Welfare effects of gas price fluctuations

Florian Kuhn* Matthias Kehrig Nicolas Ziebarth

May 14, 2017

Abstract

A large literature examines the effects of oil price shocks on aggregate output through its role in production or its impact on monetary policy. Motivated by the fact that a large share of U.S. oil consumption occurs in the household sector in the form of gasoline, in this paper we follow Kehrig and Ziebarth (2009) by explicitly modeling household gasoline consumption. We structure household behavior to replicate two patterns found in household-level data (the CEX) which show that gasoline consumption increases with income, but it decreases with income as a share of the household's budget along the intensive margin. The model includes gasoline consumption in household utility (e.g. taking road trips) on top of a fixed minimum level of gas consumption (e.g. commuting to work). This allows us to study the direct effect of an oil price shock on household welfare, as well as its distributional consequences. Calibrated to households' gasoline expenditures, the model suggests that a shock to the gas price is almost twice as costly for households in the lower half of the income distribution than for high-income households.

JEL codes:
Keywords:

*Department of Economics, Binghamton University, email: fkuhn@binghamton.edu

1 Introduction

In this paper we examine how oil price shocks can directly affect household welfare, and how these effects vary with household income. To this purpose we focus on households' gasoline consumption which is the main end to which households purchase (processed) oil. We evaluate the relationship between gas expenditures and household income empirically using data from the Consumer Expenditure Survey. In that dataset 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. We incorporate this finding in an equilibrium model with two types of households who differ in their labor income. In particular, the decreasing propensity to consume gasoline is introduced via a fixed minimum quantity of gasoline that must be consumed by all households. This inelastic part of a household's gasoline consumption can be interpreted as required for commuting to work. Any quantity of gasoline consumed beyond this minimum level enters households utility as a complement to an output good which represents the remainder of the consumption basket.

The model is then calibrated to match the differences in household consumption and gasoline expenditures between the top and the bottom half of the income distribution. We examine the welfare effects of a shock to the gasoline price on either type of household by comparing a one-time gasoline price shock to a permanent change in the steady-state labor tax. Poor households' welfare is almost twice as sensitive to the gasoline price as the welfare of rich households: For example, a temporary increase of the gasoline price from \$2 to \$3 is equivalent to a permanent labor tax hike of 0.5 percentage points, whereas for poor households it is equivalent to a hike of 0.95 percentage points.

There exists an ample literature about the effects of oil price shocks on output, inflation and the conduct of monetary policy. 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) and Barsky and Kilian (2004), there is some evidence that an important channel for the effect of the oil price 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, e.g. Blanchard and Riggi (2013) or Kilian and Vigfusson (2014).

A second strand of the literature uses a more reduced-form approach to investigate the reaction of consumer spending on energy-related and other items in response to oil-price shocks. For example, there are many estimates of gasoline demand elasticity in the 1970s and 1980s (see Brons et al. (2008) for a meta-analysis of such studies). More recently, Schmalensee and Stoker (1999) employ semi-parametric methods to estimate gasoline consumption as a function of household characteristics, whereas Edelstein and Kilian (2009) and Hughes et al. (2008) infer the price elasticity of oil demand from aggregate data. Bento et al. (2009) examine the effect of gasoline taxes on vehicle purchase decisions in a dynamic discrete choice model.

This paper adds to the literature by including the direct effect of oil price shocks on household welfare in a dynamic setting. Additionally, it takes into account household heterogeneity in gasoline usage. Finally, to think about welfare implications, it focuses not on households' absolute gasoline consumption, but rather on the share of gasoline expenditures on total expenditures.

2 Household data

We use data from the Consumer Expenditure Survey (CEX) to investigate how gasoline consumption varies with household income. Throughout this section the main focus is going to be on the "gasoline budget", i.e. gasoline expenditures as a share of total household expenditures or of household income in the given period. This means in particular that we do not refer to absolute levels of gasoline consumption unless mentioned explicitly.

The main result is that, on the intensive margin, gasoline consumption is negatively related to income. This is in contrast to the unconditional household gasoline budget which is a hump-shaped function of income. However, 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 households owning one or more automobiles yields a negative relationship between income and gasoline budget. We then focus on this intensive-margin relationship between income and gasoline consumption and estimate its linear regression coefficient on the set of households that choose positive gasoline consumption.

There are two main reason why in the structural model below we focus on the gasoline-using households (the "intensive-margin" households). The first is that this abstracts from two discrete margins and simplifies the analysis: Not owning a car is well predicted by low household income and non-employment. This suggests that households are constrained by both their discrete employment status and the non-divisibility of cars. Second, and relatedly, while gas prices may well influence the households' discrete decision of owning a car (and hence may affect non-owners' welfare), this channel is hard to identify in the CEX given that car purchase decisions are relatively rare and the dataset's panel dimension is relatively short.

2.1 Dataset

The CEX is a rotating panel of households in which households remain up to 4 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[1]. The interview data contains information on quarterly expenditures in a large number of categories, of which the main item of interest is expenditures for "Gasoline and Motor Oil". 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 is collected only in the first and in the fourth interview, 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.

We construct a measure for gasoline usage (the "gas budget") by dividing quarterly gasoline expenditures by the households total quarterly expenditures, which will be used as the main dependent variable below. Information on household income is mainly used as an independent variable, although we also consider the total expenditures as a proxy for income as a regressor because of the aforementioned concerns about the availability of the income data. Because of the limited number of observations per household along the time dimension, we pool all observations. Variation in income hence overwhelmingly comes from cross-sectional differences.

[1] Data available at http://www.bls.gov/cex/pumdhome.htm

2.2 Summary statistics and distribution of gasoline consumption

Figure 3 displays histograms of gasoline consumption with absolute gas consumption in gallons shown in the top panel and gas consumption as a share of total expenditures at the bottom (gas consumption relative to income looks very similar). The plots resemble a log-normal distribution with an added mass point at 0, which is the notable 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.

Table 1 lists how the incidence of zero gasoline expenditures varies between households. Around 85% of households consume gasoline in every observed quarter, whereas around 9% never buy gasoline. For only circa 5% of households at least one quarter with positive and one quarter of zero gasoline consumption is observed.

Table 1: Positive gasoline usage by household

never	always	sometimes
0.0896	0.8557	0.0547

Figure 4 then excludes the zero observations and replots both gasoline budgets as well as their log-transforms. Mildly interestingly, not only is gasoline consumption in gallons considerably right-skewed, so are the shares of income and total expenditures spent on gasoline. The log transforms of the budget shares are distributed more symmetrically.

This can also be observed in Table 2 which lists conditional and unconditional summary statistics for the two gasoline budgets: For the log-transformed gasoline budgets, mean and median are very close.

Table 2: Summary statistics for gasoline consumption

Statistic	N	Mean	St. Dev.	Median
Consumption, in Gal	446,114	202.3	195.6	157.3
" (positive usage)	399,301	226.0	193.4	177.2
Expenditure budget, in pct	446,025	5.0	4.6	4.0
" (positive usage)	399,281	5.6	4.5	4.4
" (log-transformed)	399,281	−3.2	0.8	−3.1
Income budget, in pct	324,902	6.5	11.0	3.4
" (positive usage)	291,268	7.3	11.4	3.9
" (log-transformed)	291,268	−3.2	1.0	−3.2

2.3 Nonparametric bivariate relationships

2.3.1 Conditional on positive usage

In this subsection we display the (bivariate) relationship of the gasoline budget with household income and a few other variables. As mentioned in section 2.1 in the analysis we mainly focus on the gas budget as the share of gas expenditures on total expenditures rather than as a share of income because total expenditures are observed every quarter and income only annually. Additionally it allows us to directly focus on how the household splits up its consumption basket between gasoline and other items.

Figure 5 displays a nonparametric relationship between the gasoline budget along the y-axis and income quantile and total expenditures, respectively, along the x-axes. Most notable is the graphs' hump shape: For low levels of income, the conditional expectation of gas expenditures increases with income. For incomes higher than approximately median or more than around $\$8,000$ of quarterly expenditures, the relationship reverses such that further increasing the income level predicts a lower gasoline budget.

Importantly, the initial increase in the gas budget for low income levels is driven entirely by the extensive margin, i.e. households that do not to buy any gasoline. Excluding such households, which is is done in Figure 6, leads to a consistently negative relationship in the nonparametric estimation. As one may expect, this is because non-usage is heavily concentrated among low-income households.

Figure 7 summarizes all observations with positive gas expenditures by binning the data into two-dimensional frequency counts where bins are defined by income quantile and gas budget. Most notable is the large variance of gas budgets for households with lower income, whereas for high-income households gas expenditures are more narrowly concentrated.

2.3.2 Conditional on demographic group

One concern is whether the observed negative relationship is a compositional effect rather than a direct association with household income. For example, if household income was positively correlated with demographic characteristic X (for example age), and X in turn is negatively correlated with gasoline expenditures this could explain the negative bivariate relationship between income and the gas budget. To this purpose, figures 8 to 12 show relationships for different subgroups.

We can generally observe the negative relationship between income and gas budget in all considered population strata. 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 more rural areas, i.e. in smaller cities or outside of metropolitan statistical areas (MSAs), tend to have a higher gas budget, as do families with more household members.

2.4 Regression

Focusing on the intensive margin of gasoline usage, we then estimate regression coefficients on the set of observations with strictly positive gasoline expenditures. We consider a number of specifications where income is proxied by income rank, annual income, or total quarterly expenditures, respectively, and where gasoline consumption is measured as the absolute expenditure budget, or its log transform. A number of household regressors and time fixed effects are also included.

Table 3 collects these OLS regression results. The controls not shown are a quarterly time dummy, 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.

For all measures of income and gas budget the relationship after controlling for the demographic characteristics is negative. The first three columns are the results from OLS regressions pooling all observations, and clustering standard errors at the household level. The effect on the gas budget is expressed in absolute terms: For example, as listed in column 1, moving up one decile in the income distribution decreases the gas budget by 0.25 percentage points. Of course the OLS estimates ignore any unobserved household heterogeneity. Although the panel dimension is short (any household remains in the dataset for at most 4 consecutive quarters), columns (4) and (5) display results from a fixed-effects regression. Since the demographic controls barely vary over the course of one year, only the time dummies are included in these specifications.

3 Model

3.1 Households

There is a continuum of mass 1 of households which are indexed with i. Households purchase gasoline for two purposes: First, they buy a fixed amount $\bar{E}$ every period in order to commute to work. Second, they buy a quantity E_{it}^H which enters the utility function when used together the durable good D_{it}. This variable "service flows from durables" S_{it} is 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 $\bar{E} + E_{it}^H$. 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.

Households differ in their labor efficiency: One hour of labor supplied by household i yields e_i units of labor services to a firm. Households' labor supply in hours is fixed at 1, so effective labor services provided by the household are e_i. There is a common labor market for all households in which one unit of labor service is paid a wage w_t. Households have to pay a proportional labor tax τ^L, a sales tax of τ^C on their consumption of the output good and a receive profits π_t from firms. The durable good D_{it} depreciates at rate d. The households' budget constraint is hence

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

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

In every period household i chooses $C_{it}, E_{it}^H, D_{i,t+1}, L_{it}^H$ to maximize lifetime utility given by

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

Table 3: Regression results

	Dependent variable: Gas expenditure share				
	OLS			Fixed effects	
	(1)	(2)	(3)	(4)	(5)
Income Rank	-0.0253^{***}			$-.0110^{***}$	
	(0.0003)			(0.0009)	
log(Income)		-0.0045^{***}			
		(0.0001)			
log(Expenditures)			-0.0244^{***}		-0.0249^{***}
			(0.0001)		(0.00025)
Constant	0.0681^{***}	0.1011^{***}	0.2638^{***}	0.0500^{***}	0.2605^{***}
	(0.0019)	(0.0020)	(0.0018)	(0.0341)	(0.0316)
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	—	—
R^2	0.1722	0.1640	0.2308	—	—

Controls not shown included in OLS regressions: Quarterly time effects, region, metropolitan area, gender, race, age, education, marital status and family size. Quarterly time effects are retained for the fixed effects regression.

subject to the sequence of budget constraints (3). 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 first-order necessary conditions (FOC) for $E_{it}^H, D_{i,t+1}, L_{it}^H$, repectively, combine with the FOC for consumption to yield

$$\frac{1-\varphi}{\varphi}(1-\gamma)\left[D_t^\gamma\left(E_t^H\right)^{1-\gamma}\right]^{\frac{\zeta-1}{\zeta}}\left(E_t^H\right)^{-1}C_{it}^{\frac{1}{\zeta}}=\frac{q_t}{1+\tau^C} \tag{4}$$

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

$$\tag{6}$$

where $[\ldots]_t \equiv \left[\varphi C_{it}^{\frac{\zeta-1}{\zeta}}+(1-\varphi)\left(D_t^\gamma\left(E_t^H\right)^{1-\gamma}\right)^{\frac{\zeta-1}{\zeta}}\right]$ is a shorthand, and E_t 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

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

such that the optimality conditions are

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

$$\kappa L_t^\alpha\left(E_t^F\right)^{\kappa-1}=q_t \tag{8}$$

3.3 Market clearing and equilibrium

There are only 2 markets in the economy: The market for labor and the market for final goods. The market for oil does not clear inside the model, instead there is an infinite supply of oil at the exogenous price q_t (an alternative way to think about this is that households and firms can "mine" oil at a constant cost of q_t in their backyards).

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

$$Y_t-E_t^F-\int_i q_tE_{it}^H di-q_t\bar{E}=\int_i\left(1+\tau^C\right)C_{it}+\left[D_{i,t+1}-(1-d)D_{it}\right]+\tau_i^Lw_te_idi \tag{9}$$

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

$$L_t=\int_i e_{it}di. \tag{10}$$

(Per Walras' Law, the two market clearing conditions are equivalent if households' budget constraints hold.)

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 (7)-(8), one of the market clearing conditions (9) or (10), and the stochastic equation for the oil price (2).

3.4 Solution, calibration and results

Solution and calibration

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. Solution of the model is straightforward since expectations about future prices are easy to compute. To find household decision rules we use a second-order approximation around the non-stochastic steady state.

Households' idiosyncratic labor productivity e_{it}, which determines their income levels, is given by an AR-1 process in logs. Its mean is normalized to 1, and we choose a persistence parameter ρ_e following XXX. To pick the variance of the shock process we observe that in the CEX households in the top half of the income distribution have mean quarterly expenditures around twice as high as the expenditures of households in the bottom half. The standard deviation of innovations to labor productivity is chosen to replicate this fact in the model. With $\alpha = 2/3$ the elasticity of production with respect to labor is set to a standard value, as is households' discount the future at a quarterly rate with $\beta = 0.99$.

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 $\varepsilon_t^q = 0.086$.

This leaves the four parameters $\tilde{q}, \bar{E}, \zeta, \gamma$ 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 1. We follow Kehrig and Ziebarth (2009) who find that consumption and oil usage enter as complements in household preferences to set the elasticity of substitution $\zeta = 0.73$.

The parameter γ determines the relative contribution of the durables stock and gasoline purchases in the 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 use this as our measure of durable expenditures. Compared to average quarterly spending on gasoline of \$533, these average vehicle expenditures of \$781 are around 50% larger. Correspondingly we calibrate the parameter $\gamma = 0.769$ in the model so that steady-state ratio of consumer expenditures on durables dD and spending on gasoline, $E^{HH} + \bar{E}$ are in the same proportion as in the data.

XXX Finally we calibrate the two remaining parameters $\varphi = 0.877$ and $\bar{E} = 0.0174$ to match the mean gas budget (as share of total expenditures) for the top and bottom half of the income distribution in the CEX, in which the observed expenditure budgets are 4.85% and

Table 4: Calibration

Parameter	Value	Description	Source
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	quarterly depreciation of cars XXX
γ	0.769	Preference weight on D_i	Spending on durables vs gas
φ	0.877	Preference weight on C_i	Gas expenditure heterogeneity
$\bar{E}$	0.0174	Minimum gas purchase	Gas expenditure heterogeneity
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	0.95	Persistence Hh labor efficiency	XXX
sd $\varepsilon_{e_{it}}$	0.062	Heterogeneity in labor efficiency	Mean consumption expenditures
τ^L	0.2	Labor income tax	Tax receipts / labor income
τ^C	0.08	Consumption tax	Average US sales tax

6.21%, respectively.

Results

Figure 1 displays the impulse response functions in percentage deviations from steady state gas price shock of 33%. Given the gas price's long-run average of $2.40 per gallon this would correspond to an increase to $3.18.

To evaluate welfare losses from oil price shocks we run counterfactuals in which instead of the one-time oil price shock we institute a permanent change in tax rates such that households are indifferent between the two. In other words, we find the tax rate such that households' expected utility without gas price shock equals expected utility upon impact of the gas price shock under the original tax rate. Table 5 contains the for equivalent variations of labor and consumption taxes. Because of the fixed amount of gas consumption $\bar{E}$, poor households are affected disproportionately, resulting in a higher relative welfare loss compared to rich households.

This effect is the stronger the larger the size of the change in the oil price. Table 6 contains the tax equivalents for rich and poor households for oil price shocks of 25%, 50%, and 75%, respectively.

Finally, the lower households' income, the large the share of their budget they have to spend on fixed gas consumption. Table 6 displays the welfare effects for different percentiles of the income distribution in addition to the 25th and 75th used in the other tables.

Robustness (XXX Tax rates?)

We also examine how different persistence in household incomes affects results. Intuitively, the longer a poor household expects remains at a low income level, the larger the loss in expected utility from an adverse oil price shock. We examine the magnitude of this effect by considering different values for the parameter ρ_e governing the persistence of the income process.

Figure 1: IRF to gas price shock

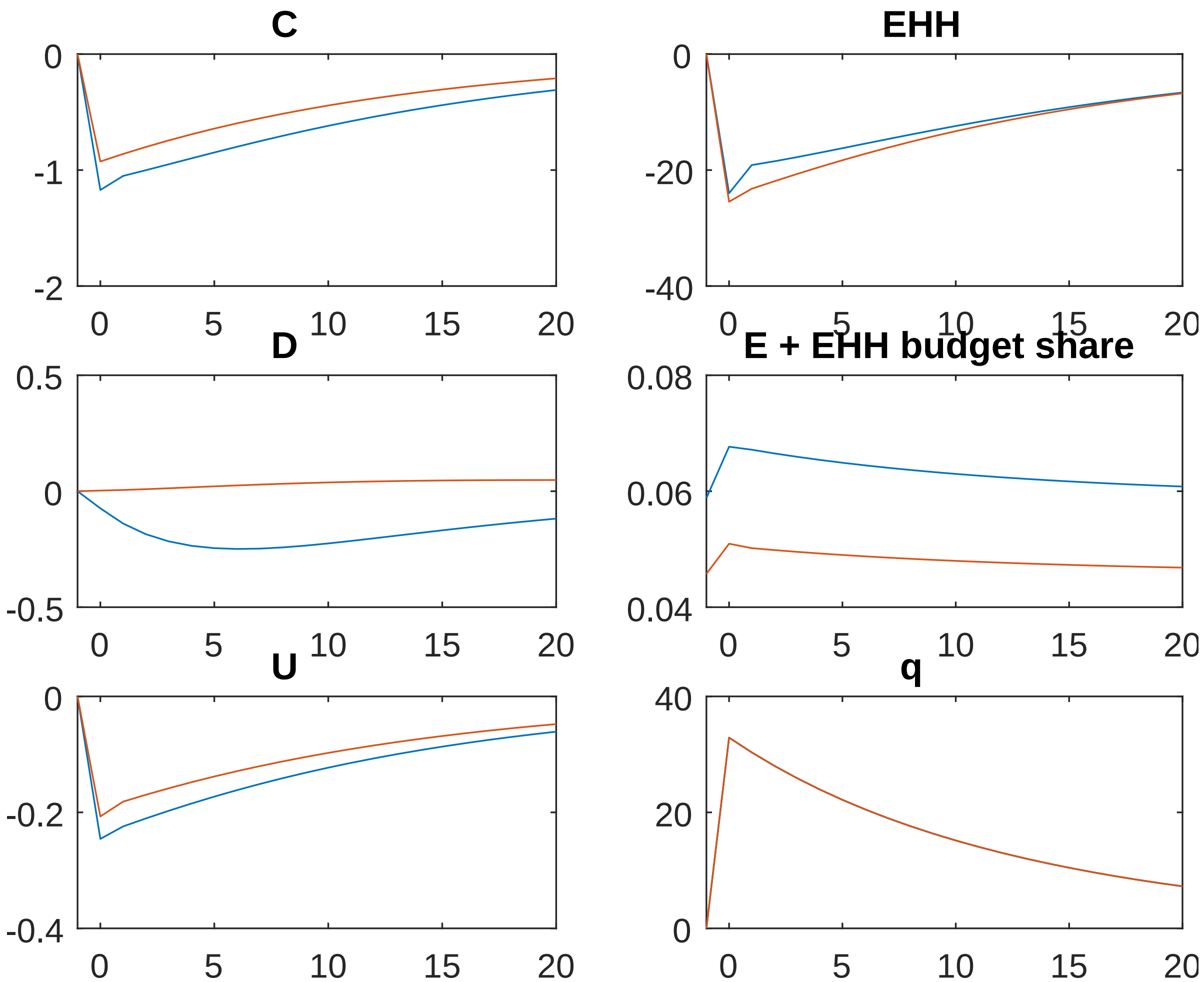

Notes: One-time shock to ε_t^q in period 0. Blue lines: Household at the 25th percentile of the income distribution. Orange lines: Household at the 75th percentile.

Table 5: Tax rates equivalent with one-time oil price shock

Household type	Equivalent tax increase		Consumption equivalents
	Labor tax	Cons. tax	
Medium sized oil-price shock (33%)			
Poor	0.2033%	0.2033%	0.9975%
Rich	0.2028%	0.08269%	0.9978%
Large oil-price shock (50%)			
Poor	0.2047%	0.2047%	0.9965%
Rich	0.204%	0.08384%	0.9969%
Small oil-price shock (25%)			
Poor	0.2026%	0.2026%	0.9981%
Rich	0.2022%	0.08208%	0.9983%

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

Size of oil price shock	Percentile of income distribution				
	10	25	50	75	90
Medium (33%)	0.2036%	0.2033%	0.203%	0.2028%	0.2026%
Large (50%)	0.205%	0.2047%	0.2043%	0.204%	0.2037%
Small (25%)	0.2027%	0.2026%	0.2024%	0.2022%	0.202%

Table 7: Equivalent tax rates by persistence of the income process

	Value of persistence parameter ρ_e			
Income Percentile	0.9	0.95	0.975	0.99
10	0.2034%	0.2036%	0.2038%	0.2041%
25	0.2032%	0.2033%	0.2034%	0.2035%
50	0.203%	0.203%	0.203%	0.203%
75	0.2029%	0.2028%	0.2027%	0.2026%
90	0.2027%	0.2026%	0.2025%	0.2023%

Figure 2: IRF to gas price shock

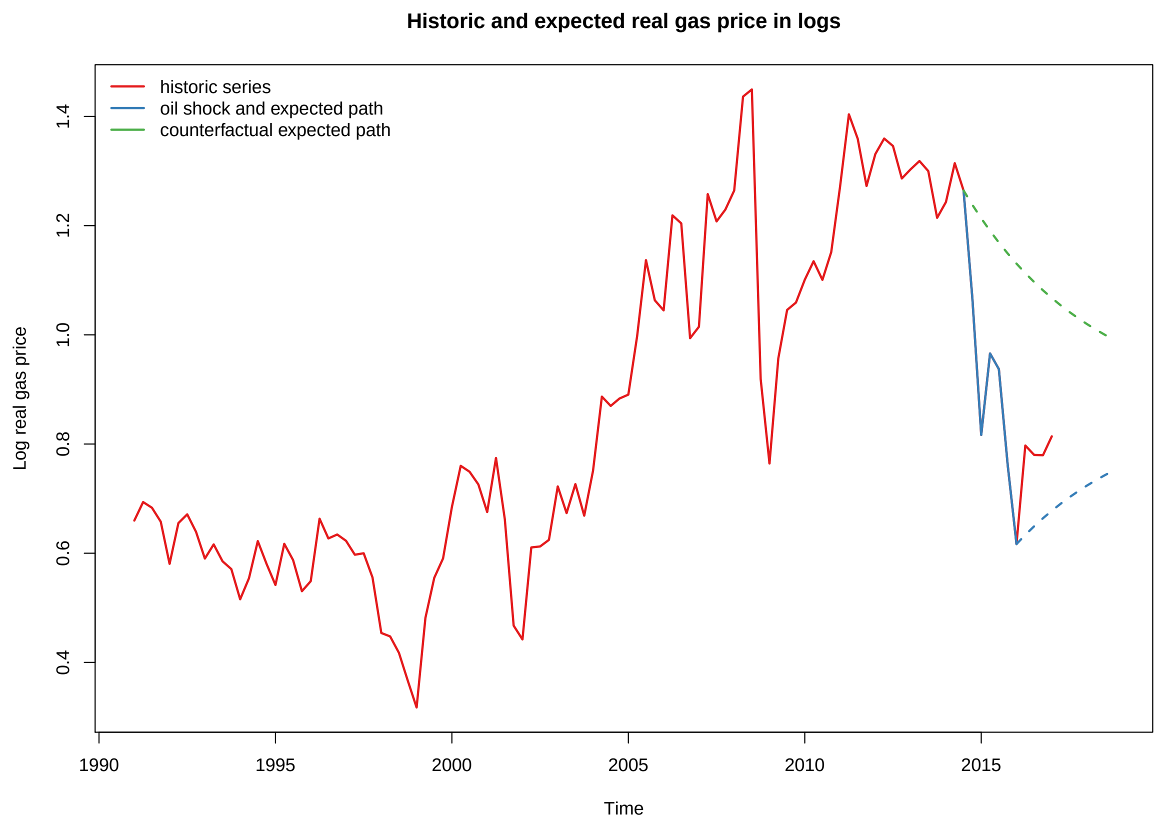

Notes – Red line: Historic gas prices. Blue line: Series of beneficial gas price shocks in 2015 (solid) and expected gas price trajectory (dotted). Green line: Expected gas prices absent the beneficial gas price shocks.

Historic counterfactual

From fall 2014 (2014Q3) to winter 2015/16 (2016Q1), the gasoline price in the US dropped significantly over the course of five quarters – in our quarterly data from \$3.54 down 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 2 displays the actual path of gas prices over this period, together 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 "burn-in" period. Then we compare 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

Table 8: Labor tax rate equivalents for 2015 drop in gas prices

Percentile of income distribution				
10	25	50	75	90
0.1928%	0.1934%	0.194%	0.1945%	0.1949%

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 finding the labor tax rate that would make households indifferent between paying this alternative tax rate not expecting any gas price shocks (i.e. households expect the gas prices to follow the green dashed path in figure 2) on the one hand and the utility they receive under the sequence of gas price shocks while paying the old tax rate of 20%.

We find that for the median household the drop in gas prices is equivalent to a cut in the labor tax to 0.194 . As expected, larger benefits accrue to those households in the lower half of the distribution. Table 8 shows that the hypothetical tax cuts offsetting the drop in gas prices are 0.1934 and 0.1945 at the 25th and 75th percentile, respectively. These welfare differences widen to equivalent tax cuts of 0.1928 and 0.1949 at the 10th and 90th percentile.

4 Summary and future work

The model of household gasoline consumption outlined above 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 with 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 rich and poor households, respectively.

References

Barsky, R. B. and L. Kilian (2004). Oil and the Macroeconomy Since the 1970s. *Journal of Economic Perspectives 18*(4), 115–134.

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

Blanchard, O. J. and M. Riggi (2013). 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.

Brons, M., P. Nijkamp, E. Pels, and P. Rietveld (2008). A meta-analysis of the price elasticity of gasoline demand. A SUR approach. *Energy Economics 30*(5), 2105–2122.

Edelstein, P. and L. Kilian (2009). How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics 56*(6), 766–779.

Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics 113*(2), 363–398.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *Energy Journal 29*(1), 113–134.

Kehrig, M. and N. Ziebarth (2009). Why Do Oil Price Shocks Matter? Transmission Mechanisms on the Supply and the Demand Side.

Kilian, L. and R. Vigfusson (2014). The Role of Oil Price Shocks in Causing U. S. Recessions.

Schmalensee, R. and T. M. Stoker (1999). Household Gasoline Demand in the United States. *Econometrica 67*(3), 645–662.

A Figures

Figure 3: Distribution of gasoline consumption and expenditures

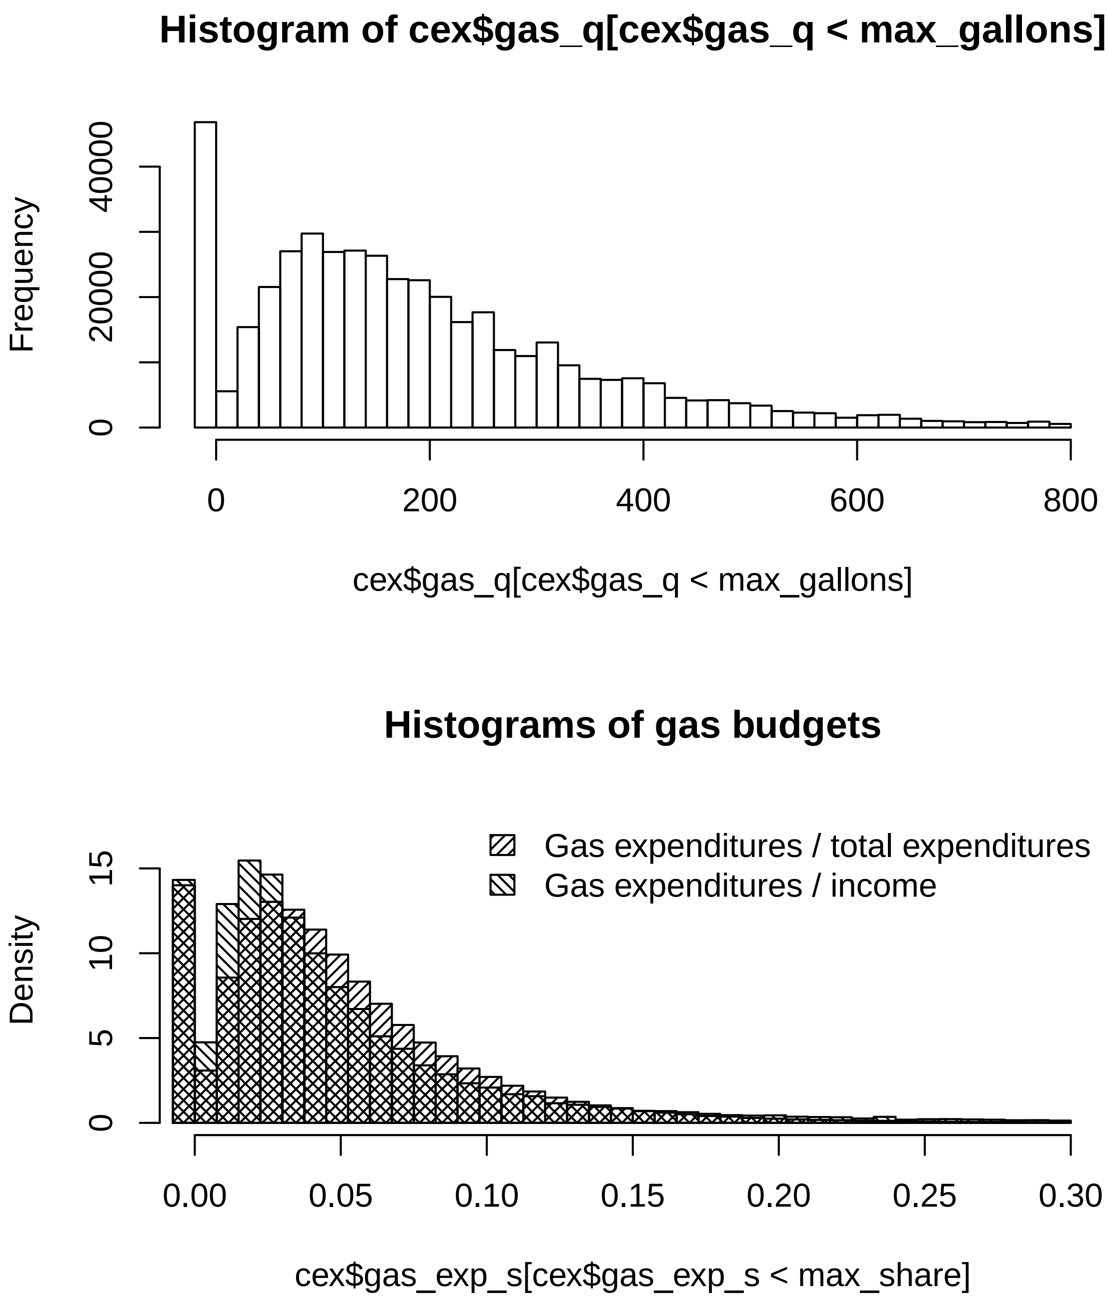

Density estimates of household gasoline consumption and gasoline budget shares pooled across households and quarters. Top panel: Approximate absolute gasoline consumption derived as quarterly gasoline expenditures divided by gasoline price. Bottom panel: Gasoline budgets derived as quarterly gasoline expenditures divided by quarterly total expenditures and divided by a quarter of annual income, respectively.

Figure 4: Distribution conditional on positive usage

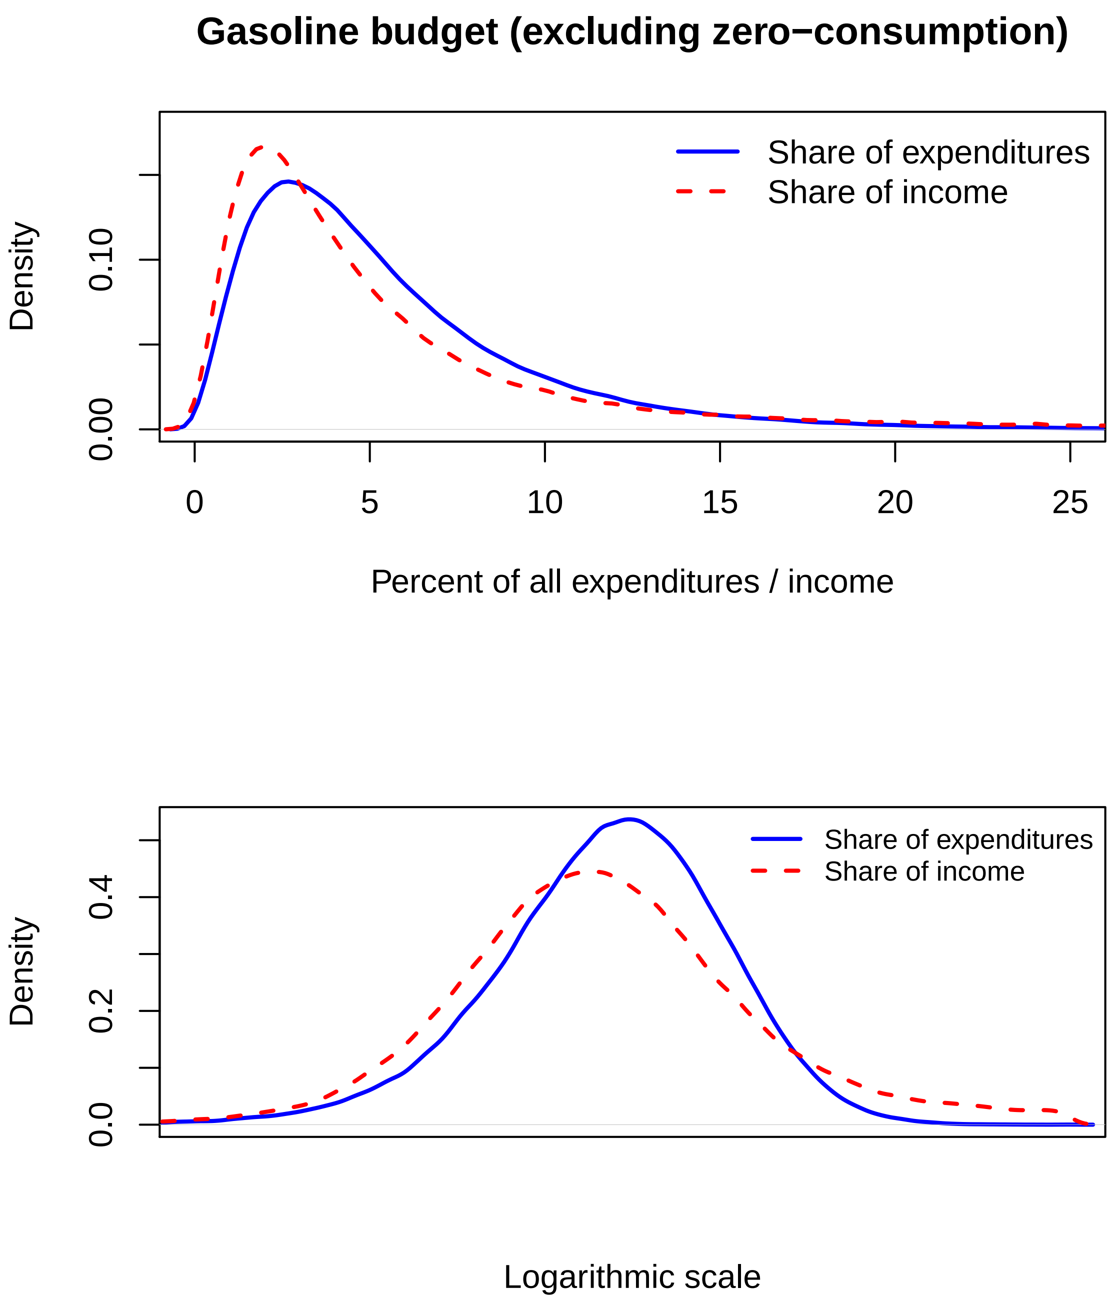

Gasoline expenditures as share of total expenditures and income and their log transforms, respectively.

Figure 5: Gas budget as function of income / total expenditures

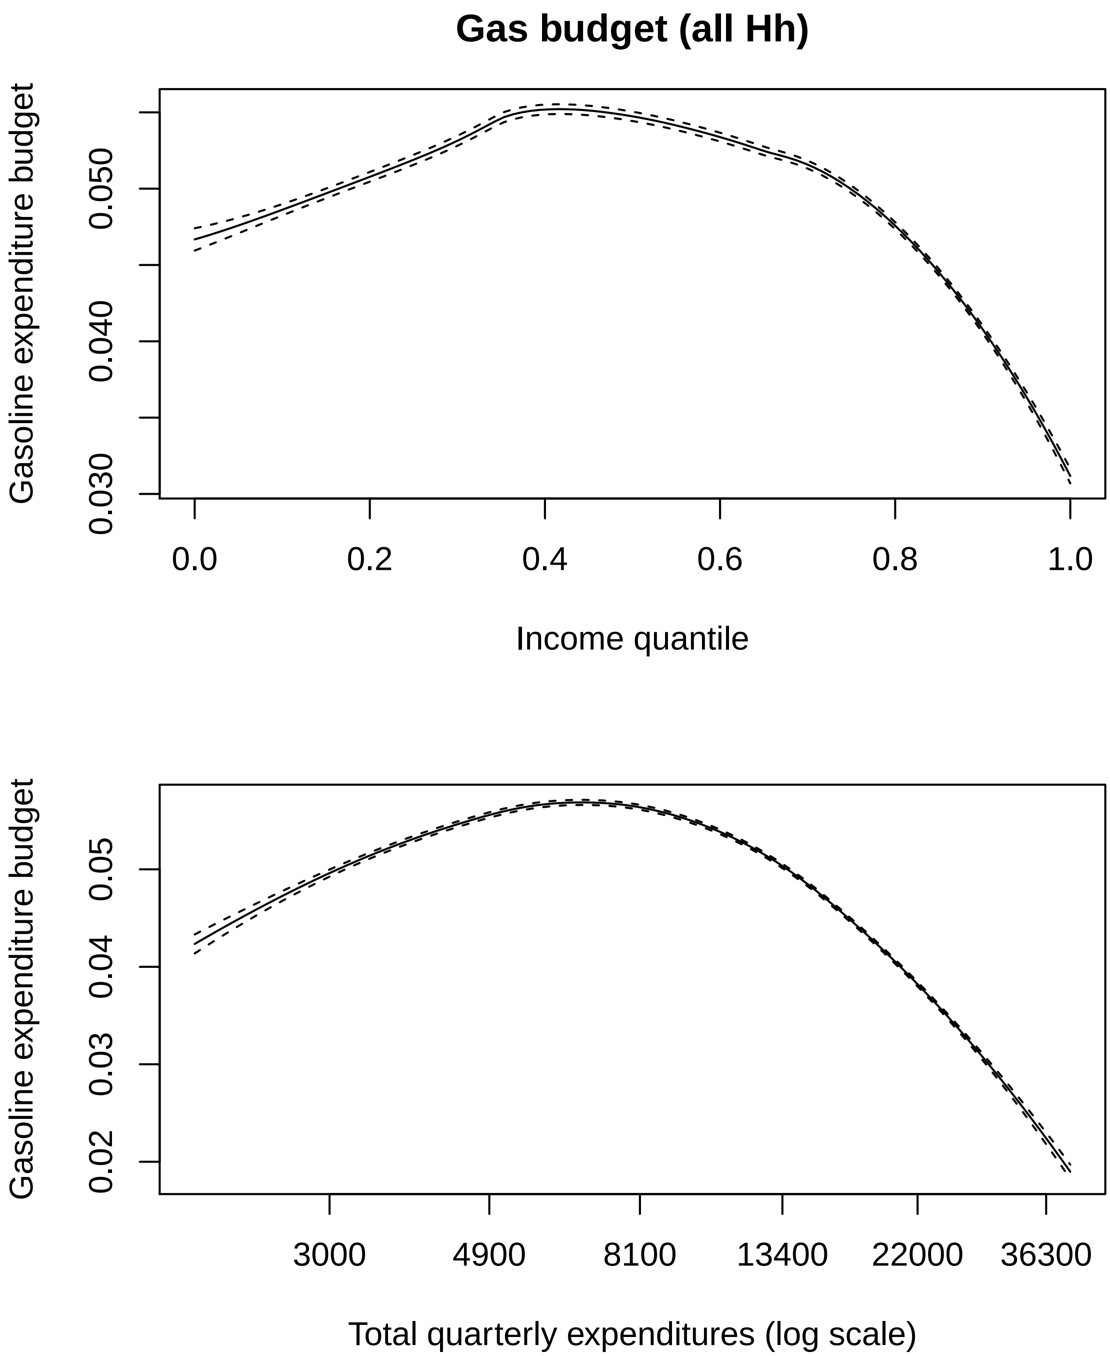

Nonparametric estimation of expected gas budget conditional on income quantile (top panel) and total expenditures (bottome panel) using local regression and including 95% confidence bands.

Figure 6: Gas budget conditional on positive usage

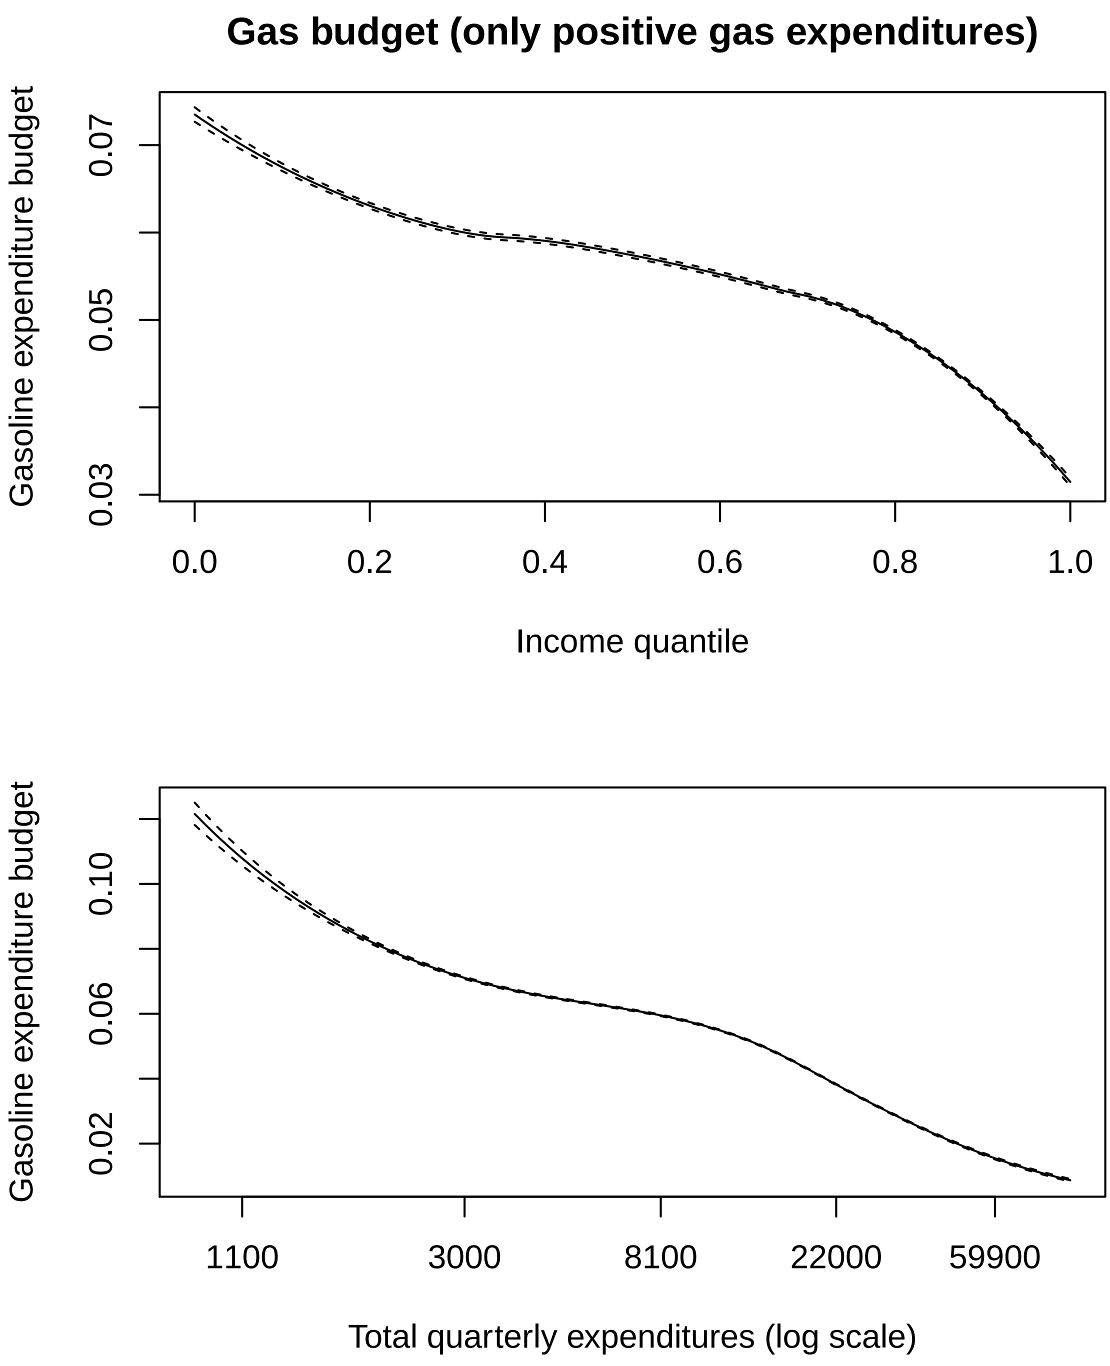

Nonparametric estimation of expected gas budget for households with positive usage conditional on income quantile (top panel) and total expenditures (bottome panel) using local regression and including 95% confidence bands.

Figure 7: Joint frequency of gas expenditure budget and income rank

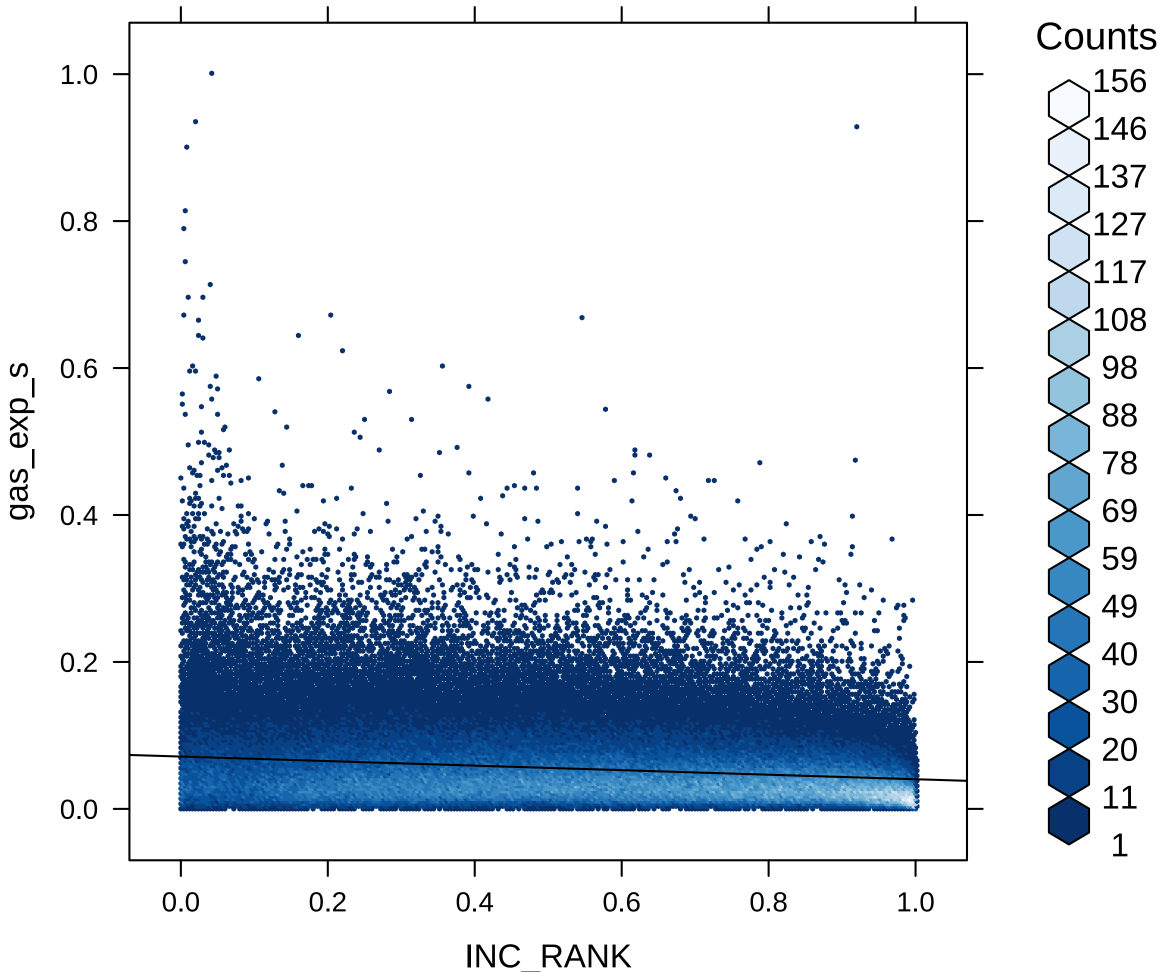

Frequency of observations for given income rank (x-axis) and gas budget (y-axis). Lighter colors correspond to more observations. Black line is linear relationship from regression of gas budget on income rank and a constant.

Figure 8: Nonparametric estimate: by age group

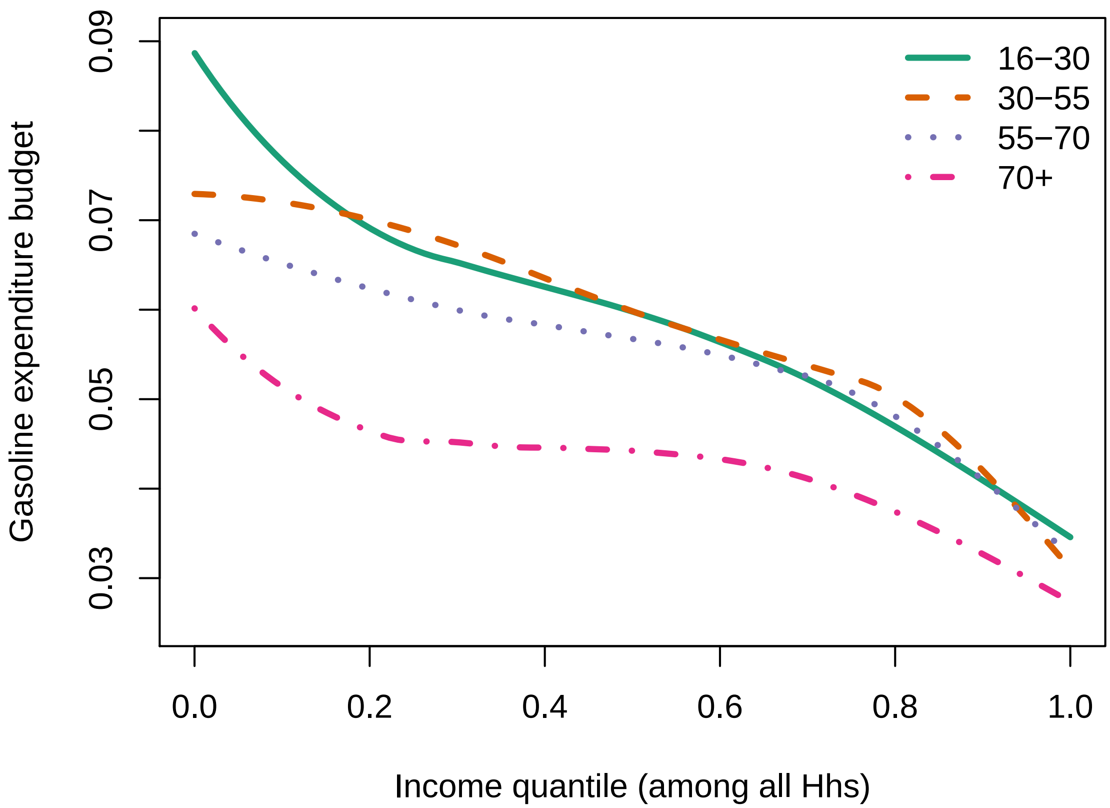

Local estimate of mean gas budget conditional on income rank for different age groups

Figure 9: Nonparametric estimate: by urban status

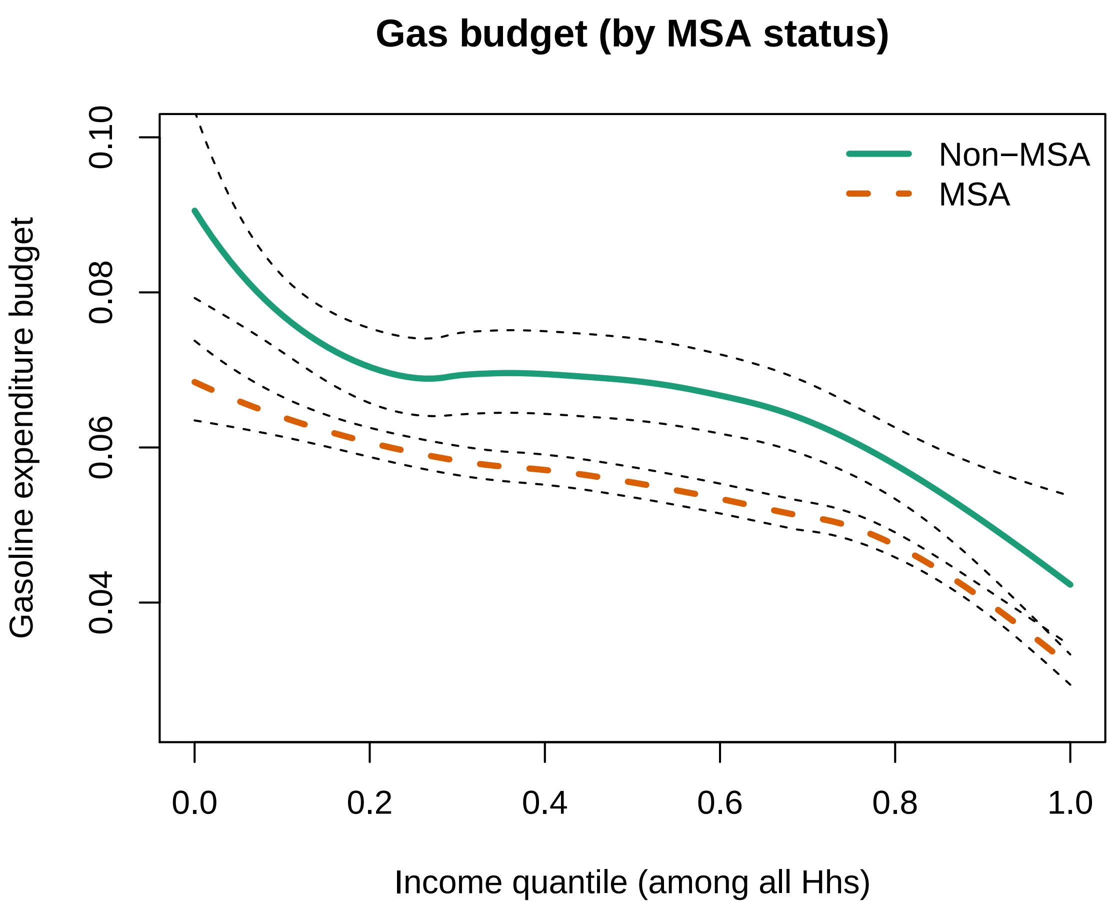

Local estimate of mean gas budget conditional on income rank by MSA status

Figure 10: Nonparametric estimate: by city size

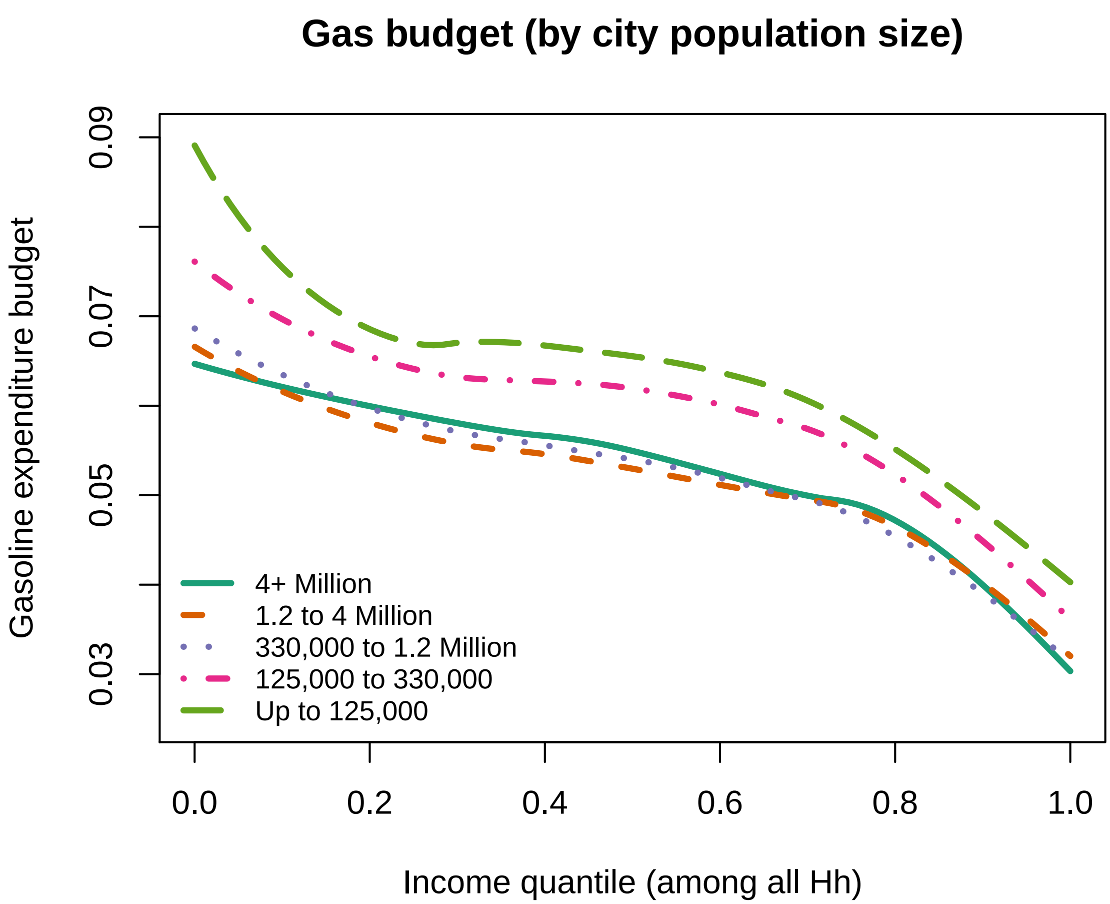

Local estimate of mean gas budget conditional on income rank by population size of city

Figure 11: Nonparametric estimate: by marital status

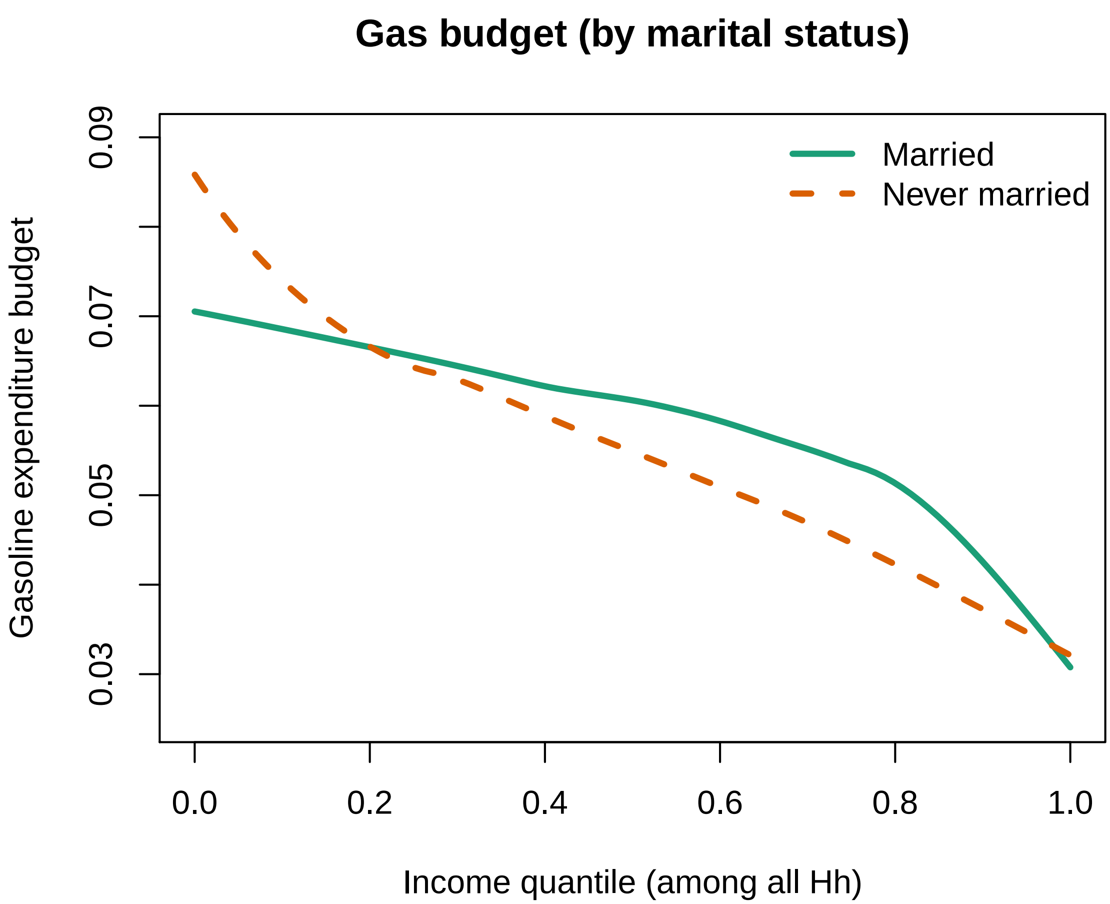

Local estimate of mean gas budget conditional on income rank by marital status

Figure 12: Nonparametric estimate: by family size

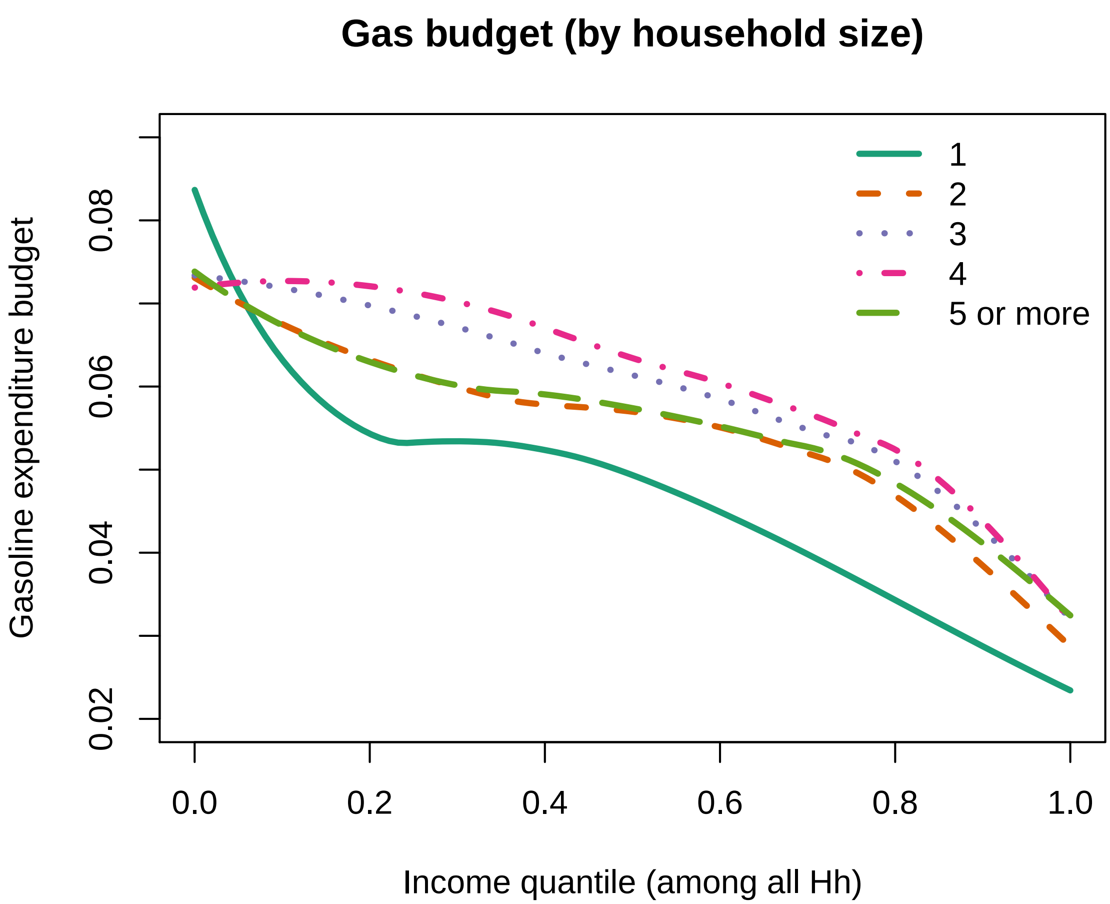

Local estimate of mean gas budget conditional on income rank by family size

Figure 13: Nonparametric estimate: working-age white males

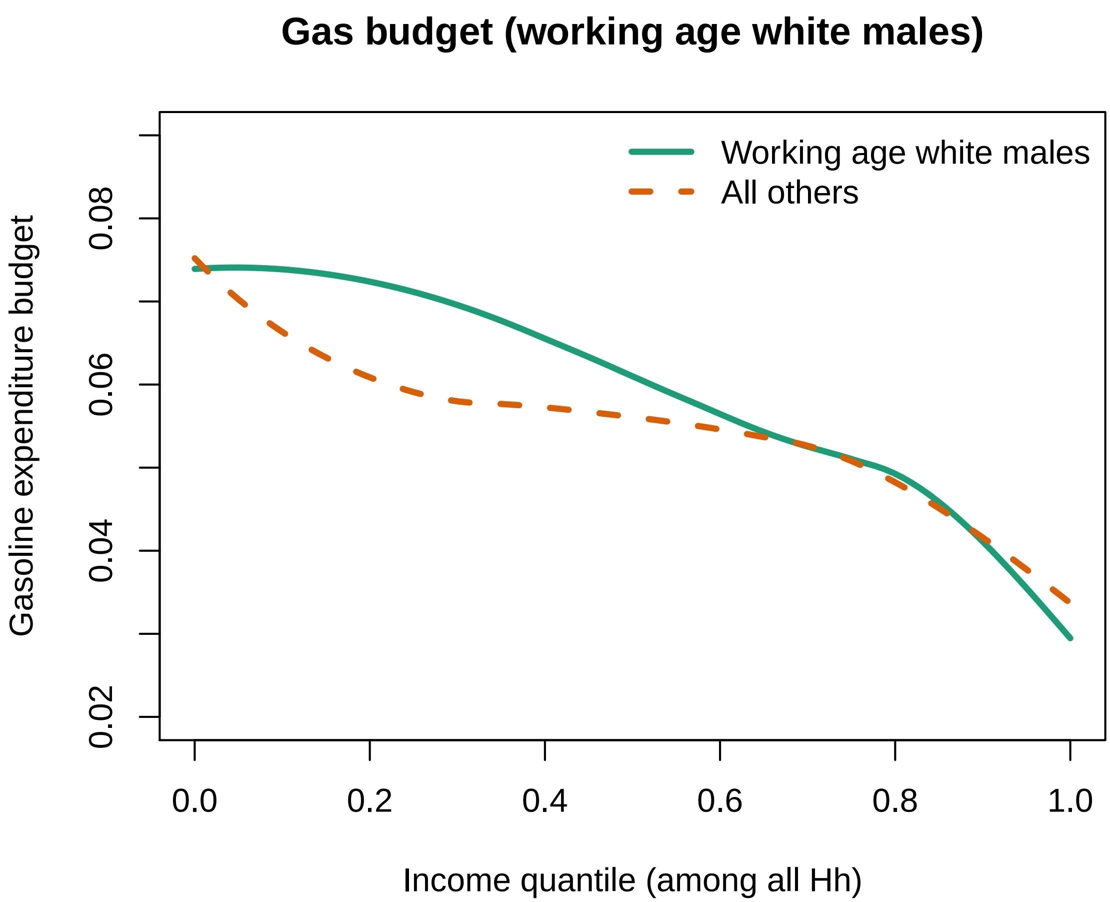

Local estimate of mean gas budget conditional on income rank for white males between 25 and 60 years old