Carpooling and Driver Responses to Fuel Price Changes: Evidence from Traffic Flows in Los Angeles

Antonio M. Bento,[1†] Jonathan E. Hughes[2] and Daniel Kaffine[3*]

December 21, 2012

Abstract

We examine carpooling and driver responses to fuel price changes. Using a simple theoretical model, we show that traffic flows in mainline lanes unambiguously decrease when fuel prices increase, and this effect is stronger when the presence of a carpool lane provides a substitute to driving alone. In contrast, in carpool (HOV) lanes flow can either increase or decrease. These predictions are tested using eight years of traffic flow data for 1,700 locations in Los Angeles. In our preferred specification, the mean elasticity of flow with respect to fuel price is 0.136 for HOV lanes. For a 10% increase in fuel price this implies 10 additional carpools per hour, \$8.8 million per year in additional congestion costs for carpoolers and \$11.3 million lower costs for mainline drivers. For mainline lanes, flow elasticities are -0.083 and -0.050 for highways with and without an HOV lane. These estimates imply that the mean highway with an HOV lane experiences a 30 percent larger decrease in hourly flow compared to the mean highway without an HOV lane. Flows in HOV lanes show an immediate decrease following a price increase but respond positively to price increases over time, which suggests time is an important input to carpool formation.

keywords: gasoline prices, traffic congestion, carpooling

[*]The authors thank Yahya Faissal Anouti, Martin Boileau, Harrison Fell, Terra McKinnish, Andrew Waxman and seminar participants at Cornell University, the Colorado School of Mines, the University of California Energy Institute and the Urban Economics Association.

[1]Charles H. Dyson School of Applied Economics and Management, Cornell University. [2]Department of Economics, University of Colorado at Boulder. [3]Division of Economics and Business, Colorado School of Mines.

†Corresponding author.

1 Introduction

In the absence of congestion pricing, by promoting carpool formation, high occupancy vehicle (HOV) lanes continue to serve as a second-best approach to mitigating traffic congestion in major urban areas. Proponents of HOV lanes also argue that carpooling in HOV lanes may help to reduce a variety of other social costs associated with driving, including local and global pollution and highway fatalities. While both HOV lanes and carpooling have received some attention in the theoretical literature,[1] empirical studies of carpooling and driver responses to fuel price changes remain sparse.[2] An important question is how carpool formation responds to changes in the cost of driving, in particular fuel costs. Increasing oil prices or policies such as renewable fuel standards, carbon taxes or cap and trade systems may impact fuel prices. Furthermore, some economists have argued for increased gasoline taxes as a measure to correct other driving-related externalities.[3] If higher fuel prices cause more drivers to form carpools, the result may be fewer vehicles on the road, lower congestion, fewer emissions and fewer fatalities.

In this paper we combine theory with empirical analysis to better understand carpool behavior. We begin by developing a simple theoretical framework for commuters' mode choice. Our approach is an equilibrium sorting model where commuters weigh fuel expenditures, time costs on congested highways, and fixed transaction costs of carpool formation against their preferences for driving and costs of alternative options. We find that, not surprisingly, higher fuel prices are associated with *decreases* in traffic flows in mainline lanes. However, in HOV lanes, higher fuel prices can result in either *increased* or *decreased* traffic flows.

The intuition for this result comes from two competing price effects on the number of drivers choosing the HOV lane. On one hand, higher fuel prices cause drivers who previously traveled alone to form carpools, increasing HOV lane flow. On the other hand, higher fuel prices raise the cost of

[1] For example, see Johnston and Ceerla (1996), Rodier and Johnston (1997) and Barth and Boriboonsomsin (2008) on HOV lane operations and emissions, and Kwon and Varaiya (2008) on HOV lanes and congestion. Lee (1984) presents a theoretical model for carpool formation. Giuliano, Levine, and Teal (1990) survey drivers to understand whether HOV lanes increase carpooling. More recently, Konishi and il Mun (2010) examine the conditions under which introducing HOV lanes is socially beneficial and consider potential welfare gains from converting HOV lanes to HOT lanes.

[2] An important exception is Ferguson (1997) who examines the declining trend in carpooling between the 1970s-1990s. In contrast to Ferguson (1997), we focus on the period of 2000-2007 which was characterized by increases in fuel prices. We exploit weekly traffic data derived from hourly measurements which greatly improves our ability to identify the relationships between fuel prices and carpooling, relative to Ferguson (1997) who used annual national aggregate data. Furthermore, the spatially disaggregated nature of our study also allows us to better capture heterogeneity in carpooling responses.

[3] These costs may be directly or indirectly related to gasoline consumption. Absent policies that price these externalities directly, Parry and Small (2005) and Parry, Walls, and Harrington (2007) propose a second-best fuel tax.

driving in general, resulting in more commuters choosing alternative options. We show that which effect dominates depends on the heterogeneity in preferences for driving relative to alternative options - increased flow in the HOV lane suggests commuters have a strong enough preference for driving such that higher prices induce substitution into the HOV lane instead of alternative options. For mainline lanes, we predict that drivers on highways with an HOV lane are more responsive to changes in fuel prices because HOV lane access increases the range of substitution options when fuel prices change. Because this decrease in mainline drivers is offset by an increase in HOV lane use, the net effect on traffic flow can be larger or smaller compared to a highway without an HOV lane, depending on preferences for driving and the number of lanes of each type on a given highway.

We test these predictions using eight years of data from Los Angeles California. We exploit detailed data on traffic flows from California's Freeway Performance Measurement System (PeMS) to construct weekly traffic flows for mainline and high occupancy vehicle (HOV) lanes at over 1,700 locations in Los Angeles.[4] We combine these data with weekly average retail gasoline prices in the region.

Across a variety of specifications and functional forms we find that higher fuel prices are associated with *higher* flows in HOV lanes (positive flow elasticity) and *lower* traffic flows in mainline lanes (negative flow elasticity). In our preferred specification, we estimate an average HOV flow elasticity of 0.136. This implies an average effect of 10 additional carpools per hour for a 10% increase in fuel price. We estimate average flow elasticities of -0.083 and -0.050 for mainline lanes on highways with and without HOV lanes.[5] For a 10% increase in price these elasticities imply decreases of 27 and 45 vehicles per hour, *i.e.* a larger decrease in flow on highways with an HOV lane, consistent with our theoretical prediction. Overall, our estimates suggest that *total* traffic flows decline more on the mean HOV highway compared with the mean highway without an HOV lane. While at first glance, these effects may appear small, even a change of 10 vehicles per hour can cause potentially significant welfare changes with important distributional consequences. For example, the implied change in the optimal toll for an average HOV lane due to a 10% increase in fuel prices is $0.013 cents per mile. This corresponds to an aggregate annual cost of $8.8 million for Los Angeles drivers who would pay this toll. If we consider only the benefit coming from increased carpooling, this equates to a decrease in congestion costs of approximately $11.3 million per year.

[4] High occupancy vehicle (HOV) lanes are highway lanes where access is restricted to vehicles carrying more than one occupant. In Los Angeles these restrictions apply 24 hours per day, 7 days per week. In LA, with the exception of I-10 which requires 3+ individuals in a carpool from 5-9AM and 4-6PM, all freeways with HOV lanes require 2+ individuals in a carpool.

[5] Across a variety of models our estimated flow elasticities range from 0.089 to 0.274 for HOV lanes and from -0.027 to -0.098 and -0.062 to -0.129 for mainline lanes on highways without and with HOV lanes.

We extend our main empirical results along several dimensions. First, we explore whether flow elasticities depend on the time of day or day of week of travel. We find that drivers are more responsive to fuel price changes during off-peak periods. We estimate flow elasticities between 0.267 and 0.372 for HOV lanes off-peak. For mainline lanes our estimates range from -0.093 to -0.108 and from -0.087 to -0.092 for mainline lanes on highways with and without an HOV lane off-peak. Flow elasticities during weekend peak and off-peak periods look surprisingly similar to weekdays, consistent with high levels of congestion in the Los Angeles area regardless of the day of the week.

Next, we investigate whether time is an input to carpool formation. We do this by estimating several models that aim to isolate the short-run and longer-run responses to fuel price changes. We find that flows in HOV lanes on average *decrease* in response to increases in the current price of fuel. However, conditional on this week's price, flows in HOV lanes tend to *increase* in response to increasing average prices. Flows in mainline lanes on highways without an HOV lane respond to changes in the current week's fuel price, but conditional on current prices do not respond to average prices. For mainline lanes on highways with HOV lanes, flows respond negatively to changes in current prices. The flow response is larger in magnitude when average prices are also rising. The effects in HOV lanes and mainline lanes on highways with HOV lanes are consistent with a short lag in drivers' responses to price changes when carpooling is an option.

Finally, we investigate heterogeneity in our mean elasticity estimates by calculating station-level flow elasticities for each location in our sample. While these estimates exhibit substantial variability, the distributions are consistent with our average parameter estimates. We attempt to explain differences in the estimated flow elasticities by correlating the point estimate for each station with observable characteristics of the freeway and surrounding area. For HOV lanes, we find that the (positive) flow elasticities increase as the distance from downtown increases. For mainline lanes we find that flow becomes less elastic the further the distance from downtown Los Angeles and that this effect may be associated with higher incomes in these areas.

Understanding the relationship between carpool formation and fuel prices is important for a number of reasons. First, carpooling reduces the number of vehicles on the road and can lead to reductions in driving-related externalities such as those described above.[6] Our analysis of Los Angeles shows that higher fuel prices do result in more carpools and fewer vehicles on the road. Second, local governments around the U.S. have made substantial investments in constructing HOV

[6] For example, price increases that result in decreased vehicle miles traveled can act to reduce local air pollution (Knittel and Sandler, 2011), traffic fatalities (Grabowski and Morrisey, 2006), or congestion (Burger and Kaffine, 2009).

lanes and policies to promote carpooling.[7] Nationwide there are over 1000 miles of operating HOV lanes in 27 metropolitan statistical areas.[8] HOV lanes act to encourage carpool formation by reducing travel time for commuters relative to drivers in mainline lanes that face higher levels of congestion and traffic delays. While an analysis of HOV lane construction and highway expansion is beyond the scope of this paper, we show that the presence of an HOV lane may make drivers *more responsive* to changes in fuel prices, which can result in larger reductions in driving-related externalities when fuel prices rise. Third, understanding carpool behavior provides insight into how commuters weigh fuel costs, time costs and the fixed costs associated with different transportation options. Our analysis suggests that commuters do take these costs into account.

This paper also contributes to the large literature estimating price elasticities of fuel demand. Dahl and Sterner (1991) and Espey (1998) survey a long literature dating back to the 1970s. More recent examples document shifts in the relationships between fuel prices and consumption (Hughes, Knittel, and Sperling, 2008) and the effects of price changes on emissions (Davis and Kilian, 2011). Because commuters' responses to fuel price changes can affect the social costs of driving, researchers have begun to focus on particular margins of response. For example, Bento et al. (2009), Busse, Knittel, and Zettelmeyer (2010), Klier and Linn (2010), and Li, Timmins, and von Haefen (2010) study fuel prices and vehicle purchase behavior, Blanchard (2009) and Currie and Phung (2007) investigate fuel prices and public transit ridership, while Puller and Greening (1999) focus on overall vehicle miles traveled. This study considers another important response to higher fuel prices in the form of carpooling.

The remainder of this paper is organized as follows. Section 2 describes our theoretical model for predicting driver responses to fuel price increases. Section 3 describes our traffic and fuel price data and Section 4 describes our empirical model. Our main results are presented in 5 and Section 6 concludes.

2 Theoretical framework

Here we present a brief theoretical model to conceptualize how changes in gas prices may affect traffic flow. Our analytical exercise utilizes a simple equilibrium sorting model where a fixed number of commuters select their cost-minimizing preferred commute option.[9] Equilibrium is determined

[7] Policy makers may view HOV lane construction as a tool to reduce traffic congestion and automobile related air pollution where congestion tolls are not feasible.

[8] HOV lane statistics are from the Federal Highway Administration at http://ops.fhwa.dot.gov/freewaymgmt/

[9] Our model of commuter choice is similar in spirit to that of Shewmake (2010). In addition to the model discussed here, a more formal model of household choice extending Parry and Bento (2001), is presented in the appendix of the

by the level of commuting where no agent has an incentive to deviate from their commute choice. We hold both overall highway capacity and the allocation of HOV lanes fixed.

We consider two freeway configurations: In the first, Scenario 1, commuters choose between traveling on L congestible mainline freeway lanes versus an uncongestible alternative commute option (i.e. public transit or telecommuting). In the second configuration, Scenario 2, commuters choose between $L-1$ congestible mainline freeway lanes, a single congestible HOV lane, and an alternative uncongestible commute option. We allow drivers to have heterogenous preferences between driving and the alternative option, where the perceived cost of the alternative option is uniformly distributed across potential commuters. We stress that the primary purpose of this section is to build intuition regarding gas price responses for the empirical exercise to follow.

We assume $\bar{N}$ total commuters choose between congestible driving and the alternative option. Drivers pay gas expenses $p \cdot g$ (where p is gas price and g is the exogenous quantity of fuel consumed to travel the commute), and incur a travel time cost T, which is a function of the number of drivers choosing to use that highway and the value of time spent commuting. Costs for the alternative option A are uncongestible and are assumed to be distributed uniformly with mean $\bar{A}$, such that $A \sim U[\bar{A}-b, \bar{A}+b]$ where b represents an arbitrary bound of the distribution.[10]

2.1 Scenario 1: L congestible mainline freeway lanes and an uncongestible outside alternative

Consider a congestible mainline highway with L lanes. Let N_m be the number of mainline drivers, C_m the number of mainline cars, and N_a the number of commuters using the alternative option. Drivers value travel time at w, and, for tractability, travel time in a lane is assumed to have the linear form $T = \alpha \cdot C_m$, such that travel time increases as more drivers use the highway.[11] The parameter α captures the effect of adding an additional car to a lane, and thus, with L mainline lanes, travel time will take the form $T = \frac{\alpha}{L} \cdot C_m$. Finally, we note that all commuters must choose a highway and mainline drivers are assumed to travel alone, such that $N_m + N_a = \bar{N}$ and $N_m = C_m$.

Commuters will compare the cost of driving in the L mainline lanes against their particular

working paper version of this paper and yields similar results. To verify the analytical results from the equilibrium sorting model, simulation models were run, confirming the derivations below (details available from the authors upon request).

[10] Costs for the alternative option can be thought of as fares and travel time cost plus any costs incurred by not commuting by car. For example, comfort or lost "face time" due to time away from the workplace.

[11] Consistent with the transportation economics literature we consider the value of time w in terms of an equivalent fraction of the wage rate.

draw of the perceived cost of the alternative option as:

$$\min\{p \cdot g + w \cdot \frac{\alpha}{L} \cdot C_m, A \sim U[\bar{A} - b, \bar{A} + b]\} \tag{1}$$

If $\hat{A}$ represents the marginal commuter who, for a given level of congestion, is indifferent between the mainline and alternative option, then commuters with $A > \hat{A}$ will sort to the mainline, while commuters with $A < \hat{A}$ will sort to the alternative option. Given this ordering and the uniform distribution over the alternative option, equilibrium in this system requires that:

$$p \cdot g + w \cdot \frac{\alpha}{L} \cdot C_m = \gamma \cdot N_a + \bar{A} - b \tag{2}$$

Given the assumption of a uniform distribution, costs of the alternative option take on a convenient linear form where $\gamma = \frac{2b}{N}$ reflects how quickly perceived costs rise across commuters. Larger b corresponds to more heterogenous costs and a steeper rise across commuters in the costs of the alternative option, while as b approaches 0, γ goes to zero and costs approach a constant $\bar{A}$. Solving for the equilibrium for each commute option and differentiating with respect to gas prices yields: $\frac{dC_m}{dp} = -\frac{gL}{w\alpha + \gamma L}$ and $\frac{dN_a}{dp} = \frac{gL}{w\alpha + \gamma L}$.

2.2 Scenario 2: one congestible HOV lane, $L-1$ congestible mainline lanes, and an outside alternative

Now suppose that one of the L freeway lanes is in fact an HOV lane, such that there is one congestible HOV lane and $L - 1$ congestible mainline lanes. We assume that carpool drivers incur a constant transaction cost of τ to form a carpool, but split the cost of gas evenly.[12] Let C_h be the number of carpools using the HOV lane, and N_h be the number of carpoolers, such that $N_h = 2C_h$. We again assume travel time in a lane takes a linear form, such that in the HOV lane, travel time is given by $T = \alpha \cdot C_h$, and in the mainline lanes, travel time is given by $T = \frac{\alpha}{L-1} \cdot C_m$. Finally, note that $N_h + N_m + N_a = \bar{N}$. Commuters will compare the cost of carpooling in the HOV lane against the cost of driving in the $L-1$ mainline lanes against their particular draw of the perceived cost of the alternative option as:

$$\min\{\frac{p \cdot g}{2} + w \cdot \alpha \cdot C_h + \tau, p \cdot g + w \cdot \frac{\alpha}{L-1} \cdot C_m, A \sim U[\bar{A} - b, \bar{A} + b]\} \tag{3}$$

[12]The inclusion of a transaction cost of carpool formation simply reflects the fact that HOV lanes typically have lower travel times than mainline lanes during congested time periods. In the absence of a transaction cost of carpool formation, that wedge in travel time could not persist.

As above, in equilibrium there will exist a commuter with perceived costs of $\hat{A}$ such that commuters with $A < \hat{A}$ will use the alternative option, while commuters with $A > \hat{A}$ will either use the HOV lane or mainline lanes. Equilibrium is determined by:

$$\frac{p \cdot g}{2} + w \cdot \alpha \cdot C_h + \tau = p \cdot g + w \cdot \frac{\alpha}{L-1} C_m = \gamma \cdot N_a + \bar{A} - b \tag{4}$$

Solving and differentiating with respect to gas prices yields: $\frac{dC_h}{dp} = -\frac{g(w\alpha+(1-L)\gamma)}{2w\alpha(w\alpha+(1+L)\gamma)}$, $\frac{dC_m}{dp} = -\frac{g(L-1)(w\alpha+\gamma)}{w\alpha(w\alpha+(1+L)\gamma)}$ and $\frac{dN_a}{dp} = \frac{gL}{w\alpha+(1+L)\gamma}$.

2.3 Summary of model predictions

Here we summarize the above results, presented in Table 1. In the case of Scenario 1 with only mainline lanes and an alternative option, increasing gas prices decreases mainline drivers and increases the use of the alternative options, as expected. The magnitude of the price effect on mainline use is influenced by the heterogeneity in the perceived cost of the outside option, reflected in the parameter γ. As γ increases, the effect of an increase in gas prices is muted as there is less substitution away from the mainline freeway towards the outside alternative when gas prices rise.

Scenario 2 changes the freeway configuration to include one HOV lane, $L-1$ mainline lanes and an alternative option. We first consider the case of a constant alternative option cost, when $\gamma = 0$, In this case, rising gas prices cause flow in both the HOV and mainline lanes to unambiguously decrease, with a smaller decline in the HOV lane. This is intuitive as the cost of using the HOV lane does not increase as much as the cost of using mainline lanes when gas prices increase. Nonetheless, with constant perceived cost of the alternative option, HOV lane use will unambiguously decrease as gas prices increase. However, as γ increases, there is less substitution towards the outside alternative when an HOV lane is present. Intuitively, we can think of the heterogenous costs of the alternative option as representing a preference for driving, and because the HOV lane represents a relatively cheaper driving alternative to the mainline, fewer total drivers substitute out of driving entirely. Furthermore, HOV flow may actually increase due to an increase in gas prices, depending on the relative magnitude of $w\alpha$ versus γ. Larger γ corresponds to a larger increase in HOV lane flows - the intuition is that if the perceived cost of the alternative option increases fast enough across commuters, then enough commuters will substitute out of the mainline lane and into the HOV lane (instead of into the alternative option) leading to an increase in HOV flow. Similarly, for sufficiently large γ, the mainline response will be larger in the presence of an HOV lane than in Scenario 1, as the presence of an HOV lane provides a substitute for the alternative option.

We can also compare the magnitude of the price effect of total flow (and thus say something about changes in externalities like emissions) when an HOV lane is present or not. Subtracting the total flow effects in Scenario 1 from Scenario 2 gives: $\frac{g(w^2\alpha^2+(1+2L)w\alpha\gamma-(L-1)L\gamma^2)}{2w\alpha(w\alpha+L\gamma)(w\alpha+(1+L)\gamma)}$. When γ is small, the negative term in the numerator is also small, and the change in total vehicles is smaller when an HOV lane is present. However, with sufficiently large γ, the numerator is negative, and the change in total vehicles is *larger* when an HOV lane is present. The intuition is that if γ is sufficiently high, then when prices rise in the absence of an HOV lane, few drivers will stop driving, leading to small changes in total vehicles. In contrast, when an HOV lane is present, commuters still want to continue driving despite higher fuel prices, but they instead choose to drive in the HOV lane. Thus, which configuration leads to a greater change in vehicles boils down to whether or not having a substitute (which still involves some vehicle use) is better than when people are limited to only substituting towards the alternative option (which does not involve vehicle use).[13]

Looking towards the empirical exercise, the model above provides several hypotheses to explore. First, if the change in HOV lane flow due to a gas price increase is positive, then this suggests that commuters do not have identical perceived costs of the alternative option - instead, commuters who prefer to drive will substitute from the mainline lane into the HOV lane rather than use the alternative option. Second, if this preference for driving is strong enough, then mainlines with an accompanying HOV lane should see a larger negative flow response to gas price changes, as the HOV lane provides a lower-priced driving option relative to the mainline lane.[14] Finally, given the ambiguous predictions of the analytical model, we can also explore how the change in total vehicles varies with freeway configuration to determine whether the decline in vehicles on mainline-only freeways is larger or smaller than the net vehicle change on mainline plus HOV freeways.

Some additional discussion is warranted in terms of relating the additional empirical results (to be presented below) with the analytical exercise above. First, the response to gas prices at different points in the city may be very heterogeneous. Longer commutes are associated with higher gasoline expenditure, and thus we would expect more distant drivers to be more responsive to higher gas prices than those with shorter commutes (all price effects in table 1 are increasing in g). Similarly, higher values of time tend to mute the response to gas price increases (all price effects in table 1 are

[13]This can also be seen from taking the limit of the price effects in table 1 as $\gamma \to \infty$. In Scenario 1, the limit of the mainline price effect is zero, while in Scenario 2, the limit of the total price effect is $g(1-L)/(w\alpha(1+L)) < 0$. Essentially as substitution towards the alternative option goes to zero, having an HOV lane available would clearly reduce total vehicles, while not having an HOV lane would not.

[14]Our model does not formally consider substitution across highways or substitution between highways and surface streets. Rising fuel prices that reduce congestion enough to entice new drivers to mainline lanes from alternate highways would bias our empirical estimates against finding an effect of fuel prices on mainline flows. For HOV lanes, rising fuel prices that cause drivers who previously traveled via surface streets to form carpools and use the HOV lane would tend to inflate the flow elasticities we estimate in the following sections.

decreasing in w). Access to public transit can also vary throughout the city, influencing commuters willingness to substitute out of driving entirely when gas prices increase. Second, in the above model, the number of commuters is held fixed, while in reality, the number of commuters varies throughout the day and over the course of the week. Intuitively, we would expect more discretionary trips to be taken outside of the main workday commute, thus increasing the response to higher gas prices (for example if higher prices encourage more people to retail online rather than drive to the shopping mall). Finally, the model presented above can be thought of as a steady-state, long-run equilibrium; however, the response by commuters to higher gas prices may not be immediate. In particular, formation and dissolution of carpools may not be an instantaneous action that can be taken in response to short-run fluctuations in gas prices.

3 Data

Our empirical analysis exploits detailed data on traffic flows from the California Freeway Performance Measurement System (PeMS). PeMS aggregates data from the state-wide network of highway vehicle sensors (detectors) and reports traffic flows and average vehicle speeds in intervals as small as 30 seconds at over 25,000 locations state-wide. We focus on highway traffic flows in Los Angeles and Ventura counties in Southern California. Los Angeles is widely know to have severe traffic congestion and also one of the most extensive networks of HOV lanes in the nation.[15] We construct a balanced panel of traffic flows at 1,727 stations over the period from July of 2000 through December of 2007.[16] Stations fall into one of three categories, HOV lanes (HV), mainline lanes (ML) or mainline lanes on sections of highway that also contain HOV lanes (ML_HV).[17] We observe 442 HV stations, 173 ML stations and 1,112 ML_HV stations.[18] Because we observe weekly fuel prices, described below, we sum hourly flows to the station-week level.[19]

[15]Other major cities, for example San Francisco and Sacramento, have many fewer HOV lanes and in many cases lanes may not be continuous along major highways. Los Angeles County alone has 36% of all HOV lane-miles in the state of California, with roughly two-thirds of all state HOV lane-miles located in the five county region of Los Angeles, Ventura, Orange, Riverside and San Bernadino counties.

[16]We drop observations for detectors located on on-ramps, off-ramps and freeway interchanges.

[17]Mainline stations are classified as ML_HV if an HOV lane exists somewhere along the highway. We classify I-5 stations as ML_HV because the presence of HOV lanes in Orange County likely affected traffic flows in our sample even though I-5 in Los Angeles did not contain HOV lanes prior to 2008. Reclassifying I-5 as a ML highway yields similar results to those presented here.

[18]In our specifications that drop imputed observations, our sample falls to 1,660 stations with 432 HV, 165 ML and 1,063 ML_HV.

[19]This approach also substantially reduces the computational burden of our analysis by reducing the sample size from approximately 114 million observations to under 700,000 observations. More importantly, this eliminates the need to model complex daily and hourly station level-trends which, using mean effects, would require the estimation of over 100,000 additional parameters. Furthermore, given that households fill their gas tank roughly once a week, it

Because we are primarily interested in transportation demand during periods when highways may be congested, our main results focus on the "peak period" from 6:00 am through 8:00 pm, Monday through Friday.[20] We do this for two reasons. First, externalities such as congestion and air pollution are likely to be more severe during these periods. Understanding carpooling behavior at these times is important for policy makers. Second, observing the population of carpoolers across any large geographic area is difficult. Here we use flow in HOV lanes as a proxy for the total number of carpools on the freeway. Because HOV lanes offer a travel time advantage relative to mainline lines during congested periods, it is likely that the vast majority of carpools use these lanes.

Table 2 presents summary statistics for the weekday peak, weekday off-peak and weekend periods. Gas prices in constant 2005 dollars range between \$1.17 and \$3.30 per gallon.[21] Mean station-level flows range from approximately 309,000 vehicles during the weekday peak period to 28,000 for the weekend off-peak periods. Since the different periods sum over different lengths of time, we also report average hourly flows as a more intuitive measure of traffic flow. Average hourly flows per lane range between 403 and 1,150 vehicles per hour. Comparing weekends and weekdays, mean hourly flow during the peak period is 1,150 vehicles per hour for weekdays and 999 for weekends. Off-peak hourly flows are also similar for weekdays and weekends at approximately 400 vehicles per hour. This suggests that traffic can still be quite congested in Los Angeles even on weekends and that the period of the day is perhaps a more important determinant of the level of congestion than the day of the week.

Figure 1 plots average hourly flows per lane across all stations of a given type versus fuel prices. Gasoline prices in Los Angeles begin relatively flat but increase steadily starting in 2002. The time-series of prices shows strong seasonal patterns with prices increasing throughout the spring, in general remaining high during the summer months, and decreasing in the fall. Aggregate traffic flows also exhibit similar seasonal patterns with greater flows during the summer months. Focusing on the later period of the sample with higher mean fuel prices, flow in mainline (ML) lanes (panel a) and mainline lanes in highways with HOV (ML_HV) lanes (panel b) appear to decrease slightly during the period from 2006 to 2008.[22] Flows in HOV (HV) lanes (panel c) increase substantially from 2006 to 2008.

seems reasonable to assume that commuters respond to changing gas prices on a weekly basis.

[20]Because different highways may have different congestion patterns depending on, for example, the direction of flow, we also experiment with off-peak and peak periods defined by the first and fourth quartiles of station-level mean hourly flows, respectively. These results are presented as a robustness check in Section 5 below.

[21]Prices are adjusted to constant 2005 dollars using CPI deflators for all goods in Los Angeles, Riverside and Orange counties from the U.S. Bureau of Labor Statistics.

[22]The drop in total mainline flows seen in January of 2005 may be the result of imputation which increased during this period. Our main specifications with year fixed effects account for level shifts across years. Estimating models for the periods before and after January 2005 produce results similar to those presented here.

We next investigate the time series properties of our data. Because both traffic flows and fuel prices may be trending over time, one may be concerned about the possibility of spurious relationships in our analysis of traffic flows and fuel prices. We would like to test whether the gasoline price and flow time series are stationary. For fuel prices, since we observe only city-level prices, we have a single time series. We conduct Elliott, Rothenberg, and Stock (1996) unit-root tests for the gasoline price series in levels and logs. Focusing on the Ng-Perron optimal lags, we reject the presence of a unit-root in gasoline prices at the 5-percent level, but fail to reject the presence of a unit-root in log gasoline prices.[23] For the station-level flow series, the problem is complicated somewhat by the panel nature of our data and the high likelihood of both serial correlation and cross-sectional dependence across stations in our sample.[24] We conduct augmented Dickey-Fuller Im, Pesaran, and Shin (2003) unit-root tests for both flow and log flow. We assume either one lag or four lags and subtract the cross-section means from the series to help mitigate the impact of cross-sectional dependence. We reject the null hypothesis that all panels contain unit roots in both cases at the 1 percent level. Similarly, we also conduct a Hadri LM test where the null assumes that all series are stationary around a linear trend. In this case, we reject the null that all series are stationary at the 1 percent level.

From these tests we conclude that at least a substantial fraction, though not all, of the station-level series are stationary. To verify this, we iterate over all the series in our panel and conduct separate Elliot-Rothenburg-Stock unit-root tests for each station. For 1,317 stations or 76 percent of our sample, these tests fail to reject the presence of a unit root in the flow series at the 5 percent level.[25]

We note that unit root tests of individual cross-sectional units are generally thought to lack power when applied in a panel setting. Therefore, the results of our Elliot-Rothenburg-Stock tests are likely an upper bound on the number of non-stationary series.[26] Furthermore, the presence of cross-sectional correlations in our traffic flow data may also reduce the power of our panel tests.[27] Because the true extent of non-stationarity in our data is unknown, we proceed with caution, treating all price and flow series as stationary. As a robustness check, Appendix A verifies that

[23] DF-GLS tau test statistics of -3.351 and -2.420, respectively.

[24] Traffic flows along highways are highly correlated. Furthermore, traffic may be correlated across alternate highways traveling in the same direction.

[25] We fail to reject a unit root for 86 percent of ML station, 82 percent of ML_HV stations and 58 percent of HOV stations.

[26] For example, Levin, Lin, and Chu (2002) investigate the asymptotic and finite-sample properties of unit root tests in a panel setting.

[27] For a thorough review of issues related to unit root tests in this setting the reader is referred to Maddala and Wu (1999).

using only the stations with stationary series produces results that are qualitatively similar to those presented below. We return to these issues in Section 5.3 where we explore the dynamic properties of traffic responses to fuel price changes and present results estimated in first differences and several alternate specifications.

A final issue relates to imputation in the PeMS data. PeMS engineers rely on two methods of imputing flows when station malfunctions result in missing data. Spatial imputation uses contemporaneous observations from nearby stations to infer missing flow data. Temporal imputation forecasts missing flows using past observations. Spatial imputation may bias the standard errors of our parameter estimates by introducing spatial correlation in the disturbance terms. Temporal imputation may bias our parameter estimates by removing any influence of fuel prices on flows in the forecasted values. We address imputation in two ways. First, our estimates below report cluster robust standard errors clustered at the highway-direction level in 5 mile increments to account for spatial common shocks.[28] Second, our preferred estimates drop observations that include on average more than 25% imputed observations in a given week. As a robustness check we estimate our model including all imputed observations. Results are qualitatively similar to the restricted sample.

4 Empirical model

In the theoretical model in Section 2, equilibrium flow in each lane is a function of gas prices p, value of time w, fuel consumption over the commute distance g, carpool formation costs τ, congestibility of the freeway α, and driving preferences γ. Because many of these features are unobserved and we are ultimately interested in the change in flow due to changes in gas prices, a reduced-form empirical approach is adopted whereby these unobserved variables are modeled as mean effects.

Our preferred empirical model regresses the log of weekly flow for station i in week t on the log of fuel price. In addition, we explore functional forms where either fuel prices alone or fuel prices and flows both enter in levels. Our main results pool stations of all three lane types. We estimate average flow elasticities for ML, ML_HV and HV lanes by interacting an indicator variable for each type with log fuel prices. In alternate specifications presented below, we also calculate station-level elasticities by interacting a set of station dummies with fuel prices, assuming common trends and seasonal effects. In a second specification, we allow each station to have unique week

[28] We are unaware of empirical evidence for the distance over which commuting choices may be correlated. Clustering at the station level, in 10 mile intervals or at the highway-direction level produce standard errors similar to those reported below.

effects and estimate station-level elasticities via separate OLS regressions. In each case, we observe mean elasticities for each lane type that are qualitatively similar to our base model.

Equation 5 below presents our preferred specification. D_l is a set of indicator variables for each lane type. The main coefficients of interest are given by the interaction $D_l \times \beta_1$. We model unobserved factors that affect traffic flows at the station-level as mean effects ϵ_i. We model unobserved trends that affect traffic flows as year effects ϵ_y and model seasonal effects using a set of 52 week dummy variables, ϵ_w.

$$ln(flow_{it}) = \beta_0 + \sum_{l=1}^{3} D_l 1(type = l)_i \times [\beta_1 ln(Pgas_t)] + \epsilon_i + \epsilon_w + \epsilon_y + \epsilon_{it} \quad (5)$$

In a series of alternate specifications below, we explore the robustness of our results to omitting year effects, week effects and adding a time trend. In addition, we explore the possibility that our time and seasonal controls do not adequately account for the potential endogeneity of fuel prices and instrument for Los Angeles gasoline prices with oil prices.

5 Empirical results

5.1 Average peak-period flow elasticities

We begin by estimating the average flow elasticities with respect to gasoline prices using our pooled data. Table 3 presents results from our preferred specification, Equation 5, as well as several alternate specifications. Model 1 includes station effects and lane-type interactions but excludes time effects. Model 2 adds week effects to capture weekly patterns in traffic flow common to all stations. Model 3 is our preferred specification which includes common week and year effects. Model 4 replaces year effects with a linear time trend.[29]

The point estimates across these four specifications are qualitatively similar.[30] Flow elasticities in the mainline lanes are *negative* and statistically significant in specifications that account for common seasonal and time effects. Flow elasticities are *larger* in magnitude for mainline lanes on highways with an HOV lane (ML_HV) compared with other mainline lanes (ML) consistent with the predictions of our theoretical model.[31] The mainline flow elasticities can be interpreted

[29] In our robustness checks we investigate higher order time trends to more flexibly capture changes in traffic flows over time. These results are very similar to the estimates presented in Table 3.

[30] Flow elasticities estimate using daily rather than weekly traffic data also yield similar results.

[31] We reject the null hypothesis that the estimated elasticities are equivalent with a *p-value* of 0.08.

as weekly VMT elasticities and are comparable to recent estimates for the short-run elasticity of gasoline demand. For example, Hughes, Knittel, and Sperling (2008) estimate a short-run price elasticity between -0.03 and -0.08.

For HOV lanes, theory suggests that the flow elasticity may be positive or negative depending on the preferences of commuters for driving. Here, we find the estimated flow elasticities are *positive* and statistically significant in each of the four specifications. One interpretation of this result is that distaste for not driving is sufficiently large that when fuel prices increase, a larger share of commuters would rather carpool than choose the alternative option.

In our main results, we attempt to limit the impact of imputation in the PeMS data by dropping all weekly observations where the mean level of imputation exceeds 25 percent. Model 5 relaxes this restriction and includes all observations regardless of the level of imputation. When all imputed observations are included, the estimated mainline elasticities are quite similar to the main results in Model 3. However, the estimated flow elasticity for HOV lanes is substantially larger.[32] While it is comforting that allowing for more imputation does not qualitatively change our findings, in what follows we will limit the level of imputation in our data to below 25 percent.[33]

Looking across the specifications, the estimated mainline elasticities are small in magnitude for models without week and time effects. This is consistent with a bias towards zero introduced by endogenous prices if we fail to account for unobserved (flow) demand shocks. Models 3 and 4 account for mean changes in flow across stations at the week and year level, which appears to at least partially account for demand shocks. To further investigate this issue, Model 6 of Table 3 presents 2SLS estimates where oil prices are used to instrument for the price of gasoline. Oil prices are plausibly orthogonal to traffic patterns in Los Angeles and also correlated with local gasoline prices.[34] The 2SLS elasticities are comparable to those in the OLS specifications, though compared with Model 3 the mainline elasticities are smaller in magnitude and the HOV elasticity is somewhat larger. Because the estimates appear quite similar and due to the computational cost of estimating 2SLS in this setting, we limit our analysis in the sections below to OLS. We focus on specifications that include year and week effects, similar to Model 3.

[32]If traffic flows only respond to current prices then we would expect temporal imputation to bias the estimated flow elasticities toward zero. However, we provide evidence below that time may be an important input into carpool formation. One possible explanation of the larger estimated HOV effect is that imputed observations capture mainly the (positive) flow response to longer run trends in prices but largely miss the (negative) short run response, introducing a positive bias.

[33]Because stations may use imputed data for short periods throughout the day, and because we are aggregating station-level data over all hours in a given week, limiting imputation completely would substantially reduce our sample size.

[34]For oil prices we use the Cushing, OK spot price for WTI crude.

Appendix A investigates the robustness of our main parameter estimates along several additional dimensions. First, we verify that our results our similar using only stations for which we can reject the presence of a unit root in the flow data. Next, we explore implications of our assumed specification in natural logarithms. We also check robustness to the inclusion of higher order time trends. Finally, we explore the effects of different HOV lane vehicle occupancy requirements and whether HOV lanes are continuous. In each case, the estimated flow elasticities are quite similar to those in our preferred specification.

5.2 Peak versus off-peak periods

Both our theoretical model and main empirical results consider the effects of changes in fuel prices on traffic flows during peak periods when traffic is congested. However, we may observe different behavior in the absence of congestion or during periods when costs of the alternative option are different. To explore behavior in other periods, we construct weekly traffic flows for peak and off-peak periods on both weekdays and weekends. Our main results assume the weekday peak period is from 6am to 8pm on Monday through Friday. We define weekday off-peak as 8pm to 6am, Monday through Friday. Weekend peak and off-peak periods use equivalent hours of the day for traffic flows on Saturday and Sunday. Because individual stations may have different peak periods depending on, for example, their distance from downtown or the direction of travel, we alternatively define weekday "peak hours" for each station as the 6 hours with highest average hourly flow and a weekday "off-peak hours" as the 6 hours of lowest average flow.[35] Table 4 summarizes the mean elasticities by lane-type during each of these periods.

We begin with the weekday peak and off-peak periods. The estimated mainline elasticities are larger in magnitude during the off-peak period compared with the peak period. For mainline lanes and mainline lanes on HOV highways, the estimated flow elasticities are -0.092 and -0.108. For HOV lanes, the estimated average flow elasticity is 0.267. This is consistent with a larger share of off-peak travel being discretionary, and therefore more able to respond to changes in fuel prices. For HOV lanes, the estimated elasticity is nearly twice as large during the off-peak period. Our finding that drivers are more responsive to fuel price changes during off-peak periods, potentially reflects a larger share of discretionary trips and lower transaction costs of carpool formation when traveling with friends or family. However, because drivers of carpools may not preferentially select

[35]Given that congestion levels are high throughout the day, the definitions in columns 1 and 2 capture behavior when highways are congested, though not necessarily at the highest levels of congestion. The quartile definitions in columns 5 and 6 can be thought of as robustness checks for the most (and least) congested periods, or peak of the peak defined by the 6 hours with highest (lowest) average flows.

HOV lanes during uncongested periods, our ability to draw conclusions based on these results is limited.[36]

The trends are similar for the weekend peak and off-peak periods. During peak times, the flow elasticities for the three lane types are comparable to the weekday estimates at -0.058 and -0.077 for ML and ML_HV, and 0.168 for HOV lanes. This is somewhat surprising given the potential differences in the types of trips made during the week and on weekends. However, the similarity of the estimates suggests that congestion and time costs during peak hours may be similar on both weekdays and weekends.[37] During the off-peak period, the estimated elasticities for mainline lanes are larger in magnitude compared to weekdays. This is again consistent with a greater portion of this travel being discretionary.

Finally, we investigate our definitions of peak and off-peak. The alternate definition of "peak hours" results in estimates comparable to our main results, though the magnitudes appear slightly smaller. Here as before we reject the null hypothesis that the estimated elasticities for ML and ML_HV lanes are equal.[38] For off-peak hours, the mainline estimates are similar to those using the common off-peak definition though the estimated flow elasticity in HOV lanes is substantially larger.

5.3 Dynamic properties of traffic responses

In this section we explore dynamic responses of traffic flows in mainline and HOV lanes to changes in fuel prices. Section 2 notes that carpool formation may require time to identify other commuters with whom to share a ride. Furthermore, Section 3 suggests both fuel prices and traffic flows, in a substantial fraction of the stations, are trending over time. Here, we further explore these issues. Our motivation is twofold. First, we hope to identify both short-run and longer-run components of the response to fuel price changes. These effects may vary by lane type. Second, we attempt to verify that persistent changes in fuel prices are important for both carpool formation and the associated decreases in flows in mainline lanes on highways with HOV lanes.

Table 5 show the results of several specifications intended to isolate the short-run and longer-run impacts of fuel prices on traffic flows. The first column contains our main specification from Section 5.1 which relies on within year variation in fuel prices to identify average flow elasticities. In Model 7, we replace weekly fuel prices with monthly average fuel prices constructed as the

[36]For instance, a large share of carpools may travel in mainline lanes off-peak.

[37]This is supported by the high flow levels we observe in the summary statistics for weekend peak hours.

[38]We reject the null hypothesis that these coefficients are equivalent with a *p-value* of 0.06.

simple average of the previous four weeks (*i.e.* excluding the current week). The intent is to limit the impact of weekly variation in prices and isolate aggregate trends that may influence long-run behavior. If consumers are responding to past trends in fuel prices we might expect larger flow elasticities. In this case, the estimated mainline elasticities are *smaller* in magnitude relative to our base model, though the difference is not statistically significant. This is consistent with a downward bias introduced by measurement error in fuel prices which one might expect if consumers are in fact responding to weekly price changes. For HOV lanes however, the estimated flow elasticity is *larger* than the base model, though again the difference is not statistically significant. This provides suggestive evidence consistent with carpool formation occurring over a period longer than one week.

The next two models attempt to investigate the short-run traffic responses to changes in fuel prices. We take two different approaches to de-trending the flow and fuel price data. Model 8 presents estimates in first-differences where we subtract from each observed price and flow, both in natural logs, the previous week's observation. Here, positive weekly changes in fuel prices are associated with *negative* changes in flows in *both* the mainline and HOV lanes. In Model 9, we de-trend fuel prices flows by regressing each series on a constant and a linear time trend. In the results shown in column 4, we regress the residual flows obtained from this procedure on the corresponding de-trended fuel prices. Accounting for trends in this manner, the estimated flow elasticities are -0.045 and -0.077 in the ML and ML_HV lanes. In HOV lanes, the estimated flow elasticity is *negative* at -0.042, though not statistically significant.

Model 10 presents parameter estimates from a model intended to capture both the short-run and longer-run responses. We include both weekly fuel prices and monthly average fuel prices. One interpretation of this approach is that of a distributed lag model where the coefficients on lagged fuel prices are constrained to have equal effects on current traffic flows.[39] In this case, the longer-run effect can be thought of as the sum of the coefficients on current price and monthly average price.

For mainline lanes on highways without HOV lanes, the short-run response captured by current week fuel prices is -0.040. The elasticity with respect to monthly average prices is small, negative and not statistically significant. This suggests that drivers on these highways respond to fuel price increases primarily in the short-run by choosing the alternative option. For HOV lanes, the short-run response is -0.094 and statistically significant. However, the elasticity with respect to monthly

[39] Specifically, using a simple average of 4 weeks' price, the interim multipliers for lags one through four are just 0.25 times the coefficient on the monthly average price. The high level of autocorrelation in fuel prices makes estimation of the lagged parameters difficult without imposing structure on the lagged response. We experimented with several weighting schemes, each of which produced qualitatively similar results.

average prices is 0.251 and statistically significant. These results are consistent with the previous specifications which attempt to estimate the short-run and longer-run responses independently. Furthermore, the estimates support the hypothesis that immediately after a fuel price increase drivers substitute away from driving. Similarly, response to average prices is consistent with carpool formation responding to earlier fuel price increases. The combined effect of the two price elasticities suggests on overall long-run response in HOV lanes of approximately 0.16. Turning to the mainline lanes on highways with HOV lanes, the estimated short-run elasticity is -0.048. The coefficient on monthly average prices is -0.028 and is statistically significant at the one percent level. Notice that the response is different compared with mainline lanes on highways without HOV lanes where monthly average prices have no statistically significant relationship with traffic flows. Furthermore, taking the distributed lag interpretation of the Model 10 results suggests a long-run effect of -0.076 which is consistent with the base model.

One interpretation of the results in Table 5 is that the aggregate HOV response consists of two parts. In the short-run, fuel price increases cause some carpool drivers to choose the alternative option resulting in a decrease in flow. This is consistent with the findings in Models 8, 9 and 10. However in the longer-run, commuters are able to form carpools, resulting in an increase in HOV lane flow. This effect is large enough such that the net effect is positive consistent with our base model and Models 7 and 10.

For mainline lanes, the estimated elasticities are fairly similar across the specifications in Table 5 and are consistent with a short-run effect of drivers choosing the alternative option when fuel prices increase. Model 10 provides some evidence that time is also a factor in the flow response in mainline lanes on highways with HOV lanes. Our estimates suggest ML_HV flows also respond to longer-run trends captured by monthly average prices. This would be expected if these drivers are forming carpools and contributing to the observed increase in HOV lane flows. We find no evidence of similar behavior in mainline lanes on highways without HOV lanes.

5.4 Station-level heterogeneity

We investigate heterogeneity in the response to higher fuel prices by estimating station-level elasticities for each of the stations in our sample. We do this in two ways. First, assuming common year effects and weekly flow patterns, we estimate station-level elasticities by interacting a dummy for each station with fuel prices and report the coefficient on the interaction term. Second, we allow stations to have unique weekly flow patterns. To do this we first de-mean the data to allow for common year effects then estimate a separate OLS regression for each station. Figure 2 and Figure

3 plot the distributions of station flow elasticities by lane type. We see that there is substantial variability in the flow response at any individual station. However, the station-level results support the findings in the previous section. For both ML and ML_HOV lanes the masses of the probability distributions lie mainly to the left of zero with means 0.00 to -0.02 for mainline lanes and -0.03 to -0.04 for mainline lanes on highways with an HOV lane. For HOV lanes, the masses of the probability distributions lie to the right of zero with mean values between 0.09 and 0.15.

Section 2.3 suggests that factors such as the length of commute, income or access to public transit may impact drivers' responses to changing fuel prices. Given the substantial heterogeneity we observe in the station-level estimates, we are interested in determining whether observable factors explain some of the variation in flow elasticities or whether these difference are more likely due to unobservable characteristics of drivers and the highway network. As motivation for the empirical exercise that follows, Figure 4 plots the estimated station-level flow elasticity versus the route-distance from downtown Los Angeles for HOV lanes on I-210 and mainline lanes on I-101.[40] The results for HOV lanes on I-210 suggest a larger positive response to changing gas prices further from downtown. This is consistent with longer commute distances and greater average gasoline expenditure per trip resulting in more elastic carpool behavior. For mainline lanes on I-101 we see different behavior. In particular the flow response to changing fuel prices becomes less elastic further from downtown. One possible explanation is that because incomes rise substantially in the suburbs, the magnitude of the flow response falls with income.[41] To illustrate this, we also plot median household income for the zip code where each station is located. The income curve increases the greater the distance from downtown.

We investigate more formally the potential sources of heterogeneity in the response to fuel price changes by correlating our station level elasticity estimates with observable characteristics of neighborhoods around each station. While an imperfect measure of driver characteristics, we believe observables near each station are likely correlated with the characteristics of drivers who use this section of highway. We focus on elasticity estimates from the specification assuming common year and week effects. Section 2 suggests the price response should depend on the characteristics of commuters, namely, fuel expenditures g, the value of time ω, variable costs of the alternative option γ, and carpool formation costs τ. We construct proxies for each of these variables using geographic data and the 2000 Census.

[40]We define "routes" by the direction of travel along a given highway, for example I-101 northbound. We use "downtown" Los Angeles as defined below, to mark one endpoint of the commute route and thus subdivide each highway into two parts. The other endpoint is unspecified but allowed to vary continuously with distance from downtown.

[41]For example due to gasoline representing a smaller share of household consumption.

We match zip-code level data to each station based on its location in Los Angeles. We approximate fuel expenditure g using each station's distance from downtown Los Angeles. Specifically, distance from "downtown" is the route distance from each station to the station on the same freeway closest to Union Station. In doing so we assume individual fuel economy and commute destinations are orthogonal to the flow elasticity. We use median household personal income as a proxy for the mean value of time ω in a neighborhood. For the alternative option γ we use the fraction of worker trips made by public transit which should be correlated with the public transit infrastructure in a given neighborhood. Finally, we use population density measured in persons per square mile as a measure of the implied transaction costs of carpool formation.

Table 6 presents the results of models which regress the estimated flow elasticities on the above neighborhood characteristics. Model 11 includes only the distance variables interacted with indicators for lane type. The estimates for elasticities in HOV lanes and mainline lanes on highways without HOV lanes are positive and statistically significant. Moving one mile away from downtown is associated with a 0.003 increase in the price elasticity for mainline lanes on highways without an HOV lane. In other words because the mean elasticity is negative, flow is *less elastic* further from downtown. For HOV lanes, moving one mile further from downtown is associated with a 0.007 increase in the flow elasticity, i.e. HOV lane flow is *more elastic* further from downtown. Model 12 adds the natural log of median household income interacted with station type. None of the income parameters are statistically significant. The distance effect for mainline lanes is no longer statistically significant, though in HOV lanes the effect is similar to Model 11 and significant. Model 13 adds the full set of neighborhood characteristics. The effect of distance on the flow response in HOV lanes is unchanged. Income appears to have a negative relationship with flow elasticity in carpool lanes suggesting that commuters from wealthy areas are less likely to form carpools when fuel prices rise. In Model 13, population density appears to have a negative and statistically significant effect on HOV lane flow elasticity. The interpretation of this estimate is unclear, but suggests commuters from more densely populated areas may be more likely to choose the alternative option when fuel prices rise.

In looking at the estimated elasticities, we would be remiss not to mention that HOV stations on highway I-210 show a particularly strong relationship between the flow response and distance from downtown. Models 14 drops observations from I-210 from the sample. We see that when these observations are excluded, the relationships between HOV flow elasticities, route distance and income are no longer statistically significant. However, we note that the signs of the estimates are consistent with estimates including I-210 stations.

Our interpretation of the results in Table 6 is that the observable characteristics of Los Angeles freeways and drivers do not seem to have a large impact on the response of traffic flows to fuel price changes. We note that none of the models in Table 6 explain more than 9 percent of the variation in flow elasticities. Modeling individual driver behavior in this context is quite complicated. Commuters may have several highways from which to choose on a particular trip. In equilibrium, they may choose the least congested highway. However, because of bottlenecks, merges or other local conditions, fuel price changes are likely to affect congestion differently along different highways. Modeling driver highway choice and the interaction between fuel price changes and local traffic congestion is well beyond the scope of this paper. However, Table 6 is helpful in elucidating some trends in the responses of drivers, in particular that income and commute distance may be important factors affecting carpool formation. We view this as an important area for future research.

6 Conclusions

Understanding how commuters respond to changes in fuel prices, whether they form carpools, and how this behavior differs if HOV lanes are present has important implications for any policy that affects fuel prices. We show that on average, flows in high occupancy vehicle (HOV) lanes increase in response to higher fuel prices, while flows in mainline lanes decrease. Our preferred elasticity estimates are 0.136 for HOV lanes and -0.083 and -0.050 for highways with and without an HOV lane. While our theoretical model predicts the net effect of higher prices on carpooling could be positive or negative, our empirical results suggest higher prices increase the number of carpools. The fact that more of these commuters choose to carpool instead of the alternative option implies preferences for driving are strong enough to prevent many commuters from opting out of driving entirely when fuel prices rise.

We also find that mainline flows are more responsive to fuel prices when an HOV lane is present. Because the locations of HOV lanes in our sample were not randomly assigned we cannot necessarily conclude that this relationship is causal.[42] However, we do take this as suggestive evidence that fuel price increases could have a larger reduction in the social costs of driving for cities with better HOV lane infrastructure. To understand whether the higher elasticity leads to fewer total vehicles on the road, consider the effects for each highway configuration. For highways without an HOV lane, higher prices result in some commuters choosing the alternative option, which reduces the number of vehicles on the road. For highways with an HOV lane, higher prices result in some

[42]For example, one might worry that HOV lanes were preferentially located on major commuting highways where drivers have systematically different preferences for driving or public transportation.

commuters choosing the alternative option, while some choose to carpool. Intuitively, the decrease in mainline drivers may be offset by an increase in HOV lane use. Therefore, the net effect on the number of vehicles on the road depends on the initial number of drivers in each lane and the configuration of each highway.

We investigate the overall response to fuel price change in Los Angeles by performing the following back of the envelope calculation. Using the estimated flow elasticities from Table 3, Model 3, we calculate the total change in hourly flow for representative highways with and without an HOV lane. For simplicity we assume a 10% increase in fuel prices and use the sample mean flows and configurations for each highway type.[43] Under these assumptions, we find that flow on the mainline only highway decreases by approximately 27 vehicles per hour. Flow in the mainline portion of the mainline highway with an HOV lane decreases by approximately 45 vehicles per hour and flow in the HOV lane increases by approximately 10 vehicles per hour. Therefore, the net effect is a *greater* decrease in the number of vehicles per hour when the HOV lane is present. Furthermore, assuming each carpool consists of on average two commuters, we note that total number of drivers choosing the alternative option is comparable on both highways, 27 in the ML case and 45 - (2×10) = 25 for ML_HV drivers.

In the spirit of Verhoef, Nijkamp, and Rietveld (1996) we should not forget that the building of HOV lanes is likely a second-best approach to congestion relief absent optimal congestion tolls. Furthermore, the possibility that HOV lanes may result in a larger decrease in vehicle-related externalities due to short-run fuel price changes does not imply that building HOV lanes is an optimal policy. Determining whether or not additional HOV lanes are socially optimal would require a more comprehensive analysis, incorporating long-run considerations such as those in Duranton and Turner (2011). Recent research has examined converting existing HOV lanes into HOT (high-occupancy toll) lanes, with Small, Winston, and Yan (2006) finding that HOT lanes are welfare improving over HOV lanes. If HOT tolls are set equal to the marginal social cost of an additional vehicle, our estimates suggest that if gas prices increased by 10%, the optimal toll would rise by 1.3 cents/mile due to increased carpooling behavior under reasonable assumptions. For comparison, the peak morning toll on nearby SR-91 in Orange County is about 40 cents per mile, with the peak afternoon toll around 40-100 cents per mile, which is also comparable to the range described by Small, Winston, and Yan (2006).[44]

[43] Specifically, we use mean hourly flows per lane of 1,326, 1,345 and 706 for ML, ML_HV and HOV lanes, respectively. We assume mainline highways have 4 lanes and mainline HOV highways have 5 lanes including one HOV lane.

[44] Specifically, we assume that travel time changes with respect to flow are unit elastic, average speeds are 45 mph in the HOV lane and 35 mph in the ML, carpool occupancy is 2, and commuters have a value of time of $21.46

While at first glance these effects appear small, even a change of 10 vehicles per hour or an increase in marginal congestion costs of 1.3 cents/mile may lead to potentially significant welfare changes. We estimate the total changes in HOV and mainline congestion costs by aggregating the marginal congestion costs across space and time.[45] We find that the total cost of increased congestion in Los Angeles HOV lanes due to a 10% increase in fuel price is approximately $8.8 million per year. On the other hand, congestion costs for mainline drivers decease due to the approximately 20 vehicles per hour that leave to form carpools. We estimate mainline congestion costs fall by approximately $11.3 million per year.

Though we do not estimate changes in other vehicle related social costs directly, the fact the we estimate a net reduction in the number of vehicles on the road implies higher prices may also reduce other external costs. For example, emissions of criteria pollutants from motor vehicles are generally thought to be proportional to miles driven (Parry, Walls, and Harrington (2007)).[46] A larger total response to fuel price increases implies a greater reduction in emissions. Similar effects are likely to hold for accidents and other driving-related social costs.

In summary, we document three effects related to fuel prices, carpooling and the social costs of driving. First, higher fuel prices lead to fewer drivers in mainline lanes. Second, higher prices lead to an increase in carpooling on highways with HOV lanes. Third, the total reduction in the number of vehicles on the road is greater on highways with an HOV lane. Based on our model in Section 2 this last result implies the presence of an HOV lane provides a better substitute to the outside option when commuters are faced with higher prices. As a result, commuters who would have continued driving alone absent an HOV option instead choose to carpool. Finally, for both highways with and without HOV lanes we show that higher prices on average reduce the number of vehicles on the road. These effects have important implications for any policy that results in changes in fuel prices. Further, they are likely to occur in many other metropolitan areas partly because, while Los Angeles ranks amongst the highest congested cities in the US (Schrank, Lomax, and Eisele, 2011), HOV lanes in other cities have typically been built along the most congested freeways. Of course, other cities may have better public transit which provides additional options for drivers to adjust to higher fuel prices. In this case, one should interpret the flow effects in HOV lanes and the difference in mainline flow estimates for mainlines with and without HOV lanes as upper bound estimates.

(Small, Winston, and Yan, 2005)

[45] Using the mean flow rate of 706 vehicles per hour in the HOV lane and 1,345 in the ML_HV, we multiply the optimal tolls by 15 peak hours, times 260 commuting days per year and assume the same effect holds for all 240 lane-miles of HOV highway in LA.

[46] This is because automakers are held to emissions standards that stipulate grams per mile emissions rates.

References

Barth, M. and K. Boriboonsomsin. 2008. "Real-world carbon dioxide impacts of traffic congestion." *Transportation Research Record: Journal of the Transportation Research Board* 2058 (-1):163–171.

Bento, A.M., L.H. Goulder, M.R. Jacobsen, and R.H. von Haefen. 2009. "Distributional and efficiency impacts of increased US gasoline taxes." *The American Economic Review* 99 (3):667–699.

Blanchard, Christopher. 2009. "The Impact of Rising Gasoline Prices on U.S. Public Transit Ridership." *Duke University, mimeo* .

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

Busse, M.R., C.R. Knittel, and F. Zettelmeyer. 2010. "Pain at the pump: The effect of gasoline prices on new and used automobile markets." *NBER Working Paper* 15590.

Currie, Graham and Justin Phung. 2007. "Transit Ridership Auto Gas Prices, and World Events: New Drivers of Change?" *Transportation Research Record: Journal of the Transportation Research Board* 1992 (2):3–10.

Dahl, C. and T. Sterner. 1991. "Analysing gasoline demand elasticities: a survey." *Energy Economics* 13 (3):203–210.

Davis, L.W. and L. Kilian. 2011. "Estimating the effect of a gasoline tax on carbon emissions." *Journal of Applied Econometrics* 26 (7):1187–1214.

Duranton, Gilles and Matthew A. Turner. 2011. "The Fundamental Law of Road Congestion: Evidence from US Cities." *American Economic Review* 101 (6):261652.

Elliott, Graham, Thomas J. Rothenberg, and James H. Stock. 1996. "Efficient Tests for an Autoregressive Unit Root." *Econometrica* 64 (4):813–836.

Espey, M. 1998. "Gasoline demand revisited: an international meta-analysis of elasticities." *Energy Economics* 20 (3):273–295.

Ferguson, Erik. 1997. "The rise and fall of the American carpool: 1970-1990." *Transportation* 24:349–376.

Giuliano, G., D.W. Levine, and R.F. Teal. 1990. "Impact of high occupancy vehicle lanes on carpooling behavior." *Transportation* 17 (2):159–177.

Grabowski, David C. and Michael A. Morrisey. 2006. "Do Higher Gasoline Taxes Save Lives?" *Economic Letters* 90:51–55.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29 (1):113–134.

Im, Kyung So, M.Hashem Pesaran, and Yongcheol Shin. 2003. "Testing for unit roots in heterogeneous panels." *Journal of Econometrics* 115 (1):53 – 74.

Johnston, R.A. and R. Ceerla. 1996. "The effects of new high-occupancy vehicle lanes on travel and emissions." *Transportation Research Part A: Policy and Practice* 30 (1):35–50.

Klier, Thomas and Joshua Linn. 2010. "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data." *American Economic Journal: Economic Policy* 2 (3):134–153.

Knittel, Christopher R. and Ryan Sandler. 2011. "Cleaning the Baby with the Bathwater: The Health Co-Benefits of Carbon Pricing in Transportation." *NBER Working Paper No. 17390* .

Konishi, Hideo and Se il Mun. 2010. "Carpooling and congestion pricing: HOV and HOT lanes." *Regional Science and Urban Economics* 40 (4):173 – 186.

Kwon, J. and P. Varaiya. 2008. "Effectiveness of Californias high occupancy vehicle (HOV) system." *Transportation Research Part C: Emerging Technologies* 16 (1):98–115.

Lee, L.W. 1984. "The economics of carpools." *Economic Inquiry* 22 (1):128–135.

Levin, Andrew, Chien-Fu Lin, and Chia-Shang James Chu. 2002. "Unit root tests in panel data: asymptotic and finite-sample properties." *Journal of Econometrics* 108 (1):1 – 24.

Li, Shanjun, Christopher Timmins, and Roger H von Haefen. 2010. "How do Gasoline Prices Affect Fleet Fuel Economy." *American Economic Journal: Economic Policy* 1 (2):113–137.

Maddala, G. S. and Shaowen Wu. 1999. "A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test." *Oxford Bulletin of Economics and Statistics* 61 (S1):631–652.

Parry, I.W.H. and A. Bento. 2001. "Revenue recycling and the welfare effects of road pricing." *The Scandinavian Journal of Economics* 103 (4):645–671.

Parry, I.W.H. and K.A. Small. 2005. "Does Britain or the United States have the right gasoline tax?" *The American Economic Review* 95 (4):1276–1289.

Parry, I.W.H., M. Walls, and W. Harrington. 2007. "Automobile Externalities and Policies." *Journal of Economic Literature* 45 (2):373–399.

Puller, Steven L. and Lorna A. Greening. 1999. "Household Adjustment to Gasoline Price Change: An Analysis Using 9 Years of US Survey Data." *Energy Economics* 21 (1):37–52.

Rodier, C.J. and R.A. Johnston. 1997. "Travel, emissions, and welfare effects of travel demand management measures." *Transportation Research Record: Journal of the Transportation Research Board* 1598 (-1):18–24.

Schrank, David, Tim Lomax, and Bill Eisele. 2011. "Urban Mobility Report." Tech. rep., Texas Transportation Institute.

Shewmake, Sharon. 2010. "Can carpooling clean the air? The economics of HOV lanes, hybrid cars and the Clean Air Act." MPRA Paper 26129, University Library of Munich, Germany.

Small, Kenneth A., Clifford Winston, and Jia Yan. 2005. "Uncovering the Distribution of Motorists' Preferences for Travel Time and Reliability." *Econometrica* 73 (4):1367–1382.

———. 2006. "Differentiated Road Pricing, Express Lanes, and Carpools: Exploiting Heterogeneous Preferences in Policy Design." *Brookings-Wharton Papers on Urban Affairs* :pp. 53–96.

Verhoef, Erik, Peter Nijkamp, and Piet Rietveld. 1996. "Second-Best Congestion Pricing: The Case of an Untolled Alternative." *Journal of Urban Economics* 40 (3):279 – 302.

7 Figures

Figure 1: Average hourly flow during weekday peak hours and fuel prices for mainline lanes, mainline lanes on freeways with an HOV lane, and HOV lanes.

(a) Mainline lanes

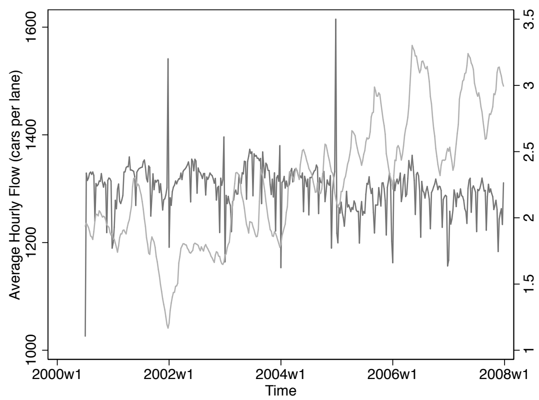

(b) Mainline lanes with HOV

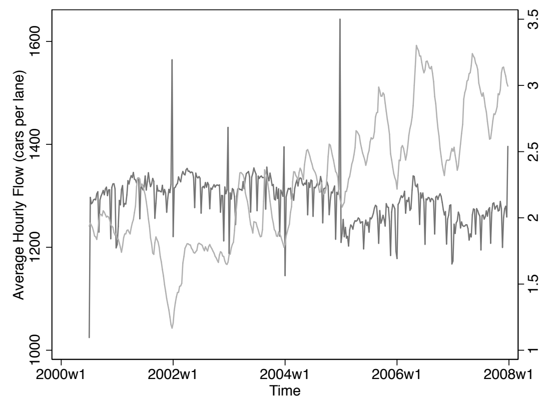

(c) HOV lanes

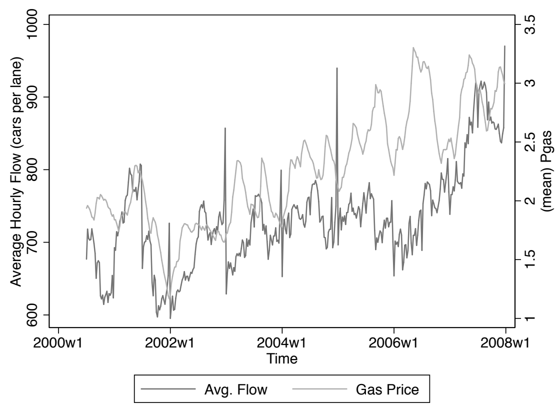

Figure 2: Distributions of station-level elasticities by lane type assuming common year and week mean-effects across stations. The distributions are truncated to include estimates between -1 and 1.

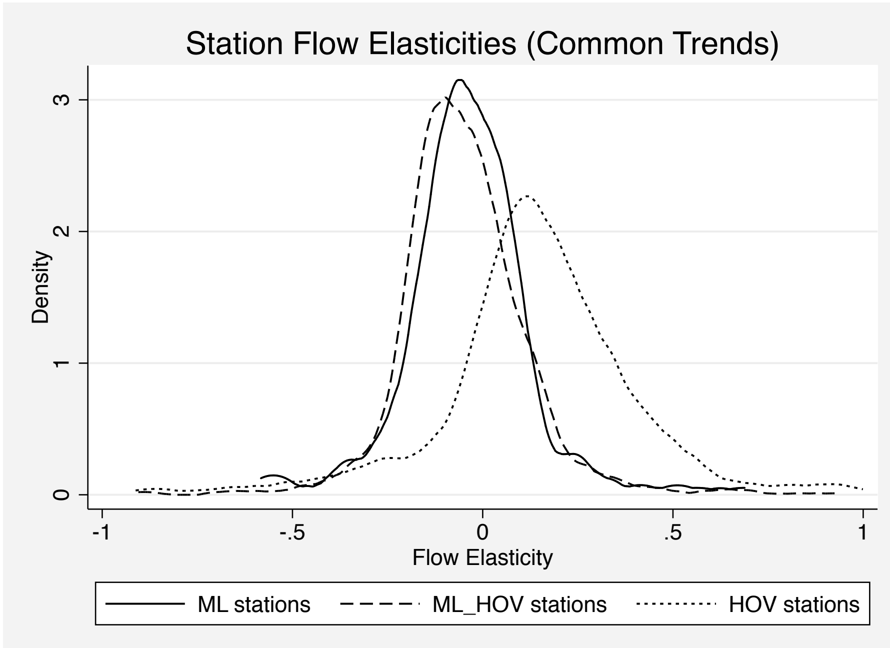

Figure 3: Distributions of station-level elasticities by lane type accounting for station-specific week mean-effects. The distributions are truncated to include estimates between -1 and 1.

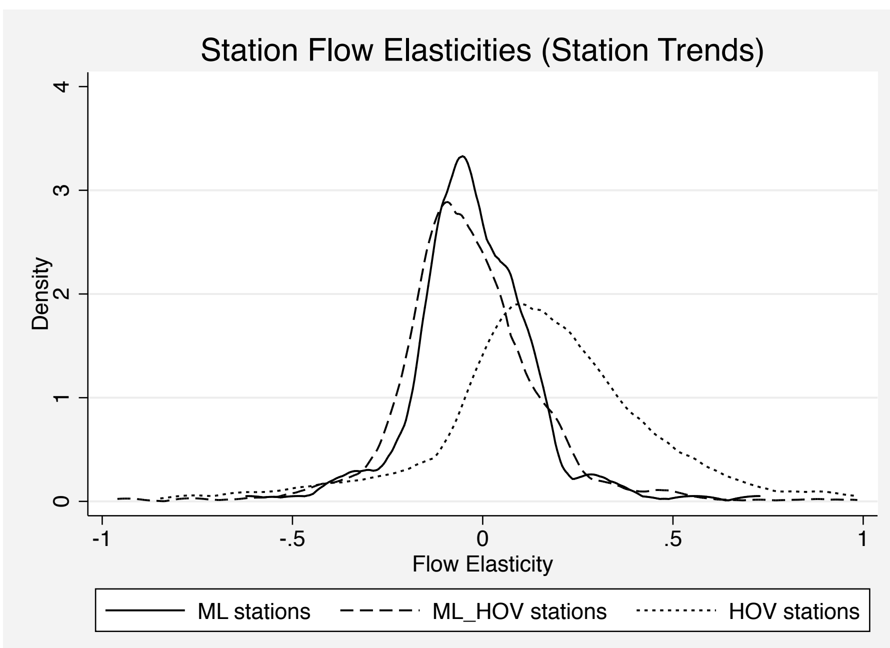

Figure 4: Station-level elasticity estimates versus distance from downtown Los Angeles. The smoothed lines are created using a weighted locally linear least squares procedure.

(a) I-210 HOV Lanes

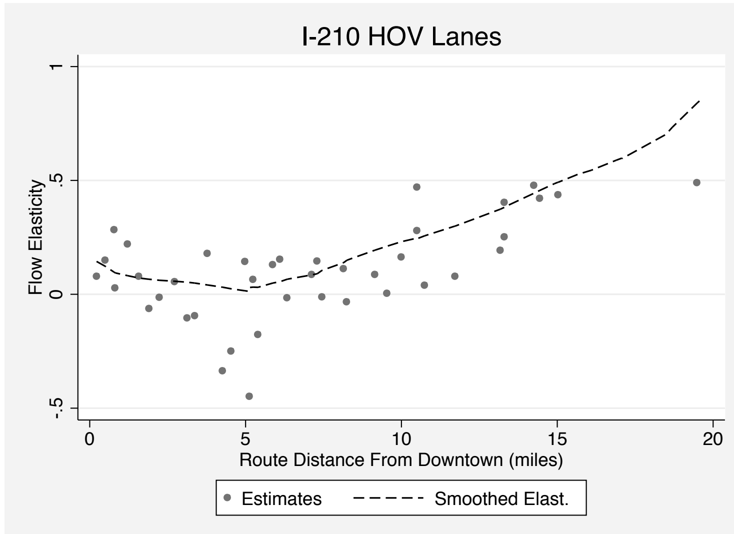

(b) I-101 Mainline Lanes

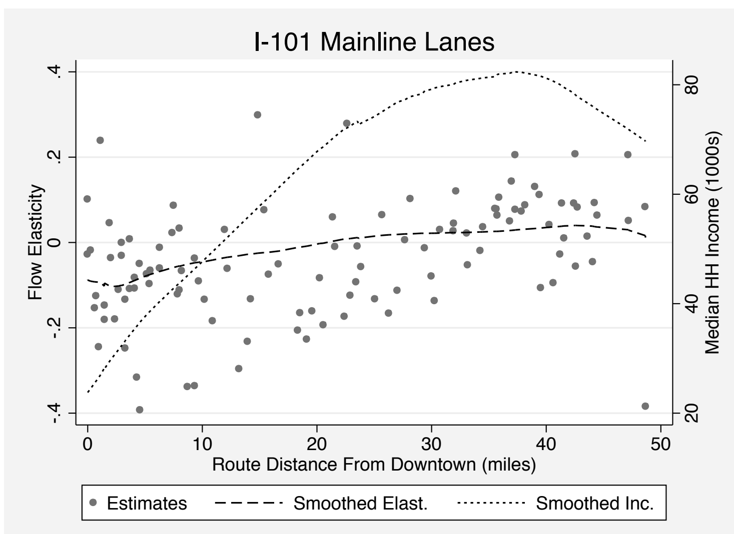

8 Tables

Table 1: Analytical results

		Price Effect	Sign
Scenario 1	ML flow	$-\frac{gL}{w\alpha+\gamma L}$	$-$
	Total flow	$-\frac{gL}{w\alpha+\gamma L}$	$-$
	Alternative users	$\frac{gL}{w\alpha+\gamma L}$	$+$
Scenario 2	HOV flow	$-\frac{g(w\alpha+(1-L)\gamma)}{2w\alpha(w\alpha+(1+L)\gamma)}$	$+/-$
	ML flow	$-\frac{g(L-1)(w\alpha+\gamma)}{w\alpha(w\alpha+(1+L)\gamma)}$	$-$
	Total flow	$-\frac{g((2L-1)w\alpha+(L-1)\gamma)}{2w\alpha(w\alpha+(1+L)\gamma)}$	$-$
	Alternative users	$\frac{gL}{w\alpha+(1+L)\gamma}$	$+$

Note: Scenario 1 represents a configuration of L mainline lanes and an alternative option. Scenario 2 consists of a configuration of $L-1$ mainline lanes, one HOV lane, and an alternative option.

Table 2: Summary statistics for weekday peak, weekday off-peak and weekend periods.

	Num. Obs.	Mean	Std. Dev.	Min.	Max.
Weekday Peak					
flow (veh.)	673,530	309,026	175,640	0	1,000,000
ln(flow)	673,530	12.33	0.94	8.74	13.86
mean hourly flow (veh./lane-hr.)	673,530	1,150	368	0	4,112
Weekday Off-Peak					
flow (veh.)	673,530	69,163	45,801	0	582,655
ln(flow)	673,530	10.61	1.34	3.00	13.28
mean hourly flow (veh./lane-hr.)	673,530	403	218	0	4,316
Weekend Peak					
flow (veh.)	673,530	106,749	61,898	0	515,274
ln(flow)	673,530	11.27	0.94	6.78	13.15
mean hourly flow (veh./lane-hr.)	673,530	999	339	0	4,100
Weekend Off-Peak					
flow (veh.)	673,530	28,210	19,128	0	276,407
ln(flow)	673,530	9.75	1.28	1.10	12.53
mean hourly flow (veh./lane-hr.)	673,530	415	224	0	5,119
All Periods					
price gasoline ($ 2005)	673,530	2.26	0.48	1.17	3.30
ln(price gasoline)	673,530	0.79	0.21	0.15	1.19
number lanes	673,530	3.29	1.46	1.00	7.00

Table 3: Flow elasticities for mainline lanes (ML), mainline lanes on highways with an HOV lane (ML HV), and HOV lanes.

Flow Elasticities by Lane Type

	Model 1	**Model 2**	**Model 3**	**Model 4**	**Model 5**	**Model 6**
ln(Gasoline Price) X D_{ML}	-0.004	-0.02	-0.050***	-0.098***	-0.049**	-0.027
	(0.0180)	(0.0180)	(0.0190)	(0.0210)	(0.0230)	(0.0250)
ln(Gasoline Price) X D_{ML_HV}	-0.036***	-0.051***	-0.083***	-0.129***	-0.075***	-0.062***
	(0.0070)	(0.0080)	(0.0100)	(0.0120)	(0.0120)	(0.0140)
ln(Gasoline Price) X D_{HV}	0.182***	0.167***	0.136***	0.089***	0.274***	0.191***
	(0.0250)	(0.0260)	(0.0200)	(0.0230)	(0.0240)	(0.0250)
Station Effects	Yes	Yes	Yes	Yes	Yes	Yes
Week Effects	No	Yes	Yes	Yes	Yes	Yes
Year Effects	No	No	Yes	No	Yes	Yes
Time Trend	No	No	No	Yes	No	No
Use all Imputed Obs.	No	No	No	No	Yes	No
Observations	339225	339225	339225	339225	673518	339225
Adj. R-Squared	0.99	0.99	0.99	0.99	0.96	

Notes: Dependent variables are the natural logarithms of total weekly flow. Standard errors clustered at the highway-direction level in five mile intervals. 2SLS estimates use world oil price IV. Station week observations with more than 25 percent imputation are dropped in each Model except model 5. ***, ** and * denote significance at the 1 percent, 5 percent and 10 10 percent levels, respectively.

Table 4: Flow elasticities for mainline lanes (ML), mainline lanes on highways with an HOV lane (ML HV), and HOV lanes during peak and off-peak periods.

Flow Elasticities by Lane Type and Period

	Weekday Peak (6am - 8pm)	**Weekday Off-Peak Period (8pm - 6am)**	**Weekend Peak Period (6am - 8pm)**	**Weekend Off-Peak Period (8pm - 6am)**	**Weekday Peak (4th Quart.)**	**Weekday Off-Peak Hours (1st. Quart)**
ln(Gasoline Price) X D_{ML}	-0.050***	-0.092***	-0.058**	-0.138***	-0.033*	-0.087**
	(0.0190)	(0.0310)	(0.0230)	(0.0320)	(0.0190)	(0.0430)
ln(Gasoline Price) X D_{ML_HV}	-0.083***	-0.108***	-0.077***	-0.144***	-0.068***	-0.093***
	(0.0100)	(0.0170)	(0.0120)	(0.0180)	(0.0100)	(0.0250)
ln(Gasoline Price) X D_{HV}	0.136***	0.267***	0.168***	0.227***	0.113***	0.372***
	(0.0200)	(0.0340)	(0.0240)	(0.0350)	(0.0200)	(0.0500)
Station Effects	Yes	Yes	Yes	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	339225	323075	336218	319731	340206	329493
Adj. R-Squared	0.99	0.98	0.98	0.98	0.99	0.98

Notes: Dependent variables are the natural logarithms of total weekly flow during each time period. Columns one through four use our preferred definition for peak and off-peak periods. Columns five and six define peak and off-peak as the 4th and 1st quartiles of hourly traffic flow for each station. Standard errors are clustered at the highway-direction level in five mile intervals. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels, respectively.

Table 5: Flow elasticities for mainline lanes (ML), mainline lanes on highways with an HOV lane (ML HV), and HOV lanes using average prices, and de-trended prices and flows.

	Dynamic Models				
	Base Model	**Model 7**	**Model 8**	**Model 9**	**Model 10**
ln(Gasoline Price) X D_{ML}	-0.050***				-0.040**
	(0.0190)				(0.0180)
ln(Gasoline Price) X D_{ML_HV}	-0.083***				-0.048***
	(0.0100)				(0.0070)
ln(Gasoline Price) X D_{HV}	0.136***				-0.094***
	(0.0200)				(0.0180)
4-week Avg. ln(Gasoline Price) X D_{ML}		-0.037*			-0.002
		(0.0190)			(0.0150)
4-week Avg. ln(Gasoline Price) X D_{ML_HV}		-0.071***			-0.028***
		(0.0100)			(0.0090)
4-week Avg. ln(Gasoline Price) X D_{HV}		0.161***			0.251***
		(0.0210)			(0.0230)
ln(Gas $Price_t$)-ln(Gas $Price_{t-1}$) X D_{ML}			-0.065***		
			(0.0110)		
ln(Gas $Price_t$)-ln(Gas $Price_{t-1}$) X D_{ML_HV}			-0.091***		
			(0.0050)		
ln(Gas $Price_t$)-ln(Gas $Price_{t-1}$) X D_{HV}			-0.168***		
			(0.0110)		
Gasoline Price X D_{ML}				-0.045***	
				(0.0120)	
Gasoline Price X D_{ML_HV}				-0.077***	
				(0.0070)	
Gasoline Price X D_{HV}				-0.042	
				(0.0260)	
Station Effects	Yes	Yes	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes	No	Yes
Observations	339225	338224	291434	339225	338224
Adj. R-Squared	0.99	0.99	0.23	0.33	0.99

Notes: Dependent variables are the natural logarithms of total weekly flow. Standard errors clustered at the highway-direction level in five mile intervals. Station-week observations with more than 25 percent imputation are dropped in each model. 4-week average prices are logged prices averaged over weeks t-1, t-2, t-3 and t-4. Flows in the first-difference model are ln($Flow_t$)-ln($flow_{t-1}$). Detrended series are calculated by regressing logged gasoline prices and logged station-level flows on a constant and a linear time trend. The results in column four are derived by regressing the residuals from the flow regressions on the residuals from the price regressions, interactions and fixed effects. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels.

Table 6: Correlations between station-level elasticity estimates and observable neighborhood characteristics.

Correlations Between Flow Elasticities and Observable Local Characteristics

	Model 11	**Model 12**	**Model 13**	**Model 14**
Route Distance X D_{ML}	0.003*	0.001	0.002	0.019
	(0.0020)	(0.0030)	(0.0030)	(0.0180)
Route Distance X D_{ML_HV}	0.001	0.002	0.002	0.000
	(0.0010)	(0.0020)	(0.0020)	(0.0020)
Route Distance X D_{HV}	0.007**	0.007**	0.007**	0.001
	(0.0030)	(0.0030)	(0.0030)	(0.0030)
ln(HH Income) X D_{ML}		0.100	0.163	0.662
		(0.0980)	(0.1850)	(0.5310)
ln(HH Income) X D_{ML_HV}		-0.045	-0.003	0.003
		(0.0370)	(0.0530)	(0.0530)
ln(HH Income) X D_{HV}		-0.033	-0.206**	-0.11
		(0.0550)	(0.0940)	(0.0970)
ln(Pop. Density) X D_{ML}			0.059	-0.05
			(0.0480)	(0.1070)
ln(Pop. Density) X D_{ML_HV}			0.004	-0.003
			(0.0190)	(0.0190)
ln(Pop. Density) X D_{HV}			-0.138***	-0.112***
			(0.0350)	(0.0360)
Transit Mode Share X D_{ML}			-0.27	3.272
			(0.7720)	(3.3980)
Transit Mode Share X D_{ML_HV}			0.194	0.172
			(0.2860)	(0.2780)
Transit Mode Share X D_{HV}			0.003	0.168
			(0.5590)	(0.5490)
Lane Type Effects	Yes	Yes	Yes	Yes
Route Effects	Yes	Yes	Yes	Yes
Include I-210 Stations	Yes	Yes	Yes	No
Observations	1651	1651	1651	1421
Adj. R-Squared	0.08	0.08	0.09	0.08

Notes: Dependent variables are station-level elasticity estimates. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels, respectively. Route distance is the distance on a given freeway from the detector closest to downtown Los Angeles. Population density is population per square mile. HH income is median household income and transit mode share is the fraction of work trips made by transit.

Appendices

A Robustness

In this section we further investigate the robustness of our results along several dimensions. First, we investigate whether our results are driven by the possibility of non-stationary series in the flow data. Second, we explore estimation in semi-log and levels specifications as alternatives to our preferred double log model. Third, we estimate average elasticities using a series of models which incorporate increasingly flexible time trends intended to capture smooth changes in average traffic flows that might confound the estimated relationships between flow and prices. Finally, we explore the effects of different HOV lane vehicle occupancy requirements and whether HOV lanes are continuous.

Appendix Table 1 shows results from our base model using the full sample of data. Column 1, shows estimates using our preferred model but limiting our sample to only stations for which an Elliot-Rothenburg-Stock test on each series rejects the presence of a unit root in the station-level flow data. The third column uses only observations where we fail to reject the presence of a unit root. The parameter estimates are qualitatively similar across the three samples. However, the magnitude of the estimated effect of fuel prices on HOV lane flow is smaller when only stationary series are used and relatively larger when the I(1) series are used. The effects on the parameter estimates are the opposite for mainline lanes where the magnitudes of the point estimates are larger when only the stationary series are used. Overall however, the signs are consistent with the full sample effects. The point estimates for ML_HV and HOV lanes are statistically significant in all samples.

Next we explore implications of our assumed functional form. While we believe changes in fuel prices are most likely to have the same proportional effect across stations with substantially different mean flows, it is important to explore whether our results rely on a particular functional form. Appendix Table 2 shows the mean elasticities estimated using our base model and two alternate specifications. In the semi-log, the dependent variable is the natural logarithm of flow. Fuel prices enter in levels and imply a semi-elasticity interpretation for the fuel price coefficients. In the levels specification, both traffic flows and fuel prices enter in levels. Qualitatively, the results across all three models look similar. Beginning with the semi-log model, the effects in the mainline are negative and statistically significant, consistent with our base model. The point estimate for ML_HV is larger in magnitude than the point estimate in the ML and statistically significant. For

HOV lanes, the relationship between fuel prices and flows is positive and statistically significant. In levels, the relationship between fuel prices and flows for ML_HV is negative and statistically significant. However, the estimate for ML is small, positive and not statistically significant. For HOV lanes, the relationship is positive and statistically significant.

At first glance, the mean flow responses in levels suggest a counterintuitive result that for a given change in fuel prices, HOV lane flow increases by approximately the same amount as mainline flow decreases. Of course if mainline drivers are forming carpools, we expect fewer than half of these vehicles to appear as additional vehicles in the HOV lane.[47] We note two characteristics of the PeMS data which may contribute to the results in levels. First, traffic detectors on mainline lanes tend to extend further from the city center compared with detectors in HOV lanes. Comparing sample means, flows for detectors in sections of the mainline where HOV flows are also monitored are approximately 5% larger than flows in outlying areas. This suggests our average flow elasticities may be biased downward if fuel prices result in proportional decreases in driving across the network. Second, PeMS engineers may locate traffic detectors more closely together on areas of freeway with high levels of flow and congestion. In addition, detectors may be spaced more closely together for HOV lanes compared with mainline lanes. To check this, we calculate the distance between detectors for each location in our sample. The correlation coefficient between spacing and flow is approximately -0.22 for HOV lanes and -0.16 for mainline lanes on highways with a HOV lane. It is approximately -0.002 for mainline lanes on other highways.

In light of these potential issues, we estimate the model in levels using a restricted sample that matches detectors in HOV lanes to detectors in mainline lanes in the same direction, on the same highway. This has the benefit of excluding observations from mainline lanes on areas of the highway where HOV lanes either do not exist or where traffic flows are not monitored. Furthermore, by matching stations one for one, we minimize complications due to the possibility that HOV flows are monitored more frequently in areas of high flow. In this restricted sample we estimate mean flow responses to fuel price changes of -7,952 in the mainline and 5,408 in HOV lanes. Furthermore, we cannot reject the hypothesis that the mainline response is twice as large and the HOV lane response.[48]

To make a more direct comparison between models and account for differences in mean flows across stations, Table 3 calculates the elasticities implied by the estimates in Table 2 evaluated

[47]All else equal. However, it may be possible that some drivers who form carpools switch from traveling on surface streets to freeways to take advantage of shorter travel times. This behavior may contribute somewhat to the size of the effects we estimate.

[48]With a test statistic of $F(1, 79) = 1.79$

at the conditional mean of the dependent variable. Comparing the log-log and semi-log models, the elasticities are strikingly similar. Compared with the preferred log-log specification, the mean elasticities implied by the levels specification are smaller in magnitude, though still qualitatively consistent with our main results.

Next, we investigate robustness to inclusion of higher-order time trends. We use as our base case Model 4 of Table 3 which includes a linear time trend. The models in Table 4 add polynomial trends to a maximum fifth-order. The average flow elasticity estimates for each of the three lane-types are quite stable across the five models. The relative sizes and ranking of elasticities by lane type remain consistent across models and the parameter estimates are in general, all statistically significant.

Finally, we explore whether are main results are robust to changes in HOV lane occupancy requirements and HOV lane continuity. Each of the highways in our sample limits HOV lane access to vehicles containing two or more commuters (2+) with one exception. HOV lanes on I-10 are limited to vehicles containing three or more commuters (3+). Intuitively, a higher occupancy requirement may raise the transaction cost of carpool formation and should result in lower HOV lane flows in equilibrium. How the higher 3+ restriction affects the responses of commuters to changes in fuel prices is less obvious. In terms of HOV lane continuity, the equilibrium level of use should depend on whether the highways are continuous, and therefore provide a larger potential savings in commute time. Whats less clear is how lane continuity would affect the marginal flow response to higher fuel prices since carpoolers could share costs regardless of whether they traveled in a HOV or ML lane.

We investigate lane restrictions by creating a dummy variable that is equal to one for HOV and mainline stations on I-10. We admit that our power to say something definitive about the effect of different lane restriction is limited by the fact that we observe only one highway with a 3+ restriction. Nevertheless, we interact the I-10 dummy with the lane type and fuel price variables to estimate separate mean flow elasticities for this highway. Regression results for this model are presented alongside our preferred specification in Appendix Table 5. We see that the estimated flow elasticities for the 2+ highways are quite similar to the pooled estimates in our preferred specification. Furthermore, none of the interaction terms are statistically significant, suggesting that behavior on the 3+ highway I-10, is not substantially different. Taken together, these results suggest our flow elasticities are not sensitive to differences in the required number of riders.

To investigate whether drivers respond differently when lanes are continuous or discontinuous, we create a dummy variable equal to one if the HOV lane on a given highway is discontinuous and interact this dummy with our lane type and price variables. Discontinuities could be due to either due to a small break in HOV lane coverage, or HOV lane coverage over only a portion of the highway. The discontinuous highways are the I-405, I-605, I-110 and I-10. The results of this regression are shown in column 3 of Appendix Table 5. We see that the flow elasticities for the continuous HOV highways are quite similar to our preferred model, though somewhat smaller in magnitude. The interaction terms suggest flows are somewhat more responsive to fuel price changes on highways where the HOV lane is discontinuous, though only the mainline effect is statistically significant. However, since we have few highways, attributing any observed difference to this particular effect is likely poorly identified given unobserved factors affecting commute choice that also vary at the highway level.

Appendix Tables

Appendix Table 1: Robustness of flow price relationships with and without non-stationary series.

	Base Model	Excl. I(1)	I(1) Only
ln(Gasoline Price) X D_{ML}	-0.050***	-0.066*	-0.032
	(0.0190)	(0.0340)	(0.0200)
ln(Gasoline Price) X D_{ML_HV}	-0.083***	-0.139***	-0.063***
	(0.0100)	(0.0230)	(0.0110)
ln(Gasoline Price) X D_{HV}	0.136***	0.053***	0.154***
	(0.0200)	(0.0190)	(0.0350)
Station Effects	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes
Observations	339225	78834	260391
Adj. R-Squared	0.99	0.99	0.99

Notes: Dependent variables are the natural logarithms of total weekly flow. Standard errors clustered at the highway-direction level in five mile intervals. Station week observations with more than 25 percent imputation are dropped in each model. Results in the second column exclude stations where Elliot-Rothenberg-Stock tests fail to reject the presense of a unit root in traffic flow. The third column uses only stations where we fail to reject the presence of a unit-root in the flow data. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels.

Appendix Table 2: Robustness of flow price relationships under log-log, semi-log and levels specifications.

Relationships Between Flow and Prices With Respect to Functional Form

	Log-Log	**Semi-Log**	**Levels**
ln(Gasoline Price) X D_{ML}	-0.050***		
	(0.0190)		
ln(Gasoline Price) X D_{ML_HV}	-0.083***		
	(0.0100)		
ln(Gasoline Price) X D_{HV}	0.136***		
	(0.0200)		
Gasoline Price X D_{ML}		-0.023***	89.52
		(0.0090)	(2956.1170)
Gasoline Price X D_{ML_HV}		-0.038***	-6263.190***
		(0.0050)	(801.0250)
Gasoline Price X D_{HV}		0.060***	6454.891***
		(0.0090)	(1220.7930)
Station Effects	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes
Observations	339225	339225	339225
Adj. R-Squared	0.99	0.99	0.98

Notes: Dependent variables in column 1 and column 2 are the natural logarithms of total weekly flow. The dependent variable in column 3 is flow in levels. Standard errors clustered at the highway-direction level in five mile intervals. Station-week observations with more than 25 percent imputation are dropped in each model. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels.

Appendix Table 3: Implied flow elasticities under log-log, semi-log and level specifications.

Flow Elasticities by Lane Type and Functional Form

	Log-Log	**Semi-Log**	**Levels**
ML Elasticity	-0.050***	-0.052***	0.001
	(0.0193)	(0.0194)	(0.0213)
ML HV Elasticity	-0.083***	-0.086***	-0.047***
	(0.0103)	(0.0104)	(0.0062)
HV Elasticity	0.136***	0.137***	0.045***
	(0.0200)	(0.0200)	(0.0082)

Notes: Flow elasticities in the semi-log and levels specifications are calculated at the conditional mean of the dependent variable. The dependent variables in the log-log and semi-log specifications are the natural logarithm of total weekly flow during weekday peak hours. Standard errors are reported in parentheses. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels, respectively

Appendix Table 4: Robustness of flow elasticity estimates for mainline lanes (ML), mainline lanes on highways with a HOV lane (ML HV), and HOV lanes for models incorporation time trends up to fifth-order.

Flow Elasticities by Lane Type Including Various Time Trends

	First	**Second**	**Third**	**Fourth**	**Fifth**
ln(Gasoline Price) X D_{ML}	-0.098***	-0.077***	-0.031	-0.031	-0.054***
	(0.0210)	(0.0190)	(0.0190)	(0.0190)	(0.0190)
ln(Gasoline Price) X D_{ML_HV}	-0.129***	-0.109***	-0.063***	-0.062***	-0.087***
	(0.0120)	(0.0110)	(0.0100)	(0.0100)	(0.0090)
ln(Gasoline Price) X D_{HV}	0.089***	0.110***	0.156***	0.157***	0.132***
	(0.0230)	(0.0230)	(0.0220)	(0.0220)	(0.0200)
Station Effects	Yes	Yes	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes	Yes	Yes
Order of Time Trend Polynomial	T	T^2	T^3	T^4	T^5
Observations	339225	339225	339225	339225	339225
Adj. R-Squared	0.99	0.99	0.99	0.99	0.99

Notes: Dependent variables are the natural logarithms of total weekly flow. Standard errors clustered at the highway-direction level in five mile intervals. Station week observations with more than 25 percent imputation are dropped in each model. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels, respectively.

Appendix Table 5: Effects of 2+ versus 3+ lane restrictions and discontinuous HOV lanes.

Additional Robustness Checks

	Base Model	**2+ vs. 3+ Lanes**	**Disc. HOV Lanes**
ln(Gasoline Price) X D_{ML}	-0.050***	-0.050***	-0.051***
	(0.0190)	(0.0190)	(0.0190)
ln(Gasoline Price) X D_{ML_HV}	-0.083***	-0.079***	-0.052***
	(0.0100)	(0.0110)	(0.0120)
ln(Gasoline Price) X D_{HV}	0.136***	0.131***	0.115***
	(0.0200)	(0.0220)	(0.0370)
ln(Gasoline Price) X D_{ML_HV} X D_{3+}		-0.024	
		(0.0160)	
ln(Gasoline Price) X D_{HV} X D_{3+}		0.066	
		(0.0420)	
ln(Gasoline Price) X D_{ML_HV} X $D_{disc.}$			-0.067***
			(0.0130)
ln(Gasoline Price) X D_{HV} X $D_{disc.}$			0.043
			(0.0480)
Station Effects	Yes	Yes	Yes
Week Effects	Yes	Yes	Yes
Year Effects	Yes	Yes	Yes
Observations	339225	339225	339225
Adj. R-Squared	0.99	0.99	0.99

Notes: Dependent variables are the natural logarithms of total weekly flow. Standard errors clustered at the highway-direction level in five mile intervals. Observations with more than 25 percent imputation are dropped. ***, ** and * denote significance at the 1 percent, 5 percent and 10 percent levels.