

NBER WORKING PAPER SERIES

ANTICIPATION, TAX AVOIDANCE, AND THE PRICE ELASTICITY OF GASOLINE
DEMAND

John Coglianese
Lucas W. Davis
Lutz Kilian
James H. Stock

Working Paper 20980
<http://www.nber.org/papers/w20980>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2015

The authors have not received any financial compensation for this project nor do they have any financial relationships that relate to this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by John Coglianese, Lucas W. Davis, Lutz Kilian, and James H. Stock. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand
John Coglianesse, Lucas W. Davis, Lutz Kilian, and James H. Stock
NBER Working Paper No. 20980
February 2015
JEL No. H23,H26,Q41,Q47

ABSTRACT

Traditional least squares estimates of the responsiveness of gasoline consumption to changes in gasoline prices are biased toward zero, given the endogeneity of gasoline prices. A seemingly natural solution to this problem is to instrument for gasoline prices using gasoline taxes, but this approach tends to yield implausibly large price elasticities. We demonstrate that anticipatory behavior provides an important explanation for this result. We provide evidence that gasoline buyers increase gasoline purchases before tax increases and delay gasoline purchases before tax decreases. This intertemporal substitution renders the tax instrument endogenous, invalidating conventional IV analysis. We show that including suitable leads and lags in the regression restores the validity of the IV estimator, resulting in much lower and more plausible elasticity estimates. Our analysis has implications more broadly for the IV analysis of markets in which buyers may store purchases for future consumption.

John Coglianesse
John F. Kennedy School of Government
Harvard University
coglianesse@fas.harvard.edu

Lucas W. Davis
Haas School of Business
University of California
Berkeley, CA 94720-1900
and NBER
ldavis@haas.berkeley.edu

Lutz Kilian
Dept.of Economics
University of Michigan
Ann Arbor, MI 48109-1220
and CEPR
lkilian@umich.edu

James H. Stock
Department of Economics
Harvard University
Littauer Center M26
Cambridge, MA 02138
and NBER
James_Stock@harvard.edu

1. Introduction

Many economists believe that U.S. gasoline consumption in the short run is largely unresponsive to fluctuations in the retail price of gasoline. Whether this view is correct, is not self-evident. Figure 1 plots U.S. gasoline consumption since January 1974. Phases of rising real U.S. gasoline prices are shown as shaded areas. Figure 1 shows clear evidence of consumption rising during periods of falling real gasoline prices and of consumption growth slowing or consumption falling during periods of rising real gasoline prices, at least until the late 1990s. More recently, this pattern has weakened somewhat (owing in part to the fact that many gasoline price increases after 2000 were associated with a thriving economy). Nevertheless, following the sustained increase in the price of gasoline, U.S. gasoline consumption declined significantly, and consumption rebounded in the second half of 2008 when gasoline prices dropped briefly, but sharply.

The central question raised by Figure 1 is how strongly U.S. gasoline consumption responds to exogenous shifts in gasoline prices. The magnitude of the short-run price elasticity of gasoline demand is of immediate policy interest. For example, the rapid decline in gasoline prices in recent months has renewed interest in the question of how U.S. gasoline consumption will respond to the decline in retail gasoline prices (see, e.g., U.S. EIA 2015). Knowledge of this elasticity also is important for gauging the macroeconomic effects of gasoline price fluctuations (see Edelstein and Kilian 2009). Moreover, the magnitude of the price elasticity of gasoline demand plays an important role in the debate about speculation in oil markets (see Hamilton 2009; Kilian and Murphy 2014). Finally, the price elasticity is an important parameter in microeconomic models of the automobile market that are used in industrial organization and in environmental economics (see, e.g., Allcott

and Wozny 2014).

Estimates of the responsiveness of gasoline consumption to changes in gasoline prices must take account of the endogeneity of the price of gasoline. It is well known that increases in the demand for gasoline cause the price of gasoline to increase, resulting in a spurious correlation between the price and the regression error, and biasing estimates of the price elasticity toward zero. As a result, traditional least squares estimates of the price elasticity of gasoline demand tend to be implausibly low. In response to this problem, several recent studies have employed instrumental variables, but it has proven difficult to find instruments that are both highly predictive of gasoline prices and uncorrelated with demand (see, e.g., Ramsey et al. 1975; Dahl 1979; Li et al. 2014; Sweeney 2015).

A seemingly natural approach is to instrument for gasoline prices using gasoline taxes. There have been more than 150 state-level gasoline tax changes since 1989 and gasoline taxes are highly predictive of gasoline prices. Moreover, tax changes are typically implemented with a considerable lag making it unlikely that tax changes are correlated with contemporaneous demand shocks. In practice, however, instrumental variable (IV) regressions using gasoline tax changes have yielded estimates that are unexpectedly large. For example, Davis and Kilian (2011) report a short-run price elasticity of gasoline demand of -1.14 under a conventional IV specification. Their preferred IV specification based on nominal tax increases yields a price elasticity of -0.46. Even the latter estimate is at the upper end of the range of elasticity values that seem economically plausible and indeed is higher than many economists would be comfortable with.

There are several potential explanations for these puzzling estimates. First, as emphasized by Davis and Kilian (2011), price changes induced by tax changes are more

persistent than other price changes and thus may induce larger behavioral responses. Second, gasoline tax changes may be more salient than typical price changes because they are widely discussed in the media (e.g., Davis and Kilian, 2011; Li et al. 2014). This argument is in line with Chetty et al.'s (2009) evidence that increasing consumers' awareness of sales taxes affects consumer demand. Third, as noted by Tiezzi and Verde (2014), consumers may suffer from tax aversion (see McCaffery and Baron 2006). These three explanations are not mutually exclusive. In response to the evidence in Davis and Kilian (2011), several studies have compared the price elasticity of gasoline demand with the tax elasticity of gasoline demand (see Li et al. 2014; Tiezzi and Verde 2014; Rivers and Schaufele 2014).

The current paper proposes an alternative and independent explanation. We provide evidence that these large elasticity estimates are an artifact of not having accounted for shifts in gasoline purchases in anticipation of gasoline tax changes. Tax changes are easily predicted, so forward-looking gasoline buyers will take future tax changes into account when deciding how much gasoline to buy. This anticipatory behavior applies not only to final consumers, but also to wholesale distributors and operators of retail gasoline stations, all of whom have some ability to store gasoline. We document large and statistically significant increases in gasoline purchases during the month leading up to gasoline tax increases. The pattern appears to be approximately symmetric in tax increases and decreases with gas purchases decreasing during the month leading up to gasoline tax decreases, though there are far fewer tax decreases in our data and this effect is not statistically significant.

This intertemporal substitution by buyers creates an endogeneity problem that

invalidates conventional IV analysis. We show that this econometric problem can be overcome by including one lead and one lag of the change in gasoline price in the regression and including one lead and one lag of the tax instrument in the first stage IV regression. This approach results in a much lower and more economically plausible point estimate of the elasticity of -0.37. Although not statistically significant, this estimate is well within the range of recent estimates of the price elasticity of gasoline demand based on alternative models and datasets that account for or minimize price endogeneity problems (see, e.g., Bento et al 2009; Tiezzi and Verde 2013; Levin et al. 2014).

The remainder of the paper is organized as follows. In Section 2, we review the problem of estimating the price elasticity of gasoline demand. Section 3 provides empirical evidence of forward-looking behavior in the gasoline market and discusses the economic mechanisms underlying this behavior. In Section 4, we analyze the econometric implications of this evidence, propose an alternative approach to implementing the IV estimation that remains valid in the presence of forward-looking behavior, and present new estimates of the short-run price elasticity of gasoline demand. Section 5 compares our preferred estimate with other elasticity estimates in the recent literature. In Section 6, we discuss implications of our analysis for future U.S. gasoline consumption. The concluding remarks are in section 7.

2. Background

There is a large literature on estimating the price elasticity of gasoline demand in the short run and in the long run. For example, Hausman and Newey (1995) estimate the U.S. long-run price elasticity of gasoline demand to be -0.8 based on pooled household data. Much of the literature has been concerned with estimating the short-run price elasticity of U.S.

gasoline demand based on monthly or quarterly data. The reviews by Dahl and Sterner (1991), Espey (1998), Greene et al. (1999), Graham and Glaister (2004) and Brons et al. (2008) identify dozens of econometric studies of this elasticity. More recent contributions include Hughes et al. (2008), Li et al. (2014), Levin, Lewis and Wolak (2014), and Tiezzi and Verde (2014).

One would expect the short-run price elasticity of gasoline demand to be much smaller than the long-run elasticity, because it takes time for consumers to fully adjust their consumption in response to higher gasoline prices (see Sweeney 1984). Typical estimates of the short-run elasticity obtained from ordinary least squares (OLS) regressions of changes in gasoline purchases on changes in gasoline prices tend to be very close to zero, however. One reason is that increases in the demand for gasoline cause the price of gasoline to increase, resulting in a spurious correlation between the price and the regression error and biasing estimates of the price elasticity toward zero. The main approach to addressing this endogeneity problem has been to instrument for gasoline prices. While this approach is appealing, the challenge has been to find instruments that are both truly exogenous and strong in the econometric sense (see Stock et al. 2002).¹

Davis and Kilian (2011) proposed the use of changes in gasoline taxes by state and month as an instrument. Even though tax legislation may respond to current prices, the implementation of tax changes typically occurs with a lag, making it reasonable to assume that changes in tax rates are uncorrelated with unobserved changes in demand. In

¹ For example, Ramsey *et al.* (1975) and Dahl (1979) use the relative prices of refinery products such as kerosene and residual fuel oil as instruments. As noted in Hughes *et al.* (2008) the problem with this approach is that the relative prices of other refinery outputs are likely to be correlated with gasoline demand shocks. Instead, Hughes *et al.* instrument using changes in global crude oil production such as a strike by oil workers in Venezuela in 2002. Li *et al.* (2014) rely on global crude oil prices as the instrument. The latter instrument is unlikely to be exogenous, however.

constructing tax instruments one has to be careful to exclude *ad valorem* gasoline taxes from the analysis because they are functionally related to price, violating the exogeneity assumption. Using U.S. data for January 1989 through March 2008, Davis and Kilian constructed a balanced state-level panel model of 50 states and the District of Columbia and estimated a regression equation of the form

$$\Delta q_{it} = \beta_0 + \beta_1 \Delta p_{it} + \rho_t + u_{it}, \quad (1)$$

where changes in log gasoline purchases in state i and month t , Δq_{it} , depend linearly on changes in the after-tax price of gasoline in logs, Δp_{it} , time fixed effects, ρ_t , and unobserved idiosyncratic state-specific time-varying factors comprising an error term, u_{it} .

The time fixed effects control for both seasonal variation and variation from year to year that is the same across states. The differencing eliminates time-invariant state-specific factors, thereby significantly mitigating concerns about omitted variables. Gasoline consumption is serially correlated, driven by the number of vehicles in circulation, the fuel efficiency of these vehicles, and the number of miles traveled. These factors evolve differentially over time across states, so the time fixed effects cannot eliminate all relevant omitted variables. The advantage of the first-differenced model is that it is identified using month-to-month changes, not longer-run before-and-after comparisons like the fixed effects model. This makes first-differenced models inherently less susceptible to unmodeled differential trends across states.

The coefficient of interest is the price elasticity parameter β_1 . For their preferred IV specification, which restricts attention to nominal tax increases, Davis and Kilian (2011) report a statistically significant estimate of the price elasticity of gasoline demand of -0.46.

This IV point estimate is much larger than the corresponding OLS estimate of -0.10 for the same data set. It is also higher than most economists would have expected. It is arguably near the upper bound of elasticity values that seem economically plausible. As noted in the introduction, other specifications yield even higher elasticity estimates.

A potential concern with this IV estimate is the possible endogeneity of changes in state gasoline taxes with respect to omitted variables. For the IV estimate to be valid, the instrument must be uncorrelated with the error term. This condition would be violated, for example, if state legislators changed gasoline taxes in response to variables that are correlated with state-level demand for gasoline, but are excluded from the regression. Davis and Kilian (2011) provide additional evidence that including several such regressors does not affect the elasticity estimate and conclude that there is no indication of the timing of the implementation of the tax instrument being endogenous.

The next section will show that there is another very different reason to be concerned with the endogeneity of the instrument that has not been considered in the existing literature. The concern is that gasoline purchases may adjust in anticipation of changes in gasoline taxes even before the tax change takes effect. Simply put, if the gas tax is scheduled to increase in January, some gasoline buyers will fill up at the end of December, shifting forward purchases that would normally occur in January; thus January sales drop both because of the higher after-tax price (the elasticity) and because sales were shifted forward to December.

Anticipatory behavior of this form would, in fact, violate the IV identifying assumption. For the tax instrument to be valid, it must be uncorrelated with the error term in equation (1), i.e.,

$$E\{u_{it} \Delta\tau_{it}\} = 0. \quad (2)$$

where the instrument, $\Delta\tau_{it}$, is the change in taxes in state i between month $t-1$ and month t . In other words, the tax instrument is allowed to affect gasoline purchases only through its effect on the price of gasoline in the first-stage regression

$$\Delta p_{it} = \alpha_0 + \alpha_1 \Delta\tau_{it} + \theta_t + v_{it}.$$

The orthogonality condition (2) is violated, when the quantity variable responds to taxes in future periods. In this case, a tax change in period t influences the change in gasoline purchases in period t not only through the change in price in period t , but also by having already affected gasoline purchases in period $t-1$. As a result, the contemporaneous change in gasoline purchases, Δq_{it} , is larger than would have been expected for a typical price change of that magnitude, and the IV estimate of the price elasticity systematically overstates the responsiveness of gasoline purchases. The next section provides evidence for the presence of such anticipatory behavior in the U.S. gasoline market.

3. The Evidence of Anticipatory Demand Shifts

3.1. An Event Study

An event study based on the data used in Davis and Kilian (2011) allows us to quantify the importance of anticipatory shifts in gasoline purchases. Figure 2a summarizes the typical patterns in the state-level data around the time of a gasoline tax increase. For expository purposes, the figure limits attention to nominal tax increases of at least 1 cent. The charts average over all 140 events of nominal state tax increases of at least 1 cent and plot gasoline taxes, gasoline prices, and gasoline purchases in the months immediately before, during, and after such a tax increase. The gasoline tax increase is normalized to occur in

month 1. Figure 2a shows that gasoline prices respond to tax increases, confirming that the instrument is strong, but it also shows an increase in gasoline purchases in the month before the tax hike. Much the same pattern is obtained in Figure 2b, after restricting attention to gasoline tax increases of at least 2 cents. In fact, the magnitude of the shifts in gasoline purchases is even more pronounced, as one might have expected. As in Figure 2a, gasoline purchases spike in the month preceding the tax increase, and then drop markedly in the month of the tax increase.

This pattern is consistent with gasoline buyers filling up their tanks one last time in the days before the higher tax goes into effect. This type of anticipatory behavior also helps explain the sharp decrease in purchases observed during the month of the tax increase, as gasoline buyers enter the month with full tanks and hence are less likely to stop at a gas station. This economic interpretation suggests that conventional IV estimates overstate the extent to which consumption falls in response to the exogenous gasoline price increase. These estimates are biased because they take as their point of departure the excessive level of purchases immediately before the tax increase rather than the normal consumption level, when computing the response of quantity to the exogenous shift in price.

So far we have focused on gas tax increases. Figures 2c and 2d provide the corresponding graphical evidence for tax decreases of at least one cent and at least two cents, respectively. Consistent with our proposed explanation, the change in gasoline purchases is negative in both figures in the month prior to the gasoline tax decrease becoming effective. This is what one would expect if gasoline buyers were strategically delaying their gasoline purchases in anticipation of the lower taxes in the following month.

However, there is much more month-to-month variation in the event study estimates overall.

3.2. Regression Evidence based on the Event Study

Table 1 provides regression estimates and standard errors describing the evolution of log gasoline purchases before, during, and month after gasoline tax changes. We regress the month-to-month change in log gasoline purchases on event study indicator variables for the month before, the month during, and the month after gasoline tax changes in order to assess the statistical significance of the changes in gasoline purchases observed in Figures 2a-2d. The point estimates in column (a) show that gasoline purchases increase by a marginally significant 1.3% in the month preceding the tax increase, followed by a statistically significant drop by 3.1% in the month of the tax increase. Focusing on tax increases of at least two cents in column (b) sharpens the pattern, with a statistically significant increase of 2.1% in the month before, and a statistically significant decrease of 3.8% in the month in which the tax increase occurs. The corresponding estimate for the month following the tax increase is not statistically significant in either specification, but, especially in column (c), is large enough to potentially affect the overall impact of the gas tax increase.

Columns (c) and (d) in Table 1 provide the corresponding evidence for gasoline tax decreases. With far fewer gasoline tax decreases, these coefficients are estimated with considerably lower precision. There are a number of reasons why one might expect less anticipatory behavior in response to tax decreases. For example, one might expect gasoline tax decreases to lack the salience of gasoline tax increases if the latter receive more media attention. Perhaps more importantly, as a practical matter, it is much easier to

fill up one's tank three days early than to wait three extra days to fill the tank. Nevertheless, the overall pattern of the point estimates is roughly symmetric in gas tax increases and decreases. In particular in column (d), which focuses on gas tax decreases of at least two cents, there appear to be non-negligible decreases in gasoline purchases during the month before as well as during the month after the tax decrease.

3.3. Discussion of Observed Magnitudes

There are several reasons why the evidence of anticipatory changes in gasoline purchases in Table 1 is economically plausible. One reason is that drivers may store gasoline in the tank of their automobile, as discussed earlier. For example, suppose that 10% of drivers were to shift forward their weekly fillup by a few days. This would amount to a shift in purchases of about 2%, which is roughly consistent with the observed effect in the quantity chart.

This is not the only mechanism at work, however. The anticipatory behavior in Table 1 likely reflects reactions not only by retail consumers, but also by gasoline station operators and gasoline distributors. This point is not immediately obvious, but related to the fact that gasoline purchases are measured upstream from actual sales to final consumers. Specifically, the conventionally used data on gasoline purchases comes from the *Petroleum Marketing Monthly Report: Prime Supplier Sales Volumes by Product and Area* issued by the U.S. Energy Information Administration (EIA). Sales volumes are collected by the EIA using the EIA-782C survey, which is a monthly survey of all prime gasoline suppliers, consisting a small group of currently 185 firms nationwide that produce, import, or transport petroleum products across state boundaries and local marketing areas. This gasoline is then purchased by distributors for delivery to gasoline stations. As a result, data

on the quantity of gasoline purchased collected by statistical agencies differs from actual gasoline purchases at the retail level (also see Levin et al. 2014).

To understand how distributors and retailers may affect the data on gasoline purchases one has to examine how gasoline taxes are collected. As Wojciech et al. (2013) discuss in a recent study, the point of taxation for fuels in the United States has been moving steadily upstream. For example, the federal gasoline tax was traditionally collected at the retail level, but in 1988 this task moved to the distributors and, in 1994, to the prime supplier. The remittance of state fuel taxes has been moving in the same direction, but with some variation across states. As a result, not only retail consumers, but also gasoline distributors and gasoline station operators have an incentive to adjust their purchase and storage choices in the days and weeks leading up to gasoline tax changes. In particular, both have an incentive to stockpile gasoline in anticipation of gasoline tax increases.

Thus, the apparent spike in the quantity of gasoline sold prior to gasoline tax increases likely reflects not only shifts in the purchases of final consumers, but also stockpiling by gasoline distributors and gasoline stations. For example, in the days leading up to a gasoline tax increase, forward-looking distributors would be expected to fill up all their tanker trucks, and gasoline station operators would be expected to have their underground storage tanks completely filled. The farther upstream the gasoline tax is collected, the greater the scope for stockpiling.

4. Econometric Implications

Our empirical findings have two important implications for the estimation of the price elasticity of gasoline demand. First, anticipatory behavior by gasoline buyers undermines

the rationale for using gasoline taxes as an instrument. Intertemporal substitution causes the gasoline tax instrument to be correlated with the error term because the presence of the tax spike is correlated with the “missing” purchases that were shifted forward in the previous period, rendering the instrument endogenous. Second, because this endogeneity is the result of lag misspecification, it can be corrected by including leads and lags of the change in prices in the regression, along with leads and lags of the tax change as instrumental variables. The lead of the change in the gas tax, in particular, is a valid instrument for the lead of the change in the gas prices because the future tax change is known to consumers, given that it already has been legislated.

We examine the implications of these two insights in Table 2. Table 2 reports coefficient estimates and standard errors for a range of alternative IV model specifications. For comparison, Table 2 also presents these specifications estimated by OLS. The dependent variable in all five specifications is the month-to-month change in log gasoline purchases in the state. All models control for month-of-sample fixed effects (time effects) and the change in state-level unemployment. Standard errors are clustered at the state level. Columns (1)-(3) include the contemporaneous change in log gasoline price only, without a lead or lag. Column (1) presents OLS estimates and column (2) instruments for the change in the log gasoline price using the contemporaneous change in the log tax, whereas column (3) adds, as additional instruments, the lead and lag of the change in the log tax. The coefficient estimates in columns (2) and (3) are nearly identical. In both cases the model yields an implausibly large price elasticity of -1.14. The estimates in columns (1)-(3) essentially replicate the findings in Davis and Kilian (2011) for a similar specification.

We already provided direct evidence that an empirical specification including only contemporaneous values would be misspecified. A complementary approach to documenting this endogeneity is to test the overidentifying restrictions for the model in column (3), which includes more instruments than endogenous regressors and hence is overidentified. Based on Hansen's J -statistic, we reject the null hypothesis that the instruments are exogenous at the 5% significance level.² When the coefficients are overidentified, the two stage least squares estimator and GMM estimator differ in general; in the case of regression (3), the GMM estimator of coefficient on the change in log price is -0.98, slightly less elastic than (and more precisely estimated than) the two stage least squares estimator, but the two estimates are within a standard error of each other and are not qualitatively different.

As discussed earlier, the endogeneity of the tax instrument is the result of lag misspecification, so it can be corrected by including leads and lags of the change in price, along with leads and lags of the change in tax as instrumental variables. Our empirical evidence in section 3 suggests that including one lead and one lag of the change in gasoline prices is sufficient. The reason is that, in the presence of forward-looking behavior, one would expect gasoline purchases during the lead month to be unusually high and purchases during the month of the tax increase to be unusually low. Only in the month after the tax change the behavioral response to the tax change would become apparent.

Results for the specification with one lead and one lag can be found in column (4) and (5) for OLS and IV estimation, respectively. The IV coefficient on the contemporaneous

² The J -statistic is computed using the two-step GMM estimator with a weight matrix of the inverse of the long-run variance matrix estimated using clustered standard errors, where the time dummy variables are projected out in an initial step.

change in log price changes very little under this augmented specification, but the coefficients for both the lead and the lag are positive. The former, in particular, is also statistically significant. Table 2 also reports estimates of the cumulative effect of the tax change on gasoline purchases, calculated by summing the estimates of the lead, contemporaneous, and lag coefficients. Including the lead and the lag in the IV regression reduces the cumulative effect of a gasoline tax change to -0.37. This estimate is much smaller than the cumulative effect in columns (2) and (3) because the sharp contemporaneous decrease in gasoline purchases is largely offset by the increases in purchases during the month before and the month after the tax change is implemented. It is also statistically insignificant. The first-stage F and Cragg-Donald weak-instrument statistics are all large, exceeding the Stock-Yogo (2005) critical values, which suggest that the results in Table 2 do not suffer from weak instrument problems.³

Table 3 reports variations on regression (5) using alternative sets of instruments as sensitivity checks. First, as a specification check, regression (6) includes two leads and two lags of the tax change as instruments. Recall that adding a lag and a lead of tax changes to regression (2) resulted in a rejection of the overidentifying restrictions in regression (3). In contrast, the analogous augmentation of the instruments in regression (6) fails to reject the overidentifying restrictions (J -statistic p -value = .296), and the two-stage least squares and GMM estimates of the cumulative effect are respectively -0.360 and -0.314, both close to the estimate in regression (5). Thus the results in regression (6) provide confirmatory

³ The first stage F and Cragg-Donald statistics are computed under the assumption of homoscedasticity and no serial correlation, so that they can be compared to the Stock-Yogo (2005) critical values, which are valid only under homoscedasticity and no serial correlation. Because there is evidence of serially correlated errors in these regressions, these statistics are only illustrative. However we note that for all the regressions discussed here, except for regression (9) in Table 3, these statistics far exceed both the Stock-Yogo critical values and the rule-of-thumb value of 10, thus suggesting that for all but regression (9) the instruments are strong.

evidence that the original specification problem is resolved by including one lead and one lag of price changes.

A reasonable conjecture is that negligible changes in the nominal gasoline tax do not matter to gasoline buyers. Column (7) considers an alternative specification in which we instead instrument using an indicator variable for state-month observations in which there was a nominal gasoline tax change of at least 2 cents. For this specification we include two separate indicators, one for tax increases of at least 2 cents, and one for tax decreases by at least 2 cents. With this model we find a slightly lower and again imprecisely estimated cumulative effect of -0.29. Columns (8) and (9) consider tax increases and tax decreases separately. With tax increases in column (8), the results are very similar to regressions (5) and (7) suggesting that the results are largely driven by tax increases over 2 cents. With tax decreases in column (9) the point estimates are again similar, with a somewhat larger cumulative effect of -0.41, but the estimates are extremely imprecise. In contrast to the other IV regressions in Tables 2 and 3, the Cragg-Donald statistic for regression (9) suggests that the instruments are weak. This finding is not surprising given that there are only 14 incidents of such gasoline tax declines in the sample. Finally, column (10) includes as instruments the interaction between the change in the log tax and an indicator variable for gasoline tax increases of at least two cents, along with its lead and lag. One might have expected there to be a differential response of prices to gasoline tax increases of at least two cents. As it turns out, however, the estimates are extremely similar with this alternative specification. The interaction term is statistically insignificant in all of the first-stage equations, and the estimated cumulative effect is very similar, with a value of -0.34 compared to -0.37. Similarly to the augmented instrument set in regression (6), the J -

statistic does not reject the overidentifying restrictions test, further supporting the adequacy of the dynamic specification including a lead and lag of price changes.

We made the case that including one lead and one lag of the change in price makes sense a priori on economic grounds. Although we know that the responsiveness of gasoline buyers may increase over time, it seems implausible that such longer-run adjustments would occur within the first months of a tax change. Consistent with this reasoning, it can be shown that adding more lags (or another lead) to the regression model reduces the cumulative effect somewhat and renders the estimates less statistically significant, but none of these additional lead and lag terms are statistically significant, consistent with the evidence in Table 1.

5. A Comparison with other Estimates of the Price Elasticity of Gasoline Demand

Our analysis adds to a growing literature that suggests that gasoline buyers are more responsive to gasoline prices than recent studies such as Hughes et al. (2008) have indicated. There is a surprising degree of agreement even among studies using very different structural estimation methodologies. Our preferred point estimate of -0.37 (Table 2, regression (5)), for example, is close to the estimate of -0.35 reported in Bento et al. (2009) for the one-year price elasticity of gasoline demand conditional on the composition of the car fleet based on a microeconomic model of the markets for new, used, and scrap vehicles.

In other related work, Tiezzi and Verde (2013) show that single-equation estimation methods based on time series data tend to underestimate price elasticities by about 20% compared with estimates from systems of demand equations using cross-sectional data.

Their system-based estimate of the one-month own price elasticity of gasoline demand is -0.50, which would imply a single-equation estimate of -0.40, only slightly higher than our point estimate of -0.37. Similar system-based estimates can be found in Nicol (2003), Oladosu (2003), and West and Williams (2004, 2007).

Finally, our estimate is also consistent with recent evidence in Levin et al. (2014). Levin et al. use daily expenditure data and prices from credit card transactions between February 1, 2006, and December 31, 2009, for 243 United States cities to estimate the price responsiveness of the daily demand for gasoline to changes in daily gasoline prices. Their approach thus addresses the measurement issues with the gasoline quantity data we discussed in section 3. For a panel of U.S. states, they arrive at monthly elasticity estimates of between -0.18 and -0.27, depending on how purchases at the pump are treated, even without accounting for price endogeneity. These estimates are two or three times as high as OLS estimates based on the aggregate data for gasoline purchases conventionally used in estimating the price elasticity of gasoline demand, consistent with our conjecture that stockpiling by gasoline stations and gasoline distributors is an important feature of the data.

There is also a link between our analysis and the related literature on the short-run price elasticity of demand for crude oil. Kilian and Murphy (2014) report a point estimate of -0.25 for the short-run price elasticity of oil demand in global oil markets. Given that crude oil accounts for about half the U.S. retail cost of gasoline, as pointed out by Hamilton (2009), one would expect the retail price elasticity of gasoline demand to be about twice that for crude oil. A reasonable conjecture is that this global price elasticity is somewhat higher than the U.S. price elasticity of oil demand. If the U.S. elasticity were two thirds of

the global elasticity, for example, this would imply a price elasticity of gasoline demand of -0.33, close to our point estimate.

6. Implications for Future U.S. Gasoline Consumption

An important question is how representative our preferred estimate of the cumulative effect in column (5) of Table 2 is for the price elasticity of gasoline demand when dealing with changes in the gasoline price reflecting exogenous variation other than gasoline tax changes. We have already discussed several reasons why one might expect consumers to respond less to typical gasoline price changes than they do to changes in gasoline taxes.

This point has received some attention in the recent literature. For example, Li et al. (2014) find much larger responses to changes in gasoline taxes than to changes in gasoline prices, and Tiezzi and Verde (2014) find that the elasticity of demand with respect to a change the gasoline tax is about 20% higher than the elasticity of demand with respect to prices.

However, without a valid instrument for gasoline prices it is difficult to know how much of this difference is due to price endogeneity, so this evidence does not help in assessing the external validity of our elasticity estimate.

We therefore provide a simple back-of-the-envelope calculation aimed at evaluating more directly how plausible our preferred elasticity estimate is in practice. A suitable natural experiment is the exogenous gasoline supply shock caused by Hurricanes Rita and Katrina in late 2005. These exogenous weather events caused a shutdown of oil refineries along the Gulf coast in September 2005, resulting in a major negative gasoline supply shock (see Kilian 2010). We make the assumption that the U.S. supply of gasoline was not affected by these events before September 2005. Given that there was little warning time

before these hurricanes made landfall, we proceed under the additional assumption that the shortage of gasoline was not anticipated in August. The observed increase in the price of gasoline in September 2005 relative to August was 15.3%. Given our elasticity estimate of -0.37, we would have expected a decline in gasoline consumption in September of -5.6%. The actual decline was -6.5%, suggesting that our one-month price elasticity appears to fit the data quite well. If anything, we slightly underestimate the decline in gasoline purchases. Of course, some of that observed decline presumably reflected a drop in final consumption in the area directly affected by the Hurricane, rather than a drop induced by the price increase.

One of the primary rationales for estimating the price elasticity of gasoline demand is the ability to make predictions about gasoline consumption given some hypothesized exogenous change in gasoline prices. There has been considerable interest in the question of how the steep decline in U.S. gasoline prices after June 2014 is expected to affect U.S. gasoline consumption. U.S. gasoline prices between June and December 2014 declined at an average monthly rate of 6%. Treating these price declines as exogenous with respect to the U.S. gasoline market, our elasticity estimate of -0.37 implies a cumulative increase in gasoline consumption by about 13% (or 1.1 million barrels per day) all else equal.

This estimate appears implausibly large. One reason is that this thought experiment does not take account of predictable variation in gasoline prices. Whereas Hurricanes Rita and Katrina were largely unpredictable shocks, the decline in gasoline prices since June 2014 was predictable in real time to a considerable extent (see Baumeister, Kilian and Lee 2015). This observation suggests that a more appropriate approach would be to focus on the unpredictable component of the change in gasoline prices for each month between July

and December 2014. For expository purposes, we rely on the real-time gasoline price predictions issued by the U.S. Energy Information Administration in its *Short-Term Energy Outlook*. Applying the elasticity estimate to the unpredictable component of the change in gasoline prices (and assuming that this component is exogenous with respect to U.S. gasoline consumption) yields a much lower and more plausible expected cumulative increase in U.S. gasoline consumption of 8% relative to June 2014 levels.

7. Concluding Remarks

The ability to store gasoline means that purchases of gasoline need not correspond to actual gasoline consumption. In this paper, we argued that this fact has important implications for the specification of gasoline demand models. Put simply, gasoline purchases respond to expected changes in gasoline prices. While most price changes are difficult to predict, gasoline tax changes are known well in advance, so forward-looking gasoline buyers will accelerate gasoline purchases in the days leading up to gasoline tax increases, and delay gasoline purchases in the days leading up to gasoline tax decreases.

This anticipatory behavior creates an endogeneity problem which undermines the validity of conventional IV estimates of the price elasticity of gasoline demand. We showed how this concern can be addressed by augmenting the regression model to include one lead and one lag of the change in log gasoline prices and augmenting the instruments with the corresponding leads and lags of the change in gasoline taxes. Using monthly data from January 1989 through March 2008, we found that the estimated price elasticity decreases from an implausibly high -1.14 to a much more plausible -0.37 after including one lead and one lag. Although this estimate is imprecise (standard error = 0.24), it is similar in

magnitude to other recent elasticity estimates that take explicit account of the endogeneity of gasoline prices. Moreover, a simple back-of-the-envelope calculation suggests that this estimate appears to fit the data quite well.

Our analysis has implications for other empirical applications as well. Many purchases other than gasoline are storable, allowing buyers to strategically time their purchases in anticipation of future price changes. This strategic behavior has the potential of undermining the validity of a host of seemingly exogenous candidate instruments. For example, a commonly used instrument is changes in the weather, but weather changes in turn are often predictable at a time horizon that could matter for monthly data. In such circumstances, allowing for additional leads and lags in the regression may help address the resulting endogeneity of the instrument. In addition, our results highlight that it may matter greatly how the quantity variable is measured when estimating the price elasticity of demand. This point also applies to other markets. A case in point is the natural gas market (see, e.g., Hausman and Kellogg, 2014). Unlike gasoline buyers, retail consumers of natural gas typically have no ability to store natural gas, so purchases and consumption occur at the same time. Throughout the rest of the supply chain, however, there is extensive storage capacity. Thus, short-run price elasticities of demand may differ substantially based on whether quantities are measured at the retail level or further upstream.

References

Allcott, H., and N. Wozny (2014), "Gasoline Prices, Fuel Economy, and the Energy Paradox," *Review of Economics and Statistics*, 96, 779-795.

- Baumeister, C., Kilian, L., and T.K. Lee (2015), "Inside the Crystal Ball: New Approaches to Predicting the Gasoline Price at the Pump," working paper, University of Michigan.
- Bento, A.M., L.H. Goulder, Jacobson, M.R., and R.H. von Haefen (2009), "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes," *American Economic Review*, 99, 667-699.
- Brons, M., Nijkamp, P., Pels, E., and P. Rietveld (2008), "A Meta-Analysis of the Price Elasticity of Gasoline Demand: A SUR Approach," *Energy Economics*, 30, 2105-2022.
- Chetty, R., Looney, A., and K. Kroft (2009), "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99, 1145-1177.
- Dahl, C. (1979), "Consumer Adjustment to a Gasoline Tax," *Review of Economics and Statistics*, 61, 427-432.
- Dahl, C., and T. Sterner (1991), "Analyzing Gasoline Demand Elasticities: A Survey," *Energy Economics*, 13, 203-210.
- Davis, L.W., and L. Kilian (2011), "Estimating the Effect of a Gasoline Tax on Carbon Emissions," *Journal of Applied Econometrics*, 26, 1187-1214.
- Edelstein, P., and L. Kilian (2009), "How Sensitive are Consumer Expenditures to Retail Energy Prices?", *Journal of Monetary Economics*, 56, 766-779
- Espey, M. (1998), "Gasoline Demand Revisited: An International Meta-Analysis of Elasticities." *Energy Economics* 20, 273-295.
- Federal Highway Administration (2008), "Highway Taxes and Fees: How They Are Collected and Distributed", Office of Highway Policy Information, <https://www.fhwa.dot.gov/policyinformation/motorfuel/hwytaxes/2008/>.
- Graham, D.J., and S. Glaister (2004), "Road Traffic Demand Elasticity Estimates: A Review,"

- Transport Reviews*, 24, 261-274.
- Greene, D.L., J.R. Kahn, and R.C. Gibson (1999), "Fuel Economy Rebound Effect for U.S. Household Vehicles," *Energy Journal*, 20, 1-30.
- Hamilton, J.D. (2009), "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 1, 215-259
- Hausman, J.A., and W.K. Newey (1995), "Nonparametric Estimation of Exact Consumers' Surplus and Deadweight Loss," *Econometrica*, 63, 1445-1476.
- Hausman, C., and R. Kellogg (2015), "Welfare and Distributional Implications of Shale Gas," working Paper, University of Michigan.
- Hughes, J.E., Knittel, C.R., and D. Sperling (2008), "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *Energy Journal*, 29, 93-114.
- Kilian, L. (2010), "Explaining Fluctuations in U.S. Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market," *Energy Journal*, 31, 87-104.
- Kilian, L. and D.P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454-478.
- Levin, L., Lewis, M.S., and F.A. Wolak (2014), "High-Frequency Evidence on the Demand for Gasoline," manuscript, Stanford University.
- Li, S., Linn, J. and E. Muehlegger (2014), "Gasoline Taxes and Consumer Behavior," *American Economic Journal: Economic Policy*, 6, 302-342.
- McCaffery, E.J., and J. Baron (2006), "Thinking about Tax," *Psychology, Public Policy, and Law*, 12, 106-135.
- Ramsey, J., Rasche, R., and B. Allen (1975), "An Analysis of the Private and Commercial

- Demand for Gasoline," *Review of Economics and Statistics*, 57, 502–507.
- Rivers, Nicholas and Brandon Schaufele (2014), "Salience and Carbon Taxes in the Gasoline Market," manuscript, *Harvard University*.
- Scott, K.R. (2012), "Rational Habits in Gasoline Demand," *Energy Economics*, 34, 1713-1723.
- Stock, J.H. and M. Yogo (2005), "Testing for Weak Instruments in Linear IV Regression," (2005), Ch. 5 in J.H. Stock and D.W.K. Andrews (eds), *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*, Cambridge University Press.
- Stock, J.H., Yogo, M., and J.H. Wright (2005), "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics*, 20, 518-529.
- Sweeney, J.L. (1984), "The Response of Energy Demand to Higher Prices: What Have We Learned?" *American Economic Review*, 74, 31-37.
- Sweeney, R.L. (2015), "Environmental Regulation, Imperfect Competition and Market Spillovers: The Impact of the 1990 Clean Air Act Amendments on the U.S. Oil Refining Industry," *Manuscript*.
- Tiezzi, S. and S.F. Verde (2014), "Overreaction to Excise Taxes: The Case of Gasoline," EUI Working Paper RSCAS 2014/54, European University Institute.
- U.S. Energy Information Administration (2015), *Short-Term Energy Outlook*, January.
- Wojciech, K., Marion, J., Muehlegger, E., and J. Slemrod (2013), "Do the Laws of Tax Incidence Hold? Point of Collection and the Pass-Through of State Diesel Taxes," NBER Working Paper No. 19410.

Table 1: Assessing the Statistical Significance of Shifts in Gasoline Purchases Before, During and after Tax Increases

	(a)	(b)	(c)	(d)
	<u>Tax Increases</u>		<u>Tax Decreases</u>	
	Of At Least One Cent	Of At Least Two Cents	Of At Least One Cent	Of At Least Two Cents
Indicator for Tax Change (One Month Lead)	0.013* (0.007)	0.021** (0.009)	-0.003 (0.011)	-0.016 (0.010)
Indicator for Tax Change (Contemporaneous)	-0.031*** (0.008)	-0.038*** (0.010)	0.023* (0.013)	0.030** (0.013)
Indicator for Tax Change (One Month Lag)	0.002 (0.005)	0.008 (0.007)	-0.017** (0.008)	-0.011 (0.009)
Time Fixed Effects	Yes	Yes	Yes	Yes
Change in Unemployment	Yes	Yes	Yes	Yes
Observations	11,730	11,730	11,730	11,730
Number of Tax Changes	140	86	22	14
R-squared	0.466	0.466	0.463	0.463

NOTE: This table reports coefficient estimates and standard errors from four separate least squares regressions. The dependent variable in all specifications is the month-to-month change in log gasoline purchases and the coefficients of interest correspond to indicator variables for the month before, the month during, and the month after state gasoline tax changes. All regressions are estimated using the complete panel of 50 states plus the District of Columbia, each observed for 230 months, for a total of 11,730 observations. Standard errors are clustered by state. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 2: IV Estimates of the Effect of Gasoline Price Changes on Gasoline Consumption

	(1)	(2)	(3)	(4)	(5)
	OLS	IV	IV	OLS	IV
		All Tax Changes, Just Identified	All Tax Changes, Over-identified	With Lead and Lag	All Tax Changes, With Lead and Lag
Change in Log Price				0.136***	0.540***
(One Month Lead)				(0.0360)	(0.171)
Change in Log Price	-0.190***	-1.135***	-1.135***	-0.192***	-1.152***
(Contemporaneous)	(0.0368)	(0.250)	(0.243)	(0.0361)	(0.250)
Change in Log Price				0.0579	0.244
(One Month Lag)				(0.0394)	(0.197)
Observations	11,730	11,730	11,628	11,628	11,628
R ²	0.468			0.473	
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Change in Unemployment	Yes	Yes	Yes	Yes	Yes
Cumulative Effect	-0.190***	-1.135***	-1.135***	0.002	-0.368
	0.0368	(0.250)	(0.243)	(0.076)	(0.239)
Instruments:					
Change in Log Tax	--	t	t-1, t, t+1	--	t-1, t, t+1
First Stage F / Cragg-Donald Statistic of Weak Identification	--	237.0 [*]	78.7 [*]	--	60.0 [*]
Hansen's J Statistic testing Overidentifying Restrictions (p-value in parentheses)	--	--	7.062 (0.029)	--	--
Cumulative Effect (GMM)	--	--	-0.977*** (0.150)	--	--


NOTE: This table reports coefficient estimates and standard errors from two ordinary least squares and six instrumental variables (IV) regressions. The dependent variable in all specifications is the month-to-month change in log gasoline purchases and the coefficients of interest correspond to the month-to-month change in log gasoline price, including a lead and a lag. Instruments vary across specifications, as indicated under the instruments section. In each column, we specify whether the contemporaneous instrument (denoted by t) is accompanied by a lead and a lag ($t+1$ and $t-1$, respectively). The statistic testing for weak identification in the IV regressions is either the first stage F or, when there are multiple included endogenous regressors, the Cragg-Donald statistic, and a “+” denotes whether the test statistic exceeds the 5% Stock-Yogo critical value assuming 10% bias. The weak identification statistics are computed under the assumption of homoskedasticity. The J -statistic is computed for the GMM estimator using clustered long-run variance weight matrix; the GMM estimator of the sum of price coefficients and its standard error are also reported if there is overidentification. Standard errors on regression coefficients are clustered by state. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 3: IV Estimates of the Effect of Gasoline Price Changes on Gasoline Consumption

	(6)	(7)	(8)	(9)	(10)
	IV	IV	IV	IV	IV
	All Tax Changes, With Additional Leads and Lags	Tax Changes Greater Than Two Cents, With Lead and Lag	Increases Greater Than Two Cents, With Lead and Lag	Decreases Greater Than Two Cents, With Lead and Lag	All Tax Changes, With Lead, Lag, and Interactions
Change in Log Price	0.553***	0.554***	0.559**	0.484	0.547***
(One Month Lead)	(0.183)	(0.211)	(0.223)	(0.501)	(0.175)
Change in Log Price	-1.152***	-1.135***	-1.131***	-1.192*	-1.129***
(Contemporaneous)	(0.251)	(0.265)	(0.261)	(0.617)	(0.243)
Change in Log Price	0.239	0.291	0.291	0.299	0.238
(One Month Lag)	(0.195)	(0.219)	(0.212)	(0.634)	(0.193)
Observations	11,526	11,628	11,628	11,628	11,628
Number of Tax Changes	N/A	100	86	14	N/A
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Change in Unemployment	Yes	Yes	Yes	Yes	Yes
Cumulative Effect	-0.360	-0.291	-0.281	-0.409	-0.344
	(0.241)	(0.292)	(0.285)	(0.950)	(0.243)
Instruments:					
Change in Log Tax	t-2, t-1, t, t+1, t+2	--	--	--	t-1, t, t+1
Indicator for Tax <u>Increase</u> of at least 2 Cents	--	t-1, t, t+1	t-1, t, t+1	--	--
Indicator for Tax <u>Decrease</u> of at least 2 Cents	--	t-1, t, t+1	--	t-1, t, t+1	--
Change in Log Tax Interacted with Indicator for Tax Increase of at least 2 Cents	--	--	--	--	t-1, t, t+1
First Stage F / Cragg-Donald Statistic of Weak Identification	35.9+	26.3+	49.4+	3.1	30.3+
Hansen's J Statistic testing Overidentifying Restrictions and (p-value in parentheses)	2.435 (0.296)	0.023 (0.999)	--	--	2.029 (0.566)
Cumulative Effect (GMM)	-0.314 (0.231)	-0.282 (0.277)	--	--	-0.452 (0.221)


NOTE: See notes to Table 2.

Figure 1: U.S. Gasoline Consumption during 1974.1-2014.11 with Phases of Rising Real Gasoline Prices Imposed as Shaded Areas




NOTES: U.S. gasoline consumption is computed as the sum of monthly gasoline consumption by the residential and commercial, industrial, and transportation sectors, as reported in the EIA's Monthly Energy Review. The series has been seasonally adjusted. Phases of rising and falling real gasoline prices are determined based on the major peaks and troughs in the evolution of the real U.S. price of gasoline, as reported by the EIA.

Figure 2a: Event Study of Gasoline Tax Increases of At Least 1 Cent




NOTES: The full sample includes the complete panel of 50 states plus Washington D.C., each observed for 230 months for a total of 11,730 observations. The event study is based on 140 gasoline tax increases in our sample of at least one cent. Time is normalized relative to the month of the change ($t=1$). We first regress month-to-month changes in the dependent variable on changes in the state-level unemployment rate and a complete set of month-of-sample fixed effects. The figures show the mean residual from this regression for each period in event time.

Figure 2b: Event Study of Gasoline Tax Increases of At Least 2 Cents




NOTES: The event study is identical to Figures 2a, except it is based on the 86 gasoline tax increases in our sample of at least two cents.

Figure 2c: Event Study of Gasoline Tax Decreases of At Least 1 Cent



NOTES: This event study is identical to that in Figures 2a and 2b, except it is based on the 22 gasoline tax decreases in our sample of at least one cent.

Figure 2d: Event Study of Gasoline Tax Decreases of At Least 2 Cents



NOTES: The event study is identical to that in Figures 2a, 2b, and 2c, except it is based on the 14 gasoline tax decreases in our sample of at least two cents.