Step on It: Approaches to Improving Existing Vehicles' Fuel Economy

Ashley Langer* and Shaun McRae†

December 31, 2014

Abstract

In this paper we attempt to understand how policy interventions could improve the fuel economy of drivers without changing the vehicles they drive or the trips they take. To do this, we use high-resolution driving data to document the large variation in actual on-road fuel economy achieved by drivers of identical vehicles. We analyze this variation using a model that links driver behavior to vehicle fuel consumption. The physical vehicle model, estimated econometrically from data on acceleration and speed choices, predicts observed fuel consumption extremely well. We combine this physical model with a model of optimizing behavior by drivers, in which drivers trade-off higher fuel consumption for shorter trip times. Using the variation in value of time for the drivers in our data set, we simulate route choice and driving behavior on a stylized set of trips. We find that drivers have little incentive to change their driving behavior to improve fuel economy, since the time cost of driving more efficiently generally outweighs the value of fuel savings. This means that gasoline taxes, while internalizing environmental externalities, are not as effective at reducing overall fuel use as policies that allow for smoother traffic flow (fewer accelerations and decelerations) such as stoplight timing, infrastructure improvements, or the increased use of vehicle automation.

*Department of Economics, University of Arizona. Email: alanger@email.arizona.edu

†Department of Economics, University of Michigan. Email: sdmcrae@umich.edu.

1 Introduction

In 2010, the United States consumed nearly 350 millions gallons of gasoline per day.[1] In addition to the direct cost of gasoline, this consumption polluted the air and contributed to climate change. To address these externalities, the US both taxes gasoline and implements Corporate Average Fuel Economy (CAFE) standards, which proscribe the average fuel economy each car manufacturer's sales must meet. While these policy approaches address the margins of which vehicles people drive and how many miles they drive, relatively little attention has been paid to policies that might change *how* people drive, given identical vehicles and trip requirements, in order to reduce fuel use. In this paper, we focus on understudied approaches to reducing gasoline consumption by focusing on ways to improve the fuel economy of a given driver on a given trip.

While some papers have attempted to understand whether drivers change their behavior in response to changing gasoline prices (Wolff (2014), Burger and Kaffine (2009)), these approaches have been hampered by the lack of data on individual drivers' behavior over time. We use an unusual dataset to begin to understand both the effect of driver behavior on gasoline consumption in real-world driving situations and the benefits to drivers and society of different policy interventions that change on-road fuel economy. Our data come from the University of Michigan's Transportation Research Institute (UMTRI). UMTRI gave 108 drivers nearly identical mid-size sedans to drive for approximately 6 weeks each.[2] These vehicles were equipped with extensive data collection systems that recorded hundreds of variables including speed, location, heading, windshield wiper use, instantaneous fuel consumption, and radar that recorded the presence of nearby vehicles. Each variable was recorded at a frequency of at least 10 times per second.

This dataset allows us to document interesting variation across drivers driving the same make and model vehicles. Drivers' average fuel economy in our sample ranges from 8.1 to 13.6 liters per 100 kilometers (L/100 km) over the entire period that they have the vehicles.[3] In our data set, drivers average 8.1 L/100 km at highway speeds (above 100 km/h) and 10.6 L/100 km at non-highway speeds. The Environmental Protection Agency (EPA) estimates that these vehicles should get 9.05 L/100 km (26 mpg) on a highway trip and 13.07 L/100 km (18 mpg) on a city trip, which is roughly in line with what we observe in the data, but

[1] Energy Information Agency

[2] All of the vehicles were well-equipped Honda Accords, but since each vehicle was given to a series of drivers, the previous use of each vehicle varried across drivers.

[3] This corresponds to a range of 16.2 to 28.6 miles per gallon (mpg).

with substantial additional variation across trips.

Of course, these differences can arise for a variety of reasons, from the speed limit and quality of the road to the congestion on a given route or the aggressiveness of the driver. It is very difficult to estimate drivers' choice of routes since we do not know the routes that the driver considers valid alternatives for a given trip or the value to drivers of completing different trips. Therefore, we start by focusing on fuel use conditional on route choice and borrow a model from the engineering literature that models fuel consumption as a non-linear function of speed and acceleration. We estimate a simplified version of this model using second-by-second data on our actual vehicle fuel consumption, speed, and acceleration data, as well as observations of other environmental factors that affect fuel usage, such as temperature and road grade. The model fits our data well, with an R^2 of 0.93, meaning that our simple model predicts a substantial majority of the variation in fuel economy over time.

We then pair this model with a straightforward behavioral model of driver decision-making, where drivers maximize utility over the time it takes to complete a trip and the fuel cost of the trip, conditional on the driver's route. This allows us to consider drivers' incentives to change how they drive in response to either changing gasoline prices or a gasoline tax. Finally, we use the combination of the behavioral model and the physical model to simulate drivers' choices over two hypothetical routes under a variety of policy interventions: an increased gasoline tax, changing speed limits, and infrastructure and technology interventions that improve urban traffic flow.

Our physical model shows that fuel economy is relatively flat over a large range of speeds from about 50 to 110 km/h; there is an extremely sharp reduction in fuel economy for lower speeds and a relatively gentle reduction in fuel economy at higher speeds. Additionally, we investigate the relationship between fuel econsumption and different rates of acceleration and deceleration. At low speeds, driving over a fixed distance, there is a U-shaped relationship between acceleration and fuel use: very slow and very fast acceleration use more fuel than moderate acceleration.

Somewhat counter-intuitively, and counter to advice on the EPA's website,[4] we find that how a driver accelerates or the speed she drives on the freeway have relatively little impact on fuel economy. In fact, at \$3.50/gallon, increasing freeway speed from 110 kilometers per hour (km/h) to 120 km/h doesn't reduce fuel costs very much compared to the increase in trip time, and is actually only cost-effective for drivers with a value of time below \$7 per hour (excluding any safety or depreciation-related savings). Similarly, differences in fuel usage

[4]http://www.fueleconomy.gov/feg/drive.shtml

across different acceleration patterns are negligible, especially compared to the potential time savings from faster acceleration and deceleration.[5] This lack of a strong monetary incentive to drive more efficiently may help explain the mixed evidence in the literature about whether drivers change their driving behavior in response to changing gasoline prices.[6]

While the model suggests that highway driving speed and the average rate of acceleration and deceleration (that is, driving "aggressiveness") do not substantially affect the variation in fuel economy across drivers of identical vehicles, we show that the number of acceleration and deceleration events contributes considerably to average driver fuel economy. This means that drivers whose routes involve a lot of stopping and starting use substantially more fuel than drivers on similar roads who are able to maintain a more constant speed. This is true of both complete stops like those at traffic lights and substantial changes in speed without a complete stop, which suggests that congestion is a major cause of wasted fuel on open roads.

Finally, we simulate drivers' decision of how to drive and which of two hypothetical routes to take between an origin and destination: one "city" route that is short but has a low speed limit and requires the driver to make a series of stops and one "highway" route that is longer but has a higher speed limit and is stop-free.[7] By pairing our model with a distribution of value-of-time generated from Michigan workers in the Current Population Survey (CPS), we are able to simulate drivers' optimal decision of not only which route to take but also how aggressively to accelerate out of any stops and the optimal speed to drive under a variety of policy counter-factuals. This allows drivers to substitute between the quicker freeway trip that uses more fuel and the slower city trip that uses less fuel. Intuitively, with an optimal Pigouvian tax on gasoline, the social cost of driving is reduced relative to the base case, with lower-income drivers choosing the more fuel efficient city route and higher-income drivers choosing the faster but less fuel efficient highway route. More interestingly, policies that improve the flow of city traffic, either by timing stoplights or installing traffic circles to reduce the amount of stopping and starting, lead to a lower overall social cost of driving than the Pigouvian tax, since drivers who were originally taking the city route have reduced time and fuel costs and drivers who switch from the highway route to the city route both improve fuel economy and save on fuel costs.

[5] At higher speeds, there is a stronger relationship between acceleration and fuel economy, with higher rates of acceleration always reducing fuel economy.

[6] While Wolff (2014) finds some evidence that drivers in rural Washington state reduce their freeway speeds when gasoline prices increase, Burger and Kaffine (2009) do not find a change in driving speeds in Los Angeles when gasoline prices change. Neither paper has access to a panel of individual-level driving data.

[7] Note that we do not model congestion on either route.

The simulations generally suggest an additional approach to reducing fuel consumption beyond the usual taxes and fuel economy standards studied by economists. Although most drivers' value of time is high enough that they will not change their driving behavior to reduce fuel use in response to a gasoline tax, policy interventions that improve traffic flow or decrease the number of vehicle stops and starts can both save drivers time and reduce fuel use. These are generally not improvements that drivers can make on their own, and require infrastructure changes or centralized reorganization of traffic signals.

The remainder of the paper is organized as follows. Section 2 describes the data on individual driver behavior and fuel use and documents the variation in fuel economy across drivers in identical vehicles. Section 3 lays out a model of driver behavior that nests a physical model of vehicle fuel consumption given drivers' acceleration and speed decisions and trip characteristics. Section 4 presents the results of our estimation of the physical model of fuel use and shows that the model fits the data well and leads to aggregate implications for the variation in average fuel economy across drivers in our sample. Section 5 uses the physical model to compute the values of time that would be necessary to convince drivers to drive more efficiently and simulates the effect of policies that decrease the number of stops on fuel use. Section 6 presents the policy simulations using the results of our physical and behavioral model. Finally, section 7 concludes.

2 Data

The data for this project come from an engineering dataset that has not been used by economists before. Between April 2009 and May 2010, the University of Michigan's Transportation Research Institute (UMTRI) conducted its Integrated Vehicle-Based Safety Systems (IVBSS) study to test modern vehicle safety equipment such as lane-departure and collision warning systems in a real-world setting. In order to do this, UMTRI provided 108 drivers with one of 16 identical vehicles to use as their primary vehicle for forty days. Data on nearly 600 variables were collected from a set of instruments in the vehicles, including fuel consumption, location, speed, radar data on nearby vehicles, and video of both the driver and the road surrounding the vehicle. This data was collected at least 10 times per second, allowing for an extremely detailed understanding of the roadway characteristics and driver behaviors that affect fuel economy.

To recruit the sample drivers, UMTRI sent out information to registered Michigan drivers in southeast Michigan with no major driving infractions. Of the respondants, UMTRI selected drivers who drove more than 12,000 miles per year and who were evenly distributed across gender-age bins. Drivers were nominally compensated for their completion of pre-and post-experiment surveys, and were given the vehicles to drive. The vehicles were given to participants with a full tank of gas, but then participants needed to purchase any additional gasoline they used. Because study participants drove more than the national average and had relatively clean driving records, we might expect them to be more efficient and/or less aggressive drivers than the U.S. population as a whole, which would bias us towards finding little variation in their fuel economy.

The original experiment allowed drivers to drive the vehicles for two weeks and then turned on the additional safety equipment. The underlying concern was that the safety equipment might startle the drivers or otherwise exacerbate dangerous driving situations, but UMTRI found that, instead, driver behavior changed little overall, while unsafe driving practices declined slightly (Sayer et al (2010)). We do not make use of the original experimental design in this study, but pool the control and treatment periods for each driver.

Table 1 provides summary statistics of the driving in our data, broken down by speed bin. In total, we observe 6,743 hours of driving across the 108 drivers who completed the study.[8] The total distance traveled exceeds 372,000 kilometers (over 230,000 miles). This corresponds to an average distance per day of 86 kilometers or 53 miles, which is about 50 percent greater than the national average.While the largest amount of time spent driving falls into the 10-20 meter per second (22.4-44.7 mile per hour) speed bin, the greater than 30 m/s (67.1 mi/hr) speed bin contains both more miles and more fuel use. Overall, drivers in our data use over 36,700 liters of gasoline (over 9700 gallons).

The vehicles used for the experiment were 2006 and 2007 Honda Accord EX 4-door sedans with a V6 engine. The EPA-reported fuel economy for these vehicles is 13.07 L/100km (18 miles per gallon) for city driving and 9.05 L/100km (26 miles per gallon) for highway driving.[9] The overall average fuel economy for driving during the experiment was 9.9 L/100km or 23.8 miles per gallon. As shown in the lower panel of table 1, the fuel economy is substantially worse at very low speeds, is best in the 20-30 m/s (22.4-44.7 mi/hr) speed bin, and deteriorates somewhat at higher speeds.

Figure 1 shows the distribution of average fuel economy across each of the 24,741 trips

[8]Some drivers were removed from the study because either they were not driving the vehicles enough or they were allowing other people to drive the vehicles.

[9]Information obtained from the EPA's fuel economy comparison website at `http://www.fueleconomy.gov/feg/Find.do?action=sbs&id=21962`. The methodology for estimating fuel economy was changed for model years 2008 and later. These estimates are based on the new methodology. There is no difference in reported fuel economy between the 2006 and 2007 model years.

in our data in L/100km.[10] There is substantial variation in fuel economy across trips, and, indeed, the figure is top-coded at 20 L/100km. However, the majority of trips do achieve an average fuel economy between the EPA's highway and city fuel economy estimates. Since the EPA's estimates of highway and city fuel economy are based upon a representative trip (or drive cycle), we would not expect them to capture the full amount of variation in fuel economy across trips, but the extent of this variation seems substantial. Figure 2 aggregates the information in Figure 1 to show the average fuel economy for each driver. This figure shows that, on average, the EPA's estimates of fuel economy are pretty good bounds on the actual fuel economy drivers achieve on the road, with only a few drivers (18%) averaging below 9 L/100 km or above 13 L/100 km. However, the variation in fuel economy across drivers is still striking, with a range from 8.1 to 13.6 L/100 km. This range is equivalent to the difference in EPA average fuel economy between the Toyota Venza (23 mpg) and the Toyota Prius (50 mpg).[11] It is also 30 percent greater than the change in the Corporate Average Fuel Economy (CAFE) standard between 1990 and 2025.

Both Figure 1 and Figure 2 are completely unconditional, meaning that there are no controls for the type of road the driver is on, the number of stops a driver faces, or the weather on a particular day. In order to understand whether policy can decrease fuel use by improving the fuel economy of the existing fleet, we need to better understand drivers' incentives to improve fuel use and the extent to which drivers' behavior actually impacts their fuel use.

3 Model of vehicle fuel consumption

To understand whether there may be a role for policy to reduce fuel use by decreasing the variation in fuel economy across drivers in identical vehicles, we need to understand why the variation in fuel economy exists. There are two reasons why one driver might get better fuel economy than another: the driver could drive in a way that better minimizes fuel use over a route or the driver could drive routes that demand less fuel. We rely on a behavioral model of driver decision-making that nests a physical model of the fuel used by a vehicle under different driving conditions to better understand the variation in fuel economy.

[10] All fuel economy measures will be given in liters per 100 kilometers because this measure is linear in fuel use while miles per gallon is nonlinear in fuel use.

[11] The difference in gasoline consumption between the most and least efficient driver was 5.5 liters per 100 kilometers. The EPA fuel economy for the 2014 Toyota Venza and 2014 Toyota Prius are 10.2 and 4.7 liters per 100 kilometers respectively (23 and 50 miles per gallon), also a difference of 5.5 liters per 100 kilometers. Fuel economy information from http://www.fueleconomy.gov/feg/findacar.shtml.

3.1 Behavioral Model

The model of driver behavior is a standard behavioral model where drivers face a trade-off between the time a particular trip takes and the amount of fuel it consumes. We assume that drivers' trips and routes are fixed, so that a driver is only deciding on the speed and acceleration pattern for the trip, conditional on the roadway characteristics. In particular, since the speed at any point during a trip is completely determined by the acceleration (or deceleration) pattern leading up to that point, we model the driver, $i = 1, ..., N$, on trip $\tau = 1, ..., T$ as making a series of acceleration decisions, a_{is} for $s = 1, ..., S_\tau$ that maximize utility:

$$U_{i\tau} = D_{i\tau} - v_i h(a_{i0}, ..., a_{iS_\tau}|x_\tau) - p_\tau f(a_{i0}, ..., a_{iS_\tau}|x_\tau) - c(a_{i0}, ..., a_{iS_\tau}|x_\tau)$$

where $D_{i\tau}$ is the value to the driver of completing trip τ, v_i is driver i's value of time, $h(a_{i0}, ..., a_{iS_\tau})$ is the time the trip takes given the driver's acceleration decisions, p_τ is the price of gasoline, and $f(a_{i0}, ..., a_{iS_\tau})$ is the amount of gasoline used on the trip given the driver's acceleration decisions. Finally, $c(a_{i0}, ..., a_{iS_\tau}|x_\tau)$ are the additional costs, such as safety and increased wear-and-tear on the vehicle, that may change with acceleration patterns and road characteristics. The characteristics of the day (e.g. temperature) and roadway (e.g. grade) are captured in x_τ.

In order to maximize utility on a given trip, the driver needs to choose the acceleration pattern that sets the marginal benefit of acceleration (decreased travel time) equal to the marginal cost of acceleration (increased fuel use):

$$v_i h'(a_{i0}, ..., a_{iS_\tau}|x_\tau) = p_\tau f'(a_{i0}, ..., a_{iS_\tau}|x_\tau) + c'(a_{i0}, ..., a_{iS_\tau}|x_\tau)$$

In this formulation, the change in time for a change in acceleration, $h'(a_{i0}, ..., a_{iS_\tau}|x_\tau)$, and the gasoline price, p_τ, are relatively straightforward to calculate or observe. The value-of-time, v_i, has been the subject of an extensive literature in economics (e.g. Ashenfelter and Greenstone (2004)), while the change in fuel use for a change in acceleration, $f'(a_{i0}, ..., a_{iS_\tau}|x_\tau)$, is a physical equation that relates the characteristics of the day and roadway (x_τ) and vehicle-specific characteristics like aerodynamic drag. The non-fuel costs of changing acceleration patterns is unknown.

From the perspective of drivers, optimal variation in fuel economy could result from differences in the value of time, differences in the fuel-savings from changing acceleration patterns, or differences in the non-fuel costs of acceleration changes. In order to better understand the fuel savings from accelerating differently, we develop a physical model of fuel use that we can estimate with our data for our particular circumstances. Since both the drivers' value of time and the non-fuel costs of acceleration are both unknown, we will use our estimates to calculate a lower bound on the value of time that would be necessary for drivers to be willing to drive more aggressively, while understanding that safety and vehicle wear-and-tear calculations will likely push these values of time up in actual driving.

3.2 Physical Model

Our model of the relationship between acceleration and fuel use is adapted from Saerens et al. (2010) and Hellström et al. (2009). We develop the theoretical structure of our model using these papers from the engineering literature, then use our observed data on fuel consumption to estimate econometrically the parameters of the model for our particular vehicles. The model provides a basic understanding that instantaneous fuel flow, $g(r, \omega | x_t au)$, is a nonlinear function of engine torque, r, and engine rotation speed, ω, and that this relationship is influenced by external factors, x_τ, like the road grade and the outside temperature.

In order to relate engine torque, acceleration, and speed (the integral of the driver's acceleration decisions) to fuel use, we need a few additional relationships. First, we recognize that the total friction force acting on the wheels of the vehicle, F_ω is given by

$$F_\omega = ma + F_a(\nu) + F_{rr}(\alpha) + F_N(\alpha)$$

where m is the mass of the vehicle and $F_a(\nu)$ is the aerodynamic resistance, which is proportional to the square of the velocity of the vehicle.[12] $F_{rr}(\alpha)$ is the rolling resistance of the vehicle, which is proportional to the roadway slope, α, and $F_N(\alpha)$ is the gravitational force, which is proportional to the sine of the road slope.

The torque required to move the vehicle is then a function of the force acting on the wheels of the vehicle, F_ω, divided by the engine efficiency, i_g, which is a function of the gear.[13] Similarly, the engine rotation speed, ω, is a function of the vehicle velocity, ν, and the engine efficiency i_g. Combining all of these components allows us to formulate an

[12]Other components of air resistance are the frontal surface of the vehicle, the vehicle drag coefficient, and density of air. These are assumed to be constant.

[13]Technically i_g is the combined transmission and final drive conversion ratio for the chosen gear g. The torque is also a function of the wheel radius and the power transmission efficiency, which are assumed to be constant.

equation for the instantaneous fuel use of the vehicle:

$$g\left(\frac{\delta_1}{i_g}(\delta_2 a + \delta_3 \nu^2 + \delta_4 cos(\alpha) + \delta_5 sin(\alpha)), \nu i_g\right) \tag{1}$$

Where $\delta_1, ..., \delta_5$ are constants that are specific to vehicle models. This model demonstrates that fuel use is a non-linear function of gear, acceleration, speed, and road grade. Since the vehicles have automatic transmissions, we avoid modeling gear changes by assuming that gear changes are instantaneous and that the engine efficiency is a function of velocity: $i_g = e(\nu)$.

3.3 Empirical framework

We take this physical model to our data to understand how drivers' acceleration decisions affect fuel consumption, and analyze the incentives facing drivers in our behavioral model. Equation (1) shows us that the rate of fuel use is a nonlinear function of velocity, acceleration, and road grade. We will approximate this function using the interation of sixth-order polynomials in velocity, fourth-order polynomials in acceleration, and road grade. By separately modelling positive and negative acceleration, we allow additional flexibility. Combining all of these components with our second-by-second data gives us our estimation equation:

$$y_t = \alpha + \sum_{i=0}^{6} v_t^i \left[\sum_{j=0}^{4} \beta_{ij0}(a_t^+)^j + \sum_{k=0} \beta_{i0k}(a_t^-)^k + \delta_i \alpha_t\right] + \gamma \mathbf{z_t} + \varepsilon_t \tag{2}$$

In equation (2), y_t is the fuel use for a given second of driving in our sample, as measured in liters per 100 kilometers. v_t is the mean speed during the second-of-sample t, a_t^+ is the mean acceleration in meters per second squared if this is positive (and zero otherwise), and a_t^- is the mean acceleration in meters per second squared if this is negative (and zero otherwise), and α_t is the mean road grade measured in radians. z_t contains other explanatory variables that are not interacted with the polynomial in speed: the outside temperature in degrees Celsius and a measure for the use of the air conditioner.

Because the dependent variable of equation (2) is fuel use as measured in liters per 100 kilometers, this variable approaches infinity for extremely low speeds or idling. For this reason, we estimate a separate model for fuel consumption at zero or very low speeds (less than 2 meters per second or 4.5 miles per hour). This low-speed model is identical to equation

(2) except that the dependent variable is measured in mililiters per second.

4 Driver behavior and fuel consumption results

4.1 Estimates

Table 2 shows the estimation results for equation (2), estimated for all observations with mean speed above two meters per second and liters per 100 kilometers as the dependent variable. Column 1 shows the results for a quadratic in speed and exluding all acceleration terms, while column 2 adds linear terms in acceleration, with positive and negative acceleration entering separately. Column 3 adds the interaction of the linear terms in acceleration and speed and column 4 shows a selection of coefficients from the full set of estimation results for equation (2), including the full interactions between a sixth-order polynomial in speed and a fourth-order polynomials in the two acceleration terms.

The first thing to notice is that the full model in column 4 fits the data extremely well and substantially better than other models, with an R^2 of 0.933. The largest increase in R^2 comes from adding the linear acceleration terms to the model, suggesting that acceleration is critically important in understanding fuel use, even controlling for speed. This means that speed sensor data will be substantially weaker at understanding on-road fuel use than the panel data that we use in this study.

The coefficient estimates are fairly consistent across models at least in terms of sign. Increased speed decreases fuel use, but at a decreasing rate. The implied minimum constant-speed fuel use occurs 90 kilometers per hour (55.9 miles per hour) for the full model, at which point the vehicle is using 7.14 liters per 100km (32.9 miles per gallon). Figure **??** makes this relationship clear by showing both the observed and predicted fuel economy for constant speed driving at different speeds over 2.5 kilometer per hour speed bins. The gray bars are the observed fuel economy in our data over all one-second observations in our data where acceleration is zero and the vehicle is driving on level road. The black line is the fitted relationship using the coefficients from column 4 of table 2, with acceleration and grade set to zero and all other non-speed terms set to their mean values in the sample. The model fits the data very well and shows that there is a substantial improvement in fuel use as speed increases from very low speeds (and therefore inefficient low gears) and a more gradual increase in fuel use at speeds over 100 kilometers per hour (62 mph). The minimum predicted fuel use from our model is somewhat higher than the observed minimum fuel use at 72.5-75 kilometers per hour (45-46.6 mph), but both the predicted curve and the observed fuel use are very flat through this entire region.

Positive acceleration increases fuel use substantially across all of the models in table 2, and negative acceleration decreases fuel use in the full model in column 4. The interactions between acceleration and speed show that his effect is largest at low speeds and diminishes at higher speeds. Uphills increase fuel use and downhills decrease fuel use, as shown by the positive coefficient on the sine of road grade, and air conditioner use always increases fuel use. Higher outdoor temperatures generally decrease fuel use, as the engine runs more efficiently at higher temperatures.

Table 3 shows the same set of estimation results, but on observations at speeds less than 2 m/s and a dependent variable of mL/s. In this speed range, increasing speed increases fuel use per second, but at a decreasing rate (although in each second the vehicle is covering a greater distance). Positive acceleration again increases fuel use and negative acceleration decreases fuel use, but this relationship is strongly tied to the speed of the vehicle as the interaction term is large and negative for positive acceleration interacted with speed and large and positive for negative acceleration interacted with speed. As before, air conditioner use increases fuel use and higher outside temperatures decrease fuel use. This model does not fit the data nearly as well as it fits the higher-speed data, with the R^2 only reaching 0.494 for the full model in column 4.

4.2 Model fit

We have seen that the model fits constant speed driving quite well, as evidenced by figure 3. We conduct an additional check of model fit by taking a single trip and looking at the predicted and observed fuel use given the characteristics of the trip.

Figure 4 shows the observed and predicted fuel use over a short trip of just under 6 kilometers. The top panel of figure 4 shows the actual fuel use in one-tenth of a kilometer bins with the gray bars. The black diamonds display the predicted fuel use in that one-tenth of a kilometer bin using the characteristics of this particular trip. The bottom panel of figure 4 shows the speed in kilometers pre hour over the trip. The obvious first take-away is that the model predicts the variation in fuel consumption over the trip extremely well. The black diamonds are generally very close to the tops of the gray bars, although there are occasionally some small differences. The second thing to notice about figure 4 is that fuel consumption is much higher during acceleration events. The fuel consmption when the vehicle is accelerating is substantially higher than the fuel consumption when the speed is either constant or decreasing. This figure does not make it clear whether different acceleration patterns drastically affect fuel use, but we will explore these questions further in our counterfactual exercises.

4.3 Determinants of overall fuel use

While the second-by-second estimation of the physical model gives us the ability to strongly predict nearly instantaneous fuel use, it is useful to confirm that the lessons learned from these models aggregate up appropriately to understand overall fuel use by a driver. In order to do this, we did a simple, cross-sectional analysis of the mean driver fuel consumption over the entire sample as measured in liters per 100 kilometers as a function of the important components of the physical model. Table 4 provides the results of a regression of this mean driver fuel consumption on different characteristics of the type of driving that the driver does over the entire period she has the vehicle. In order to standardize the interpretation of the coefficients, all independent variables are converted into z-scores over the distribution of the variable across the 108 drivers, such that a one unit increase in the independent variable should be thought of as a one standard deviation increase in that variable.

Column 1 of table 4 shows the results when we only include characteristics of the period that the driver had the vehicle that are somewhat out of the driver's control, includingthe driver's age and gender and the mean outside air temperature and the share of time the driver has the air conditioning on. These variables explain 21.1 percent of the variation across the 108 drivers in mean fuel use.

In columns 2 and 3 of table 4 we first add variables describing the speed the driver chooses over the period in column 2 and then the acceleration of the driver in column 3. Speed variables include the average speed, the percent of time the driver spends idling, the percent of time the driver spends at speeds over 100 km/h, and the average speed conditional on being over 100 km/h. Finally, in column 3, we add the number of acceleration events per kilometer of driving,[14] the average acceleration rate during acceleration events, and the average deceleration rate during deceleration events. Adding the speed variables increases the R^2 to 0.779 and additionally including the acceleration variables increases the R^2 to a somewhat surprising 0.897.

The first thing to notice is that while the demographics of the driver appear to be important in the first column, these effects nearly disappear once we control for the decisions about how to drive in column 3. Additionally, in column 2, before we control for acceleration,

[14] Acceleration events are defined as an increase in speed of at least 5 m/s with no more than a 0.001 m/s reduction in speed over any one second interval during the event.

it appears that higher speeds drastically decrease fuel use conditional on being under 100 km/h and idle time has almost no effect. Once we control for acceleration characteristics in column 3, it is clear that the most important determinant of average driver fuel use in a per-standard-deviation sense is actually the number of acceleration events per kilometer of driving, with idle time, speed if greater than 100 km/h, and average acceleration rate trailing behind. Average speed and the average time spent driving over 100 km/h become insignificant and small. In this full model, outside temperature and air conditioning use are still substantial determinants of fuel use.

This analysis of the aggregate fuel use by drivers shows that the lessons learned in the second-by-second physical model do aggregate up to all driving. Acceleration events use a lot of fuel and are extremely important in explaining the variation in fuel use across drivers, and may swamp decisions about how fast to accelerate or how fast to drive. In combination with the physical model estimates, this motivates our analysis of the value of time that drivers would have to have in order to decide to accelerate more aggressively and the potential for policy to intervene in driving to improve fuel use.

5 Drivers' Incentives to Improve Fuel Economy

Now that we have estimated a model of fuel use for our sample, we can look at the incentives of drivers to unilaterally improve fuel economy. First, we will look at drivers' incentive to drive at a faster constant speed. Then we will look at acceleration events to understand the values of time for which drivers have an incentive to accelerate more aggressively. Finally we will analyze the cost of accelerating and decelerating around an average speed rather than maintaining that speed in order to think about the role of cruise control or other computer-assisted driving tools in improving fuel economy.

As we explained in the introduction, the EPA's website suggests that drivers should decrease their speed on the freeway and accelerate less aggressively in order to improve fuel economy. The estimates of our physical model suggest that this advice is somewhat true: fuel consumption is minimized at 90 km/h (55.9 mph), which is substantially slower than most people drive on the freeway, and the linear acceleration term does have a positive relationship with fuel use. However, neither of these facts take into account the trade-offs that are central to the behavioral model: increasing fuel consumption by increasing speed or acceleration may be optimal if a driver has a high enough value of time, taking into account any changes in safety. In this section we explicitly solve for the minimum combination of value of time and value of safety changes that implies that drivers should ignore the EPA advice.

Constant speed driving

As we saw in figure 3, driving faster decreases fuel consumption at less-than-freeway speeds and increases fuel consumption above 90 km/h (55.9 mph). In table 5, we calculate the cost of driving 100 km at different speeds in terms of fuel consumption, fuel cost (at $3.50 per gallon), and time, and then use these numbers to calculate the cost of time in $ per hour that makes the driver indifferent between driving that speed or 10 km/h (6.1 mph) slower if the safety and depreciation costs are zero. At speeds below 90 km/h, increasing speed actually decreases fuel use, so all drivers would prefer to drive faster. Above 90 km/h, increasing speed 10 km/h is optimal if the driver's value of time net of non-fuel costs is quite low: 81 cents for the increase to 100 km/h (61 mph), $3.84 for the increase from 100 km/h to 110 km/h (68.4 mph), and $7.88 for the increase from 110 km/h to 120 km/h (74.6 mph). Estimating the changes in safety or depreciation from increasing speeds by this amount is beyond the scope of this paper, but these values seem well below the value of time for most American workers.[15] Even for speed increases above 120 km/h, many drivers will find that their values of time net of safety are above the $11.67/hour cost of increasing to 130 km/h (80.8 mph) or the $13.61/hour cost of increasing to 140 km/h (87 mph). This suggests that individuals, left to their own devices, have very little incentive to decrease their freeway speed to the fuel-economy-maximizing level.

There is one dimension on which it makes a lot of sense for drivers to adjust their driving behavior to decrease fuel consumption. If a driver at freeway speeds drives a constant speed over a distance, she will use less fuel than if she covers the same distance at the same average speed, but varies her speed over the distance. Table 8 compares the fuel consumption and fuel economy of a driver averaging 30 m/s (67.1 mph) over a 1 km stretch of road driving at a constant speed versus accelerating and then decelerating at a constant rate (2 m/s^2) and then finishing the 1 km drive at 30 m/s. While the time spent completing each of these kilometers is the same, the fuel use is strictly increasing in the amount of variation in speed that occurs over the kilometer. This means that there is no value of time that makes it optimal to drive at varying speeds on the highway rather than an average, consistent speed. This suggests that, at least on flat road, cruise control is useful at improving fuel economy,

[15] An annual income of $42,500 is equivalent to an hourly wage of $21.25. Using the standard assumption that the value of time spent driving is equal to one-half of the wage, this puts the average value of time of American workers at $10.63.

and congestion that forces the driver to vary speed is costly even if the average speed the driver can travel is unaffected.

Acceleration events

Since our physical model showed that acceleration is at least as important as speed in determining fuel economy, we also look at how drivers' acceleration decisions affect their fuel use and find the minimum value of time for which a driver would choose to accelerate more aggressively. In order to understand general acceleration patterns, we find all acceleration events in our data where drivers either accelerate from 2-15 m/s (4.5-33.6 mph) or from 20-30 m/s (44.7-67.1 mph). We think of the first set of acceleration events as representing acceleration from something close to a stop and the second set as representing merging onto a freeway.

Table 6 displays descriptive statistics for the acceleration and deceleration events in our data, including the fifth and ninety-fifth percentiles. Intuitively, accelerations at higher speeds take up more distance than accelerations at lower speeds. They also take up more time and fuel than the accelerations from very low speeds. Additionally, as expected, deceleration events use very little fuel even though they last nearly as long in terms of both distance and time as acceleration events.

In figure 5 we show how different choices about acceleration from 2-15 m/s (4.5-33.6 mph) affect fuel consumption. There are two conflicting effects to keep in mind. First, accelerating more aggressively requires more torque on the wheels, which can increase fuel consumption. However, by accelerating more aggressively, the driver gets into the higher, more efficient gears more quickly, which decreases fuel consumption. Additionally, since accelerating more aggressively means that the driver hits the 15 m/s speed in a shorter distance, we standardize our comparison by looking at the fuel consumed over the longest acceleration distance, 250 m, which occurs with $0.5m/s^2$, assuming that drivers that accelerate more aggressively maintain the constant 15 m/s speed until they reach 250 m. In figure 5, the dark bars represent the amount of fuel used during the acceleration period at different acceleration rates, and the corresponding diamonds represent the predicted fuel use for these accelerations from our model. The light colored bars add the constant-speed driving that allows the driver to reach 250m.

There are two important things to take away from figure 5. First, accelerating aggressively (up to 2.75 m/s^2) uses substantially less fuel during the acceleration phase than accelerating very slowly, although this fuel savings diminishes with acceleration. This is because, although the fuel consumption at any second is higher at higher acceleration rates, the time spent accelerating up to 15 m/s is much shorter with more aggressive acceleration. The second point is that this fuel savings is offset by the fact that the top cruising speed of 15 m/s is reached in a shorter distance, so the vehicle needs to drive at a constant speed of 15 m/s for a longer distance. The combined effect is that accelerating more aggressively uses slightly more fuel than accelerating less aggressively, but the difference is quite small.

Of course, accelerating at a slower rate means that it takes substantially longer to cover 250 m than accelerating quickly. Table 7 shows, for different acceleration rates, the fuel consumption, fuel cost and time for accelerating from 2-15 m/s over 250 m (top panel) and for accelerating from 15-25 m/s over 500m (bottom panel). Table 7 also shows the minimum value of time net of non-fuel costs that would be required to make accelerating at that level preferable to accelerating one level more slowly. For accelerations from a near-stop, the minimum value of time net of non-fuel costs never exceeds \$4.76, which means that drivers would need to have an extraordinarily low value of time or be extremely safety- and depreciation- conscious for accelerating less aggressively to be preferable to more aggressive acceleration. For accelerations onto a highway, the optimal approach is to either accelerate somewhat slowly or extremely aggressively. Of course, both of these results assume that the driver is allowed to drive freely after the acceleration event, allowing the decreased time during the acceleration event to translate into a decrease in the total trip time.

Our overall take-away from our value-of-time calculations is that most individual drivers have very little incentive to drive more efficiently in a given vehicle on a given route. The exception, of course, is that drivers should refrain from unnecessarily varying their speed over a flat road, which increases fuel use without simultaneously decreasing trip time. However, often drivers are faced with choices about which route to take between a given origin and destination, and our policy simulations seek to understand whether fuel use can play a role in these choices and whether that may open up an additional avenue for policy to decrease overall fuel consumption without changing the vehicles on the road or the trips that drivers need to take.

6 Policy Simulations

In this section we use our econometric estimates to develop a stylized model for simulating gasoline consumption. This model incorporates two levels of optimal decision-making by drivers. First, on any route drivers are assumed to choose their optimal speed and acceleration, given the value of their time, and subject to legal constraints on maximum speed. Second, drivers choose optimally between alternative routes that differ in the time taken and the gasoline used, in order to minimize the total cost of their travel. Using this model we then compare the effectiveness of alternative policies aimed at reducing gasoline consumption.

6.1 Simulation model

We assume that drivers behave optimally to minimize the total cost of their travel between fixed start and end locations. This total cost TC_i comprises two components: the gasoline cost and the value of time. Driver i chooses both the route j and their optimal pattern of acceleration and speed for that route, $a(j)$.

$$\min_{j,a(j)} TC_i = PQ(j, a(j)) + VOT_i t(j, a(j)) \tag{3}$$

In this expression P is the price of gasoline in dollars per liter, $Q(j, a(j))$ is the quantity of gasoline in liters consumed on route j given the acceleration sequence $a(j)$, VOT_i is the value of time for driver i in dollars per minute, and $t(j, a(j))$ is the time in minutes to drive route j with acceleration sequence $a(j)$.

The only heterogeneity across drivers in our model is in their assumed value of time, VOT_i. Drivers with low value of time will be more likely to choose a slower route and to drive in a way that reduces gasoline consumption. For our simulation, we use the distribution of personal hourly earnings for Michigan adults from the Current Population Survey (CPS) in 2012. Following the standard approach in the travel cost literature, we assume that VOT_i is half the hourly earnings. Because the CPS data is reported in 40 bins, we only need to calculate the optimal travel behavior once for each of the possible bins. We then weight each of the observations by the percentage of Michigan adults with earnings in that bin in order to calculate the distribution of simulated gasoline consumption for the population.

The two possible routes that we consider for our simulation correspond, in broad terms, to a "city" and a "highway" route between the same origin and destination points. Figure 6 illustrates the differences between the two possible routes. The city route has a shorter distance (10 kilometers) but the maximum allowed speed is lower (60 kilometers per hour in the base case). The city route also has several required stops (five in the base case) where the car is required to come to a complete stop before being able to accelerate again. By comparison, the highway route has a longer distance (15.7 kilometers) but a higher maximum allowed speed (110 kilometers per hour). There are no intermediate stops for the highway route.

We characterize the driving "style" on each route by three parameters: the maximum speed, the constant rate of acceleration from a stop, and the constant rate of deceleration to a stop. These parameters are sufficient to calculate the entire time path of speed and acceleration along the route. The calculation immediately gives the total time for the trip. Using the time path of speed and acceleration (in steps of one hundredth of a second), combined with the vehicle fuel consumption model from Section [3], we predict the gasoline consumption over the trip. Total trip cost can then be calculated using an assumed gasoline price and value of time.

For each possible value of time, and for the two possible routes, we choose the values for the speed and acceleration parameters that minimize the total travel cost. We do this using a grid search over 21 possible maximum speeds (in 0.5 m/s increments), 14 acceleration rates (in 0.2 m/s^2 increments) and 14 deceleration rates (also in 0.2 m/s^2 increments). After calculating the minimum travel cost on each route, each driver is assumed to choose the route (city or highway) with the lowest travel cost.

The usefulness of the simulation model is that it allows us to compare various policies aimed at reducing gasoline consumption. These include increases in the gasoline tax (modelled by increasing the gasoline price), reductions in the speed limit (modelled by changing the constraint on maximum speed), and reductions in the number of city stops. Although the model is highly stylized, it does incorporate several realistic features: the heterogeneity across drivers in value of time, the trade-off between time and gasoline consumption estimated from real-world vehicle data, and the ability of drivers to choose between alternative routes to reach the same destination.

6.2 Simulation results and policy comparisons

The first column of Table 9 shows the base case results for the simulation. Out of the representative population of Michigan adult drivers, 40.2% chose the city route and 59.8% chose the highway route. With a lower speed and more stops, the city route took a longer time than the highway route: 10.80 minutes compared to 8.75 minutes. However, the advantage of the city route is that gasoline consumption is lower: 1.065 liters compared to 1.330 liters on the highway. The drivers choosing the city route are the ones for whom it is optimal to take an extra two minutes in order to save about a quarter of a liter of gasoline. These are the drivers with the lowest value of time. The drivers with the highest value of time will always choose the faster highway route that uses more gasoline.

Interestingly, although the city route uses less gasoline, the observed fuel economy in liters per 100 kilometers is lower for the city than for the highway. This is because the additional fuel consumption on the highway route is less than the proportional increase in trip distance. This suggests that a focus on comparison of fuel economies across drivers can be misleading. Because the benefit that drivers obtain from either route is identical—transportation between the same origin and destination—the most relevant comparison is the fuel use for the entire trip, not the fuel use per unit of distance.

The bottom of the table shows the calculation of the overall average cost to the drivers and society of the trip. The driver cost includes the gasoline cost (including gasoline tax if included) and the value of the driver's time. The average cost of the trip is \$2.92, comprising \$1.16 of gasoline cost and \$1.76 for the value of time. The social cost also includes the value of the carbon dioxide externality from the gasoline consumption, assuming a social cost of carbon of \$40 per ton of carbon dioxide.[16] In the base case, the carbon externality is about 11.4 cents per trip.

The remaining four columns of Table 9 show the effect of imposing different levels of a gasoline tax. The first tax considered is exactly the Pigouvian tax corresponding to the social cost of carbon in gasoline. The remaining three cases consider higher taxes: 25, 50 and 75 cents per liter.

All of the gas taxes cause drivers to shift from the highway to the city route. The drivers that change their route are exactly those drivers whose value of time leaves them only slightly better off on the highway (with shorter time but more gasoline consumption). A slight increase in the gasoline price causes these marginal drivers to switch to the route with lower gasoline consumption.

The higher gas tax leads to only a small change in driving behavior (the choice of speed and acceleration on a given route). On the city route, gasoline consumption declines from 1.065 liters for the trip to 1.063 liters for the trip, when the gas tax increases to 75 cents per liter. This understates the change in behavior of individual drivers because the composition of the city drivers is also changing: the drivers with higher value of time who switch from the highway route are less likely to reduce their speed or acceleration in order to save fuel.

Overall the gas taxes reduce the average fuel consumption for the trip: from 1.223 liters to 1.144 liters. However, because more drivers are choosing the slower route (and, to a lesser extent, are choosing to drive slower), the average time for the trip increases from 9.58

[16] In the current analysis we do not consider other externalities such as local pollution, traffic congestion, or accidents. Additionally, we do not currently consider the fact that a single gallon of gasoline could have very different overall emissions of local air pollutants as discussed in Fullerton and West (2010).

minutes to 10.25 minutes. This increase in trip time more than offsets the reduction in gasoline consumption: social cost for the trip is \$3.055 with a 75 cents per liter gas tax, compared to \$3.035 in the base case. Imposing the Pigouvian gasoline tax reduces the social cost by 0.1 cents.

Table 10 shows the effect of changing the city or highway speed limits.[17] The first column in the table repeats the base case results from Table 9. Reducing the city or the highway speed limit by 10 kilometers per hour increases the average trip cost for the drivers by 2 cents (for the lower city limit) and 11 cents (for the lower highway limit). The lower speed limit in the city causes marginal drivers to switch to the highway route, which is faster but consumes more gasoline. Conversely, the lower speed limit on the highway causes marginal drivers to switch to the city route, which uses less gasoline but is much slower. The result that lower speed limits makes drivers worse off in our model is not surprising. Drivers are already assumed to choose their optimal speed, acceleration and route, and could have already chosen the slower speed if it had a lower cost. Any additional constraints on the driver's behavior will necessarily make them worse off, but the difference between a speed limit change on the city route rather than the highway route is that the welfare losses are in terms of increased fuel use rather than increased time.

The table also shows the effect of increasing the city or the highway speed limits. Either policy change would make the drivers in our model better off. A higher city speed limit would lead to a large increase in the number of drivers on the city route. Average fuel consumption for the city drivers would increase by 0.039 liters but this is more than offset by a reduction in the average city travel time by 1.24 minutes. A higher highway speed limit would not change the number of drivers on the highway route, because the marginal driver on the city route has a value of time sufficiently low that they would not drive at the (new or old) highway speed limit anyway. Instead, the higher speed limit on the highways benefits the existing highway drivers who save more than half a minute at the expense of an additional 0.08 liters of fuel consumption. Again, this reinforces the idea that policies that make the more efficient city route more attractive to drivers (by making it take less time or use less fuel) leads to an overall decrease in fuel use relative to policies that make the highway route more attractive.

Interestingly, there is little change in the value of the carbon externality for any of the possible changes in speed limit. It increases by, at most, 1 cent per trip. This means that

[17] Of course, since these analyses do not account for changes in accident rates, these changes in speed limits should not be seen as a full welfare analysis of changing speed limits. For a more thorough discussion of the safety impacts of speed limit changes, see van Benthem (2012).

the overall change in social cost is dominated by the changes in the drivers' private costs: the gasoline cost and the value of time. In the context of this model, reducing speed limits makes society unambiguously worse off, wherease raising speed limits makes society better off.

Our final policy simulation is to analyze the effect of either reducing the number of stops on the city route or turning all of the stops into traffic circles or "rotaries" where drivers must decrease their speed to 25 km/h in order to pass through the circle, but are not required to stop. Column 1 of Table 11 again presents the same base case as before, with 5 full stops on the city route. Columns 2 through 4 of Table 11 reduce the number of stops to 4, 3, or 2 respectively. As the number of stops on the city route decreases, both the fuel consumption and the trip time on the city route decrease, making the city route more attractive to marginal drivers. This means that the overall private and social trip cost declines, even though the average time cost increases as drivers switch from the quicker highway route to the slower city route. The final column of Table 11 shows the effect of converting all of the city stops to traffic circles, which is roughly similar in aggregate to reducing the number of stops from 5 to 3.

7 Conclusion

Overall, our results point strongly to the fact that Pigouvian gasoline taxes and other driver incentives can only go so far to improve fuel economy. With increasingly fuel efficient vehicles, the incentive for drivers to individually improve their fuel economy in their existing vehicles is substantially mitigated by the time cost of doing so. However, there is still the potential for win-win policies that reduce fuel use and substantially reduce the private cost of driving. These policies require focusing attention on aggregate traffic flows and investing in improved technology (connected vehicles, intelligent traffic signals, and potentially self-driving vehicles) and infrastructure. While no driver is incentivized to reduce fuel use for a given vehicle on a given trip individually, the new frontier for environmentally-focused traffic research may be on how to collectively save fuel while saving time.

References

Ashenfelter, Orley and Michael Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy*, 2004, *112* (1), 226–267.

Burger, Nicholas E and Daniel T Kaffine, "Gas prices, traffic, and freeway speeds in Los Angeles," *The Review of Economics and Statistics*, 2009, *91* (3), 652–657.

Fullerton, Don and Sarah West, "Tax and Subsidy Combinations for the Control of Car Pollution," *B. E. Journal of Economic Analysis and Policy – Advances*, 2010, *10* (1).

Hellström, Erik, Maria Ivarsson, Jan Åslund, and Lars Nielsen, "Look-ahead control for heavy trucks to minimize trip time and fuel consumption," *Control Engineering Practice*, 2009, *17* (2), 245–254.

Saerens, Bart, Moritz Diehl, and Eric Van den Bulck, "Optimal control using Pontryagins maximum principle and dynamic programming," in "Automotive Model Predictive Control," Springer, 2010, pp. 119–138.

van Benthem, Arthur, "Do we need speed limits on freeways?," 2012.

Wolff, Hendrik, "Value of time: Speeding behavior and gasoline prices," *Journal of Environmental Economics and Management*, 2014, *67* (1), 71–88.

Table 1: Summary statistics for all driving, by speed category

	Speed bin						Total
(m/s)	0	0–2	2–10	10–20	20–30	≥ 30	
(mi/hr)	0	0–4.5	4.5–22.4	22.4–44.7	44.7–67.1	≥ 67.1	
Total by speed bin							
Time (hours)	1,025	191	1,087	1,752	1,277	1,021	6,352
% of total	16.1	3.0	17.1	27.6	20.1	16.1	100.0
Distance (km)	0	290	21,638	97,389	112,660	121,427	353,404
% of total	0.0	0.1	6.1	27.6	31.9	34.4	100.0
Fuel consumption (L)	1,814	414	4,643	9,413	8,540	10,000	34,824
% of total	5.2	1.2	13.3	27.0	24.5	28.7	100.0
Mean by speed bin							
Fuel economy (L/100 km)	.	143.0	21.5	9.7	7.6	8.2	9.9
Speed (km/h)	0.0	1.5	19.9	55.6	88.2	119.0	55.6
Acceleration > 0 (m/s^2)	0.00	0.22	0.42	0.22	0.10	0.06	0.17
A/C usage	0.27	0.26	0.24	0.20	0.21	0.22	0.23
Outside temperature (°C)	13.9	14.8	14.5	13.9	13.5	13.6	13.9

Note: Each

Table 2: Estimates for fuel consumption in liters per 100 kilometers, at speeds above 1 meter per second

	(1)	(2)	(3)	(4)
Speed (m/s)	-2.23	-1.62	-1.26	-7.77
	(0.002)	(0.001)	(0.001)	(0.123)
Speed squared	0.04	0.03	0.03	0.82
	(0.000)	(0.000)	(0.000)	(0.021)
Acceleration > 0 (m/s^2)		23.41	28.51	51.04
		(0.006)	(0.014)	(2.308)
Acceleration < 0 (m/s^2)		3.87	-0.03	-11.94
		(0.003)	(0.008)	(2.136)
(Acceleration > 0) $\times$ speed			-0.47	-7.33
			(0.001)	(1.408)
(Acceleration < 0) $\times$ speed			0.34	8.31
			(0.001)	(1.087)
Sin of road grade	34.86	51.72	50.63	6.77
	(0.235)	(0.213)	(0.208)	(1.714)
Air conditioner (0–1)	1.83	1.67	1.66	1.62
	(0.007)	(0.003)	(0.003)	(0.002)
Outside temperature (°C)	-0.02	-0.03	-0.03	-0.03
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	34.30	22.88	19.37	37.41
	(0.018)	(0.008)	(0.010)	(0.231)
Higher order interactions	N	N	N	Y
Minimum fuel speed (km/h)	91.5	83.6	80.2	90.0
Minimum fuel usage (L/100 km)	5.97	4.00	5.25	7.16
Observations	18,037,794	18,037,794	18,037,794	18,037,794
Adjusted R^2	0.263	0.881	0.892	0.932

Note: Each observation is the fuel consumption, in liters per 100 kilometers, during one second of driving at a speed greater than one meter per second. The regression in the fourth column includes the full set of interactions between a sixth-order polynomial in speed and fourth-order polynomials in positive and negative acceleration as well as the sine of the road grade. The speed that minimizes fuel usage per kilometer, assuming zero acceleration and zero grade, is shown at the bottom of the table, along with the fuel consumption at this speed. Robust standard errors are shown in parentheses.

Table 3: Estimates for fuel consumption in milliliters per second, at speeds below 1 meter per second

	(1)	(2)	(3)	(4)
Speed (m/s)	0.27	-0.05	0.01	0.60
	(0.001)	(0.001)	(0.001)	(0.042)
Speed squared	-0.07	0.05	-0.00	-1.45
	(0.001)	(0.001)	(0.001)	(0.307)
Acceleration > 0 (m/s^2)		0.54	0.26	0.08
		(0.001)	(0.001)	(0.044)
Acceleration < 0 (m/s^2)		-0.06	-0.13	-0.32
		(0.000)	(0.001)	(0.031)
(Acceleration > 0) $\times$ speed			0.24	-3.67
			(0.001)	(0.134)
(Acceleration < 0) $\times$ speed			0.06	3.92
			(0.001)	(0.187)
Air conditioner (0–1)	0.12	0.12	0.12	0.12
	(0.000)	(0.000)	(0.000)	(0.000)
Outside temperature (°C)	-0.00	-0.00	-0.00	-0.00
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.51	0.51	0.51	0.51
	(0.000)	(0.000)	(0.000)	(0.000)
Higher order interactions	N	N	N	Y
Observations	4,828,216	4,828,216	4,828,216	4,828,216
Adjusted R^2	0.174	0.456	0.478	0.492

Note: Each observation is the fuel consumption, in milliliters, during one second of idling or driving at a speed less than one meter per second. The regression in the fourth column includes the full set of interactions between a sixth-order polynomial in speed and fourth-order polynomials in positive and negative acceleration. Robust standard errors are shown in parentheses.

Table 4: Estimates for mean driver fuel consumption in liters per 100 kilometers

	(1)	(2)	(3)
Speed (km/h)		-1.29	-0.05
		(0.194)	(0.209)
Idle time (0–1)		-0.03	0.35
		(0.113)	(0.097)
Speed > 100 km/h		0.30	-0.06
		(0.139)	(0.106)
Speed | Speed > 100 km/h		0.38	0.20
		(0.074)	(0.054)
Acceleration events per km			0.71
			(0.095)
Acceleration (m/s^2)			0.17
			(0.046)
Deceleration (m/s^2)			0.03
			(0.048)
Outside temperature (°C)	-0.32	-0.39	-0.43
	(0.155)	(0.088)	(0.060)
Air conditioning (0–1)	0.05	0.29	0.27
	(0.160)	(0.087)	(0.059)
Female (0/1)	0.73	0.31	0.15
	(0.198)	(0.122)	(0.085)
Age 40–50 (0/1)	-0.69	0.04	-0.01
	(0.259)	(0.154)	(0.106)
Age 60–70 (0/1)	-0.64	0.02	0.05
	(0.247)	(0.162)	(0.111)
Constant	10.19	9.94	10.02
	(0.205)	(0.131)	(0.090)
Observations	108	108	108
Adjusted R^2	0.211	0.779	0.897

Note: Each observation is a driver, and the dependent variable is the mean fuel economy for the driver over the six week driving period. All explanatory variables (except the demographic indicators) are standardized z-scores for the mean value of the variable over the six week driving period. Means and standard deviations for the variables are shown in Table **??**.

Table 5: Fuel cost and value of time for constant-speed highway driving

	Constant speed (km/h)							
	70	80	90	100	110	120	130	140
Cost per 100 km								
Fuel consumption (L)	7.44	7.27	7.16	7.25	7.63	8.27	9.08	9.90
Fuel cost ($)	6.87	6.72	6.62	6.71	7.05	7.65	8.40	9.16
Time (minutes)	85.7	75.0	66.7	60.0	54.5	50.0	46.2	42.9
Effect of increasing speed								
Δ Fuel cost ($)	.	-0.16	-0.10	0.09	0.34	0.60	0.75	0.76
Δ Time (minutes)	.	-10.7	-8.3	-6.7	-5.5	-4.5	-3.8	-3.3
Cost of time ($/hour)	.	-0.88	-0.70	0.77	3.78	7.86	11.76	13.78

Note: The first block shows the fuel consumption, fuel cost, and time taken to drive 100 kilometers at constant speeds. Fuel cost assumes a gasoline price of $3.50 per gallon. The second block shows the effect of increasing speed from the previous column by 10 km/h. For example, increasing speed from 100 to 110 km/h will increase the fuel cost by $0.38 but reduce the trip time by 5.5 minutes.

Table 6: Summary statistics for acceleration and deceleration events

	Acceleration events (m/s)			Deceleration events (m/s)		
	All	2–15	20–30	All	15–2	30–20
Distance (m)						
P5	20.9	47.0	125.1	20.2	39.0	94.2
Mean	140.8	82.8	309.9	93.2	70.4	214.6
P95	381.2	137.6	568.7	227.0	114.3	375.9
Time (s)						
P5	5.1	5.4	5.0	4.2	4.5	3.7
Mean	12.0	8.9	12.2	9.1	7.9	8.4
P95	22.7	14.0	22.2	16.8	12.5	14.9
Fuel used (mL)						
P5	10.0	27.0	58.0	0.0	0.8	0.0
Mean	36.3	32.0	73.3	3.0	2.4	0.2
P95	80.6	38.8	95.2	7.0	5.4	1.0
Acceleration (m/s^2)						
P5	0.48	0.93	0.45	-1.96	-2.88	-2.68
Mean	1.03	1.60	1.00	-1.25	-1.80	-1.41
P95	1.67	2.44	2.02	-0.69	-1.04	-0.67
Observations	234119	50074	2020	229251	40137	999

Note:

Table 7: Fuel cost and value of time for acceleration events

	Acceleration (m/s^2)					
	0.5	1.0	1.5	2.0	2.5	3.0
2–15 m/s over 250 m						
Fuel consumption (mL)	47.12	45.26	45.16	45.66	46.20	46.74
Fuel cost (cents)	4.36	4.18	4.18	4.22	4.27	4.32
Time (seconds)	27.9	22.3	20.4	19.5	18.9	18.5
Δ Fuel cost (cents)	.	-0.17	-0.01	0.05	0.05	0.05
Δ Time (seconds)	.	-5.6	-1.9	-0.9	-0.6	-0.4
Cost of time ($/hour)	.	-1.10	-0.18	1.77	3.18	4.76
15–25 m/s over 500 m						
Fuel consumption (mL)	71.31	76.71	79.80	80.96	80.54	79.10
Fuel cost (cents)	6.59	7.09	7.38	7.49	7.45	7.31
Time (seconds)	24.0	22.0	21.3	21.0	20.8	20.6
Δ Fuel cost (cents)	.	0.50	0.29	0.11	-0.04	-0.13
Δ Time (seconds)	.	-2.0	-0.7	-0.3	-0.2	-0.1
Cost of time ($/hour)	.	9.00	15.35	11.63	-6.89	-34.27

Note: Each block shows the simulated fuel consumption and time to drive a fixed distance, with an initial acceleration period. Each column illustrates the effect of a different value of acceleration for this initial period.

Table 8: Additional fuel cost from highway speed fluctuations, for constant average speed

	Variation in speed (m/s)					
	0	2	4	6	8	10
Distance (km)	1.00	1.00	1.00	1.00	1.00	1.00
Time (min)	0.56	0.56	0.56	0.56	0.56	0.56
Fuel consumption (L)	0.075	0.085	0.095	0.105	0.117	0.129
L/100 km	7.46	8.48	9.50	10.54	11.65	12.91

Note: The table shows fuel consumption for driving on a 1 km highway segment at an average speed of 30 m/s. There is an initial period of acceleration and deceleration (at a rate of 2 m/s^2) followed by the constant speed of 30 m/s for the remaining distance. The total difference between the maximum and minimum speed is shown at the top of the column.

Table 9: Policy simulation results: Effect of gasoline tax

	Base	Gas tax (cents/liter)			
		9.4	25	50	75
Share on City Route	0.402	0.456	0.517	0.619	0.699
Fuel consumption (L)	1.223	1.209	1.193	1.165	1.144
City	1.065	1.065	1.065	1.064	1.063
Highway	1.330	1.330	1.330	1.330	1.330
Time (minutes)	9.58	9.69	9.83	10.07	10.25
City	10.80	10.80	10.83	10.88	10.89
Highway	8.75	8.75	8.75	8.75	8.75
Fuel economy (L/100km)	9.343	9.460	9.592	9.809	9.980
City	10.651	10.649	10.645	10.637	10.632
Highway	8.753	8.753	8.753	8.753	8.753
Cost components ($)					
(1) Gasoline cost	1.16	1.15	1.13	1.11	1.09
(2) Gasoline tax	0.00	0.11	0.30	0.58	0.86
(3) Social cost of carbon	0.11	0.11	0.11	0.11	0.11
(4) Value of time	1.76	1.77	1.79	1.83	1.86
Driver cost (1+2+4)	2.921	3.034	3.221	3.516	3.805
Social cost (1+3+4)	3.035	3.034	3.035	3.043	3.055

Note: Base case uses average gas price in Michigan in 2012 ($3.60 per gallon). The other four columns show the effect of a gasoline tax of 9.4, 25, 50 and 75 cents per liter. The first of these (9.4 cents per liter) corresponds to the social cost of the carbon content in one liter of gasoline (based on $40 per ton of carbon dioxide).

Table 10: Policy simulation results: Changes to speed limits

	Base	City (km/h)		Hwy (km/h)		Hwy 100
		50	70	100	120	City 50
Share on City Route	0.402	0.177	0.828	0.565	0.402	0.216
Fuel consumption (L)	1.223	1.270	1.143	1.150	1.272	1.207
City	1.065	1.024	1.104	1.066	1.065	1.024
Highway	1.330	1.323	1.330	1.258	1.411	1.258
Time (minutes)	9.58	9.52	9.42	10.29	9.23	10.35
City	10.80	12.81	9.56	10.76	10.80	12.81
Highway	8.75	8.81	8.75	9.67	8.18	9.68
Fuel economy (L/100km)	9.343	8.745	10.598	9.509	9.654	8.490
City	10.651	10.234	11.042	10.661	10.651	10.235
Highway	8.753	8.812	8.753	9.673	8.175	9.676
Cost components ($)						
(1) Gasoline cost	1.16	1.21	1.09	1.09	1.21	1.15
(2) Gasoline tax	0.00	0.00	0.00	0.00	0.00	0.00
(3) Social cost of carbon	0.11	0.12	0.11	0.11	0.12	0.11
(4) Value of time	1.76	1.74	1.78	1.94	1.66	1.91
Driver cost (1+2+4)	2.921	2.942	2.868	3.028	2.866	3.060
Social cost (1+3+4)	3.035	3.061	2.975	3.136	2.985	3.173

Note: Base case (identical to previous table) assumes a city speed limit of 60 km/h and a highway speed limit of 110 km/h. The four city and highway columns show the effect of lowering and raising the speed limit for either city or highway, in 10 km/h increments, from these base levels. The final column shows the effect of lowering the speed limit on both the city and highway at the same time.

Table 11: Policy simulation results: Reducing city stops

	Base	Number of stops			Traffic
		4	3	2	circles
Share on City Route	0.402	0.517	0.619	0.699	0.565
Fuel consumption (L)	1.223	1.168	1.107	1.044	1.110
City	1.065	1.018	0.970	0.921	0.942
Highway	1.330	1.330	1.330	1.330	1.330
Time (minutes)	9.58	9.73	9.82	9.89	9.95
City	10.80	10.64	10.47	10.38	10.87
Highway	8.75	8.75	8.75	8.75	8.75
Fuel economy (L/100km)	9.343	9.350	9.228	8.982	9.003
City	10.651	10.178	9.699	9.206	9.418
Highway	8.753	8.753	8.753	8.753	8.753
Cost components ($)					
(1) Gasoline cost	1.16	1.11	1.05	0.99	1.05
(2) Gasoline tax	0.00	0.00	0.00	0.00	0.00
(3) Social cost of carbon	0.11	0.11	0.10	0.10	0.10
(4) Value of time	1.76	1.79	1.81	1.83	1.81
Driver cost (1+2+4)	2.921	2.895	2.864	2.826	2.865
Social cost (1+3+4)	3.035	3.004	2.968	2.924	2.969

Note: Base case (identical to the previous tables) assumes there are five stops on the city route (and none on the highway route). The middle three columns shows the effect of reducing the number of city stops to four, three and two. The final column shows the effect of replacing the five city stops with five traffic circles, in which traffic slows to about 25 km/h to pass through the circle.

Figure 1: Distribution of mean fuel economy across trips

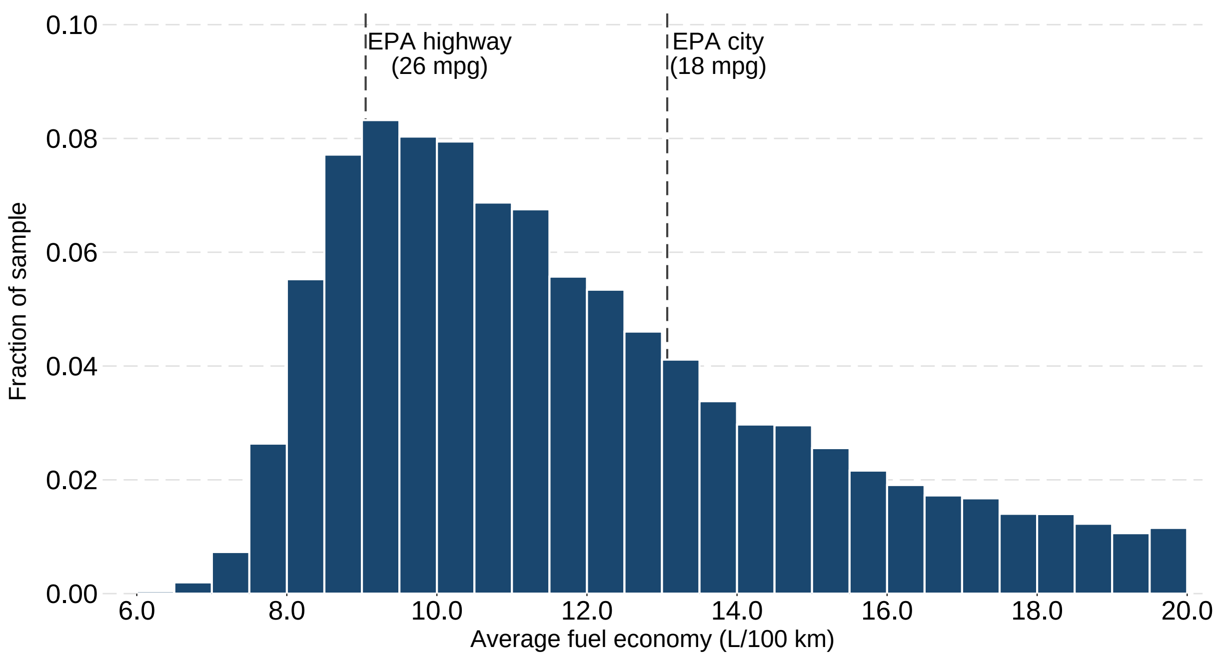

Note:

Figure 2: Distribution of mean fuel economy across drivers

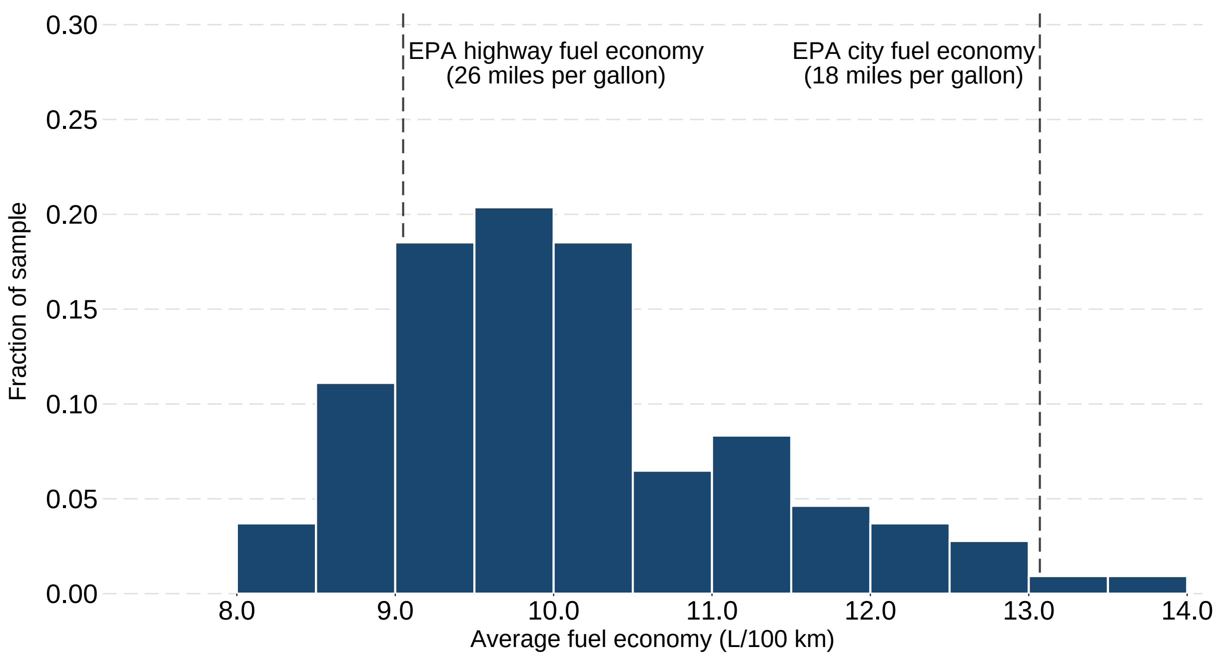

Note:

Figure 3: Observed and Predicted Relationship between Fuel Consumption and Speed

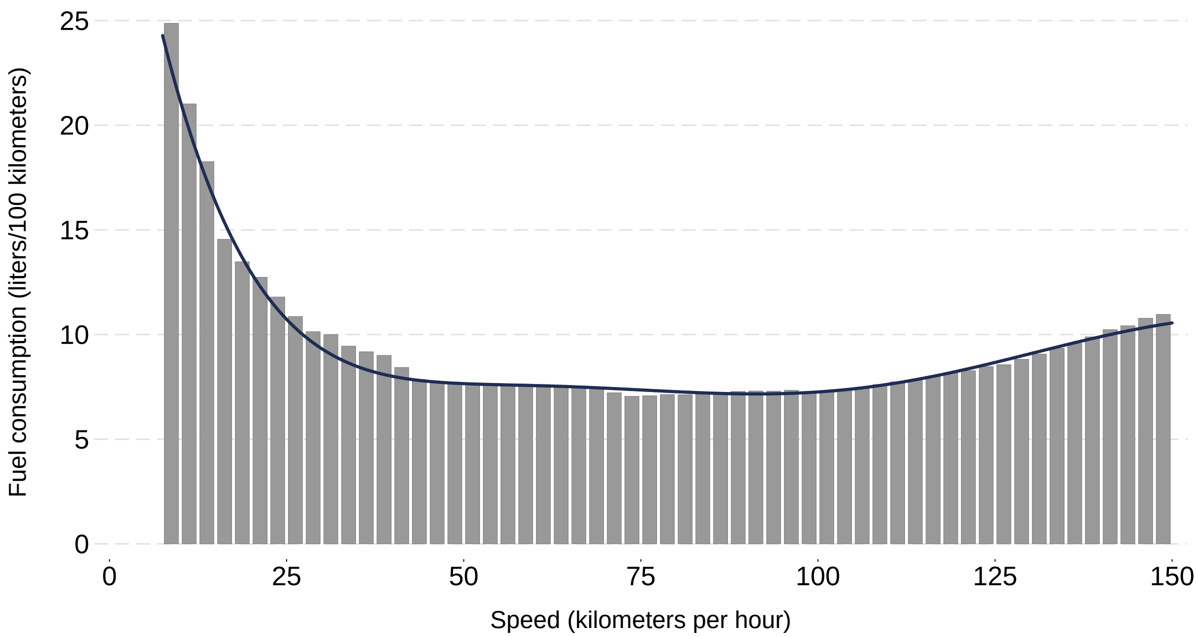

Notes: The gray bars show the mean fuel consumption for level driving at constant speed (defined as the absolute value of the grade less than 0.01 radians and the absolute value of acceleration less than 0.25 m/s^2), using 2.5 kilometer/hour speed bins. The function plot shows the predicted relationship between speed and fuel consumption, assuming zero acceleration and zero grade, based on the estimates in Table **??**.

Figure 4: Observed and predicted fuel consumption and observed speed for a single trip

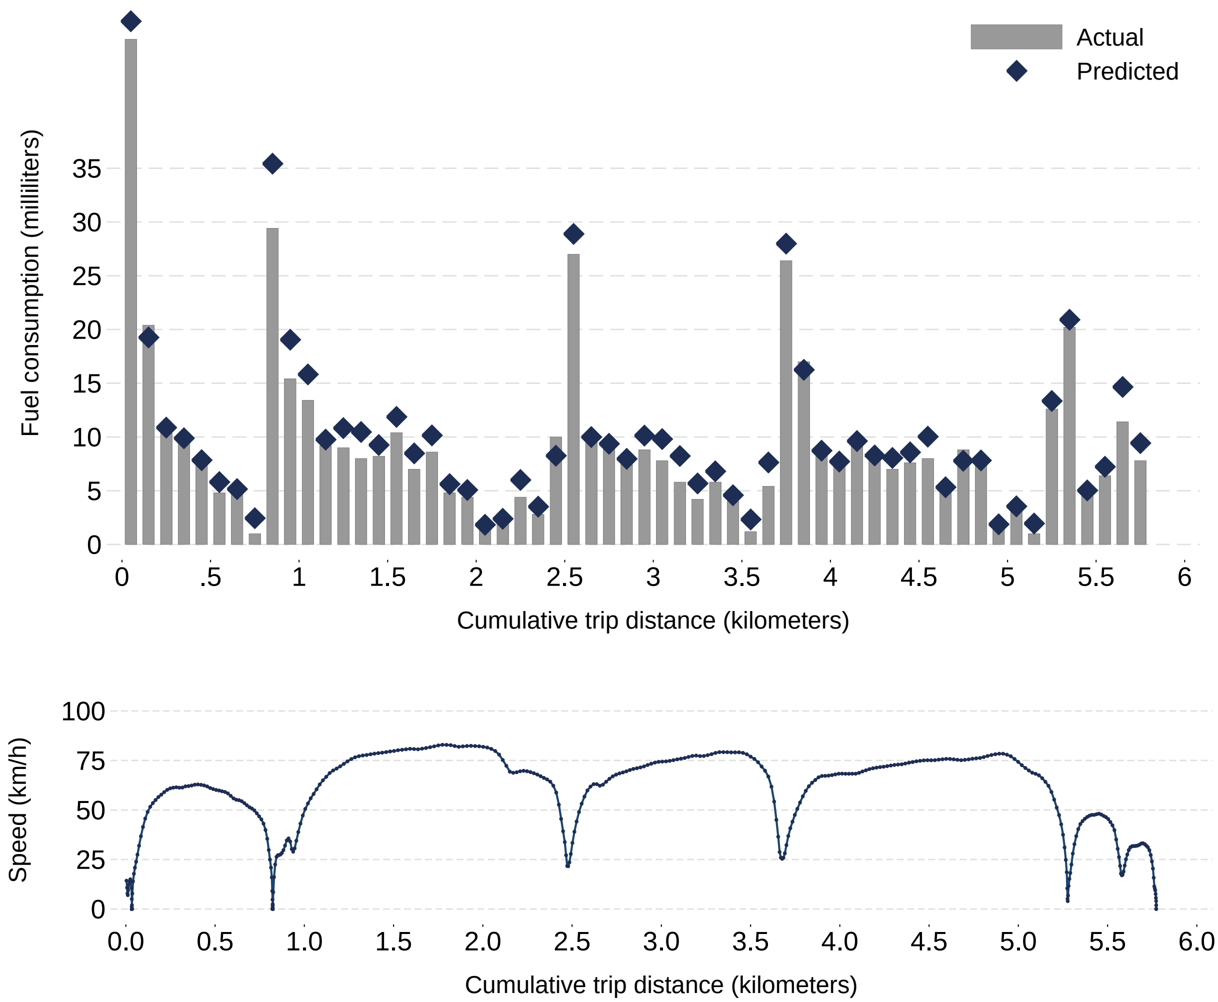

Note: The top figure shows the fuel consumption during one particular trip. The gray bars show the actual fuel consumption for each 100 meter segment of the trip. The diamond markers show the predicted fuel consumption based on the estimates in Table ??. The vertical axis shows fuel consumption in milliliters for the 100 meter segment; this also corresponds to fuel economy in L/100 km. The bottom figure shows the trace of speed (in kilometers per hour) for the same trip.

Figure 5: Observed and Predicted Relationship between fuel consumption and 2–15 m/s acceleration

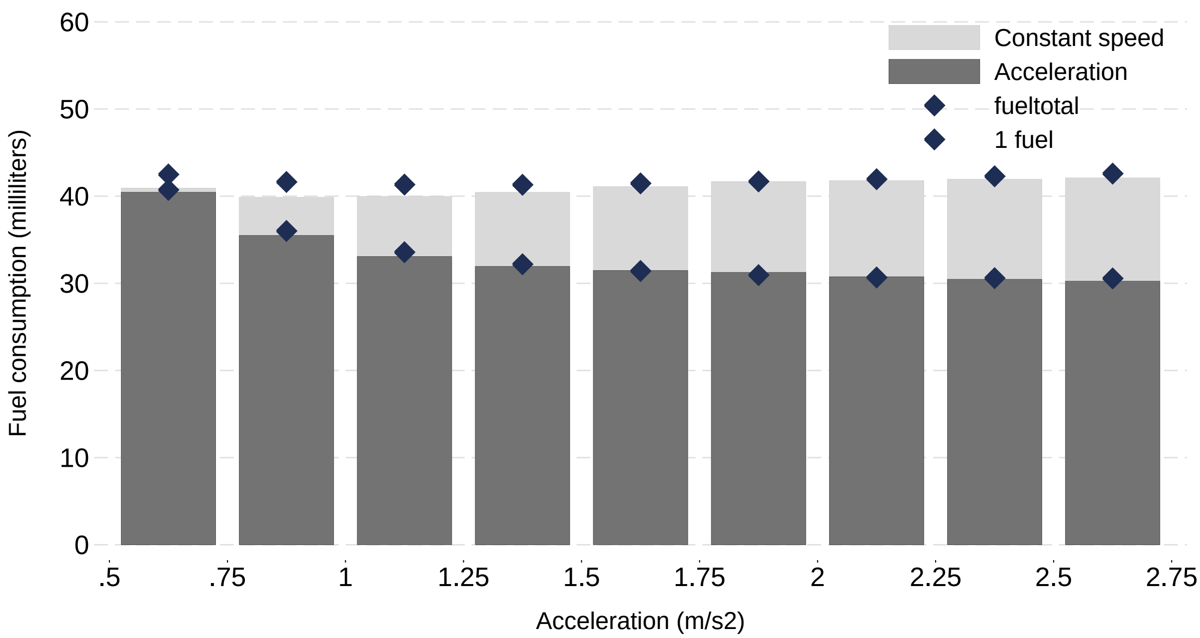

Notes: The dark bars show the mean fuel consumption for different rates of acceleration from 2 to 15 m/s. The lighter bars show the mean predicted fuel consumption for driving at a constant speed of 15 m/s until a total distance (including the acceleration period) of 250 meters has been covered. The diamond markers show the predicted fuel consumption, for the acceleration period and for the entire 250 meters, for acceleration rates at the midpoint of each bar.

Figure 6: Summary of driving simulation model for policy comparisons