly positive and significant. Specifications (A) demonstrates a significant impact for the exchange rate variables. We also find a positive effect from the crude oil price, but this effect is consistently weaker than the effect from taxes. We interpret our first stage results as credible evidence for demand leakage from gasoline taxes, as cross-border trips are not only affected by exchange rates (and thus the demand for cross-border shopping and travel), but also by opportunities for refuelling. In other words, the carbon tax, as one element in the list of fuel taxes, has contributed to increased cross-border travel. This effect is tantamount to demand-side *carbon leakage* in the sense that a carbon policy leads to the substitution of foreign with domestic products (i.e., US gasoline replaces Canadian gasoline).

The second first-stage regressions instruments for the wholesale price of gasoline (dependent variable: log wholesale gasoline price). The variables capturing exchange rates and the crude oil price are essential. Clearly, the effect of crude oil prices on wholesale gasoline prices is positive, but the elasticity is less than unity because the "petrodollar effect" compensates through appreciation of the Canadian Dollar. The exchange rate elasticity is therefore negative (as a rising CAD/USD exchange rate amounts to a depreciation of the Canadian Dollar).

Our second-stage results (preferred regression column A) indicate that both gasoline prices and taxes negatively and statistically significantly influence gasoline demand. A 1% increase in (gasoline taxes/net price) lower gasoline demand by approximately 1.4%. Thus our preferred estimate for the fuel tax elasticity of fuel demand $\eta_F = -1.439$. However, while we see evidence of carbon leakage in the first stage, in the second stage we find that border crossings to the US do not affect overall fuel demand (the coefficient is positive and statistically insignificant). This implies that while people tend to cross the border due to increases in gasoline taxes, the number of people crossing, and the amount of fuel displaced due to this crossing is likely not large enough to influence overall demand. On balance, our IV results underline the effectiveness of fuel and carbon taxes.

6.2 Fuel Efficiency and Vehicle Fleet Composition

In this subsection we analyze how gasoline prices and taxes influence the fuel efficiency of new vehicles sold across provinces—the investment effect. Our aim is to derive an estimate of η_f (see Section 3), the elasticity of fuel efficiency in the vehicle fleet from gasoline taxes.

We analyze the impact of gasoline prices and taxes on the sales-weighted fuel efficiency of the average vehicle sold in each province. Let the subscript $m \in \{1, ..., M\}$ index vehicle models, superscript $i \in \{1, ..., 9\}$ index provinces,

Table 2: Instrumental Variables Extension of LLM Model

	(A)		(B)	
First-stage (OLS) — Dependent variable: log border crossings				
Log all taxes	1.471^c	(7.73)	1.741^c	(10.0)
Log per-capita income	1.275^c	(8.35)	1.474^c	(17.8)
Log CAD-USD rate	-0.863^a	(2.39)	-0.690^b	(3.07)
CAD strength	3.293^a	(2.16)	4.009^c	(4.10)
CAD weakness	-2.825^a	(2.13)	0.671	(1.01)
Log WTI oil price	0.516^c	(7.70)	0.533^c	(10.6)
Regression fit (R^2)	0.988		0.987	
First-stage (OLS) — Dependent variable: log wholesale gasoline price				
Log all taxes	-1.683^c	(32.0)	-1.635^c	(26.2)
Log per-capita income	-0.355^c	(8.42)	0.456^c	(15.3)
Log CAD-USD rate	-0.291^b	(2.93)	-1.018^c	(12.6)
CAD strength	0.054	(.128)	2.778^c	(7.91)
CAD weakness	0.972^b	(2.65)	-1.764^c	(7.38)
Log WTI oil price	0.305^c	(16.5)	0.427^c	(23.6)
Regression fit (R^2)	0.977		0.959	
Second-stage (IV) — Dependent variable: log per-capita fuel demand				
Log net price	-0.485^b	(2.65)	-0.363^c	(4.25)
Log all taxes	-1.439^b	(3.08)	-0.920^c	(3.77)
Log per-capita income	0.142	(.710)	0.202^c	(3.58)
Log border crossings	0.123	(1.24)	0.070	(1.48)
Regression fit (R^2)	0.780		0.787	
Observations	1,259		1,259	
Province Fixed Effects	yes		yes	
Month Fixed Effects	yes		yes	
Year Fixed Effects	yes		no	

Note: The dependent variable in the models above is the log of per-capita gasoline demand in each region and each month. Statistical significance at the 95%, 99%, and 99.9% confidence levels is indicated by the superscripts ^{a}, ^{b}, and ^{c}, respectively. Absolute values of t-statistics are reported in parentheses. All models include a constant, which is not reported. The border crossing variable has been instrumented with the two exchange rate variables.

and superscript $t \in \{2001, ..., 2014\}$ index years. Let n_{mit} denote the sales count for model m sold in province i in year t. Finally, let N_{it} be the sum of sales of all models sold in province i in year t. The sales-weighted average fuel consumption in litres per 100 kilometres FE_{it} for new vehicles sold in province i in year t is thus,

$$\mathrm{FE}_{it} = \left[\frac{\sum_{m=1}^{M} n_{mit} \mathrm{FE}_{mt}}{N_{it}} \right]. \tag{14}$$

where FE_{mt} is the published fuel consumption for model m in year t in litres per 100 kilometres traveled (our data is from the US Environmental Protection Agency). We use the sales-weighted fuel efficiency (and its geometric and log transformed variant) as dependent variables in the following cross-provincial panel specification,

$$\ln(\mathrm{FE}_{it}) = \alpha + \beta p_{it}^{0} + \beta \tau_{it} + \gamma v_{it} + \theta y_{it} + \mu_i + \nu_t + \epsilon_{it} \tag{15}$$

where p_{it}^{0} is the tax-exclusive net price for gasoline, τ_{it} is the gasoline tax in cents, v_{it} is vehicles per capita, y_{it} is real per capita income, μ_i are a set of provincial fixed effects, ν_t are a set of year fixed effects, and ϵ_{it} is idiosyncratic error.

Results corresponding to equation (15) are shown in Panel I of Table 3. In the columns (A) and (B) we include only provincial fixed effects, while in the last column (C) we also include year fixed effects in the regression. The coefficient on gasoline taxes is negative and significant through all three specifications. According to column B, our preferred specification, a one-cent increase in the gasoline tax lowers the sales weighted fuel consumption in a province by 0.05%. For context, Canada's sales-weighted new vehicle fuel consumption in the year 2014 was 10.18 litres/100km.

We also estimate the log-log version that follows the Li et al. (2014) specification we employed earlier in our paper. Thus,

$$\ln(\mathrm{FE}_{it}) = \alpha + \beta \ln(p_{it}^{0}) + \delta \ln\left(1 + \frac{\tau_{it}}{p_{it}^{0}}\right) + \gamma \ln(v_{it}) + \theta \ln(y_{it}) + \mu_i + \nu_t + \epsilon_{it} \tag{16}$$

Results are presented in Panel II of Table 3. In this specification both gasoline price (net of taxes), and gasoline taxes have a consistently significant and negative effect on the fuel consumption of the provincial new vehicle fleet. Once again we consider column B—with provincial effects alone—as our preferred specification. According to this specification, the elasticity of provincial new fleet fuel consumption to changes in the net gasoline price is –0.28, that is a 1% increase in the net price of gasoline reduces new fleet fuel consumption by 0.28%. In contrast the same number for all fuel taxes is –0.79, that is a 1% increase in gasoline tax reduces new fleet fuel consumption by 0.79%. Note that the tax effect is over three times that deriving from equivalent but transitory price changes. [11]

[11] There is an alternative and indirect method to explore the effect of fuel taxes on vehicle purchases following the approach pioneered by Klier and Linn (2010, 2011). This involves estimating

Table 3: Fuel Efficiency Regressions

	(A)		(B)		(C)	
Panel I: Linear Specification						
Net gasoline price [¢/L]	-0.001^c	(7.19)	-0.001	(1.47)	-0.000	(.867)
All fuel taxes [¢/L]	-0.004^c	(3.34)	-0.005^c	(4.33)	-0.001^a	(2.51)
Real household income			-0.000	(1.35)	0.000^c	(4.61)
Vehicles per capita			-0.246	(1.80)	0.125^b	(2.78)
Constant	2.642^c	(91.2)	2.858^c	(39.1)	2.337^c	(58.0)
R^2	0.850		0.863		0.989	
Panel II: Logarithmic Specification						
log(Net Gas Price)	-0.269^c	(8.03)	-0.275^c	(8.39)	-0.090^a	(2.54)
log(All Gas Taxes)	-0.759^c	(6.45)	-0.792^c	(6.87)	-0.135^b	(2.65)
log(household income)	-0.190^c	(3.70)	-0.087	(1.37)	0.123^c	(5.20)
log(vehicles/population)			-0.209^b	(2.66)	0.039	(1.45)
Constant	5.778^c	(10.8)	4.670^c	(7.00)	1.665^c	(5.21)
R^2	0.863		0.871		0.990	
Observations	126		126		126	
Provincial Fixed Effects	yes		yes		yes	
Year Fixed Effects	no		no		yes	

Note: The dependent variable in the models above is the sales-weighted fuel efficiency (in L/100km) in Panel I, and its log transformation in Panel II. Statistical significance at the 95%, 99%, and 99.9% confidence levels is indicated by the superscripts ^{a}, ^{b}, and ^{c}, respectively. Absolute z-scores are reported in parentheses.

Now we are ready to use the decomposition from equation 7. Our preferred estimate for η_F—the overall effect is -1.439, and our preferred estimate for η_f, the investment effect, is -0.792. Given our assumptions, the fuel tax elasticity of VMT η_m is then $\eta_F - \eta_f = -0.647$. This is the fuel tax elasticity for VMT in Canada. The proportion of overall reduction from a reduction of VMT is approximately 0.45. This is similar to that quoted by Parry and Small (2005), who cite other studies to argue that "Empirical studies suggest, however, that probably less than half of the long-run price responsiveness of gasoline consumption is due to changes in VMT."

---

the impact of taxes on the market share for individual models sold in the region. This approach estimates the average effect of fuel taxes on market shares, and also interacts this with the fuel consumption of vehicle makes and models. The two estimates, intercept and slope of the tax effect, imply that a fuel tax increase lowers the market share of fuel-inefficient vehicles and increases the market share of fuel-efficient vehicles. There are caveats to this approach that make it less useful than the direct fleet-wide estimate reported above. It is usually necessary to aggregate models or makes into sufficiently "wide" bins to estimate coefficients reliably, and this binning may obscure minor shifts within each bin in response to tax changes. Furthermore, estimating market shares faces econometric challenges in particular in the presence of spatial heterogeneity. We are not reporting results of these approaches in this paper as these approaches will be part of a related forthcoming paper.

7 Conclusions

We present a simple method estimating the elasticity of fuel tax on vehicles miles traveled. This elasticity is probably the most important component of an optimal gasoline tax, and its estimation is crucial for setting policy. Our method allows researchers to use aggregate data for estimating this elasticity. Thus far, individual driving data has been used to estimate this elasticity. However, due to the high cost of collection, representative individual data is typically unavailable for large jurisdictions. Further, there are significant privacy concerns around its use, making access complicated. A simple decomposition allows us to demonstrate the relationship between the aggregate fuel demand fuel tax elasticity, fleet fuel efficiency elasticity, and the vehicle miles traveled fuel tax elasticity. The aggregate fuel demand elasticity, and fleet fuel efficiency elasticity are typically much easier to estimate. The data series needed to estimate them are usually easier to obtain and aren't subject to the same temporal and geographical constraints that individual driving has.

For such an estimation, our ideal data requirements are: fuel consumption data, and fleet turnover data for coincident years. In our illustration, we use fuel consumption and new vehicle sales data from Canadian provinces for years 2001-2014. However, having coincident data is not necessary, especially if we have reason to believe that the elasticities do not change much across time. Ideally the researcher would have access to fleet turnover data, new vehicle and old vehicle sales data, and all retirements from the fleet. This would give an accurate picture of fleet fuel efficiency in the jurisdiction of interest. However, we discuss assumptions that allow us estimate a long run elasticity using only new vehicle sales data. We assume that eventually, old vehicles exit the fleet and new vehicle sales make up the fleet in the long run. In other words, one could make do with less than ideal fleet data as well.

References

Allcott, H., Wozny, N., 2014. Gasoline prices, fuel economy, and the energy paradox. The Review of Economics and Statistics 96, 779–795. doi:10.1162/REST_a_00419.

Banfi, S., Filippini, M., Hunt, L.C., 2005. Fuel tourism in border regions: The case of switzerland. Energy Economics 27, 689–707. doi:10.1016/j.eneco.2005.04.006.

Barla, P., Couture, E., Samano, M., 2016. Gasoline prices and fuel economy of new vehicles in quebec. Canadian Public Policy 42, 181–193. doi:10.3138/cpp.2015-076.

Barla, P., Lamonde, B., Miranda-Moreno, L.F., Boucher, N., 2009. Traveled distance, stock and fuel efficiency of private vehicles in canada: price elasticities and rebound effect. Transportation 36, 389–402. doi:10.1007/s11116-009-9211-2.

Basso, L.J., Oum, T.H., 2007. Automobile fuel demand: A critical assessment of empirical methodologies. Transport Reviews 27, 449–484. doi:10.1080/01441640601119710.

Brons, M., Nijkamp, P., Pels, E., Rietveld, P., 2008. A meta-analysis of the price elasticity of gasoline demand. a sur approach. Energy Economics 30, 2105–2122. doi:10.1016/j.eneco.2007.08.004.

Busse, M.R., Knittel, C.R., Zettelmeyer, F., 2013. Are consumers myopic? evidence from new and used car purchases. American Economic Review 103, 220–256. doi:10.1257/aer.103.1.220.

Chandra, A., Head, K., Tappata, M., 2014. The economics of cross-border travel. Review of Economics and Statistics 96, 648–661. doi:10.1162/REST_a_00404.

Dahl, C., Sterner, T., 1991. Analysing gasoline demand elasticities: a survey. Energy economics 13, 203–210.

Davis, L.W., Kilian, L., 2011. Estimating the effect of a gasoline tax on carbon emissions. Journal of Applied Econometrics 26, 1187–1214.

Dixit, A., Pindyck, R., 1994. Investment under uncertainty. Princeton University Press.

Dorval, J., Barla, P., 2017. Does quebec have the right gasoline tax? an empirical investigation. Canadian Public Policy 43, 350–362. doi:10.3138/cpp.2016-065.

Feng, Y., Fullerton, D., Gan, L., 2005. Vehicle choices, miles driven, and pollution policies. Technical Report. National Bureau of Economic Research.

Ferraro, D., Rogoff, K., Rossi, B., 2015. Can oil prices forecast exchange rates? an empirical analysis of the relationship between commodity prices and exchange rates. Journal of International Money and Finance 54, 116–141.

Fullerton, D., West, S.E., 2002. Can taxes on cars and on gasoline mimic an unavailable tax on emissions? Journal of Environmental Economics and Management 43, 135–157.

Gillingham, K., 2014. Identifying the elasticity of driving: evidence from a gasoline price shock in california. Regional Science and Urban Economics 47, 13–24.

Graham, D.J., Glaister, S., 2002. The demand for automobile fuel: A survey of elasticities. Journal of Transport Economics and Policy 36, 1–25. URL: http://www.jstor.org/stable/20053890.

Klier, T., Linn, J., 2010. The price of gasoline and new vehicle fuel economy: Evidence from monthly sales data. American Economic Journal: Economic Policy 2, 134–153. doi:10.1257/pol.2.3.134.

Klier, T., Linn, J., 2011. Fuel prices and new vehicle fuel economy in europe. Resources for the Future Discussion Paper .

Li, S., Linn, J., Muehlegger, E., 2014. Gasoline taxes and consumer behavior. American Economic Journal: Economic Policy 6, 302–342. doi:10.1257/pol.6.4.302.

Li, S., Timmins, C., von Haefen, R.H., 2009. How do gasoline prices affect fleet fuel economy? American Economic Journal: Economic Policy 1, 113–137. doi:10.1257/pol.1.2.113.

Nicol, C., 2003. Elasticities of demand for gasoline in canada and the united states. Energy Economics 25, 201–214. doi:10.1016/S0140-9883(03)00002-1.

Parry, I.W., Small, K.A., 2005. Does britain or the united states have the right gasoline tax? The American Economic Review 95, 1276–1289. doi:10.1257/0002828054825510.

Pindyck, R.S., 1999. The long-run evolution of energy prices. The Energy Journal 20, 1–27. doi:10.5547/ISSN0195-6574-EJ-Vol20-No2-1.

Schwartz, E.S., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. Journal of Finance 52, 923–973. doi:10.1111/j.1540-6261.1997.tb02721.x.

Sipes, K.N., Mendelsohn, R., 2001. The effectiveness of gasoline taxation to manage air pollution. Ecological Economics 36, 299–309. doi:10.1016/S0921-8009(00)00230-5.

A Price Variance Decomposition

We follow Li et al. (2014) in decomposing the final sales price p of gasoline into a wholesale price p^0 and a tax-mark up $(1 + \tau/p^0)$ so that $p = p^0(1 + \tau/p^0)$ and thus $\ln(p) = \ln(p^0) + \ln(1 + \tau/p_0)$. We are specifically interested in how this decomposition affects the attribution of variance, where we assume that p^0 is a random variable and $\tau = \tau^f + (\tau^f + p^0)\tau^s$ can be decomposed into a fixed part τ^f that is independent of the wholesale price, and a sales tax part τ^s that varies the total amount of taxes as the wholesale price changes. Discrete changes in the tax rates τ^f and τ^s also introduce variation in the data, but we assume that this process is not stochastic but deterministic.

For non-linear transformations of random variables there are often no closed-form transformations for the variance. Instead, it is common to use first-order or second-order Taylor series approximations so that for a transformation $f(x)$ of random variable x, the variance is given by the approximation

$$\mathrm{Var}(f(x)) \approx [f'(\mu)]^2 \mathrm{Var}(x) + \frac{1}{4}[f''(\mu)]^2 \mathrm{Var}^2(x) + \cdots$$

where $\mu = \mathrm{E}(x)$ is the expected value of x. We will ignore second-order and all higher-order effects in what follows. Therefore, for the log transformation of the wholesale price, we find that

$$\mathrm{Var}(\ln(p^0)) \approx \frac{\mathrm{Var}(p^0)}{(\bar{p}^0)^2}$$

where $\bar{p}^0$ is the mean wholesale price. Furthermore, note that

$$\ln\left(1 + \frac{\tau^f + (\tau^f + p^0)\tau^s}{p^0}\right) = \ln(1 + \tau^s) + \ln\left(1 + \frac{\tau^f}{p^0}\right)$$

and therefore

$$\mathrm{Var}(\ln(1 + \tau/p^0)) \approx \frac{\mathrm{Var}(p^0)}{(\bar{p}^0)^2}\left[\frac{\bar{\tau}^f}{\bar{p}^0}\right]^2$$

where $\bar{\tau}^f$ and $\bar{p}^0$ are the mean fixed tax rate the mean wholesale price. Combining the previous results, the Li et al. (2014) decomposition ensures that

$$\frac{\mathrm{Var}(\ln(1 + \tau/p^0))}{\mathrm{Var}(\ln(p^0))} \approx \left[\frac{\bar{\tau}^f}{\bar{p}^0}\right]^2$$

Consider a typical wholesale price of 76 cents, and a fixed tax component of 38 cents, which implies a ratio of $\bar{\tau}^f/\bar{p}^0 = 1/2$. Then the variance of the tax measure is approximately only one quarter of the variance of the wholesale price. In other words, the method of decomposition assures that the tax component, which is meant to capture the permanent price element, has much less variance than the price component.

B Fuel Taxes and and Fuel Consumption Elasticities

In this mathematical appendix we develop the details of our theoretical model of repeated purchases of a vehicle with fuel consumption f and amenity level a, expressed in the value function equation (8), combined with a price shock with mean reversion process (9). First we introduce a series of lifetime multipliers (all of which are positive) that will simplify our calculations:

$$\psi^{\gamma} \equiv \frac{e^{(\gamma-\rho)T}-1}{\gamma-\rho} \quad \text{and} \quad \psi^{0} \equiv \frac{1-e^{-\rho T}}{\rho} \quad \text{and} \quad \psi^{\lambda} \equiv \frac{1-e^{-(\rho+\lambda)T}}{\rho+\lambda} \tag{17}$$

Note in particular that $\psi^0 > \psi^\lambda$ as long as $\lambda > 0$; this can be seen most easily in the limit as $\lim_{T\to\infty}(\psi^0 - \psi^\lambda) = 1/(\rho(1+\rho/\lambda)) > 0$. Solving equation (8) for V and using the above definitions yields

$$\rho V = fm\left[\bar{p} + (\hat{p}-\bar{p})\frac{\psi^\lambda}{\psi^0}\right] + \frac{K + c\psi^\gamma}{\psi^0} \tag{18}$$

The rising fuel and maintenance cost will trigger vehicle replacement at time

$$T^* = \text{argmin}_{T>0}(V) \tag{19}$$

which is thus endogenous. Further define

$$\kappa \equiv \frac{K(f,a)}{fm\bar{p}(\psi^0 - \psi^\lambda)} \tag{20}$$

as the ratio of fixed costs to lifetime fuel costs. We are interested in the fuel consumption elasticity $\hat{\phi}$ and $\bar{\phi}$ with respect to changes in the temporary price $\hat{p}$ and the permanent price $\bar{p}$, respectively. We define the two corresponding elasticities $\hat{\phi}$ and $\bar{\phi}$ in (11). We also introduce the elasticity ζ of vehicle price with respect to fuel efficiency, holding vehicle amenities constant at $\bar{a}$:

$$\zeta \equiv \frac{\partial K(f,\bar{a})}{\partial f}\frac{f}{K(f,\bar{a})} < 0 \tag{21}$$

Making a vehicle more fuel efficient comes at an extra cost while holding the vehicle's amenities $\bar{a}$ constant. Through implicit differentiation of $V(f,p) = \bar{V}$ it follows that

$$\bar{\phi} = -\left[\frac{\partial V/\partial \bar{p}}{\partial V/\partial f}\right]\left[\frac{\bar{p}}{f}\right] = -\frac{1}{1+\kappa\zeta+\delta} < 0 \tag{22}$$

and

$$\hat{\phi} = -\left[\frac{\partial V/\partial \hat{p}}{\partial V/\partial f}\right]\left[\frac{\bar{p}}{f}\right] = -\frac{\delta}{1+\kappa\zeta+\delta} < 0 \tag{23}$$

where we introduce $\delta \equiv (\hat{p}/\bar{p})(\psi^\lambda/(\psi^0-\psi^\lambda)) > 0$. Dividing the two elasticities and setting permanent and transitory prices equal produces the positive ratio

$$\omega \equiv \left.\frac{\bar{\phi}}{\hat{\phi}}\right|_{\hat{p}=\bar{p}} = \left.\frac{1}{\delta}\right|_{\hat{p}=\bar{p}} = \left[\frac{\psi^0}{\psi^\lambda} - 1\right] > 0 \tag{24}$$

The elasticity ratio thus depends on three variables only: the discount rate ρ, the price-shock decay rate λ, and the vehicle lifespan T.

Quantifying κ and ζ are essential for gaining insights into the effective response of consumers. It is quite reasonable to assume that $|\zeta| \ll 1$ for most vehicles; making them 10% more fuel efficient increases purchase prices far less than 10%. Furthermore, the ratio κ depends obviously on driver type. It is low for high-mileage motorists and high for low-mileage motorists. This effect may be exaggerated if high-mileage motorists drive cheaper vehicles than low-mileage motorists. On the other hand, the size of ζ may also be correlated with κ so that the elasticity ζ is smaller in magnitude for higher-price vehicles with higher κ. The magnitude of ϕ depends crucially on the magnitude of $\kappa\zeta$, the product of the capital-cost to fuel-cost ratio, and the fuel-price elasticity of the vehicle's fuel consumption. Note that $\zeta < 0$, and for the denominator to be positive, we require $1 + \delta > \kappa|\zeta|$. Consider a simple numerical example where m is 12,500km, f is 10L/100km, the fuel price is $1/L, and the capital cost is $36,000. Further assume that $\zeta = -0.1$ and $\delta = 0.3$, while $\psi^0 - \psi^\lambda = 6$. Then $\kappa = 4.8$. This generates an elasticity of $\bar{\phi}$ of –1.2. This is to show that plausible parameter assumptions get us into the vicinity of parameter estimates that we observe in our data.

Our theoretical model hinges on the empirical magnitude of λ, the mean-reversion speed. There is an extensive literature on estimating mean reversion in commodity prices; see for example Pindyck (1999). It is possible to estimate the mean-reversion speed by discretizing the underlying stochastic Ornstein-Uhlenbeck process $dp_t = \lambda(\mu - p_t)dt + \sigma dW$, where dW is the Wiener process. Discretization yields the first-order autoregressive process

$$p_t - p_{t-1} = \beta_0 + \beta_1 p_{t-1} + \epsilon \tag{25}$$

Estimation by ordinary least squares uncovers $\hat{\lambda} = -\beta_1/\Delta t$ and $\hat{\mu} = -\beta_0/\beta_1$, where $\hat{\mu}$ is the estimate of $\bar{p}$. In regressions for weekly gasoline wholesale price data in Canadian cities for the 2001-2015 period we find $\hat{\mu}$ in the 70¢-region and $\hat{\lambda}$ in the 0.9 range, sufficiently high to justify the conclusion that a large ω is entirely plausible.