

FUNGIBILITY AND CONSUMER CHOICE: EVIDENCE FROM COMMODITY PRICE SHOCKS*

JUSTINE S. HASTINGS AND JESSE M. SHAPIRO

First Version: March 2011

This Version: April 2013

We formulate a test of the fungibility of money based on parallel shifts in the prices of different quality grades of a commodity. We embed the test in a discrete-choice model of product quality choice and estimate the model using panel microdata on gasoline purchases. We find that when gasoline prices rise consumers substitute to lower octane gasoline, to an extent that cannot be explained by income effects. Across a wide range of specifications, we consistently reject the null hypothesis that households treat “gas money” as fungible with other income. We compare the empirical fit of three psychological models of decision-making. A simple model of category budgeting fits the data well, with models of loss aversion and salience both capturing important features of the time series. Keywords: fungibility, income effects, consumer psychology, gasoline. *JEL* Codes: D12, L15, Q41, D03

*We are grateful for comments from Nick Barberis, Matt Lewis, Erich Muehlegger, Justin Sydnor, and seminar audiences at the NBER, Yale University, the University of Chicago, Northwestern University, Cornell University, UC Berkeley, and Columbia University. This paper benefited greatly from conversations with Keith Chen during the project’s early stages. This work was supported by the Centel Foundation/Robert P. Reuss Faculty Research Fund at the University of Chicago Booth School of Business, the Yale University Institution for Social and Policy Studies, and the Brown University Population Studies Center. We thank Eric Chyn, Sarah Johnston, Phillip Ross, and many others for outstanding research assistance. Atif Mian and Amir Sufi generously provided cleaned zipcode-level income data originally obtained from the IRS. Wharton Research Data Services (WRDS) was used in preparing this study. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers. E-mail: justine_hastings@brown.edu, jesse.shapiro@chicagobooth.edu.

I. INTRODUCTION

Neoclassical households treat money as fungible: a dollar is a dollar no matter where it comes from. But many households keep track of separate budgets for items like food, gas, and entertainment (Rainwater, Coleman, and Handel 1959; Zelizer 1993; Antonides, de Groot, and van Raaij 2011). In hypothetical choices, participants routinely report different marginal propensities to consume out of the same financial gain or loss depending on its source (Heath and Soll 1996). Violations of fungibility matter for the evaluation of public policies such as income tax withholding (Feldman 2010), tax-deferred retirement accounts (Thaler 1990), and fiscal stimulus (Sahm, Shapiro, and Slemrod 2010). Yet there is little empirical evidence measuring the importance of violations of fungibility outside the laboratory, and even less evidence on how best to incorporate violations of fungibility into models of consumer choice.

This paper has two goals. The first is to formulate and execute a test of the fungibility of money using real-world data on consumer choice. Having found evidence of such a violation, the second goal is to compare the performance of a set of psychologically rich models of consumer decision-making in explaining the data.

Our empirical test for the fungibility of money is based on the following thought experiment (Fogel, Lovallo, and Caringal 2004). Consider a household with income M . The household must purchase one indivisible unit of a good that comes in two varieties: a low-quality variety with price P_L and a high-quality variety with price P_H , where $P_H > P_L$. Now consider two scenarios. In the first scenario, the prices of the two varieties each increase by Δ dollars to $P_L + \Delta$ and $P_H + \Delta$ while household income remains constant at M . In the second scenario, the household's income declines by Δ dollars to $M - \Delta$ while prices remain constant (at P_L, P_H). Both scenarios lead to the same budget constraint and hence to the same utility-maximizing behavior. However, both laboratory evidence (Fogel, Lovallo, and Caringal 2004; Saini, Rao, and Monga 2010) and an existing body of theory (discussed in more detail below) predict that households will exhibit more substitution towards the low-quality variety in the price-increase scenario than in the income-decrease scenario.

We develop an estimable discrete-choice model that captures the logic of the thought experiment. In the model, households trade off the added utility of more expensive varieties of a good against the marginal utility of other consumption goods. As the household gets poorer, either through a loss of income or an increase in the level of prices, the marginal utility of other consumption goods rises relative to the marginal utility of higher-quality varieties, leading to substitution towards lower quality levels. Under standard utility-maximization, the model implies fungibility: a parallel shift in the prices of all varieties is behaviorally

equivalent to an appropriately scaled change in income. We translate this implication into a formal statistical test of the null hypothesis that households treat money from different sources as fungible.

We apply our empirical framework to the choice of gasoline grade. We use aggregate data from the Energy Information Administration for 1990-2009, and panel microdata from a retailer covering over 10.5 million transactions from 61,494 households. Gasoline comes in three octane levels—regular, midgrade, and premium—which differ in price and perceived quality. When global supply and demand conditions cause an increase in the price of oil, households (in both aggregate and microdata) substitute substantially towards lower-octane grades of gasoline. Substitution to lower grades occurs despite the fact that, if anything, the price gaps among different gasoline grades tend to fall when the price of gasoline is high, a finding that further corroborates a shift in demand towards lower octane levels.

Two facts suggest that the relationship between gasoline prices and octane choice cannot be explained by income effects. First, in the second half of 2008 gasoline prices fell due to the deepening of the financial crisis and associated recession. During this period, although almost all indicators of consumer spending and well-being were plummeting, households substituted to higher-octane gasoline. Second, the cross-sectional relationship between income and octane choice implies that a loss of \$1000 in household income increases the propensity to buy regular gasoline by less than one tenth of one percentage point. Yet we find that a \$1 increase in the price of gasoline—equivalent to a loss of income of about \$1200 for a typical household—increases the propensity to buy regular gasoline by 1.4 percentage points.

Consistent with these descriptive facts, our econometric model consistently rejects the statistical null of fungibility across a range of specifications. Baseline estimates imply that the marginal utility of money increases 15 times more in response to an increase in the gasoline price than in response to an equivalent decrease in income from other sources. We reject fungibility in models allowing for household preference heterogeneity and in models exploiting cross-sectional, time-series, or aggregate variation in income. We consider a number of alternative explanations for the observed pattern, including changes over time in the composition of households buying gasoline, changes in vehicle use, correlation between gasoline prices and other prices, and measurement error in income. None of these alternatives can account for the large deviations from fungibility that we observe.

To further check our identification strategy, we conduct a placebo exercise in which we test whether gasoline money and other money are treated as fungible when households make a quality choice in a non-gasoline domain. We show that poorer households buy less expensive brands of orange juice and milk, but that gasoline prices exert a weak (and statistically insignificant) positive effect on the quality of brands chosen in these categories. We cannot reject the null hypothesis that consumers treat gasoline money and

other money as fungible when choosing among milk or orange juice brands. The failure to reject the null is not driven by a lack of power: we can statistically rule out violations of fungibility as large as those that we observe for gasoline. We do not, however, observe sufficient variation in orange juice or milk prices to use these categories for an independent test of fungibility.

Having found evidence that households do not treat money as fungible, we turn next to an evaluation of three alternative models of decision-making that might plausibly account for a violation of fungibility in our setting. First, we formulate an ad-hoc model of category budgeting motivated by psychological evidence (Heath and Soll 1996). Second we estimate loss-aversion model based on Köszegi and Rabin (2006), but with a non-stochastic referent. Finally we estimate a salience model based on Bordalo, Gennaioli, and Shleifer (2012).

All three models can generate the finding that higher gasoline prices lead households to substitute to regular gasoline. In the category budgeting model, substitution arises as households try to avoid deviating from their typical level of gasoline expenditure. In the loss aversion model, higher gasoline prices lead to lower-than-expected non-gasoline consumption, making households more price-sensitive. In the salience model, an increase in prices causes households to focus on price rather than octane when evaluating gasoline grades.

We estimate each model and investigate its empirical fit. The category budgeting model fits the data better than the loss aversion and salience models because rational-expectations assumptions in the latter two models predict rapid adaptation to higher prices, and because both models effectively allow for only two levels of the marginal utility of money. We show that both the loss aversion and salience models can replicate the empirical patterns almost perfectly when these two restrictions are relaxed.

The first contribution of this paper is to provide evidence of violations of fungibility “in the wild.” Most extant evidence comes from hypothetical choices or incentivized laboratory behaviors (Thaler 1999; Read, Loewenstein, and Rabin 1999). Important exceptions include Kooreman’s (2000) study of child care benefits in the Netherlands, Milkman and Beshears’ (2009) study of the marginal propensity to consume out of a coupon in an online grocery retail setting, and related work by Abeler and Marklein (2008) and Feldman (2010). To our knowledge, ours is the first paper to test for violations of fungibility in the response to price changes.

The second contribution of this paper is to compare alternative psychological explanations for a violation of fungibility. To our knowledge, ours is the first paper to estimate Köszegi and Rabin’s (2006) and Bordalo, Gennaioli, and Shleifer’s (2012) models using data on retail purchases. In that sense, our paper contributes to a growing literature that structurally estimates the parameters of psychological mod-

els of decision-making using field data (Laibson, Repetto, and Tobacman 2007; Conlin, O’Donoghue, and Vogelsang 2007; Barseghyan et al. forthcoming; Crawford and Meng 2011; Grubb and Osborne 2012; DellaVigna, List, and Malmendier 2012). Ours is among the first of these papers to compare the predictions of more than one psychological model.

Methodologically, we follow Allenby and Rossi (1991), Petrin (2002) and Dubé (2004) in enriching the role of income effects in discrete-choice models of household purchase decisions. We show that allowing for violations of fungibility significantly improves model fit. In that sense, we also contribute to a literature in marketing that incorporates psychological realism into choice models with heterogeneity (Chang, Siddarth, and Weinberg 1999).

Two existing literatures predict the opposite of what we find. First, a literature following Barzel (1976) exploits tax changes to test the Alchian-Allen conjecture that higher category prices result in substitution to *higher* quality varieties (Sobel and Garrett 1997). In the context of gasoline, Nesbit (2007) and Coats, Pecquet, and Taylor (2005) find support for the Alchian-Allen conjecture; Lawson and Raymer (2006) do not. Second, a literature in psychology and economics examines “relative thinking” in which consumers focus on ratios when normative decision theory implies that they should focus on differences (Azar 2007 and 2011).

The remainder of the paper is organized as follows. Section II provides background information on grades of gasoline. Section III describes our data. Section IV presents our model of consumer choice and discusses our empirical strategy for testing fungibility. Section V presents a descriptive analysis of gasoline grade choice. Section VI presents estimates of our model. Section VII presents evidence on the empirical fit of alternative psychological models of decision-making. Section VIII concludes.

II. BACKGROUND ON GASOLINE GRADE CHOICE

Gasoline typically comes in three grades, with each grade defined by a range of acceptable octane levels: regular (85-88), midgrade (88-90), and premium (90+) (EIA 2010). A higher octane level increases gasoline’s combustion temperature so that it can be used in high-compression engines (which yield higher horsepower for a given engine weight) without prematurely igniting (also known as “knocking”).

Typically, a gasoline retailer maintains a stock of regular and premium gasoline on site, and midgrade is produced by mixing regular and premium at the pump. Regular and premium gasoline are, in turn, produced at refineries by blending intermediate product streams with different chemical properties to achieve the desired output. There are generally multiple ways to arrive at an acceptable final product, and refineries

use programming models to decide on the profit-maximizing mix given spot prices for various input, intermediate, and output streams. Changing the output of the refinery to include, say, more premium and less regular gasoline would involve changing the mix of intermediate streams used in gasoline production (Gary and Handwerk 2001), which can be achieved seamlessly for small changes in the product mix.

A large proportion of high-octane gasoline sales go to cars that do not require it, with most consumers justifying their purchase of premium gasoline on “vague premises” (Setiawan and Sperling 1993). Most modern cars have sensors that prevent knocking at any octane level. Perhaps because auto makers often recommend premium gasoline for sports cars, the most frequently stated reason for using high-octane gasoline is a performance gain, for example in the time to accelerate from 0 to 60 miles per hour (Reed 2007). *Consumer Reports* (2010) and other consumer advocates have questioned whether such performance gains are real. Buyers of high-octane gasoline may also believe that using above-regular grades helps promote long-term engine cleanliness and health, but because detergents are required for all grades of gasoline, using above-regular grades does not in fact help an engine “stay clean” (Reed 2007). In addition, any supposed gains in fuel economy from using high-octane grades are “difficult to detect in normal driving conditions” (API 2010; see also Click and Clack 2010). Thus, according to Jake Fisher at *Consumer Reports*, “There are two kinds of people using premium gas: Those who have a car that requires it, and the other kind is a person who likes to waste money” (Carty 2008).

It is well known that higher octane gasolines tend to lose market share when the price of gasoline goes up (Lidderdale 2007), a phenomenon that gasoline retailers call “buying down” (Douglass 2005). Due to their association with good performance, high-octane varieties are perceived as a luxury good that the consumer can do without. Industry analyst Jessica Caldwell notes that buying down is surprising in light of the small stakes involved: “It really doesn’t add up to very much... It’s more of a psychological thing. You’re at the pump, and it seems like every time you hit a certain threshold, you cringe” (quoted in Lush 2008). Caldwell’s account seems to capture a psychological intuition related to the models we consider in section VII below.

III. DATA

III.A. Panel Microdata

Our main data source is a transaction-level file from a large U.S. grocery retailer with gasoline stations on site. The data include all gasoline and grocery purchases made from January 2006 through March 2009 at 69 retail locations, located in 17 metropolitan areas in 3 different states.

For each gasoline transaction, the data include the date, a store identifier, the number of gallons pumped, the grade of gasoline (regular, midgrade, or premium), and the amount paid. We use these data to construct a price series by store, grade, and date equal to the modal price across all transactions, where transaction prices are calculated as the ratio of amount spent to number of gallons, rounded to the nearest tenth of a cent. The majority of transactions are within one tenth of one cent of the daily mode, and 88 percent of transactions are within one cent of the daily mode.

The data allow us to match transactions over time for a given household using a household identifier linked to a retailer loyalty card. Approximately 87 percent of gasoline purchases at the retailer can be linked to a household identifier through the use of a loyalty card.

Our main measure of household income is supplied by the retailer, and is based on information given by the household to the retailer when applying for the loyalty card, supplemented with data, purchased by the retailer from a market research firm, on household behaviors (e.g., magazine subscriptions) that are correlated with income.

For comparison and sensitivity analysis we also make use of two geography-based measures of income. For the large majority of households in our sample, the retailer data include the census block group of residence. We use this to obtain 2000 U.S. Census income data at the block group level. We further match block groups to zipcodes using 2000 Census geography files provided by the Missouri Census Data Center (2011). For each zipcode, we obtain annual measures for 2006, 2007, and 2008 of the mean adjusted gross income reported to the IRS (Mian and Sufi 2009).

For estimation we use a subsample comprised of purchases by households that make at least 24 gasoline purchases in each year of 2006, 2007, and 2008, and for whom we have a valid household income measure. We exclude some outlier cases from the estimation sample.¹ The final sample we use in estimation includes 10,548,175 transactions by 61,494 households. In the appendix we show that our results are not sensitive to excluding the top and bottom 10% of households according to gasoline purchase frequency.

Our analysis exploits the fact that our data allow us to match gasoline transactions to grocery transactions by the same household. We construct a time-varying proxy for total household consumption by computing for each household the total spending on non-gasoline (grocery) items in the four weeks prior to each gasoline purchase.

In a placebo exercise, we examine the effect of gasoline prices on purchases of refrigerated orange juice

¹These are: households that purchase more than 665 times over the length of the sample, households that ever purchase more than 210 times in a given year, households that ever purchase more than 10 times in a given week, and a small number of transactions that involve multiple gasoline purchases. We also exclude from the sample a small number of store-days in which reported prices are too large by an order of magnitude, and a small number of store-days in which stockouts or reporting errors mean that only one grade of gasoline is purchased. Together, these exclusions represent 4.78 percent of transactions.

and milk. We focus on these two categories as they are perishable, relatively high in volume, and involve clear quality and price delineations (for example, between conventional and organic varieties). We aggregate individual UPCs in these categories into products grouped by size and brand and construct a weekly price series for each store and product. The online appendix contains additional details on how we group UPCs into products and how we construct the price series. For estimation, we use data on households that purchase at least once in the category in each sample year. We exclude households that purchase 200 or more times in a given category in any sample year. In the online appendix, we present estimates of our key results using even tighter restrictions on frequency of purchase and show that our substantive conclusions are unchanged.

III.B. Aggregate Data

To confirm that the key patterns in the retailer panel are representative, we use monthly data from 1990-2009 on retail prices and sales volume by grade of gasoline for the 50 states (and the US total) obtained from the Energy Information Administration (EIA) at eia.doe.gov in June 2010. Portions of our analysis also make use of national and regional weekly price series obtained from the EIA in April 2012. The EIA collects price and volume data from a sample survey of retailers and a census of prime suppliers, essentially large firms that deliver a significant volume of petroleum products to “local distributors, local retailers, or end users” (EIA 2009).

We supplement the EIA data with data from the Consumer Expenditure Survey (CEX) Interview Files, 2006-2009. We use the Consumer Expenditure Survey data to evaluate the representativeness of grocery expenditures in our sample and to project the total annual expenditures of sample households.

To analyze the effect of nationwide economic shocks we use monthly US-level data on personal income per capita, personal consumption expenditures per capita, total US population, and the price level from the National Income and Product Accounts (NIPA), obtained from bea.gov in July 2010. We convert per capita amounts to per household amounts using annual data on the number of US households from the Census Historical Time Series Tables, obtained from census.gov in March 2011. We impute the number of people per household as the ratio of a five-year moving average of population and a five-year moving average of the number of households. We use a moving average to avoid jumps in the series at annual updates. We supplement these with monthly data on the S&P 500 index obtained from the Center for Research in Security Prices (CRSP) via the Wharton Research Data Service (wrds.wharton.upenn.edu) in December 2012. We use the S&P index as a proxy for aggregate trends in household wealth. We convert household income, wealth, and expenditures to 2008 US dollars using the price level for personal consumption expenditures.

III.C. Sample Representativeness

Table I evaluates the representativeness of our sample on key dimensions of interest. The first column presents statistics for all households in the retailer database who ever buy gasoline. The second column presents statistics for households in our estimation sample. The third column presents representative state-level statistics for the three states our retail sites are located in. Thus comparing columns (1) and (2) reveals differences between all households purchasing gasoline and those purchasing gasoline at least 24 times per year during our 3-year period, and comparing columns (1) and (3) reveals differences between the retailer's customers and state populations.

Given our requirement that households in the estimation sample purchase gasoline at the retailer at least 24 times per year for a little over 3 consecutive years, the majority of gasoline-buying households are excluded from our estimation sample. During our sample period, households could move, stop in to one of our retail stores even if they live in other areas, discard their loyalty cards, or purchase their gasoline primarily at other gasoline retailers. However, while the households in our estimation sample are a minority of the gasoline-buying households in the full retailer database, on most dimensions the two samples look similar. The main points of distinction between estimation sample households and those in the full sample result directly from our selection rule. Relative to households in the full sample, estimation sample households make more gasoline trips per purchase month, buy more groceries at the retailer, and live closer to the retailer's store.

The third column of the table shows means for all households in the three states from which we draw our retailer data, with each state weighted according to its number of households in the full retailer database. Relative to the average household, households from the retailer data live in higher-income block groups. Households in the retailer sample buy slightly less regular gasoline than reported in the EIA data for the same states, and also pay about 4-5 cents less per gallon of gasoline than the state average as reported by the EIA.² Sample households spend less on groceries at the retailer than the average household in the state spends on groceries overall, presumably reflecting the fact that sample households buy some groceries at other retailers.

III.D. Validity of Income Measures

The geographic variation in our main household income measure corresponds well with data from other sources. The median of our household income measure at the Census block group level has a correlation of

²The lower average price per gallon at retailer sites presumably arises because the retailer does not sell a major brand of gasoline, whereas the EIA average price series is based on data that include (higher) major-brand prices.

0.82 with median household income from the 2000 Census. The mean of our household income measure at the zipcode level has a correlation of 0.77 with mean adjusted gross income in the zipcode, as reported to the IRS in 2008.

A drawback of our main household income measure is that it is only available at a single point in time. To address this limitation, we use our measure of household grocery expenditures to proxy for time-varying shocks to household income. Existing literature shows that food expenditure responds to variation in income in the cross-section and over time, predicting about 40 percent of the cross-sectional variation in total expenditure (Skinner 1987) and responding significantly to shocks to current and future household income (Stephens 2001, 2004; Japelli and Pistafferi 2010).

Table II shows that, in our data, food expenditures are related to income variation in the cross-section and over time. Across households, we estimate an income elasticity of grocery expenditure of 0.17, which closely matches the analogous estimate of 0.17 from the Consumer Expenditure Survey. Across zipcodes, we estimate an elasticity of 0.14. Importantly, the zipcode-level relationship remains similar in magnitude (at 0.09) and marginally statistically significant in a model with zipcode fixed effects, indicating that changes in income at the zipcode level are correlated with changes in food expenditure at our retailer. These findings lend credibility to food expenditures as a proxy for shocks to income over time, especially in light of the large existing literature establishing the responsiveness of food expenditures to shocks.

IV. AN ECONOMETRIC TEST OF FUNGIBILITY

In this section we lay out a discrete-choice model of consumer behavior with standard, neoclassical foundations. We use this model to construct an econometric test of fungibility that we can take to the data.

IV.A. *Model*

A household i must buy $q_{it} > 0$ gallons of gasoline in period t , and may choose among grades of gasoline indexed by $j \in \{0, \dots, J\}$.³ Here $j = 0$ denotes regular gasoline. Grade j has a price p_{jt} per gallon at time t .

The household has total per-period expenditures m_{it} . Money not spent on gasoline is spent on other goods, delivering indirect utility $\Lambda_i(m_{it} - q_{it}p_{jt})$, which we normalize so that $\Lambda_i(m_{it} - q_{it}p_{0t}) \equiv 0$.

³Implicitly we assume a unitary household. Because violations of fungibility can arise from strategic behavior within the household (Lundberg and Pollak 1993), in the appendix we show that our estimates are similar when we restrict the sample to households with only one adult member.

A household purchasing grade j at time t obtains total utility U_{ijt} where

$$U_{ijt} = v_{ijt}q_{it} + \Lambda_i(m_{it} - q_{it}p_{jt}). \quad (1)$$

Here, v_{ijt} is a taste parameter. Observe that the model exhibits fungibility: an increase in gasoline prices of \$1 is equivalent to a decrease in non-gasoline expenditures of q_{it} dollars.

Taking a first-order approximation to $\Lambda_i(m_{it} - q_{it}p_{jt})$ around $(m_{it} - q_{it}p_{0t})$, we can write per-gallon utility $u_{ijt} = U_{ijt}/q_{it}$ as

$$u_{ijt} = v_{ijt} - \lambda_{it}(p_{jt} - p_{0t}) \quad (2)$$

where λ_{it} is household i 's marginal utility of non-gasoline expenditures at time t .

In general the marginal utility λ_{it} is a function of $(m_{it} - q_{it}p_{0t})$. We assume that λ_{it} takes a linear form:

$$\lambda_{it} = \mu_i - \eta(m_{it} - q_{it}p_{0t}). \quad (3)$$

Here, μ_i determines the level of a household's marginal utility of money, and η determines the extent to which the marginal utility of money diminishes with greater non-gasoline consumption.

Our specification of λ_{it} corresponds to quasi-linearity in money when $\eta = 0$, and to a quadratic indirect utility $\Lambda_i(\cdot)$ up to a second-order term in the price gaps between grades.⁴ In the online appendix, we show that our conclusions are robust to approximating $\Lambda_i(\cdot)$ with a flexible polynomial.

We assume that tastes are given by

$$v_{ijt} = \alpha_{ij} + \varepsilon_{ijt}, \quad (4)$$

where α_{ij} is a household-specific, time-invariant taste intercept and ε_{ijt} is an unobservable, i.i.d. taste shock distributed type I extreme value independently of the other terms. In the appendix, we present estimates from models in which v_{ijt} includes an aggregate preference shock.

We estimate the model via maximum likelihood under alternative assumptions about α_{ij} and μ_i . To test the hypothesis that households treat money as fungible, we estimate an unrestricted model:

$$\lambda_{it} = \mu_i - \eta^M m_{it} + \eta^G q_{it} p_{0t}. \quad (5)$$

We then test the restriction that $\eta^M = \eta^G = \eta$.

⁴Formally, suppose that $\Lambda_i(m_{it} - q_{it}p_{0t}) = K_i - \frac{\eta}{2} \left(\frac{\mu_i}{\eta} - (m_{it} - q_{it}p_{0t}) \right)^2$ for some constant K_i . Then from equation (1) it follows that $U_{ijt}/q_{it} = v_{ijt} - \lambda_{it}(p_{jt} - p_{0t}) - \frac{\eta}{2} q_{it} (p_{jt} - p_{0t})^2$ where λ_{it} is as defined in equation (3).

It is worth noting that although we follow Houde (2012) in taking gasoline quantities q_{it} as exogenous, our specification is formally consistent with a model in which gasoline quantities are endogenous but appropriately separable from the choice of octane.⁵ In our descriptive analysis, we discuss and rule out several reasons for non-separability between gasoline quantity and octane, such as changes in composition of households or cars on the road. In the appendix, we show that our results survive controlling directly for gallons purchased.

IV.B. Identification and Implementation

Following the assumptions above we can write per-gallon utility as:

$$u_{ijt} = \alpha_{ij} - (\mu_i - \eta^M m_{it} + \eta^G q_{it} p_{0t}) (p_{jt} - p_{0t}) + \varepsilon_{ijt} \quad (6)$$

To build intuition for identification, consider a special case with no heterogeneity in tastes or in gasoline consumption, and with only two grades whose price differs by a constant that we normalize to unity. Then our model can be represented as a binary logit with utility

$$u_{it} = (\alpha - \mu) + \eta^M m_{it} - \eta^G q p_{0t} + \varepsilon_{it}. \quad (7)$$

We will identify η^M from variation in total expenditures m_{it} and η^G from variation in the price p_{0t} of regular gasoline.⁶ We will reject the null hypothesis of fungibility when a parallel increase of \$1 in the price of all grades increases the propensity to purchase regular gasoline by more than a decrease of \$ q in total expenditure m_{it} .

To estimate the model we construct two measures of per-period expenditures m_{it} . Our main measure, m_i , does not vary over time. To construct it, we estimate a regression of total annual expenditure on total annual family income using the 2006-2009 Consumer Expenditure Survey interview files. We use the parameters from this regression to predict each household's total annual expenditure from the household income measure supplied by the retailer. We adjust all standard errors for the use of this two-step procedure following Murphy and Topel (1987).

By estimating a first-stage model that predicts expenditures from reported income, we limit attenuation

⁵In particular, when $\eta = 0$ our model corresponds to the “cross-product repackaging” model of Willig (1978) and Hanemann (1984).

⁶As equation (6) shows, in practice η^M and η^G are also identified by the relationship between income and the sensitivity of purchase probabilities to variation in $p_{jt} - p_{0t}$. During our sample period the retailer engaged in significant experimentation with grade price gaps, providing a credible source of identification of the effect of the price gaps $p_{jt} - p_{0t}$ on purchase behavior. We show in the appendix that our results survive on a subsample in which the price gaps do not vary at all.

bias in η^M due to measurement error or transitory shocks to income. In the online appendix, we lay out this argument formally, and show that our results are stronger in a specification in which we allow explicitly for measurement error and transitory shocks to income. In the appendix, we show that our results are robust to using Census block group income to predict total expenditures, and to allowing the parametrization of marginal utility to differ across households of different income levels (Petrin 2002).

Our second measure of per-period expenditures, m_{it} , is time-varying. To construct it, we estimate a regression of total annual expenditure on total monthly expenditure on food at home using the 2006-2009 Consumer Expenditure Survey interview files. We use the parameters from this regression to predict each household's total annual expenditure from the household's total grocery spending in the four weeks prior to a given gasoline transaction.

We measure total annual gasoline expenditures (at the regular-grade price) $q_{it}p_{0t}$ as follows. We measure q_{it} as average annual US gasoline consumption during our sample period (from the EIA), divided by the number of US households in 2006. We use an aggregate measure of q_{it} because measuring purchases at a single retailer would tend to understate a household's total gasoline consumption, leading us to overstate η^G .⁷ In the appendix, we confirm that our results strengthen when we measure q_{it} using data on purchases at the retailer, and we show that our results are robust to allowing that the prices of energy goods other than gasoline are correlated with the price of gasoline.

We measure p_{0t} as the weekly average national retail price of gasoline (from the EIA). Variation in p_{0t} is a credible source of identification of parameter η^G because national gas prices are affected by global supply and demand shocks that are plausibly unrelated to tastes for octane levels. To the extent that gasoline prices are driven by income shocks, we will tend to understate η^G and hence our test for fungibility will be conservative. In the appendix, we show results from a specification in which we isolate the variation in p_{0t} that is attributable to fluctuations in the spot price of crude oil.

V. DESCRIPTIVE EVIDENCE ON FUNGIBILITY

V.A. Gasoline Grade Choice

Figure I plots, separately by decade, the regular-grade share of total US gasoline sales as well as the (real) US average price for regular unleaded gasoline, from the EIA data. Figure II plots the regular-grade share and average price by week for transactions in our retailer panel. Both figures show a clear pattern: the share

⁷The number of gallons of gasoline per household that we estimate (1183) is greater than average annual purchases in our panel for all but 4.7 percent of households.

of regular gasoline tends to increase (at the expense of premium and midgrade) when the price of gasoline rises, and to fall when the price of gasoline falls. Estimates presented in the online appendix show that the effect of a one-time increase in the price of gasoline persists for several months, with no evidence of decay.

Figure III uses a fixed-effects regression to decompose the series in figure II into between- and within-household components. The figure shows that the changes over time in the market share of regular gasoline result almost entirely from households switching between grades of gasoline, rather than from changes over time in the types of households buying gasoline. In the online appendix, we show that, among households who switch gasoline grades over time, a majority exhibit a tendency to buy regular gasoline more often when the price of gasoline is high.

Income effects do not provide a good explanation of the time-series relationship between octane choice and the price of gasoline. The market share of regular gasoline fell by almost 4 percentage points in the second half of 2008, a time when the worsening of the financial crisis led households to tighten their belts, with plummeting automobile and retail sales (Linebaugh and Dolan 2008; Zimmerman, Saranow, and Bustillo 2008) and slowing growth in luxury items such as organic foods (NielsenWire 2009). The switch to high-octane gasoline during this period is a puzzle in light of falling incomes, but is consistent with the decline in the price of gasoline, itself brought on by the impact of the global economic slump on world oil demand (Taylor 2009).

Income effects are also too small to account for much of a relationship between the price of gasoline and the choice of octane level. During the price spike from January to June of 2008, gasoline prices increased from \$2.98 to \$4.10 per gallon, generating an annual income loss of \$1313 for a typical household with annual gasoline consumption of 1183 gallons. During that same period, the share of transactions going to regular gasoline increased by 1.4 percentage points, from 80.2 percent to 81.6 percent. Figure IV shows the cross-sectional relationship between household income and the propensity to buy regular gasoline. An OLS regression fit to the plot implies that an income loss of \$1313 would increase the share of regular gasoline by only 0.02 percentage points, and that explaining a change in the regular share of 1.4 percentage points requires an income loss of almost \$100,000.

In the online appendix, we present estimates of a linear probability model in which both gasoline prices and income affect the propensity to buy regular gasoline. Using various measures of income and successive controls for time effects, we find that the effect of gasoline prices is not quantitatively consistent with income effects. Coupled with the descriptive evidence above, these linear probability models serve to illustrate the patterns in the data that identify our formal econometric model.

V.B. Price Variation and the Supply Side

The thought experiment in the introduction assumes that the price gap between high and low quality grades remains constant. In practice the assumption of constant price gaps between regular, midgrade, and premium gasoline is a good approximation but does not hold exactly. Our formal model explicitly allows for variation in price gaps, and in the appendix we show that our findings are unchanged if we estimate on the subsample of transactions in which the price gaps between grades are 10 cents each.

Price gaps provide an additional source of evidence on changes in household demand, because an increase in the demand for regular gasoline should reduce the price gap between premium and regular gasoline, at least until refineries can adjust their output. We show in the online appendix that this prediction is borne out: an increase in the price of regular gasoline induces a small and temporary decline in the price gap between premium and regular gasoline. The fact that the relative quantity and relative price of regular grade move together in response to a change in the price of regular gasoline is not consistent with an explanation for our findings based solely on shocks to the relative supply of different octane levels.

If higher prices induce greater price sensitivity among consumers, this price sensitivity should not only affect choice among octane levels but also choice among retailers. In the online appendix, we present an analysis of the dynamics of retailer markups based on Lewis (2011) that shows that retailer behavior is qualitatively consistent with a positive relationship between the price of gasoline and the price sensitivity of consumers.

V.C. Alternative Explanations

The descriptive evidence above shows that when gasoline prices increase, households increase their demand for regular gasoline in a manner that is difficult to reconcile with income effects and hence with the fungibility of money. Before turning to formal econometric evidence, we pause to consider some alternative explanations for our findings.

Vehicle substitution. When gasoline prices rise, households substitute toward driving more fuel-efficient vehicles. If fuel economy is correlated with octane recommendations, vehicle substitution could explain a portion of the time-series variation in octane choice. Empirically, the correlation between fuel economy and octane recommendations is weak: across vehicles in model years 2003-2008, the correlation between combined miles per gallon and an indicator for recommending or requiring premium gasoline is -0.11 (Environmental Protection Agency 2011). Using benchmark estimates of the effect of gasoline prices on vehicle purchases, scrappage, and driving intensity, in the online appendix we estimate that a \$1 increase in

the gasoline price should increase the share of regular gasoline by only 0.03 percentage points in the short run due to vehicle substitution.

Vehicle maintenance. An owner of a fuel-inefficient vehicle who perceives premium gasoline as an investment in vehicle maintenance may be less willing to make that investment when the price of gasoline is high and hence the expected lifespan of the vehicle is low. In the appendix, we proxy for a vehicle's fuel efficiency with an estimate of its tank size, and show that the effects we estimate are similar between households with more or less fuel efficient vehicles, with a slightly larger effect for households with more fuel efficient vehicles.⁸

Driving behavior. If drivers perceive higher octane levels as complementary to sporty driving, and if sporty driving is more expensive when gasoline prices are high, then households might substitute to regular gasoline when gasoline prices are high. In the online appendix we present evidence from vehicle accident data on the relationship between driving speeds and the price of gasoline. We find no evidence of a relationship between the two.

Learning. Over time consumers may be learning that regular gasoline involves little or no sacrifice in performance for their vehicles. Such learning may well accelerate when the gasoline price is high, say due to greater media attention to gasoline. However, because learning is cumulative, a learning-based explanation for the time series of octane shares cannot explain why the share of regular gasoline falls when gasoline prices fall. In the online appendix we formalize this intuition by estimating a stylized model of learning and showing how the model's predictions diverge from the empirical data series.

VI. FORMAL TESTS OF FUNGIBILITY

VI.A. Main Results

Table III presents our main results. For each specification we present estimates of the effect on marginal utility of a \$1000 decrease in gasoline expenditures or a \$1000 increase in total expenditures (parameters η^G and η^M , respectively). We also present the average marginal effect on regular share of three experiments: increasing the price of regular gasoline by \$1, decreasing gasoline expenditures by \$1000, or increasing total expenditures by \$1000. We present the ratio η^G/η^M as a quantitative summary of the extent to which households violate fungibility, and we present the p-value from a Wald test of the hypothesis that $\eta^M = \eta^G$.

⁸Using data from Reuters (2007) and www.fueleconomy.gov, we estimate that for the top 20 selling vehicle models in January-July 2007, the correlation between tank size (in gallons) and combined fuel efficiency (in miles per gallon) is -0.76 . Calculations kindly performed for us by staff at fueleconomy.gov show that among all vehicles in 2010 the correlation between tank size and combined fuel efficiency is -0.73 .

In column (1), we present our baseline specification. In this model we use our cross-sectional measure of household expenditures m_i and we assume that there is no heterogeneity in taste parameters α_{ij} and μ_i . This model is a conditional logit model (McFadden 1973).

In our baseline specification in column (1) we find that a \$1 increase in the price of regular gasoline increases the regular share by 1.4 percentage points, which, in turn, implies that a \$1000 decrease in household gasoline expenditures decreases the regular share by 1.2 percentage points. By contrast, a \$1000 increase in total household expenditures decreases the regular share by 0.08 percentage points. The Wald test rejects the equality of the effects of gasoline and total expenditures with a high level of confidence.

In column (2), we use our time-varying measure of household expenditure m_{it} . We allow that μ_i is a linear function of m_i to ensure that η^M is identified by variation in m_{it} . If anything, using time variation to identify η^M tends to weaken the estimated income effect, strengthening our rejection of fungibility.

In column (4) we present a specification in which we allow for unobservable variation in α_{ij} . We assume that α_{ij} are normally distributed independently across households and octanes, and independently of m_i . For computational reasons we estimate the model on a subsample consisting of every 10th transaction for each household. In column (3) we re-estimate the model from column (1) on the subsample to illustrate its comparability to the full sample, and in the online appendix we present results from a specification with heterogeneity in α_{ij} estimated on the full sample. Comparing columns (3) and (4), we find that allowing for household-specific unobservable tastes tends, if anything, to strengthen the estimated effect of the gasoline price level on the propensity to buy regular-grade gasoline. We continue to confidently reject the null hypothesis of fungibility.

In the appendix, we show that the estimates in column (1) of table III are robust to identifying the model using variation in world crude oil prices, splitting the sample into high- and low-income households, allowing for a correlation between gasoline prices and other energy prices, using several alternative estimates of household gasoline and total expenditures, and allowing for aggregate preference shocks. In the online appendix, we present estimates from a model in which we allow for unobserved heterogeneity in μ_i and a model in which we allow for heterogeneity in both α_{ij} and μ_i without imposing distributional assumptions. Across these specifications we consistently reject the null hypothesis of fungibility.

The appendix also presents sensitivity analysis of the estimates from time-varying expenditures in column (2) of table III. Our results are robust to improving the measurement of m_{it} by purging seasonal variation in grocery expenditures, by restricting the sample to regular grocery buyers, and by eliminating data from the early part of the sample. Our results also survive further efforts to purge cross-sectional identification of η^M from the model.

VI.B. Interpretation of Magnitudes

The violation of fungibility that we estimate is economically significant. Across the specifications in table III, we find that households respond 15 to 33 times more to a reduction in income due to an increase in gasoline prices than to equivalent variation in income from other sources. The point estimate in column (1) of table III implies that a \$1 increase in the price of gasoline would have to reduce a typical household's non-gasoline expenditures by more than \$17,000 per year to reconcile the observed increase in the propensity to purchase regular gasoline.

Figure V illustrates the violation of fungibility in a different way. The figure shows weekly averages for three series. The first is the observed share of transactions going to regular gasoline. The second is the predicted share of transactions going to regular gasoline from our baseline model. The third series is the predicted share of transactions going to regular gasoline from a model estimated with the constraint that $\eta^G = \eta^M$, equivalent to imposing fungibility. The first two figures track each other closely: our model fits the large swings in the market share of regular gasoline fairly well. But the third figure, which imposes fungibility, fits very poorly, predicting almost no variation over time in the market share of regular gasoline.

We can also evaluate the magnitude of the violation of fungibility by asking how often households would choose differently if they were forced to obey fungibility. To perform this calculation, for each transaction in our dataset we randomly draw utility shocks ε_{ijt} from their assumed distribution. We then compute the utility-maximizing choice of octane level according to both our baseline model and an alternative model in which we impose $\eta^G = \eta^M$ and adjust μ_i so that each household's mean marginal utility of income is the same as in the baseline model. We compute statistics of interest averaged over five such simulations.

We estimate that 60.8 percent of households make at least one octane choice during the sample period that they would have made differently if forced to obey fungibility. Forcing households to treat gas money as fungible with other money would change octane choices in 0.6 percent of transactions overall.

VI.C. Placebo Tests

We interpret our findings as evidence that, when purchasing gasoline, consumers are especially sensitive to their total gasoline expenditures, and therefore treat changes in gasoline prices as equivalent to very large changes in income when deciding which grade of gasoline to purchase. A prediction of this interpretation is that the effect of gas prices on non-gasoline purchases should be commensurate with income effects. That is, we would expect that gasoline and other income would be fungible in decisions about non-gasoline purchases.

Table IV presents an estimate of our model applied to sample households' choice of orange juice and milk rather than gasoline grade. Here consumers choose between brand-size combinations in each category instead of grades of gasoline. We allow the marginal utility of money to vary separately with gasoline prices and income, just as we did in our baseline model estimated on gasoline purchases.

We find that higher incomes result in a shift in demand away from the private label and towards higher-quality brands. We find that higher gasoline prices tend, if anything, to induce shifting towards higher-quality brands, although the effect is not statistically significant. The counterintuitive sign may result from the fact that some gasoline price shocks are themselves due to income variation (such as the recession), which is a source of conservative bias in our main tests.

In contrast to our findings for gasoline grade choice, we cannot reject the equality of gasoline and total expenditure effects in these cases. That is, consistent with Gicheva, Hastings, and Villas-Boas (2007), we find that gasoline and other income are fungible in decisions about grocery purchases. In the online appendix, we show that our findings are similar even when we restrict attention to orange juice or milk purchases that occur on the same day as a gasoline purchase, when the salience of gasoline prices is presumably at its greatest. The online appendix also presents a visual representation of our findings, showing that when gasoline prices rise, consumers act much poorer when buying gasoline but not when buying orange juice or milk.

The lack of evidence of a violation of fungibility in these placebo categories does not result from a lack of power. Table IV presents p-values from a test that the ratio η^G/η^M for brand choice in the placebo category is equal to the analogous ratio for gasoline grade choice (using the baseline parameters for gasoline estimated in table III). For both orange juice and milk we confidently reject the hypothesis that the ratio for the placebo category is equal to that for gasoline. In this sense, we can statistically reject the hypothesis that fungibility is violated as much in placebo categories as in gasoline grade choice.

Note, however, that power would be an issue if we were to attempt to test whether an increase in, say, the price of milk (as opposed to gasoline) causes substitution to lower-quality milk varieties. Milk and orange juice prices do not exhibit the large swings that gasoline prices do, and the prices of different brand-size combinations do not move in close parallel. Milk and orange juice purchases therefore do not afford a good laboratory for testing the effect of own-category price variation on quality substitution, although they do serve as a valid test of the specification of our gasoline models.

VI.D. Evidence from the Aggregate Time Series

The period of our sample coincides with the onset of a major recession and therefore provides significant aggregate variation in income and wealth that we can exploit to validate our inferences on the effect of income shocks on gasoline grade choice.

Table V presents the results of such an exercise. For each of a set of three indicators of aggregate economic well-being, we compute a household- and time-specific expenditure measure m_{it} equal to the product of m_i and the ratio of the indicator's value to its value in January 2008. We then estimate our model, allowing μ_i to depend directly on m_i so that all identification comes from variation in expenditures over time.

For all three aggregate indicators—proxies for household income, wealth, and consumption—we confidently reject the null of fungibility. In each case the estimated income effect is below that in our baseline specification, and in the case of aggregate income it has the wrong sign.

While the assumption that each household's permanent income varies in exact proportion to a single aggregate index is implausible, the fact that we estimate such small income effects means that even the large changes in permanent income brought about by the onset of recession in 2008 did not have much impact on grade choice once gasoline prices are accounted for. This finding seems difficult to reconcile with a model in which the effect of gasoline prices on grade choice are driven by large income effects.

Panel A of figure VI illustrates our findings for the 2008-2009 recession in a different way. The figure plots the observed (retailer) series for regular share alongside total real expenditures per US household. The figure also plots the predictions of a simple linear model in which we assume that each \$1000 increase in real income has the same effect on the probability of buying regular gasoline as the estimated effect of gasoline expenditures in column (1) of table III. This amounts to assuming—counter to our preferred interpretation—that the effect of gasoline prices on grade choice should be interpreted as an income effect.

The figure shows that such an assumption leads to a wildly counterfactual prediction for the evolution of the market share of regular grade. With consumption steadily rising in 2006 and 2007, the share of regular gasoline is predicted to fall substantially, whereas in fact it fluctuates with the price of regular gasoline. In early 2008, when the price of gasoline rises significantly, the observed share of regular gasoline rises. However, because real consumption expenditures did not fall during this period, the model predicts further declines in the regular share. Finally, and perhaps most strikingly, during the onset of recession in autumn 2008, a fall in the gasoline price coincides with a fall in the share of regular. However, falling real consumption leads to a prediction of rising regular share.

Panel B of the figure illustrates the same patterns during the recession of 1990-1991. Here, steadily rising consumption levels outside the heart of the recession lead to a predicted slow-moving decline in the regular share, when in fact regular share fluctuates significantly in line with variation in oil prices brought about by the events surrounding the first Persian Gulf War. When the recession eats into household consumption, the model predicts the regular share to rise slightly but in fact it falls significantly, coinciding with a fall in gasoline prices.

These plots show that in order to reconcile the aggregate time series with the hypothesis that gasoline prices matter solely due to income effects, we must believe that falling gasoline prices caused an increase in economic well-being during the onset of recession in the autumn of 2008 and the summer of 1990, contrary to national accounts. Our preferred explanation is that falling gasoline prices prompted a temporary feeling of greater economic well-being at the gas pump, leading to substitution to higher-octane grades.

In the online appendix, we take our analysis of aggregate data further, using quarterly state-level data to estimate a version of our model using both cross-section and time-series variation in income. Again, we find that we can confidently reject the null hypothesis of fungibility.

VII. PSYCHOLOGICAL UNDERPINNINGS OF A VIOLATION OF FUNGIBILITY

In this section we consider three models that might account for our findings: an ad hoc model of category budgeting, a model of loss aversion based on Köszegi and Rabin (2006), and a model of salience based on Bordalo, Gennaioli, and Shleifer (2012). We estimate each model, show its predictions for octane choice, and discuss how our findings depend on subjective features of the environment such as beliefs or framing that are not pinned down by our data.

VII.A. Model Specification and Estimation

We will assume throughout that the utility function is a special case of:

$$U_{ijt} = q_{it}\alpha_j + \mu(m_{it} - q_{it}p_{jt}) + \Gamma_{ijt} + q_{it}\epsilon_{ijt} \quad (8)$$

where Γ_{ijt} captures psychological factors that shift preferences for different grades of gasoline and where the other notation is inherited from section IV.

Equation (8) imposes that households have homogeneous preferences up to the random utility shock

ε_{ijt} and that utility is quasilinear in non-gasoline expenditure. In the preceding analysis we find that preference heterogeneity and income effects cannot account for the time-series pattern in octane choice that we document. Excluding those complications makes it easier to add psychological realism to the model.

We will not allow Γ_{ijt} to depend on ε_{ijt} ; in this sense we treat the random utility shock ε_{ijt} as “outside” the psychological models that we consider below. Substantively, this means that we will think of ε_{ijt} as representing random influences on choice that are not perceived by the consumer as a dimension of consumption (in the sense of Köszegi and Rabin 2006) or as a product attribute (in the sense of Bordalo, Gennaioli, and Shleifer 2012).

Some of the models we estimate require attaching cardinal units to the utility from octane. Let g_j denote the octane level of grade j , with $g_0 = 87$ (regular), $g_1 = 89$ (midgrade), and $g_2 = 91$ (premium). It will be helpful to write that $\alpha_j = \alpha g_j + \xi_j$ where α is a parameter and ξ_j is a grade-specific utility shock. We will treat ξ_j as a fixed effect in estimation, so that, absent Γ_{ijt} , this representation does not change the model’s empirical content. In parallel with our treatment of the random utility term ε_{ijt} , we will treat ξ_j as outside the psychological models we estimate.

We estimate all models by maximum likelihood under different assumptions about Γ_{ijt} , which we detail below.

Category Budgeting Model. Existing evidence shows that households keep track of category-specific budgets (Heath and Soll 1996; Antonides, de Groot, and van Raaij 2011) and try to maintain category spending at a target level, perhaps in order to economize on optimization or memory costs (Gilboa and Gilboa-Schechtman 2003; Gilboa, Postlewaite, and Schmeidler 2010).

Because our baseline estimates imply that octane switching offsets only 0.2 cents of every \$1 increase in the price of gasoline, we conclude that category budgets, if they exist, must be flexible. We therefore specify a model in which the household experiences disutility from spending an atypical amount on gasoline. Formally, letting r_i denote a household’s sample mean transaction expenditure, we assume that:

$$\Gamma_{ijt} = -\gamma(q_{it}p_{jt} - r_i)^2 \quad (9)$$

for some constant γ . Note that the category budgeting model is equivalent to our econometric model up to a second-order term in the price difference across grades.⁹

⁹More precisely, equations (8) and (9) together imply the following per-gallon utility function:

$$u_{ijt} = \alpha_j - (\mu - 2\gamma r_i + 2\gamma q_{it} p_{0t})(p_{jt} - p_{0t}) - \gamma q_{it} (p_{jt} - p_{0t})^2 + \varepsilon_{ijt}$$

which, up to the second-order term $(p_{jt} - p_{0t})^2$, is a special case of the choice model in section IV with $\alpha_{ij} = \alpha_j \forall i$, $\mu_i =$

Although the idea that households maintain category budgets is grounded in existing evidence, the assumption that “gasoline” represents a mental category is ad hoc. We do not know of a model that delivers clear predictions about which products will be grouped together in mental budgets, or of how mental budgets should be modeled in a discrete-choice setting. Absent deeper foundations, the model in equation (9) should be thought of as a point of reference rather than as a meaningful alternative to standard choice theory.

Loss Aversion Model. We estimate a model of loss aversion based on Köszegi and Rabin (2006). In the model, households experience both direct utility from consumption and “gain-loss” utility from departures from a stochastic reference level of consumption.

We assume that consumption has two dimensions, gasoline and non-gasoline consumption. Total consumption utility is the sum of these two components, as in equation (8). The term Γ_{ijt} represents the gain-loss utility to household i from buying grade j at time t . It follows from Köszegi and Rabin (2006) that

$$\Gamma_{ijt} = \int [\gamma(\alpha g_j q_{ijt} - \alpha r_{it}^g) + \gamma(\mu(m_{it} - q_{it} p_{jt}) - \mu r_{it}^m)] dG_{it}(r_{it}^g, r_{it}^m) \quad (10)$$

where $\gamma(\cdot)$ is a universal gain-loss function exhibiting loss aversion and $G_{it}(r_{it}^g, r_{it}^m)$ is a probability measure that defines the distributions of the reference level of gasoline consumption (r_{it}^g) and non-gasoline consumption (r_{it}^m).

We follow Crawford and Meng (2011) in making two important simplifications to the model in equation (10). First, we focus on the special case in which $\gamma(\cdot)$ is piecewise linear with a kink at 0, thus ruling out diminishing sensitivity. Second, we assume that the reference point is a point—formally, that $G_{it}(r_{it}^g, r_{it}^m)$ is a degenerate distribution with value equal to the expected consumption level in period t . The online appendix shows the fit of a model with diminishing sensitivity and of a model allowing for a stochastic referent.

The final step in operationalizing the model is to determine the reference points. We follow Köszegi and Rabin (2006) in assuming that the reference points are consistent with rational expectations about both prices and consumption decisions. We assume that the information set on which expectations are based includes q_{it} and m_{it} but that households do not know prices p_{jt} in advance. Rather, households forecast future consumption decisions and prices based on current prices.

We follow the literature on estimation of discrete games (e.g., Pakes, Ostrovsky, and Berry 2007) and estimate households’ expectations in a first stage, estimating the remaining parameters in a second stage via

$\mu - 2\gamma r_i \forall i, \eta^M = 0$ and $\eta^G = 2\gamma$.

maximum likelihood.¹⁰ We estimate the expected octane level and transaction price as the predicted values from regressions of realized values on a cubic polynomial in the national regular price as of one week prior to purchase. The current gasoline price is a good proxy for both the objective and subjective expectation of the future price (Anderson, Kellogg, and Sallee 2011). Using national prices rather than household-specific purchase prices avoids conflating loss aversion with household heterogeneity (Bell and Lattin 2000).

We expect the time horizon for expectation formation to be important because gain-loss utility is relevant only when there is surprise.¹¹ We choose a one-week horizon because the average household in the sample buys gasoline 4.6 times in the average purchase month. In the online appendix we present results using alternative time horizons.

Saliency Model. We estimate a model of salience based on Bordalo, Gennaioli, and Shleifer (2012). In the model, households place greater weight on product attributes that are salient, where salience is determined by the degree to which an attribute varies within an “evoked set” of options brought to mind by the purchase occasion. The evoked set typically includes both the current choice set and the historical choice set.

We assume that the choice set consists of the three grades of gasoline, each characterized by its octane level and its price. Primitive utility weights on octane and price are as in equation (8).

The household will overweight more salient attributes. Only the ordinal ranking of attribute salience matters, so with only two attributes it is sufficient to define an indicator z_{ijt} equal to one when price is more salient than octane for household i , grade j , and time t , and zero otherwise. The term Γ_{ijt} represents the difference in utility assigned to grade j by household i at time t due to local thinking (Bordalo, Gennaioli, and Shleifer 2012):

$$\Gamma_{ijt} = \gamma(1 - z_{ijt}) q_{ijt} \alpha g_j - \gamma(z_{ijt}) \mu q_{it} p_{jt} \quad (11)$$

where $\gamma(\cdot)$ is a function with $\gamma(1) \geq \gamma(0)$.¹²

It remains to specify what determines the salience indicator z_{ijt} . We specify this in two steps, both following Bordalo, Gennaioli, and Shleifer (2012). First, we define the salience function for an attribute x_{jt}

¹⁰With our assumptions, the estimating equation for Γ_{ijt} is given by

$$\Gamma_{ijt} = \tilde{\gamma} q_{it} \left(\alpha (g_j - E_{it}(g_j)) \mathbf{1}_{g_j \leq E_{it}(g_j)} - \mu (p_{jt} - E_{it}(p_{jt})) \mathbf{1}_{p_{jt} \geq E_{it}(p_{jt})} \right)$$

where $\tilde{\gamma}$ is a parameter and $\mathbf{1}_x$ is the indicator function. Only losses show up here because gains and losses cannot both be identified separately from consumption utility. Because households expect to buy regular gasoline, households never experience losses in the octane dimension, so it is variation in prices that identifies the parameter $\tilde{\gamma}$.

¹¹Formally, if a household knows its budget set in advance with certainty, then the household never experiences any gain-loss utility, $\Gamma_{ijt} = 0$, and the model collapses to a benchmark discrete-choice model with fungibility of money (Kőszegi and Rabin 2006, Proposition 3).

¹²In the notation of Bordalo, Gennaioli, and Shleifer (2012), $\gamma(1) + 1 = \frac{1}{1+\delta}$ and $\gamma(0) + 1 = \frac{\delta}{1+\delta}$ for some constant $\delta \in (0, 1]$, so that, when added to the primitive utility weights, the weights are $\frac{1}{1+\delta}$ on the more salient attribute and $\frac{\delta}{1+\delta}$ on the less salient attribute, respectively.

of grade j at time t to be $\sigma(x_{jt}, \bar{x}_{it}) = \frac{|x_{jt} - \bar{x}_{it}|}{|x_{jt}| + |\bar{x}_{it}|}$, where \bar{x}_{it} is the mean of the attribute in the evoked set of household i at time t . Second, we define the evoked set to consist of the current choice set plus the three grades of gasoline at historical national mean prices one week prior to purchase. In the online appendix we present results using a one-month horizon for the evoked set.

The economic content of these assumptions is as follows. First, an attribute's salience is judged by how much its value differs (in relative terms) from its mean value in the evoked set. Second, the evoked set includes options available in the past, so that, in particular, a difference between current and past prices can make price a more salient attribute.

We expect the specification of the evoked set and of the salience function to be important for the model's predictions. If we define the evoked set to contain only the current choice set, then a parallel increase in the price of all gasoline grades makes prices *less* salient and induces substitution to higher-octane gasolines. If we further modify the specification so that salience is determined by the difference between the maximum and minimum utility values of each attribute (Köszegi and Szeidl 2013), then parallel price changes will have no effect on salience and hence no effect on octane choice.

VII.B. Results and Discussion

Figure VII presents our findings. The first plot shows the observed regular share and its predicted path under a benchmark model with no psychological factors ($\Gamma_{ijt} = 0 \forall i, j, t$). The remaining plots show the predictions of the category budgeting, loss aversion, and salience models, respectively. The online appendix contains parameters estimates and formal measures of goodness-of-fit for each model.

Figure VII shows that both the loss aversion and salience models capture some of the empirical dynamics and that the two models exhibit a similar fit to the data. The figure also shows that neither model fits as well as the category budgeting model.

In the case of loss aversion, when prices rise, households are spending more than expected on gasoline, leading to a higher marginal utility of money and hence a tendency to switch to lower grades. Two features of the model prevent a better fit. First, the model predicts that if prices go up and stay up, the resulting surprise will last only for a week, so octane choice will quickly go back to its baseline level. Second, the model predicts that if prices rise continually for an extended period, the regular share will increase but level off, because once all households are in the losses region there is no channel by which further price increases can affect the marginal utility of money.

In the case of salience, when prices rise, the gap between present and past prices increases, making price more salient and leading households to put more weight on price relative to octane in decision-making and

hence to switch to regular grade. Two features of the model parallel those of loss aversion and prevent a better fit. First, because salience depends on deviations between the present and the past, during a period of consistently high (but not rising) prices, the regular share returns to its baseline level. Second, for the same reason, continual price increases do not produce continual increases in the regular share. A final important model feature is that falling prices can induce a shift to regular gasoline, because the salience function puts weight on attributes whose values differ from typical values, whether positively or negatively.

In summary, the fit of both models is limited because the models are too adaptive (expectations adjust too quickly) and too discrete (utility weights can only be high or low). To make these intuitions more precise, in figure VIII we show that both models fit the data almost perfectly when we assume that expectations are static and allow that utility weights can be a continuous function of gains and losses or of salience. Although these modifications are post hoc, they suggest potentially large gains in empirical validity from moving away from rational expectations and from allowing more continuous responses around reference points.

We omit analysis of some models whose predictions do not fit the primary patterns in the data.¹³ For example, models with “relative thinking” (Azar 2007 and 2011) predict that, when all prices increase, price differences become smaller in magnitude and therefore less salient, leading to quality upgrading. We find the opposite.¹⁴

VIII. CONCLUSIONS

We formulate a test of the fungibility of money based on parallel movements in prices of substitute goods of varying quality. We implement the test using panel data on household gasoline purchases. Households substitute from higher to lower octane levels when gasoline prices rise, to an extent that cannot be accounted for by income effects. A formal discrete-choice model rejects the null hypothesis that consumers treat gasoline expenditure and other income as fungible. Placebo tests using choices of non-gasoline products show that gasoline prices do not exert a disproportionate effect on purchases in non-gasoline domains.

We evaluate the performance of a set of psychologically rich models of decision-making in explaining the violation of fungibility in our data. We estimate an ad-hoc model of category budgeting, a model of loss

¹³We also omit analysis of models of social preferences in which octane switching serves to punish the retailer for unfair price increases. Because the retailer is not an oil major, such dynamics are unlikely in our setting (Kahneman, Knetsch, and Thaler 1986).

¹⁴Saini, Rao, and Monga (2010) offer a possible reconciliation of relative thinking evidence with our findings. They employ a hypothetical choice methodology and find that relative thinking appears to hold when hypothetical price changes are expected, but loss-aversion (“referent-thinking” in their model) holds instead when hypothetical price changes are unexpected. Because households cannot predict the path of future gasoline prices (Anderson, Kellogg, and Sallee 2011), it is reasonable to assume that referent thinking would dominate relative thinking in our context. Indeed, Saini, Rao, and Monga (2010) employ a gasoline-related vignette in their study, and show that higher-than-expected gasoline prices can induce consumers to state that they will drive further to seek out discounts.

aversion based on Köszegi and Rabin (2006), and a model of salience based on Bordalo, Gennaioli, and Shleifer (2012). We find that all three models can replicate features of the data that the neoclassical model cannot, and that the category budgeting model fits better than the other two. We also show that the fit of the loss aversion and salience models improves significantly when we assume static expectations (instead of rational expectations) and when we allow marginal utilities to be continuous in the gains and losses or in salience, suggesting possible avenues for improving these models' performance in future empirical applications.

APPENDIX: ADDITIONAL ROBUSTNESS CHECKS

Appendix table I presents a series of variants on the specification in column (1) of table III. In each case we present the marginal effect on the regular share of a \$1 change in gasoline prices, a \$1000 change in total expenditure, and a \$1000 change in gasoline expenditure, as well as the p-value from a test of the null hypothesis that $\eta^G = \eta^M$. Row (1) reproduces the baseline specification for comparison.

In rows (2) and (3) we drop the top and bottom 10% of households according to the number of total gasoline purchases they make in our sample period.

In row (4) we restrict attention to households that have a single adult member according to demographic information supplied by the retailer.

In row (5) we allow grade preferences to depend on the amount of gasoline purchased. We include in v_{ijt} an interaction between grade and the difference between the household's gasoline purchases in the transaction month and the household's mean monthly purchases over the sample period.

In rows (6) and (7) we estimate the model separately for households whose income is below or above the median in our sample. Following Petrin (2002), this approach allows for greater flexibility in the parametrization of the marginal utility function λ_{it} .

In row (8) we construct m_i as the predicted value from a regression of our baseline measure of m_i on per capita income in the household's Census block group.

In row (9) we adjust estimated gasoline expenditures $q_{it}p_{0t}$ to account for the correlation of gasoline prices with other energy prices. We do this by rescaling our estimate of household gasoline expenditures so that its ratio to household energy expenditures over the period we study is equal to the coefficient in an OLS regression of the annual change in log energy prices on the annual change in log gasoline prices. We measure energy prices using the price level for energy consumption from the NIPA, obtained from bea.gov in July 2010.

In row (10) we estimate gasoline expenditures $q_{it}p_{0t}$ as the product of the household's average annual gasoline purchases (in gallons) at the retailer and the price of regular gasoline at the retail location at the time of purchase.

In row (11) we identify η^G from variation in the spot price of oil instead of the US average price of regular gasoline. We do this by running a first-stage regression of gasoline expenditures $q_{it}p_{0t}$ on the oil price and allowing both v_{ijt} and λ_{it} to contain a linear term in the residual from the first-stage regression. We measure the oil spot price using a monthly series on the spot price of oil at Cushing, OK from the EIA, obtained at eia.doe.gov in December 2010.

In row (12) we restrict attention to purchases in which the price gap between midgrade and regular gasoline is 10 cents and the price gap between premium and regular gasoline is 20 cents, rounded to the nearest cent.

In row (13) we aggregate the data to the store-week level. We estimate our model letting all variables equal their store-week mean. We transform the market share of each grade of gasoline into the mean utility of that grade in a given store-week, relative to the share of regular gasoline (Berry 1994). We estimate via OLS, allowing for a store-week-grade utility shock that is mean zero conditional on the included variables.

In row (14) we allow that tastes v_{ijt} have an additive component distributed independently normal across choices and store-weeks. We estimate via maximum likelihood on a one percent sample of the data, approximating the likelihood via sparse grid integration (Heiss and Winschel 2008) with accuracy four.

In rows (15) and (16) we estimate the model separately for households with below- and above-median gas tank size, where we proxy for gas tank size with the maximum amount of gasoline purchased by the household across all transactions in the sample period.

Appendix table II presents a series of variants on the specification in column (2) of table III. The format follows that of appendix table I and row (1) reproduces the baseline specification for comparison.

In row (2) we adjust our time-varying expenditure measure for seasonality. To do this, we regress grocery expenditures in the four weeks prior to each gasoline transaction on a set of week-of-year dummies (defined so that the first dummy variable is equal to one in the first week of each year, and so on). We extract residuals from this model and center them at the overall sample mean four-week expenditure to obtain a seasonally adjusted measure of grocery expenditure. We then predict annual total expenditures from the resulting seasonally adjusted grocery expenditure measure following the procedure described in section IV.B.

In rows (3) and (4) we restrict attention to regular grocery shoppers, defined in two ways. In row (3) we restrict the sample to households who spend at least \$50 on non-gasoline purchases at the retailer in the four weeks prior to every gasoline transaction outside of the first sample month. In row (4) we restrict the sample to households who purchase milk at the retailer in more than half of our sample months.

In row (5) we drop transactions in the first sample month when our expenditure proxy (based on purchases over the prior four weeks) is poorly measured.

In row (6) we allow that the marginal utility of money μ_i is a linear function of the mean value of expenditure m_{it} for household i . Doing this eliminates any remaining identification from cross-sectional variation in grocery expenditures.

REFERENCES

- Abeler, Johannes, and Felix Marklein, “Fungibility, Labels and Consumption,” IZA Discussion Paper No. 3500, 2008.
- Allenby, Greg M., and Peter E. Rossi, “Quality Perceptions and Asymmetric Switching Between Brands,” *Marketing Science*, 10, no. 3 (1991), 185-204.
- Anderson, Soren, Ryan Kellogg, and James M. Sallee, “What Do Consumers Believe About Future Gasoline Prices?” NBER Working Paper No. 16974, 2011.
- Antonides, Gerrit, I. Manon de Groot, and W. Fred van Raaij, “Mental Budgeting and the Management of Household Finance,” *Journal of Economic Psychology*, 32, no. 4 (2011), 546-555.
- API, “Gasoline – Is it All the Same? What About Octane?” American Petroleum Institute (2010), accessed at <<http://www.consumerreports.org/cro/cars/tires-auto-parts/car-maintenance/premium-gas-a-waste-of-money/overview/index.htm>> on November 4, 2010, posted on December 10, 2008.
- Azar, Ofer H., “Relative Thinking Theory,” *Journal of Socio-Economics*, 36, no. 1 (2007), 1-14.
- , “Does Relative Thinking Exist in Real-World Situations? A Field Experiment with Bagels and Cream Cheese,” *Economic Inquiry*, 42, no. 2 (2011), 564-572.
- Barseghyan, Levon, Francesca Molinari, Ted O’Donoghue, and Joshua C. Teitelbaum, “The Nature of Risk Preferences: Evidence from Insurance Choices,” *American Economic Review*, forthcoming, 2013.
- Barzel, Yoram, “An Alternative Approach to the Analysis of Taxation,” *Journal of Political Economy*, 85, no. 6 (1976), 1177-1197.
- Bell, David R., and James M. Lattin, “Looking for Loss Aversion in Scanner Panel Data: The Confounding Effect of Price Response Heterogeneity,” *Marketing Science*, 19, no. 2 (2000), 185-200.
- Berry, Steven, “Estimating Discrete Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25 (1994), 242-262.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, “Salience and Consumer Choice,” Harvard University Working Paper, 2012.
- Byrd, Richard H., Jorge Nosedal, and Richard A. Waltz, “KNITRO: An Integrated Package for Nonlinear

- Optimization,” In *Large-Scale Nonlinear Optimization*, Gianni di Pillo and Massimo Roma, eds. (New York: Springer-Verlag, 2006).
- Carty, Sharon S., “Luxury Cars Switch from Premium to Regular Gas Diet,” *USA Today*, accessed at <http://www.usatoday.com/money/autos/2008-08-06-premium-gas-regular_N.htm> on November 3, 2010, posted on August 7, 2008.
- Chang, Kwangpil, S. Siddarth, and Charles B. Weinberg, “The Impact of Heterogeneity in Purchase Timing and Price Responsiveness on Estimates of Sticker Shock Effects,” *Marketing Science*, 18, no. 2 (1999), 178-192.
- Click and Clack, “Premium Versus Regular: Myths of Premium Gas,” accessed at <<http://www.cartalk.com/content/features/premium/myths.html>> on November 4, 2010.
- Coats, R. Morris, Gary M. Pecquet, and Leon Taylor, “The Pricing of Gasoline Grades and the Third Law of Demand,” Nicholls State University Working Paper, 2005.
- Conlin, Michael, Ted O’Donoghue, and Timothy J. Vogelsang, “Projection Bias in Catalog Orders,” *American Economic Review*, 97, no. 4 (2007), 1217-1249.
- Consumer Reports, “Paying for Premium Gas Can Be a Waste of Money,” accessed at <<http://www.consumerreports.org/cro/cars/tires-auto-parts/car-maintenance/premium-gas-a-waste-of-money/overview/index.htm>> on November 4, 2010, posted in March 2010.
- Crawford, Vincent P., and Juanjuan Meng, “New York City Cab Drivers’ Labor Supply Revisited: Reference-Dependent Preferences with Rational Expectations Targets for Hours and Income,” *American Economic Review*, 101, no. 5 (2011), 1912-1932.
- DellaVigna, Stefano, John List, and Ulrike Malmendier, “Testing for Altruism and Social Pressure in Charitable Giving,” *Quarterly Journal of Economics*, 127, (2012), 1-56.
- Douglass, Elizabeth, “Octane’s Allure Hurt by High Cost,” *Los Angeles Times*, accessed at <<http://articles.latimes.com/2005/nov/28/business/fi-premium28>> on November 3, 2010, posted on November 28, 2005.
- Dubé, Jean-Pierre, “Multiple Discreteness and Product Differentiation: Demand for Carbonated Soft Drinks,” *Marketing Science*, 23, no. 1 (2004), 66-81.
- EIA, *Energy Information Administration Petroleum Marketing Annual: Explanatory Notes*. Energy Information Administration, Department of Energy. accessed at <http://www.eia.doe.gov/oil_gas/petroleum/data_publications/petroleum_marketing_annual/pma.html> on March 21, 2011, posted in 2009.
- EIA, *Definitions of Gasoline Grades*. Energy Information Administration, Department of Energy, accessed

- at <http://www.eia.doe.gov/petroleum/data_publications/wrgp/gasoline_grades.html> on November 4, 2010, posted in 2010.
- Environmental Protection Agency, *Fuel Economy Data*, accessed at <<http://www.fueleconomy.gov/feg/download.shtml>> on May 31, 2011, posted in 2011.
- Feldman, Naomi E., "Mental Accounting Effects of Income Tax Shifting," *Review of Economics and Statistics*, 92, no. 1 (2010), 70-86.
- Fogel, Suzanne, Dan Lovallo, and Carmina Caringal, "Loss Aversion for Quality in Consumer Choice," *Australian Journal of Management*, 29, no. 1 (2004), 45-63.
- Gary, James H., and Glenn E. Handwerk, *Petroleum Refining: Technology and Economics*, 4th ed. (New York: Marcel Dekker, Inc., 2001).
- Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas, "Revisiting the Income Effect: Gasoline Prices and Grocery Purchases," NBER Working Paper No. 13614, 2007.
- Gilboa, Itzhak, and Eva Gilboa-Schechtman. "Mental Accounting and the Absentminded Driver," In *The Psychology of Economic Decisions, Vol. 1: Rationality and Well-Being*, Isabelle Brocas and Juan D. Carrillo, eds. (Oxford: Oxford University Press, 2003).
- Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler. "The Complexity of the Consumer Problem and Mental Accounting," University of Pennsylvania Working Paper, 2010.
- Grubb, Michael D., and Matthew Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," MIT Working Paper, 2012.
- Hanemann, W. Michael, "Discrete/Continuous Models of Consumer Demand," *Econometrica*, 52, no. 3 (1984), 541-561.
- Heath, Chip, and Jack B. Soll, "Mental Budgeting and Consumer Decisions," *The Journal of Consumer Research*, 23, no. 1 (1996), 40-52.
- Heiss, Florian, and Viktor Winschel, "Likelihood Approximation by Numerical Integration on Sparse Grids," *Journal of Econometrics*, 144, no. 1 (2008), 62-80.
- Houde, Jean-François, "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, 102, no. 5 (2012): 2147-82.
- Jappelli, Tullio, and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2 (2010), 479-506.
- Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler, "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review*, 76, no. 4 (1986), 728-741.
- Kooreman, Peter, "The Labeling Effect of a Child Benefit System," *American Economic Review*, 90, no. 3

- (2000), 571-583.
- Kőszegi, Botond, and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 71, no. 4 (2006), 1133-1165.
- Kőszegi, Botond, and Adam Szeidl, "A Model of Focusing in Economic Choice," *Quarterly Journal of Economics*, 128, no. 1 (2013), 53-107.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices Over the Lifecycle," Harvard University Mimeograph, 2007.
- Lawson, Robert, and Lauren Raymer, "Testing the Alchian-Allen Theorem: A Study of Consumer Behavior in the Gasoline Market," *Economics Bulletin*, 4, no. 35 (2006), 1-6.
- Lewis, Matthew, "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market," *Journal of Economics and Management Strategy*, 20, no. 2 (2011), 409-449.
- Lidderdale, Tancred, "Gasoline Demand Trends," *EIA Energy Outlook Modeling and Data Conference March 8, 2007*, accessed at <<http://www.eia.doe.gov/oiaf/aeo/conf/pdf/lidderdale.pdf>> on December 14, 2010, posted in 2007.
- Linebaugh, Kate, and Matthew Dolan, "US Auto Sales Plunged in October," *Wall Street Journal*, accessed at <<http://online.wsj.com/article/SB122573166905093595.html>> on October 18, 2011, posted on November 4, 2008.
- Lundberg, Shelly, and Robert A. Pollak, "Separate Spheres Bargaining and the Marriage Market," *Journal of Political Economy*, 101, no. 6 (1993), 988-1010.
- Lush, Tamara, "Premium Gas Sales Tank as Fuel Prices Rise," Associated Press, accessed at <http://www.usatoday.com/money/economy/2008-06-19-1622609293_x.htm> on November 3, 2010, posted on June 19, 2008.
- McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," In *Frontiers in Econometrics*, Paul Zarembka, ed. (New York: Academic Press, 1973).
- Mian, Atif, and Amir Sufi, "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *Quarterly Journal of Economics*, 124, no. 4 (2009), 1449-1496.
- Milkman, Katherine L., and John Beshears, "Mental Accounting and Small Windfalls: Evidence from an Online Grocer," *Journal of Economic Behavior and Organization*, 72, no. 2 (2009), 384-394.
- Missouri Census Data Center, "Geographic Correspondence Engine with Census 2000 Geography," accessed at <<http://mcdc2.missouri.edu/websas/geocorr2k.html>> on July 22, 2011.
- Murphy, Kevin M., and Robert H. Topel, "Estimation and Inference in Two-Step Econometric Models," *Journal of Business and Economic Statistics*, 3, no. 4 (1987), 370-379.

- Nesbit, Todd, "Taxation and Product Quality: The Gasoline Market," *Economic Issues*, 12, no. 2 (2007), 1-14.
- NielsenWire, "Growth of Organic Sales Slows with Recession," accessed at <http://blog.nielsen.com/nielsenwire/consumer/sales-of-organic-products-dive-with-recession/> on October 28, 2011, posted on April 23, 2009.
- Pakes, Ariel, Michael Ostrovsky, and Steven Berry, "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *RAND Journal of Economics*, 38, no. 2 (2007), 373-399.
- Petrin, Amil, "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110 (2002), 705-729.
- Rainwater, Lee, Richard P. Coleman, and Gerald Handel, *Workingman's Wife: Her Personality, World and Life Style*, (New York, NY: McFadden Books, 1959).
- Read, Daniel, George Loewenstein, and Matthew Rabin, "Choice Bracketing," *Journal of Risk and Uncertainty*, 19, no. 1-3 (1999), 171-197.
- Reed, Philip, "Do You Really Need Premium? And Answers to Other Gasoline Questions," Edmunds.com, accessed at <http://www.edmunds.com/advice/fueleconomy/articles/106293/article.html> on November 3, 2010, posted on August 30, 2007.
- Reuters, "Top-20 Selling Vehicles in U.S. Through July," accessed at <http://www.reuters.com/article/2007/08/01/autos-usa-chart-idUSN0140862420070801> on April 21, 2011, posted on August 1, 2007.
- Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod, "Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How it is Delivered?" *American Economic Journal: Economic Policy*, forthcoming, 2013.
- Saini, Ritesh, Raghunath Singh Rao, and Ashwani Monga, "Is that Deal Worth My Time? The Interactive Effect of Relative and Referent Thinking on Willingness to Seek a Bargain," *Journal of Marketing*, 74, no. 1 (2010), 34-48.
- Setiawan, Winardi, and Daniel Sperling, "Premium Gasoline Overbuying in the US: Consumer-Based Choice Analysis," University of California Transportation Center Working Paper No. 457, Berkeley, CA, 1993.
- Skinner, Jonathan, "A Superior Measure of Consumption from the Panel Study of Income Dynamics," *Economics Letters*, 23, no. 2 (1987), 213-287.
- Sobel, Russell S., and Thomas A. Garrett, "Taxation and Product Quality: New Evidence from Generic Cigarettes," *Journal of Political Economy*, 105, no. 4 (1997), 880-887.

- Stephens Jr., Melvin, "The Long-Run Consumption Effects of Earnings Shocks," *Review of Economics and Statistics*, 83, no. 1 (2001), 28-36.
- , "Job Loss Expectations, Realizations, and Household Consumption Behavior," *Review of Economics and Statistics*, 86, no. 1 (2004), 253-269.
- Taylor, John B., "The Financial Crisis and the Policy Responses: An Empirical Analysis of What Went Wrong," *Critical Review*, 21, no. 2-3 (2009), 341-364.
- Thaler, Richard H., "Anomalies: Saving, Fungibility, and Mental Accounts," *Journal of Economic Perspectives*, 4, no. 1 (1990), 193-205.
- , "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 12 (1999), 183-206.
- Willig, Robert D., "Incremental Consumer's Surplus and Hedonic Price Adjustment," *Journal of Economic Theory*, 17, no. 2 (1978), 227-253.
- Zelizer, Viviana A., "Making Multiple Monies," In *Explorations in Economic Sociology*, Richard Swedberg, ed. (New York: Sage Foundation, 1993).
- Zimmerman, Ann, Jennifer Saranow, and Miguel Bustillo, "Retail Sales Plummet," *Wall Street Journal*, accessed at <<http://online.wsj.com/article/SB123025036865134309.html>> on October 18, 2011, posted on December 26, 2008.

TABLE I: DESCRIPTIVE STATISTICS FOR RETAILER DATA

Sample	(1) All retailer households	(2) Estimation sample	(3) All households in same state
Household income measure provided by retailer (in 2008 \$US)	\$86,968	\$97,173	
In household's Census block group:			
Average household income (in 2008 \$US)	\$95,421	\$98,355	\$81,252
Average commute time among workers	26.292	26.941	26.780
Fraction of workers commuting using public transportation	0.028	0.024	0.038
Number of gasoline trips per month (conditional on at least one trip)	1.664	4.601	
Average gallons per purchase occasion	11.740	12.323	
Average distance from block group centroid to most visited store (in miles)	20.579	4.100	
Fraction of gallons purchased which are regular grade	0.800	0.792	0.822
Average retail price paid per gallon	\$2.836	\$2.848	\$2.892
Average 2008 grocery expenditure	\$506	\$2554	\$4295
Number of households	1,306,748	61,494	—

Notes: In columns (1) and (2) the table shows the mean across households in each sample. Column (1) shows statistics for all households in the retailer database who ever purchase gasoline at one of the retailer's outlets. Column (2) shows statistics for the estimation sample. Column (3) shows the mean across US states, where states are weighted by the proportion of households in the full sample who reside in each state. Census block group characteristics are missing for 5.5 percent of households in the full sample and 5.0 percent of households in the estimation sample. Distance to most visited store is treated as missing for households living in a different state from the most visited store (<8% of all households and <1% of households in estimation sample). Characteristics of the Census block group (average household income, average commute time, and fraction of workers using public transportation) are from the 2000 US Census and are averaged (with population weights) at the state level to compute the statistics in column (3). Average distance from block group centroid to most visited store is computed using data from the 2000 US Census on the latitude and longitude of block groups and data from the retailer on the latitude and longitude of each store location. The state equivalent measure for fraction of gallons purchased of regular grade and average price paid is based on 2006-2009 EIA data for the state of the household's most visited store. The state equivalent measure for average grocery expenditure is from the 2008 Consumer Expenditure Survey Interview Files for the state of the household's most visited store.

TABLE II: INCOME ELASTICITY OF GROCERY EXPENDITURES

Dependent variable: Log(Average monthly grocery expenditures)				
	(1)	(2)	(3)	(4)
Log(Household income)	0.1686 (0.0014)	0.1682 (0.0076)		
Log(Average adjusted gross income in zipcode-year)			0.1379 (0.0207)	0.0893 (0.0528)
Unit of Analysis	Household	Household	Zipcode-year	Zipcode-year
Data source	CEX	Retailer	Retailer	Retailer
Year dummies?			X	X
Zipcode dummies?				X
Number of observations	122483	61333	3820	3820
R^2	0.1072	0.0079	0.0151	0.9725

Note: Data for specification (1) are from the Consumer Expenditure Survey 2006-2009 Interview Files and use family income before tax as the household income concept. Data for specifications (2) through (4) are from the retailer. All specifications exclude households/zipcode-years with zero expenditure on groceries during sample period. Regressions at the zipcode-year level are weighted by the number of sample households in the zipcode-year. Specifications (3) and (4) use data on the average adjusted gross income in each zipcode from the IRS, described in Mian and Sufi (2009).

TABLE III: MODEL OF GASOLINE GRADE CHOICE

Dependent variable: Choice of gasoline grade				
	(1)	(2)	(3)	(4)
Effect on marginal utility of:				
\$1000 increase in gasoline expenditures (Parameter η^G)	0.4306 (0.0314)	0.4327 (0.0305)	0.4132 (0.0335)	0.7145 (0.0317)
\$1000 decrease in total expenditures (Parameter η^M)	0.0293 (0.0008)	0.0127 (0.0005)	0.0297 (0.0010)	0.0416 (0.0042)
Average marginal effect on regular share of:				
\$1 increase in price of regular gasoline	0.0142 (0.0010)	0.0143 (0.0010)	0.0136 (0.0011)	0.0140 (0.0006)
\$1000 decrease in gasoline expenditures	-0.0120 (0.0009)	-0.0121 (0.0009)	-0.0115 (0.0009)	-0.0118 (0.0005)
\$1000 increase in total expenditures	-0.0008 (0.0000)	-0.0004 (0.0000)	-0.0008 (0.0000)	-0.0007 (0.0001)
Ratio of effects on marginal utility (η^G/η^M)	14.68	33.99	13.90	17.17
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.0000	0.0000	0.0000	0.0000
Sample	All	All	1/10th	1/10th
Time-varying expenditure measure?		X		
Household-level random coefficients?				X
Number of transactions	10548175	10548175	1082486	1082486
Number of households	61494	61494	61494	61494

Note: Data are from retailer. Table reports estimates of the model described in section IV. Standard errors in parentheses allow for correlation in residuals by month. Models are estimated via maximum likelihood. In specification (1) we measure total expenditures with our time-constant measure m_i . We assume that α_{ij} and μ_i are constant across households. In specification (2) we measure total expenditures with our time-varying measure m_{it} and allow that μ_i is a linear function of m_i . Specification (3) repeats specification (1) on a sample of every 10th transaction for each household. Specification (4) uses the sample in specification (3) and allows that α_{ij} are distributed independently normal across households and choices. To estimate the mixed logit model in specification (4), we approximate the likelihood using sparse grid integration with accuracy 9 (Heiss and Winschel 2008) and maximize the likelihood using KNITRO's active-set algorithm for unconstrained problems (Byrd et al. 2006). We validated our implementation of the mixed logit model by replicating benchmark Monte Carlo exercises reported in Heiss and Winschel (2008).

TABLE IV: PLACEBO MODEL OF NON-GASOLINE CHOICE
Dependent variable: Choice of brand

	(1)	(2)
Effect on marginal utility of:		
\$1000 increase in gasoline expenditures (Parameter η^G)	-0.0141 (0.0250)	-0.0128 (0.0197)
\$1000 decrease in total expenditures (Parameter η^M)	0.0044 (0.0002)	0.0034 (0.0001)
Average marginal effect on private label share of:		
\$1 increase in price of regular gasoline	-0.0169 (0.0299)	-0.0055 (0.0084)
\$1000 decrease in gasoline expenditures	0.0143 (0.0252)	0.0046 (0.0071)
\$1000 increase in total expenditures	-0.0045 (0.0002)	-0.0012 (0.0000)
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.4571	0.4115
p-value of test that η^G/η^M for category is equal to η^G/η^M for gasoline grade choice	0.0000	0.0000
Category	OJ	Milk
Number of transactions	411161	2210312
Number of households	13493	34128

Note: Data are from retailer. Table reports estimates of the model described in section IV but applied to choice of orange juice or milk brand rather than choice of gasoline grade. See online appendix for details on the construction of choice sets for milk and orange juice. Standard errors in parentheses allow for correlation in residuals by month. We assume that α_{ij} and μ_i are constant across households. The test that η^G/η^M for the category is equal to is equal to η^G/η^M for gasoline grade choice is performed via a jackknife over months, accounting for the correlation in η^G/η^M across models. The test uses our baseline specification from column (1) of table III.


TABLE V: EXPLOITING TIME-SERIES VARIATION IN INCOME

	Average marginal effect on regular share of:			Fungibility
	Increase in gas price (\$1)	Increase in gas expenditure (\$1000)	Decrease in total expenditure (\$1000)	p-value
Baseline	0.0142 (0.0010)	-0.0120 (0.0009)	-0.0008 (0.0000)	0.0000
Total expenditure evolves in proportion to aggregate per household:				
Income	0.0140 (0.0010)	-0.0119 (0.0008)	0.0002 (0.0004)	0.0000
Wealth (S&P 500 index)	0.0151 (0.0013)	-0.0128 (0.0011)	-0.0001 (0.0001)	0.0000
Consumption expenditures	0.0143 (0.0010)	-0.0120 (0.0009)	-0.0000 (0.0005)	0.0000


Notes: The baseline specification is from column (1) of table III. Each subsequent specification assumes that total expenditure m_{it} is equal to the product of m_i and the ratio of the specified aggregate indicator to its value in January 2008 and allows that μ_i is a linear function of m_i . Standard errors in parentheses allow for correlation in residuals by month. See section III.B for details on the construction of measures of aggregate household income, wealth, and consumption expenditures.

FIGURE I: Regular share and price of regular gasoline (aggregate data)

Panel A: 1990-1999




Panel B: 2000-2009




Notes: Data are from the EIA. Each panel plots the monthly US market share of regular gasoline and the monthly US average price of regular gasoline (in 2005 US dollars). The level shift in the share of regular gasoline at the beginning of 1996 coincides with a change in the EIA survey instrument. Prices are converted to 2005 dollars using the NIPA price index for personal consumption expenditures excluding food and energy.

FIGURE II: Regular share and the price of regular gasoline (retailer data)




Notes: Data are from the retailer. The plot shows the weekly share of transactions that go to regular gasoline and the weekly average transaction price of regular gasoline (in current US dollars).

FIGURE III: The role of compositional change




Notes: Data are from the retailer. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline. The next two lines are based on a transaction-level regression of an indicator for purchase of regular gasoline on household fixed effects. The line labeled “no composition” is the weekly average residual from the regression, normalized to have the same mean as the observed series. The line labeled “composition only” is the weekly average predicted value from the regression.

FIGURE IV: Regular share and household income



Notes: Data are from the retailer. The plot shows the average share of transactions that go to regular gasoline against the average household income for different groups of households. Households are grouped into bins of width \$10,000, i.e. into bins \$20,000-29,999, \$30,000-\$39,999, etc. Figure excludes households with income below the 5th percentile or above the 95th percentile. The area of each symbol is proportional to the number of sample households in each income bin. The line shown is the prediction from an OLS regression of average regular share on average household income, weighted by the number of sample households in each bin.


FIGURE V: Model fit, with and without fungibility




Notes: Data are from the retailer. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline. The line labeled “predicted: unconstrained” shows the average predicted probability of buying regular gasoline from the baseline model in column (1) of table III. The line labeled “predicted: constrained” shows the average predicted probability of buying regular gasoline from the same model, re-estimated imposing the constraint that $\eta^G = \eta^M$.

FIGURE VI: Permanent income shocks and the business cycle

Panel A: 2006-2009 (retailer data)

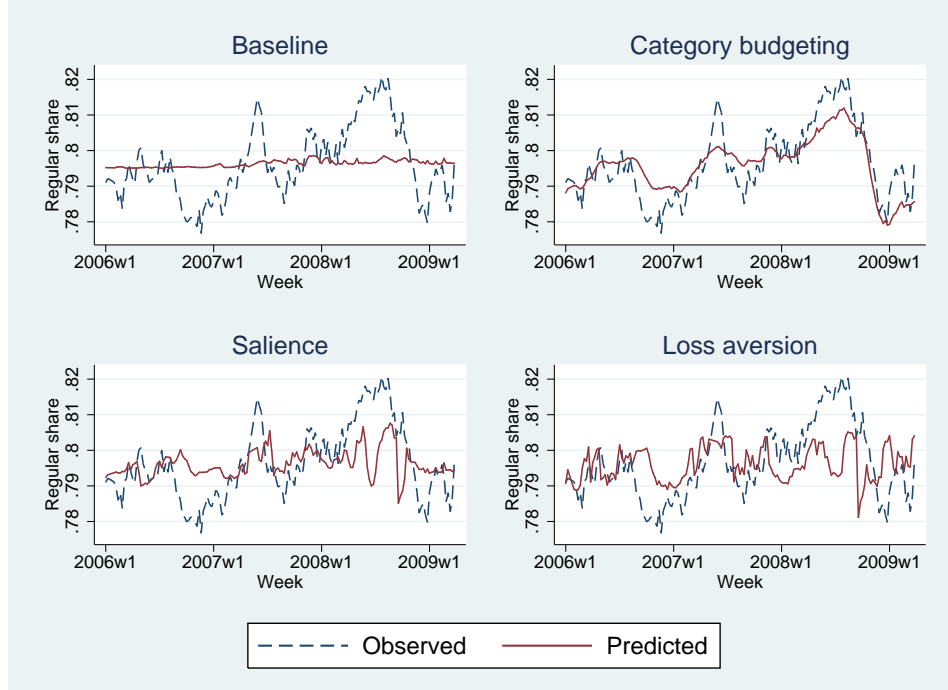


Panel B: 1990-1991 (aggregate data)



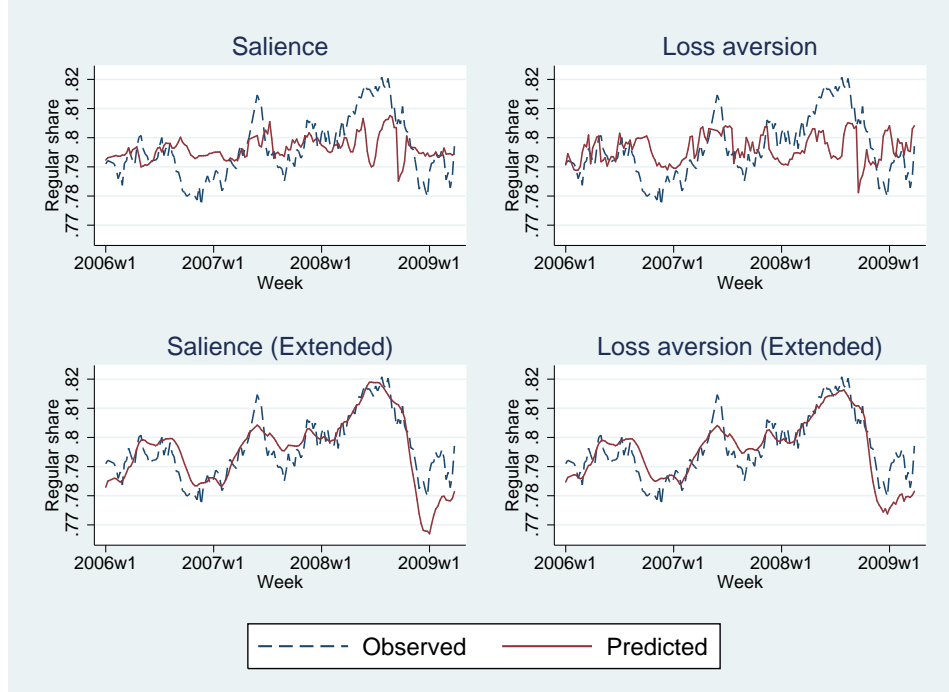
Notes: The line labeled “household expenditure” is total consumption expenditure per household (in 2008 US dollars), computed from the National Income and Product Accounts and the US Census Historical Time Series Tables. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline, measured using retailer data in Panel A and aggregate (EIA) data in Panel B. The line labeled “predicted: income effect” shows the average predicted probability of buying regular gasoline from a model in which the regular share is a linear function of household expenditure with a slope chosen to match the average marginal effect of gasoline expenditure in the baseline model in column (1) of table III and an intercept chosen to center the prediction at the empirical mean during the period shown.

FIGURE VII: Fit of psychological models of decision-making



Notes: Data are from the retailer. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline. The line labeled “predicted” shows the average predicted probability of buying regular gasoline from estimates of the model in equation (8) with different specifications of Γ_{ijt} . In the baseline model we set $\Gamma_{ijt} = 0 \forall i, j, t$. In the category budgeting model we specify Γ_{ijt} as in equation (9). In the loss aversion model we specify Γ_{ijt} as in equation (10) and we assume that household expectations are based on national gasoline prices one week prior to purchase. In the saliency model we specify Γ_{ijt} as in equation (11) and we assume that the evoked set consists of both the current choice set and the set of all gasoline grades at their national mean prices one week prior to purchase. See section VII for additional details.

FIGURE VIII: Extensions to psychological models



Notes: Data are from the retailer. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline. The line labeled “predicted” shows the average predicted probability of buying regular gasoline from estimates of the model in equation (8) with different specifications of Γ_{ijt} . The first row of plots are reproduced from figure VII. The second row of plots are from extended models specified as follows. In the extended loss aversion model we specify Γ_{ijt} as

$$\Gamma_{ijt} = \tilde{\gamma} q_{it} (\alpha (g_j - E_{it}(g_j)) \logit^{-1}(E_{it}(g_j) - g_j) - \mu (p_{jt} - E_{it}(p_{jt})) \logit^{-1}(p_{jt} - E_{it}(p_{jt})))$$

and we assume that household expectations of price or octane are given by the variable’s sample mean. In the extended salience model we specify Γ_{ijt} as

$$\Gamma_{ijt} = \tilde{\gamma}_0 \frac{\sigma(g_j, \bar{g})}{\sigma(g_{jt}, \bar{g}) + \sigma(p_{jt}, \bar{p}_{it})} q_{ijt} \alpha g_j - \tilde{\gamma}_1 \frac{\sigma(p_{jt}, \bar{p}_{it})}{\sigma(g_{jt}, \bar{g}_{it}) + \sigma(p_{jt}, \bar{p}_{it})} \mu q_{it} p_{jt}$$

and we specify the evoked set to consist of the current choice set plus the three grades of gasoline at prices of \$1.00, \$1.10 and \$1.20. See section VII for additional details.

APPENDIX TABLE I: ALTERNATIVE SPECIFICATIONS
(TIME-CONSTANT EXPENDITURE MEASURE)

		Average marginal effect on regular share of:			Ratio of	Fungibility
		Increase in gas price (\$1)	Increase in gas expenditure (\$1000)	Decrease in total expenditure (\$1000)	gas effect to total effect	p-value
(1)	Baseline	0.0143 (0.0010)	-0.0121 (0.0009)	-0.0008 (0.0000)	15.18	0.0000
(2)	Drop 10% of households who buy gas least often	0.0136 (0.0019)	-0.0115 (0.0016)	-0.0012 (0.0000)	9.40	0.0000
(3)	Drop 10% of households who buy gas most often	0.0142 (0.0010)	-0.0120 (0.0009)	-0.0008 (0.0000)	14.68	0.0000
(4)	Single-adult households	0.0149 (0.0010)	-0.0126 (0.0009)	-0.0013 (0.0000)	9.99	0.0000
(5)	Control for gallons purchased	0.0146 (0.0010)	-0.0123 (0.0009)	-0.0008 (0.0000)	15.04	0.0000
(6)	Households with below-median income	0.0151 (0.0010)	-0.0128 (0.0008)	-0.0010 (0.0000)	13.31	0.0000
(7)	Households with above-median income	0.0138 (0.0012)	-0.0117 (0.0010)	-0.0010 (0.0000)	11.35	0.0000
(8)	Predict total expenditure from Census block income	0.0143 (0.0010)	-0.0120 (0.0009)	-0.0014 (0.0000)	8.55	0.0000
(9)	Account for correlation with energy prices	0.0142 (0.0010)	-0.0104 (0.0008)	-0.0008 (0.0000)	12.70	0.0000
(10)	Estimate gasoline expenditure from retailer data	0.0151 (0.0012)	-0.0211 (0.0017)	-0.0008 (0.0000)	25.18	0.0000
(11)	Identify from variation in world oil price	0.0137 (0.0012)	-0.0116 (0.0010)	-0.0008 (0.0000)	14.00	0.0000
(12)	Restrict to transactions with 10-cent price gaps	0.0162 (0.0023)	-0.0137 (0.0020)	-0.0008 (0.0000)	16.28	0.0000
(13)	Aggregate to store-week level	0.0147 (0.0011)	-0.0125 (0.0009)	-0.0012 (0.0000)	10.57	0.0000
(14)	Allow store-week-level preference shock	0.0121 (0.0023)	-0.0103 (0.0020)	-0.0008 (0.0001)	13.41	0.0000
(15)	Below-median tank size	0.0150 (0.0011)	-0.0127 (0.0009)	-0.0010 (0.0000)	12.65	0.0000
(16)	Above-median tank size	0.0133 (0.0010)	-0.0113 (0.0009)	-0.0008 (0.0000)	14.03	0.0000

Notes: The baseline specification is from column (1) of table III. Other specifications are variants on the baseline specification. See appendix for details.

APPENDIX TABLE II: ALTERNATIVE SPECIFICATIONS
(TIME-VARYING EXPENDITURE MEASURE)

		Average marginal effect on regular share of:			Ratio of	Fungibility
		Increase in gas price (\$1)	Increase in gas expenditure (\$1000)	Decrease in total expenditure (\$1000)	gas effect to total effect	p-value
(1)	Baseline	0.0143 (0.0010)	-0.0121 (0.0009)	-0.0004 (0.0000)	33.99	0.0000
(2)	Adjust grocery expenditures for seasonal variation	0.0142 (0.0010)	-0.0120 (0.0009)	-0.0004 (0.0000)	33.98	0.0000
(3)	Restrict to households making regular grocery purchases	0.0161 (0.0012)	-0.0136 (0.0010)	-0.0003 (0.0000)	46.95	0.0000
(4)	Restrict to households making regular milk purchases	0.0148 (0.0010)	-0.0125 (0.0008)	-0.0003 (0.0000)	48.99	0.0000
(5)	Drop first sample month	0.0144 (0.0010)	-0.0122 (0.0009)	-0.0004 (0.0000)	34.42	0.0000
(6)	Control for household's mean grocery expenditures	0.0143 (0.0010)	-0.0121 (0.0009)	-0.0001 (0.0000)	93.65	0.0000

Notes: The baseline specification is from column (2) of table III. Other specifications are variants on the baseline specification. See appendix for details.