

ESTIMATING THE EFFECT OF A GASOLINE TAX ON CARBON EMISSIONS

LUCAS W. DAVIS^{a,b,*} AND LUTZ KILIAN^{c,d}

^a *Haas School of Business, University of California, Berkeley, CA, USA*

^b *National Bureau of Economic Research, Cambridge, MA, USA*

^c *Department of Economics, University of Michigan, Ann Arbor, MI, USA*

^d *Centre for Economic Policy Research, London, UK*

SUMMARY

Recently the proposal has been made to raise gasoline taxes in the United States to curb carbon emissions. The existing literature on the sensitivity of gasoline consumption to changes in price may not be appropriate for evaluating the effectiveness of such a tax. First, most of these studies fail to address the endogeneity of gasoline prices. Second, the responsiveness of gasoline consumption to a change in tax may differ from the responsiveness of consumption to an average change in price. We address these challenges using a variety of methods including traditional single-equation regression models, estimated by least squares or instrumental variables methods, and structural vector autoregressions. Our preferred approach exploits the historical variation in US federal and state gasoline taxes. Our most credible estimates imply that a 10-cent per gallon increase in the gasoline tax would reduce carbon emissions from vehicles in the United States by about 1.5%. Copyright © 2010 John Wiley & Sons, Ltd.

Received 5 April 2009; Revised 9 September 2009

1. INTRODUCTION

Several policymakers and economists have proposed the adoption of a carbon tax in the United States.¹ In the United States 33.8% of carbon dioxide emissions are derived from the transportation sector, so the responsiveness of gasoline consumption to tax changes plays a significant role in determining the evolution of overall carbon dioxide emissions in response to policy interventions.² It is widely recognized that a carbon tax in practice must take the form of a tax on the consumption of energy products such as gasoline (e.g. Fullerton and West, 2002). A tax of \$10.00 per ton of carbon dioxide, as suggested by Nordhaus (2007), for example, would increase gasoline taxes by approximately 9 cents.³ In this paper, we evaluate how effective a gasoline tax increase would be in reducing gasoline consumption.

* Correspondence to: Lucas W. Davis, Haas School of Business, Berkeley, University of California, CA 94720-1900, USA. E-mail: ldavis@haas.berkeley.edu

¹ For example, in September 2007 Representative John Dingell, then chairman of the House Energy and Commerce Committee, proposed a \$50 tax per ton of carbon (\$13.64 per ton of carbon dioxide). See H. Josef Hebert 'Plan Uses Taxes to Fight Climate Change', Associated Press (2 October 2007) for details. More recently, in June 2009 the US House of Representatives approved the American Clean Energy and Security Act (also known as the Waxman–Markey Bill), which would establish a cap-and-trade program for greenhouse gas emissions, effectively placing a price on carbon emissions in the United States.

² US Department of Energy, 'Energy Outlook 2008', DOE/EIA-0383(2008), Table A18, Carbon Dioxide Emissions by Sector and Source, 2006 and 2030.

³ A carbon tax of \$35 per ton from Nordhaus (2007) implies a tax of approximately \$10 per ton of carbon dioxide because the atomic weight of carbon is 12 atomic mass units, while the weight of carbon dioxide is 44, so one ton of carbon

Our most credible estimates imply that a 10-cent tax increase would decrease US carbon emissions from the transportation sector by about 1.5% and decrease total US carbon emissions by about 0.5%. To put this estimate in context, total US carbon dioxide emissions increased by 1.1% annually between 1990 and 2007, so a 10 cent gasoline tax increase would approximately offset half a year of growth in total US emissions.⁴ This estimate captures only the short-run response resulting from reduced discretionary driving and reduced driving speeds, for example. The long-run response is likely to be considerably larger as drivers substitute toward more fuel-efficient vehicles.

Despite the policy relevance of the question to be addressed in this paper, empirical evidence on the effectiveness of gasoline taxes on carbon emissions is virtually nonexistent. A large related literature has examined the sensitivity of gasoline consumption to changes in prices. Reviews by Dahl and Sterner (1991) and Greene *et al.* (1999) identify dozens of relevant econometric studies containing estimates of price elasticities. These estimates, however, are not appropriate for evaluating the effectiveness of gasoline taxes. First, most of these studies do not address 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. Second, the response of gasoline consumption to a change in tax is likely to differ from its response to an average change in price. Price changes induced by tax changes are more persistent than other price changes and thus may induce larger behavioral changes. In addition, gasoline tax increases are often accompanied by media coverage that may have an effect of its own.

In this paper we address these challenges using a variety of methods including traditional single-equation regression models, estimated by least squares or instrumental variables methods, and structural vector autoregressions. We compare the results from these approaches, highlighting the advantages and disadvantages of each. Our preferred approach exploits the historical variation in US federal and state gasoline taxes. Where possible, lessons learned from applying one model are used in specifying another. Although all models have their limitations, the overall exercise serves to document what can and what cannot be learned from the available data. Such a discussion seems warranted before adopting any policy proposal regarding gasoline taxes.

We begin by modeling gasoline consumption at the national level using monthly data from 1989 to 2008.⁵ The use of monthly data facilitates the identification of the causal effects of a gasoline tax increase. To allow a direct comparison with the previous literature, we first focus on the construction of price elasticities. Using a single-equation specification in log differences we find a price elasticity of -0.10 , similar to results from a recent study by Hughes *et al.* (2008). Although we are sympathetic to the idea that the price elasticity is low, standard concerns about price endogeneity with respect to quantity suggest that this estimate may be biased toward zero. Next we model gasoline consumption at the state level using panel data methods. The advantage of the panel specification is that we can include time fixed effects to control for unobserved time-varying consumption factors that may have obscured the true elasticity in the national data. With this specification we find a price elasticity of -0.19 —still small, but significantly larger than the estimate based on the national data.

equals 44/12 tons of carbon dioxide. To avoid confusion we use carbon dioxide throughout. The carbon dioxide content of gasoline is approximately 0.0088 metric tons per gallon, so a \$10.00 tax per ton implies 8.8 cents per gallon.

⁴ US Department of Energy, Energy Information Administration, *Emissions of Greenhouse Gases Report 2007*, DOE/EIA-0573, Table V, 'U.S. Carbon Dioxide Emissions from Energy and Industry, 1990, 1995, and 2000–2007'.

⁵ A complete series of monthly data by state for gasoline consumption is not available prior to 1989.

Even this panel approach, however, does not fully address the issue of price endogeneity. An alternative approach to this endogeneity problem has been the use of instruments 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).⁶ In this paper we use 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 believe that changes in tax rates are uncorrelated with unobserved changes in demand. In constructing our instrument we are careful to exclude *ad valorem* gasoline taxes (used in many states) because they are functionally related to price, violating the endogeneity assumption.

In our preferred instrumental variable (IV) specification that restricts attention to dates of nominal state tax increases, we find a statistically significant price elasticity of -0.46 . We examine a variety of alternative IV specifications including specifications that control for factors potentially correlated with gasoline tax changes and we consider alternative estimators. The results are remarkably similar across specifications. In addition, we contrast the IV estimates to alternative elasticity estimates obtained from recursively identified vector autoregressions in which the percentage change in gasoline prices (or alternatively the percentage change in gasoline taxes) is ordered first and the percentage change in gasoline consumption is ordered second. With this approach the elasticity is identified under the assumption that changes in gasoline prices (or taxes) are predetermined with respect to all factors driving US gasoline consumption.⁷

Overall, our results indicate that gasoline consumption is more sensitive to gasoline taxes than would be implied by recent estimates of the gasoline price elasticity. Even under the largest plausible estimates, however, gasoline tax increases of the magnitude that have been discussed would have only a moderate short-run impact on total US gasoline consumption and carbon emissions based on our estimates. A natural conjecture is that the long-run elasticities will be larger, but standard econometric models based on historical data do not allow the prediction of such long-run effects. The paper is organized as follows. Section 2 describes the data used for the analysis and describes the evolution of gasoline taxes in the United States over the period 1989–2008. Section 3 presents least squares and IV estimates from log differenced single-equation specifications. Section 4 presents results from structural VAR models. Section 5 assesses the effects of the proposed gasoline tax on carbon emissions, and Section 6 summarizes the policy implications.

2. DATA

Figure 1 describes national gasoline prices, consumption, and taxes over the period 1989–2008 from the US Department of Energy and US Department of Transportation. The data are monthly.

⁶ 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, Knittel and Sperling instrument using changes in global crude oil production such as a strike by oil workers in Venezuela in 2002.

⁷ The assumption of predetermined prices rules out instantaneous feedback from quantities to prices. In other words, gasoline prices respond to changes in demand only with a delay. This identifying assumption permits the consistent estimation of the expected response of consumption to innovation in prices. In conjunction with the assumption that there are no other exogenous events that are correlated with the exogenous gasoline price innovation, these responses can be interpreted as the causal effect of the price innovation. Similar impulse response models have been estimated by a variety of authors in related contexts (see, for example, Edelstein and Kilian, 2009; Hamilton, 2009).

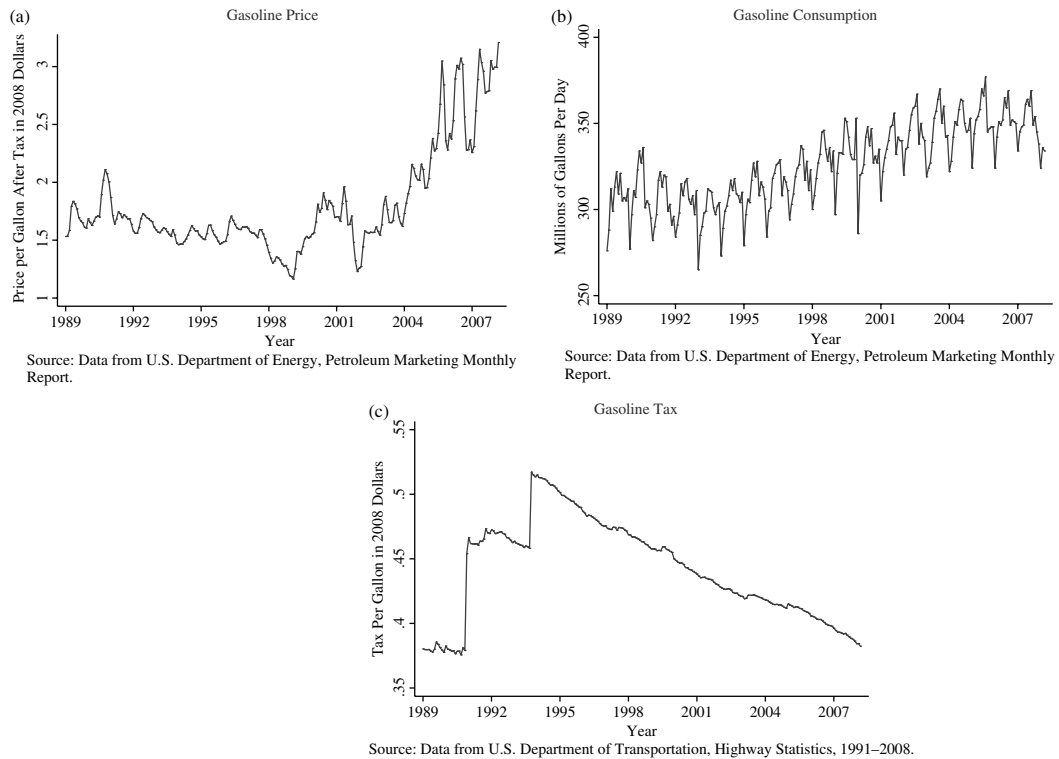


Figure 1. National gasoline prices, consumption and taxes

Gasoline prices and taxes have been deflated using the US consumer price index. Our sample period covers January 1989 through March 2008. We focus on this period because a continuous time series of gasoline consumption by state is not available prior to 1989. Moreover, recent evidence from Hughes *et al.* (2008) indicates that gasoline price elasticities have declined since the mid 1980s, making the 1990s and 2000s the most relevant period for understanding current demand behavior. Both the price and quantity time series exhibit a clear seasonal pattern. Real prices fluctuate between \$1.50 and \$2.00 between 1989 and 2004 before increasing to \$3.00 and above in 2005.⁸ Gasoline consumption increases steadily during the 1990s and then appears to level off after 2005.⁹ Tax per gallon increases with federal tax increases at the end of 1990 and in 1993, and then steadily erodes in real terms.

⁸ Gasoline prices are from US Department of Energy, Energy Information Administration, 'Petroleum Marketing Monthly Report: Gasoline Prices by Formulation, Grade, Sales Type'. Prices are collected by the EIA using the EIA-782B, a monthly survey of approximately 2200 gasoline resellers and retailers and reflect average prices for sales to end users of all grades of motor gasoline. Here and throughout, where appropriate, we calculate national average prices weighting by volume.

⁹ Sales volumes come from US Department of Energy, Energy Information Administration, 'Petroleum Marketing Monthly Report: Prime Supplier Sales Volumes by Product and Area'. Sales volumes are collected by the EIA using the EIA-782C, a monthly survey of all prime suppliers, a small group of currently 185 firms nationwide that produce, import, or transport petroleum products across state boundaries and local marketing areas. The EIA-782C measures sales volumes within each state. Accordingly, respondents are asked to report gasoline sales to local distributors, local retailers, and local end users, excluding sales to other prime suppliers and to other companies that transport gasoline across state lines.

Prices from the US Department of Energy exclude federal, state, and local taxes. Accordingly, for this study we constructed a comprehensive panel of state-level gasoline tax rates and after-tax prices.¹⁰ In constructing taxes we were careful to distinguish between ‘per unit’ and *ad valorem* taxes. The most common method for taxing gasoline is to establish a ‘per unit’ tax per gallon. This is the method used for the federal gasoline tax and most states. However, a handful of states use *ad valorem* taxes in addition to ‘per unit’ taxes. Gasoline sales in California, for example, are subject to the state’s sales tax. As gasoline prices have increased, these states have seen large increases in tax receipts. We include *ad valorem* taxes in our measure of after-tax prices. However, it is important to exclude this variation in taxes when we use tax as an instrumental variable because these changes fail the exogeneity test; indeed the gasoline tax in these states is functionally related to price. Finally, as an overall test of the validity of our measure of the after-tax price of gasoline, we compare our measure (weighted by volume) to the city average price of gasoline from the Bureau of Labor Statistics in Figure 2. It is reassuring that the two series are very similar.

The advantage of this data source is that it provides data for all states and years. An alternative to using state-level aggregate prices and consumption would be to use available household surveys such as the Consumer Expenditure Survey (CEX) or the Residential Transportation and Energy Consumption Surveys (RTECS). For example, Schmalensee and Stoker (1999) use the RTECS to describe non-parametrically the relationship between gasoline consumption and household income. Whereas the RTECS is valuable for cross-sectional studies such as Schmalensee and Stoker (1999), it is not appropriate in our context because the survey is carried out only once every 3 years. Furthermore, whereas the RTECS measure of gasoline consumption is

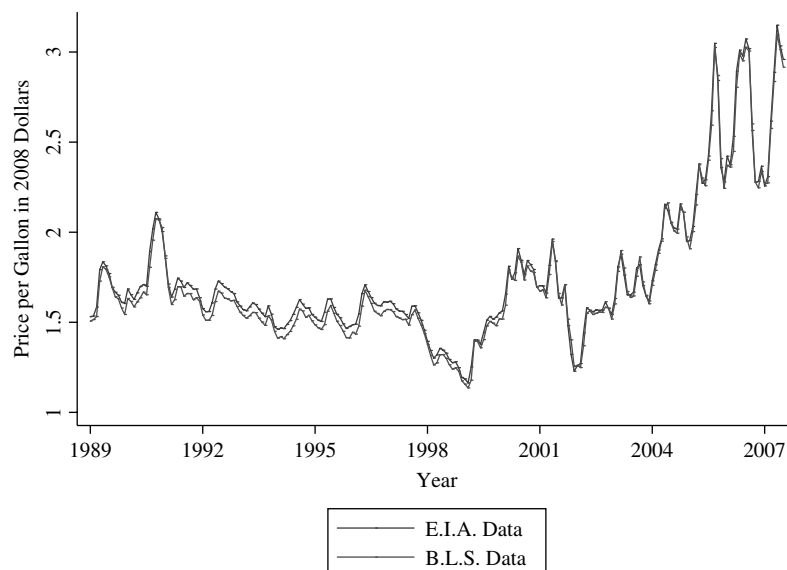


Figure 2. Comparison of BLS data and tax-adjusted EIA data on gasoline prices

¹⁰ US Department of Transportation, ‘Highway Statistics’, 1990–2008, publishes state tax rates for gasoline by year as well as the exact dates of all changes.

based on odometer readings, the CEX measure is calculated using self-reported expenditure on gasoline and thus is likely subject to more reporting error than the Department of Energy data.

Figures 3 and 4 describe the variation in tax rates across states and over time. States vary both in the size and frequency of tax rate changes. In real terms there is more variation during the first half of the period because the overall level of taxes is higher. Finally, Figure 5 shows the timing of tax increases and decreases that exceed 2 cents per gallon. The number of tax increases far exceeds the number of tax decreases and since 2005 there have been no changes to the gasoline tax in any state. An obvious concern is that the timing of these tax changes may be correlated with state-level macroeconomic conditions and other factors that influence gasoline demand. We will address potential concerns about the exogeneity of tax changes below.

3. SINGLE-EQUATION MODELS OF GASOLINE CONSUMPTION

3.1. Least Squares Estimates

A key feature of our analysis is the use of monthly rather than yearly observations. Most importantly, this allows us, in later specifications, to describe the entire temporal path of the response of gasoline consumption to changes in gasoline prices. The monthly measures also reveal variation in price and consumption that may be obscured in annual data.

We begin by modeling gasoline consumption at the national level. As a starting point, consider the following commonly used specification (see, for example, Hughes *et al.*, 2008). National gasoline consumption in month t in logs y_t is postulated to depend linearly on the volume-weighted average inflation-adjusted after-tax price of gasoline (in logs) p_t , month-of-the-year dummies λ_t , and unobserved idiosyncratic time-varying factors ε_t :

$$y_t = \alpha_0 + \alpha_1 p_t + \lambda_t + \varepsilon_t$$

The month-of-the-year dummies are important because the gasoline market is highly seasonal.

Whereas most previous studies have estimated gasoline consumption equations in log levels, we estimate the national aggregate model using log differences, since both y_t and p_t are highly persistent and trending:

$$\Delta y_t = \alpha_0 + \alpha_1 \Delta p_t + \lambda_t + \varepsilon_t \quad (1)$$

Thus the response of gasoline consumption is identified using price changes rather than price levels.

In addition to this national specification, we estimate a state panel using our balanced panel of 50 states and the District of Columbia:

$$\Delta y_{it} = \beta_0 + \beta_1 \Delta p_{it} + \rho_t + \omega_{it} \quad (2)$$

where gasoline consumption in logs y_{it} for state i and month t depends linearly on the after-tax price of gasoline in logs p_{it} , time fixed effects ρ_t , and unobserved idiosyncratic state-specific




Figure 3. Variation in gasoline tax 1989–2008, by state




Figure 4. Variation in state tax rates across states, by year

time-varying factors ω_{it} .¹¹ Here the time fixed effects control for both seasonal variation (as with the month-of-the-year dummies) and variation from year to year that is the same across states.

The coefficients of interest are α_1 in the national specification and β_1 in the state panel.¹² With the log-log specification these coefficients describe the 1-month price elasticity. From a policy standpoint the 1-month elasticity is of limited interest. Rather we would like to know how households respond over a longer period of time such as 1 year. As we will show in Section 4, the estimates of 1-month and 1-year elasticities obtained from alternative estimation approaches are quite similar, suggesting that much of the adjustment takes place on impact, and that the 1-month price elasticities are likely to be representative for horizons as long as 1 year.

Table I presents parameter estimates and standard errors from the national aggregate and state panel specifications. In the national specification, the price elasticity is -0.10 and statistically significant at the 1% level. As a point of comparison, Hughes *et al.* (2008) using a level specification and similar data for 2001–2006 find a price elasticity of -0.04 , which is not much different from our estimate after accounting for sampling uncertainty. In contrast, in the state panel specification, the price elasticity is almost twice as large. The estimate of -0.19 is statistically significant at the 1% level. In order to understand the difference between these two estimates it is helpful to recall the key identifying assumptions in each model. The estimate in the national

¹¹ It is important to emphasize that the log differences specification controls implicitly for time-invariant factors, rendering state fixed effects unnecessary. One disadvantage of this specification is that it makes it difficult to estimate the income elasticity, because income changes very slowly over time. Since the focus of our paper is the price elasticity, this is of little practical concern.

¹² It can be shown that including additional lags in models (1) and (2) does little to change the estimates of α_1 and β_1 .




Figure 5. Timing of tax changes, 1989–2008

specification is consistent if price changes are uncorrelated with ε_t :

$$E(\Delta p_t \varepsilon_t) = 0 \quad \forall t$$

This orthogonality condition is unlikely to hold because of standard price endogeneity considerations. Increases in gasoline consumption cause prices to increase, leading to a spurious correlation between Δp_t and ε_t . This correlation will, in general, cause estimates of the price elasticity to be biased toward zero as all of the predicted change in consumption is attributed to the change in price.

While this concern is of first-order importance in model (1), the state panel is less likely to suffer from standard price endogeneity concerns because of the time fixed effects that control for unobserved changes in demand over time. This helps explain the increase in the price elasticity estimate in Table I. Even the panel approach, however, is an imperfect solution because state-level prices may still reflect state-specific differential changes in

Table I. The effect of gasoline prices on gasoline consumption: Least squares estimates

	National aggregate	State panel
Δp_t	-0.10	-0.19
(Standard error)	(0.04)	(0.04)
Month of the year indicators	Yes	Yes
Time fixed effects	No	Yes
R^2	0.71	0.50
Observations	230	11,730

Note: This table reports the price elasticity of gasoline consumption from the national aggregate and state panel specifications. In both specifications the dependent variable is the change in log monthly gasoline consumption. The sample includes monthly observations from January 1989 to March 2008. The state panel includes observations from all states and District of Columbia. The national specification includes month of the year indicators and the state panel includes separate indicators for all 230 months in the sample. In the state panel the standard error is clustered by state.

demand.¹³ In the following subsection we address the problem of price endogeneity more directly using an instrumental variable estimation strategy. Specifically, we are interested in identifying movements in gasoline prices driven by tax changes.

3.2. Instrumental Variables Estimates

In both the national specification and in the state panel we instrument for price using inflation-adjusted changes in the log of the tax per gallon, Δtax_t and Δtax_{it} , $i = 1 \dots 51$. As discussed earlier, we are careful to exclude *ad valorem* taxes in our measure of tax because this component of taxes is functionally related to price. For example, gasoline sales in New York state in January 1995 were subject to an 18.4-cent federal tax and a 22.6-cent state tax. Our measure of the tax is therefore 41 cents, which we express in 2008 dollars. Gasoline sales in New York are also subject to the state's sales tax, but we exclude this *ad valorem* component in calculating our measure of the unit gasoline tax.

Table II reports results the IV estimates of models (1) and (2). In the national specification, the IV estimate of the effect of price on gasoline consumption is -2.38 . This point estimate is dramatically larger than the OLS estimate of -0.10 reported in Table I, though not statistically significant. The credibility of this estimate hinges on the instrument being valid and strong. Since the instrument F -statistic is near zero, we cannot reject the null hypothesis of a zero coefficient on Δtax_t in the first-stage equation, indicating a possible weak-instrument problem. It is perhaps not surprising that gasoline taxes are not a particularly strong instrument at the national level. Our measure of taxes in the national specification is an average of tax rates across states, weighted by the volume of gasoline consumption in each state, and thus includes both federal and state taxes. Although federal taxes clearly influence the national price, tax changes in individual states are typically too small to provide meaningful variation at the national level. The lack of correlation between tax changes and price changes at the national level causes the standard error on Δp_t to be large in magnitude and raises concerns about asymptotic bias in the estimate of α_1 . Accordingly,

¹³ In addition to standard concerns about price endogeneity, oil refineries in some markets may be able to exercise market power and the degree to which they are able to increase prices above marginal costs may depend on local demand. Refining margins may be particularly sensitive to demand conditions in states such as California with requirements for gasoline that exceed national standards. See, for example, Borenstein *et al.* (2004).

Table II. The effect of gasoline prices on gasoline consumption: Instrumental variables estimates

	National aggregate	State panel
Δp_t	-2.38	-1.14
(Standard error)	(6.88)	(0.25)
Month of the year indicators	Yes	Yes
Time fixed effects	No	Yes
Instrument F -statistic	0.11	242.6
(p -value)	(0.74)	(0.00)
Observations	230	11,730

Note: This table reports the price elasticity of gasoline consumption from the national aggregate and state panel instrumental variables specifications. In both specifications the dependent variable is the change in log monthly gasoline consumption. The measure of gasoline taxes that we use as an instrument excludes *ad valorem* taxes. The sample includes monthly observations from January 1989 to March 2008. The state panel includes observations from all states and District of Columbia. The national specification includes month of the year indicators and the state panel includes separate indicators for all 230 months in the sample. In the state panel the standard error and instrument F -statistic are clustered by state.

we have also examined the confidence region for the coefficient on Δp_t based on the conditional likelihood ratio approach developed by Moreira (2003). The confidence set includes essentially all real numbers. We are left with the conclusion that IV estimates provide little information about the parameter of interest in the national data.

In addition to these weak-instrument concerns, there are concerns about the exogeneity of the instrument. For the IV estimator to be consistent, changes in taxes must be uncorrelated with the error term. This assumption cannot be tested, and we must make the case that $\text{cov}(\Delta \text{tax}_t, \varepsilon_t) = 0$ based on economic arguments. At both the federal and the state level, gasoline tax legislation is made by democratically elected legislators, and it stands to reason that policy decisions also reflect current economic conditions. If, for example, gasoline tax increases tend to be implemented during times when gasoline demand is decreasing, this would introduce negative correlation between Δtax_t and ε_t , biasing the IV estimate away from zero. These concerns would be particularly important if gasoline tax increases could be implemented quickly. In practice, however, even though tax legislation may respond to current conditions, the implementation of tax changes typically occurs with a considerable lag.¹⁴ This delay strengthens the case for the validity of tax changes as an instrument.

Although there is little information in national IV estimates, a more promising IV strategy may be based on the additional variation in gasoline tax changes at the state level. In the state panel specification of Table II, the IV estimate is -1.14 and statistically significant at the 1% level. While smaller than the IV estimate from the national specification, this point estimate is still much larger than the panel OLS estimate of -0.19 reported in Table I. The standard error of the elasticity estimate is large. After accounting for sampling variation one cannot rule out at the 95% confidence level the possibility that the true elasticity is only -0.58, for example.

Again it is important to evaluate the strength of the instrument. Unlike in the national specification, the instrument F -statistic in the state panel is large, with a p -value near zero.

¹⁴ For example, a 3-cent gasoline tax was approved in Arkansas on 1 April 1999 and implemented over 3 years beginning 1 July (see 'Gas Tax Rise Among 656 State Laws Effective This Week', *Arkansas Democrat-Gazette*, 29 June 1999). In May 2005, the Governor of Washington signed a bill that increased the gas tax by 9.5 cents per gallon over 4 years beginning 2 months after the bill was signed (see 'Governor Signs Gas Tax Increase, \$8.5 Billion Transportation Package', Associated Press, 10 May 2005).

Thus one normally would not suspect a weak-instrument problem. Nonetheless, we have also examined the confidence region for the coefficient on Δp_{it} based on the conditional likelihood ratio approach developed by Moreira (2003). The implied 95% confidence interval of $[-1.45, -0.86]$ is somewhat narrower than the panel estimate in Table II, but in the same ballpark. It has to be noted, however, that Moreira's approach does not allow for clustering at the state level in accounting for serial correlation, unlike the standard error reported in Table II. Overall, the state panel does not appear to suffer from a weak-instrument problem. This highlights the usefulness of exploiting across-state variation in gasoline taxes in constructing the instrument.

The state panel IV estimate is considerably smaller than the national IV estimate, but still much larger than most previous estimates of the short-run price elasticity of gasoline demand reported in the literature. To understand better the mechanics underlying this result, in Table III we report additional least squares estimates for both the national aggregate and state panel specifications. First, changes in gasoline prices are regressed on changes in gasoline taxes. Second, changes in gasoline consumption are regressed on changes in gasoline taxes. Because the IV estimator is exactly identified, the IV estimates reported in Table II are a simple ratio of the estimates presented in Table III. In the state panel the elasticity of gasoline prices with respect to gasoline taxes is 0.28. Evaluated at the average value of the after-tax gasoline price of \$1.81 and the average value of the gasoline tax of \$0.45 per gallon, this implies a near one-for-one relationship between taxes and prices, consistent with evidence from previous studies that have examined the incidence of gasoline taxes.¹⁵ In the state panel the elasticity of gasoline consumption with respect to gasoline taxes is -0.32 . Taken literally, this estimate implies that a doubling of the gasoline tax from \$0.45

Table III. Least squares estimates of the effect of gasoline taxes on gasoline prices and gasoline consumption

	National aggregate	State panel
<i>Dependent variable: changes in gasoline prices (Δp_t)</i>		
Δtax_t	0.07	0.28
(Standard error)	(0.20)	(0.02)
<i>F</i> -statistic Δtax_t	0.11	242.6
(<i>p</i> -value)	(0.74)	(0.00)
R^2	0.19	0.81
<i>Dependent variable: changes in gasoline consumption (Δy_t)</i>		
Δtax_t	-0.16	-0.32
(standard error)	(0.11)	(0.07)
<i>F</i> -statistic Δtax_t	2.13	19.26
(<i>p</i> -value)	(0.15)	(0.00)
R^2	0.70	0.47

Note: This table reports least squares estimates of the effect of changes in gasoline taxes on changes in gasoline prices and changes in gasoline consumption. Both national aggregate and state panel results are presented. Our measure of gasoline taxes excludes *ad valorem* taxes. Because the IV estimator is exactly identified, the IV estimates reported in Table II are a simple ratio of the estimates presented here. The sample includes monthly observations from January 1989 to March 2008. The state panel includes observations from all states and District of Columbia. The national specification includes month of the year indicators and the state panel includes separate indicators for all 230 months in the sample. In the state panel standard errors and *F*-statistics are clustered by state.

¹⁵ Using evidence from gasoline sales tax moratorium in Illinois and Indiana, Doyle and Samphantharak (2008) find that 70% or more of state tax is borne by buyers. Studies that have used aggregate state-level data have also found near full

per gallon to \$0.90 per gallon would reduce gasoline consumption by 32%.¹⁶ This appears too large to be credible. There are several reasons to be skeptical of this result. In addition to the concerns discussed earlier about the possible endogeneity of the timing of gasoline tax changes, one might be concerned about variation in real gasoline taxes driven by changes in the price index. Identification in the IV model comes from changes in taxes over time. If consumers are responding to announcements of nominal tax changes rather than actual changes in real taxes, the IV estimate may be biased. This distinction seems plausible, as gradual real changes are difficult to detect for consumers. This reasoning suggests that we restrict attention to months in which discrete changes in nominal gasoline taxes occurred at the state level.

Table IV reports IV estimates of the price elasticity of gasoline consumption from six alternative specifications. Although the first-stage F -statistic declines with the sample size, none of the six specifications appears to suffer from a weak-instrument problem. In column 1, the sample is restricted to include states and months in which increases or decreases in nominal state gasoline taxes occurred. Note that out of 11,730 total state by month observations, only 268 were strict changes. In that specification $\hat{\beta}_1$ drops to -0.74 . It may seem that larger tax changes would induce disproportionately larger responses, but further restricting the sample in columns 2 and 3 to include only nominal changes of at least 1 and 2 cents has little impact on the estimate. One might also

Table IV. The effect of gasoline prices on gasoline consumption. State panel instrumental variables estimates: Nominal tax changes only

	Nominal increases and decreases			Nominal increases only		
	All (1)	Greater than 1 cent (2)	Greater than 2 cents (3)	All (4)	Greater than 1 cent (5)	Greater than 2 cents (6)
Δp_t	-0.74	-0.70	-0.70	-0.46	-0.42	-0.28
(Standard error)	(0.38)	(0.45)	(0.40)	(0.23)	(0.23)	(0.38)
Instrument F -statistic	73.5	68.3	58.3	30.6	27.7	12.8
(p -value)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	268	162	100	210	140	86

Note: This table reports the price elasticity of gasoline consumption from six alternative state panel instrumental variables specifications. In all specifications the dependent variable is the change in log monthly gasoline consumption. The measure of gasoline taxes that we use as an instrument excludes *ad valorem* taxes. Since there are few overlapping tax increases it is not practical to use time fixed effects in these specifications. Column (1) includes state by month observations during the period January 1989 to March 2008 in which there was a strict change in nominal state gasoline taxes. The specification in column 4 restricts the sample to include only strict increases. The specifications in columns 2 and 3 restrict the sample to include only nominal changes of at least 1 and 2 cents, respectively. The specifications in columns 5 and 6 restrict the sample to include only nominal increases of at least 1 and 2 cents, respectively. Standard errors and instrument F -statistics are clustered by state.

shifting of state gasoline taxes. See, for example, Chouinard and Perloff (2004) and Alm *et al.* (2009). Our elasticity implies that $(0.28) \times \frac{1.81}{0.45} \times 100 = 113\%$ of the tax is shifted to prices with a standard error of 7%.

¹⁶ Between January 2004 and March 2008 average after-tax gasoline prices increased from \$1.73 to \$3.21 (see Figure 1). This is a 62% price increase. Abstracting from sampling uncertainty, a mechanical interpretation of the estimated price elasticities in Table I might suggest that this price increase would be associated with between a 6% and 12% drop in gasoline consumption. It would be misleading, however, to draw the conclusion that based on the IV estimate of the same elasticities in Table II, this price increase should have been associated with a $(-1.14)(62\%) = 71\%$ predicted drop in gasoline consumption. The reason is that the IV estimate identifies only those movements in gasoline prices that are driven by tax increases, whereas the observed increase in price was not associated with any increase in taxes.

expect that increases and decreases would have different impacts. Since the closest analogy to the gasoline tax increases recently contemplated by policymakers would be a nominal tax increase, column 4 restricts the sample to include only nominal increases. This reduces the sample size to 210 and causes the elasticity to fall to -0.46 . That coefficient is statistically significant at the 5% level.¹⁷

We view this estimate as the most credible estimate of the likely effect of proposed gasoline tax increases because restricting the sample to include only nominal increases is closest to the thought experiment of raising gasoline taxes as proposed in the carbon tax debate. In addition, it is worth emphasizing that this estimate is similar in magnitude to recent estimates by Bento *et al.* (2009). Using a multimarket model of the new, used, and scrap vehicle markets, they simulate a discrete–continuous choice model that allows them to predict not only the responsiveness of gasoline consumption to changes in gasoline taxes, but also to simulate changes in fleet composition. Although their approach is very different, their baseline estimate of -0.35 for the price elasticity of gasoline conditional on fleet composition is similar to our estimate of -0.46 . Our estimate is somewhat larger. One possible reason is that it also reflects short-run substitution across vehicles (e.g. within multiple-vehicle households).

Finally, in the final two columns of Table IV we investigate whether the elasticity increases disproportionately in the magnitude of the tax increase. In columns 5 and 6 the sample is restricted further to include only increases of at least 1 cent and at least 2 cents, respectively. As we restrict the sample to large increases, the point estimates of the price elasticity actually decline, while the standard errors increase. There is no evidence of disproportionate adjustment.

Alternative Specifications

While the preferred parameter estimate in Table IV is considerably smaller than the state panel IV parameter estimate reported in Table II, the fact that the elasticity estimate is still large relative to the panel OLS estimate in Table I calls for additional robustness checks. Table V presents estimates from a number of alternative IV specifications based on the nominal tax increase specification in column 4 of Table IV. One possible concern is the presence of outliers. Column 2 in Table V shows that allowing for outliers does not alter the basic results. The median IV estimator $\hat{\beta}_1$ is almost identical to the baseline IV estimate.

Another possible concern is the possible endogeneity of changes in state gasoline taxes with respect to omitted variables. This concern did not arise in the national specification and may account for the difference between the panel IV estimates and the panel OLS estimate. It is instructive to review the typical rationales for state gasoline tax increases provided by policymakers. There are two main justifications that are used time and again. First, gasoline tax increases are often implemented in order to finance investments in states' transportation infrastructure such as highways and bridges. Such tax increases are plausibly exogenous, though one might be concerned that infrastructure improvements take place primarily during economic booms. Second, gasoline tax increases are often implemented in response to state budget deficits which are likely to be correlated with macroeconomic conditions at the state level, such as a state's unemployment rate.¹⁸ If state legislators are responding to unobserved changes in factors that are correlated with state-level demand for gasoline, this will violate the key identifying assumption for IV. In order to

¹⁷ If, in addition, we include month-of-the-year fixed effects this estimate remains virtually unchanged with a point estimate of -0.45 and a standard error of 0.10.

¹⁸ Specifically, high levels of employment and economic activity are associated with high levels of state revenue and low welfare expenditures. High levels of employment and economic activity also may be associated with high gasoline consumption. Thus failing to control for employment may bias estimates of the price elasticity away from zero.

Table V. The effect of gasoline prices on gasoline consumption, alternative specifications. State panel instrumental variables estimates: Nominal increases only

	Baseline estimate (1)	Median IV regression (2)	Including unemployment rate (3)	Including population density (4)	All state-level controls (5)
Δp_t (Standard error)	-0.46 (0.23)	-0.47 (0.41)	-0.49 (0.24)	-0.43 (0.24)	-0.45 (0.24)

Note: This table reports the price elasticity of gasoline consumption from the state panel instrumental variables specification with nominal increases only and four alternative specifications. In all specifications the dependent variable is the change in log monthly gasoline consumption. The measure of gasoline taxes that we use as an instrument excludes *ad valorem* taxes. The sample includes monthly observations from January 1989 to March 2008. The baseline specification in column 1 is identical to the specification in column 4 of Table IV. The specification in column 2 reports results from median IV regression. The specification in column 3 includes changes in state-level unemployment rates. The specification in column 4 includes the percentage change in state population density. Finally, the specification in column 5 allows for both changes in the unemployment rate and the population density. Standard errors in columns 1 and 3–5 are clustered by state. Standard errors in column 2 are block bootstrap by state with 100 replications.

assess the validity of this concern, Table V presents results from three additional IV specifications. The specification in column 3 controls for changes in state-level unemployment rates.¹⁹ Another possible confounding factor is heterogeneity in population growth across states. The specification in column 4 controls for the percentage change in the state population density.²⁰ As a final check, the specification in column 5 allows for both changes in the unemployment rate and in the population density. In all cases the results are similar to the baseline IV estimate. The additional controls in columns 3, 4 and 5 have virtually no effect on $\hat{\beta}_1$.

Table V demonstrates that omission of macroeconomic conditions at the state level does not provide an explanation for the relatively high IV estimates. Yet another potential concern is that the timing of gasoline tax changes at the state level may depend on the overall level of gasoline prices. One reason is that at times state policymakers have chosen to cut gasoline taxes in an effort to assist consumers, when gasoline prices are unusually high. This was the case, for example, when Connecticut temporarily decreased their gasoline tax by 7 cents between July 2000 and January 2001 in response to anticipated higher than usual summer gasoline prices.²¹ Another reason is that policymakers may be reluctant to call for gasoline tax increases when prices are unusually high, but are likely to resume their efforts when gasoline prices have fallen. An obvious concern is that such behavior would induce a negative correlation between changes in gasoline taxes and changes in gasoline prices, causing a downward bias in the estimate of the effect of gasoline tax changes

¹⁹ We also estimated a specification that includes both changes in state-level unemployment rates and lagged changes in state-level unemployment rates. The results were similar.

²⁰ Population density was constructed by the authors using population statistics from US Census Bureau, Population Division, 'Table II of Current Population Reports Series, Intercensal Estimates of the Total Resident Population of States: 1980 to 1990 (P25-1106)'; 'Time Series of Intercensal State Population Estimates: 1990 to 2000 (CO-EST2001-12-00)'; and 'Table I: Annual Estimates of the Resident Population for the United States, Regions, States, and Puerto Rico: 2000 to 2008 (NST-EST2008-01)'. Population was then divided by the total land area for each state from 'Census 2000 Summary File 1'. Population statistics are available in July of each year from 1988 to 2008. We created monthly estimates of population density using linear interpolation by state.

²¹ See, for example, Daniela Altamari and Matthew Daly, 'Rowland Warns of "Gridlock" Without Gas Tax Cut', *Hartford Courant* (1 April 2000) and Christopher Keating, 'The Tax Cuts are Served \$12.3 Billion Budget Deal Reduces Gas Levy, Aids Hospitals', *Hartford Courant* (28 April 2000).

on gasoline price changes. It is unclear how pervasive this phenomenon is in the data, but the fact that the estimates in Table III imply a nearly proportionate relationship between taxes and prices suggests that the resulting bias must be negligible in practice.²²

It would be tempting to relate these estimates to well-known estimates of the price elasticity of gasoline consumption from dynamic models. Relative to that benchmark, these estimates may seem surprisingly large. This analogy is flawed, however, because IV methods are designed for the estimation of structural models. If we interpret regression model (2) as a structural model of the effect of gasoline tax changes on gasoline consumption, then β_1 captures changes in gasoline consumption caused by changes in gasoline prices that are induced by changes in gasoline taxes. In contrast, the same parameter in a reduced-form model to be estimated by OLS corresponds to the expected change in gasoline consumption given a change in gasoline prices whether induced by taxes or not. This distinction matters because gasoline price changes caused by tax changes are much more persistent than gasoline price changes caused by most other factors and may, for example, lead households to replace a fuel-inefficient car when they would not have made that choice in response to a typical price change. In addition, gasoline tax changes are typically accompanied by extensive media coverage which may amplify the impact on gasoline consumption. For these reasons, it does not make sense to interpret our IV estimates as price elasticities and the comparison with conventional price elasticity estimates is misleading.

Policy Implications

When comparing the results from the OLS and IV approaches it is important to consider what the different estimates imply for the effect of a given change in gasoline taxes on gasoline consumption. The percentage reduction in gasoline consumption resulting from a tax increase of τ cents evaluated relative to a base after-tax price of p cents is

$$\hat{\eta}_1 \left(\frac{\tau}{p} \right) 100,$$

where $\hat{\eta}_1$ corresponds to the estimates $\hat{\alpha}_{1,LS}$, $\hat{\beta}_{1,LS}$, and $\hat{\beta}_{1,IV}$. For expository purposes, we consider a 10-cent tax increase because it yields effects that are easily scalable and large enough to be of economic interest, yet within the range of changes for which there is some historical precedent. We evaluate the effect at the volume-weighted mean after-tax price of \$3.21 in March 2008.

Table VI reports the effect on gasoline consumption implied by our elasticity estimates. The OLS estimates from Table I imply that a 10-cent tax increase is associated with less than a 1% decrease in consumption. The IV estimate from column 4 of Table IV implies much larger effects, with a 1.4% decrease in gasoline consumption. As a point of comparison, Bento *et al.* (2009) find that a 25-cent gasoline tax increase is associated with approximately a 5% decrease in gasoline consumption. It is worth highlighting that during the second half of 2008 gasoline prices decreased dramatically, reaching \$1.75 in December 2008. If instead the effects are evaluated at this lower base level, the effect of a 10-cent tax increase is approximately twice as large.

In calculating the effects in Table VI, we are assuming implicitly that the incidence of the tax is entirely on buyers. This is consistent with the state panel estimates in Table III that imply a near

²² In the state panel in Table III the elasticity of gasoline prices with respect to gasoline taxes is 0.28. Evaluated at the average value of the after-tax gasoline price of \$1.81 and the average value of the gasoline tax of \$0.45 per gallon, this implies a near one-for-one relationship between taxes and prices. For example, a 10-cent (22%) tax increase is associated with an 11-cent (6%) price increase.

Table VI. The effect of a 10-cent gasoline tax increase on gasoline consumption: Traditional regression estimates

	(%)
OLS estimate, national aggregate	-0.31
(Standard error)	(0.11)
OLS estimate, state panel	-0.59
(Standard error)	(0.12)
IV Estimate, state panel	-1.43
(Standard error)	(0.72)

Note: The effect of a 10-cent gasoline tax is evaluated at the volume-weighted mean after-tax price of \$3.21 in March 2008. The implied effects for OLS are based on the -0.10 and -0.19 elasticities in Table I. The implied effect for state panel IV is based on the -0.46 estimate in column 4 of Table IV.

one-to-one relationship between tax increases and price increases. However, even if the supply of gasoline for individual states is perfectly elastic, one might be concerned that the supply curve for the entire United States is not. Indeed, although the national aggregate estimate in Table III does not rule out a one-to-one relationship between tax increases and price increases at conventional significance levels, the point estimate implies less than full shifting. To the extent that supply is less than fully elastic, part of the tax would be borne by the sellers. This would tend to lower the reduction in gasoline consumption associated with such a tax.

4. VECTOR AUTOREGRESSIVE MODELS

4.1. Price Specification

It is widely recognized that gasoline prices are endogenous with respect to macroeconomic aggregates that drive the consumption of gasoline (see, for example, Kilian, 2008). An alternative to using instruments in identifying exogenous movements in gasoline prices is to focus on unpredictable changes (or innovations) in gasoline prices. To the extent that these innovations are predetermined with respect to the economic aggregates that drive gasoline consumption, the response of gasoline consumption to such an innovation can be estimated. This alternative approach is particularly appealing when working with monthly data, as we do in this paper, but less suitable for data measured at lower frequency. For example, at the annual frequency one would expect feedback from unexpected changes in gasoline consumption to the price of gasoline violating the assumption of predetermined gasoline prices.

Let Δp_t denote the percentage change in inflation-adjusted after-tax prices of gasoline and Δy_t the percentage change in real gasoline consumption. Consider the stylized structural model

$$\begin{aligned} \Delta p_t &= \theta \Delta y_t + \varepsilon_{1t} \\ \Delta y_t &= \delta \Delta p_t + \varepsilon_{2t} \end{aligned} \Leftrightarrow \underbrace{\begin{bmatrix} 1 & -\theta \\ -\delta & 1 \end{bmatrix}}_{A_0} \underbrace{\begin{pmatrix} \Delta p_t \\ \Delta y_t \end{pmatrix}}_{x_t} = \underbrace{\begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix}}_{\varepsilon_t}$$

where ε_{1t} and ε_{2t} are mutually uncorrelated structural errors. All deterministic terms and the dependence of both variables on lagged observations have been suppressed. Predeterminedness in this model means that $\theta = 0$ (see Cooley and LeRoy, 1985). In other words, we rule out

contemporaneous feedback from Δy_t to Δp_t . In conjunction with the assumption that there are no other exogenous innovations outside of the model that are correlated with the innovation in the real price of gasoline, the coefficient δ may be interpreted as the contemporaneous causal effect of the price innovation on real gasoline consumption.

The same logic applies in the more general VAR(p) model that allows for additional unrestricted delayed feedback between real gasoline consumption growth and real gasoline price changes up to some lag order p . Let $x_t \equiv (\Delta p_t \ \Delta y_t)'$ and $\varepsilon_t \equiv (\varepsilon_{1t} \ \varepsilon_{2t})'$. Then the structural representation of the VAR(p) model is

$$A_0 x_t = \sum_{i=1}^p A_i x_{t-i} + \varepsilon_t.$$

In the benchmark model, we set the lag order p to 12. We also report alternative VAR results using estimated lag orders obtained by formal model selection using tools such as the Akaike Information Criterion (AIC), as suggested by Ivanov and Kilian (2005). Predeterminedness in this VAR model implies an exclusion restriction on the upper right element of A_0 . The reduced-form representation of this model is

$$x_t = \sum_{i=1}^p B_i x_{t-i} + e_t,$$

where $B_i = A_0^{-1} A_i$, $i = 1, \dots, p$, and $e_t = A_0^{-1} \varepsilon_t$. The reduced form may be estimated by the least squares method. Applying a lower triangular Cholesky decomposition to the estimate of the 2×2 variance-covariance matrix $\Sigma_{e_t} = A_0^{-1} A_0^{-1'}$ of the reduced-form VAR errors allows us to estimate A_0^{-1} subject to the identifying restriction of predetermined changes in gasoline prices. Given an estimate of A_0^{-1} , it is straightforward to construct the dynamic responses of gasoline consumption to an exogenous change in the price of gasoline from this model.²³

How reasonable the exactly identifying assumption of predetermined changes in gasoline prices is, by construction, cannot be tested within the context of this VAR model. Rather, the credibility of this assumption rests on an economic argument. Following Kilian (2010), we postulate that gas distributors charge a fixed markup on the wholesale price of gasoline. We postulate that the price of gasoline reflects without delay gasoline supply shocks arising at the refining stage and changes in the price of imported crude oil representing costs shocks for US refiners. The model stipulates that the retail supply curve for gasoline is perfectly elastic in the short run. The presumption is that gas distributors have enough gasoline to supply the required quantities of gasoline at the current retail price in response to an unexpected increase in demand. Given that the consumption of gasoline typically evolves smoothly and predictably, we abstract from the possibility that gas distributors may run out of gas at the posted price. While unanticipated shifts in gasoline demand in the model do not move the price of gasoline instantaneously, they are allowed to affect the price of gasoline with a delay of 1 month. The rationale for this assumption is the difficulty of distinguishing a temporary blip in gasoline demand from a change in trend. Only if a change in demand is sustained and hence is expected to persist will a gas distributor respond by raising the price of gasoline. Recognizing such persistent shifts in expected demand by construction requires

²³ Essentially the same identification strategy has been used, for example, by Edelstein and Kilian (2009) and Hamilton (2009). Both papers focused on a broader measure of retail energy prices. Similar identifying assumptions are used routinely in the literature on estimating the dynamic responses to oil price shocks (see, for example, Rotemberg and Woodford, 1996; Blanchard and Galí, 2009).

time, justifying the delayed price response. Under these assumptions, the model can be used to estimate the dynamic response of US gasoline consumption to an unanticipated 1% gasoline price increase, as discussed above.

It may seem that the VAR model should include in addition any number of macroeconomic aggregates since shifts in US gasoline consumption are driven in part by US macroeconomic events. The inclusion of such additional variables is not required for the identification of the gasoline price shock, however, as long as changes in gasoline prices are predetermined with respect to all omitted variables. To the extent that US gasoline prices are predetermined with respect to these macroeconomic aggregates, there is no loss of generality in restricting the analysis to a bivariate model in prices and quantities, since we can view the bivariate model as a marginalized representation in which all additional variables other than the percentage change in gasoline consumption have been integrated out. Kilian and Vega (2009), on the basis of daily data, have recently shown that there is no evidence of feedback from exogenous US macroeconomic news to the US retail price of gasoline (or for that matter the price of crude oil) within 1 month of such shocks, suggesting that the assumption that US gasoline prices are predetermined with respect to US macroeconomic aggregates is a good approximation.

Our VAR analysis focuses on aggregate data at the national level.²⁴ All models include seasonal dummies. First, consider the bivariate VAR(12) model. The upper panel of Figure 6 shows the time path of the response of US gasoline consumption to a 1% increase in volume-weighted gasoline prices. The impulse response coefficients have been normalized such that the coefficient estimates can be interpreted as elasticities. One-standard-error bands have been computed based on the recursive-design wild bootstrap that allows for conditional heteroskedasticity of unknown form (see Goncalves and Kilian, 2004). The use of a lower significance level makes sense in the context of unrestricted VAR models since these models involve many more parameters than the models underlying Tables I and II. The elasticity estimate implied by the VAR is small and statistically insignificant on impact, but at longer horizons the price elasticity reaches near -0.12 . This 1-year price elasticity is similar in magnitude to the least squares estimate of the impact elasticity in column 1 of Table I, but that comparison is misleading since the VAR impact elasticity estimate of -0.02 is much smaller than the estimate of -0.10 reported in Table I.

An important question is to what extent our estimates would change if the assumption of a flat short-run gasoline supply curve were not literally correct. To address this issue we experimented with VAR models that impose a weak degree of feedback from demand shifts to the price of gasoline. In practice, we impose a value for θ that is greater than zero and solve numerically for the remaining unrestricted parameter of A_0 conditional on that identifying assumption (for a similar approach see Abraham and Haltiwanger, 1995). We investigated $\theta \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$. The choice of θ has implications for the supply curve. If θ is equal to zero, as assumed in our baseline model, the supply curve is flat. For $\theta > 0$ the supply curve is upward sloping. Relative to the benchmark of no feedback, the 1-year price elasticity is remarkably robust to these alternative assumptions. Table VIII shows that the estimates at the 12-month horizon increase in absolute terms only slightly from -0.12 for $\theta = 0$ to -0.14 for $\theta = 0.5$. Thus the additional evidence in Table VIII strengthens the case for our original identifying assumption. There is a more marked increase in absolute terms in the impact price elasticity from -0.02 to -0.09 , as θ is increased,

²⁴ We have also examined elasticities separately by state. Average elasticities, weighted by volume, are very similar to the aggregate results.




Figure 6. (a) Response of US aggregate gasoline consumption to a positive gasoline price shock with one-standard-error bands, VAR(12) model. (b) Response of US aggregate gasoline consumption to a positive tax shock with one-standard-error bands, VAR(12) model

but that point is inconsequential for the policy question of interest in this paper. We conclude that the 1-year price elasticity estimates are robust to alternative identifying assumptions.

4.2. Tax Specification

Although these initial VAR estimates provide a crude estimate of the response of US gasoline consumption to a tax increase, they ignore the fact that tax increases may cause more persistent price changes than other shocks. We would expect a larger consumption response to a gasoline tax increase than to a gasoline price increase driven, for example, by a refinery outage (see Kilian, 2010). A more direct estimate of the tax elasticity may be obtained by replacing Δp_t in the bivariate VAR(12) model by Δtax_t such that $x_t \equiv (\Delta tax_t \ \Delta y_t)'$. As before, we focus on unit taxes and exclude *ad valorem* taxes, and all taxes have been deflated using the US consumer price index. The model treats tax changes as predetermined with respect to changes in gasoline consumption. Since innovations in gasoline taxes arise from a slow legislative process,




Figure 7. (a) Response of US aggregate gasoline consumption to a positive gasoline price shock with one-standard-error bands, VAR model based on AIC lag order estimate. (b) Response of US aggregate gasoline consumption to a positive tax shock with one-standard-error bands, VAR model based on AIC lag order estimate

the identifying assumption of no contemporaneous feedback from US gasoline consumption (or macroeconomic aggregates correlated with US gasoline consumption) to innovations in gasoline taxes is uncontroversial, providing a particularly clean identification of the causal effects, provided consumers do not anticipate the tax increase. Of course, to the extent that consumers exhibit forward-looking behavior, one may question the appropriateness of a vector autoregressive model for identifying the causal effects in general. In our case, it is not clear in what sense forward-looking behavior would have contaminated the identification. Households have few margins of adjustment through which they can respond to anticipated gasoline tax increases beyond filling up their gas tank one last time prior to the tax change. While it is conceivable that a household might choose to buy a more fuel-efficient car in anticipation of a pending gasoline tax increase, it seems unlikely that this type of behavior would have been quantitatively important in our data. The lower panel of Figure 6 shows the time path of the response of US gasoline consumption to a 1% tax increase. The tax elasticity is much larger on impact than the corresponding price elasticity in the upper panel and statistically significant. While the dynamic response is more

Table VII. Price and tax elasticities estimated from structural VAR models

	Horizon in months					
	0	1	3	6	9	12
Price elasticity (12 lags)	-0.02	-0.10	-0.12	-0.12	-0.09	-0.12
Tax elasticity (12 lags)	-0.14	0.03	-0.06	-0.06	-0.15	-0.13
Price elasticity (AIC)	-0.02	-0.09	-0.10	-0.09	-0.06	-0.07
Tax elasticity (AIC)	-0.14	0.02	-0.07	-0.07	-0.17	-0.15

Note: The sample period is 1989.1–2007.3 and the data frequency is monthly. Boldface indicates statistical significance based on one-standard error bands. All models include seasonal dummies. The AIC lag orders were selected based on an upper bound of 12 lags. The AIC selects nine lags for the model of price elasticity and 11 lags for the model of tax elasticity.

Table VIII. Price elasticities estimated from structural VAR(12) models under alternative assumptions about the strength of feedback from gasoline consumption to the price of gasoline as measured by θ

θ	Horizon in months					
	0	1	3	6	9	12
0	-0.02	-0.10	-0.12	-0.12	-0.09	-0.12
0.1	-0.03	-0.11	-0.12	-0.12	-0.09	-0.12
0.2	-0.05	-0.11	-0.13	-0.12	-0.10	-0.13
0.3	-0.06	-0.11	-0.13	-0.13	-0.10	-0.13
0.4	-0.08	-0.11	-0.14	-0.13	-0.11	-0.14
0.5	-0.09	-0.12	-0.15	-0.13	-0.11	-0.14

Note: The sample period is 1989.1–2007.3 and the data frequency is monthly. Boldface indicates statistical significance based on one-standard-error bands. All models include seasonal dummies.

erratic and the error bands are wider, the 1-year tax elasticity is about -0.13 and is statistically significant.

A natural conjecture is that the elasticity estimates may be further improved upon by letting the data help us in the selection of the lag order. We follow Ivanov and Kilian (2005) in selecting the VAR lag order based on the AIC that trades off the improvement in the fit of the model against the profligacy of parameters. Allowing for up to 12 lags, the AIC selects nine lags for the model of the price elasticity and 11 lags for the model of the tax elasticity. The implied price and tax elasticity on impact are essentially unchanged relative to the model with a fixed number of 12 lags, but the 1-year price elasticity falls to -0.07 , whereas the 1-year tax elasticity rises to -0.15 , suggesting a greater responsiveness to tax changes than to average price changes (see Table VII). Figure 7 shows the time path of the response of US gasoline consumption to a 1% increase in volume-weighted gasoline prices and to a 1% tax increase with the VAR model based on the AIC lag order estimate. Regarding the sensitivity of the estimates to the degree of contemporaneous feedback from gasoline consumption to the price of gasoline for this final specification, the results are qualitatively the same as in the case of the VAR model with a fixed number of 12 lags. Increasing θ to 0.5 increases the 1-year price elasticity to -0.09 , which still falls short of the direct estimate of the 1-year tax elasticity.

4.3. Policy Implications

As in the case of the single-equation models, it is important to translate these elasticity estimates into the effect of a gasoline tax increase of τ cents on gasoline consumption. In calculating these effects it is critical to distinguish between price elasticities and tax elasticities. For example, given a base price of p cents, the 12-month price elasticity of -0.12 implies a change in consumption of

$$-0.12 \left(\frac{\tau}{p} \right) 100,$$

whereas the 12-month tax elasticity of -0.13 implies a much larger effect because it is relative to the base tax level, tax :

$$-0.13 \left(\frac{\tau}{tax} \right) 100.$$

As before, we evaluate these effects for a 10-cent tax increase at the volume-weighted mean after-tax price of \$3.21 in March 2008. For the base tax level we use 38.4 cents, the mean volume-weighted tax level in March 2008.

Table IX reports the effect of a 10-cent gasoline tax increase on gasoline consumption implied by the VAR estimates. The estimates based on the model with Δp_t yield less than a 1% decrease in consumption. These implied changes in gasoline consumption are similar in magnitude to the changes implied by the OLS estimates described in Table VI. The estimates from the model with Δtax_t yield considerably larger effects. The AIC model, for example, implies that a 10-cent tax increase would decrease gasoline consumption by 3.9%. When these effects are evaluated at the lower price levels in place in late 2008, the effects based on tax elasticities are unchanged but the results based on price elasticities are approximately twice as large.

There are two clear patterns in the estimates presented in Tables VI and IX. First, estimates based on methods that attempt to control for endogeneity, namely the state panel IV estimates and the VAR estimates, are distinctly larger than other estimates. Second, the tax VAR estimates are much larger than the preferred IV estimate based on using nominal tax increases. The apparent reason is that the VAR analysis does not restrict the sample to months in which discrete changes in nominal gasoline taxes occurred. In fact, the tax VAR estimates are very similar to the IV

Table IX. The Effect of a 10-cent gasoline tax increase on gasoline consumption at 12-month horizon according to the structural VAR estimates in Table VII

	(%)
Based on price elasticity, VAR(12)	-0.37
(Standard error)	(0.07)
Based on tax elasticity, VAR(12)	-3.39
(Standard error)	(3.17)
Based on price elasticity, VAR(AIC)	-0.22
(Standard error)	(0.13)
Based on tax elasticity, VAR(AIC)	-3.91
(Standard error)	(2.91)

Note: The effect of a 10-cent gasoline tax is evaluated at the volume-weighted mean after-tax price of \$3.21 and mean tax of 38.4 cents in March 2008.

estimate implied by column 2 of Table II based on the unrestricted sample. In the latter case we would have obtained a reduction in gasoline consumption of -3.91 compared with a reduction of between -3.39 and -3.91 (depending on the lag specification) for the tax VAR. These results underscore the importance of restricting attention to dates of nominal tax increases in evaluating policy proposals for gasoline tax increases. Although the VAR approach provides an alternative to identifying the effects of tax changes, it is not clear how to incorporate such nonlinearities into the VAR framework (see, for example, Kilian and Vigfusson, 2009). This suggests caution in interpreting the tax VAR estimates. For the remainder of the paper we therefore focus on our preferred IV elasticity estimate of -0.46 .

5. THE EFFECT OF A GASOLINE TAX ON CARBON EMISSIONS

Estimating the effect of raising gas taxes on gasoline consumption, while interesting for other purposes, is only a first step in computing the effect of such a policy on carbon emissions. The percentage change in total carbon dioxide emissions in the United States is calculated by multiplying the gasoline consumption effect by 0.338, the fraction of carbon dioxide emissions in the United States derived from the transportation sector:²⁵

$$(-1.43\%)0.338 = -0.48\%$$

Here -1.43% is the change in gasoline consumption from a 10-cent gasoline tax increase based on the estimate from Table VI. A 10-cent gasoline tax increase reduces carbon emissions in the United States by 0.48% (see row 2 of Table X).²⁶ Although not negligible, this is small when compared to, for example, recent annual increases in carbon emissions. Sampling variation implies that the true effect could be much smaller. The effect could also be much larger, though it is unlikely to exceed a few percentage points. For example, one can rule out carbon emission decreases in excess of 1.0% with 95% confidence.

Carbon emissions are a global problem. In order to put these results into perspective, the third and fourth rows of Table X report the implied change in OECD and world emissions. These percentage changes are calculated by multiplying the percentage effect in the United States (shown in row 2) by 0.45 and 0.21, respectively, the fraction of OECD and world emissions represented by the United States.²⁷ A 10-cent gasoline tax in the United States reduces OECD and world emissions by considerably less than 1%. This reflects the facts that vehicles in the United States represent a small and decreasing fraction of total OECD and global carbon dioxide emissions. Carbon dioxide emissions in the United States are growing less quickly than emissions in other countries, most notably China, India and Brazil. By 2030, according to predictions from the US

²⁵ US Department of Energy, *Energy Outlook 2008*, DOE/EIA-0383, Table A18, 'Carbon Dioxide Emissions by Sector and Source, 2006 and 2030'.

²⁶ This estimate presumes that substitution away from gasoline does not raise greenhouse gas emissions. To the extent that consumers substitute from gasoline to other carbon-producing goods, the aggregate reductions will be overestimated by our approach.

²⁷ US Department of Energy, Energy Information Administration, *International Energy Outlook 2008*, DOE/EIA-0484, Table A10. 'World Carbon Dioxide Emissions by Region, Reference Case, 1990-2030'.

Table X. The effect of a 10-cent U.S. gasoline tax increase on carbon emissions

	(%)
Gasoline consumption in the United States (Standard error)	-1.43 (0.72)
Total carbon dioxide emissions in the United States (Standard error)	-0.48 (0.24)
Total carbon dioxide emissions in the OECD (Standard error)	-0.22 (0.11)
Total carbon dioxide emissions worldwide (Standard error)	-0.10 (0.05)

Note: These results are based on the -0.46 price elasticity from column 4 of Table IV and the volume-weighted mean after-tax price of \$3.21 in March 2008. We assume that 33.8% of carbon dioxide emissions in the United States are derived from the transportation sector following US Department of Energy, *Energy Outlook 2008*, DOE/EIA-0383, Table A18, 'Carbon Dioxide Emissions by Sector and Source, 2006 and 2030'. Moreover, we assume that the United States represents 44.7% of total OECD carbon dioxide emissions and 21.0% of total world carbon dioxide emissions following predictions for 2010 from US Department of Energy, *International Energy Outlook 2008*, DOE/EIA-0484, Table A10, 'World Carbon Dioxide Emissions by Region, Reference Case, 1990–2030'.

Department of Energy, the United States will represent only 16.1% of total world carbon dioxide emissions.²⁸

It is worth re-emphasizing that these estimates capture only the short-run response. In the short run, drivers can adjust discretionary driving patterns, drive slower, for example, or improve fuel efficiency by increasing tire pressure. The long-run price elasticity is likely to be larger as agents may employ additional margins of adjustment. For example, one would expect to see widespread substitution toward more fuel-efficient vehicles, some of which may not even be available when the tax is first implemented. Because the vehicle stock turns over slowly, the full impact of a tax change will not be realized for many years. Likewise, households may choose to relocate closer to their workplace in an effort to cut down on commuting or they may demand improved public transportation. Predicting such long-run effects is beyond the scope of any econometric model based on historical data.

These results highlight three mechanisms by which the global reduction in carbon emissions might be amplified. First, widespread international adoption of gasoline tax increases may have considerably larger effects. Our analysis is focused on the United States and differences in income levels and other factors suggest that the responsiveness of gasoline consumption to changes in gasoline taxes is likely to differ across countries. Moreover, many European countries already impose high gasoline taxes. Nonetheless, if one takes our elasticity estimate as representative for the world, a global 10-cent gasoline tax increase would decrease global emissions by 0.48%. This calculation, of course, abstracts from the general equilibrium effects of such a concerted action.

Second, a broad-based carbon tax on all sources of carbon dioxide rather than just gasoline would reduce carbon emissions further. It is important to keep in mind, however, that emissions from transportation represent a substantial fraction of total emissions and that there is no reason to expect demand in other sectors to be substantially more elastic. Consider, for example, the

²⁸ US Department of Energy, *International Energy Outlook 2008*, DOE/EIA-0484, Table A10. 'World Carbon Dioxide Emissions by Region, Reference Case, 1990–2030'. Between 2005 and 2030, total energy-related carbon dioxide emissions are forecast to increase by 83% in non-OECD countries compared to only 14% in OECD countries.

electric power sector which is responsible for 40% of total carbon dioxide emissions in the United States. Although there is scope for substitution toward low-carbon fuels within the power sector, much of the existing generating capacity is built around coal and natural gas and the transition to renewable technologies is likely to be a slow process spanning decades.

Third, one could induce larger emission decreases with larger gasoline tax increases. Indeed, some of the policies currently being proposed would increase the gasoline tax by as much as \$1.00 per gallon or more.²⁹ Our results should be applied with extreme caution when extrapolating to such larger tax changes. Our estimates are based on examining the responsiveness of gasoline consumption to a series of relatively small tax changes implemented during the period 1989–2008. There is no historical precedent for tax increases of more than 10 cents and it is unclear how well linear econometric models will fare when subject to variation so clearly beyond the available support, nor would it be possible to estimate such nonlinear effects from historical data.

6. CONCLUDING REMARKS

Although interest in carbon taxes has quieted down recently as a result of the rapidly deteriorating global economic conditions, that situation is likely to be temporary. As the economy recovers, it will be only a question of time before these issues reemerge. Combating climate change is widely known to be one of the priorities of the Obama administration. The current respite provides an opportunity to reflect impassionately on the merits of a gasoline tax increase, as the leading example of how a carbon tax would be implemented in practice.

In this paper we provided a careful analysis of what can be learned from historical data about the likely effect of higher gasoline taxes on US gasoline consumption. Such an analysis seems warranted before adopting any proposal regarding gasoline taxes. A natural approach to this question is to study the effects of historical variation in gasoline tax rates both over time and across states. We explored a variety of alternative econometric methods designed to account for the endogeneity of gasoline prices as well as the unique nature of the policy experiment in question. Our most credible estimates imply larger effects than suggested by previous estimates based on less appropriate econometric methodologies.

Nevertheless, our estimates imply that a gasoline tax increase of the magnitude currently contemplated by policymakers would have only a modest short-run impact on carbon emissions. For example, a 10-cent increase in gasoline taxes would lower US emissions by about half of 1%. Another way of putting these results in perspective is to observe that this is roughly equal to one-half of the typical annual increase in US carbon emissions.

While our analysis makes full use of the available data, there are two obvious caveats. First, it is conceivable that at longer horizons a gasoline tax would reduce carbon emissions by more. For example, in the long run households could respond to a gasoline tax by moving closer to work or changing jobs to work closer to home, or cities could choose to build mass transit systems. In addition, the introduction of more energy-efficient vehicles, some of which might not even be

²⁹ Gregory N. Mankiw, 'Raise the Gas Tax,' *Wall Street Journal* (20 October 2006) advocates increasing the gas tax by \$1 per gallon, phased in gradually by 10 cents per year over the next decade. Mankiw has formed the 'Pigou Club', a group of 'economists and pundits with the good sense to have publicly advocated higher Pigouvian taxes, such as gasoline taxes or carbon taxes'. Current members include Bill Nordhaus, Martin Feldstein, Gary Becker, and Alan Greenspan. See also Daniel Gross, 'Raise the Gasoline Tax? Funny, It Doesn't Sound Republican', *New York Times* (8 October 2006) and Kenneth Rogoff, 'What America Must Do: Step on the Gas', *Foreign Policy* (January/February 2008).

available at the time when the policy is enacted, could amplify the long-run effect of a gasoline tax increase. The extent to which the fuel efficiency of cars would be increased, however, and the extent to which the introduction of electric cars, for example, would reduce carbon emissions, is an open question because of the carbon emissions from the generation of electricity. Predicting such long-run effects is beyond the scope of standard econometric models based on historical data.

Second, one possible response to our estimates of likely carbon emission reductions would be to raise the gasoline tax even further to generate larger effects on carbon emissions. Indeed some pundits have recommended raising gasoline taxes cumulatively by as much as 1 dollar per gallon. Still, it has to be borne in mind that the resulting reduction in carbon emissions may not grow proportionately. Once we consider gasoline tax increases far larger than any gasoline tax increase in US history, there is reason to doubt the accuracy of predictions generated from linear econometric models; nor is it possible to estimate such nonlinear effects from historical data.

ACKNOWLEDGEMENTS

Comments from Soren Anderson, Severin Borenstein, Ryan Kellogg and two anonymous referees substantially improved the paper.

REFERENCES

- Abraham KG, Haltiwanger JC. 1995. Real wages and the business cycle. *Journal of Economic Literature* **33**: 1215–1264.
- Alm J, Sennoga E, Skidmore M. 2009. Perfect competition, urbanization, and tax incidence in the retail gasoline market. *Economic Inquiry* **47**: 118–134.
- Bento AM, Goulder LH, Jacobsen MR, von Haefen RH. 2009. Distributional and efficiency impacts of increased U.S. gasoline taxes. *American Economic Review* **99**: 667–699.
- Blanchard OJ, Galí J. 2009. *The Macroeconomic Effects of Oil Shocks: Why are the 2000s so different from the 1970s?*. Forthcoming: J. Galí and M. Gertler (eds) *International Dimensions of Monetary Policy*, University of Chicago Press: Chicago, IL.
- Borenstein S, Bushnell J, Lewis M. 2004. *Market power in California's gasoline market*. CSEM Working Paper 132.
- Chouinard H, Perloff JM. 2004. Incidence of federal and state gasoline taxes. *Economics Letters* **83**: 55–60.
- Cooley TF, LeRoy S. 1985. Atheoretical macroeconometrics: a critique. *Journal of Monetary Economics* **16**: 283–308.
- Dahl C. 1979. Consumer adjustment to a gasoline tax. *Review of Economics and Statistics* **61**: 427–432.
- Dahl C, Sterner T. 1991. Analysing gasoline demand elasticities: a survey. *Energy Economics* **13**: 203–210.
- Doyle J Jr, Samphantharak K. 2008. \$2.00 Gas! Studying the effects of a gas tax moratorium. *Journal of Public Economics* **92**: 869–884.
- Edelstein P, Kilian L. 2009. How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics* **56**: 766–779.
- Fullerton D, West SE. 2002. Can taxes on cars and on gasoline mimic an unavailable tax on emissions? *Journal of Environmental Economics and Management* **43**: 135–157.
- Goncalves S, Kilian L. 2004. Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. *Journal of Econometrics* **123**: 89–120.
- Greene DL, Kahn JR, Gibson RC. 1999. Fuel economy rebound effect for U.S. household vehicles. *The Energy Journal* **20**: 1–30.
- Hamilton JD. 2009. Causes and consequences of the oil shock of 2007–2008. *Brookings Papers on Economic Activity*. 1: 215–261.
- Hughes JE, Knittel CR, Sperling D. 2008. Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal* **29**: 93–114.

- Ivanov V, Kilian L. 2005. A practitioner's guide to lag order selection for VAR impulse response analysis. *Studies in Nonlinear Dynamics and Econometrics* **9**: March 2005, Article 2.
- Kilian L. 2008. The economic effects of energy price shocks. *Journal of Economic Literature* **46**: 871–909.
- 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*. Forthcoming: *Energy Journal* 31(2): 105–130.
- Kilian L, Vega C. 2009. *Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices*. Forthcoming: *Review of Economics and Statistics*.
- Kilian L, Vigfusson R. 2009. *Pitfalls in estimating asymmetric effects of energy price shocks*. Mimeo, Department of Economics, University of Michigan.
- Moreira M. 2003. A conditional likelihood ratio test for structural models. *Econometrica* **71**: 1027–1048.
- Nordhaus W. 2007. A review of the Stern review on the economics of global warming. *Journal of Economic Literature* **45**: 686–702.
- Ramsey J, Rasche R, Allen B. 1975. An analysis of the private and commercial demand for gasoline. *Review of Economics and Statistics* **57**: 502–507.
- Rotemberg J, Woodford M. 1996. Imperfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit, and Banking* **28**(Part 1): 550–577.
- Schmalensee R, Stoker TM. 1999. Household gasoline demand in the United States. *Econometrica* **67**: 645–662.
- Stock JH, Wright JW, Yogo M. 2002. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics* **20**: 518–529.