Welfare effects of gas price fluctuations[*]

Florian Kuhn[†] Matthias Kehrig[‡] Nicolas L. Ziebarth[§]

March 9, 2021

Abstract

The share of gasoline consumption in household expenditures decreases with income, gasoline demand being least elastic for low-income households. Based on this empirical evidence for non-homotheticities, we develop a quantitative heterogeneous-agent general equilibrium model to quantify the distributional consequences of oil price shocks. Although oil price shocks have small aggregate effects, they hurt low-income households considerably with costs to lifetime utility two to three times larger for those in the bottom decile of income relative to those in the top decile. Additionally, the 2014/15 oil glut depressed gasoline prices, which delivered comparable welfare benefits to the 2018 tax cuts.

Keywords: oil prices, welfare costs, consumer durables, commuting
JEL classification: D12, E22, H22.

---

[*] We would like to thank Nick Roussanov, Stephie Fried, Yuriy Gorodnichenko, Cristina Arellano and Fabian Lange and seminar participants at Duke, the Annual Meeting of German Economists Abroad and the Empirical Macro Workshop for helpful comments. We thank Elessar Chen for excellent research assistance. All errors are our own.

[†] Department of Economics, Binghamton University. Email: fkuhn@binghamton.edu.

[‡] Department of Economics, Duke University, NBER, and CEPR. Email: matthias.kehrig@gmail.com

[§] Department of Economics, Auburn University and NBER. Email: nicolas.ziebarth@me.com

1 Introduction

Short-run fluctuations in oil prices and gasoline or energy, more broadly, draw considerable attention from policymakers and the popular press for what those changes portend for consumer well-being. The public debate commonly pictures a household that cannot change its oil consumption much—for example, they must commute to work—which leaves them vulnerable to any oil price change. Unlike the public perception of the effects of oil price shocks, macroeconomic research has largely focused on the *indirect* or supply-side welfare effects of oil prices through their impact on (market) production and inflation. However, the evidence for the link between oil prices and the incidence of a recession is mixed.

We argue that by focusing on the supply-side effects on aggregates like output or inflation, the macroeconomic literature has not addressed the distributional consequences of the direct or demand-side effects of oil price shocks. To illustrate these distributional impacts, we use data from the Consumer Expenditure Survey (CEX) to study households' gasoline consumption—the main form in which households purchase (processed) oil. We find that among gasoline-using households, there is a robust negative relationship between income and the share of a household's budget spent on gasoline. Furthermore, gasoline demand appears to be less elastic at the low end of the income distribution. The non-homotheticities of gasoline demand highlights the way in which exogenous fluctuations in the price of oil act as a regressive income tax, disproportionately hurting low-income households.

We then incorporate these empirical findings into a dynamic heterogeneous-agent general equilibrium model, in which households differ in their labor income and in their energy usage. In particular, the decreasing propensity to consume gasoline is introduced via a fixed minimum quantity of gasoline that must be consumed by all households in addition to variable gasoline consumption. This inelastic part of a household's gasoline consumption can be interpreted as being required for commuting to work. Any quantity of gasoline consumed beyond this minimum level enters a household's utility as a complement to an output good, which represents the remainder of the consumption basket. Furthermore, to capture the direct effects of oil price shocks on households, we include in the utility function home production that requires both energy and consumer durables.

We calibrate the model to match the joint distribution in household income and gasoline expenditures. Using the estimated model, we can examine the welfare effects of a shock to the gasoline price for any household in the income and gasoline usage distribution by comparing a one-time oil price shock to a permanent change in the steady-state labor tax.

We find that the welfare of low-income households is almost twice as sensitive to the gasoline price as that of high-income households. For example, a temporary increase of the gasoline price from \$2 to \$3 per gallon is equivalent to a permanent labor tax hike of 0.5 percentage points for high-income households, whereas for poor households, it is equivalent to a hike of 0.95 percentage points. Expressed in permanent labor tax changes, households at the bottom decile of the income distribution would be willing to accept a labor tax increase that is 4.5 times as high as those agreed to by households at the top decile of the income distribution. This finding shows the important distributional consequences of the direct effects of oil price shocks even in the short run.

As a case study, we consider the period in 2014–2015 when many oil-exporting countries in the Middle East drastically expanded supply to put pressure on the US fracking industry. This caused a sharp though, in the end, temporary decline in oil and gasoline prices. We quantify the welfare benefits of this oil glut for almost all households to be at least as large as the labor income tax cuts implemented with the Tax Cuts and Jobs Act of 2017. This improvement is particularly strong at the lower end of the income distribution. In particular, we find that the bottom decile benefited three times as much from that oil glut compared to the 2017 tax reform.

Our paper adds to the macroeconomic literature on the effects of oil price shocks. Extensive reviews of this literature can be found, for example, in Barsky and Kilian (2004) and Hamilton (2003). As emphasized by Edelstein and Kilian (2009), Barsky and Kilian (2004) and Baumeister and Kilian (2014), there is some evidence that an important channel for the oil price's effect on output is through aggregate demand and, specifically, consumer spending. Several New Keynesian models have taken this approach by combining a role for oil in firm or household demand with price rigidities and fluctuations in the oil price, for example, Blanchard and Riggi (2013) or Kilian and Vigfusson (2017). We too focus on the direct, household-side effects of these shocks. In our model, oil price shocks will have longer-lasting effects because they depress investment in consumer durables that must be powered with energy.

Our paper connects to a growing literature on the costs and benefits of energy taxes in response to climate change, which takes as a starting point that a carbon tax is necessary to internalize the negative externality from carbon emissions. For example, Golosov et al. (2014) analyze a dynamic stochastic general equilibrium with an externality from fossil energy and find that the optimal tax should be a bit higher than the most well-known estimates of Nordhaus and Boyer (2000) and Stern (2007). An important question is whether such a tax is regressive and, if so, how that regressivity can be offset through rebates of the proceeds or by reducing other distortionary taxes (the so-called double dividend). For example, Metcalf (1999) finds that a (regressive) carbon tax can be combined with changing other tax rates to leave the overall income distribution unchanged, while Hassett et al. (2007) argues that these carbon taxes are not that regressive to begin with when compared to relative lifetime versus annual income. More recent work considers distributional consequences intergenerationally. Fried et al. (2018) show there is a disagreement between current and future generations on how the dividends of a carbon tax should be rebated, with current generations preferring that the proceeds be refunded in a lump sum manner and future generations preferring that other distortionary taxes be decreased. Similar intergenerational disagreements are present in Leach (2009), Carbone et al. (2013), and Rausch (2013). Ready et al. (2020) study the labor supply response embedded in commuting changes after gasoline price shocks in a two-sector model. Ready (2018) studies the impact of oil supply uncertainty on asset prices when oil enters the household utility function.

Many of these studies on the costs of energy taxes, while careful about distinguishing between short- vs. long-run responses, are only able to compare welfare across different steady states corresponding to different energy prices. For example, Bento et al. (2009) offer an elaborate model of household demand for automobiles and gasoline consumption. Households make a forward looking decision as to whether scrap their existing vehicle by comparing its present discounted value relative to its scrap value. However, households assume that "future rental values will be the same as the current-period rental values of older vintages of the same vehicle type" (pg. 674). This assumption of myopic expectations is not problematic for simulating the long-run effects of permanent change in gasoline prices. It is problematic, as they write on pg. 674, "[i]f the policy involved government committing to a path of varying gasoline taxes in the future, for example, a more complex modeling of expected future prices might be called for." Accurately capturing the short-, medium-, and long-run effects of the time-varying gasoline prices is one of our primary areas of concern and is why we consider a forward looking household with rational expectations at the expense of a less elaborate consumer durable choice.

Papers closest to ours in this environmental literature, including Chiroleu-Assouline and Fodha (2014) and Williams et al. (2015), focus on these cross-sectional differences in the welfare effects of carbon taxes. However, there are a few important differences between those and ours. Chiroleu-Assouline and Fodha (2014) is a purely theoretical exercise showing there exists a Pareto-improving environmental tax reform, while Williams et al. (2015) and Goulder et al. (2019) use a computable general equilibrium model without any dynamics or consumer durables the generate utility. In addition, we use consider the joint distribution of income and gasoline usage to discipline the household heterogeneity our model.

Finally, a microeconomic literature has attempted to measure the income and price elasticities of demand for gasoline or energy. Brons et al. (2008) and Labandeira et al. (2017) conduct meta-analyses of these studies with the former focusing on gasoline and the latter on energy, more broadly.[1] They tend to find short-run elasticities of around –0.34 to –0.22 and longer-run elasticities two to three times as large, with the short-run elasticities in line with what Levin et al. (2017) estimate using daily city-level data. As for the income elasticity, Havranek and Kokes (2015) in a meta-study find a short-run elasticity of around 0.28 with long-run elasticities over two times larger. While the papers in this literature tend to use individual- or household-level data, Edelstein and Kilian (2009) and Hughes et al. (2008) infer the price elasticity of oil demand from aggregate data.

Much of this empirical literature has assumed constant price and income elasticities that are independent of household characteristics. One exception is Schmalensee and Stoker (1999), who employ semi-parametric methods to estimate gasoline consumption as a function of household characteristics. They find non-constant income and price elasticities though household characteristics enter in an additively separable way from the price term.[2] More recently, West and Williams (2004), Wadud et al. (2010), and Gillingham (2014) find evidence that the price and income elasticity of gasoline declines with income.

Our paper proceeds as follows: Section 2 documents the empirical gasoline usage throughout the income distribution; Section 3 develops a heterogeneous-agent model with gasoline consumption by households and firms that is geared to replicate the key empirical patterns documented. Section 4 studies the quantitative magnitudes of welfare costs throughout the income distribution and uses the model to carry out several counterfactuals: we compare oil price shocks and gluts to the recent tax reform and consider the welfare impact of higher carbon taxes for high- and low-income households. Section 5 concludes.

2 Evidence on Household Gasoline Demand

We now provide some motivating evidence for some of our model's key features using the CEX. In particular, we study how gasoline demand varies by household income. We are not the first to consider household-level differences in gasoline demand empirically. For example, both West and Williams (2004) and Wadud et al. (2010) also use this data source to study gasoline demand due to the rich information it provides on household expenditures and characteristics. Our focus is different from these earlier works with a specific focus on characterizing the joint distribution of income and gasoline consumption. This joint distribution is critical for calibrating the model we present in section 3 as well as for the magnitude of the welfare effects of gasoline price shocks.

---

[1] Note that there are many other meta-studies in the literature on this topic. We cite the most recent ones we are aware of.

[2] Yatchew and No (2001) estimates a similar model using better data from Canada.

The CEX is a rotating panel of households in which households remain for up to four consecutive quarters. In any given quarter, there are observations for around $6,800$ households. In this analysis we use the publicly available part of the dataset ranging from 1999 to 2013.[3] The interview data contain information on quarterly expenditures in a large number of categories, of which the main item of interest is expenditures for "Gasoline and Motor Oil." Throughout this section, the main focus is on the "gasoline budget," that is, gasoline expenditures as a share of total household expenditures, or of household income, in the given period.[4] We construct this measure by dividing quarterly gasoline expenditures by the household's total quarterly expenditures.

The CEX also contains data for the household's annual income before and after taxes as well as its income rank (among the set of CEX households). The income data are collected only once for most households such that for the second and third interview, the income data from the first interview is used. There is a considerable number of missing values (around 27%) for the income data. The survey also contains a number of other demographic characteristics.

2.1 Gasoline Expenditures and Earnings

Figure 1 displays histograms of gas consumption as a share of total expenditures and as a share of income in the top panel. These histograms show an approximately log-normal distribution with an added mass point at zero, which is the share of households that did not have any gasoline expenditures in the previous quarter. Out of all $446,114$ household-quarter observations, zero gasoline expenditures are reported in $46,813$ cases or 10.4% of observations. As shown in Table 1, around 85% of households consume gasoline in every observed quarter, and for only around 5% of households there is at least one quarter with positive gasoline consumption and another quarter of zero gasoline consumption. Around 9% of households never buy gasoline.

---

[3] Data are available at www.bls.gov/cex/pumdhome.htm.

[4] Given that households usually use additional types of energy products beyond gasoline (like natural gas, fuels and electricity), and the fact that these energy prices can be correlated, one may be concerned that only looking at gasoline consumption does not capture the full impact that gasoline price fluctuations have on households. In appendix B we use all energy expenditures as a robustness check and show that there is also a strong negative relationship between income and total energy expenditure shares.

Table 1: Summary statistics for gasoline consumption

	Average	StD	Q_{10}	Median	Q_{90}	N
Gasoline consumption (gal)	202	196	0	157	426	446,114
(positive usage)	226	193	57	177	448	399,301
Gasoline expenditure share (%)	5	4.59	0	3.98	10.5	446,025
(positive usage)	5.59	4.5	1.52	4.44	10.96	399,281
Gasoline income share (%)	6.54	11.04	0	3.43	13.5	324,902
(positive usage)	7.29	11.42	1.23	3.9	14.69	291,269
Share of permanent gasoline users	85.6%					
Share of occasional gasoline users	5.5 %					

Author calculations based on data from the CEX; see text and appendix A for description.

The literature has effectively ignored households with zero gasoline consumption by estimating log-linear models of demand. This effectively drops these households with zero gasoline consumption and estimates gasoline demand on the *intensive* margin, which has major effects on estimated income and price elasticities, a point also emphasized by West (2004). In fact, the relationship between income and the household gasoline budget including households with no gasoline consumption is a hump-shaped function of income, a pattern documented by Poterba (1991). However, Poterba's work effectively aggregates across different types of low-income households. This misses the fact that the increase in the gasoline budget at low income levels is driven entirely by the extensive consumption margin (i.e., the decreasing likelihood that households consume zero gasoline): restricting the sample to households with positive gasoline usage (or as an alternative, households owning one or more automobiles) yields a negative relationship between income and gasoline budget, as shown in the middle and bottom panels of Figure 1.

One concern is whether this negative link between income and gasoline composition is a compositional effect rather than a direct association. For example, if household income was positively correlated with some demographic characteristic, like age, that negatively correlated with gasoline expenditures, this could explain the negative bivariate relationship between income and the gas budget. To allay this concern, we now examine this link conditional on observable characteristics as shown in Figure 2. This figure shows the relationships for different subgroups by age, family size, education levels, and city size, respectively.

Across all these demographic categories, we find a negative relationship between income

Figure 1: Distribution of gasoline consumption and expenditures

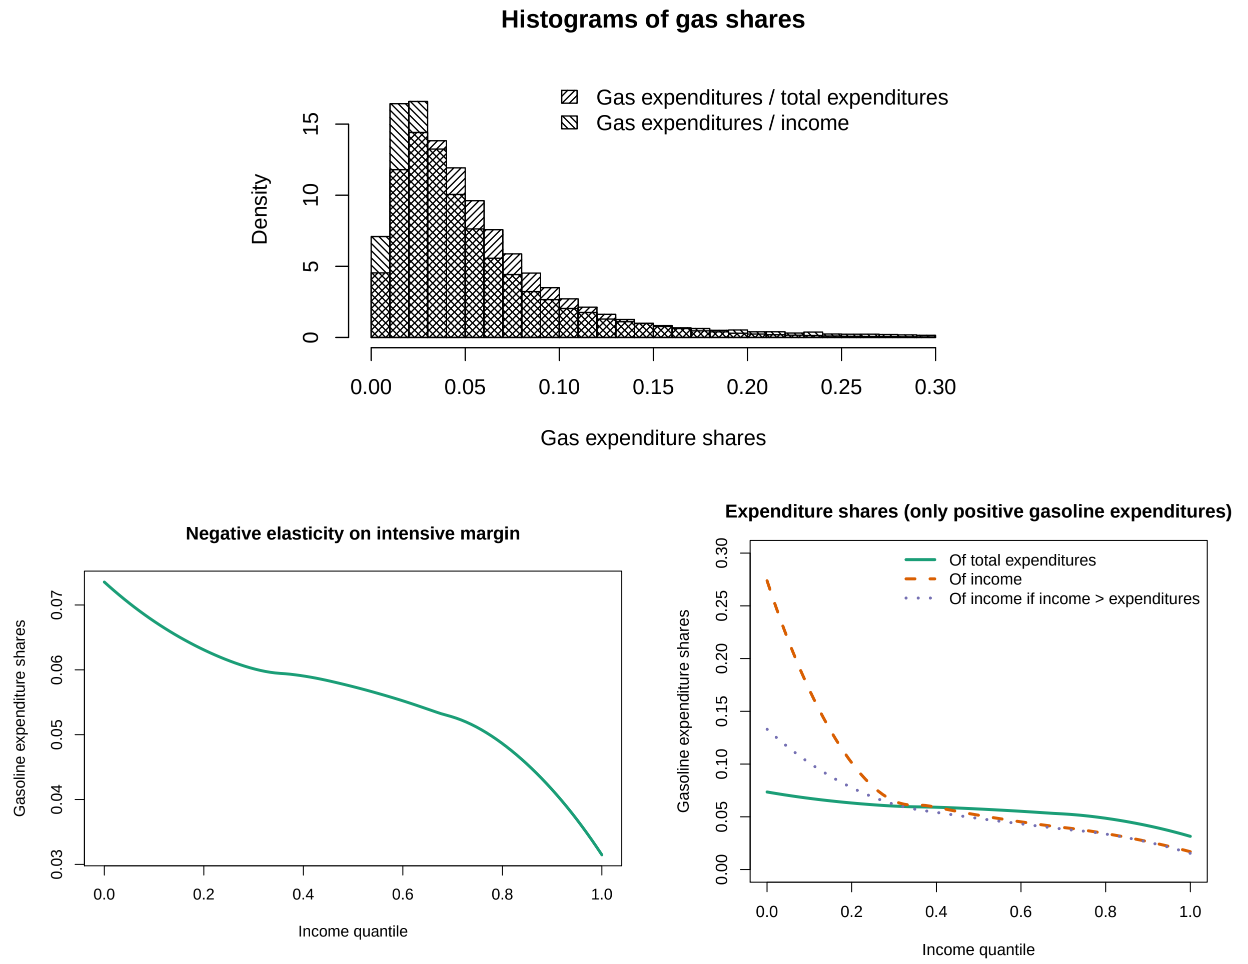

Top: Histogram of expenditure shares as part of total household expenditures and part of income, respectively. Middle: Expenditure shares of all households and conditional on positive gasoline usage. Bottom: Expenditure shares conditional on positive gasoline usage as part of all expenditures and as part of income, respectively.

and gas budget. Additionally, the graphs confirm some expected relationships between demographic subgroup and the *level* of gas consumption. For example, seniors spend considerably less on gas than younger people for all income levels. Similarly, households living in smaller cities tend to have a higher gas budget, as do families with more household members.

Figure 2: Relationship between gasoline expenditure and income by demographics

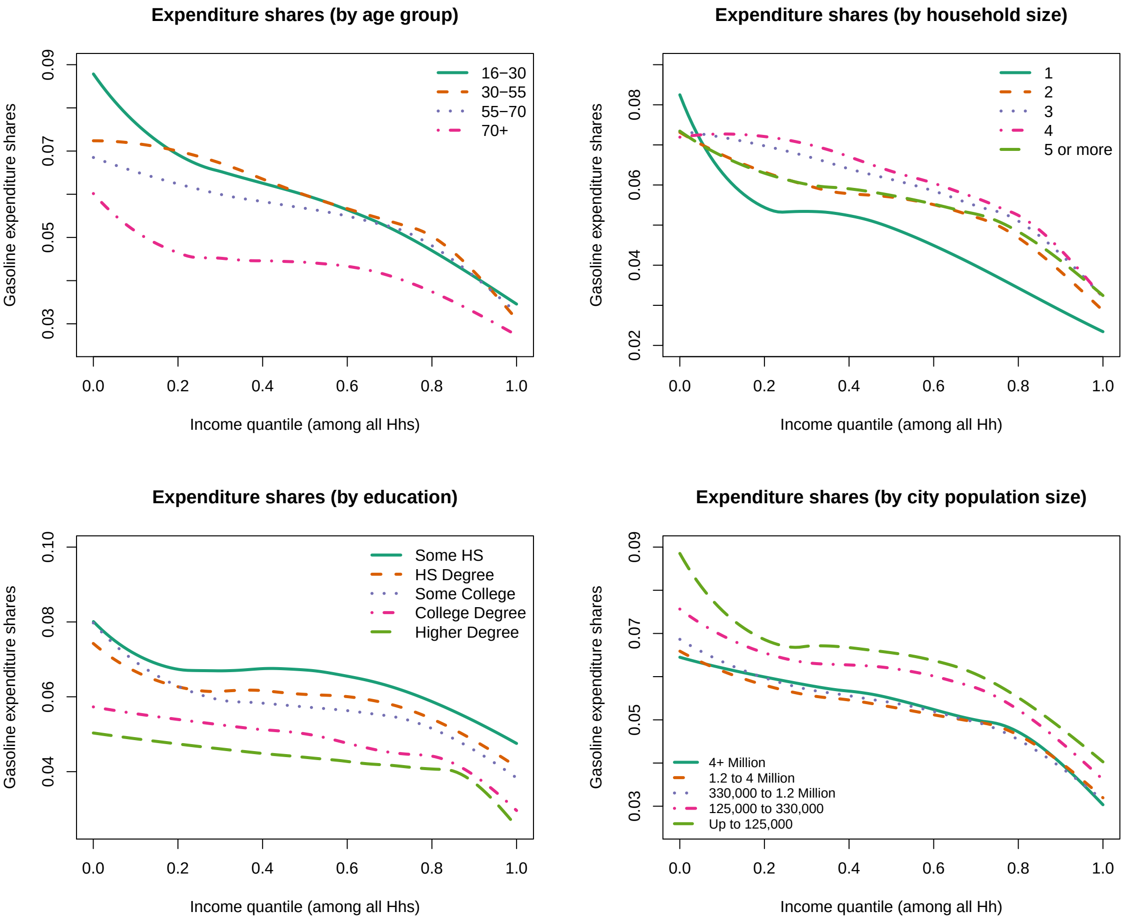

2.2 Estimating Gasoline Engel Curves

We now estimate Engel curves or income elasticities for gasoline demand, focusing on the households with strictly positive gasoline expenditures. A key question in many of these studies of demand for gasoline (or any good generally) is how to measure income. Theory would suggest that permanent income is the right variable to include, but that is never available. We thus consider a number of proxies including income percentile rank, annual income, and consumption expenditure. We control in addition for a number of household characteristics as well as time fixed effects. The household characteristics include a region dummy (Midwest, Northeast, South, West), a categorical variable for size of the metropolitan area, gender, race, age, and education of the household's reference person as defined in the CEX as well as marital status and family size. We then run the regression

$$y_{it} = \gamma \, \mathrm{inc}_{it} + \beta X_{it} + d_t + \varepsilon_{it},$$

where inc_{it} is a measure of income of household i at time t and d is the time dummy. We cluster standard errors at the household level.

Table 2 shows these regression results. For all measures of income, the income elasticity after controlling for the demographic characteristics is negative. The first three columns are the results from OLS regressions pooling all observations, while the last two include household fixed effects. As listed in column (1), moving up one decile in the income distribution decreases the gas budget by 0.25 percentage points. This finding of a negative income elasticity of the gasoline expenditure share is line with the estimates in papers by Kayser (2000) and Sipes and Mendelsohn (2001). Going beyond these earlier works, we provide evidence that the income-gasoline expenditure relationship is monotonic and not driven by compositional effects. Both of these observations are important for our our choice of modeling households' gasoline consumption as non-homothetic.

Of course, the OLS estimates ignore any unobserved household heterogeneity: it could be that there is an unobserved household characteristic driving gasoline consumption that is correlated with income. Although the time dimension of the panel is short (any household remains in the dataset for at most four consecutive quarters), columns (4) and (5) display results from a regression with an added household fixed effect. Since households' demographic controls barely vary over the course of one year, only the time dummies are included in these specifications. Even with these caveats, the income elasticities are similar to the OLS estimates.

Another concern is that we only observe households' current income. But economic theory predicts that expected lifetime income affects consumption choices to some degree beyond current income. For example, if households face transitory income shocks around a long-run average, then using current income will underestimate the response of gasoline consumption to changes in permanent income. A similar issue arises when income measures are subject to measurement error, which will lead to attenuation bias, biasing estimates towards 0. In columns (6) and (7), we therefore follow the literature (e.g. as in Dynarski et al. (1997)) and use education level and lagged income, respectively, as an instrument for current income. As expected, this leads to an increase in the magnitude of the respective coefficients.

Table 2: Estimates of Income Elasticities for Gasoline

	Dependent variable: Gas expenditure share						
	OLS			Fixed effects		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income Rank	-0.0253^{***}			$-.0110^{***}$		-0.0401^{***}	
	(0.0003)			(0.0009)		(0.0012)	
log(Income)		-0.0045^{***}					-0.0076^{***}
		(0.0001)					(0.0002)
log(Expenditures)			-0.0244^{***}		-0.0249^{***}		
			(0.0001)		(0.00025)		
S.e. FE (σ_u)				0.0405	0.0429		
S.e. iid noise (σ_e)				0.0284	0.0288		
Observations	281,167	298,549	398,139			58,983	61,894
R^2	0.1722	0.1640	0.2308				

Controls not shown included in OLS and IV regressions: Quarterly time effects, region, metropolitan area, gender, race, age, education, marital status and family size. Columns (6) and (7) use the log of lagged income as an instrument.

3 Model

3.1 Households

We now write a model aimed at replicating the empirical findings that allows us to quantify the welfare effects of oil price shocks. The key aspect of our model will be the non-homothetic behavior of gasoline expenditures by income rank. Consistent with our empirical analysis, we will only focus on households with positive gasoline consumption. One way to interpret this assumption is that households make a (costly) decision at the start of life on whether to consume gasoline or not. This could be interpreted as a decision of whether to live in a city that has public transportation or not. Our implicit assumption in the subsequent analysis is that the oil price shocks we study are not large enough to cause a sizable numbers of household to reconsider their extensive margin decisions. In this case, it makes sense to abstract away those who do not consume gasoline at all and calibrate the model to those who are making an intensive margin decision of how much gasoline to consume.

There is a continuum of mass 1 of households that are indexed with i and can buy consumption goods C_{it}, durable goods D_{it}, and gasoline. Households purchase gasoline for two purposes: first, they inelastically buy an exogenously given amount $\bar{E}_{it}$ every period, which can be interpreted as needed to commute to work. Second, they buy a quantity E_{it}^H that enters the utility function when used together with the durable good D_{it}. The durable good is the composite variable "service flows from durables" S_{it}, defined as

$$S_{it} \equiv D_{it}^{\gamma} \left(E_{it}^H\right)^{1-\gamma} \tag{1}$$

and captures the flexible part of gasoline consumption, for example, households going on road trips. A household's total gas consumption in gallons in period t is therefore $\bar{E}_{it} + E_{it}^H$. Durable goods are purchased in period t for period $t+1$. They have the same unit as consumption goods C and depreciate at a quarterly rate d.

The oil price q_t evolves exogenously according to an autoregressive process given by

$$\log q_t = (1-\rho_q)\log \tilde{q} + \rho_q \log q_{t-1} + \varepsilon_t^q. \tag{2}$$

The parameter $\tilde{q}$ sets the steady-state level of the gasoline price.

Labor productivity of household i is given by the exogenous variable e_{it}. Because labor supply is fixed at 1, e_{it} also represents the effective labor supply a household offers on the labor market. This labor market is common for all households so that each effective unit of labor hired is paid the same wage w_t. Households have to pay a proportional labor tax τ^L

and a sales tax of τ^C on their consumption C_{it} of the output good, and they receive profits π_t from firms.

Households can differ along two dimensions: income level e_{it} and exogenous gasoline usage level. In line with much prior literature, we model the income process with a log-normal AR-1. Because we do not have longer panels on households gasoline usage, we split households' exogenous gas requirements $\bar{E}_{it}$ as having a permanent and a transitory component so that $\log \bar{E}_{it} = \log \bar{E}_i + \varepsilon_{it}^{\bar{E}}$. We describe the calibration of both processes in more detail in Section 3.4.

The households' budget constraint in real terms is

$$\left(1-\tau^L\right) w_t e_{it} + (1-d) D_{it} + \pi_t = \left(1+\tau^C\right) C_{it} + D_{i,t+1} + q_t \left(E_{it}^H + \bar{E}_{it}\right), \tag{3}$$

where C_{it} is consumption of the output good that serves as the economy's numeraire.

Household i chooses $C_{it}, E_{it}^H, D_{i,t+1}, S_{it}$ to maximize lifetime utility given by

$$E\left[\sum_{t=0}^{\infty} \beta^t \log\left[\varphi C_{it}^{\frac{\zeta-1}{\zeta}} + (1-\varphi) S_{it}^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}\right],$$

subject to the sequence of budget constraints (3) and the definition of durables services (1). While this utility function specifies that households have a constant elasticity of substitution between consumption of the output good and the service flow from durable goods, the fixed component $\bar{E}_{it}$ will introduce a non-homotheticity to households' gasoline demand, causing it to fall as a share of total expenditures with increasing income.

It is useful to compare our modeling of households to papers such as Bento et al. (2009) and Dubin and McFadden (1984) that have combined both a durable discrete choice with a continuous choice of how much gasoline (or energy) to use. Households in the model of Bento et al. (2009) have preferences over various automobile attributes including fuel efficiency. They then make a choice of which vehicle to buy (if any) as well as whether the car should be new or old. Finally, households can make a decision to scrap their current vehicle. This last margin of adjustment allows Bento et al. (2009) to interpret the price and income elasticities they estimate as being "long-run" elasticities. Furthermore, by including the scrap decision, they can study how the fuel efficiency of the car fleet evolves following a change in the gasoline tax. In contrast, while our model does not have an explicit choice of fuel efficiency, the continuous household decision to accumulate durable goods capital could be interpreted as just a such a decision. That is, the stock of durables owned by a household controls the productivity of turning the (physical) amount of fuel consumed into actual energy services that households derive utility from.[5]

The first-order necessary conditions (FOC) for E_{it}^{H} and $D_{i,t+1}$, respectively, combine with the FOC for consumption to yield

$$\frac{1-\varphi}{\varphi}\left(1-\gamma\right) S_{it}^{\frac{\zeta-1}{\zeta}} \left(E_{it}^{H}\right)^{-1} C_{it}^{\frac{1}{\zeta}} = \frac{q_t}{1+\tau^{C}}, \tag{4}$$

$$E_t\left[\beta \tilde{S}_{t+1}^{-1}\left\{\left(1+\tau^{C}\right)\frac{1-\varphi}{\varphi}\gamma S_{it}^{\frac{\zeta-1}{\zeta}} D_{i,t+1}^{-1} + C_{i,t+1}^{-\frac{1}{\zeta}}\left(1-d\right)\right\}\right] = \tilde{S}_{it}^{-1} C_{it}^{-\frac{1}{\zeta}}, \tag{5}$$

where $\tilde{S}_{it} \equiv \left[\varphi C_{it}^{\frac{\zeta-1}{\zeta}} + \left(1-\varphi\right) S_{it}^{\frac{\zeta-1}{\zeta}}\right]$ is a shorthand and E_t is the expectations operator conditional on information at time t.

3.2 Firms

On the economy's production side, a representative firms produces output goods Y_t according to

$$Y_t = L_t^{\alpha}\left(E_t^{F}\right)^{\kappa},$$

where L_t is the firm's labor input and E_t^{F} is oil used by the firm in production.

The firm's repeated static problem is to maximize

$$\max_{L_t, E_t^{F}} L_t^{\alpha}\left(E_t^{F}\right)^{\kappa} - w_t L_t - q_t E_t^{F}$$

such that the optimality conditions are

$$\alpha L_t^{\alpha-1}\left(E_t^{F}\right)^{\kappa} = w_t, \tag{6}$$

$$\kappa L_t^{\alpha}\left(E_t^{F}\right)^{\kappa-1} = q_t. \tag{7}$$

3.3 Market Clearing and Equilibrium

We do not model the supply of oil. Instead we assume there is a perfectly elastic supply of oil at the exogenous price q_t. One way to think about this is that the price is set by OPEC and they commit to producing any amount at that spot price.

The market clearing condition for final goods corresponds to the economy's resource

---

[5] We do not think the fact that difference between the continuous durable good choice in our model vs. the discrete one in their model is important.

constraint

$$Y_t - E_t^F - q_t \bar{E} - \int_i q_t E_{it}^H di = \int_i \left[\left(1+\tau^C\right) C_{it} + D_{i,t+1} - (1-d) D_{it} + \tau_i^L w_t e_i. \right] di, \qquad (8)$$

Labor market clearing requires that households' supply of labor efficiency units corresponds to the firm's labor demand

$$L_t = \int_i e_{it} di. \qquad (9)$$

An equilibrium is defined as endogenous quantities $\left\{C_{it}, D_{i,t+1}, E_{it}^H, L_t, E_t^F\right\}$, the real wage $\{w_t\}$, and a stochastic process for the exogenous state variable $\{q_t\}$ such that optimality and market clearing conditions are satisfied. This is summarized by the system of equations consisting of households' FOC (4)–(5) and budget constraints (3), the firm FOC (6)–(7), one of the market clearing conditions (8) or (9), and the stochastic equation for the oil price (2).

3.4 Calibration and Solution

While the economy features aggregate uncertainty, the source of this uncertainty is directly in the oil price itself and does not propagate into endogenous aggregate state variables. The model's solution is then straightforward since expectations about future prices, depending only on the future path of oil prices, are easy to compute.

Households' idiosyncratic labor productivity e_{it}, which determines their income levels, is given by a normal AR-1 process in logs with the cross-sectional mean normalized to one; specifically,

$$\log e_{it} = \alpha_i + \rho_e \log e_{i,t-1} + \varepsilon_{it}^e.$$

The empirical literature estimating household income processes (e.g., Guvenen et al. (2019), Guvenen and Ozkan (2014), Heathcote et al. (2010), Storesletten et al. (2004)) finds large values for the persistence ρ_e estimated on annual data ranging from 0.958 up to values exceeding unity. We follow Storesletten et al. (2004) who estimate a specification similar to ours, finding an annual persistence of 0.977, which in our quarterly setup corresponds to $\rho_e = 0.994$. Based on the variance of innovations they find, we use a quarterly standard deviation of $sd\left(\varepsilon_{it}^e\right) = 0.014$. Finally, we set the cross-sectional variance of the household-specific fixed effects α_i in a way to match the long-run average of dispersion in household incomes of around 0.9 (see, e.g., Nichols and Zimmerman (2008) and Diaz-Gimenez et al. (2011)).

With $\alpha = 2/3$, the elasticity of production with respect to labor is set to a standard value, as is households' quarterly discount rate with $\beta = 0.99$. We also follow Kehrig and Ziebarth (2009) who find that consumption and oil usage enter as complements in household preferences and set the elasticity of substitution $\zeta = 0.73$.

The labor tax rate $\tau^L = 0.2$ roughly targets total receipts of income tax and payroll tax in the US as a share of total labor income. For the exogenous process describing the gas price, running an AR-1 estimation on quarterly retail gasoline prices from 1976Q1 to 2014Q4 yields a persistence parameter of $\rho_q = 0.95$ and a residual standard deviation of $\varepsilon_t^q = 0.14$. (Interestingly, these estimates are robust to starting the sample in 1987 after the large shocks of 1979 and 1986, respectively.)

This leaves the parameters $\tilde{q}, \bar{E}, \text{sd}\left(\varepsilon_i^{\bar{E}} t\right), \text{sd}\left(\bar{E}_i\right), \gamma,$ and φ governing the household's gasoline consumption. Of these, both the steady-state gas price $\tilde{q}$ and the preference weight φ determine the steady-state level of the gas expenditure budget and are not separately identified, so we normalize the long-run gas price $\tilde{q}$ to one. In the data, we observe considerable heterogeneity in gas purchases conditional on income, which in our model is governed by the distribution of $\bar{E}_{it}$ around its mean $\bar{E}$ which in turn consists of the two components $\bar{E}_i + \varepsilon_{it}^{\bar{E}}$. A random-effects regression in the CEX of households' gasoline expenditure share on income reveals that (with 51% and 49%, respectively) the random effect and the transitory shock contribute equally to the residual variance of gasoline consumption. We therefore fix $\text{sd}\left(\varepsilon_i^{\bar{E}} t\right) = \text{sd}\left(\bar{E}_i\right)$.

We set the remaining four parameters in a way so that the following four moments related to gasoline consumption in the model match their empirical counterparts. While all four parameters affect these outcomes jointly, we give some intuition for why we use these particular moments to identify the parameters. First, the random-effects regression of gasoline expenditure shares on income leaves a residual variance of 0.04. Intuitively, the parameter $\text{sd}\left(\varepsilon_i^{\bar{E}} t\right)$ strongly influences this moment in the model. Second, the parameter γ has a strong effect on the relative contribution of the durable stock and gasoline purchases in the utility flow they provide as defined in equation (1). We use the CEX averaging across all households to find mean quarterly expenditures for vehicle purchases, financing, and insurance and use this as our measure of durable expenditures. Compared to an average quarterly spending on gasoline of \$533, these average vehicle expenditures of \$781 are around 50% larger. Correspondingly, we calibrate the model parameter γ so that the steady-state ratio of consumer expenditures on durables D and spending on gasoline $E^{HH} + \bar{E}$ are in the same proportion as in the data. Third, we match the mean gas expenditure share of 5.5% in the CEX. Intuitively, mean fixed gasoline consumption $\bar{E}$ significantly affects this moment in the model. Finally, we make the model match the average price elasticity of gasoline of –0.28 as estimated by recent empirical studies. Intuitively, due to non-elastic gasoline need $\bar{E}_i t$, the preference weight on consumption in the utility function φ impacts how much the household substitutes out of gasoline as relative prices change.

Table 3: Calibration

Parameter	Value	Description	Source
Set outside the model			
α	$\frac{2}{3}$	Production elasticity	Standard value
β	0.99	Hh discount factor	Standard value
ζ	0.73	EOS gas / consumption	Kehrig and Ziebarth (2009)
d	0.04	Depreciation rate durables	20% annual depreciation of cars
ρ_q	0.95	Persistence gas price shock	Gas price data
sd ε_t^q	0.08	S.d. gas price shock	Gas price data
ρ_e	0.994	Persistence Hh labor efficiency	Storesletten 2004
sd ε_{it}^e	0.014	Heterogeneity in labor efficiency	Storesletten 2004
sd α_i	0.0048	Heterogeneity in labor efficiency	$sd(\log e_i) = 0.9$
τ^L	0.2	Labor income tax	Tax receipts / labor income
τ^C	0.08	Consumption tax	Average US sales tax
Targeting moments (jointly):			
γ	0.895	Preference weight on D_i	Spending on durables vs gas
φ	0.901	Preference weight on C_i	Price elasticity of gasoline
$\bar{E}$	0.0237	(Mean of) minimum gas purchase	Mean expenditure share 5.5%
sd $\varepsilon_{it}^{\bar{E}}$	0.5496	(Mean of) minimum gas purchase	Gas expenditure heterogeneity

4 Quantitative Analysis

We now consider the quantitative effects of oil price fluctuations, starting with the impact on households' consumption bundles and then turning to the effects on welfare and the heterogeneity by income. Later, we will quantitatively assess the impact of oil gluts and carbon taxes proposed in the current policy debate.

4.1 Dynamic Responses to Oil Price Shocks

Figure 3 displays the aggregate impulse response functions[6] following an oil price shock of 33%, corresponding to an increase in the gas price from its long-run average of \$2.40 per gallon to \$3.18 per gallon. This increase causes households to lower their purchases of gasoline, but because the reduction is less than proportional, gas *expenditures* as a share of the overall budget increase. Overall, the magnitude of our results are in line with the estimated response of US aggregates to an oil supply disruption in the literature (see Figure 5 in Kilian (2009) and the more detailed work by Baumeister and Hamilton (2019)).

---

[6] The non-homotheticity in gas consumption exacerbates the non-linearity in households' policy functions, making impulse response functions state-dependent. We therefore obtain the impulse response functions through simulation, see notes to figure 3 for details.

Figure 3: Aggregate impulse response functions

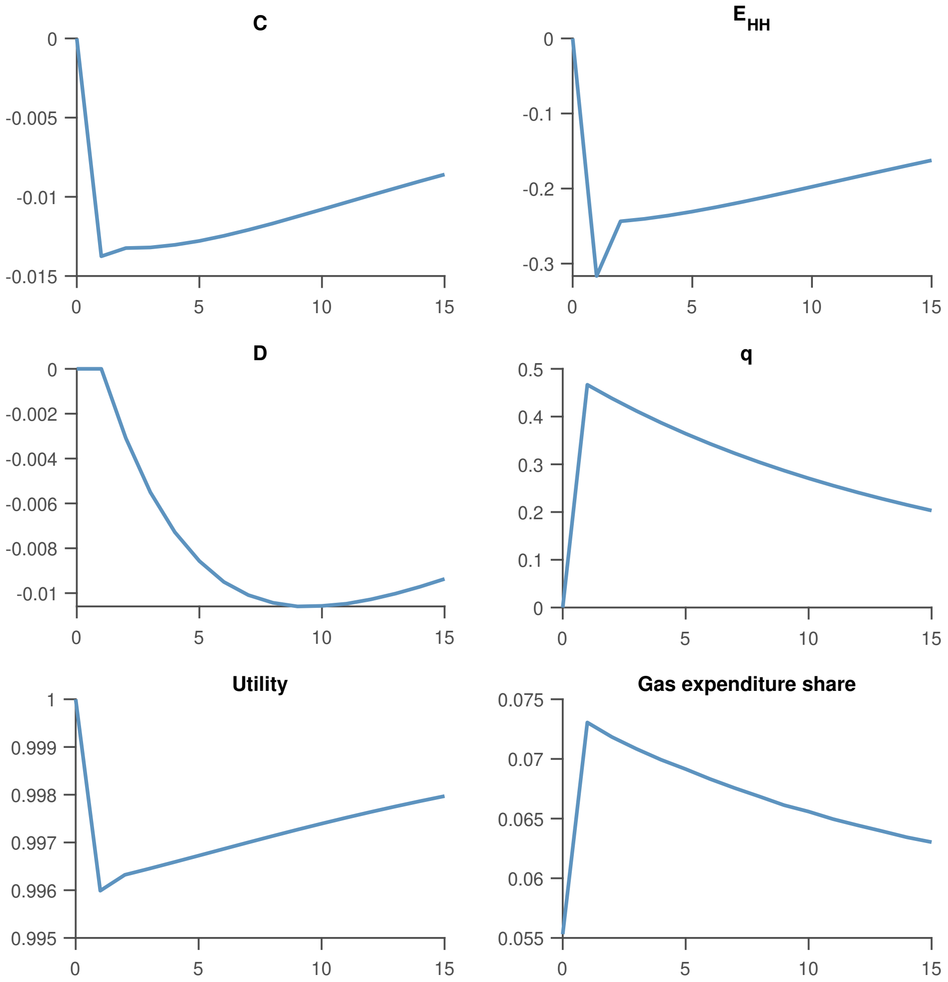

Aggregate impulse response functions following a 33% oil price shock. Impulse response functions were generated by simulating a panel of households subject to idiosyncratic and aggregate shocks ("baseline"), and then repeating the simulation with the same shock sequences plus an added 33% innovation to the aggregate shock ("impulse"). Differencing the two simulations gives the impulse response to an aggregate shock at a particular aggregate state. Because households' policy functions are state-dependent we repeat this procedure many times and average the aggregate responses together to obtain the displayed aggregate impulse response functions.

The economy's behavior in the aggregate masks considerable heterogeneity in the responses. Both a lower income and a higher baseline level of gas usage increases a household's share of gas expenditures and amplifies the effect of a given gas price shock. Figure 4 shows that while all household types reduce the utility-providing component of gas spending E^{HH} to a similar degree, their overall gas budget $E^{HH} + \bar{E}$ responds differently because of the fixed usage component. The households with high gas expenditure shares also have to reduce their consumption and durable stock drastically.

Figure 4: Impulse response functions by household types

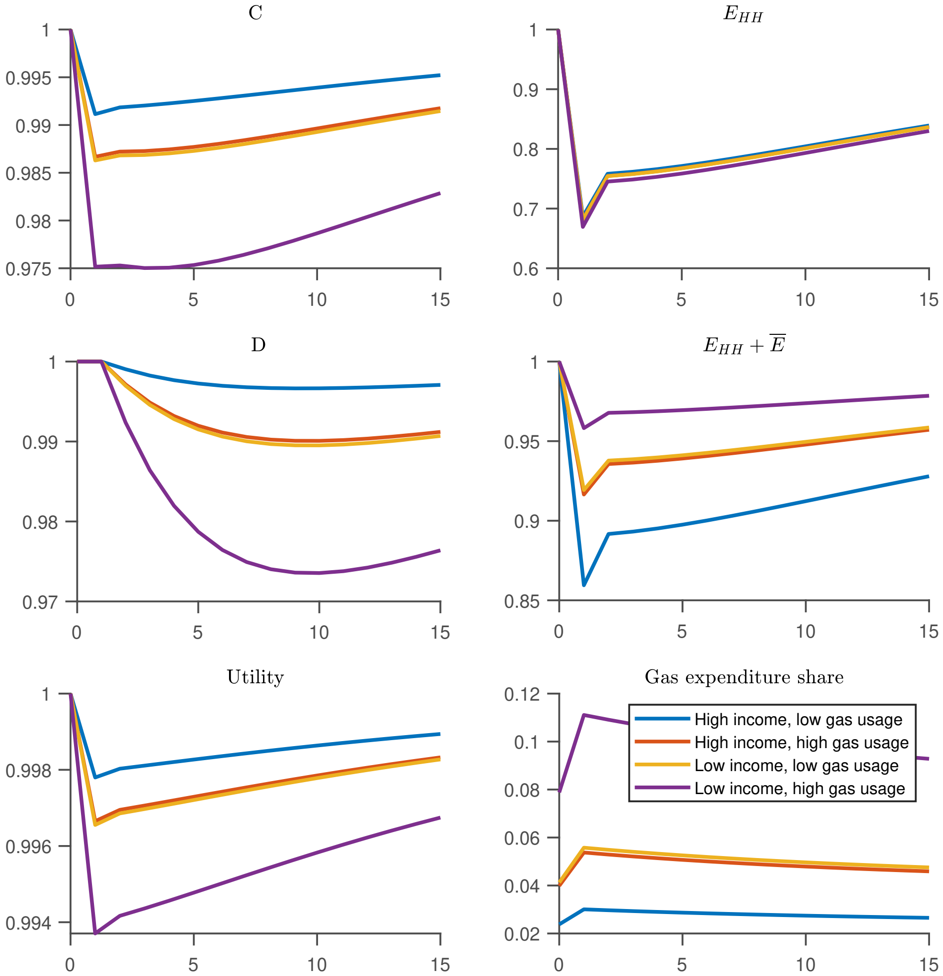

Impulse response functions for households at different points of the distribution of income and gas usage. "High" income and high usage households are at the respective 75th percentiles of the distribution over income and usage $\bar{E}_i$, and households with "low" income and usage are at the respective 25th percentiles.

4.2 The Long-run Welfare Cost of Oil Price Shocks

To evaluate the welfare effect of an oil price shock, we run a counterfactual in which we compare a one-time increase in the gas price to an increase in households' tax rate. To this end, we consider—starting from steady-state—a scenario in which a gas price shock hits as well as a counterfactual in which gas prices remain at steady state and in which a permanent tax hike is instituted such that households are indifferent between the two scenarios. In other words, we find the tax rate such that households' expected utility without a gas price shock equals the expected utility upon impact of the gas price shock under the original tax rate. Tables 5 and 6 contain these welfare-equivalent tax increases for labor and consumption taxes. Because in the following we will often evaluate households at those percentiles of income and gasoline usage distributions, respectively, Table 4 contains descriptive information about the joint distribution of those households.

Table 4: Descriptive statistics about model households at percentiles of distribution

	Percentile of income e_{it}				
Percentile of gas usage $\bar{E}_{it}$	10	25	50	75	90
10	0.27 (2.83%) [4.57%]	0.52 (3.47%) [3.94%]	1.00 (4.23%) [3.22%]	1.93 (3.47%) [2.55%]	3.73 (2.83%) [2.03%]
25	0.27 (3.47%) [5.97%]	0.52 (4.25%) [5.07%]	1.00 (5.18%) [4.05%]	1.93 (4.25%) [3.10%]	3.73 (3.47%) [2.36%]
50	0.27 (4.23%) [7.94%]	0.52 (5.18%) [6.68%]	1.00 (6.32%) [5.23%]	1.93 (5.18%) [3.88%]	3.73 (4.23%) [2.83%]
75	0.27 (3.47%) [10.73%]	0.52 (4.25%) [8.95%]	1.00 (5.18%) [6.90%]	1.93 (4.25%) [4.98%]	3.73 (3.47%) [3.50%]
90	3.73 (2.83%) [14.68%]	3.73 (3.47%) [12.16%]	3.73 (4.23%) [9.26%]	3.73 (3.47%) [6.55%]	3.73 (2.83%) [4.45%]

Notes: Relative income: Income relative to median-income household. Mass of Households in parentheses: Share of population close to the percentile (number indicates share of population for which given percentile is "nearest" percentile listed). Gas budget in brackets: Share of income spent on gasoline.

Table 5: Labor tax rate equivalents at different percentiles of the income distribution

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	0.93% (2.83%)	0.58% (3.47%)	0.38% (4.23%)	0.26% (3.47%)	0.20% (2.83%)
25	1.14% (3.47%)	0.71% (4.25%)	0.45% (5.18%)	0.30% (4.25%)	0.22% (3.47%)
50	1.44% (4.23%)	0.89% (5.18%)	0.54% (6.32%)	0.35% (5.18%)	0.25% (4.23%)
75	1.87% (3.47%)	1.13% (4.25%)	0.68% (5.18%)	0.43% (4.25%)	0.30% (3.47%)
90	2.47% (2.83%)	1.48% (3.47%)	0.88% (4.23%)	0.54% (3.47%)	0.36% (2.83%)

Notes: Hypothetical permanent changes in labor income tax that households would be willing to accept to avoid one-time gas price shock. In parentheses: Share of population close to the percentile (number indicates share of population for which given percentile is "nearest" percentile listed).

Table 6: Consumption tax rate equivalents at different percentiles of the income distribution

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	0.39%	0.36%	0.31%	0.27%	0.23%
25	0.49%	0.44%	0.37%	0.31%	0.26%
50	0.64%	0.56%	0.46%	0.37%	0.30%
75	0.85%	0.73%	0.59%	0.45%	0.35%
90	1.17%	0.99%	0.78%	0.58%	0.43%

An alternative way to interpret welfare losses is to consider proportional equivalent cuts to the utility-providing goods, which here consist of the bundle $\left(C, D, E^{HH}\right)$. Table 7 contains the permanent proportional cuts to all present and future consumption bundles that different households would be willing to accept instead of the one-time oil price shock in order to maintain their welfare.

Table 7: Proportional reductions in (C, D, E^{HH})

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−0.33%	−0.30%	−0.26%	−0.22%	−0.19%
25	−0.41%	−0.36%	−0.31%	−0.25%	−0.22%
50	−0.53%	−0.46%	−0.38%	−0.30%	−0.25%
75	−0.70%	−0.61%	−0.49%	−0.38%	−0.29%
90	−0.96%	−0.82%	−0.64%	−0.48%	−0.36%

4.3 The Importance of Heterogeneity

Along both dimensions of income rank and gas usage fixed effects, the absolute magnitude of welfare effects increases more than proportionally as a household moves up in the distribution; in other words, the welfare costs of an oil price increase are convex in the households' rank in the income and gas usage distributions. To see this, consider Tables 5–7. For example, for the median household, the gas price shock corresponds to a labor tax increase of 0.54%, but the mean of the distribution is 0.72%. Similarly, the median consumption equivalent is −0.38%, whereas the mean is −0.42%: low-income households and high-usage households would be willing to pay a particularly high price to avoid an increase in the cost of gas.

To see exactly how the welfare effects fall on different households, we also display their unconditional distribution in Figure 5. To construct it, we run a simulation over a long time horizon in which we allow a panel of households to be subject to aggregate and idiosyncratic shocks. Displayed in the graph are consumption equivalents for a potential large oil price shock hitting across time and across households, which approximates the unconditional distribution of welfare effects. Most notable is the long left tail: while the right half of the distribution is close to the median of −0.41%, the left half is spread out with a larger magnitude of welfare effects, suggesting that the utility impact of gas price shocks is distributed unevenly across the population.

4.4 Sensitivity Analysis

To see how the model results and, particularly, the welfare costs respond to individual assumptions about parameters, in this section we consider how they are affected by a) extending the duration of gas price shocks, b) including all energy products (like heating) in households' energy consumption instead of only gasoline, and c) varying the persistence of the individual income process.

For the first alternative calibration, we increase the persistence of the oil price shock to be close to unity with $\rho_q = 0.99$. All else equal, a higher persistence of gas prices will increase the welfare costs of oil price shocks across the distribution. While there is mixed evidence on whether oil prices follow a unit root process (e.g., Pestana Barros et al. (2014) and Elder and Serletis (2008)), Gelman et al. (2019) demonstrate that households' responses to gas price shocks are consistent with households *perceiving* gas prices as having a unit root. Relative to this evidence, our main analysis took a conservative approach to calibrating this parameter as it relates to the welfare costs of oil price shocks. Assuming this higher persistence, the consumption equivalent of a one-time gasoline shock for the median household increases in magnitude to −1.23% (Tables 8 and 9), a substantial change relative to the baseline.

Table 8: High gas price persistence: Labor tax equivalents to gas price shock

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	3.34%	2.12%	1.38%	0.97%	0.75%
25	3.90%	2.44%	1.56%	1.06%	0.80%
50	4.65%	2.87%	1.80%	1.20%	0.88%
75	5.64%	3.45%	2.12%	1.37%	0.97%
90	6.94%	4.20%	2.54%	1.60%	1.11%

Table 9: High gas price persistence: Proportional consumption equivalents to gas price shock

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−1.16%	−1.05%	−0.93%	−0.80%	−0.71%
25	−1.37%	−1.23%	−1.06%	−0.89%	−0.77%
50	−1.66%	−1.46%	−1.23%	−1.01%	−0.84%
75	−2.06%	−1.79%	−1.47%	−1.17%	−0.95%
90	−2.61%	−2.23%	−1.80%	−1.38%	−1.08%

We also consider including all energy products like natural gas, electricity, and heating oil in household preferences instead of only focusing on gasoline usage. As Fried (2018) argue, prices of other energy products are strongly correlated with the oil price, and they show that the share of all energy consumption still decreases in income. We follow their approach and calibrate energy consumption across households to average 12%, roughly double the expenditures on gasoline that we find for our main calibration focused only on gasoline expenditures. Tables 10 and 11 show that increasing this share also sizably raises the welfare costs of oil price shocks relative to the baseline.

Table 10: All energy products: Labor tax equivalents to gas price shock

	Percentile of income distribution				
Percentile of usage	10	25	50	75	90
10	3.11%	1.90%	1.16%	0.74%	0.52%
25	3.26%	1.99%	1.21%	0.77%	0.53%
50	3.42%	2.08%	1.26%	0.79%	0.55%
75	3.59%	2.18%	1.32%	0.83%	0.57%
90	3.77%	2.28%	1.38%	0.86%	0.59%

Table 11: All energy products: Proportional consumption equivalents to gas price shock

	Percentile of income distribution				
Percentile of usage	10	25	50	75	90
10	−1.19%	−1.02%	−0.83%	−0.64%	−0.51%
25	−1.26%	−1.08%	−0.87%	−0.67%	−0.53%
50	−1.33%	−1.14%	−0.91%	−0.70%	−0.54%
75	−1.42%	−1.20%	−0.96%	−0.73%	−0.56%
90	−1.50%	−1.27%	−1.01%	−0.76%	−0.58%

Both of these sensitivity analyses have highlighted two ways in which we view our calibration as conservative. As a final sensitivity analysis, we consider lowering the persistence of idiosyncratic shocks to households. Most studies find autocorrelation in the persistent component of income shocks very close to (or above) unity. One of the lowest estimates for the (annual) persistence parameter is Guvenen et al. (2019)'s 0.958, which translates to a quarterly value of 0.984 compared to 0.994 in our baseline. At the same time, the (quarterly) innovation to the persistent component in that specification is 0.09, higher than our baseline value of 0.014. Both lower persistence and higher variability of shocks to household income tend to reduce the heterogeneity of welfare effects since a currently low-income household can expect to revert toward a higher income soon. This in turn mitigates the negative impact of a gas price shock (and vice versa for a currently high-income household). Tables 12 and 13 show that this qualitative mechanism is at play. Quantitatively, this change in the calibration has only minor effects on the welfare costs and we conclude that our results are robust along this dimension.

Table 12: Less persistent/more variable shocks: Labor tax equivalents to gas price shock

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	0.71%	0.50%	0.36%	0.27%	0.21%
25	0.87%	0.61%	0.43%	0.31%	0.24%
50	1.10%	0.76%	0.53%	0.38%	0.28%
75	1.42%	0.98%	0.67%	0.47%	0.34%
90	1.88%	1.28%	0.86%	0.59%	0.42%

Table 13: Less persistent/more variable shocks:: Proportional consumption equivalents to gas price shock

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−0.31%	−0.28%	−0.25%	−0.23%	−0.21%
25	−0.39%	−0.34%	−0.30%	−0.26%	−0.23%
50	−0.50%	−0.44%	−0.38%	−0.32%	−0.28%
75	−0.66%	−0.57%	−0.48%	−0.40%	−0.33%
90	−0.90%	−0.77%	−0.64%	−0.52%	−0.42%

4.5 Quantitative Exercise 1: The Welfare Benefits of the 2014–2016 Oil Glut

From fall 2014 (2014Q4) to winter 2015/2016 (2016Q1), US gas prices dropped significantly over the course of five quarters, from \$3.54 to \$1.85. What were the welfare effects of this beneficial series of shocks for households? We feed the sequence of historic gasoline prices into our model and compare that baseline to a counterfactual scenario in which there is no beneficial oil price shock. In the counterfactual we use a sequence of zeros as gas price innovations, meaning that oil prices start reverting back to their mean over the same time horizon. Figure 6 displays the actual path of gas prices over this period together with the expected course of oil prices just before and after the period in fall 2014 and the beginning of 2016, respectively.

We use the historic gasoline prices from 1993 to the second quarter of 2014 as a "burn-in" period. Then we compare the utility of households at different income percentiles under both scenarios. To do this, we consider households that have stayed constant at the nth percentile of income until 2014. We then sum up realized flow utility derived from the six quarters 2014Q3 to 2016Q1 under the respective gas price shocks and add the expected future continuation utility from 2015Q4 onward. The drop in the oil price hence has two positive effects on households' utility: greater consumption during the time of low prices and higher expected continuation utility because prices are expected to remain low for a while after the shock.

We again evaluate the increase in welfare by considering the welfare-equivalent tax rates; that is, we are finding hypothetical tax cuts that would make households indifferent between paying the lower tax rate while not expecting any gas price shocks (i.e., they expect the gas prices to follow the green dashed path in Figure 6) and the utility they receive under the actually realized sequence of gas price shocks while paying the old tax rate. Tables 14 and 15 contain the tax cuts for labor and consumption taxes. Table 16 shows the equivalent increases in the consumption bundles of $\left(C, D, E^{HH}\right)$.

As a way of comparison, we report in the bottom row the change in average tax rates following the passage of the 2017 Tax Cut and Jobs Act (TCJA). We note here a few complications in making this comparison: (1) the Treasury reports we use do not report rates for particular percentiles but *averages within groups of percentiles*; (2) the changes here are not specific to the labor tax but reflect changes in tax rates on all types of income;[7] (3) changes in average rates are not solely due to changes in statutory rates but also to changes in the pre-tax income distribution itself; and (4) the reductions in income tax rates are only temporary. With these caveats in mind, we report the percentage point change in these average rates. Focusing on the median gas user, we find that the welfare effects of the gas glut expressed in labor tax rate cuts are larger than the tax cuts of the TCJA for all incomes up to 50,000 USD (roughly the bottom 40 percent of the income distribution). At the median income level, the gas glut benefits are still more than half of the benefits of the TCJA. For higher income levels, this comparison starts to pale as the TCJA was more generous toward high earners, who are less affected by gasoline price changes.

---

[7] The changes in payroll taxes, the closest analog to our labor tax, were minor and hence not an interesting point of comparison.

Table 14: Historic counterfactual: Labor tax rate equivalents

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−1.17%	−0.75%	−0.49%	−0.34%	−0.26%
25	−1.44%	−0.90%	−0.58%	−0.39%	−0.29%
50	−1.81%	−1.12%	−0.70%	−0.46%	−0.33%
75	−2.34%	−1.43%	−0.88%	−0.56%	−0.39%
90	−3.07%	−1.87%	−1.13%	−0.70%	−0.47%
TCJA	−0.5%	−0.9%	−1.3%	−1.5%	−1.5%

Note: The TCJA row reports percentage point changes in average tax rates between 2017 and 2018. Incomes at the appropriate precentiles were computed from the Current Population reports (Fontenot et al. (2018)) and average labor tax rates were obtained from Table 2 in Joint Committee on Taxation (2017).

Table 15: Historic counterfactual: Consumption tax rate equivalents

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−0.50%	−0.46%	−0.40%	−0.35%	−0.31%
25	−0.63%	−0.56%	−0.48%	−0.40%	−0.34%
50	−0.80%	−0.71%	−0.59%	−0.48%	−0.39%
75	−1.06%	−0.92%	−0.76%	−0.59%	−0.47%
90	−1.46%	−1.25%	−0.99%	−0.75%	−0.57%

Table 16: Historic counterfactual: Proportional increases in (C, D, E^{HH})

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−0.42%	−0.38%	−0.34%	−0.29%	−0.26%
25	−0.53%	−0.47%	−0.40%	−0.34%	−0.29%
50	−0.68%	−0.60%	−0.50%	−0.40%	−0.33%
75	−0.90%	−0.78%	−0.63%	−0.50%	−0.39%
90	−1.23%	−1.05%	−0.84%	−0.63%	−0.48%

4.6 Quantitative Exercise 2: The Welfare Cost of Higher Gasoline Taxes

There is a policy debate in US on how to reduce carbon emissions to help stem climate change. A key ingredient in the so called "New Green Deal" proposal is a carbon tax that increases the price of gasoline. Such taxes are common in Europe. Given the non-homothetic nature of gasoline demand, gasoline taxes would be regressive in nature and thus hurt low-income households the most. What would be the welfare cost of higher gasoline taxes? Studying this question is difficult for a number of reasons: First, and in contrast to the cyclical gas price fluctuations we have considered in the previous sections, for permanent increases in the gas price, we expect stronger adjustments in household gasoline consumption. In other words, the price elasticity of gasoline is likely to be higher than at business cycle frequency, because households in the long run are more likely to, for example, switch to less energy-intensive modes of transportation or consider transportation costs in their location choices. Second, policymakers may be aware of the distributional impact of higher gasoline taxes and accompany them with compensating measures that may be progressive such as reductions in taxes for low-income households, offsets in other taxes or fees (like per-vehicle taxes), or transportation infrastructure or subsidies. Our model abstracts from either feature. With that limitation in mind, we do not consider gasoline tax levels from high-income European countries[8], which some US policy makers may have in mind. Rather, we take the average European gasoline tax, which corresponds to a gasoline price increase of about 35 cents/gallon or an increase of 14.5% in the long-run average $\tilde{q}$, and study what labor tax rate hikes households would be willing to pay to avert such a New Green Deal. This is a more conservative exercise than what is being discussed in the US, which diminishes direct welfare effects.[9] On the other hand, the absence of substitution for alternative means of transportation and redistributive subsidies in our model make the quantitative results larger. We hence consider the values listed in Tables 17 and 18 as a first pass to assess the welfare effects of permanent gas price increases.

---

[8] For example, Germany charges a gasoline tax of about 2.50 Euros per gallon of gasoline.

[9] Bento et al. (2009) consider a similar 25 cents/gallon tax increase.

Table 17: Permanent gasoline tax: Labor tax rate equivalents

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	2.17%	1.35%	0.85%	0.57%	0.43%
25	2.73%	1.67%	1.03%	0.67%	0.48%
50	3.52%	2.12%	1.28%	0.81%	0.56%
75	4.66%	2.76%	1.63%	1.00%	0.67%
90	6.29%	3.68%	2.13%	1.27%	0.82%

Notes: Labor tax changes equivalent to an increase in the long-run average of the gas price by $0.35/gal from $2.40 to $2.75.

Table 18: Permanent gasoline tax: Proportional decreases in (C, D, E^{HH})

	Percentile of income distribution				
Percentile of gas usage	10	25	50	75	90
10	−0.77%	−0.68%	−0.59%	−0.49%	−0.42%
25	−0.98%	−0.85%	−0.71%	−0.57%	−0.47%
50	−1.29%	−1.10%	−0.89%	−0.70%	−0.55%
75	−1.75%	−1.47%	−1.16%	−0.87%	−0.66%
90	−2.46%	−2.03%	−1.56%	−1.13%	−0.83%

Notes: Consumption equivalents to an increase in the long-run average of the gas price by $0.35/gal from $2.40 to $2.75.

Our welfare calculations and our qualitative about the regressive incidence are comparable to the other papers. Goulder et al. (2019) quantify the welfare cost of carbon taxes in general equilibrium, but in our quantitative analysis we consider that such taxes directly affect consumer durables, which will propagate the effects and make them more persistent. We also differ in that we consider richer household heterogeneity in incomes and gasoline usage. As the above tables show, welfare costs between the top and bottom decile of commuting needs are more than three-fold, thus highlighting the importance of using the joint empirical distribution of gasoline and income heterogeneity to discipline our model.

5 Summary and Future Work

The model of household gasoline consumption outlined in this paper mirrors the empirical finding that high-income households allocate a lower share of their budget to gasoline consumption. Calibrated to the CEX data on household income and expenditures, the model predicts that oil price shocks have direct effects on welfare that are larger for households in the lower half of the income distribution than for those in the upper half. For example, a one-time increase in the gas price by 33% from \$2.40 to \$3.18 would be equivalent to a permanent tax increase of 0.26 percentage points for a household at the 90th income percentile but equivalent with a hike of 0.36 percentage points for a household at the 10th percentile. Similarly, the drop in the gas price during the year 2015 that extended from fall 2014 to early 2016 was approximately equivalent to a permanent tax cut of 0.51 and 0.72 percentage points for these high- and low-income households, respectively.

Future work may consider differences in fuel efficiency of consumer durables powered by gasoline and associated resale frictions in secondary markets. For example, households may want to sell gas guzzling vehicles during times of high gasoline prices such as 2008, but this will be exactly when demand for such vehicles is low. This means that fuel inefficient cars are least liquid when it is most costly to operate them. This will likely exacerbate the welfare cost of oil price shocks.

References

Robert B Barsky and Lutz Kilian. Oil and the macroeconomy since the 1970s. *Journal of Economic Perspectives*, 18(4):115–134, 2004. ISSN 0895-3309. doi: 10.1257/0895330042632708. (Cited on page 3.)

Christiane Baumeister and James D. Hamilton. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910, May 2019. (Cited on page 17.)

Christiane Baumeister and Lutz Kilian. Do oil price increases cause higher food prices? *Economic Policy*, 29(80):691–747, October 2014. (Cited on page 3.)

Antonio M. Bento, Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. Von Haefen. Distributional and efficiency impacts of increased US gasoline taxes. *American Economic Review*, 99(3):667–699, 2009. ISSN 00028282. doi: 10.1257/aer.99.3.667. (Cited on pages 4, 13, and 28.)

Olivier J Blanchard and Marianna Riggi. Why are the 2000s so different from the 1970s? a structural interpretation of changes in the macroeconomic effects of oil prices. *Journal of the European Economic Association*, 11(5):1032–1052, 2013. ISSN 15424766. doi: 10.1111/jeea.12029. (Cited on page 3.)

Martijn Brons, Peter Nijkamp, Eric Pels, and Piet Rietveld. A meta-analysis of the price elasticity of gasoline demand. a SUR approach. *Energy Economics*, 30(5):2105–2122, 2008. ISSN 01409883. doi: 10.1016/j.eneco.2007.08.004. (Cited on page 4.)

Jared Carbone, Richard Morgenstern, Roberton Williams III, and Dallas Burtraw. Deficit reduction and carbon taxes: Budgetary, economic, and distributional impacts. Considering a Carbon Tax: A Publication Series from RFF's Center for Climate and Electricity Policy, 2013. (Cited on page 4.)

Mireille Chiroleu-Assouline and Mouez Fodha. From regressive pollution taxes to progressive environmental tax reforms. *European Economic Review*, 69:126–142, 2014. (Cited on page 4.)

Javier Diaz-Gimenez, Andy Glover, and Jose-Victor Rios-Rull. Facts on the distributions of earnings, income, and wealth in the united states: 2007 update. *Federal Reserve Bank of Minneapolis Quarterly Report*, 34(1), 2011. (Cited on page 15.)

Jeffrey A. Dubin and Daniel L. McFadden. An econometric analysis of residentail electric appliance holdings and consumption. *Econometrica*, 52:345–362, 1984. (Cited on page 13.)

Susan Dynarski, Jonathan Gruber, Robert A Moffitt, and Gary Burtless. Can families smooth variable earnings? *Brookings Papers on Economic Activity*, 1997(1):229–303, 1997. (Cited on page 10.)

Paul Edelstein and Lutz Kilian. How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics*, 56(6):766–779, 2009. ISSN 03043932. doi: 10.1016/j.jmoneco.2009.06.001. URL `http://dx.doi.org/10.1016/j.jmoneco.2009.06.001`. (Cited on pages 3 and 5.)

John Elder and Apostolos Serletis. Long memory in energy futures prices. *Review of Financial Economics*, 17(2):146–155, jan 2008. ISSN 10583300. doi: 10.1016/j.rfe.2006.10.002. URL `http://doi.wiley.com/10.1016/j.rfe.2006.10.002`. (Cited on page 23.)

Kayla Fontenot, Jessica Semega, and Melissa Kollar. Income and poverty in the United States: 2017. *Current Population Reports, P60-263*, September 2018. (Cited on page 27.)

Stephie Fried. Climate policy and innovation: A quantitative macroeconomic analysis. *American Economic Journal: Macroeconomics*, 10(1):90–118, January 2018. (Cited on page 23.)

Stephie Fried, Kevin Novan, and William B. Peterman. The distributional effects of a carbon tax on current and future generations. *Review of Economic Dynamics*, 30(2016):30–46, 2018. ISSN 10942025. doi: 10.1016/j.red.2018.02.001. (Cited on page 4.)

Michael Gelman, Yuriy Gorodnichenko, Shachar Kariv, Dmitri Koustas, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. The response of consumer spending to changes in gasoline prices. *NBER Working Paper No. 22969*, 2019. URL `http://www.nber.org/papers/w22969`. (Cited on page 23.)

Kenneth Gillingham. Identifying the elasticity of driving: Evidence from a gasoline price shock in california. *Regional Science and Urban Economics*, 47:13–24, July 2014. (Cited on page 5.)

Mikhail Golosov, Josh Hassler, Per Krusell, and Aleh Tsyvinski. Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82:41–88, 2014. (Cited on page 3.)

Lawrence H. Goulder, Marc A. C. Hafstead, GyuRim Kim, and Xianling Long. Impacts of a carbon tax across US household income groups: What are the equity-efficiency trade-offs? *Journal of Public Economics*, 175:44–64, July 2019. (Cited on pages 4 and 29.)

Fatih Guvenen and Serdar Ozkan. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, 2014. (Cited on page 15.)

Fatih Guvenen, Fatih Karahan, Serdar Ozkan, and Jae Song. What do data on millions of U.S. workers reveal about life-cycle earnings risk? *Working Paper*, 2019. doi: 10.2139/ssrn.2646033. (Cited on pages 15 and 24.)

James D. Hamilton. What is an oil price shock? *Journal of Econometrics*, 113(2):363–398, April 2003. (Cited on page 3.)

Kevin A. Hassett, Aparna Matur, and Gilbert E. Metcalf. The incidence of a U.S. carbon tax: A lifetime and regional analysis. NBER Working Paper 13554, 2007. (Cited on page 3.)

Tomas Havranek and Ondrej Kokes. Income elasticity of gasoline demand: A meta-analysis. *Energy Economics*, 47:77–86, 2015. (Cited on page 5.)

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy*, 118 (4):681–722, 2010. ISSN 00223808. URL `http://www.nber.org/papers/w14052`. (Cited on page 15.)

Jonathan E. Hughes, Christopher R. Knittel, and Daniel Sperling. Evidence of a shift in the short-run price elasticity of gasoline demand. *Energy Journal*, 29(1):113–134, 2008. ISSN 01956574. doi: 10.3386/w12530. (Cited on page 5.)

Joint Committee on Taxation. Distributional effects of the conference agreement for H.R. 1, the "Tax Cuts and Jobs Act". *JCX-68-17*, December 2017. URL `https://www.jct.gov/CMSPages/GetFile.aspx?guid=45580637-252a-4c6f-b434-a2085f417920`. (Cited on page 27.)

Hilke A. Kayser. Gasoline demand and car choice: Estimating gasoline demand using household information. *Energy Economics*, 22(3):331–348, 2000. ISSN 01409883. doi: 10.1016/S0140-9883(99)00043-2. (Cited on page 10.)

Matthias Kehrig and Nicolas Ziebarth. Why do oil price shocks matter? transmission mechanisms on the supply and the demand side. *Working Paper*, 2009. (Cited on pages 15 and 17.)

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, June 2009. (Cited on page 17.)

Lutz Kilian and Robert Vigfusson. The role of oil price shocks in causing U. S. recessions. *Journal of Money, Credit and Banking*, 49, December 2017. (Cited on page 3.)

Xavier Labandeira, Jose M. Labeaga, and Xiral Lopez-Otero. A meta-analysis on the price elasticity of energy demand. *Energy Policy*, 102:542–568, 2017. (Cited on page 4.)

Andrew J. Leach. The welfare implications of climate change policy. *Journal of Environmental Economics and Management*, 57:151–165, 2009. (Cited on page 4.)

Laurence Levin, Matthew S. Lewis, and Frank A. Wolak. High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–347, August 2017. (Cited on page 5.)

Gilbert E. Metcalf. A distributional analysis of green tax reforms. *National Tax Journal*, 52: 655–681, 1999. (Cited on page 3.)

Austin Nichols and Seth Zimmerman. Measuring trends in income variability. *Working Paper*, pages 1–58, May 2008. (Cited on page 15.)

William D. Nordhaus and Joseph Boyer. *Warming the World: Economic Modeling of Global Warming*. MIT Press, 2000. (Cited on page 3.)

Carlos Pestana Barros, Luis Alberiko Gil-Alana, and James E. Payne. Long range dependence and breaks in energy prices. *Energy Sources, Part B: Economics, Planning, and Policy*, 9 (2):196–206, apr 2014. doi: 10.1080/15567249.2012.753959. (Cited on page 23.)

James M. Poterba. Is the gasoline tax regressive? *Tax Policy and the Economy*, 5:145–164, 1991. (Cited on page 7.)

Sebastian Rausch. Fiscal consolidation and climate policy: An overlapping generations perspective. *Energy Economics*, 40:S134–S148, 2013. (Cited on page 4.)

Robert Ready, Nikolai Roussanov, and Evelina Zurowska. Why does oil matter? commuting and aggregate fluctuations. *Working Paper*, 2020. (Cited on page 4.)

Robert C. Ready. Oil consumption, economic growth, and oil futures: The impact of long-run oil supply uncertainty on asset prices. *Journal of Monetary Economics*, 94:1–26, April 2018. (Cited on page 4.)

Richard Schmalensee and Thomas M. Stoker. Household gasoline demand in the United States. *Econometrica*, 67(3):645–662, 1999. ISSN 0012-9682. doi: 10.1111/1468-0262.00041. (Cited on page 5.)

Kristin N Sipes and Robert Mendelsohn. The effectiveness of gasoline taxation to manage air pollution. *Ecological Economics*, 36(2):299–309, 2001. (Cited on page 10.)

Nicholas Stern. *The Economics of Climate Change: The Stern Review*. Cambridge University Press, 2007. (Cited on page 3.)

Kjetil Storesletten, Christopher I. Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, apr 2004. ISSN 0304-3932. doi: 10.1016/J.JMONECO.2003.06.005. URL `https://www.sciencedirect.com/science/article/abs/pii/S0304393204000182`. (Cited on page 15.)

Zia Wadud, Daniel J. Graham, and Robert B. Noland. Gasoline demand with heterogeneity in household responses. *Energy Journal*, pages 47–74, 2010. (Cited on page 5.)

Sarah E. West. Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics*, 88(3-4):735–757, March 2004. (Cited on page 7.)

Sarah E. West and Roberton C. Williams. Estimates from a consumer demand system: Implications for the incidence of environmental taxes. *Journal of Environmental Economics and Management*, 47:535–558, 2004. (Cited on page 5.)

Roberton Williams, Hal Gordon, Dallas Burtraw, and Jared Carbone. The initial incidence of a carbon tax across income groups. *National Tax Journal*, 68:195–214, 2015. (Cited on page 4.)

Adonis Yatchew and Joungyeo Angela No. Household gasoline demand in Canada. *Econometrica*, 69:1697–1907, 2001. (Cited on page 5.)

A Data

We use the public-use portion of the Bureau of Labor Statistics' Consumer Expenditure Survey. It contains quarterly expenditures for "Gasoline and Motor Oil" as well as a measure of households' total quarterly expenditures (expenditures for food, housing, apparel, transportation, health, entertainment, personal care, education, cash contributions and insurance) and we combine these two to construct our main measure of the households gasoline expenditure shares by dividing gasoline expenditures over all expenditures. We remove observations for which households spend more than 70% of their quarterly on gasoline (this adjustment has no impact on results).

The CEX also provides some income data, albeit only on an annual frequency, and not for all households. We use this income data in two ways: we construct gasoline expenditures as a share of after-tax income (instead of as share of total expenditures) as a robustness check. Second, we use the provided information on households' income rank as the main way to graph the negative relationship between income and gas expenditure shares. While we use income rank instead of absolute income because it is more robust to measurement error, results are very similar when using absolute income instead.

We use several additional variables as controls. We use dummies for geographic region (Midwest, Northeast, South, West), metropolitan area status, race, gender, education (some high school, completed high school, some college, completed college, more than college), marital status (married, never married, separated, divorced, widowed) and family size $(1, 2, 3, 4, 5+)$. Where applicable, these attributes refer to the household's reference person as defined by the CEX. We also include a quadratic in the reference person's age, as well as quarterly time dummies.

B Robustness check: All energy products

Because prices across different types of energy products (like gasoline, natural gas, and fuel oil) are correlated, and because households may to some degree substitute between different energy products, as a robustness check we consider how income is related to the share of expenditures spent on a comprehensive measure of energy products. We find the results under this broader measure of energy usage to be very consistent with the our main specification of using only gasoline. Specifically, in the following we repeat some of the key analyses from section 2 using the CEX information on total combined expenditures on gasoline, natural gas, heating fuels, and electricity (the "all-energy" specification). We again exclude observations with no energy expenditures reported (2.7% of the sample). Table 19 shows summary statistics for expenditures on "all energy". These expenditures are on average roughly twice as large as the expenditures on gasoline alone (a mean expenditure share of 10.6% compared to 5.6%, respectively).

Table 19: Summary statistics for all energy expenditures

	Average	StD	Q_{10}	Median	Q_{90}	N
Energy expenditure share (%)	10.3	6.9	3.1	9.0	19.1	446,016
(only positive usage)	10.6	6.8	3.6	9.2	19.3	433,805
Gasoline expenditure share (%)	5	4.59	0	3.98	10.5	446,025
(positive usage)	5.59	4.5	1.52	4.44	10.96	399,281

Author calculations based on data from the CEX. Values for gasoline expenditures repeated from table 1 for comparison.

Next, we verify that the negative association between income and gasoline expenditure shares holds when using all energy expenditures. Figure 7 plots a bivariate local line of best fit between income and energy expenditures showing a robust negative relationship. This is analogous to the relationship between income and gasoline expenditures displayed in figure 1.

Finally, we repeat the regressions of table 2 using the all-energy expenditure share as dependent variable. Across specifications we find the coefficient of interest to be approximately twice as large as before using only gasoline expenditures.

Figure 5: Distribution of welfare effects

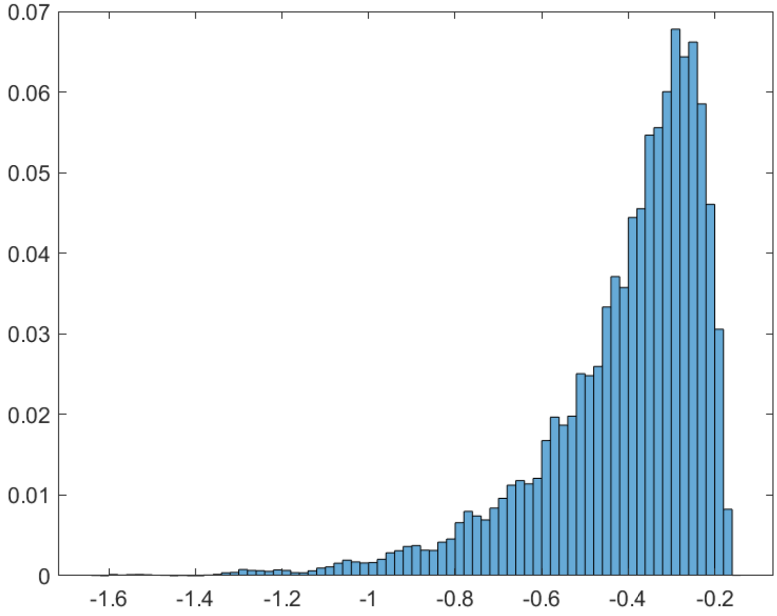

Distribution of welfare effects in a simulation of a panel of households.

Figure 6: Historic gas prices

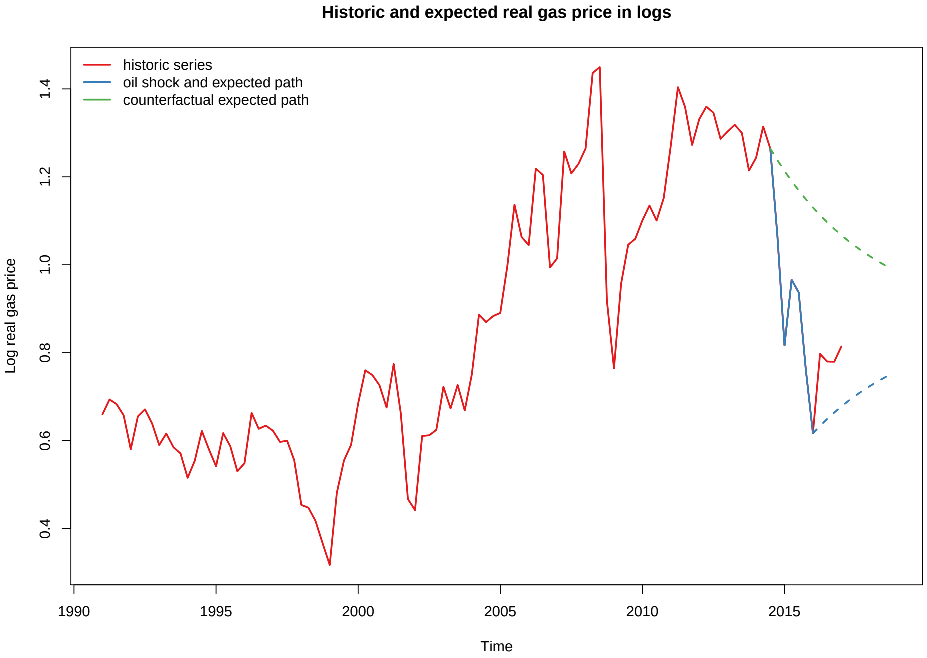

Notes—Red line: Historic gas prices. Blue line: Series of beneficial gas price shocks in 2015 (solid) and expected gas price trajectory (dotted). Green line: Expected gas prices absent the beneficial gas price shocks.

Figure 7: Negative Income Elasticity of All Energy Expenditures

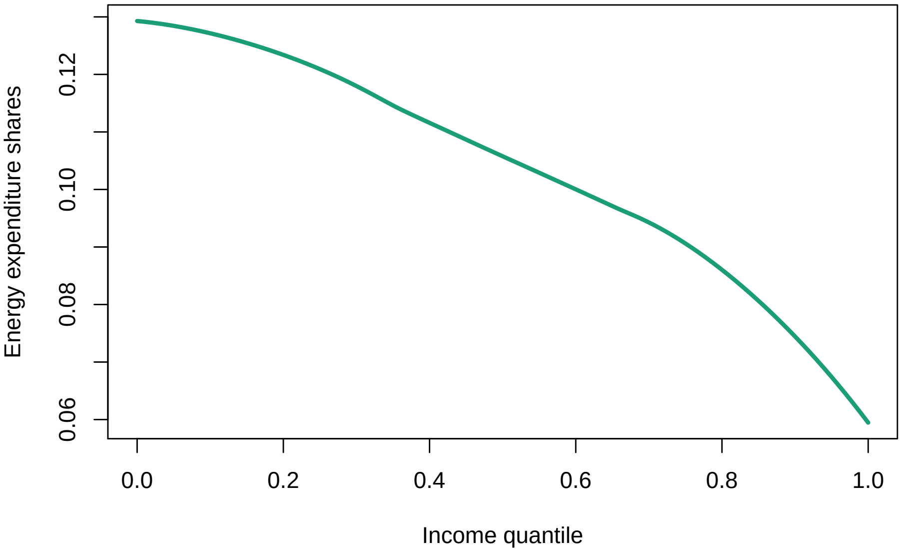

Expenditure shares of all energy consumption conditional on positive usage.

Table 20: Estimates of Income Elasticities for All Energy Expenditures

	Dependent variable: Energy expenditure share						
	OLS			Fixed effects		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income Rank	-0.0498^{***} (0.0005)			$-.0277^{***}$ (0.0011)		-0.0792^{***} (0.0016)	
log(Income)		-0.0086^{***} (0.0001)					-0.0193^{***} (0.0004)
log(Expenditures)			-0.0453^{***} (0.0001)		-0.0613^{***} (0.0002)		
Observations	303,510	323,334	432,592	304,623	433,715	63,567	67,256
R^2	0.2116	0.1992	0.3108				

Controls not shown included in OLS and IV regressions: Quarterly time effects, region, metropolitan area, gender, race, age, education, marital status and family size. Columns (6) and (7) use the log of lagged income as an instrument.