

ONLINE APPENDIX: FUNGIBILITY AND CONSUMER CHOICE*

JUSTINE HASTINGS AND JESSE M. SHAPIRO

April 2013

I. ADDITIONAL FINDINGS AND SPECIFICATIONS

Online appendix figure I shows the dynamics of the effect of a change in the price of regular gasoline on the share of gasoline that is regular grade. The plot shows that a permanent increase in the regular price would increase the share of regular gasoline for at least six months.

Online appendix figure II shows the distribution across households of the correlation between buying regular gasoline and the price of regular gasoline. This correlation is positive for 59 percent of households.

Online appendix figure III shows that the premium-regular price gap tends to compress when the regular-grade price rises. Online appendix figure IV shows evidence on the dynamics of this compression. After an increase in the regular price, the gap in prices falls, but the decrease dissipates within two months.

Online appendix figure V illustrates the findings of our placebo analysis. In each panel we estimate a logit model with a product-specific constant, allowing the marginal utility of money to depend on household income and to vary flexibly across weeks in the data. Using this model we can calculate the household income implied by purchasing behavior in each week: the household income level that would rationalize the average household's behavior in a given week with the model estimated in a baseline week. For gasoline, the implied income series moves inversely with prices. This is simply another representation of our main finding: as the gas price rises, households act far poorer than can be reconciled with the plausible income effect of gasoline. For orange juice and milk purchases, we do not observe such behavior. Indeed, during the large fall in gasoline prices in the autumn of 2008, households act, if anything, slightly poorer when choosing orange juice or milk brands, consistent with the onset of the financial crisis but not with their behavior at the pump.

*E-mail: justine_hastings@brown.edu, jesse.shapiro@chicagobooth.edu.

Online appendix figure VI plots the average automobile speed against the price of regular gasoline. We use data from the National Automotive Sampling System General Estimates System (NASS-GES) to construct a monthly time series of traveling speeds for 1990-2009 (National Highway Traffic Safety Administration 2009). The data come from a nationally representative sample of police-reported motor vehicle accidents. We construct a monthly series of the average traveling speed for automobiles traveling more than 0 miles per hour and at or below 95 miles per hour that were struck by another vehicle. (We exclude automobiles that struck other vehicles because these automobiles are likely to be less representative of vehicles on the road.) A regression in first differences reveals no statistically significant relationship between gasoline prices and traveling speeds.

Online appendix table I presents tests of fungibility under successively weaker assumptions about the functional form of the indirect utility $\Lambda_i(\cdot)$ from non-gasoline expenditure. The first specification shows our baseline model, which assumes that the marginal utility is linear in expenditure. The second, third, and fourth specifications assume that $\Lambda_i(\cdot)$ is quadratic, cubic, and quartic, respectively. Each specification includes an interaction between gasoline expenditure $q_{it}p_{0t}$ and the grade price p_{jt} . The coefficient on this interaction is a measure of the extent to which the marginal utility of money responds more to gas price changes than to comparable changes in other income. In all specifications the “excess response” to gas price changes is comparable to that in our baseline specification, and in all specifications we confidently reject the null hypothesis of fungibility.

Online appendix table II presents linear probability models of the effect of total expenditures and gasoline expenditures on the propensity to buy regular gasoline with various specifications and controls. In all cases we can confidently reject the null hypothesis that gasoline expenditures and total expenditures exert the same effect on the propensity to buy regular gasoline, analogous to rejecting the null hypothesis of fungibility in our baseline econometric model.

Online appendix table III presents alternative specifications of preference heterogeneity. The first specification is our baseline model, in which we assume that α_{ij} and μ_i are constant across households. The second, also presented in the body of the paper, assumes that α_{ij} are independently normally distributed, and is estimated with quadrature accuracy 9 on a subsample of the data. The third specification repeats the second using the full sample and accuracy 3. (Accuracy and sample size trade off due to memory constraints.) The fourth specification allows for unobserved heterogeneity in μ_i instead of α_{ij} . The fifth specification relaxes the distributional assumptions on α_{ij} and μ_i by estimating the model separately for each household and estimating parameters η^M and η^G via FGLS from the household-specific estimates, under the assumption that μ_i is uncorrelated with m_{it} conditional on α_{ij} . This specification restricts attention

to households that buy all three grades of gasoline at some point in the sample period. (For households that do not buy all grades, the constants α_{ij} are unidentified.)

Online appendix table IV presents our placebo models of orange juice and milk purchases for both the full sample of transactions and a subsample of transactions that occur on the same day as a gasoline purchase. Results are similar in the two samples.

Online appendix table V presents our placebo models of orange juice and milk purchases for both the full sample of transactions and a subsample of transactions that are made by frequent buyers, defined as households that buy at least 12 times in the respective category in each of the full years in our sample (2006, 2007, and 2008). In the case of orange juice, the point estimate on η^G changes meaningfully but remains well within the confidence interval of the original estimate, and the estimate on η^M becomes less precise due to the smaller number of households in the subsample. In the case of milk, point estimates are similar between the full sample and the subsample. In both cases our substantive conclusion is unchanged: we cannot reject the null hypothesis of fungibility in any specification.

II. ALLOWING FOR MEASUREMENT ERROR AND TRANSITORY INCOME SHOCKS

Let z_i be a household's true current income and let \hat{z}_i be the household's self-reported income. Let m_i be a household's true permanent income, which we assume is equivalent to its total expenditure on goods and services. Let \hat{m}_i be a household's reported total expenditures.

For households in our retailer panel we measure reported current income \hat{z}_i . For households in the Consumer Expenditure Survey we measure reported current income \hat{z}_i and reported expenditure \hat{m}_i . Suppose that the measurement errors in both m_i and z_i are classical in the sense that these errors are normally distributed independently of one another and of m_i , z_i , and the other exogenous variables in our model. Suppose further that m_i and z_i are jointly normally distributed with some nonzero covariance.¹

Then from a regression of \hat{m}_i on \hat{z}_i in the Consumer Expenditure Survey we obtain an estimate of the conditional expectation $E(m_i|\hat{z}_i)$. The marginal utility of money is then

$$\lambda_{it} = \tilde{\mu}_i - \eta^M E(m_i|\hat{z}_i) + \eta^G q_{it} p_{0t} \tag{1}$$

¹Implicitly, we also assume that the extent of measurement error (and transitory shocks) in household income does not differ between our retailer panel and the Consumer Expenditure Survey. The results in the paper's data discussion provide important evidence that this assumption is a reasonable approximation.

where $\tilde{\mu}_i = \mu_i - \eta^M(m_i - E(m_i|\hat{z}_i))$ is a function both of underlying structural heterogeneity in marginal utility μ_i and of deviations between the true permanent income m_i and the econometrician’s conditional expectation $E(m_i|\hat{z}_i)$. By construction, the deviation $(m_i - E(m_i|\hat{z}_i))$ is orthogonal to the expectation $E(m_i|\hat{z}_i)$.

The model in equation (1) is formally equivalent to one in which the econometrician perfectly measures permanent income ($E(m_i|\hat{z}_i) = m_i$) and the marginal utility parameter μ_i is normally distributed independently of \hat{z}_i . Specification (4) in online appendix table III corresponds to such a model. If anything, our results are stronger in this specification than in our baseline estimates.

III. PRODUCT AGGREGATION IN GROCERY CATEGORIES

Transaction data in grocery categories include a universal product code (UPC) for each item purchased. Our data include 91 orange juice UPCs and 129 milk UPCs purchased during the sample period. We exclude small-format and single-serving UPCs and those that sell fewer than 100 units over the sample period.

We aggregate the UPCs into products for estimation purposes in two steps.

First, we recode as identical any two UPCs that are substantially the same. In orange juice, this means recoding as identical any two UPCs that have the same brand, size, pulp content, and enrichment (e.g., calcium added). In milk, this means recoding as identical any two UPCs that have the same brand, size, and fat content.² UPCs that are recoded in this way typically do not differ on any attributes recorded by the retailer; in many cases differences will reflect technical changes in the retailer’s database or small differences in packaging.

Second, we aggregate UPCs into products by size and brand. For UPCs in the same size, brand, and store, the Pearson correlation in the change in log prices is 0.89 for orange juice and 0.88 for milk. Following the composite commodity theorem (Deaton and Muellbauer 1980), the high correlation in prices within brand and size gives some support to the level of aggregation we choose in our analysis.

We construct a price series for each product as follows. We rescale prices to be in units of dollars per half gallon. For each product and each store, we construct the overall market share of each component UPC and use these market shares to construct a fixed-weight average price in each calendar week. We use the universe of store purchases (rather than only those purchases made by sample households) to construct the product price series.

In orange juice, there are 19 products (brand-size groups). The most popular is half-gallon Tropicana,

²Note that in the case of milk, common brands for conventional milk vary regionally. For estimation purposes we treat all non-private-label conventional milk as a single brand. Because most of the variation in conventional milk brands is across stores rather than within stores, this aggregation should not greatly distort our picture of the consumer’s choice set.

which has a market share of 22 percent. Prices range from \$2.30 per half gallon for one-gallon private label juice to \$4.06 per half gallon for 89 oz format Simply Orange organic juice. (A few smaller formats have even higher prices when measured in price per half gallon.)

In milk, there are 9 products. The most popular is the one-gallon private label with a market share of 70 percent. The one-gallon private label is also the least expensive option, costing \$1.48 per half gallon as against \$3.66 per half gallon for the half-gallon format Horizon organic milk.

IV. EFFECTS OF VEHICLE SUBSTITUTION ON OCTANE CHOICE

In this section we discuss the extent to which changes in which vehicles are driven could explain the time-series variation in octane choice that we observe.

There are three main channels through which vehicle substitution might occur. The first is a relative increase in the market share of fuel-efficient vehicles among new purchases. In an average week in 2006, new car sales represented about one-tenth of one percent of the automobile stock (United States Census 2009). Applying Busse, Knittel, and Zettelmeyer's (2010) estimate of the change in market share by quartile of fuel economy to the estimated fraction of vehicles in each quartile that recommend regular gasoline, we estimate that a \$1 increase in the price of gasoline increases the share of the vehicle stock recommending regular gasoline by less than one-hundredth of one percentage point over one week.

The second channel is disproportionate scrappage of less-fuel-efficient vehicles. As with new car purchases, the share of the vehicle stock scrapped in any given week is too small to allow for a significant change in the stock of vehicles on the road. In addition, Knittel and Sandler (2011) find that vehicle age is a more important determinant of scrappage rates than fuel economy per se.

The third channel is changes in the intensity of driving of different types of vehicles, which is likely to be especially important for households with multiple vehicles. Knittel and Sandler (2011) find that, at annual horizons, less fuel efficient cars are driven less than fuel efficient cars when gasoline prices rise, at least partly due to within-household substitution in vehicles driven. Adjusting their estimates to apply to short-run changes by matching to the short-run elasticity estimates in Hughes, Knittel, and Sperling (2008), we estimate that a \$1 increase in the price of gasoline increases the (mileage-weighted) share of vehicles recommending regular gasoline by two-hundredths of one percentage point.

Taking these three channels together, we estimate that due to vehicle substitution alone a \$1 increase in the price of gasoline would increase the share of regular gasoline by 0.03 percentage points in the short run,

far below the empirical effects that we report in the body of the paper.³

V. MODEL OF LEARNING ABOUT GASOLINE QUALITY

In this section we outline a stylized model of learning about the quality of gasoline. We argue that the dynamics implied by this model differ in important ways from those that we observe in our data and therefore that learning is unlikely to account for our findings.

Consider the special case of our model with only two grades of gasoline, no income effects ($\eta = 0$), and a constant price gap between grades. Normalize the utility from the premium grade to 0 and let the utility from regular gasoline for household i at time t be given by

$$u_{it} = \alpha \omega + \varepsilon_{it} \quad (2)$$

where α is a parameter and $\omega \in [0, 1]$ is an unknown state of the world indicating the impact of using regular gasoline on vehicle performance.

The household begins with a uniform prior over ω . At each time t , each household obtains N_t draws from a binomial distribution with success probability equal to ω and updates its posterior. The household evaluates the expectation of u_{it} given its current and past draws $\{N_k\}_{k=0}^t$. Suppose that $\omega = 1$ so that all draws are successes. Then it follows from Bayes' Rule that the household's expected utility from regular gasoline is

$$\begin{aligned} E(u_{it} | \{N_k\}_{k=0}^t) &= \alpha E(\omega | \{N_k\}_{k=0}^t) + \varepsilon_{it} \\ &= \alpha \left(\frac{1 + \sum_{k=0}^t N_k}{2 + \sum_{k=0}^t N_k} \right) + \varepsilon_{it} \end{aligned} \quad (3)$$

The household's expectation of ω approaches the true value of 1 as more data arrives.

Variation in the gasoline price may affect the arrival rate of information. When gasoline prices are high, gasoline prices are more likely to be in the news and households may engage in search effort to learn more about fuel-saving tips, which may in turn lead them to acquire information about the relative merits of different gasoline grades. We therefore allow that the arrival of information is a function of the

³To aggregate the three channels together, we assume that the second (scrappage) channel is equal in magnitude to the first (market share of new vehicles). Because Knittel and Sandler (2011) find that age is more important than fuel economy in determining the response of scrappage to gasoline prices, we regard this approach as conservative. Note that Levin, Lewis, and Wolak (2012) find a much larger short-run elasticity of demand for gasoline than do Hughes, Knittel, and Sperling (2008). Using Levin, Lewis, and Wolak's (2012) largest estimate of 0.48, we estimate that the aggregate effect of the three channels of vehicle substitution is 0.16 percentage points, still well below the effect of 1.42 percentage points that we report in our baseline models.

prevailing price of regular gasoline p_{0t} , formally that $N_t = N(p_{0t})$ for some function $N(\cdot)$. (Here we treat N_t as continuous even though, strictly speaking, it is a natural number.)

We specify three possible functional forms for $N(\cdot)$. The first assumes that $N(\cdot)$ is a positive constant so that there is a steady arrival of information over time. The second assumes that $N(\cdot)$ is proportional to p_{0k} with a positive constant of proportionality. The third assumes that $N(\cdot)$ is equal to a positive constant when $p_{0k} > \$3$ and 0 otherwise, i.e. information arrives when gas prices exceed a threshold. We estimate α and the parameter of the function $N(\cdot)$ in each case using nonlinear least squares, using weekly data on the price of regular gasoline and the share of transactions going to regular grade. We treat time $t = 0$ as the beginning of our sample period.

Online appendix figure VII shows the predictions of each model. None is a good match to the observed series. The model with a constant arrival rate of information essentially predicts a linear trend in the share of regular gasoline. The model in which the arrival of information is proportional to the gasoline price predicts a trend that accelerates when the price is high. The model in which information arrives only when the price of regular gasoline exceeds \$3/gallon predicts steady learning during high-price periods and no learning otherwise. None of these models predicts the empirical fact that the regular share rises and falls with the price of regular gasoline.

The limitation shared by all of these models is that, since learning is cumulative, while they can predict that the share of regular rises—or rises more quickly—when gasoline prices rise, they fail to predict the fact that the share of regular falls when gasoline prices fall. Once households have learned that regular gasoline is of good quality, there is no reason for them to “unlearn” this when the price falls again. This makes it difficult for a model of learning to serve as a complete explanation for the patterns we observe.

VI. ESTIMATES FROM STATE-LEVEL DATA

For each state and quarter we obtain state aggregate personal income from the Bureau of Economic Analysis (BEA), which we convert into aggregate expenditures by scaling by the US ratio of quarterly consumer expenditures to quarterly personal income from the National Income and Product Accounts (NIPA). We obtain state aggregate gasoline gallons purchased and the average retail price of regular gasoline from the Energy Information Administration (EIA). We estimate the number of households in each state and quarter by multiplying BEA midyear population estimates by the US ratio of households to population (from the Census). We transform all incomes, expenditures and prices into 2005 dollars using the NIPA price index for personal consumption expenditures excluding food and energy.

To estimate our model we assume that expenditures, gasoline quantities, and preferences are identical (up to taste shocks ε_{ijt}) across households within a state. We measure m_{it} as total expenditures divided by the number of households. We measure $q_{it}p_{0t}$ as the product of state gasoline gallons purchased and state average retail price of regular gasoline, divided by the number of households. We transform the market share of each grade of gasoline into the mean utility of that grade in a given state and quarter, relative to the share of regular gasoline (Berry 1994). We estimate the model via two-stage least squares, treating state per-capita income and the national price of gasoline as excluded instruments and treating the key interactions between total and gasoline expenditures and the price gap as endogenous. (This approach ensures that our key parameters are identified only by variation in income and in national prices.) We allow for a state-quarter-grade utility shock that is mean zero conditional on instruments and controls.

Online appendix table VI presents our results. The format is identical to the presentation of our main results in the paper. Column (1) presents our baseline results and column (2) presents a specification including state fixed effects. In both specifications we confidently reject the null hypothesis of fungibility.

VII. IMPLICATIONS FOR RETAIL MARKETS

Existing evidence suggests that the retail markup on gasoline tends to fall when the oil price rises (Peltzman 2000; Chesnes 2010; Lewis 2011). To illustrate, Panel A of figure VIII reproduces figure 1 of Lewis (2011), which shows the pre-tax retail price and wholesale (spot) price of regular reformulated gasoline in Los Angeles in 2003 and 2004 as measured by the EIA. When the spot price rises, the markup—the gap between the wholesale and retail prices—compresses.

Lewis (2011) provides a search-based account of this effect. When prices rise, consumers cannot tell how much of the increase is retailer-specific, so they increase the intensity with which they search for better prices at other retailers, thus putting downward pressure on retailer margins.

Our findings offer a complementary explanation. We show that when prices rise, consumers act as if they have a high marginal utility of money in the gasoline domain. If this force operates when consumers decide which retailer to purchase from, it will result in greater price sensitivity and hence lower retail markups.

To illustrate, consider the following toy model of retail pricing. The market consists of a large number of identical retailers selling regular grade gasoline to a unit mass of households. (Formally, we consider the limit case as the number of retailers grows large.) Each household's utility is quasilinear in money with marginal utility ρ_t and is subject to an additive type-I extreme value error i.i.d. across households and retailers. Retailers set prices simultaneously, taking the marginal utility ρ_t as given, and face a common and

exogenous wholesale price c_t . Then in the unique equilibrium (Anderson, de Palma, and Thisse 1992) all retailers charge the same price p_{0t}^* :

$$p_{0t}^* = c_t + \frac{1}{\rho_t}. \quad (4)$$

Given this model of equilibrium pricing, we can ask how much of the tendency of markups to fall when wholesale prices rise can be explained by the variation in marginal utility that we estimate in our model of grade choice. We assume that

$$\rho_t = \bar{\rho} (\mu - \eta^M m + \eta^G q p_{0t}^*) \quad (5)$$

where m is average annual household expenditure, q is average annual gallons of gasoline per household, and $\{\bar{\rho}, \mu, \eta^G, \eta^M\}$ are preference parameters. Equations (4) and (5) uniquely define p_{0t}^* as a function of parameters. We estimate $\bar{\rho}$ via nonlinear least squares using the EIA data shown in figure VIII, matching the predicted markup to the observed series. We take the other preference parameters from our baseline model of grade choice.

Panel B of figure VIII shows the observed markup, the markup predicted using preference parameters from our baseline model, and the markup predicted from a model of grade choice in which we constrain $\eta^G = \eta^M$. The markup tends to increase when the spot price decreases. Using the preference parameters from our baseline model, our model of retail pricing explains some, though not all, of this pattern. By contrast, using preference parameters that impose fungibility ($\eta^G = \eta^M$), the retail pricing model predicts essentially no variation in the markup.

VIII. ADDITIONAL EVIDENCE ON PSYCHOLOGICAL MODELS

Online appendix table VII presents, for each of the psychological models that we estimate in the body of the paper, an estimate of the model's key parameter and its standard error. The table also presents two measures of goodness of fit. The first measure is the mean of the log likelihood of the model. The second measure is the gain in log likelihood for the model relative to a baseline model with no psychological factors (i.e. with $\Gamma_{ijt} = 0$), scaled by the gain in log likelihood from the category budgeting model.

Online appendix figure IX shows that the salience model fits better when we assume a one-week horizon than when we assume a four-week horizon for the evoked set. The reason is that prices will be more salient than octane when the prices in the evoked set vary more, in percentage terms, than the octane levels. With octane levels of 87, 89, and 91, an absolute deviation of a few percent between past and present prices is sufficient to make prices more salient than octanes. Such deviations are the norm over four-week horizons,

so there is little time-series variation in the indicator z_{ijt} when a four-week horizon is used.

Online appendix figure X shows that the loss aversion model fits better when we assume a four-week horizon for expectation formation than when we assume a one-week horizon. The reason is that if expectations adapt quickly then persistently high prices only impact octane choice for a short period. The figure also shows the fit of the loss aversion model when we assume that the expected price paid and octane purchased are equal to their sample averages for all households and time periods.

Online appendix figure XI shows two additional variants on the loss aversion model estimated in the body of the paper. Panel A presents a variant allowing for diminishing sensitivity. Formally, we assume that the universal gain-loss function is a power function (Tversky and Kahneman 1992):

$$\gamma(x) = \gamma_+ x^\nu \mathbf{1}_{x \geq 0} + \gamma_- |x|^\nu \mathbf{1}_{x < 0} \quad (6)$$

where γ_+ , γ_- , and ν are parameters. We set $\nu = 0.88$ (Tversky and Kahneman 1992) and estimate the remaining separately identified parameters.

Panel B of online appendix figure XI allows for a stochastic referent. Formally, we assume that the price paid for gasoline is lognormally distributed with a mean that is a linear function of the national gasoline price four weeks prior to purchase. As most of the variation in price paid is due to variation in the oil price and oil spot prices are approximately stationary lognormal (Gibson and Schwartz 1989), we expect this distributional assumption to be a reasonable approximation. Following Köszegi and Rabin (2006), we assume that households evaluate non-gasoline gain-loss utility by integrating over the (lognormal) cdf of the referent. We estimate the model in two steps. In the first step, we estimate the parameters of the lognormal distribution of prices paid. In the second step, we estimate preference parameters assuming household expectations follow the estimated lognormal distribution.⁴

BROWN UNIVERSITY AND NBER

CHICAGO BOOTH AND NBER

REFERENCES (NOT CITED IN PAPER)

Anderson, Simon P., André de Palma, and Jacques-François Thisse, *Discrete Choice Theory of Product Differentiation*, (Cambridge, MA: MIT Press, 1992).

⁴Note that we continue to assume a point-mass referent for octane levels. Allowing for a stochastic referent in octane consumption using a two-step procedure would mean allowing preference to be a function of the observed relationship between prices and grade choice, which would lead to a good model fit for purely mechanical reasons.

Busse, Meghan, Christopher R. Knittel, and Florian Zettelmeyer, "Pain at the Pump: The Differential Effect of Gasoline Prices on New and Used Automobile Markets," NBER Working Paper No. 15590, 2010.

Chesnes, Matthew, "Asymmetric Pass-Through in US Gasoline Prices," US Federal Trade Commission Bureau of Economics Working Paper No. 302, 2010.

Deaton, Angus, and John Muellbauer, *Economics and Consumer Behavior*, (Cambridge, UK: Cambridge University Press, 1980).

Gibson, Rajna, and Eduardo S. Schwartz, "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, 45, no. 3 (1989), 959-976.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling, "Evidence of a Shift in the Short-Run Elasticity of Gasoline Demand," *Energy Journal*, 29, no. 1 (2008), 93-114.

Knittel, Christopher, and Ryan Sandler, "Cleaning the Bathwater with the Baby: The Health Co-Benefits of Carbon Pricing in Transportation," NBER Working Paper No. 17390, 2011.

Levin, Laurence, Matthew S. Lewis, and Frank A. Wolak, "High Frequency Evidence on the Demand for Gasoline," Ohio State University Mimeograph, 2012.

National Highway Traffic Safety Administration, *National Automotive Sampling System General Estimates System*, accessed at <<http://www.nhtsa.gov/NASS>> on July 13 2011, posted in 2009.

Peltzman, Sam, "Prices Rise Faster Than They Fall," *Journal of Political Economy*, 108, no. 3 (2000), 466-502.

Tversky, Amos, and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5 (1992), 297-323.

United States Census, *Statistical Abstract of the United States 2009*, accessed at <<http://www.census.gov/compendia/statab/2009/>> on September 17, 2011, posted in 2009.

ONLINE APPENDIX TABLE I: FUNCTIONAL FORM OF NON-GASOLINE UTILITY

Dependent variable: Choice of gasoline grade

	Excess effect of gas money ($\eta^G - \eta^M$)	P-value of test for fungibility ($\eta^G = \eta^M$)
Specification of $\Lambda(\cdot)$:		
(1) Baseline (linear approximation)	-0.4013 (0.0312)	0.0000
(2) Quadratic	-0.4009 (0.0311)	0.0000
(3) Cubic	-0.4054 (0.0313)	0.0000
(4) Quartic	-0.4060 (0.0313)	0.0000

Note: Data are from retailer. Table reports estimates of the model described in the body of the paper with alternative specifications of the indirect utility function $\Lambda_i(\cdot)$ for non-gasoline expenditures. Standard errors in parentheses allow for correlation in residuals by month. In all cases we estimate the model by including both $\Lambda_i(\cdot)$ and a separate interaction of total gasoline expenditure $p_{0t}q_{it}$ with the grade price p_{jt} . The coefficient we report is the coefficient from the interaction term, which is zero when gas money and other money are treated as fungible. Specification (1) is our baseline specification; the coefficient reported corresponds to $(\eta^G - \eta^M)$ in our main model estimates. Specifications (2), (3), and (4) assume that $\Lambda_i(\cdot)$ is quadratic, cubic, and quartic, respectively. In those specifications, the coefficient reported is analogous to $(\eta^G - \eta^M)$ in that it measures how much the marginal utility of money responds to the gasoline price above what is predicted by the extent of diminishing utility in $\Lambda_i(\cdot)$. We assume that α_{ij} and the function $\Lambda_i(\cdot)$ are constant across households.

ONLINE APPENDIX TABLE II: LINEAR PROBABILITY MODELS OF GASOLINE GRADE CHOICE

Dependent variable: Household chooses regular gasoline

	(1)	(2)	(3)	(4)	(5)	(6)
Marginal effect on probability of buying regular:						
\$1000 decrease in gasoline expenditures	-0.0129 (0.0011)	-0.0130 (0.0011)	-0.0123 (0.0014)	-0.0116 (0.0013)	-0.0108 (0.0008)	-0.0092 (0.0009)
\$1000 increase in total expenditures	-0.0006 (0.0000)	-0.0003 (0.0000)	-0.0001 (0.0000)	-0.0001 (0.0000)	-0.0001 (0.0000)	-0.0001 (0.0000)
p-value of test for null that gasoline and total expenditures have equal effects	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Time-varying expenditure measure?		X	X	X	X	X
Household fixed effects?			X	X	X	X
Controls for price gaps between grades?				X	X	X
Linear time trend?					X	X
Dummies for calendar month (seasonality)?						X
Number of transactions	10548175	10548175	10548175	10548175	10548175	10548175
Number of households	61494	61494	61494	61494	61494	61494

Note: Data are from retailer. Table reports estimates of a linear probability model whose dependent measure is an indicator for purchasing regular gasoline. Standard errors in parentheses allow for correlation in residuals by month. In all specifications we measure gasoline expenditures with our empirical proxy for $q_{it}p_{0t}$ as defined in the paper. In specification (1) we measure total expenditures with our time-constant measure m_i . In specifications (2) through (6) we measure total expenditures with our time-varying measure m_{it} . In specification (2) we control for m_i . In specification (3) we add household fixed effects. In specification (4) we add controls for the difference in price between premium and regular and between midgrade and regular gasoline. In specification (5) we add a linear time trend. In specification (6) we add dummies for (seasonal) calendar month, i.e. 11 dummies total excluding a baseline month.

ONLINE APPENDIX TABLE III: ADDITIONAL MODELS WITH UNOBSERVED HETEROGENEITY

Dependent variable: Choice of gasoline grade					
	(1)	(2)	(3)	(4)	(5)
Effect on marginal utility of:					
\$1000 increase in gasoline expenditures (Parameter η^G)	0.4306 (0.0314)	0.7145 (0.0317)	0.7434 (0.0188)	0.7721 (0.0340)	1.0477 (0.0133)
\$1000 decrease in total expenditures (Parameter η^M)	0.0293 (0.0008)	0.0416 (0.0042)	0.0250 (0.0036)	0.0204 (0.0087)	0.0159 (0.0040)
Average marginal effect on regular share of:					
\$1 increase in price of regular gasoline	0.0142 (0.0010)	0.0140 (0.0006)	0.0335 (0.0008)	0.0217 (0.0010)	0.0349 (0.0004)
\$1000 decrease in gasoline expenditures	-0.0120 (0.0009)	-0.0118 (0.0005)	-0.0283 (0.0007)	-0.0184 (0.0008)	-0.0295 (0.0004)
\$1000 increase in total expenditures	-0.0008 (0.0000)	-0.0007 (0.0001)	-0.0010 (0.0001)	-0.0005 (0.0002)	-0.0004 (0.0001)
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.0000	0.0000	0.0000	0.0000	0.0000
Unobservable variation in...		α_{ij}	α_{ij}	μ_i	α_{ij}, μ_i
Quadrature accuracy	—	9	3	9	—
Sample	All	1/10th	All	1/10th	Buy all grades
Number of transactions	10548175	1082486	10548175	1082486	4556308
Number of households	61494	61494	61494	61494	24643

Note: Data are from retailer. Table reports estimates of the model described in the body of the paper. Standard errors in parentheses allow for correlation in residuals by month. Sample “1/10th” is a sample of every 10th transaction for each household. Sample “Buy all grades” restricts to households who buy each grade of gasoline at least once during the sample period. In specification (1) we assume that α_{ij} and μ_i are constant across households and estimate via maximum likelihood. In specifications (2) and (3) we assume that α_{ij} are distributed independently normal across households and choices. In specification (4) we assume that μ_i are distributed independently normal across households and choices. We estimate the models in specifications (2), (3) and (4) via maximum likelihood, approximating the likelihood using sparse grid integration with given accuracy (Heiss and Winschel 2008). In specification (5) we estimate the model separately for each household and estimate parameters η^M and η^G via FGLS from the household-specific parameter estimates, under the assumption that μ_i is uncorrelated with m_{it} conditional on α_{ij} . The Wald test for specification (5) assumes no covariance in the estimates of η^G and η^M .

ONLINE APPENDIX TABLE IV: PLACEBO MODELS: GAS DAYS VS. ALL DAYS

Dependent variable: Choice of brand				
	(1)	(2)	(3)	(4)
Effect on marginal utility of:				
\$1000 increase in gasoline expenditures (Parameter η^G)	-0.0141 (0.0250)	-0.0039 (0.0275)	-0.0128 (0.0197)	-0.0117 (0.0225)
\$1000 decrease in total expenditures (Parameter η^M)	0.0044 (0.0002)	0.0044 (0.0003)	0.0034 (0.0001)	0.0036 (0.0001)
Average marginal effect on private label share of:				
\$1 increase in price of regular gasoline	-0.0169 (0.0299)	-0.0047 (0.0327)	-0.0055 (0.0084)	-0.0050 (0.0096)
\$1000 decrease in gasoline expenditures	0.0143 (0.0252)	0.0039 (0.0276)	0.0046 (0.0071)	0.0042 (0.0081)
\$1000 increase in total expenditures	-0.0045 (0.0002)	-0.0044 (0.0003)	-0.0012 (0.0000)	-0.0013 (0.0000)
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.4571	0.7617	0.4115	0.4977
Category	OJ	OJ	Milk	Milk
Sample	All	Gas day	All	Gas day
Number of transactions	411161	97684	2210312	514449
Number of households	13493	12760	34128	33756

Note: Data are from retailer. Table reports estimates of the model described in the paper but applied to choice of orange juice or milk brand rather than choice of gasoline grade. See online appendix section III 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. “Gas day” means that the purchase was made on the same day as a gasoline purchase.

ONLINE APPENDIX TABLE V: PLACEBO MODELS: RESTRICTING TO FREQUENT BUYERS

Dependent variable: Choice of brand				
	(1)	(2)	(3)	(4)
Effect on marginal utility of:				
\$1000 increase in gasoline expenditures (Parameter η^G)	-0.0141 (0.0250)	0.0082 (0.0219)	-0.0128 (0.0197)	-0.0191 (0.0192)
\$1000 decrease in total expenditures (Parameter η^M)	0.0044 (0.0002)	0.0046 (0.0003)	0.0034 (0.0001)	0.0033 (0.0001)
Average marginal effect on private label share of:				
\$1 increase in price of regular gasoline	-0.0169 (0.0299)	0.0100 (0.0266)	-0.0055 (0.0084)	-0.0076 (0.0076)
\$1000 decrease in gasoline expenditures	0.0143 (0.0252)	-0.0084 (0.0225)	0.0046 (0.0071)	0.0064 (0.0064)
\$1000 increase in total expenditures	-0.0045 (0.0002)	-0.0047 (0.0003)	-0.0012 (0.0000)	-0.0011 (0.0000)
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.4571	0.8676	0.4115	0.2443
Category	OJ	OJ	Milk	Milk
Sample	All	Frequent buyer	All	Frequent buyer
Number of transactions	411161	164742	2210312	1659129
Number of households	13493	2287	34128	17044

Note: Data are from retailer. Table reports estimates of the model described in the paper but applied to choice of orange juice or milk brand rather than choice of gasoline grade. See online appendix section III 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. “Frequent buyer” means that the household made a purchase in the category at least 12 times in each of 2006, 2007, and 2008.

ONLINE APPENDIX TABLE VI: ESTIMATES FROM STATE-LEVEL DATA

Dependent variable: $\log(\text{share of grade})$ minus $\log(\text{share of regular})$		
	(1)	(2)
Effect on marginal utility of:		
\$1000 increase in gasoline expenditures (Parameter η^G)	0.7877 (0.0842)	0.8327 (0.1679)
\$1000 decrease in total expenditures (Parameter η^M)	0.0719 (0.0190)	0.0608 (0.0707)
Average marginal effect on regular share of:		
\$1 increase in price of regular gasoline	0.0293 (0.0031)	0.0309 (0.0062)
\$1000 decrease in gasoline expenditures	-0.0232 (0.0025)	-0.0245 (0.0049)
\$1000 increase in total expenditures	-0.0021 (0.0006)	-0.0018 (0.0021)
p-value of Wald test for fungibility ($\eta^G = \eta^M$)	0.0000	0.0000
State fixed effects?		X
Number of state-quarter-grades	7377	7377
Number of states	50	50


Note: See section VI of the online appendix for details of data construction. All specifications include decade fixed effects, quarter (season) fixed effects, and decade-specific linear time trends, all interacted with gasoline grade fixed effects. Specification (2) adds state fixed effects interacted with gasoline grade fixed effects. All models are estimated via two-stage least squares, with state per-capita income and the national average price of regular gasoline treated as excluded instruments and the interaction between total expenditures and the grade price gap, and between gasoline expenditures and the grade price gap, treated as endogenous.

ONLINE APPENDIX TABLE VII: ESTIMATES OF KEY PARAMETERS OF PSYCHOLOGICAL MODELS

Model	Key parameter	Interpretation	Estimate (Standard error)	Mean log likelihood (Relative pseudo- R^2)
Baseline		—	—	-0.6407 (0.0000)
Category budgeting	γ	Utility weight on deviation from mean expenditure	0.0240 (0.0002)	-0.6397 (1.0000)
Loss aversion (weekly horizon)	$1 + \tilde{\gamma}$	Ratio of sensitivity to price in losses to sensitivity to price in gains	10.8460 (2.5533)	-0.6403 (0.3465)
Saliency (weekly horizon)	$\frac{(1+\kappa(1))}{(1+\kappa(0))}$	Ratio of sensitivity to price when price is salient to sensitivity to price when octane is salient	1.1702 (0.0096)	-0.6404 (0.2617)

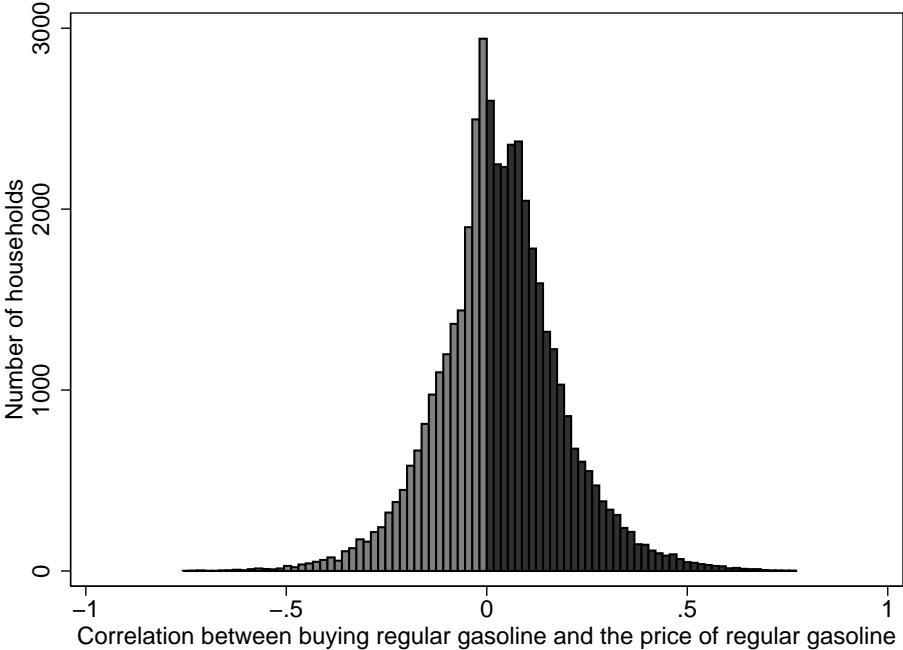
Note: See body of paper for details on estimating equations. Each row shows the given parameter and its standard error, which allows for correlation in residuals by month. Each row also shows the model's mean log likelihood and the "Relative pseudo- R^2 ," which is the difference between the model's log likelihood and that of the baseline model, divided by the difference between the category budgeting model's log likelihood and that of the baseline model. The baseline model corresponds to a case with $\Gamma_{ijt} = 0$, i.e. no psychological factors.

ONLINE APPENDIX FIGURE I: Persistence of the effect of the regular price



Notes: Data are from the EIA. The plot shows the cumulative effect of a \$1 increase in the regular price on the share of regular gasoline. The bars show a 95 percent confidence interval. The plot is based on a regression of the change in the regular share on the change in the regular price and six lags of the change in the regular price. The regression uses US-level monthly data from 1990-2009. All prices are converted to 2005 US dollars using the NIPA price index for personal consumption expenditures excluding food and energy.


ONLINE APPENDIX FIGURE II: Correlation between buying regular gasoline and the regular price




Notes: Data are from the retailer. For each household, we compute the Pearson correlation between an indicator for buying regular gasoline and the national price of regular gasoline in the transaction week. The plot shows the distribution of the correlation across households. The darkly shaded region corresponds to households for which the estimated correlation is nonnegative. The plot excludes households who always buy regular gasoline and households who never buy regular gasoline.

ONLINE APPENDIX FIGURE III: Price gap and price of regular gasoline

Panel A: 1990-1999




Panel B: 2000-2009



Notes: Data are from the EIA. Each panel plots the monthly US average price of regular gasoline (in 2005 US dollars) and the monthly average price difference between premium and regular gasoline (in 2005 US dollars). Prices are converted to 2005 dollars using the NIPA price index for personal consumption expenditures excluding food and energy.


ONLINE APPENDIX FIGURE IV: Decay in the effect of regular price on the price gap




Notes: Data are from the EIA. The plot shows the cumulative effect of a \$1 increase in the regular price on the price gap between premium and regular gasoline. The bars show a 95 percent confidence interval. The plot is based on a regression of the change in the price gap between premium and regular gasoline on the change in the regular price and six lags of the change in the regular price. The regression uses US-level monthly data from 1990-2009. All prices are converted to 2005 dollars using the NIPA price index for personal consumption expenditures excluding food and energy.

ONLINE APPENDIX FIGURE V: Income implied by purchasing behavior: Gasoline vs. placebo categories


Panel A: Gasoline grade choice



Panel B: Orange juice brand/size choice




Panel C: Milk brand/size choice



Notes: Data are from the retailer and exclude stores with significant changes to the milk assortment during the sample period. We estimate a logit model allowing for a product-specific constant, a product-specific linear time trend (for orange juice and milk), and a store-week-product-level utility shock that is mean zero conditional on the included variables. We assume that the marginal utility of money is a linear function of income and includes a week-level shock common to all households. We define the “income implied by purchasing behavior” in a given week as the household income such that the marginal utility of money in a baseline week is equal to the marginal utility of money in the given week. (We choose the baseline week so that mean implied income is equal to the sample mean income.) To estimate the model, we aggregate the data to the store-week level and assume household income is equal to mean household income in the store-week. We transform the market share of each product (grade of gasoline, brand/size of orange juice or milk) into the mean utility of that grade in a given store-week, relative to the share of a base product (regular gasoline, private label one-gallon orange juice or milk) and estimate via OLS (Berry 1994).

ONLINE APPENDIX FIGURE VI: Traveling speed of automobiles in collisions

Panel A: 1990-1999




Panel B: 2000-2009



Notes: Each panel plots the average traveling speed for automobiles in the US that are struck during a motor vehicle accident (from the NASS-GES) and the monthly US average price of regular gasoline (from the EIA, in 2005 US dollars). Prices are converted to 2005 dollars using the NIPA price index for personal consumption expenditures excluding food and energy.


ONLINE APPENDIX FIGURE VII: Predictions of a learning model



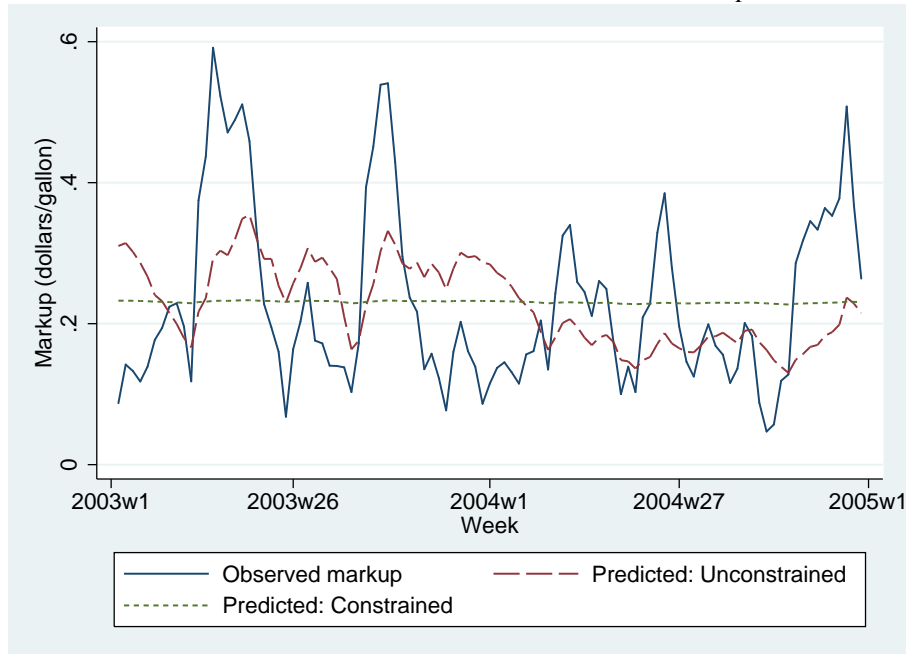
Notes: Data are from the retailer. The line labeled “price of regular” shows the mean price of regular gasoline by week. The remaining lines show the predictions of the models of learning about the quality of regular gasoline specified in section V of the online appendix. Each model assumes a different process governing the arrival of information. The line labeled “constant learning” shows the predictions of a model with a constant arrival rate of information. The line labeled “learning when price>3” shows the predictions of a model in which learning takes place at a constant rate whenever the price of regular gasoline exceeds \$3/gallon. The line labeled “learning proportional to price” shows the predictions of a model in which learning takes place at a rate proportional to the price of gasoline.

ONLINE APPENDIX FIGURE VIII: Implications for retail prices

Panel A: Wholesale and retail prices




Panel B: Predicted and actual variation in markup



Notes: Data are from the EIA. Panel A shows the weekly average pre-tax retail price of regular reformulated gasoline in Los Angeles and the previous week's average daily spot price of reformulated gasoline. Panel B shows the observed markup (equal to the difference between the pre-tax retail price and the spot price) and the markup predicted by the model of retailer behavior described in section VII under two different assumptions about the determinants of household marginal utility of money. The line labeled "predicted: unconstrained" uses the estimates of the marginal utility function from our baseline model as presented in the paper. The line labeled "predicted: constrained" uses the estimates of the marginal utility function from same model, re-estimated imposing the constraint that $\eta^G = \eta^M$.


ONLINE APPENDIX FIGURE IX: Saliency model: Alternative horizons




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: no salience” shows the average predicted probability of buying regular gasoline from an estimate of the model in section 7 of the paper with $\Gamma_{ijt} = 0 \forall i, j, t$. The line labeled “predicted: salience” shows the average predicted probability of buying regular gasoline from an estimate of the model in equation section 7 of the paper with Γ_{ijt} specified according to the salience model. The evoked set is assumed to consist of both the current choice set and the set of all gasoline grades at their national prices one week or four weeks prior to purchase. See text for additional details.

ONLINE APPENDIX FIGURE X: Loss aversion model: Alternative horizons


Panel A: One-week horizon



Panel B: Four-week horizon




Panel C: Static expectations



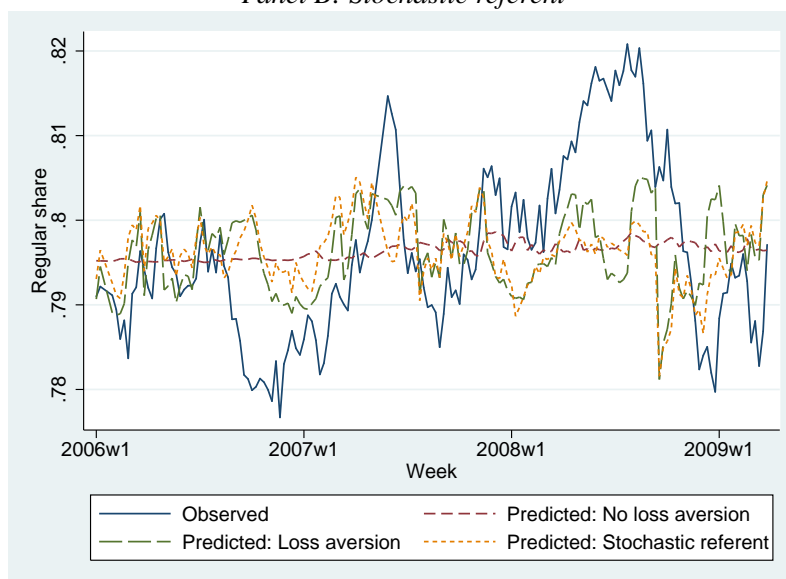
Notes: Data are from the retailer. The line labeled “observed” shows the weekly share of transactions that go to regular gasoline. The line labeled “predicted: no salience” shows the average predicted probability of buying regular gasoline from an estimate of the model in section 7 of the paper with $\Gamma_{ijt} = 0 \forall i, j, t$. The line labeled “predicted: loss aversion” shows the average predicted probability of buying regular gasoline from an estimate of the model in section 7 of the paper with Γ_{ijt} specified according to the loss aversion model. In panels A and B, household expectations are assumed to be based on national gasoline prices [one week/four weeks] prior to purchase. In panel C, household expectations of price and octane are given by the variable’s sample mean. See text for additional details.

ONLINE APPENDIX FIGURE XI: Loss aversion: Model variants

Panel A: Diminishing sensitivity



Panel B: Stochastic referent



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: no loss aversion” shows the average predicted probability of buying regular gasoline from an estimate of the model in section 7 of the paper with $\Gamma_{ijt} = 0 \forall i, j, t$. The line labeled “predicted: loss aversion” shows the average predicted probability of buying regular gasoline from an estimate of the model in section 7 of the paper with Γ_{ijt} specified according to the loss aversion model with household expectations formed on a one-week horizon. In panel A, the line labeled “predicted: diminishing sensitivity” shows the average predicted probability of buying regular gasoline from the model in equation (6) in the online appendix with $\nu = 0.88$ and a one-week horizon. In panel B, the line labeled “predicted: stochastic referent” shows the average predicted probability of buying regular gasoline from a model in which we allow that the reference gasoline price is distributed lognormally with mean linear in the national gasoline price one weeks prior to purchase.