



Federal Reserve
Bank of Dallas

Heterogeneity in the Pass-Through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand

Lutz Kilian and Xiaoqing Zhou

Working Paper 2301

January 2023

Research Department

<https://doi.org/10.24149/wp2301>

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Heterogeneity in the Pass-Through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand*

Lutz Kilian[†] and Xiaoqing Zhou[‡]

December 24, 2022

Abstract

We propose a new instrument for estimating the price elasticity of gasoline demand that exploits systematic differences across U.S. states in the pass-through of oil price shocks to retail gasoline prices. We show that these differences are primarily driven by the cost of producing and distributing gasoline, which varies with states' access to oil and gasoline transportation infrastructure, refinery technology and environmental regulations, creating cross-sectional gasoline price shocks in response to an aggregate oil price shock. Time-varying estimates do not support the view that the gasoline demand elasticity has declined in absolute value to near zero since the 1980s. The elasticity was stable near -0.3 until the end of 2014. It rose to about -0.2 in 2015-16, but has remained stable since 2016. Gasoline demand is more responsive in states with lower personal income, higher unemployment rates and lower urban population shares. There is no evidence for an asymmetry in the elasticity with respect to positive and negative gasoline price shocks. We illustrate how these elasticity estimates inform the recent policy debate about the impact of gasoline tax holidays on consumers' discretionary income, about the demand destruction from the spike in gasoline prices after the invasion of Ukraine and about the impact of rising gasoline prices on carbon emissions.

Keywords: Price elasticity of gasoline demand, pass-through, gasoline tax, gasoline supply, identification, IV, cross-section.

JEL Codes: D12, Q41.

* We thank John Coglianesse, Catie Hausman and David Rapson for helpful comments and discussions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

[†]Lutz Kilian, Federal Reserve Bank of Dallas, 2200 N. Pearl St., Dallas, TX 75201, lkilian2019@gmail.com.

[‡]Xiaoqing Zhou, Federal Reserve Bank of Dallas, 2200 N. Pearl St., Dallas, TX 75201, xqzhou3@gmail.com.

1 Introduction

The price elasticity of gasoline demand has received substantial policy and research interest in the past decade (e.g., [Allcott and Wozny \(2014\)](#); [Bento et al. \(2009\)](#); [Davis and Kilian \(2011\)](#); [Doyle and Samphantharak \(2008\)](#); [Edelstein and Kilian \(2009\)](#); [Hughes et al. \(2008\)](#); [Levin et al. \(2017\)](#); [Knittel and Tanaka \(2021\)](#); [Li et al. \(2014\)](#)).¹ Reliable estimates of this elasticity are important for understanding the response of fuel consumption and carbon emissions to fluctuating gasoline prices, for modeling automobile demand, for designing regulatory policies that aim to correct negative externalities from vehicle use, and, from a macroeconomic perspective, for evaluating the impact of federal or state gasoline tax holidays, a tool often considered by policy makers to raise household discretionary income. Moreover, this elasticity helps measure the demand destruction caused by unexpectedly surging fuel prices and the macroeconomic impact of gasoline price shocks.

The identification of this elasticity in empirical work requires the researcher to control for unobserved factors that affect both gasoline prices and consumption. Traditionally, the literature has relied on aggregate time series data with commonly cited estimates between -0.08 and -0.03 (e.g., [Hughes et al. \(2008\)](#)). At this level of aggregation, shifts in U.S. demand are likely to cause gasoline prices and consumption to move in the same direction, biasing the elasticity estimate toward zero. Hence, a growing literature has turned to disaggregate data at the state, city or individual level to exploit cross-sectional variation. While the identification issue is widely acknowledged, nearly all of these studies employ OLS estimation due to the difficulty of finding suitable instruments.²

One exception is the work of [Coglianese et al. \(2017\)](#), who rely on excise gasoline taxes as the instrument, taking advantage of the fact that state gasoline tax changes are highly correlated with retail gasoline price changes, but are predetermined due to legislative and implementation lags. Their instrumental-variable (IV) estimate is -0.37, indicating that demand is far more responsive than suggested by traditional estimates, but this IV estimate

¹Early contributions to this literature include [Dahl \(1979\)](#), [Dahl and Sterner \(1991\)](#), [Hausman and Newey \(1995\)](#), [Ramsey et al. \(1975\)](#) and [Sweeney \(1984\)](#), among others.

²Some cross-sectional studies instrument retail gasoline prices with global oil prices (or spot gasoline prices) and obtain elasticity estimates similar to or even more attenuated than the OLS estimates. This approach, as shown in Appendix A, solely relies on time series variation in the cross-sectional setting and is subject to the same endogeneity concern as using the aggregate time series data.

is not statistically significant. This result raises a number of important questions. First, with most fluctuations in monthly retail gasoline prices driven by the pre-tax price rather than gasoline taxes, is it possible to develop instruments that exploit exogenous variation in the pre-tax price? Second, how do these elasticity estimates compare to tax-based estimates? Third, are these estimates more precise than the tax-based IV estimates? Fourth, does the gasoline demand elasticity vary with demographic and economic conditions and with positive and negative gasoline price shocks?

We address these questions based on a new instrument that exploits cross-sectional variation in the systematic pass-through from global oil price shocks to U.S. state-level retail gasoline prices. Historically, when the oil price experiences large fluctuations, retail gasoline prices have been more responsive in some states relative to others. For example, the Iraqi invasion of Kuwait on August 2, 1990 caused a surge in the real oil price of 39% over the month, which was passed through to state-level gasoline prices to varying degrees. Whereas Oklahoma experienced a 20% increase in the retail gasoline price in that month, Arizona only saw a 6% increase (see panel (a) of Figure 1). As we will show, this cross-sectional variation in gasoline price responses induced by an oil price shock helps identify the price elasticity of gasoline demand.

We start by estimating the systematic pass-through from oil to gasoline prices using state-level data from January 1989 to March 2008. This estimation period is chosen to facilitate comparisons with the literature on tax instruments. The pass-through from oil to pre-tax gasoline prices ranges from 9% to 65% with a mean of 42%. We then show that this variation is largely driven by the cost of supplying gasoline, which varies with the geographical location. The provision of gasoline to retail customers involves the transportation of crude oil to refiners, the reformulation of gasoline to satisfy environmental regulations, and the distribution of gasoline to city terminals and retail locations. In states where these costs account for smaller shares of retail gasoline prices, retail gasoline prices are more sensitive to oil price shocks, implying a higher degree of pass-through. In addition, the market power of local retailers also contributes to the difference in the pass-through. Together, the costs associated with producing and distributing gasoline and the retail market structure explain 74%-90% of the variation in the systematic pass-through across states.

Our proposed instrument consists of the interaction between the state-specific systematic pass-through and the monthly change in the real price of oil.³ The idea is that an increase in the global oil price induces cross-sectional gasoline price shocks as a result of differential cost structures in the gasoline supply chain, which are exogenous to demand conditions. One may be concerned that the systematic pass-through we estimate to some extent correlates with the state’s industry composition (e.g., the importance of the oil sector), such that state income varies with oil prices, biasing the elasticity estimate. We purge these effects by controlling for two-digit industry shares, their interactions with oil price changes and other income-related variables. In addition, we control for the differential effects of macroeconomic variables that may comove with oil prices. Given this rich set of controls and state and time fixed effects, it is unlikely that some unobserved variable would exist that both correlates with monthly oil price changes and differentially affects the change in state gasoline consumption.

With this strategy, we obtain a statistically significant gasoline demand elasticity of -0.31 for the period of January 1989 to March 2008. Our estimate is in line with the tax-based IV estimate in [Coglianese et al. \(2017\)](#), but more precisely estimated. We document that consumers are equally responsive to gasoline price changes driven by the tax component and by the pre-tax component. The reason why our conclusion differs from earlier studies such as [Li et al. \(2014\)](#) is that we account for the anticipation effect of gasoline tax changes and the endogeneity of pre-tax gasoline price changes. Our analysis suggests that reliable estimates of the price elasticity of gasoline demand are also informative about the effect of gasoline price changes driven by gasoline tax changes.

Since policy makers often care about the distributional effect of a nation-wide policy, it is important to examine heterogeneity in the gasoline demand elasticity along several dimensions. We find that consumers living in states where real personal income is lower tend to be more responsive to gasoline price changes. Consistent with this pattern, in states where the unemployment rate is higher, the gasoline demand elasticity tends to be higher in absolute terms. We also find that states with lower urban population shares, less commuting

³Instruments with a similar shift-share structure have also been employed in other recent studies. For example, [Nakamura and Steinsson \(2014\)](#) use this approach to estimate fiscal multipliers, [Guren et al. \(2021\)](#) use it to estimate the housing wealth effects, and [Kilian and Zhou \(2022\)](#) to estimate regional labor and housing market responses to aggregate demand shocks.

by public transit, and more registered motor vehicles per capita have more elastic gasoline demand. In contrast, we do not find evidence for an asymmetry in the elasticity with respect to positive and negative gasoline price shocks.

We then turn to the question of whether the price elasticity of gasoline demand has changed over time. To address this question, we extend the monthly data until March 2022. Based on rolling windows of 20-year length, we find a robust pattern. The IV estimate of the elasticity was stable near -0.3 until the end of 2014, rose to -0.2 in 2015-16, and has remained near -0.2 since 2016. Our analysis shows that, even allowing for the change in the elasticity over the last decade, demand remains more responsive than sometimes thought. Interestingly, we do not find evidence that the Covid-19 pandemic has altered the gasoline demand elasticity, even though indicators of mobility (such as time spent away from home) as of March 2022 remain below their pre-pandemic levels.

In the last part of the paper, we consider several empirical applications and policy experiments for which accurate estimates of gasoline demand elasticities matter. First, we evaluate the aggregate and distributional effects of the gasoline tax holidays proposed in 2022 amid surging gasoline prices. We find that a three-month federal gasoline tax holiday, as proposed by President Biden, would increase the discretionary income of the average household by only \$26, all else equal. In contrast, a seven-month state gasoline tax holiday (as implemented in New York state) would save households on average \$95 and in some states up to \$200. Second, the short-run demand elasticity may be used to assess the reduction in U.S. gasoline consumption caused by surging gasoline prices after the invasion of Ukraine in late February 2022. We estimate that U.S. gasoline consumption declined 3.5% cumulatively by April 2022 as a result. Lastly, we estimate that a 10% increase in retail gasoline prices would lower U.S. (global) carbon emissions by only 0.4% (0.05%) in the short run, suggesting a limited role for gasoline price increases and carbon taxes in curbing carbon emissions.

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 introduces our empirical strategy, presents estimates of the systematic pass-through, and links these estimates to the costs incurred along the gasoline supply chain. Section 4 presents our IV estimates and compares them with tax-based IV estimates. Section 5 examines the

heterogeneity in the elasticity estimate and the asymmetry between positive and negative gasoline price shocks. Section 6 examines the time-varying pattern of the price elasticity of gasoline demand based on rolling-window regressions. Section 7 provides additional analysis and robustness checks. Section 8 illustrates how reliable elasticity estimates can help answer policy relevant questions. The concluding remarks are in Section 9.

2 Data

We obtain monthly state-level retail gasoline prices from two sources. For the period of January 1989 to March 2008, prices of gasoline sold to end users in each state are obtained from the Energy Information Administration (EIA). These prices are collected through the EIA-782B form (suspended in 2011), a monthly survey of retailers and resellers, reflecting the sales-volume weighted prices of gasoline of all grades. For the period of April 2008 to March 2022, we obtain retail prices of regular, mid-grade and premium gasoline by state from the Oil Price Information Services (OPIS), which collects gasoline prices from fleet card transactions, consumer apps and direct reports from gas stations. Given the OPIS grade-specific prices, we construct the average gasoline price using the EIA state-level monthly sales volume for each grade as the weight. Figure 2 shows that the OPIS price at the aggregate level is very similar to that published by the EIA. At the state level, the two measures are also very close for the overlapping period. Figure C1 in the appendix plots the EIA-782B and OPIS prices for selected states, illustrating the similarity of the two series.⁴

State-level gasoline sales volume data are obtained from the EIA Prime Supplier Sales Volumes Series, collected through a monthly survey of all U.S. prime suppliers (refiners, gas plant operators, importers, retailers and resellers) who produce, import, or transport refined products within and across state boundaries and sell the product to local distributors, local retailers, or end users. The advantage of these gasoline consumption data over alternatives such as card transaction data and gas station surveys is the availability of a longer history, their broad coverage and the consistent method of data construction.

For state-level gasoline taxes, we extend the analysis in [Davis and Kilian \(2011\)](#) for the period after March 2008. State gasoline taxes as of December 31 each year are available from

⁴Our empirical results are also unaffected when we interact the state fixed effect with a dummy for the post-March-2008 period to account for potential changes in the gasoline price data.

the Federal Highway Administration (FHWA) Annual Highway Statistics. The effective date for a change in the state gasoline tax is documented in the Department of Transportation (DOT) Monthly Motor Fuel Reports. The two sources allow us to construct state-level monthly tax rates. As in [Davis and Kilian \(2011\)](#), we include *ad valorem* taxes in the measure of after-tax prices, but exclude them when using tax changes as the instrument.

Our measures of the costs incurred along the gasoline supply chain come from various sources. From the EIA, we obtain the sulfur content and API gravity of the oil input used by refiners at the Petroleum Administration for Defense District (PADD) level, the dealer tank wagon (DTW) price and the rack price at the state level, and wholesale gasoline prices by state. From the Environmental Protection Agency (EPA), we obtain the list of metropolitan areas where the sale of reformulated gasoline is mandated according to the 1990 amendments to the Clean Air Act. From the DOT Hazardous Liquid Annual Reports, we obtain oil and refined petroleum product pipeline miles by state. County distance data are available from the NBER County Distance Database. The number of retail gasoline stations by state and year is obtained from the Department of Energy Alternative Fuels Data Center.

To measure economic and demographic conditions at the state level, we collect the unemployment rate from the Bureau of Labor Statistics (BLS), personal income and population from the Bureau of Economic Analysis (BEA), employment by industry from the BLS Quarterly Census of Employment and Wages (QCEW) database, and the percentage of the urban population by state from the U.S. Census Bureau. To explore how the gasoline demand elasticity varies by location and reliance on motor vehicles, we collect data on the share of workers commuting by public transit from the American Community Survey and the number of motor vehicle registrations from the FHWA. All price and income variables are adjusted for CPI inflation.

3 Empirical Strategy

In this section, we discuss the identification of the price elasticity of gasoline demand in a classic simultaneous-equation system. The interaction between the oil price change and the differential pass-through from oil to retail gasoline prices will be a valid instrument if (i) global oil price fluctuations are unrelated to state economic conditions affecting gasoline

demand, and (ii) the differential pass-through from oil to gasoline prices is exogenous. We provide evidence supporting the validity of our instrument and discuss how our empirical strategy addresses possible concerns regarding these identifying assumptions.

3.1 Identification

We observe the gasoline price ($P_{i,t}$) and the volume consumed ($Q_{i,t}$) with i indexing a state and t the time period. The objective is to estimate the gasoline demand elasticity using a panel of I states and T periods. Gasoline demand and gasoline supply in logs can be described by the structural equations

$$\begin{aligned} q_{i,t}^d &= \beta_0 + \beta_1 p_{i,t} + \epsilon_{i,t}^d \\ q_{i,t}^s &= \delta_0 + \delta_1 p_{i,t} + \epsilon_{i,t}^s \\ q_{i,t}^d &= q_{i,t}^s = q_{i,t}, \end{aligned} \tag{1}$$

where β_1 is the price elasticity of gasoline demand, δ_1 is the price elasticity of gasoline supply, and $\epsilon_{i,t}^d$ and $\epsilon_{i,t}^s$ are structural error terms. In this system, price and quantity are jointly determined through shifts in demand and supply. Since $p_{i,t} = \frac{\beta_0 - \delta_0}{\delta_1 - \beta_1} + \frac{\epsilon_{i,t}^d - \epsilon_{i,t}^s}{\delta_1 - \beta_1}$, β_1 is not identified.

Identification of β_1 requires an instrument ($Z_{i,t}$) that captures exogenous variation in gasoline supply. In that case,

$$q_{i,t}^s = \delta_0 + \delta_1 p_{i,t} + Z_{i,t} + \epsilon_{i,t}^s, \tag{2}$$

and the equilibrium price is $p_{i,t} = \frac{\beta_0 - \delta_0}{\delta_1 - \beta_1} - \frac{Z_{i,t}}{\delta_1 - \beta_1} + \frac{\epsilon_{i,t}^d - \epsilon_{i,t}^s}{\delta_1 - \beta_1}$. Since $Z_{i,t}$ is orthogonal to $\epsilon_{i,t}^d$, β_1 is identified. Our proposed instrument exploits variation in the systematic component of the pass-through from oil price shocks to retail gasoline prices across states. Crude oil is a major input in gasoline production. The price of gasoline depends the price of crude oil and other costs. Whereas the price of crude oil is determined in the global market, the other costs may vary across states, for example, due to excise taxes, transportation costs or markups.

We decompose the percent change in the gasoline price into the systematic pass-through from oil prices ($\theta_i \Delta p_t^O$), an idiosyncratic component ($\tilde{\mu}_{i,t}$) and a constant (γ_i):

$$\Delta p_{i,t} = \theta_i \Delta p_t^O + \tilde{\mu}_{i,t} + \gamma_i, \tag{3}$$

where θ_i denotes the systematic pass-through of a 1% increase in the oil price to the retail gasoline price in state i . Substituting equation (3) into model (1) and expressing the variables

in differences yields

$$\begin{aligned}
\Delta q_{i,t}^d &= \beta_1 \Delta p_{i,t} + \Delta \epsilon_{i,t}^d \\
\Delta q_{i,t}^s &= \delta_1 \theta_i \Delta p_t^O + \delta_1 \tilde{\mu}_{i,t} + \delta_1 \gamma_i + \Delta \epsilon_{i,t}^s \\
\Delta q_{i,t}^d &= \Delta q_{i,t}^s = \Delta q_{i,t}.
\end{aligned} \tag{4}$$

Given model (4), as long as $\theta_i \Delta p_t^O$ is orthogonal to $\Delta \epsilon_{i,t}^d$ conditional on the state fixed effect (γ_i), $\theta_i \Delta p_t^O$ is a valid instrument for estimating β_1 . This implies two sufficient conditions for identification. One condition is that global oil price fluctuations, Δp_t^O , are unrelated to state economic conditions that affect gasoline demand. This condition is weaker than the assumption that global oil price fluctuations are unrelated to U.S. aggregate economic conditions. We provide extensive evidence supporting the validity of this assumption. The other condition is the exogeneity of the systematic pass-through, θ_i . As shown in Sections 3.3 and 3.4, variation in this pass-through is primarily driven by states' geographic attributes, refining technology, environmental regulations and retail market structure, which all contribute to the cost of the gasoline sold at the pump.

3.2 Implementation

Using monthly state-level panel data, the price elasticity of gasoline demand may be estimated from

$$\Delta q_{i,t} = \zeta_t + \gamma_i + \beta_1 \Delta p_{i,t} + \mathbf{b}\mathbf{x}_{i,t} + \varepsilon_{i,t}, \tag{5}$$

where $\Delta q_{i,t}$ is the change in log gasoline consumption per capita in state i and month t , $\Delta p_{i,t}$ the change in the log real retail gasoline price including taxes, and $\mathbf{x}_{i,t}$ a set of controls. ζ_t and γ_i are the month and state fixed effects. The standard errors are clustered at the state level. β_1 is the one-month price elasticity of gasoline demand.

As illustrated in Section 3.1, the OLS estimate of β_1 is likely to be biased due to unobserved demand shocks that move $\Delta p_{i,t}$ and $\Delta q_{i,t}$ in the same direction. To address this identification challenge, we exploit cross-sectional variation in the systematic pass-through from oil price shocks to retail gasoline prices. There are two ways of making the IV strategy operational. One is to use $\hat{\theta}_i \Delta p_t^O$ as the instrument, with $\hat{\theta}_i$ estimated from

$$\Delta p_{i,t} = \xi_t + \gamma_i + \theta_i \Delta p_t^O + \alpha \mathbf{x}_{i,t} + u_{i,t}, \tag{6}$$

where $\theta_i \Delta p_t^O$ is short for $\sum_{k \in I} \theta_k \mathbb{I}_{k=i} \Delta p_t^O$ in the panel regression, with $\mathbb{I}_{k=i}$ the indicator

for state i , and ξ_t denotes the month-of-the-year and year fixed effects.⁵ Panel (b) of Figure 1 shows the change in gasoline prices predicted by this instrument for Oklahoma and Arizona around the time of the invasion of Kuwait in 1990. The figure highlights the higher pass-through from oil to gasoline prices in Oklahoma after purging confounding factors and idiosyncratic shocks that affect the actual gasoline price in panel (a) of Figure 1.⁶

An alternative and cleaner approach is to exploit variation in the systematic pass-through from an oil price shock to the pre-tax gasoline price by estimating

$$\Delta p_{i,t}^{EX} = \xi_t + \gamma_i + \eta_i \Delta p_t^O + \alpha \mathbf{x}_{i,t} + \nu_{i,t}, \quad (7)$$

where $\Delta p_{i,t}^{EX}$ denotes the change in the log real retail gasoline price excluding state and federal taxes. The instrument $\hat{\eta}_i \Delta p_t^O$ is then used to estimate β_1 in equation (5). We report the results from both approaches. Which approach is used makes little difference in practice.

To ensure the orthogonality between Δp_t^O and $\varepsilon_{i,t}$, we control for a number of state-level income measures that could potentially bias $\hat{\beta}_1$, including the growth of real personal income per capita, the unemployment rate, and two-digit industry shares. In addition, we control for the differential effects of macroeconomic variables that may comove with oil prices, such as the 1-year treasury yield, aggregate unemployment rate and CPI inflation.⁷ As a result, it is unlikely that there are unobserved factors that are both correlated with oil price changes and that differentially affect states with higher systematic pass-through from oil to gasoline prices. Moreover, we control for the full set of interactions between the oil price change and six-month moving-average industry shares. This alleviates the concern that oil price shocks may affect gasoline demand through channels other than gasoline price changes. For example, in states where the oil sector is relatively more important, income and employment are likely to be more responsive to oil price shocks. Failing to control for the interaction

⁵The panel estimation of the systematic pass-through does not allow for the inclusion of the full set of month fixed effects, because $\sum_{k \in I} \mathbb{I}_{k=i} \Delta p_t^O$ would be perfectly collinear with Δp_t^O , which is subsumed in the month fixed effect. Thus, θ_k would not be identifiable. However, conditional on the systematic pass-through instrument $\hat{\theta}_i \Delta p_t^O$, we include the full set of month fixed effects in the IV estimation of the price elasticity of gasoline demand.

⁶Unlike with generated regressors, we need not correct the standard errors for $\hat{\beta}_1$ when using $\hat{\theta}_i \Delta p_t^O$ as the instrument. This is because under mild assumptions, the asymptotic distribution of $\hat{\beta}_1$ in equation (5) is the same whether θ_i or $\hat{\theta}_i$ is used in constructing the instrument (see Wooldridge (2001)).

⁷We estimate these differential effects in a way similar to estimating the systematic pass-through from oil to gasoline prices. For a macroeconomic variable y_t , we obtain $\hat{\kappa}_i$ from estimating $\Delta p_{i,t} = \xi_t + \gamma_i + \kappa_i y_t + \alpha \mathbf{x}_{i,t} + u_{i,t}$. We then include $\hat{\kappa}_i y_t$ as a control variable.

between the oil price change and the oil sector share would lead to attenuation bias, because gasoline demand increases with income.⁸

A more general concern regarding this identification assumption is that economic conditions in some U.S. states could affect global oil prices. In Section 7, we show that dropping individual states (or oil states altogether) does not affect our IV estimates. More importantly, since our strategy can be easily applied to alternative panel-data settings, we show, using data for 241 metropolitan statistical areas (MSAs), that the MSA-level estimates, for which such feedback seems implausible, are very similar to the state-level estimates.

Before establishing the exogeneity of θ_i and η_i , it is useful to relate our strategy to the shift-share instrument design that has been widely used in empirical studies. Our estimator may be interpreted as a shift-share estimator with the shifts corresponding to aggregate oil price shocks and the shares corresponding to θ_i and η_i . The identification of and inference in shift-share models is discussed in [Goldsmith-Pinkham et al. \(2020\)](#), [Borusyak et al. \(2022\)](#) and [Adao et al. \(2019\)](#) who show that the consistency of the shift-share estimator requires either the shifts or the shares to be exogenous. Since our design involves aggregate shocks to a specific industry, the exogeneity of the shares is key to our identification. Sections 3.3 and 3.4 provide empirical evidence that variation in θ_i and η_i is primarily driven by exogenous differences in the cost of gasoline. Appendix D reports additional tests of the plausibility of this identifying assumption recommended by [Goldsmith-Pinkham et al. \(2020\)](#).⁹

⁸Including these interaction terms also addresses the concern that unobserved global demand shocks driving the oil price could affect U.S. states differentially. One example is China’s rise after the 1990s, which raised global demand for oil and negatively impacted U.S. states where imports increased the most. To the extent that consumers in these states cut back on gasoline consumption not only because of higher gasoline prices but also due to lower income, the value of the gasoline demand elasticity would be overstated in absolute terms.

⁹Whereas [Goldsmith-Pinkham et al. \(2020\)](#) treat exogeneity of the shares as a sufficient condition for consistency, [Borusyak et al. \(2022\)](#) derive two conditions for identification: (i) shifts are quasi-randomly assigned and (ii) there are many exogenous and independent shocks (to ensure that the bias resulting from non-random shares averages out to zero). Since we focus on oil price shocks, the framework of [Borusyak et al. \(2022\)](#) does not apply to our setting. For the same reason, the standard errors derived in [Adao et al. \(2019\)](#) under similar assumptions cannot be applied to our setting. However, we follow [Adao et al. \(2019\)](#) in conducting a placebo test, in which $\hat{\beta}_1$ is estimated with randomly generated oil price shocks and actual gasoline consumption and systematic pass-through. Our results show that conventional clustered standard errors do not pose a severe over-rejection problem. Our elasticity estimate remains highly statistically significant even after adjusting the critical values accordingly.

3.3 What Explains Variation in the Systematic Pass-Through?

Panel (a) of Figure 3 shows the heat map of the systematic pass-through from oil prices to pre-tax gasoline prices, $\hat{\eta}_i$, estimated using equation (7). This pass-through rate ranges from 9% to 65%, with a mean of 42%. The South Central (e.g. Oklahoma) and Midwest (e.g., Ohio, Kansas, Missouri, Michigan and Indiana) regions have the highest pass-through (greater than 60%), whereas the West (e.g., Arizona, Washington and Nevada) and Northeast (e.g. New York and New Jersey) have the lowest pass-through (less than 30%). The pass-through from oil to post-tax gasoline prices, $\hat{\theta}_i$, displays a similar cross-sectional pattern (panel b), but has less dispersion (ranging from 9% to 48%). This is because the tax component does not change frequently and its change is nearly uncorrelated with that in the pre-tax component.

Given the complex supply chain involved in transforming crude oil into retail gasoline at the pump, variation in the systematic pass-through reflects the buildup of costs along the gasoline supply chain. The chain starts with U.S. refineries buying crude oil from domestic and foreign producers and having it shipped to their facilities via pipeline, truck, rail or tanker.¹⁰ Refiners process this oil into gasoline and other refined products.¹¹ Nearly all of the gasoline sold in the U.S. is produced domestically. Refiners sell large quantities of generic gasoline directly from the refinery to distributors and other refiners in spot transactions. Generic gasoline is then moved to large storage terminals near population centers mainly by product pipelines, but also to a limited extent by truck and barge. From large storage terminals, gasoline is shipped by truck to smaller blending terminals for processing into finished motor gasoline (i.e., conventional, reformulated, or oxygenated gasoline). Tanker trucks pick up the finished motor gasoline at the blending terminal and deliver it to the underground storage tanks at gas stations. Gas station owners set the retail prices at the pump.¹²

¹⁰In the U.S., 70% of crude oil and petroleum products are shipped by pipeline, 23% by tankers and barges, 4% by trucks and only 3% by rail.

¹¹One barrel of crude oil (42 gallons) yields about 19 to 20 gallons of motor gasoline. The composition of refined products can be altered by changing the refining process, but the scope for output substitution is limited (see [Borenstein et al. \(1997\)](#)).

¹²Oil companies and refiners do not play an important role in setting retail gasoline prices, as the vast majority of branded stations are owned and operated by independent retailers licensed to represent that brand by paying franchise fees. According to the National Association of Convenience Stores, convenience

Retail prices at the pump reflect the costs added along this chain. We classify these costs into three categories: refining costs, distribution costs, and retail markups. We use available data to measure cross-state variation in these costs. Our cost measures explain 74%-90% of variation in the systematic pass through. Figure B1 in the appendix shows the heatmaps for selected cost measures which display a clear regional pattern.

3.3.1 Variation in Refining Costs

Transporting crude oil. Refiners' costs of obtaining oil are higher, the further they are located from the source of oil supply, since transportation costs increase with the distance. Cushing, Oklahoma, plays a central role in the trading of oil. This is not only because the city is the delivery location for the benchmark light sweet crude oil futures contract (known as WTI), but also because of its massive storage capacity and expansive inbound and outbound pipeline infrastructure that transports oil produced in U.S. shale areas and in Canada and supplies oil to the main refining centers in the Midwest and on the Gulf coast. Refiners in the coastal areas, in contrast, have very limited pipeline infrastructure for transporting domestic oil and historically have had to import foreign oil to meet demand.

This suggests that refiners located near Cushing and connected to Cushing by oil pipelines, all else equal, have better access to crude oil supply and pay less for oil transportation. We therefore use the distance between Cushing and a state's largest population center as a proxy for the state's refinery costs of transporting crude oil. We also use state-level historical average crude oil pipeline miles normalized by the state's area as a proxy for pipeline connectivity to domestic oil.

Processing crude oil. Not all crude oil is the same. U.S. refiners are configured to process different grades of crude oil classified by the density (known as the API gravity) and sulfur content of the oil. Light sweet oil such as WTI crude (high API gravity and low sulfur content) is desirable for producing gasoline because it is easier to refine, distill, and transport. In contrast, heavy sour crude oil requires additional and more expensive processing to produce high-value products. This process lowers the sensitivity of the finished gasoline price to oil price changes. We use the historical average API gravity of oil inputs

stores sell approximately 80% of the U.S. motor fuels, and about 60% of the retail stations are owned by an individual or family that owns a single store.

reported by U.S. refiners at the PADD level as a proxy for the cost of processing crude oil.¹³ The data show that East Coast, Midwest and Rocky Mountain refiners tend to use lighter crude oil, whereas Gulf Coast and West Coast refiners use heavier oil.

Meeting environmental regulations. Some areas in the U.S. are required to use reformulated gasoline (RFG) that includes additives to help reduce carbon monoxide, smog, and toxic air pollutants. The 1990 amendments to the Clean Air Act mandated the sale of RFG in the nine metropolitan areas with the highest smog levels.¹⁴ Any area reclassified as suffering from severe ozone nonattainment also becomes a mandated RFG program area.¹⁵ In addition to the federal RFG program, states may implement their own gasoline programs. In California, for example, all counties implement another version of the RFG program. These programs add to the cost of producing, storing, and distributing gasoline. To capture variation in the cost of fulfilling mandatory RFG programs, we construct a state-level RFG exposure measure that takes the value of one if the population in the areas covered by mandatory RFG programs exceed 50% of the state population in 2000. Our empirical results are robust to changing this threshold.

3.3.2 Variation in Gasoline Distribution Costs

From refineries to terminals. Since the majority of refined products are transported to city terminals by pipeline, we use state-level historical average product pipeline miles normalized by the state area as a proxy for the cost of transporting gasoline to city terminals. A denser pipeline network implies higher transportation capacity and lower distribution costs. Figure B2 in the appendix shows the map of refined product pipelines in 2022 (similar to maps created using data in the 2010s), suggesting substantial variation in the density of the pipeline network across states.

From terminals to stations. As in [Deltas \(2008\)](#), we use the difference between the DTW

¹³We focus on the API gravity because heavier oil also tends to have a higher sulfur content. We also find that the API gravity has a higher explanatory power for the systematic pass-through than the sulfur content. In regressions that include both API gravity and sulfur content as explanatory variables for the pass-through, the sulfur content does not have a significant effect.

¹⁴These metropolitan areas are Baltimore, Chicago, Hartford, Houston, Los Angeles, Milwaukee, New York City, Philadelphia, and San Diego.

¹⁵Areas currently or previously designated as not attaining ozone limits may be included in the RFG program at the request of the governor of the state. This is known as the RFG opt-in program. There are also procedures that allow states or areas to opt out the RFG program.

price and rack price to measure local transportation costs. The DTW price is the wholesale price including delivery, whereas the rack price is for truckload excluding delivery. The historical average DTW-rack margin is about 9%. It is highest in Washington D.C. (26%), and lowest in Utah (2%).

3.3.3 Variation in Retail Market Power

Some studies suggest that in areas where retailers have greater market power, gasoline prices respond more slowly to changes in wholesale prices (e.g., [Borenstein et al. \(1997\)](#)), which implies a lower pass-through within one month. We construct two measures for the market power of retailers at the state level. One is the retail markup (i.e., the percent deviation of nominal retail prices from nominal wholesale prices), and the other is the density of gasoline stations (i.e., the number of retail gas stations per 1,000 people).¹⁶

3.4 How Much Variation Does the Supply Side Explain?

Columns (1)-(8) of Table 1 show results from regressions of the pass-through, $\hat{\eta}_i$, on each of the supply-side cost measures, one at a time. Except for the RFG exposure indicator, all explanatory variables are standardized to show the effect of a one standard deviation increase in the variable. The coefficients have the expected signs. The pass-through is higher for states (i) located near the oil trading center in Cushing, Oklahoma, (ii) having denser oil and product pipeline networks, (iii) using lighter crude oil inputs, (iv) less exposed to RFG programs, (v) having smaller local transportation costs, and (vi) featuring a more competitive retail market (with lower retail markup and higher gas station density).

Since these measures may be correlated with each other, including them in a multiple regression is more appropriate in assessing their joint explanatory power. As shown in column (9), these variables together explain 74% of variation in $\hat{\eta}_i$. The distance to Cushing, API gravity, RFG exposure, product pipeline density and retail markup are significant at the 5% level. Their usefulness in explaining variation in $\hat{\eta}_i$ can also be seen from the R^2 of the joint regression leaving out one variable at a time. It falls below 73% for the specifications where the variable left out is significant in column (9).

Column (10) is our baseline regression restricted to include only statistically significant

¹⁶The retail markup in principle captures local transportation costs, retail operational costs, and market power. By controlling for the first two in the regression, we are able to capture variation in market power.

regressors. Together, these regressors explain 73% of the variation in $\hat{\eta}_i$. In Figure 4, the left panel plots $\hat{\eta}_i$ against its predicted value using the specification in column (10). The R^2 is 73%. Similar results are obtained for the pass-through from oil to retail gasoline prices, $\hat{\theta}_i$, as shown in the right panel, with an R^2 of 72%.

In column (11) of Table 1, we also include PADD fixed effects to capture common supply-side constraints faced by refiners in that PADD. This specification explains more than 90% of the variation in $\hat{\eta}_i$. We are cautious in interpreting these fixed effects, however, because they may also capture common demand factors. Finally, we add an explicit demand control, the historical average level of log real income per capita, which does not affect the coefficient estimates or R^2 in column (10).

Our analysis suggests that much of variation in the pass-through is explained by differential costs incurred along the gasoline supply chain, rather than demand factors. In later sections, our empirical strategy for estimating the price elasticity of gasoline demand will mainly use $\hat{\eta}_i \Delta p_t^O$ and $\hat{\theta}_i \Delta p_t^O$ as the instruments, but as shown later, using the pass-through predicted by the set of cost measures in column (10) of Table 1 gives similar estimates.

4 Empirical Results

In this section, we address three key questions. First, what price elasticity of gasoline demand is implied by our identification strategy and is this estimate robust to alternative specifications? Second, how does our IV estimate compare to the estimate using tax-based instruments and is our estimate robust to explicitly accounting for tax changes? Third, are consumers more responsive to changes in gasoline taxes than to changes in the pre-tax gasoline price (e.g., Li et al. (2014); Rivers and Schaufele (2015)). For our results to be comparable with related studies, we restrict the sample to January 1989-March 2008 for now. In Section 6, we address the question of whether the price elasticity of gasoline demand has changed over time using data ending in March 2022.

4.1 IV Estimate Based on the Systematic Pass-Through

Table 2 shows the price elasticity of gasoline demand estimated using the IV strategies detailed in Section 3.2. For comparison, the OLS estimate is -0.19, which is likely biased toward zero due to the endogeneity concerns discussed earlier. Some studies attempt to

address this concern by using global oil price changes or oil price shocks associated with supply disruptions as the instrument. As shown in Appendix A, these methods either do not address the endogeneity concern or suffer from a weak IV problem.

Column (2) in Table 2 is our baseline specification using $\hat{\eta}_i \Delta p_t^O$ as the instrument. The estimate of the gasoline demand elasticity is -0.31, strongly statistically significant and larger than the OLS estimate in absolute value. Column (3) including additional controls shows a slightly larger elasticity estimate, -0.37, but it is statistically indistinguishable from the estimate in column (2). Column (4) shows the estimate using $\hat{\theta}_i \Delta p_t^O$ as the instrument with the full set of controls, which is almost unchanged from column (3). Columns (5) and (6) present estimates that exploit variation in the systematic pass-through predicted by our cost measures. Specifically, we use the predicted value from the regression in column (10) of Table 1 interacted with the oil price change as the instrument. The resulting estimates (-0.33 and -0.34) are similar to the baseline estimate. Column (7) uses the interactions between the observed cost measures and the oil price change as instruments and again yields a similar estimate (-0.36). In all specifications, the instruments pass the Montiel Olea and Pflueger (2013) robust weak IV test, with the effective first-stage F-statistics strongly exceeding the 5% critical value for rejection.¹⁷

Finally, we consider another variant of our IV estimator based on the average pass-through in adjacent states, defined as Δp_t^O multiplied by $\bar{\hat{\eta}}_{-i} = \sum_{k=1}^{N_i} \hat{\eta}_k / N_i$, where N_i is the number of neighboring states, as the instrument for state i 's gasoline price change. This exercise is motivated by the observation that, when cost shocks are correlated across states but demand shocks are not, focusing on neighboring states' pass-through helps address concerns that unobserved demand shocks in state i may be correlated with $\hat{\eta}_i$ (see Hausman (1996)). Indeed, we find a strongly positive relationship between a state's own pass-through and the pass-through in its neighboring states (see Appendix Figure B3). Column (8) in Table 2 shows that this alternative estimate is almost identical to the estimate in column (3), adding further credence to our analysis. Analogous results hold when applying this strategy

¹⁷To account for possible co-integration between oil and gasoline prices, we also estimate the pass-through rates with state-specific error correction terms included in equation (7). This term is the residual from estimating $p_{i,t}^{EX} = \xi_t + \gamma_i + \eta_i p_t^O + \alpha \mathbf{x}_{i,t} + u_{i,t}$. The resulting elasticity estimate is -0.31 (s.e.=0.065), almost identical to the baseline estimate.

to the baseline specification.¹⁸

4.2 Comparison with the Tax-Based IV Estimate

Whereas we exploit exogenous variation in the pre-tax component of retail gasoline prices, Coglianesi et al. (2017) proposed using monthly changes in state excise gasoline taxes as the instrument. Changes in state gasoline taxes are strongly correlated with retail gasoline price movements, but due to long lags in passing and implementing legislation, these tax changes are predetermined with respect to the current price movements, making them valid instruments. In Coglianesi et al.’s approach, not only the contemporaneous percent change in the gasoline price, $\Delta p_{i,t}$, is instrumented, but also the lead and the lag ($\Delta p_{i,t-1}$ and $\Delta p_{i,t+1}$) are instrumented to account for the anticipation and avoidance effects of pre-announced tax changes. The resulting cumulative response gives a price elasticity of gasoline demand of -0.37, as shown in column (1) of Table 3. This elasticity, however, is not precisely estimated and is not statistically distinguishable from zero at conventional significance levels. The elasticity obtained using our approach is of a similar magnitude but is highly statistically significant (column 2, Table 3).

Our approach exploits a distinct source of variation in retail gasoline prices that is independent of taxes changes. To see this, in column (3) of Table 3, we control for monthly changes in state gasoline taxes and their lead and lag to allow for the direct impact of tax changes on gasoline consumption, while instrumenting gasoline price growth with the systematic pass-through IV, $\hat{\eta}_i \Delta p_t^O$. The point estimate is -0.32 (essentially unchanged from our baseline specification) and is highly statistically significant. To conclude, our IV strategy generates a very similar point estimate to the tax-based IV approach, despite using a different source of variation, but our strategy has the advantage of a higher statistical precision, which facilitates the investigation of cross-sectional heterogeneity and time-varying patterns. An additional advantage of our approach is that it allows IV estimation even in the absence of state- or city-level tax data. It only requires access to oil price data.

¹⁸A potential concern with our approach is that motorists in response to differential gasoline price shocks may cross state lines to fill their gas tanks, biasing the elasticity estimate. This effect is likely to be quantitatively unimportant for two reasons. First, states’ population centers are typically not located close to state borders, according to census data. Second, our evidence shows that the systematic pass-through is highly correlated across neighboring states, diminishing the incentive for shopping across the border.

4.3 Are Consumers More Responsive to Gasoline Tax Changes?

Our strategy helps answer the question of whether consumers are more responsive to changes in the tax component of retail gasoline prices than changes in the pre-tax component. Previous studies using state-level data such as [Li et al. \(2014\)](#) argued that the tax elasticity of gasoline demand exceeds the price elasticity for two reasons. One is the salience of gasoline tax changes that are often extensively covered by the media around the time of implementation. The other reason is that consumers may perceive gasoline price changes driven by tax changes to be permanent, but not other gasoline price changes. [Li et al. \(2014\)](#) therefore caution against the use of tax elasticities when analyzing generic gasoline price shocks.

On the other hand, [Coglianese et al. \(2017\)](#) reported an external validity exercise which showed that their tax elasticity estimate accurately predicts the gasoline consumption response to the gasoline price surge triggered by Hurricanes Rita and Katrina. Moreover, our more direct evidence suggests that tax and price elasticities are quite close in practice. The question is how to reconcile these facts with the empirical pattern documented in [Li et al. \(2014\)](#) of a higher responsiveness of gasoline consumption to tax shocks than to non-tax price shocks.

We first replicate this pattern in column (1) of Table 4, using a specification similar to [Li et al. \(2014\)](#) that regresses $\Delta q_{i,t}$ on $\Delta p_{i,t}^{EX}$ and the normalized tax change, $\Delta \tau_{i,t} / p_{i,t-1}^{EX}$. This normalization ensures that the resulting changes in the tax-inclusive price are of the same magnitude (in dollar amounts). The result suggests that gasoline consumption is much more responsive to the contemporaneous change in the state gasoline tax (-0.86 vs. -0.13), and that the hypothesis of equal effects is strongly rejected with a p-value of 0.001.

This specification, however, suffers from two shortcomings in light of more recent research. First, anticipation of tax changes may affect gasoline consumption before the actual implementation of the change, so the inclusion of the lead and lag of tax changes is necessary for correctly estimating the tax effect. Second, $\Delta p_{i,t}^{EX}$ may be endogenous to gasoline consumption, necessitating the use of an IV estimator.

Column (2) of Table 4 addresses the first concern by including one lead and one lag of

the normalized tax change, as in [Coglianese et al. \(2017\)](#). As a result, gasoline consumption increases significantly in the month before the tax change, and falls sharply when the change is implemented. Ignoring the anticipation effect exaggerates the response to the tax change. The correct tax effect of -0.16 is given by the cumulative response. The p-value for testing the null of equal effects now is 0.9, so we fail to reject the null. In column (3), we address the second concern by instrumenting $\Delta p_{i,t}^{EX}$ with $\hat{\eta}_i \Delta p_t^O$. The gasoline consumption response to a change in the pre-tax component is larger than the OLS estimate, while the tax effect is similar. The p-value for testing the null of equal effects is 0.7, again providing no evidence for a larger response to the tax change. This conclusion also holds in the extended sample ending in March 2022 (see Appendix Table C2).

In short, Table 4 supports the view that consumers are equally responsive to gasoline price changes driven by the tax component and the pre-tax component. This result is not surprising for two reasons. First, gasoline price changes unrelated to gasoline tax changes are likely to be equally salient to consumers. Not only are these price changes reflected in posted prices at gas stations, but extensive media coverage, especially during periods of heightened geopolitical tensions in oil markets (e.g., the 2022 invasion of Ukraine) and periods of unexpectedly large shocks hitting the economy (e.g., the Covid-19 outbreak in 2020), tends to draw further attention to these price changes. Second, since most observed gasoline price changes reflect changes in the pre-tax component, the fact that a no-change forecast fits household expectations of the real gasoline price consistently well with only rare exceptions, as documented in [Anderson et al. \(2013\)](#), argues against the view that consumers differentiate between alternative sources of retail gasoline price fluctuations.

5 Heterogeneity in Gasoline Consumption Responses

Our estimates so far inform us about the average price elasticity of gasoline demand. There are reasons to believe that households may respond differently to the same price change due to differences in income, local economic conditions, urbanization and commuting patterns. Moreover, it is important for policymakers to know whether asymmetry exists in the responses to positive and negative gasoline price shocks. Understanding these dimensions of heterogeneity helps assess the distributional effect of aggregate policies (such as a federal

gasoline tax holiday) and evaluate regional policies (such as state gasoline tax changes). Existing studies of heterogeneous effects using disaggregate data rely on OLS estimators and have not adequately addressed the identification issue. In this section, we study these effects building on the analysis in Section 4. In Appendix C, we show that these patterns also hold in the extended sample period.

5.1 By Income, Location and Reliance on Motor Vehicle

We focus on five key dimensions of heterogeneity: (i) the average level of real personal income per capita, as a measure of household wealth, (ii) the unemployment rate in the preceding 12 months, as a measure of state business cycles, (iii) the share of urban population, as a measure of urbanization, (iv) the share of workers commuting by public transit, and (v) the number of motor vehicle registrations per capita, as a proxy for vehicle ownership. For each measure of heterogeneity, we estimate the gasoline demand elasticity separately for states above and below the median of the distribution using the baseline IV specification. The estimates are robust to using alternative specifications in Table 2.

Figure 5 presents the results. States with lower personal income are more responsive to gasoline price changes. Their gasoline demand elasticity, -0.43 , is almost twice as high as that of higher-income states. This heterogeneity supports the view that low-income households are more exposed to gasoline price shocks due to a higher fraction of their income spent on this category, and hence are more sensitive to these shocks. Higher income households, in contrast, tend to have higher liquid wealth, which helps them smooth consumption in response to expenditure shocks. For the same reason, we expect the gasoline demand elasticity to vary over the state business cycle. During economic booms, demand for gasoline should be less elastic, compared to periods of economic downturns. This is supported by the data. In states where the unemployment rate has been high in the past 12 months, the gasoline demand elasticity is -0.41 , compared to -0.2 in periods of low unemployment rates.

We find that states with higher urban population shares have less elastic gasoline demand, with an elasticity estimate of -0.1 . In contrast, the elasticity for states with lower urban population shares is -0.5 . Consistent with this finding, states with higher shares of workers commuting to work by public transit have less elastic gasoline demand (-0.17) than others

(-0.47). Finally, splitting the sample based on the number of motor vehicles per capita suggests that states above the median have more elastic gasoline demand (-0.46).

5.2 Responses to Positive and Negative Gasoline Price Shocks

There is no consensus in the literature on the existence of an asymmetry in the response of gasoline consumption to gasoline price changes. [Levin et al. \(2022\)](#), for example, using city-level daily data provide evidence that gasoline demand is more elastic when gasoline prices rise above their average over the previous year than when prices fall below this average. [Kilian and Vigfusson \(2011\)](#), in contrast, using augmented monthly structural VAR models do not find evidence for asymmetry. Our empirical strategy allows us to examine this issue for monthly data in the panel setting. A related debate is about the asymmetric pass-through from oil to retail gasoline prices. We proceed in two steps, first testing asymmetry in the response of retail gasoline prices to oil price shocks, and then testing asymmetry in the price elasticity of gasoline demand.

We test for asymmetry in the pass-through from oil to pre-tax gasoline prices as follows. For each state, we estimate

$$\Delta p_{i,t}^{EX} = \xi_t + \gamma_i + \eta_i \Delta p_t^O + \eta_i^+ \Delta p_t^{O,+} + \alpha \mathbf{x}_{i,t} + u_{i,t} \quad (8)$$

and test the null of $\eta_i^+ = 0$. The nonlinear term, $\Delta p_t^{O,+} \equiv \max\{0, \Delta p_t^O\}$, accounts for any additional effect of an oil price increase. We compute Newey-West standard errors with three lags. A similar procedure is used for testing asymmetry in the pass-through from oil prices to after-tax gasoline prices.

The first row of [Table 5](#) shows the distribution of the p-value for testing $\eta_i^+ = 0$. Using conventional critical values, we reject the null of symmetric pass-through for only 4 out of 51 states at the 5% significance level. None of the states shows significant asymmetry at the 1% level. For the pass-through from oil to after-tax gasoline prices, the results are similar. Overall, we do not find compelling evidence for the existence of asymmetric pass-through. This result is consistent with studies using richer dynamic models that find no support for asymmetric energy price responses at the one-month horizon (e.g., [Venditti \(2013\)](#); [Lewis and Noel \(2011\)](#); [Borenstein et al. \(1997\)](#)).

To test for asymmetry in the gasoline demand response, we estimate a modified version

of equation (5) in the second stage:

$$\Delta q_{i,t} = \zeta_t + \gamma_i + \beta_1 \Delta p_{i,t} + \beta_1^+ \Delta p_{i,t}^+ + \mathbf{b}x_{i,t} + \varepsilon_{i,t}, \quad (9)$$

with $\Delta p_{i,t}$ and $\Delta p_{i,t}^+$ instrumented by $\hat{\eta}_i \Delta p_t^O$ and $\hat{\eta}_i \Delta p_t^{O,+}$ (or $\hat{\theta}_i \Delta p_t^O$ and $\hat{\theta}_i \Delta p_t^{O,+}$). $\hat{\eta}_i$ and $\hat{\theta}_i$ are estimated using the linear model, since we do not find evidence for asymmetric pass-through. However, our test of the asymmetric gasoline demand elasticity does not rely on this condition. We also estimate the specification that allows for asymmetric pass-through from oil to gasoline prices. In this case, $\Delta p_{i,t}$ and $\Delta p_{i,t}^+$ are instrumented with $\hat{\eta}_i \Delta p_t^O$ and $\hat{\eta}_i^+ \Delta p_t^{O,+}$ (or $\hat{\theta}_i \Delta p_t^O$ and $\hat{\theta}_i^+ \Delta p_t^{O,+}$). The results, shown in Table 6, provide no evidence for asymmetric gasoline consumption responses. The coefficient of the non-linear term is statistically insignificant regardless of the identification strategy (OLS vs IV) and regardless of whether we allow for asymmetric pass-through from oil to retail gasoline prices.

6 Has the Gasoline Demand Elasticity Changed Over Time?

Since Hughes et al. (2008), an important question has been whether the gasoline demand elasticity has declined over time. Data limitations and identification challenges are the biggest hurdles in answering this question. Our data and approach offer several advantages compared to previous studies. First, we employ data from January 1989 to March 2022. This unusually long panel data set at the monthly frequency allows us to obtain stable and reliable estimates of time-varying demand elasticities. Second, our IV strategy can be easily applied in this context to address the identification issue. Third, we are in a position to examine whether the Covid-19 pandemic has altered consumers' gasoline demand elasticity substantially. We first present the full-sample estimates for the extended estimation period and then report rolling window estimates of the time-varying elasticity.

6.1 Estimates Using the Extended Data

Table 7 shows the estimates using the extended data from January 1989 to March 2022. Both the OLS estimate (-0.17) and the IV estimate (-0.26) are somewhat smaller (in absolute value) than for the data ending in March 2008, suggesting a moderation in the responsiveness of gasoline demand in more recent years. To see if this change is driven by the Covid-19 pandemic, in column (3) we drop the pandemic period (starting from March 2020). The estimate does not change much; if anything, the elasticity is smaller in magnitude, suggesting

that the moderation in the demand responsiveness already started before the pandemic.

The impact of the pandemic on energy prices and driving behavior was extreme in the initial months of the Covid-19 outbreak. The real price of oil cumulatively fell by more than 50% in March and April, before surging by 55% in May. Gasoline prices also fluctuated considerably, but not as dramatically as the oil price. More importantly, temporary lock-down measures and social distancing substantially reduced time spent away from home between March and May (by as much as 25% relative to February 2020, according to the *Opportunity Insights Economic Tracker*) and hence driving. Given these facts, the unusually large shocks to the economy in the first few months of the pandemic are likely to introduce outliers in our analysis that may affect the estimates of the systematic pass-through and the gasoline demand elasticity. We are reluctant to discard all pandemic observations, which would prevent us from gaining insights about the post-pandemic behavior. In our main analysis, we therefore only drop observations for March 2020 through May 2020. The results are very similar when we instead include separate fixed effects for the pandemic slump period (March and April 2020), the swift rebound period (May and June 2020), and the slow recovery period (July 2020 and onwards).

Column (4) of Table 7 shows that the elasticity estimate, after dropping the three initial months, is -0.2, similar to the pre-pandemic estimate. Column (5) shows that the estimate is robust to exploiting the pass-through from oil to after-tax gasoline prices. This result is also robust to various alternative specifications as in Table 2 (not shown to conserve space).

6.2 Rolling Window Estimates

When did the gasoline demand elasticity start to decline in absolute value? To address this question, we apply our IV strategy to rolling windows of data. Before we turn to the time-varying estimates of the elasticity, we first examine whether the systematic pass-through from oil prices to pre-tax gasoline prices displays any time variation.

Figure 6 shows the distribution of $\hat{\eta}_i$ over time using a 20-year rolling window. At each point in time, the distributional statistics (the median, the 25th-75th percentile range, and the 10th-90th percentile range) are obtained from estimating the panel regression (7) using 20 years of data ending at that point. There is no evidence for a time-varying pattern in the

systematic pass-through. While the median increased slightly from 47% in December 2008 to 55% in March 2022, the interquartile and interdecile ranges remain unchanged. We find a similar time-invariant pattern for the pass-through from oil to after-tax gasoline prices ($\hat{\theta}_i$).

Given that there is no evidence for time variation in the pass-through from oil to gasoline prices, we construct a rolling-window estimate of the price elasticity of gasoline demand, with $\hat{\eta}_i \Delta p_t^O$ the instrument and $\hat{\eta}_i$ estimated using the full sample (dropping March-May 2020). Figure 7 shows that this elasticity was stable around -0.3 until the end of 2014, rose to -0.2 in 2015-16, and has remained near -0.2 since 2016. Despite the increase in the elasticity, the magnitude still implies more responsive gasoline demand than traditionally thought. The Covid-19 pandemic has not substantially changed this elasticity.

The moderate rise in the gasoline demand elasticity since 2015 is robust to alternative model specifications. In Appendix Figure C4, we show that this pattern persists (i) when we also allow for time-varying pass-through from oil to retail gasoline prices and construct the instruments accordingly, (ii) when we drop the first three months of 2015, (iii) when we use a narrower rolling window (16 years), and (iv) even for the OLS specification.

The decline in the absolute value of the price elasticity of gasoline demand is not unexpected, given the increase in the average fuel efficiency of light vehicles since 2006, the decline in gasoline expenditures as a share of total consumer expenditures, and the rise in urbanization. It should be noted, however, that these changes have tended to be smoother than the decline in the rolling window elasticity estimates. In contrast, there is no empirical support for other possible explanations of lower gasoline demand elasticities discussed in the literature such as an increase in the share of dual income households or a substantial increase in the fraction of households with multiple vehicles.

While we are not the first to document the time variation in the price elasticity of gasoline demand, ours is the most comprehensive analysis of this question to date. Our estimates inform the debate about a secular decline in the gasoline demand elasticity to near zero in the early 2000s. Whereas Hughes et al.'s (2008) analysis based on aggregate U.S. gasoline price and consumption data until 2006 has been commonly interpreted as suggesting that the gasoline demand elasticity since the 1980s has dropped to near zero in absolute value,

our panel data IV estimates do not support that view.

7 Additional Evidence and Robustness

In this section, we provide additional evidence and robustness checks that support our identification strategy and empirical results. We show that individual states or major oil producing states collectively do not affect our estimates. Nor does excluding periods of oil and gasoline supply disruptions change our estimates. This helps alleviate the concern that demand or supply shifts in certain U.S. states may drive global oil price fluctuations and at the same time directly affect gasoline consumption. More importantly, we show that applying our identification strategy to U.S. city-level data gives very similar elasticity estimates, supporting the validity and broad applicability of our approach.

MSA-level evidence. To further alleviate the concern that state-level data may not adequately address the concern of endogeneity due to the large influence of some states on the national economy, we provide additional evidence using city-level data. We obtain gasoline prices and citywide gasoline expenditures for 241 MSAs from 2006 to 2009 from [Levin et al. \(2017\)](#).¹⁹ In their data, retail gasoline prices are provided by OPIS, and the expenditure data are from card transactions of Visa debit and credit card users.

These expenditure data have the advantage of directly measuring consumers' purchases at the gas station, but they are not without drawbacks. For example, gasoline prices and consumption conditional on Visa card users may not be representative for city-level gasoline prices and consumption. Moreover, unlike the state-level analysis, there are limited data at the city level that allow us to link the systematic pass-through to the cost structure of retail gasoline, complicating the interpretation. However, these data help with the identification because it is even harder to argue that global oil price fluctuations are driven by economic conditions in U.S. cities where the systematic pass-through is high.

For the data to be comparable with our state-level data, we transform the original daily data to monthly frequency by taking the average of daily gasoline prices in a city within a month. Monthly gasoline consumption is constructed by averaging daily quantities purchased

¹⁹The data are from the [Levin et al. \(2017\)](#) replication package, which is available from <https://www.aeaweb.org/articles?id=10.1257/pol.20140093>. The cities and states in the MSA-level data provided by the authors have been anonymized.

at the pump per transaction (i.e., by dividing expenditures by the daily gasoline price). We then estimate the systematic pass-through from oil prices to tax-included gasoline prices (the only available gasoline price measure in the data) for city c located in state i , $\hat{\theta}_{c,i}$. The corresponding instrument is $\hat{\theta}_{c,i}\Delta p_t^O$.

Table 8 presents the results. The OLS estimate controlling for city and month fixed effects is -0.29 (column 1), and the IV estimate is -0.28 (column 2). The similarity between the two suggests that endogeneity is not a big concern in MSA-level data. In column (3), we control for city and state-by-month fixed effects and obtain an elasticity estimate of -0.36. The estimate is essentially unchanged when we weight the regression by the population of cardholders in a city to account for heterogeneity across cities. Overall, we obtain very similar IV estimates to those in the state-level analysis using a very different data source and a shorter estimation period, suggesting that our strategy has broad applicability and that our elasticity estimates are robust.

U.S. supply disruptions. One important identifying assumption underlying our empirical strategy is that monthly fluctuations in global oil prices are not driven by economic conditions in U.S. states where gasoline prices are more responsive to oil price changes. One may be concerned that supply disruptions in major oil- or gasoline-producing states can affect global oil prices and, at the same time, create temporary shortages that reduce gasoline consumption, biasing the elasticity estimate downward (in absolute value). As shown in Appendix A, supply disruptions do not have much explanatory power for oil price fluctuations. The type of disruptions faced by U.S. producers (mainly shutdowns due to hurricanes) have even lower explanatory power. One way to explicitly address this concern is to drop periods of U.S. oil and gasoline supply disruptions. Column (2) in Table C1 presents the IV estimate excluding periods of hurricanes that made landfall in Texas or Louisiana and caused significant supply disruptions, as identified by Coyle et al. (2012).²⁰ The estimates are unchanged from our baseline estimates.

The role of oil states. Another concern is that oil producing states may change their production levels as an endogenous response to their state economy. If these changes affect

²⁰These events include Hurricane Andrew in 1992, Hurricane Bret in 1999, Hurricane Lili in 2002 and Hurricanes Katrina and Rita in 2005.

the global oil price, our elasticity estimate could be biased in either direction. We address this concern by assessing the impact of excluding major oil states. According to the EIA, until 2011, the largest oil producing states were Texas, Alaska and California. Since the shale boom started, production in North Dakota and New Mexico has grown dramatically. These two states together with Texas have been the top three oil producers since 2018. In addition, Louisiana has also been a major oil producing state historically. Columns (3)-(5) in Table C1 present the results excluding the largest oil state (TX), the largest three oil states until 2017 (TX, AK and CA), and all major oil states (TX, AK, CA, NM, ND, and LA). The estimates are essentially unchanged from our baseline estimate.

8 Applications and Policy Experiments

The price elasticity of gasoline demand matters for a range of economic applications. In this section, we consider a number of empirical applications and policy experiments: the effects of federal and state gasoline tax holidays, the demand destruction caused by surging gasoline prices after the invasion of Ukraine, and the impact of higher gasoline prices in general, and a carbon tax in particular, on carbon emissions.

8.1 Gasoline Tax Holiday

Reliable estimates of the short-run price elasticity of gasoline demand are useful for evaluating federal and state tax holidays. These policies are often proposed when gasoline prices rise dramatically. For example, in June 2002, President Biden proposed to temporarily suspend the federal gasoline tax.²¹ Some states also implemented their own gasoline tax holidays in 2022 (e.g., New York, Maryland, Georgia and Connecticut).

We evaluate the impact of a three-month suspension of the federal gasoline tax of 18.4 cents (as proposed by President Biden) starting from June 2022, and the impact of a seven-month suspension of the state gasoline tax (as implemented by New York) for all states starting from June 2022. The key to our evaluation is the change in retail gasoline prices due to the tax suspension. In a perfectly competitive market, the fraction of the tax change that is passed on to consumers is given by $e^s/(e^s - e^d)$, where e^s denotes the own-price elasticity of gasoline supply and e^d the own-price elasticity of gasoline demand

²¹See the White House statement “*President Biden Calls for a Three-Month Federal Gas Tax Holiday*”, June 22, 2022.

(see [Doyle and Samphantharak \(2008\)](#)). The only formal estimate of e^s in the literature we are aware of is the IV estimate of 0.29 in [Coyle et al. \(2012\)](#). Given e^s , we show how three alternative estimates of e^d change the expected aggregate and distributional impacts of these tax policies: (1) the midpoint of the widely cited estimates of [Hughes et al. \(2008\)](#), $\hat{e}^d = -0.05$, (2) our IV estimate using the 20-year rolling window ending in March 2022, $\hat{e}^d = -0.2$, and (3) our IV estimates for high-income states, $\hat{e}_{high\ inc}^d = -0.18$, and for low income states, $\hat{e}_{low\ inc}^d = -0.24$.

8.1.1 Federal Gasoline Tax Holiday

Table 9 presents the expected impacts of the federal tax holiday proposal for alternative demand elasticity estimates. The difference is substantial. Whereas the traditional elasticity estimate of [Hughes et al. \(2008\)](#) implies that 85.3% of the tax change (or 15.7 cents) is passed on to retail gasoline prices, our IV estimate implies a much more modest pass-through of 59% (or 11 cents). Given the gasoline price of \$4.97 per gallon at the time, this implies a 3.2% and a 2.2% fall, respectively, in the gasoline price, and a 0.16% and a 0.44% increase in the quantity consumed. According to the FHWA, the average car in the U.S. consumes about 600 gallons of motor fuel per year (or 50 gallons per month) and the average household owns two cars. These statistics imply that the expected change in gasoline expenditures for the average household resulting from a three-month federal tax holiday is \$45, if $\hat{e}^d = -0.05$, and is only \$26, if $\hat{e}^d = -0.2$.

Overall, our analysis shows that the proposed federal tax holiday generates little savings for households. Of course, there will be differences across states. In lower-income states, the change in gasoline consumption in response to the tax stimulus is larger, the pass-through of the tax reduction is lower, and the resulting savings are even smaller, since the gasoline demand elasticity tends to be higher. In contrast, states with higher real personal income have less elastic gasoline consumption, suggesting that households in these states would save more from the tax suspension. For example, using the elasticity estimates based on the extended data, the savings in high-income states would be about \$28 per household, whereas in low-income states, the savings from gasoline expenditures would only be \$23.

8.1.2 State Gasoline Tax Holiday

Figure 8 shows the expected effect of a state gasoline tax holiday on consumers' discretionary income for alternative demand elasticity values, given each state's gasoline price and taxes in June 2022. On average, households save \$162 on gasoline spending with $\hat{\epsilon}^d = -0.05$, and \$95 with $\hat{\epsilon}^d = -0.2$. These effects are larger than under the proposed federal tax holiday because of the longer implementation period we assumed and because of the larger amounts of state gasoline taxes (28.5 cents on average).

While the average effect of state gasoline tax holidays is still moderate using our IV estimate, the distributional effect varies substantially. Households in states where gasoline taxes are high are expected to save more. In Pennsylvania, for example, the average household could save \$194, equivalent to one month of spending on gasoline for a typical car (if gasoline costs \$4 per gallon). In contrast, households in Alaska, Hawaii, Virginia, Missouri and Arizona benefit the least from a state gasoline tax holiday (less than \$60). Based on the traditional estimate of the gasoline demand elasticity of -0.05, one would have concluded that the state tax holiday is much more effective in boosting discretionary income than it really is. For example, one would have concluded that households in Pennsylvania save as much as \$330, and for other states (except for Alaska), households save \$90-\$290.

8.2 How Much Demand Destruction Has Been Associated with Rising Gasoline Prices Since February 2022?

Following the invasion of Ukraine in late February 2022, U.S. retail gasoline prices surged by 20% in March 2022. It is widely understood that these price increases, all else equal, must have reduced U.S. gasoline consumption. Accurate estimates of the short-run gasoline demand elasticity are important for approximating the causal effect of higher gasoline prices. Here we focus on the demand destruction taking place in March and April 2022, which most clearly can be associated with the spike in oil prices caused by the invasion.

As of February 2022, according to the EIA, 8,228 thousand barrels of gasoline were consumed by the U.S. transportation sector. By weighting the monthly percent changes in the retail price of regular gasoline in March and April, as reported by the EIA, by the time-varying IV elasticity estimate of -0.2 obtained using the data ending in March 2022, we

can infer how much gasoline consumption must have declined in response to higher gasoline prices. We estimate a cumulative decline by 286 thousand barrels or 3.48% by April 2022.

It should be noted that there is no reason for actual U.S. gasoline consumption to have declined as much as predicted or, for that matter, at all. Indeed, the actual data show a cumulative increase in U.S. gasoline consumption by 149 thousand barrels or 0.18% from February to April, consistent with a largely flat path for gasoline consumption. The reason is that one would have expected a strong seasonal increase in gasoline demand in the first half of the year. For example, during 2019, gasoline consumption increased by 3.96% from February to April largely due to seasonal demand fluctuations. Assuming a similar pattern in 2022, we would expect a modest increase in gasoline consumption of 0.48% (by subtracting the predicted decline of 3.48% from the expected seasonal increase of 3.96%). This prediction is quite close to the observed increase of 0.18%, indicating that other shocks did not play an important role in driving U.S. gasoline consumption in March and April 2022.

8.3 How Much Would A 10% Increase in Retail Gasoline Prices Lower U.S. and Global Carbon Emissions?

There has also been interest in understanding the impact of higher gasoline prices on U.S. carbon emissions. We address this question using an approach similar to the analysis of higher gasoline taxes in [Davis and Kilian \(2011\)](#). Consider a 10% increase in retail gasoline prices. Then the impact on U.S. emissions may be approximated as the percent change in gasoline prices times the price elasticity of gasoline demand, weighted by the share of the carbon emissions from motor gasoline in total U.S. carbon emissions. This share is 0.21, as may be inferred from data in the EIA's *Monthly Energy Review* for October 2021 (Tables 11.2-11.5). The impact on global carbon emissions may be approximated by further weighting this impact by the share of U.S. carbon emissions in global carbon emissions, which in 2020 was 13%, according the EIA's *International Energy Outlook 2021*.

Our earlier analysis shows that the short-run price elasticity of gasoline demand using data ending in March 2022 is -0.2. This implies that, in the short run, a 10% gasoline price increase would lower U.S. carbon emissions by 0.42% and global carbon emissions by 0.05%, which is negligible compared to the rate of growth in emissions.

A closely related question is what the effect of a carbon tax would be on carbon emissions. Given our evidence that the tax elasticity of gasoline demand is close to the own price elasticity of gasoline demand, we may also use our elasticity estimate to explore how the imposition of a carbon tax would affect carbon emissions from motor gasoline. Many representatives and senators in the U.S. Congress have recently proposed legislation authorizing a federal carbon tax. A study of ten such policy proposals by the Center on Global Energy Policy at the School of International and Public Affairs at Columbia University concludes that a \$145/ton carbon tax would raise the U.S. price of gasoline by \$1.27 per gallon.²² Imposing such a carbon tax would raise the gasoline price by 32%, given a \$4 per gallon gasoline price. The resulting fall in gasoline consumption would lower U.S. carbon emissions by 1.3% on impact and global carbon emissions by 0.2%.

9 Conclusion

Our analysis is part of a large literature on estimating the responsiveness of gasoline consumption to changes in gasoline prices that dates back to the 1970s. This question has always been considered important from a policy point of view. The recent debate about gasoline tax holidays is a case in point. Knowledge of this elasticity also matters for modeling the transmission of gasoline price shocks to the domestic economy and for structural models of automobile demand. The growing interest in carbon emissions from gasoline-powered engines and in the design of regulatory and tax policies that aim to correct externalities from vehicle use have further strengthened interest in this key parameter.

At the same time, there has not been much progress in developing methods for estimating the price elasticity of gasoline demand in recent years. With few exceptions, researchers continue to use methods that were designed many years ago. In this paper, we proposed a new class of instruments that differs fundamentally from earlier estimation methods and is easy to implement in applied work. Compared with the IV approach in [Coglianese et al. \(2017\)](#), our IV estimator tends to produce elasticity estimates with much lower standard errors, and it may be applied even in the absence of state or city-level data on excise taxes.

Our baseline estimate of -0.31 is close to recent estimates that address the endogeneity

²²See, <https://www.energypolicy.columbia.edu/what-you-need-know-about-federal-carbon-tax-united-states>.

of gasoline price changes, but has a standard error of 0.07 only. The corresponding elasticity estimates based on the monthly averages of the daily city-level data used in [Levin et al. \(2017\)](#) are between -0.28 and -0.37, depending on the specification. We illustrated the implications of these elasticity estimates for policy analysis. We also provided evidence that gasoline demand is more responsive in states with lower personal income, higher unemployment rates and lower urban population shares, for example, suggesting caution in applying aggregate elasticity estimates indiscriminately to all regions.

There has been much debate about the gasoline demand elasticity having declined in absolute value to near zero since the 1980s. Our analysis does not support this view. Rolling windows estimates show that the one-month price elasticity of gasoline demand remained stable near -0.3 from the 2000s to the end of 2014. Although the elasticity has come down in absolute terms since then, it is still near -0.2 in early 2022, showing more responsiveness in gasoline demand than traditional estimates. Nor did we find compelling evidence for an asymmetric response to positive and negative gasoline price shocks in monthly data, suggesting that the linear specification is adequate.

While we focused on estimating the price elasticity of gasoline demand, our approach also has direct implications for the literature on understanding the responsiveness of vehicle miles traveled to gasoline prices. Understanding the response of consumers' driving behavior to price signals is of first-order importance for modeling the transition to electric vehicles for example. Finally, our analysis addressed and corrected the view that the tax elasticity of gasoline demand systematically exceeds the pre-tax price elasticity of gasoline demand.

References

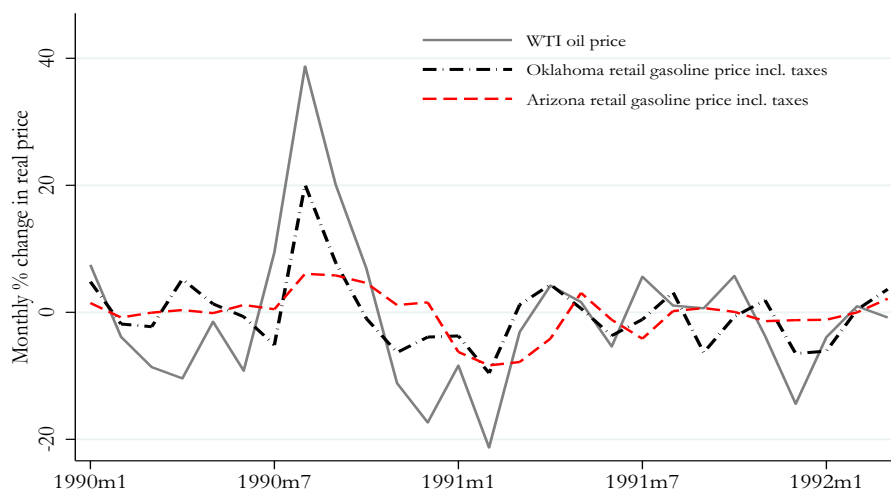
- ADAO, R., M. KOLESAR AND E. MORALES, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics* 134(4) (2019), 1949–2010.
- ALLCOTT, H. AND N. WOZNY, "Gasoline Prices, Fuel Economy, and the Energy Paradox," *Review of Economics and Statistics* 96(5) (2014), 779–795.
- ANDERSON, S., R. KELLOGG AND J. SALLEE, "What Do Consumers Believe about Future Gasoline Prices?," *Journal of Environmental Economics and Management* 66(3) (2013), 383–403.
- BENTO, A. M., L. H. GOULDER, M. R. JACOBSEN AND R. H. VON HAEFEN,

- “Distributional and Efficiency Impacts of Increased US Gasoline Taxes,” *American Economic Review* 99(3) (2009), 667–699.
- BORENSTEIN, S., A. C. CAMERON AND R. GILBERT, “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?,” *Quarterly Journal of Economics* 112(1) (1997), 305–339.
- BORUSYAK, K., P. HULL AND X. JARAVEL, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies* 89(1) (2022), 181–213.
- COGLIANESE, J., L. W. DAVIS, L. KILIAN AND J. H. STOCK, “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand,” *Journal of Applied Econometrics* 32(1) (2017), 1–15.
- COYLE, D., J. DEBACKER AND R. PRISINZANO, “Estimating the Supply and Demand of Gasoline Using Tax Data,” *Energy Economics* 34(1) (2012), 195–200.
- DAHL, C., “Consumer Adjustment to a Gasoline Tax,” *Review of Economics and Statistics* 61(3) (1979), 427–432.
- DAHL, C. AND T. STERNER, “Analysing Gasoline Demand Elasticities: A Survey,” *Energy Economics* 13(3) (1991), 203–210.
- DAVIS, L. W. AND L. KILIAN, “Estimating the Effect of A Gasoline Tax on Carbon Emissions,” *Journal of Applied Econometrics* 26(7) (2011), 1187–1214.
- DELTAS, G., “Retail Gasoline Price Dynamics and Local Market Power,” *Journal of Industrial Economics* 56(3) (2008), 613–628.
- DOYLE, J. J. J. AND K. SAMPHANTHARAK, “\$2.00 Gas! Studying the effects of a Gas Tax Moratorium,” *Journal of Public Economics* 92(3-4) (2008), 869–884.
- EDELSTEIN, P. AND L. KILIAN, “How Sensitive Are Consumer Expenditures to Retail Energy Prices?,” *Journal of Monetary Economics* 56(6) (2009), 766–779.
- GOLDSMITH-PINKHAM, P., I. SORKIN AND H. SWIFT, “Bartik Instruments: What, When, Why, and How,” *American Economic Review* 110(8) (2020), 2586–2624.
- GUREN, A. M., A. MCKAY, E. NAKAMURA AND J. STEINSSON, “Housing Wealth Effects: The Long View,” *Review of Economic Studies* 88(2) (2021), 669–707.
- HAUSMAN, J. A., *Valuation of New Goods under Perfect and Imperfect Competition* (in T. Bresnahan and R. Gordon eds. *The Economics of New Goods*, University of Chicago Press, Chicago, chapter 5, 1996).
- HAUSMAN, J. A. AND W. K. NEWEY, “Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss,” *Econometrica* 63(6) (1995), 1445–1476.
- HUGHES, J. E., C. R. KNITTEL AND D. SPERLING, “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand,” *The Energy Journal* 29(1) (2008), 113–134.

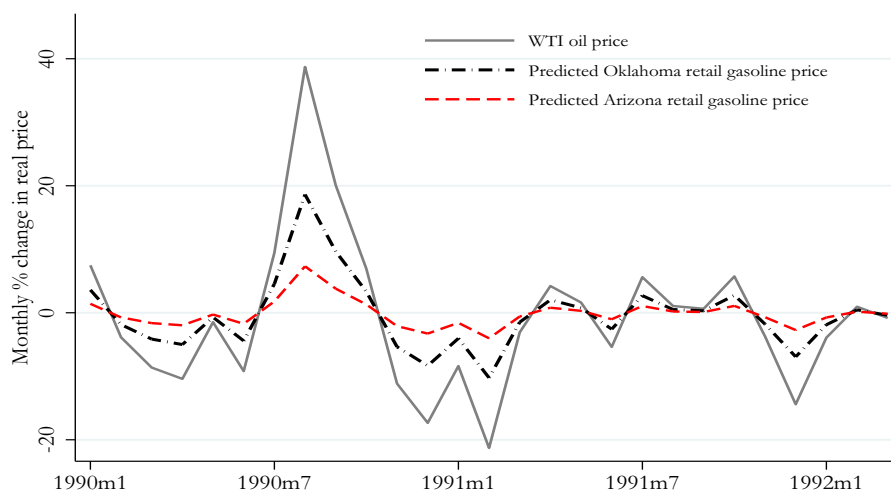
- KILIAN, L., “The Economic Effects of Energy Price Shocks,” *Journal of Economic Literature* 46(4) (2008a), 871–909.
- , “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?,” *Review of Economics and Statistics* 90(2) (2008b), 216–240.
- , “Understanding the Estimation of Oil Demand and Oil Supply Elasticities,” *forthcoming: Energy Economics* (2022).
- KILIAN, L. AND R. J. VIGFUSSON, “Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?,” *Quantitative Economics* 2(3) (2011), 419–453.
- KILIAN, L. AND X. ZHOU, “The Propagation of Regional Shocks in Housing Markets: Evidence from Oil Price Shocks in Canada,” *Journal of Money, Credit and Banking* 54(4) (2022), 953–987.
- KNITTEL, C. R. AND S. TANAKA, “Fuel Economy and the Price of Gasoline: Evidence from Fueling-Level Micro Data,” *forthcoming: Journal of Public Economics* 202 (2021).
- LEVIN, L., M. S. LEWIS AND F. A. WOLAK, “High Frequency Evidence on the Demand for Gasoline,” *American Economic Journal: Economic Policy* 9(3) (2017), 314–347.
- , “Reference Dependence in the Demand for Gasoline,” *Journal of Economic Behavior and Organization* 197 (2022), 561–578.
- LEWIS, M. AND M. NOEL, “The Speed of Gasoline Price Response in Markets with and without Edgeworth Cycles,” *Review of Economics and Statistics* 93(2) (2011), 672–682.
- LI, S., J. LINN AND E. MUEHLEGGGER, “Gasoline Taxes and Consumer Behavior,” *American Economic Journal: Economic Policy* 6(4) (2014), 302–347.
- MONTIEL OLEA, J. L. AND C. PFLUEGER, “A Robust Test for Weak Instruments,” *Journal of Business and Economic Statistics* 31(3) (2013), 358–369.
- NAKAMURA, E. AND J. STEINSSON, “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review* 104(3) (2014), 753–792.
- RAMSEY, J., R. RASCHE AND B. ALLEN, “An Analysis of the Private and Commercial Demand for Gasoline,” *Review of Economics and Statistics* 57(4) (1975), 502–507.
- RIVERS, N. AND B. SCHAUFELE, “Salience of Carbon Taxes in the Gasoline Market,” *Journal of Environmental Economics and Management* 74 (2015), 23–36.
- SWEENEY, J. L., “The Response of Energy Demand to Higher Prices: What Have We Learned?,” *American Economic Review* 74(2) (1984), 31–37.
- VENDITTI, F., “From Oil to Consumer Energy Prices: How Much Asymmetry along the Way?,” *Energy Economics* 40(C) (2013), 468–473.
- WOOLDRIDGE, J. M., *Econometric Analysis of Cross Section and Panel Data* (The MIT Press, Cambridge, Massachusetts, 2001).

Figure 1: Actual and predicted gasoline price changes around the invasion of Kuwait

(a) Data

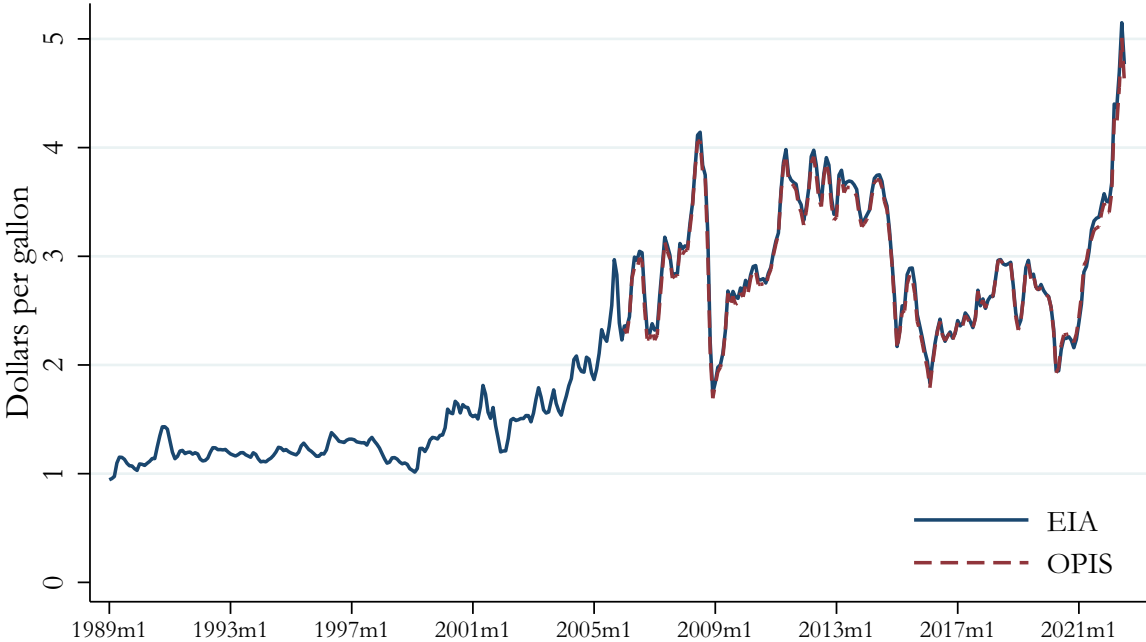


(b) Prediction



Notes: The predicted gasoline price change is constructed by multiplying the percent change in the real oil price by the systematic pass-through from oil to tax-included gasoline prices, $\hat{\theta}_i$.

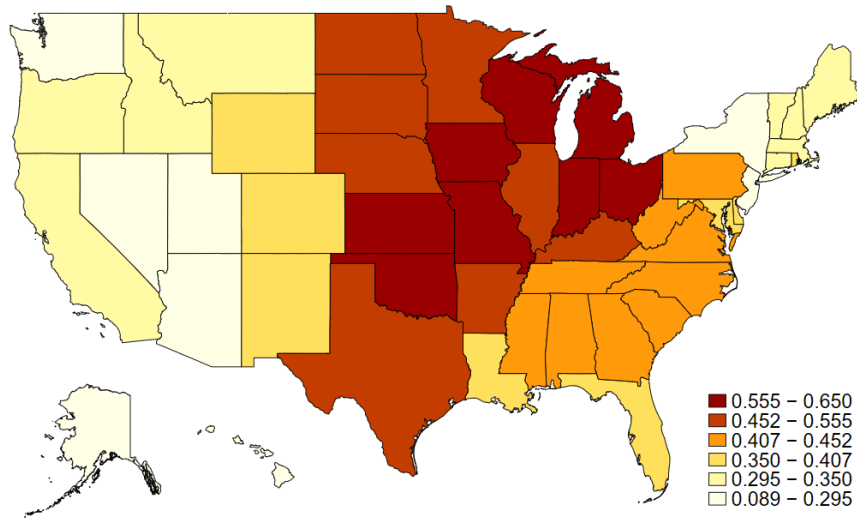
Figure 2: U.S. retail gasoline prices including taxes



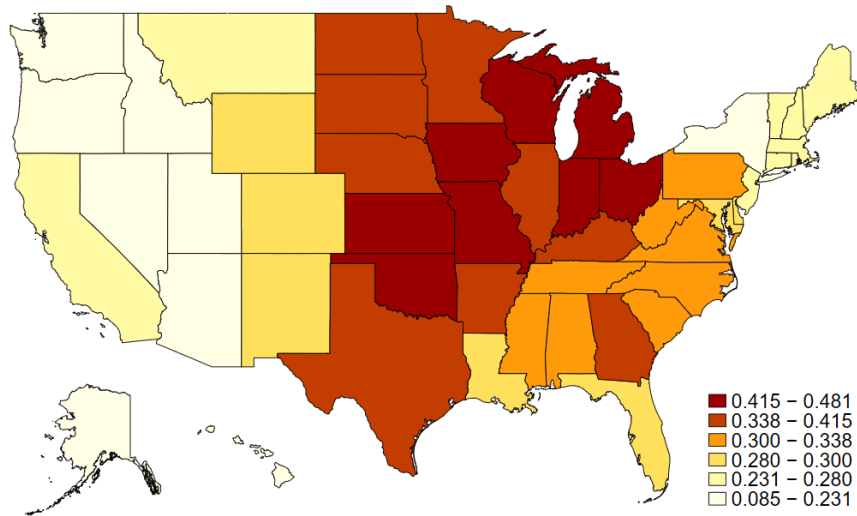
Source: EIA *Monthly Energy Review* U.S. city average retail price of all grades of gasoline; OPIS.

Figure 3: Systematic pass-through from oil to gasoline prices

(a) Pre-tax gasoline price response, $\hat{\eta}_i$

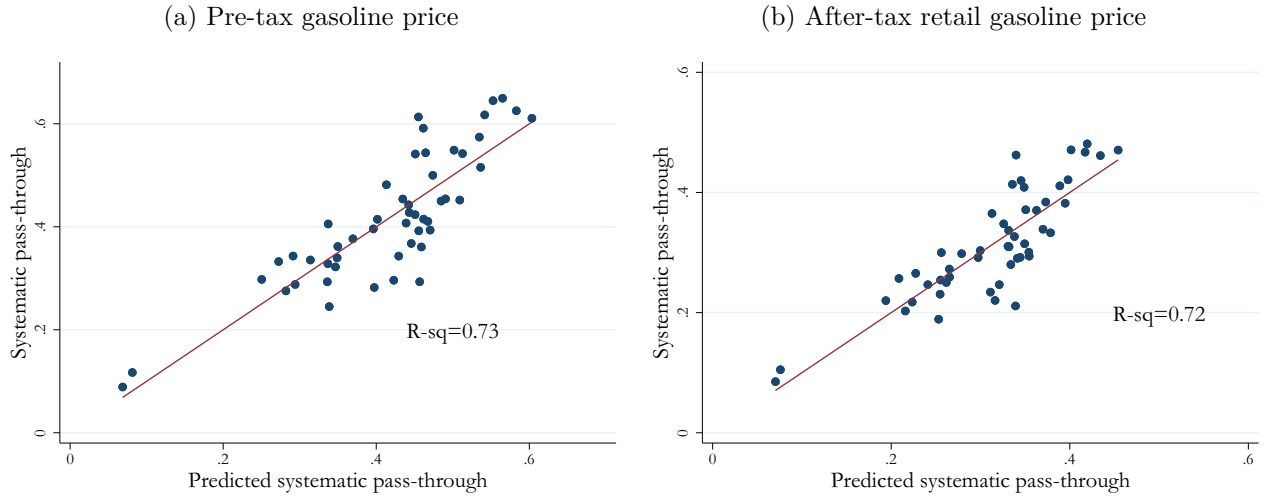


(b) After-tax retail gasoline price response, $\hat{\theta}_i$



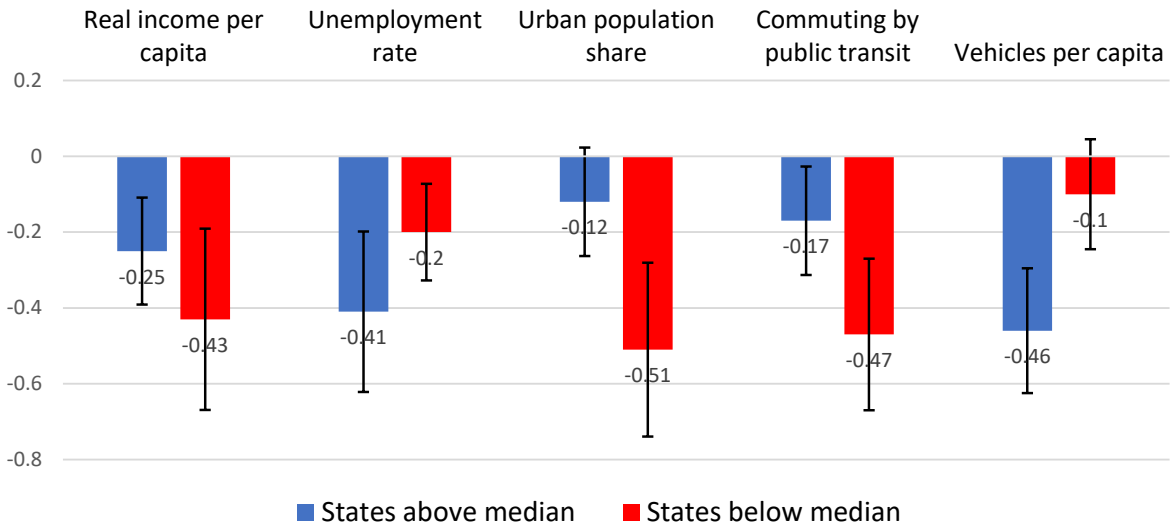
Notes: $\hat{\eta}_i$ and $\hat{\theta}_i$, $i \in I$, are estimated using equations (7) and (6), respectively.

Figure 4: Systematic pass-through and the predicted value using cost measures



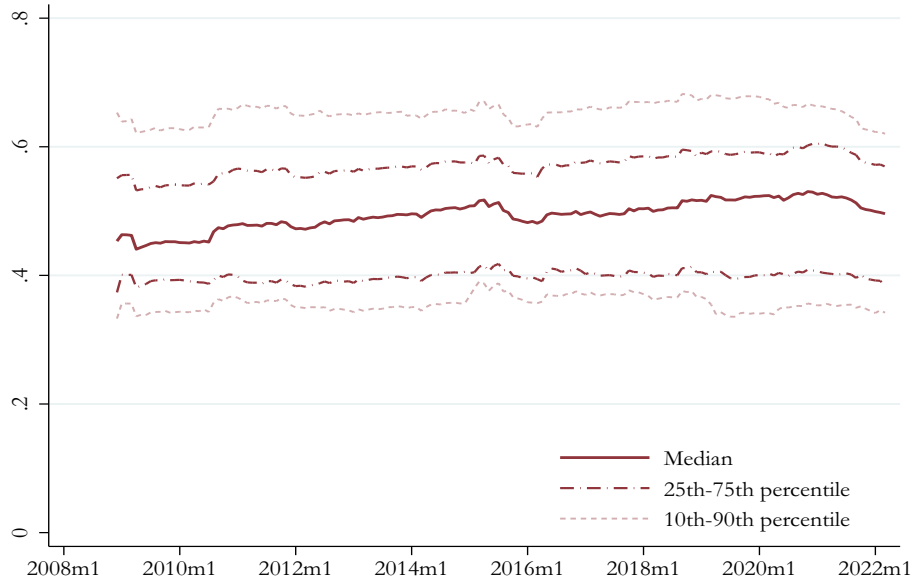
Notes: The predicted values of the systematic pass-through ($\hat{\eta}_i$ and $\hat{\theta}_i$) are obtained by estimating the specification in column (10) of Table 1.

Figure 5: Heterogeneity in the price elasticity of gasoline demand



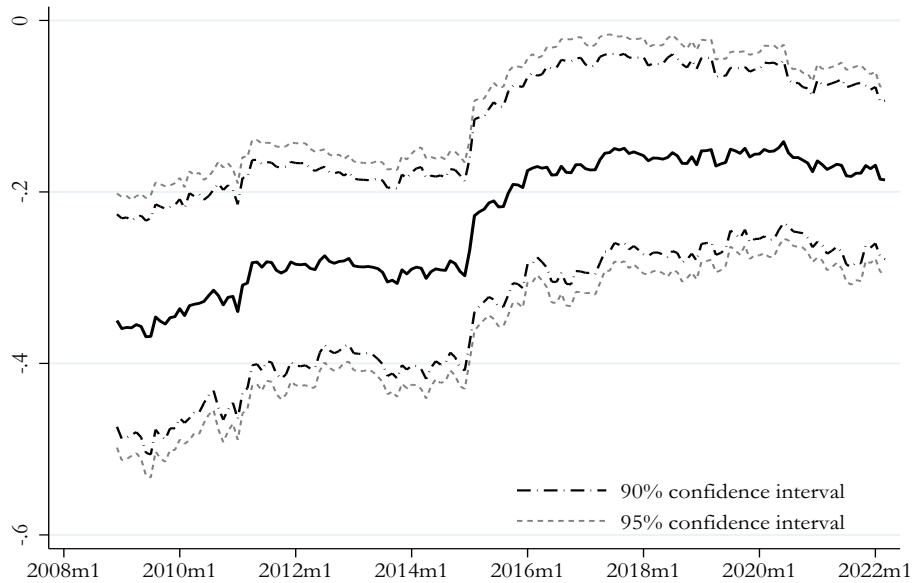
Notes: Point estimates and 95% confidence intervals under the same model specification as for the baseline using the pass-through from oil to pre-tax retail gasoline prices.

Figure 6: Time variation in the systematic pass-through, $\hat{\eta}_{i,t}$



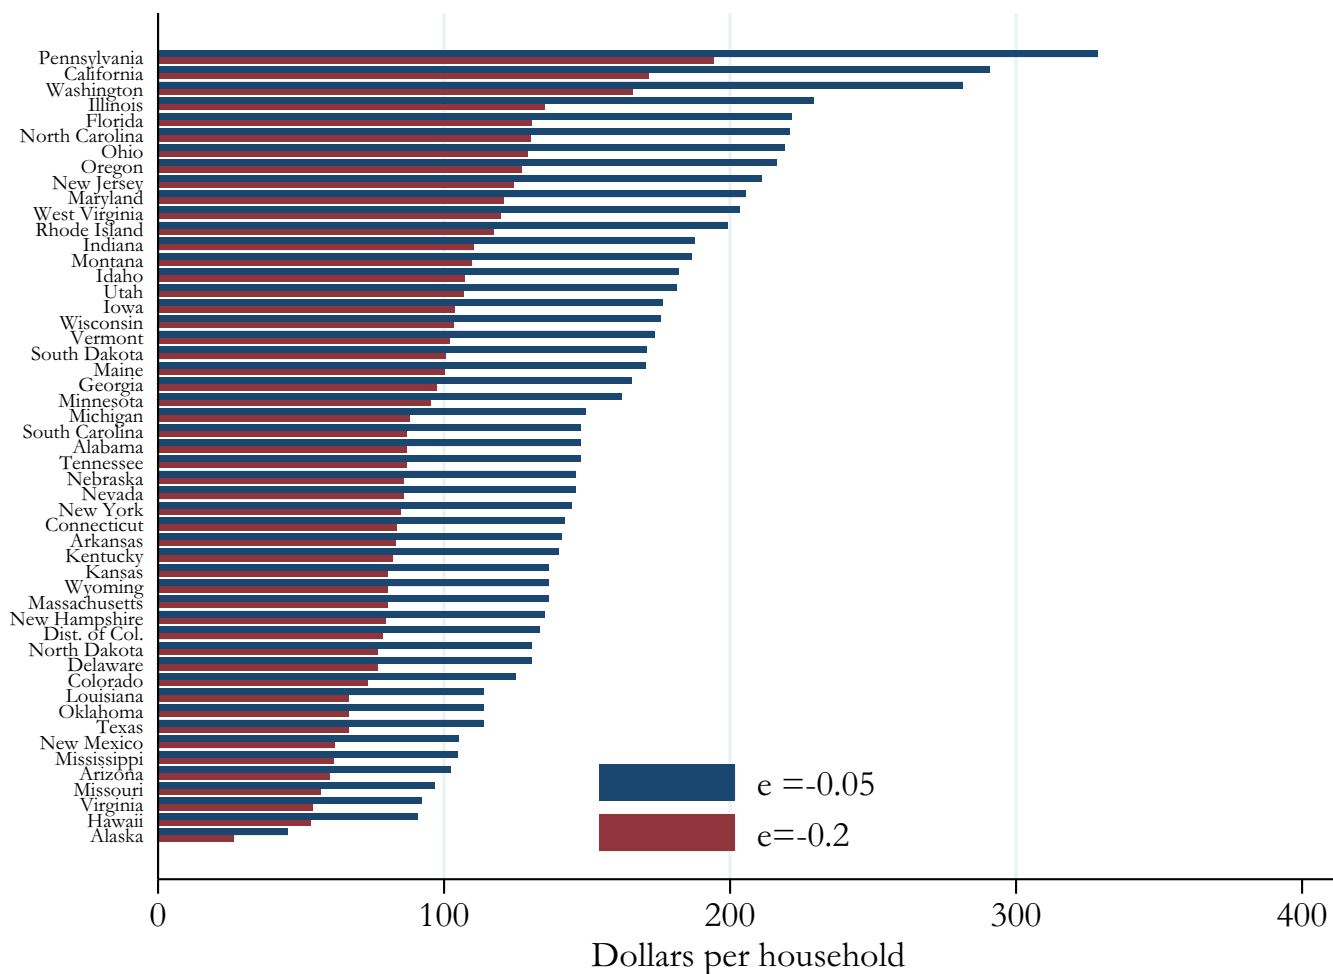
Notes: Rolling-window estimates. At each point of the horizontal axis, the distributional statistics (the median, 25th-75th percentile range, and 10th-90th percentile range) are obtained from estimating the panel regression (equation 7) using the 20 years of data ending at that point.

Figure 7: Time-varying price elasticity of gasoline demand



Notes: Rolling-window estimates. At each point of the horizontal axis, the point estimate and confidence intervals are obtained from estimating the baseline IV specification using the 20 years of data ending at that point.

Figure 8: Effect of a state tax holiday on household discretionary income



Notes: The figure shows the effect of a seven-month state gasoline tax holiday starting from June 2022 on the discretionary income of the average household who has two cars.

Table 1: Determinants of the systematic pass-through from oil to pre-tax gasoline prices ($\hat{\eta}_i$)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Distance to Cushing	-9.403*** (1.384)								-5.694*** (1.236)	-6.475*** (1.045)	-5.426*** (0.794)	-7.406*** (1.458)
Oil pipeline density		5.250*** (1.679)							1.115 (1.238)			
API gravity (light oil)			8.030*** (1.430)						3.085** (1.469)	3.014** (1.274)		2.620* (1.321)
RFG exposure				-3.370 (3.868)					-6.252** (2.962)	-7.426*** (2.764)	-4.991** (2.027)	-9.848*** (3.577)
Product pipeline density					6.415*** (1.882)				4.282*** (1.284)	4.324*** (1.138)	2.188** (0.825)	4.042*** (1.172)
Local transportation costs						-1.286 (2.025)			-0.674 (1.720)			
Retail-wholesale margin							-0.962 (1.894)		-3.056** (1.244)	-2.524** (1.057)	-2.635*** (0.646)	-2.398** (1.150)
Retail station density								3.426*** (1.604)	0.430 (1.274)			
PADD fixed effects	No	No	No	No	No	No	No	No	No	No	Yes	No
Demand control	No	No	No	No	No	No	No	No	No	No	No	Yes
R^2	0.54	0.17	0.40	0.01	0.25	0.01	0.01	0.07	0.74	0.73	0.90	0.74
R^2 leave-out	0.65	0.73	0.70	0.72	0.68	0.74	0.70	0.74	-	-	-	-
# Obs.	51	51	51	51	51	51	51	51	51	51	51	51

Notes: ** and *** denote significance at the 5% and 1% level, respectively.

Table 2: IV estimates using variation in the systematic pass-through from oil to gasoline prices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	$\hat{\eta}_i \Delta p_t^O$ as IV	$\hat{\eta}_i \Delta p_t^O$ as IV	$\hat{\theta}_i \Delta p_t^O$ as IV	$\hat{\eta}_i \Delta p_t^O$ predicted by cost measures as IV	$\hat{\theta}_i \Delta p_t^O$ predicted by cost measures as IV	Δp_t^O interacted with cost measures as IV	$\bar{\eta}_{-i} \Delta p_t^O$ as IV
$\Delta p_{i,t}$	-0.190*** (0.037)	-0.314*** (0.066)	-0.373*** (0.100)	-0.379*** (0.102)	-0.334*** (0.105)	-0.338*** (0.105)	-0.357*** (0.108)	-0.379*** (0.111)
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes
1st stage effective F-stat	-	626	423	423	242	240	61	305
5% critical value for MOP weak IV test	-	37.4	37.4	37.4	37.4	37.4	21.5	37.4
# Obs.	11,730	11,730	10,863	10,863	10,863	10,863	10,863	10,437

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level. The effective F-stat is computed according to [Montiel Olea and Pflueger \(2013\)](#).

Table 3: Tax-based IV and systematic pass-through-based IV estimates

	(1)	(2)	(3)
	CDKS tax change as IV	$\hat{\eta}_i \Delta p_t^O$ as IV	$\hat{\eta}_i \Delta p_t^O$ as IV
$\Delta p_{i,t}$	-1.152*** (0.250)	-0.314*** (0.066)	-0.320*** (0.067)
$\Delta \log(\tau_{i,t}^g)$			-0.230*** (0.068)
$\Delta \log(\tau_{i,t+1}^g)$			0.156*** (0.052)
$\Delta \log(\tau_{i,t-1}^g)$			0.061 (0.055)
State fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
# Obs.	11,628	11,730	11,628
Implied price elasticity of gasoline demand	-0.368 (0.239)	-0.314*** (0.066)	-0.320*** (0.067)

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table 4: Gasoline consumption responses to tax shocks and pre-tax price shocks

	(1)	(2)	(3)
	OLS	OLS	$\hat{\eta}_i \Delta p_t^O$ as IV
$\Delta p_{i,t}^{EX}$	-0.130*** (0.027)	-0.130*** (0.027)	-0.230*** (0.049)
$\Delta \tau_{i,t} / p_{i,t-1}^{EX}$	-0.860*** (0.218)	-0.856*** (0.217)	-0.849*** (0.216)
$\Delta \tau_{i,t+1} / p_{i,t}^{EX}$		0.493*** (0.148)	0.488*** (0.143)
$\Delta \tau_{i,t-1} / p_{i,t-2}^{EX}$		0.206 (0.157)	0.209 (0.158)
State fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
# Obs.	11,730	11,628	11,628
P-value for testing equal effects	0.001	0.877	0.690

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table 5: Distribution of p-values for testing asymmetric pass-through

	10th	25th	50th	75th	90th	# of states w. p-value < 0.05	# of states w. p-value < 0.01
P-value: $\eta_i^+ = 0$	0.051	0.071	0.127	0.385	0.852	4	0
P-value: $\theta_i^+ = 0$	0.055	0.078	0.122	0.369	0.833	3	0

Notes: This table shows the distributional statistics for the p-values testing $\eta_i^+ = 0$ and testing $\theta_i^+ = 0$.

Table 6: No evidence for asymmetry in the price elasticity of gasoline demand

	(1) OLS	(2) IV based on $\hat{\eta}_i$	(3) IV based on $\hat{\theta}_i$	(4) IV based on $\hat{\eta}_i$	(5) IV based on $\hat{\theta}_i$
$\Delta p_{i,t}$	-0.238*** (0.048)	-0.443*** (0.103)	-0.447*** (0.102)	-0.280 (0.145)	-0.327** (0.144)
$\Delta p_{i,t}^+$	0.097 (0.058)	0.303 (0.184)	0.314 (0.182)	-0.170 (0.329)	-0.039 (0.334)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes
Asymmetric pass-through	-	No	No	Yes	Yes
# Obs.	11,730	11,730	11,730	11,730	11,730

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table 7: IV estimates using variation in the systematic pass-through from oil to gasoline prices, extended sample from 1989m1 to 2022m3

	(1) OLS	(2) $\hat{\eta}_i \Delta p_t^O$ as IV	(3) $\hat{\eta}_i \Delta p_t^O$ as IV	(4) $\hat{\eta}_i \Delta p_t^O$ as IV	(5) $\hat{\theta}_i \Delta p_t^O$ as IV
$\Delta p_{i,t}$	-0.172*** (0.031)	-0.259*** (0.051)	-0.186*** (0.061)	-0.198*** (0.053)	-0.197*** (0.053)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes
Drop 2020m3 onwards	No	No	Yes	No	No
Drop 2020m3-m5	No	No	No	Yes	Yes
# Obs.	20,298	20,298	19,023	20,145	20,145

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table 8: MSA-level estimates of the price elasticity of gasoline demand

	(1)	(2)	(3)	(4)
	OLS	$\hat{\theta}_{c,i}\Delta p_t^O$ as IV	$\hat{\theta}_{c,i}\Delta p_t^O$ as IV	$\hat{\theta}_{c,i}\Delta p_t^O$ as IV
$\Delta p_{c,i,t}$	-0.286*** (0.006)	-0.278*** (0.021)	-0.359*** (0.070)	-0.367*** 0.087
MSA fixed effects	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	No	No
State-by-month fixed effects	No	No	Yes	Yes
Weighted by cardholder population	No	No	No	Yes
# Obs.	11,086	11,086	10,810	10,810

Data source: [Levin et al. \(2017\)](#). Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the MSA level.

Table 9: Impact of a three-month federal gasoline tax holiday

	$e^d = -0.05$	$e^d = -0.2$	$e_{high\ inc}^d = -0.18$	$e_{low\ inc}^d = -0.24$
Pass-through to consumers	85.3%	59.2%	61.7%	54.7%
Effect on retail gasoline price	-15.7 cents	-10.9 cents	-11.4 cents	-10.1 cents
Effect on gasoline consumption	0.16%	0.44%	0.41%	0.49%
Effect on discretionary income per household	\$44.8	\$26.3	\$28.1	\$23.1

Notes: The federal gasoline tax is 18.4 cents per gallon. Our computations assume that $e^s = 0.29$ based on evidence in [Coyle et al. \(2012\)](#). When computing the effect on discretionary income we assume a two-car household.

Appendices to “Heterogeneity in the Pass-Through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand”

A Conventional Strategies for Estimating the Price Elasticity of Gasoline Demand

This section reviews conventional strategies for estimating the price elasticity of gasoline demand and potential drawbacks of these strategies. The estimation results are summarized in Table A1. Traditionally, the literature has relied on aggregate data and time series variation for estimating this elasticity. Column (1) of panel I in Table A1 shows that the OLS estimate of the elasticity using aggregate data for 1989m1-2008m3 is -0.09.

One proposal for addressing the identification problem is to use oil price changes as the instrument for gasoline price changes on the grounds that oil prices are determined by global demand and supply rather than U.S. economic conditions. This argument ignores the fact that shifts in U.S. demand are often correlated with demand shifts in other advanced economies and hence can affect global oil prices. Indeed, this proposal yields an insignificant elasticity estimate of -0.06, even more attenuated than the OLS estimate, suggesting that unobserved demand shocks are likely biasing the elasticity estimate.

A more refined approach is to use oil supply shocks as the instrument. Kilian (2008a,b) and Kilian (2022), however, have shown that measures of exogenous oil supply shocks developed in the literature have low explanatory power for oil price fluctuations, which makes them weak instruments. Our analysis confirms this conclusion. In column (3), we use one example of these shock measures in applied work: separate dummies for periods of unexpected geopolitical events that led to global oil price swings. These events include the August 1990 invasion of Kuwait, the Venezuelan oil strike in December 2002-February 2003, and the March 2003 invasion of Iraq.²³ Column (3) shows that the elasticity estimate is -0.13, but not statistically significant. The first-stage effective F-statistic is only 8.4, below the 10% critical value for rejecting the null of weak instruments. Some studies also propose

²³We note that these measures capture not only physical disruptions in the oil supply, but also changes in expectations about future oil supply and demand, and hence may have higher explanatory power for oil price changes than measures only capturing supply disruptions.

to include major refinery disruptions caused by hurricanes. In column (4), we expand the set of instruments to include separate dummies for the hurricane events identified by [Coyle et al. \(2012\)](#). The elasticity estimate of -0.11 is again close to zero and the instrument is still weak.

These approaches have also been adapted to panel data settings. A number of studies use global oil price changes (or aggregate wholesale gasoline price changes) as an instrument for state-level or city-level gasoline price changes. The resulting estimate for our sample period is -0.04 (column 2, panel II). There are two obvious problems. First, this strategy cannot control for month fixed effects, and, second, the variation used for identification comes from the time-series dimension rather than the cross-sectional dimension. Hence, this approach is subject to the same identification concerns as using aggregate data. In periods when state-level gasoline consumption is collectively high, it is likely driven by underlying aggregate demand shocks that may also push up global oil prices.

A better approach appears to be the use of oil and gasoline supply shocks as an instrument in the cross-sectional setting. As shown in columns (3) and (4) of panel II, these estimates are precisely estimated and the null of weak instruments is firmly rejected. It is unclear, however, whether the high effective F-statistic is driven by the high explanatory power of these instruments, or the inclusion of state fixed effects in the panel regression. In panel III, we provide one solution to this question. Specifically, we demean state-level gasoline consumption changes and price changes (i.e., $\Delta q_{i,t}$ and $\Delta p_{i,t}$) to account for state fixed effects, and then take the average of the demeaned variables across states in a given month (since the variation used for identification comes from the time-series dimension). We then regress the average demeaned gasoline consumption growth on the average demeaned gasoline price growth, using measures of oil and gasoline supply shocks as the instrument. By construction, this procedure yields the same point estimates as their counterparts in panel II, but it helps evaluate the relevance of the instrument after purging the state fixed effects. As panel III shows, these instruments turn out to be weak as well. Another concern with these oil supply/gasoline supply shock measures is that they may capture unobserved demand shocks due to shifts in expectations or other events at the same time, violating the exclusion restriction.

Table A1: Gasoline demand elasticity estimates using conventional strategies

Panel I. Aggregate data	(1)	(2)	(3)	(4)
	OLS	Oil price changes as IV	Oil supply shocks as IV	Oil/gasoline supply shocks as IV
Δp_t	-0.0934** (0.023)	-0.064 (0.060)	-0.133 (0.068)	-0.112** (0.047)
Month-of-year fixed effects	Yes	Yes	Yes	Yes
1st stage effective F-stat	-	51.1	8.7	13.1
10% critical value	-	23.1	20.2	19.8
5% critical value	-	37.4	33.0	33.0
# Obs.	230	230	230	230
Panel II. State panel data				
		(2)	(3)	(4)
		Oil price changes as IV	Oil supply shocks as IV	Oil/gasoline supply shocks as IV
Average demeaned $\Delta p_{i,t}$		-0.042** (0.021)	-0.126*** (0.024)	-0.124*** (0.017)
Month-of-year fixed effects		Yes	Yes	Yes
State fixed effects		Yes	Yes	Yes
1st stage effective F-stat		629	331	405
# Obs.		11,730	11,730	11,730
Panel III. State average data				
			(3)	(4)
			Oil supply shocks as IV	Oil/gasoline supply shocks as IV
$\Delta p_{i,t}$			-0.126 (0.081)	-0.124** (0.056)
Month-of-year fixed effects			Yes	Yes
1st stage effective F-stat			15.6	21.3
10% critical value			15.7	18.3
5% critical value			25.5	31.0
# Obs.			230	230

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level for state panel regressions. For aggregate data and state average data, the standard errors are corrected for heteroskedasticity and autocorrelation using the Newey-West estimator with 3 lags. The effective F-stat is computed according to [Montiel Olea and Pflueger \(2013\)](#).

B Additional Evidence on the Pass-Through

Figure B1: Heatmaps of selected cost measures

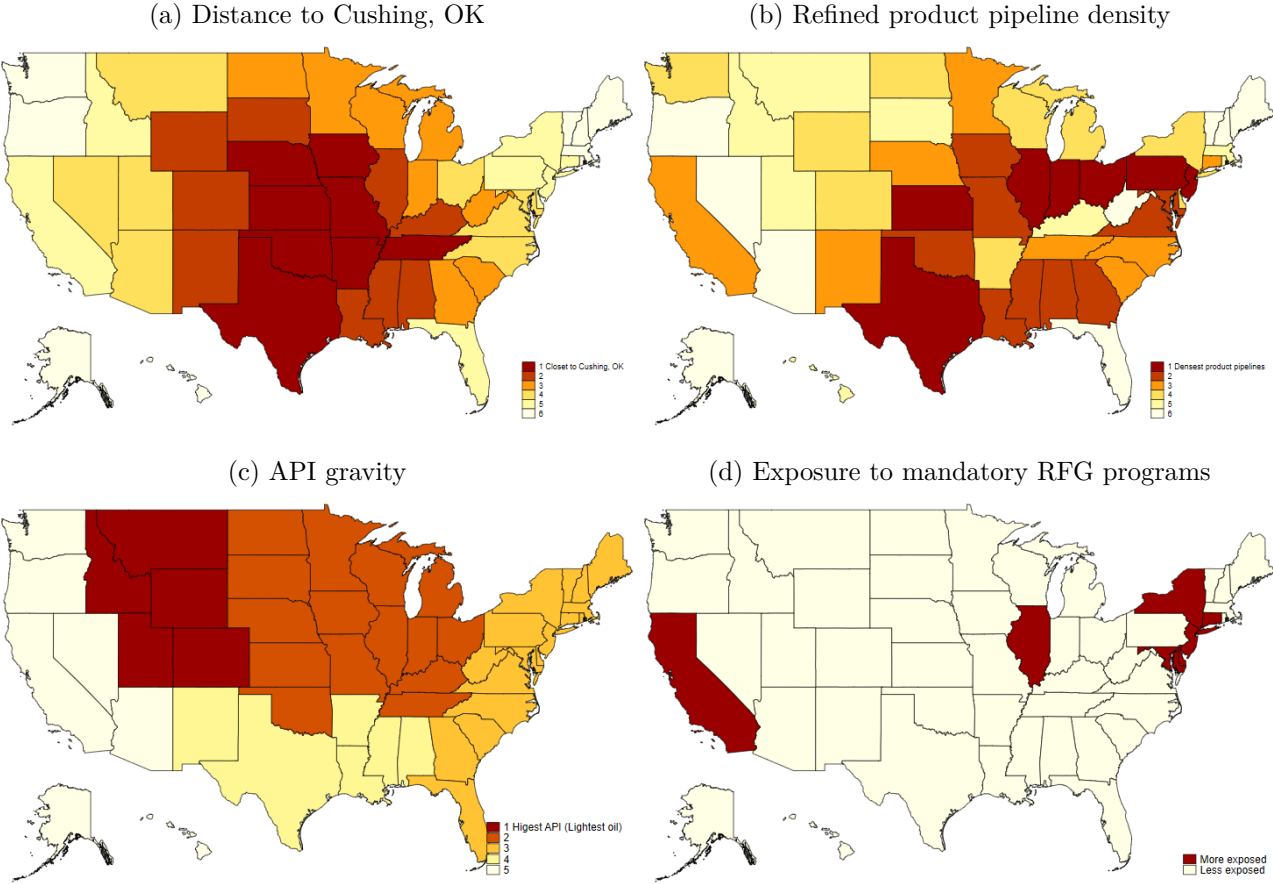
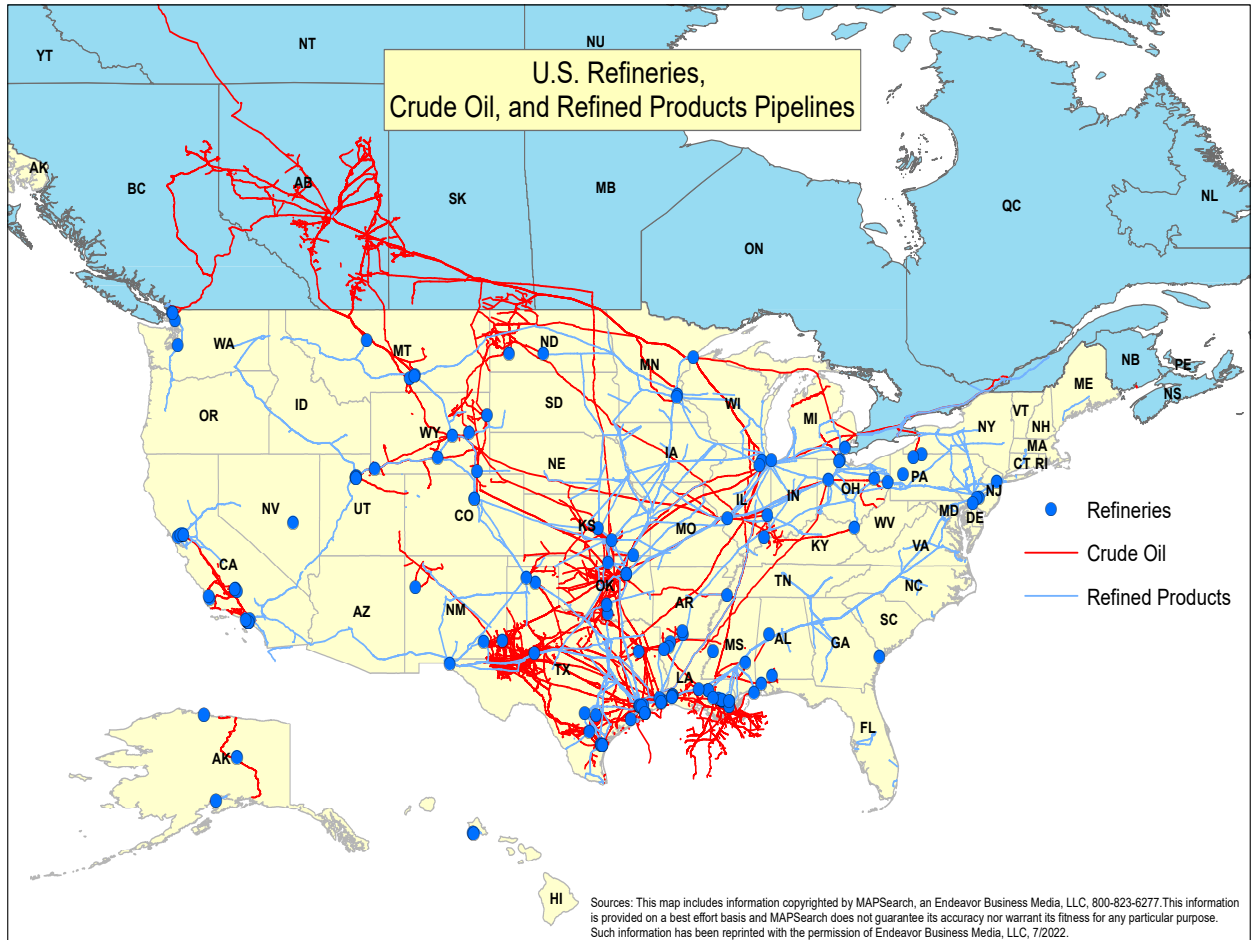
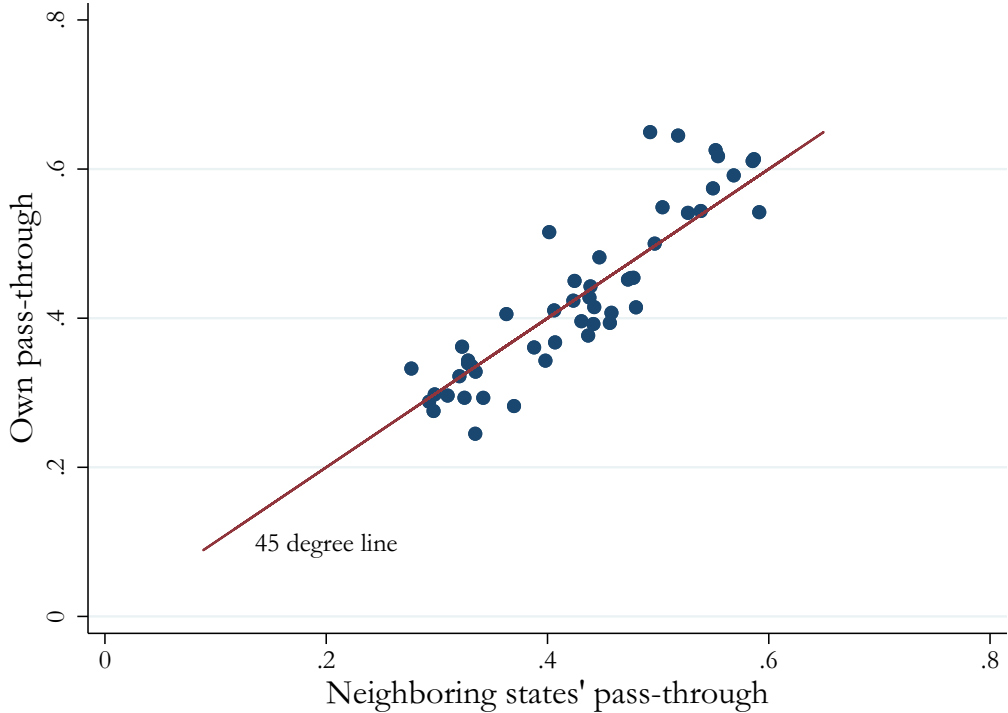


Figure B2: Map of crude oil and refined product pipelines



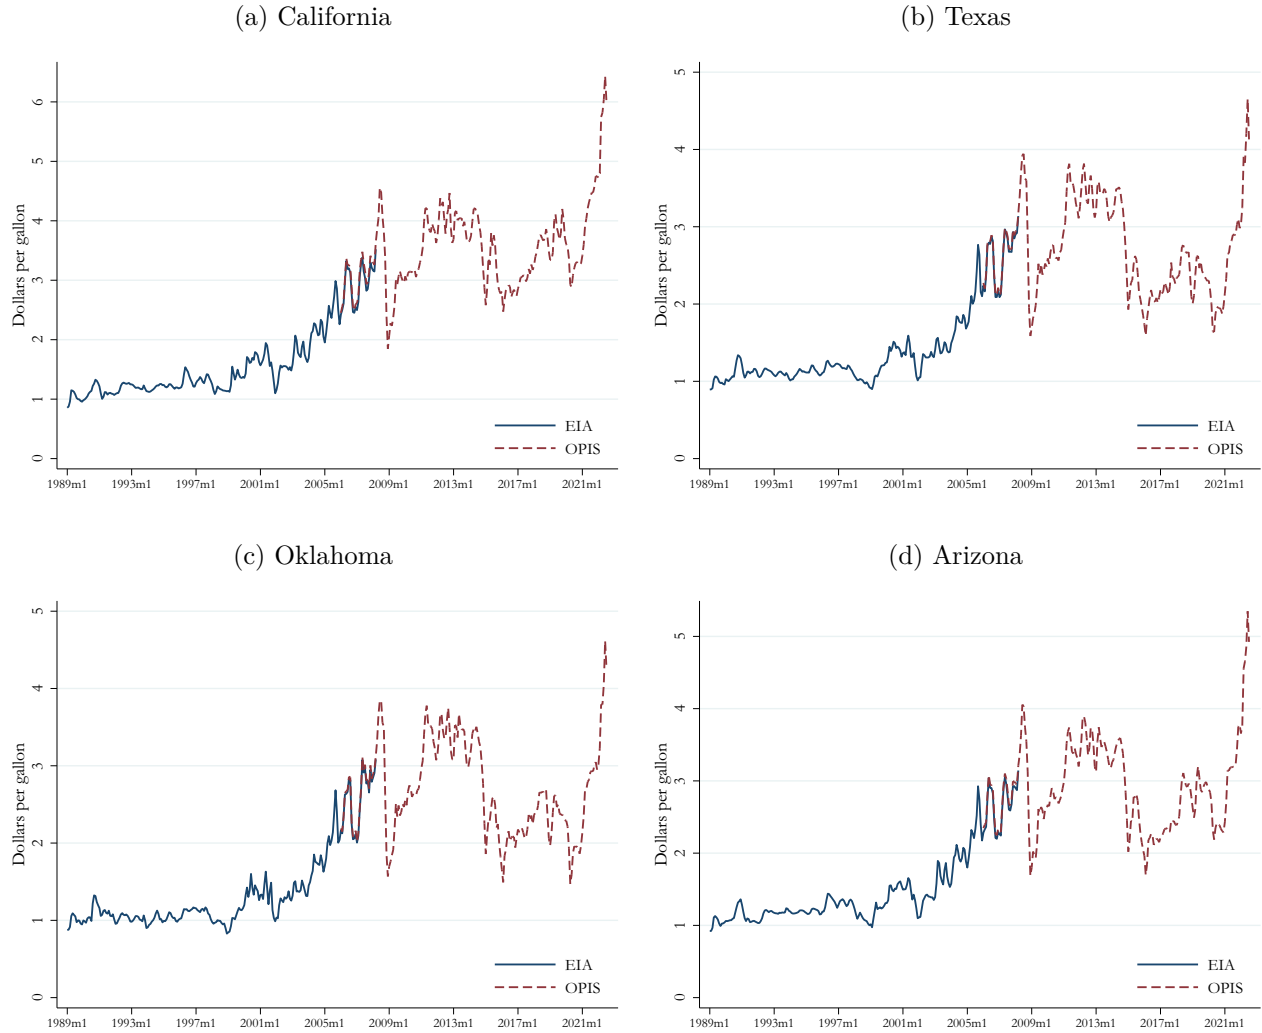
Sources: MAPSearch. This map includes information copyrighted by MAPSearch, an Endeavor Business Media, LLC, 800-823-6277. This information is provided on a best effort basis and MAPSearch does not guarantee its accuracy nor warrant its fitness for any particular purpose. Such information has been reprinted with the permission of Endeavor Media, LLC, 7/2/2022.

Figure B3: Relationship between states' own pass-through and average pass-through in adjacent states



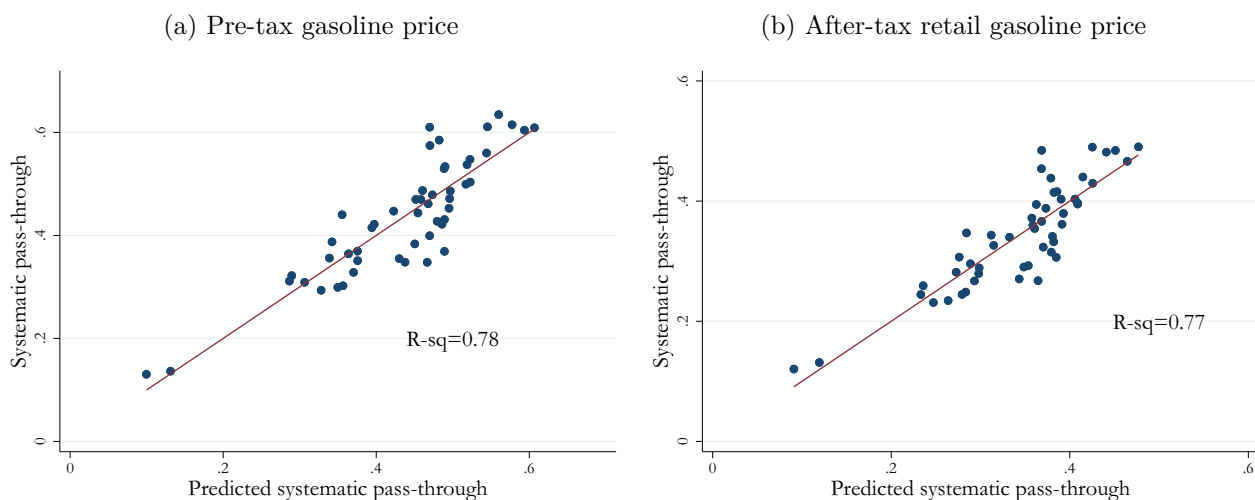
C Additional Empirical Results Using Extended Sample Period

Figure C1: Retail gasoline prices including taxes for selected states



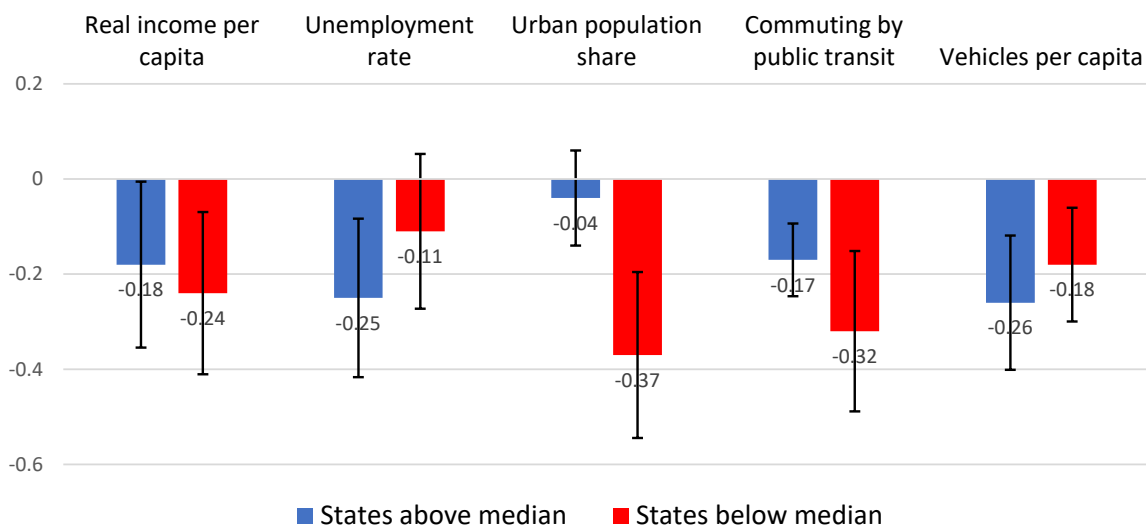
Sources: EIA state-level monthly prices of gasoline sold to end users (collected through EIA-782B form); OPIS.

Figure C2: Systematic pass-through and the predicted value using cost measures, extended data



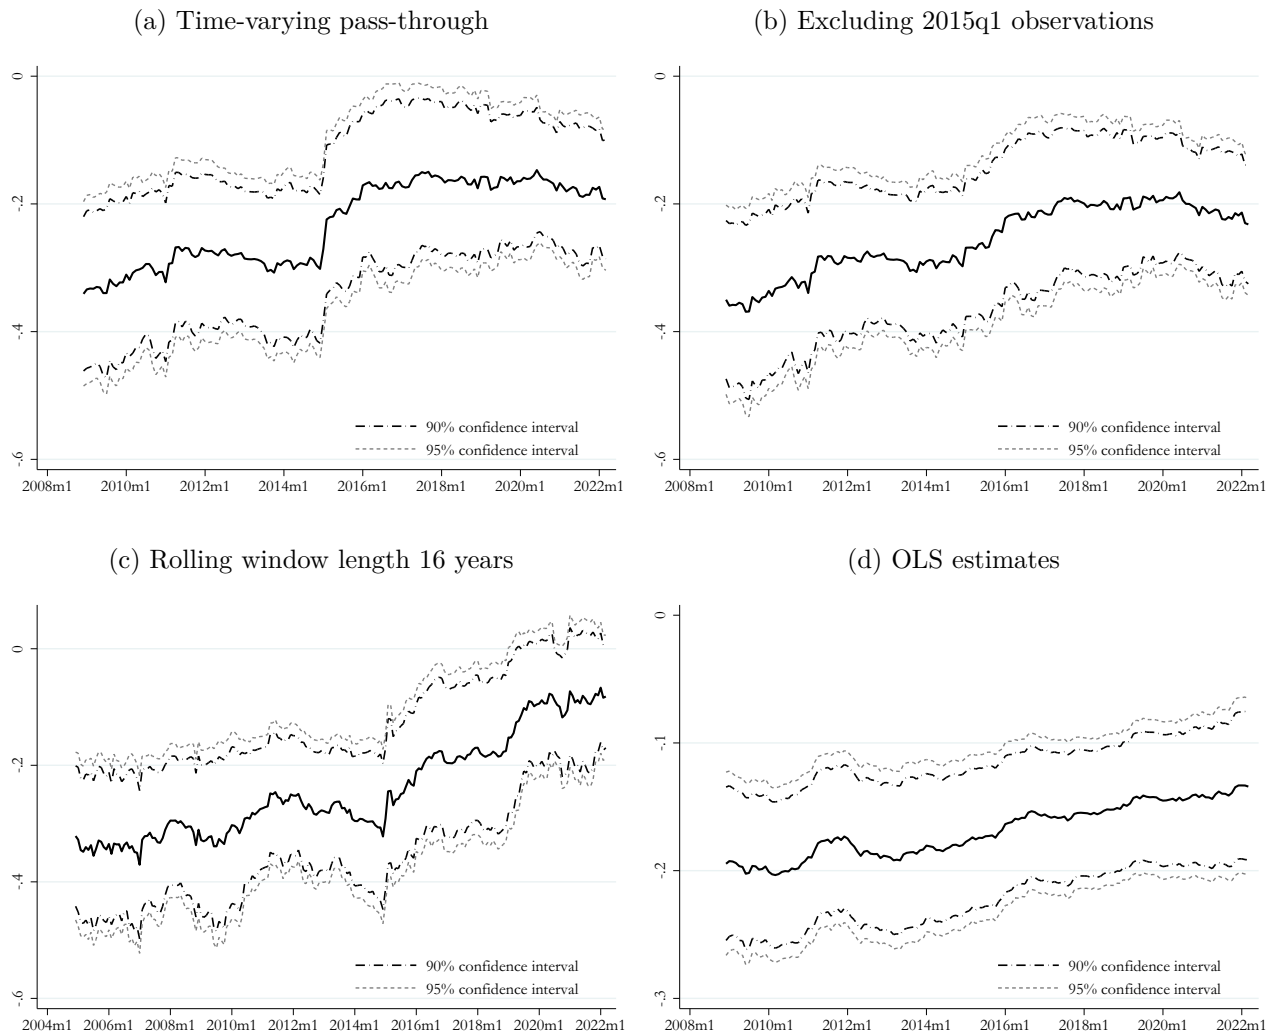
Notes: The predicted values of the systematic pass-through ($\hat{\eta}_i$ and $\hat{\theta}_i$) are obtained by estimating the specification in column (10) of Table 1.

Figure C3: Heterogeneity in the price elasticity of gasoline demand, extended data



Notes: Point estimates and 95% confidence intervals under the same model specification as for the baseline using the pass-through from oil to pre-tax retail gasoline prices.

Figure C4: Robustness: Time-varying price elasticity of gasoline demand



Notes: Rolling-window estimates. In each panel, at each point of the horizontal axis, the point estimate and confidence intervals are obtained using 20 years of data ending at that point, except for panel (c) that uses 16 years of data.

Table C1: Robustness checks

Panel I. 1989m1-2008m3	(1)	(2)	(3)	(4)	(5)
	Baseline IV	Excl. supply disruptions	Excl. TX	Excl. TX, AK, CA	Excl. TX, AK, CA, LA, NM, ND
$\Delta p_{i,t}$	-0.314*** (0.066)	-0.318*** (0.082)	-0.316*** (0.066)	-0.314*** (0.077)	-0.306*** (0.078)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes
# Obs.	11,730	10,863	11,500	11,040	10,350

Panel II. 1989m1-2022m3	(1)	(2)	(3)	(4)	(5)
	Baseline IV	Excl. supply disruptions	Excl. TX	Excl. TX, AK, CA	Excl. TX, AK, CA, LA, NM, ND
$\Delta p_{i,t}$	-0.198*** (0.053)	-0.190*** (0.063)	-0.200*** (0.053)	-0.189*** (0.058)	-0.191*** (0.058)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes
# Obs.	20,145	19,278	19,750	18,960	17,775

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table C2: Gasoline consumption responses to tax shocks and pre-tax price shocks, extended data

	(1)	(2)	(3)
	OLS	OLS	$\hat{\eta}_i \Delta p_t^O$ as IV
$\Delta p_{i,t}^{EX}$	-0.123*** (0.025)	-0.122*** (0.024)	-0.158*** (0.040)
$\Delta \tau_{i,t} / p_{i,t-1}^{EX}$	-0.672*** (0.156)	-0.670*** (0.156)	-0.678*** (0.156)
$\Delta \tau_{i,t+1} / p_{i,t}^{EX}$		0.307** (0.123)	0.307** (0.122)
$\Delta \tau_{i,t-1} / p_{i,t-2}^{EX}$		0.183 (0.114)	0.184 (0.115)
State fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
# Obs.	20,145	20,043	20,043
P-value for testing equal effects	0.001	0.694	0.856

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

Table C3: No evidence for asymmetry in the price elasticity of gasoline demand, extended data

	(1)	(2)	(3)
	OLS	IV based on $\hat{\eta}_i$	IV based on $\hat{\theta}_i$
$\Delta p_{i,t}$	-0.183*** (0.038)	-0.175** (0.068)	-0.185** (0.071)
$\Delta p_{i,t}^+$	0.039 (0.040)	-0.051 (0.146)	-0.027 (0.103)
State fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes
# Obs.	20,145	20,145	20,145

Notes: ** and *** denote significance at the 5% and 1% level, respectively. Standard errors are clustered at the state level.

D Testing the Plausibility of the Assumption of Exogenous Shares

We provided evidence in Section 3 that most of variation in θ_i and η_i is explained by differential costs in producing and distributing gasoline, supporting the assumption that the systematic pass-through from oil price shocks to retail gasoline prices is exogenous to innovations in gasoline demand. The recent literature on the shift-share design has recommended additional tests for the plausibility of this assumption. [Goldsmith-Pinkham et al. \(2020\)](#) recommend three tests. Their proposal to test for parallel pre-trends does not apply in our context. Since oil prices fluctuate continuously, unlike in program evaluation studies, there is no well-defined pre-period in our setting, making it impossible to tests for parallel pre-trends without further assumptions. It is possible, however, to implement two other tests recommended by [Goldsmith-Pinkham et al. \(2020\)](#).

- *Correlates of the systematic pass-through.* We first examine whether θ_i and η_i are correlated with innovations in gasoline demand. We consider three correlates: the employment share of the oil sector, the unemployment rate, and log per capita personal income. As recommended by [Goldsmith-Pinkham et al. \(2020\)](#), we measure these characteristics at the beginning of the estimation period (1989-1992). Table [D1](#) shows that the size of the oil sector is not correlated with θ_i or η_i , alleviating the concern that the industrial structure may cause the differential response to oil price shocks, biasing the elasticity estimate. Nor do we find a statistically significant correlation between the systematic pass-through and state unemployment rate or personal income.
- *Alternative estimators.* Table [D2](#) examines how close the estimate based on an alternative specification suggested by [Goldsmith-Pinkham et al. \(2020\)](#) is to our baseline estimate. That alternative estimator involves using the vector of the systematic pass-through, $\hat{\eta}_i$, multiplied by each time dummy as the instrument. Column (2) shows the Limited Information Maximum Likelihood (LIML) estimate that is widely recommended for dealing with many instruments. That estimate is similar to the baseline estimate. A test of the null that these estimates are identical fails to reject at conventional significance levels. As noted by [Goldsmith-Pinkham et al. \(2020\)](#), if these alternative estimates agree, researchers can be more confident in their

identifying assumptions.

Table D1: Relationship between systematic pass-through and demand-side characteristics

	(1)	(2)	(3)	(4)
	$\hat{\eta}_i$	$\hat{\eta}_i$	$\hat{\theta}_i$	$\hat{\theta}_i$
Initial oil-sector share (1989-1992)	-0.009 (0.009)	-0.009 (0.008)	-0.006 (0.006)	-0.006 (0.006)
Initial unemployment rate (1989-1992)		-0.004 (0.016)		-0.002 (0.011)
Initial log p.c. personal income (1989-1992)		-0.063 (0.091)		-0.042 (0.065)
R^2	0.02	0.04	0.02	0.03
# Obs.	51	51	51	51

Notes: ** and *** denote significance at the 5% and 1% level, respectively.

Table D2: Alternative estimators and overidentification tests

	(1)	(2)
	Baseline IV	LIML
β_1	-0.314*** (0.066)	-0.242*** (0.032)
State fixed effects	Yes	Yes
Month fixed effects	Yes	Yes
# Obs.	11,730	11,730
P-value for equal elasticity		0.22

Notes: ** and *** denote significance at the 5% and 1% level, respectively.