Flexible fuel vehicles, less flexible minded consumers: Price salience experiments at the pump

Alberto Salvo

Department of Economics, National University of Singapore

December 2016*

Abstract

This paper addresses a puzzle: why do many of Brazil's energy consumers driving "flexible fuel" gasoline-ethanol vehicles forgo energy savings at the pump, choosing the fuel that yields the *lower* mileage per dollar of spending, in some cases by a substantial margin? In a large-scale set of randomized experiments with 10,400 consumers at the pump—the first of its kind—I raise the salience of the price difference across both fuels, just as the consumer pulls up. The largest treatment effect I obtain is to shift one-tenth of consumers, who absent the intervention would have chosen expensive gasoline, to instead choose very favorably priced ethanol. While statistically significant, this shift is small compared with the higher likelihood that the favorably priced fuel is chosen among college-educated subjects relative to their less schooled counterparts. I estimate the increase in consumer welfare from mandating higher price salience at the pump to be equivalent to a 1 to 3% (general) reduction in fuel prices, depending on the relative price point.

JEL classification: D12, D64, L62, L71, Q21, Q41, Q42, Q48, R41

Keywords: Product differentiation, limited information, limited attention, price salience, information disclosure, gasoline, alternative fuels, biofuels, energy efficiency gap, flexible fuel vehicles, discrete-choice models

*I gratefully acknowledge João Alberto Abreu, Marcos Lutz, Carlos Resende and, in particular, Fábio Henkes, from Raízen (Shell-Cosan), as well as their numerous sales executives, for arranging access to fueling stations operating under the firm's brands, and their customers. I acknowledge comments from Hunt Allcott, Alcir Bandeira, Angela Lee, Teresa Meirelles, Luiz Perez, Florian Zettelmeyer, and audiences at the Energy Camp at UC Berkeley, Frontiers of Behavioral Economics (USC/NUS), and NUS. I am grateful to David Austen-Smith at the Dean's Office, Kellogg School of Management for partly supporting the field experiments. I thank Elizabeth Lehman, Fernando Luco and Guillermo Marshall for excellent research assistance. Northwestern University's Institutional Review Board determined that this study qualified for Exemption under 45 CFR 46.101(b). The usual disclaimer applies.

1 Introduction

Several policies to displace consumption of gasoline and diesel, the standard road vehicle fuels worldwide, emphasize the supply side, seeking to raise the distribution of alternative power sources and vehicles. A prominent example is given by the Corporate Average Fuel Economy (CAFE) regulations in the United States, by which the sale of bi-fuel gasoline-ethanol vehicles—known as "flexible fuel vehicles"—helps carmakers meet their average fleet mileage standards irrespective of whether these vehicles end up utilizing the alternative fuel during their lifetime (Anderson and Sallee 2011). A common assumption on the demand side is that consumers will view alternative energies as close substitutes to gasoline, each product offering homogeneous vehicle miles traveled. By this assumption, consumers will be quick and eager to adopt the alternative energy and vehicle when this alternative reaches the market at a monetary price per mile traveled that is similar to that offered by gasoline (Holland et al. 2009, Salvo and Huse 2011).

In a unique market setting where a majority share of car owners already enjoy an established alternative to gasoline when they shop for fuel, Salvo and Huse (2013) showed that Brazil's consumers substitute between gasoline and sugarcane ethanol over a very wide range of relative price variation. From a sample of 2160 *bi-fuel vehicle* motorists making choices at the pump, Salvo and Huse found that when gasoline was priced 20% above ethanol in \$/mile,[1] about 20% of consumers chose gasoline. Ethanol was not only available at the station, it was typically only one nozzle away from gasoline at the same pump. Similarly, when ethanol was priced 20% above gasoline, 20% of consumers chose ethanol.[2] Examining substitution between gasoline and corn ethanol in the US midwest, Anderson (2012) similarly identified households who were paying a premium for ethanol.

Such findings matter for policy design. If consumers exhibit strong and heterogeneous tastes over real or perceived non-price attributes when choosing between gasoline and an arguably similar alternative—and it may be difficult to find an alternative as close to gasoline as ethanol when it comes to distribution, physical properties, engine technology and performance, etc—it then follows that substitution to less similar energies and technologies, such as natural gas and electricity, is likely to be, at best, gradual and limited. Unless, of course, policymakers are prepared and able to spend heavily and successfully on public campaigns that seek to shift preferences or perceptions. Gasoline may be here to stay for a second century of motorized road travel.

In this paper I ask: To what extent do Brazil's bi-fuel vehicle consumers, who in principle can "flexibly" substitute to the lower priced fuel among gasoline and ethanol at the pump, choose to purchase the more expensive fuel because they value its non-price

[1] This is equal to the fuel's monetary price at the pump, in \$/gallon, divided by the consumer's vehicle mileage in miles/gallon (mpg).

[2] The quick transition by carmakers in Brazil, starting in 2003, from offering single-fuel engines to *only* bi-fuel engines on the vast majority of models, makes consumer selection into owning vehicles with dual gasoline-ethanol capability less of a concern (Salvo and Huse 2013).

characteristics, as they perceive them? These attributes might include, for example: (i) extra range on a full tank, which favors gasoline, (ii) varying engine performance—concern for this attribute pushes some consumers to gasoline, others to ethanol, and (iii) externalities (including pecuniary) such as environmental damage and the local economy (e.g., jobs), with typical perceptions favoring ethanol (Salvo and Huse 2013).

Alternatively, some consumers may pick the expensive fuel from a choice of gasoline and ethanol not because they value different non-price attributes but because *price differences are not well understood or salient.* In particular, since it is a partially oxidized hydrocarbon, a gallon of ethanol provides a lower yield in terms of miles traveled than a gallon of gasoline. Importantly, both in Brazil and in the US, these liquid fuels are priced to the end consumer per unit of volume (liters or gallons) and there is little guidance, if any, at the point of sale as to which fuel yields the highest amount of miles per dollar purchase. Given the local composition of fuels,[3] the rule of thumb which Brazil's motorists are constantly exposed to on the radio and other media channels is that when a liter of ethanol is priced at or below 70% relative to a liter of gasoline, then ethanol offers the lower energy-adjusted price, in \$/km terms.[4] Observing consumers not realize savings per fuel km purchased could then reveal an inability or unwillingness to do the cost conversion or process the available information. Such limited information, limited attention, or bounded rationality would have very different policy implications to taste heterogeneity (Goeree 2008, Clerides and Courty 2016, DellaVigna 2009).

Since unobserved taste heterogeneity can be challenging to establish directly over alternative explanations, I follow an indirect approach. Via a field experiment, the research design is to "shock" consumers at the point of sale with increased relative price salience, and measure the extent to which choice shifts to the energy efficient—lowest \$/km—fuel, away from the alternative with the highest monetary cost. Assuming the intervention succeeds in raising the salience of effective price differences, a low treatment effect would reveal strong taste heterogeneity over fuels. On the other hand, if treatment raises the price elasticity of demand for one fuel relative to the other, this would indicate a role for instruments to better inform consumers about energy-adjusted prices, or to make these more salient.

This study is the first to conduct a real-world randomized experiment to provide

[3]Ethanol retailed in Brazil is E100, pure but "hydrated," i.e., containing up to 4% of water. In the sample period, gasoline retailed in Brazil was either E25 (a 25% ethanol blend) or, after November 2011, E20 (20% ethanol). I refer throughout to E25/E20 and E100 as "gasoline" and "ethanol," respectively. For comparison, gasoline and ethanol available in the US are typically E5/E10 and E85, respectively.

[4]The following example aired on a leading São Paulo radio channel, *CBN Notícias*, on March 24, 2011, during the sample period: (Reporter) "So it's not worth fueling with ethanol?" (A Station Attendant) "No way, you end up spending more..." (Reporter) "For the consumer it makes sense to fuel with ethanol if its price is less than or equal to 70% of the price of gasoline. For example, here in the state capital this Thursday... one liter of gasoline retails for about 2.60 (Brazil Real). Ethanol should be priced at 1.80 to be worth it. As ethanol is being sold for 2.20, this calculation shows that gasoline is favorably priced. The calculation is simple. You take the price of gasoline and multiply by 0.7..." (A Consumer) "I used to fuel with ethanol and now only gasoline..."

consumers with information on price differences across alternative energies.[5] As bi-fuel vehicle motorists pulled up at the pump and before they placed their order with the station attendant,[6] they were informed both verbally and through a flyer about which fuel was comparatively favorably priced at their station on the day. The magnitude of the effective price difference between gasoline and ethanol was also made salient—and it is along this margin that two separate treatments varied. In a sample of participating stations in several cities over one year, the fuel choices of 10,400 subjects were observed in treatment and control groups, as wholesale ethanol prices fluctuated on the back of exogenous variation in world sugar prices, while wholesale gasoline prices remained controlled by the federal government (Salvo and Huse 2011). The control group consisted of random arrivals—also bi-fuel vehicle motorists, but not subject to salient price information—at the same stations and at similar times to treated subjects.

I use a combination of reduced-form and structural methods to analyze these fuel choices. In experiments in which ethanol was very favorably priced relative to gasoline, mainly in São Paulo ($N = 1512$), not long after ethanol prices had fallen sharply across the city, I obtain an average treatment effect that is statistically significant in shifting consumers out of gasoline and into ethanol. In these experiments, gasoline was at least 10% more expensive than ethanol in \$/km. Yet, even in such experiments, the estimated salience treatment effects are small in magnitude: in a population of consumers equipped with technology that in principle allows them to transition seamlessly between gasoline and ethanol, the choice of very favorably priced ethanol increases from about 57% in the control to 63% in the treatment, i.e., by 6 percentage points or by one-tenth.

In experiments in which gasoline was very favorably priced relative to ethanol, and had been so for at least six months, in Recife and Belo Horizonte ($N = 4050$), one treatment increased the adoption of regular-grade gasoline by only 1.4 percentage points, compared to 83% of the control choosing the favorably priced fuel, and this increase is statistically insignificantly different than zero. Faced with such asymmetric prices, that a sizable mass of consumers (37% and 16%) chooses the low yield fuel—offering low mileage per dollar of spending—despite being reminded of this at the point of sale suggests strong fuel preferences or large perceived switching costs.

I then pool the experimental samples and fit two alternative structural models of demand to quantify, over the full range of price variation, the impact of relative price salience on fuel substitution. These allow me to do counterfactual simulations and consumer welfare analysis. I model the treatment effect as either: (model 1) shifting the consumer's sensitivity to price differences, as a "reduced form" in the indirect utility function; or, in a separate model, (model 2) shifting the likelihood that the consumer is considering both gasoline and ethanol rather than just one or the other. In this second model, I allow the consumer's "mindset" to be single-fuel, or "inflexible," despite

[5] Dranove and Jin (2010) review the empirical literature on information (quality) disclosure.

[6] As in New Jersey, consumers in these markets do not pump fuel themselves.

driving a vehicle equipped with "flexible fuel" capability. Model 1 follows most of the empirical IO demand literature, in which advertising and promotion (e.g., store features and displays) enter the utility function. It provides a coherent yet straightforward way to combine the experimental samples, at different price points, within a random utility framework. Model 2 lends itself more naturally to a welfare analysis of the intervention, since treatment affects a consumer's utility insofar as it shifts choice—fixing the choice, treatment does not shift the disutility from a given price level. Model 2 avoids the common assumption that consumers choose from the entire set of alternatives, following Goeree (2008) who models information and advertising as expanding limited choice sets. The modeling is motivated by surveys of consumers who quite literally invoke "inertia," "habit" and "not considering the other fuel" as the basis for their fuel choice (Salvo and Huse 2013), meaning that they are not disposed to considering a substitute when at the pump (or when at the shelf, e.g., Pires 2016), despite knowing of its existence.

The structural analysis shows that my information intervention raises overall price sensitivity—whether in reduced form (model 1) or cognitively by relaxing limited choice sets (model 2). However, the effect is rather small when compared to shifting observable individual characteristics, including education, wealth, and city of residence. For example, college-educated drivers are significantly more price sensitive than drivers with no more than primary education, and the difference in price sensitivity between these two demographic groups is more than twice the price salience treatment effect.

To illustrate, when ethanol is 10% more expensive than gasoline per fuel-km, the probability that gasoline is chosen in the most effective price salience treatment is 76% compared to 74% in the control—only 2 percentage points apart. By contrast, the probability that a college-educated consumer chooses ethanol at this price point is 75% compared to 70% for a consumer with no more than primary education—a larger 5 percentage point variation. More education may correlate with greater financial literacy or weaker perceived differences in non-price attributes across the fuels.[7]

Estimates of the limited choice model imply that only about one-half the population of consumers equipped with bi-fuel vehicle technology actually considers bi-fuel choice sets when making fuel purchases. Treatment can raise this proportion from one-half to six in 10 consumers.[8] Estimates further suggest that in the presence of price asymmetry, policies to increase price salience at the pump could raise aggregate consumer welfare by as much as a 3% drop in fuel prices.

My finding that energy demand is not very price sensitive along the extensive margin of which fuel to buy even when consumers are reminded of the magnitude of relative price

[7]Bronnenberg et al. (2015) find that better-informed or better-educated consumers are less likely to pay extra to buy "premium" brands of physically homogeneous products (e.g., aspirin).

[8]As an analogy, in a price search setting for a specific product, the Halo Reach Xbox game listed by different vendors on eBay, Dinerstein et al. (2014) find a 14% density for *single*-listing "consideration sets" before an information intervention (Table 3). They infer the existence of singleton consideration sets despite many competing offerings (vendors) appearing on the same screen.

differences, is consistent with an empirical literature, particularly on energy efficiency, documenting a gap between choices that appear privately beneficial and those that are actually pursued (Allcott and Greenstone 2012, Gillingham and Palmer 2014). For example, in a field experiment on home "weatherization," Fowlie et al. (2015) observe low take-up of an energy savings program even after the researchers took "extraordinary efforts to inform households—via multiple channels—about the sizeable benefits and zero monetary costs" (p.204). No more than 15% of households in the information ("encouragement") treatment submitted a weatherization application. My paper contributes to an expanding literature, both with and without a behavioral orientation, that considers how information at the point of consumption can shift household energy use (Allcott and Mullainathan 2010, Jessoe and Rapson 2014).[9]

Beyond the energy setting, Chetty et al. (2009) conduct a field experiment that shares features with mine. They find that posting tax-inclusive price tags on the shelf at a grocery store reduces demand by 8%, concluding that "consumers underreact to taxes that are not salient." Much as consumers might underreact to a sales tax that "is not included in the posted price... is added at the register (and hence is less salient)" (p.1146), consumers may underreact to price differences stemming from differences in energy content that are not posted at the pump (and are "added" only on the road).[10] The "left-digit bias" in the valuation of car mileage identified by Lacetera et al. (2012) and Busse et al. (2013) suggests that salience may matter even when information is staring consumers in the face, namely that "final customers... find the first digit of the odometer more salient than trailing digits" (p.575). Goeree (2008) uses observational advertising data—not experimental information exposure shocks as I do—to shift "random" choice sets. Similarly, in her model "imperfect substitutability between different brands may arise from limited consumer information about product offerings as well as from idiosyncratic brand preferences" (p.1018). A recent line of research investigates how price (or financial) comparisons impact consumer choice, including Choi et al. (2011), Kling et al. (2012) and Handel (2013).

The balance of the paper is as follows. Section 2 presents the experimental design and setting, including price variation. I estimate treatment effects in Section 3. Section 4

[9] For perspective, Allcott and Greenstone (2012) survey the "energy efficiency gap" literature and state: "An additional example of testing for imperfect information using equilibrium outcomes is to examine whether information disclosure increases the elasticity of energy-saving technical change with respect to energy prices... (this) approach to assessing the magnitude of imperfect information is to test for the effects of information disclosure on purchase decisions. This approach has the benefit of being based on observed choices in the marketplace, instead of beliefs stated on a survey. We are not aware of any large-scale randomized evaluations of energy efficiency information disclosure" (p.20).

[10] Chetty et al. (2009) distinguish between consumers "uninformed about the sales tax rate or which goods are subject to sales tax" and an "alternative hypothesis... that salience matters" (p.1147). They call on surveys finding that a majority of grocery shoppers knew the tax status of products on the survey to argue that "our empirical findings are driven by salience effects." Motivating my field study, Salvo and Huse (2013) report telephone surveys of bi-fuel vehicle owners that found that most were aware of differences in energy content between equal volumes of gasoline and ethanol.

specifies and estimates two alternative structural models in which raising salience at the point of sale shifts demand, and conducts counterfactual analysis. Section 5 concludes.

2 Experimental setting and design

The experiments were conducted during 193 "visit-days" to 52 different fueling stations, or about four visits per station. Stations from an authorized set of Shell, Cosan or Esso branded stations were repeatedly visited by enumerators, on spaced out dates, as the price of ethanol varied relative to that of gasoline.[11] Visits took place in four large cities between May 2011 and March 2012, and comprised 10,422 subjects across control and treatment groups.

A station visit. The agreed procedure for visiting a station to conduct experiments was as follows.[12] Late afternoon prior to the day of the visit, the market research firm's regional office would phone the station manager to confirm the visit by an enumerator and to obtain pump prices for the next day. (Shell-Cosan's sales executives would already have pre-arranged the visits with the station manager, to facilitate access.) After entering regular-grade gasoline and ethanol prices into a protected spreadsheet that I had provided, office staff would print the visit-specific price salience flyers, for two different treatments, and hand them over to the enumerator for use in the field visit.

An enumerator would begin a visit at a varying time of the day (e.g., early morning or early afternoon), Monday through Saturday, and the visit ended later that same day. By design, the first 18 consumers to pull up were assigned to control, the next 18 consumers were assigned to one treatment group, and the subsequent 18 consumers were assigned to another treatment group. While we randomized on the sequence in which the two treatments were applied, the market research firm recommended that the control lead the treatment. This was to reduce the possibility that the experimenter, having implemented a treatment immediately before, also mistakenly treat control subjects with price information.[13]

[11] These station brands account for one-fifth to one-quarter of all stations across the cities.

[12] This procedure was agreed between myself, as the Principal Investigator, the market research firm I hired (CNPBrasil), its regional field offices and enumerators, as well as Shell-Cosan's pricing manager in the central office and sales executives in the field, who liaise with station managers and franchisees. Thus many parties were involved. The research design involved repeated visits to each station due in part to the administrative cost of enlisting the station in the experiment and gaining access to its customer pool. It is unlikely, however, that a same consumer was sampled more than once. Consecutive visits to a same station were spaced out by at least one fortnight, and often by more than this.

[13] The market research firm, experienced in this particular environment, argued that the sampling times across treatments and control would be similar, and that an alternative design whereby information provision would randomly vary one consumer arrival to the next might be poorly implemented, in particular, due to station attendants also administering treatment (see below). Importantly, while the retailer generously granted access to its workers and consumers, it was critical to avoid disrupting a high-frequency sales operation. The balancing tests reported below are generally favorable.

In the control group, once the consumer had placed his order with the station's attendant and the vehicle was being serviced, the enumerator would briefly survey the consumer to complement the observed fuel choice with demographic characteristics. Consumers mostly agreed to participate, thanks to the low opportunity cost of their time as they waited while fuel was being pumped, as well as the survey's non-commercial research purpose. Whenever a consumer refused to participate in the interview, the enumerator would comply and step back, but would record the vehicle make and model. The design suggested that a consumer, having pulled up, would be about as likely to agree to participate irrespective of whether he was in the control or in the treatments. On completing an observation, the enumerator would typically wait for the next consumer to arrive, as queues at the pump rarely form in these markets.

After completing 18 control observations, the enumerator would proceed to the treatment stages of the visit. Subjects were again assigned by order of arrival. There were two treatment groups, again, each consisting of 18 consumers. In each case, subjects were verbally informed, prior to placing their order, of the favorably priced fuel—gasoline or ethanol—at the given station they had chosen to pull up at that given day.

While the verbal information was the same, the treatments differed in the flyer that the subject was shown, also prior to ordering. In the first treatment, subjects were shown the ethanol-to-gasoline per-liter price ratio, p_e/p_g, relative to a 70% "parity" threshold widely reported in the media as being the ratio at which \$ per kilometer traveled for a typical vehicle equalizes across fuels (more below). In the second treatment, subjects were shown a table that compared the expected distance to be driven on 50 Brazil Real (R\$) of gasoline purchased at the station to the distance from alternatively spending R\$ 50 on ethanol.[14] As with the control group, after placing his order and while idling inside the vehicle, each treated subject was interviewed to collect observables. For some field visits, the "price ratio relative to 70% " treatment was applied first, while for other visits the "km per R\$ 50" treatment was applied first.

Perhaps the most critical aspect of price salience experiments such as these is how information is made available to the consumer, just after he arrives at the point of sale and before placing his order (Chetty et al. 2009). Besides the alternative flyers, over time I experimented with either the station attendant or the enumerator sharing price information with the subject. A potential upside of having the attendant treat the subject is that he is the person who normally interacts with the consumer at the pump, and with whom the consumer is often times familiar. Salvo and Huse (2013) found that motorists tended to fuel at the same station, i.e., in equilibrium, substitution across stations was low in their sample. The downside of using the attendant is that the subject

[14] R\$ 50 amounted to about US\$ 25 and, in a market where most consumers paid in cash, this ticket value was popular. The mean purchase value in Salvo and Huse's (2013) survey was R\$ 47. These values correspond to regular-grade fuels and, in the km/\$ treatment, representative vehicle engines. Compared to the US (Hastings and Shapiro 2013), the share of midgrade fuel was low.

may become suspicious as to why he is unusually being offered price information by someone who works for the seller.[15] Unless asked, attendants do not usually recommend one fuel over another, though consumers do occasionally request their advice. Moreover, attendants may not be clear communicators.

On the other hand, enumerators tend to be better communicators—the ability to communicate is presumably a characteristic that determines their selection into the job. However, they are not part of the standard shopping experience. For all experiments, both the attendant and the enumerator were immediately visible to subjects, including control, with the attendant dressed in his branded, colored work outfit, whereas the enumerator wore a shirt and identification card labeled "research" and held a clipboard.

An alternative design might be one in which salient price information—which in the present design was printed on the flyers—were automatically displayed by the pump and would be visible to the treated subjects prior to their ordering, without being guided by a representative of the seller or of the market research firm (e.g., Chetty et al. 2009). In this alternative design, the control might then be shoppers at other display-free but otherwise similar outlets—rather than, as in the present experiments, same-station consumers at similar times.

Fuel price paths. Figures 1A and 1B report price paths prepared from weekly samples of stations that are representative of market prices in the four cities. The top panels report the ethanol-to-gasoline per-liter price ratio, p_e/p_g, for regular-grade fuels, with thick curves indicating medians and thin curves indicating the 5th, 25th, 75th, and 95th percentiles in these large cross-sections of stations surveyed by an external source. The 70% ethanol-gasoline parity threshold (again, commonly reported in the media) is marked by a horizontal line. The bottom panels report p_g and p_e separately, and 70% of p_g—a popular consumer heuristic (see note 4)—is also shown for comparison with p_e. During the period, gasoline prices varied significantly less than ethanol prices, and correlation between the series is in large part due to gasoline fuel retailed in Brazil in fact containing a one-fifth (E20) to one-quarter (E25) volume fraction of ethanol. The vertical lines in the figures indicate weeks with field activity: dashed lines denote visits in which treated subjects received price information by a station attendant, whereas solid lines denote visits in which price information was rendered by an enumerator.

In the southeastern/southern cities of São Paulo, Curitiba, and Belo Horizonte, experiments were conducted right after ethanol prices had fallen sharply, between May and July 2011. After several months of high prices, ethanol was now favorably priced

[15] Allcott and Sweeney (2015) run an experiment in which an appliance retailer's sales agents offer consumers information on the energy efficiency of alternative water heaters in their choice set. The information being disclosed in their setting is arguably more complex and less verifiable than a simple comparison of prices at the pump—a comparison that Brazil's motorists may be more familiar with given the media (e.g., radio) attention and the repeated nature of fuel purchases. Also, motorists are often times familiar with their neighborhood station's attendants (Salvo and Huse 2013).

relative to gasoline in São Paulo and Curitiba, but ethanol remained more expensive in Belo Horizonte, even if less so than before. Again, to see this, compare p_e/p_g against the 70% parity line in the top panels, or compare p_e against 70% of p_g in the bottom panels. Experiments were also carried out around mid 2011 in Recife, in the country's northeast, home to a sugarcane industry that is less integrated with southeastern markets: there, ethanol prices did not fall and gasoline remained favorably priced relative to ethanol.

After a two-month window which allowed for preliminary analysis of the "first wave" of experimental work and repeated instruction of the field teams in the different cities, fieldwork resumed in September 2011 and continued until March 2012. Price paths during this "second wave" of field activity, from September 2011 on, were quite stable in all four cities compared to the earlier wave up to July 2011. Ethanol and gasoline were similarly priced (in $/km) in São Paulo (Figure 1A, left). Gasoline was favorably priced relative to ethanol in Curitiba, Belo Horizonte, and Recife (Figures 1A and 1B).

In terms of relative price points, Table 1 summarizes, for each price informant by wave combination implemented in the field, the distribution of p_e/p_g faced by subjects in each city: (i) station attendant informing prices in wave 1, (ii) attendant informing prices in wave 2, and (iii) enumerator informing prices in wave 2. In terms of recent price history, a distinction can be made between: (i) fieldwork in São Paulo, Curitiba, and Belo Horizonte during wave 1, in which ethanol prices had fallen sharply, and (ii) fieldwork during later dates or in Recife, characterized by flatter price paths.

Two treatments by station visit. A treated subject would, prior to placing his order, hear one of the following statements, either from the attendant or from the enumerator. The statement was consistent with actual price levels for regular-grade gasoline and ethanol posted at the station's pump on the day of the visit:

a. *Hoje a gasolina está mais vantajosa, veja aqui*, loosely translated as "Today gasoline is more advantageous / the better deal, see here,"

b. *Hoje o álcool está mais vantajoso, veja aqui*, translated as "Today ethanol is more advantageous / the better deal, see here,"

c. *Hoje a gasolina e o álcool estão com rendimento parecido, veja aqui*, translated as "Today gasoline and ethanol offer similar yields / similar deals, see here."

As he heard the verbal statement on relative price conditions at the station that day, a subject would be handed a flyer similar to that shown in Figure 2A, for the "price ratio relative to 70% " treatment, or Figure 2B, for the "km per R$ 50" treatment. The illustrated flyers correspond to a station visited in São Paulo on June 13, 2011 in which $(p_e, p_g) = (1.649, 2.499)$. To be clear, some treated consumers were handed a price-ratio flyer and others were treated with a km-per-R$50 flyer, and in either case the flyer was consistent with the verbal statement that introduced it.

The price-ratio flyer (Figure 2A) gave "thumbs up"—alongside *Mais vantagem*—

status to gasoline when $p_e/p_g \geq 0.705$, thumbs-up to ethanol when $p_e/p_g < 0.695$, and stated that both fuels offered similar yields (*Rendimento parecido*) otherwise. The flyer also reminded subjects of the media-reported parity threshold, stating that "specialists advise that when this price ratio (between ethanol and gasoline) is: (i) lower than 70%, ethanol is more advantageous; and (ii) higher than 70%, gasoline is more advantageous."

The km-per-R$50 flyer (Figure 2B) gave thumbs-up to gasoline when the distance to be traveled on R$ 50 of regular gasoline purchased at the station was expected to exceed, by some margin, the distance to be traveled on R$ 50 of regular ethanol, thumbs-up to ethanol for the opposite situation, and was neutral when gasoline and ethanol offered similar yields. Specifically, a table in the flyer compared the expected distances across ethanol and gasoline for three most-popular vehicle engine sizes: 1.0, 1.4, and 1.8 liter (absolute fuel economy generally declines with engine size). Appendix A.1 explains how I calculated these distances from pump prices and fuel efficiencies for a sample of vehicles obtained from urban driving simulations in the laboratory.

To illustrate, fuel economy for one such 1.0-liter engine vehicle, a GM Celta Life model-year 2009, is $k_e = 7.7$ km/liter on ethanol (E100) and $k_g = 11.1$ km/l on gasoline (E25). I used this GM Celta 1.0, along with other tested vehicles, to report distances $50k_e/p_e$ and $50k_g/p_g$ shown in the first row of the table printed on the km-per-R$50 flyer. For the illustrated station visit, this amounts to 233 km on R$ 50 of ethanol and 222 km on R$ 50 of gasoline; averaged across lab-tested 1.0-liter vehicles, respective distances would be 227 and 218 km as indicated (Figure 2B).

3 Average treatment effects

3.1 Covariate balance

Table 2 reports on balancing tests for observables that vary at the consumer level, since covariates that are invariant within station visit are balanced across control and treatments by design. Overall, evidence in support of balance is strong. One small difference is in the age composition of subjects, with the treatment groups containing slightly lower proportions of consumers aged over 65 years (3.5% and 3.7% in the treatments versus 4.5% in the control). Broadly speaking, the sample appears to have the statistical properties of a randomized experiment.

It is reassuring that the demographic characteristics of the sample are very similar to that of Salvo and Huse (2013). For example, females, subjects aged 25 to 40 years, subjects who completed college, and subjects driving Fiat vehicles account, respectively, for 34%, 47%, 55% and 30% of the present sample ($N = 10,422$), compared to 34%, 46%, 50% and 28% in the Salvo and Huse sample ($N = 2,160$, Table A2).

Compliance. Across station visits, the mean reported number of subjects during the control stage ("assigned to control") who declined to participate in the interview is 3.3, out of 18 "compliers" who agreed to participate. The reported number of non-compliers reduces to 2.7 and 2.5 (again out of 18 subjects) in the price-ratio and km-per-R$ 50 treatments, respectively. While not large, these differences are statistically significant at the 10% and 5% levels, respectively—see the last row of Table 2. However, there is suggestive evidence that some enumerators may erroneously have tallied non-compliers in the subsequent treatments as non-compliers in the control, so this difference between treatments and control may be overstated.[16]

Coupled with the favorable balancing test results, the low incremental cost of treatment in these experiments, both to the subject and to the person delivering treatment, suggests that any drivers of non-compliance would not only be limited in magnitude but also similar across control and treatment. During fieldwork there was no indication that subjects, who had already driven up to the pump to communicate with the station attendant for service, were more or less likely to decline to participate on the basis of any price information that was shared with them. Nor was there any indication that treatment and control subjects would be differentially induced to drive away without fueling based on such information.

The vehicle make and model recorded for non-compliers, who declined to participate, is comparable to that for compliers. For example, estimated vehicle prices average R$ 30,793 for non-compliers ($N = 1,819$) against R$ 29,043 for compliers ($N = 10,422$, control and treatments). While statistically significant, this 6% difference in vehicle price is not large, and there is suggestive evidence that some enumerators may erroneously have tallied as "non-compliers" consumers who drove more expensive single-fuel imported vehicles—and who were thus, correctly, left outside the subject pool.

3.2 Fuel choices: Control and treatments

Tables 3 and 4 report average treatment effects for experiments conducted at fueling stations grouped by different relative price points, namely when ethanol was very or somewhat favorably priced relative to gasoline and the other way round, when gasoline was very or somewhat favorably priced relative to ethanol. Appendix A.2 reports results for experiments conducted in markets where both fuels were similarly priced. Appendix A.3 provides descriptive plots of fuel choices at the different price points and distinct price histories in the sample, separately by control and treatment groups. Appendix A.3 also plots average treatment effects, at the station-visit level, by treatment type, i.e., the type of information flyer and the type of agent disclosing information to the consumer.

[16] Enumerators were asked to tally non-compliers separately by control and treatment group in a table on the "lot cover page" (and record their vehicle characteristics on the following page). The first row of this table was meant only for the control group, but may have been used in the treatments. In hindsight, the design or enforcement of this aspect of fieldwork could have been tighter.

3.2.1 Favorably priced ethanol experiments

Experiments where $p_e/p_g < 0.7/1.1$ [EE]. Ethanol was "very" favorably priced relative to gasoline mostly for first-wave visits in São Paulo between May and July 2011, after ethanol prices had fallen sharply, with the attendant informing prices (Figure 1A and Table 1). For clarity, the condition $p_e/p_g < 0.7/1.1$ should be interpreted as gasoline being at least 10% more expensive than ethanol in \$/km (recall that $k_e/k_g \simeq 0.7$).

The top panel of Table 3 reports that for such very-favorable-ethanol experiments, 57% of subjects in the control group chose ethanol over gasoline, with this proportion rising by 6 percentage points to 63% of subjects in the price-ratio treatment. This increase is significantly different than zero at the 5% level. Similarly, a higher proportion of subjects in the km-per-R\$ 50 treatment chose very favorably priced ethanol over gasoline compared with the control, 62% versus 57%, and this 5 percentage point difference is significant at the 10% level.

The panel also reports that the mean value of fuel purchased is statistically similar across treatments and control, suggesting that treatment did not induce subjects to change their allocated expenditure. Also statistically indistinguishable across treatment and control is the mean weekly vehicle usage stated by consumers, among those consumers who during the brief demographic survey were able to state their typical usage.

Two further points should be noted. First, the 57% share of consumers in the control group choosing ethanol when this fuel was very favorably priced compared to the gasoline substitute (by 10% or more), may strike one as low. This finding may owe partly to the fact that in such markets ethanol prices had just fallen and the depth of the price drop may not have been salient to many consumers. The time path of prices thus provided a unique setting to conduct price salience experiments. Second, the 5 to 6 percentage increase, upon treatment, in the adoption of a very favorably priced fuel, while statistically significant, seems economically limited: 37-38% of bi-fuel vehicle motorists who could switch to a substantially lower priced fuel and were reminded of this at the pump did not do so.

It may be that attendants were unwilling or unable to communicate effectively, but the effectiveness of enumerators in later experiments, as measured by treatment effects, was similar. Rather, it is possible that many consumers had strong tastes for non-price characteristics. Consumers may pull up at the pump having already decided which fuel to purchase, and their choice was largely not affected by my attempt to make effective prices more salient. This is the paper's main finding. While automakers market bi-fuel vehicles as "flexible," their owners may not be as flexible or homogeneous when it comes to price substitution across fuels. As shown below, greater price salience at the point of sale did not make a fuel's demand curve significantly more elastic.

Experiments where $0.7/1.1 \leq p_e/p_g < 0.7/1.05$ **[E].** Ethanol was "somewhat" favorably priced relative to gasoline on several other first-wave visits in São Paulo and Curitiba, in which effective prices were made salient to treated subjects mostly by attendants (Table 1). In these experiments, gasoline was between 5% and 10% more expensive than ethanol in km-equivalent units.

As the bottom panel of Table 3 shows, despite ethanol not being as favorably priced as in experiments [**EE**], a slightly higher 59% of the control chose ethanol over gasoline (compared to 57% in experiments [**EE**]). It is conceivable that the drop in ethanol prices was less recent at this point in time compared to when experiments [**EE**] were undertaken. Ethanol prices fell sharply in May/June 2011 and then rose somewhat, particularly in Curitiba (Figure 1A). Indeed, for first-wave Curitiba experiments, the median date of station visits with somewhat favorably priced ethanol is ten days later than the median date with very favorably priced ethanol.

For these mostly first-wave São Paulo and Curitiba experiments, the average treatment effect is statistically insignificant, and the point estimate goes the "wrong" way.[17] Again, the mean purchase value and stated vehicle usage are similar for treatments and control, which is indicative of both balancing between control and treatment groups as well as insignificant treatment effects on these additional outcome variables.

3.2.2 Favorably priced gasoline experiments

Experiments where $p_e/p_g \geq 0.7 \times 1.1$ **[GG].** Gasoline was "very" favorably priced relative to ethanol during most experiments in Recife and some experiments in Belo Horizonte, across both first and second waves and with treatment rendered variably by attendants or enumerators (Table 1). Again for clarity, the condition $p_e/p_g \geq 0.7 \times 1.1$ requires ethanol to be at least 10% more expensive than gasoline in \$/km. Importantly, relative prices had been largely stable in these markets (Figure 1B).

During these experiments, as the top panel of Table 4 indicates, 88% of subjects in the control group chose gasoline over ethanol, compared to slightly higher adoption of very favorably priced gasoline in the price-ratio and km-per-R\$ 50 treatments, respectively, of 90% and 89%. While positively signed, i.e., in the direction of the favorably priced fuel, differences for either treatment versus control are statistically insignificantly different from zero. This is despite the number of subjects in these experiments being substantially higher compared to the other experiments, as a power calculation would indicate (the favorably priced fuel was chosen by almost 9 in 10 control subjects).

The vast majority of stations retailed midgrade gasoline alongside the regular grade.

[17] Perhaps treatment in Curitiba alerted some consumers that ethanol prices had reversed course and were now on the rise. Restricting to [**E**] experiments in São Paulo only, the point estimate for the average treatment effect no longer goes the "wrong" way: the proportion of subjects choosing ethanol is 0.544 in the treatment groups ($N = 540$) compared with 0.526 in the control ($N = 270$), though the difference is statistically insignificant. For brevity, such results are not reported.

Thus, Table 4 also reports average treatment effects on the adoption of regular-grade gasoline, which are also insignificant; point estimates are 0 to 1 percentage point. Regular-grade gasoline consumers outnumbered midgrade gasoline ones by 15 to 1, in control and treatment groups alike.

In light of the history of relative price stability running up to these experiments, as well as the highly-competitive-gasoline price point, it is plausible that any remaining consumers of ethanol in these markets had a strong taste for the fuel. They would thus be unlikely to adopt gasoline simply by virtue of being treated with more salient price information. These holdouts are, in the words of Salvo and Huse (2013), "ethanol fans."

Experiments where $0.7 \times 1.05 \leq p_e/p_g < 0.7 \times 1.1$ **[G].** Experiments with "somewhat" favorably priced gasoline relative to ethanol are reported in the bottom panel of Table 4. Similar to experiments [**GG**] in which gasoline was even more favorably priced, treatment yielded statistically insignificant increases in adoption of gasoline, of 0 to 2 percentage points. 84% of subjects in the control group chose gasoline over ethanol, a lower proportion than control subjects in experiments [**GG**] (88%). Two-thirds of experiments at this price point took place in Belo Horizonte during first-wave visits, after ethanol prices had fallen but still remained high relative to gasoline, with treatment administered by attendants (Table 1).

3.2.3 Further statistics and analysis

Comparing treatments. A tentative point can be made from Tables 3 and 4 that average effects for the price-ratio treatment, while small and often themselves insignificant, tend to be slightly larger than average effects for the km-per-R$ 50 treatment. In view of the 70% benchmark widely reported in the media (see note 4), it is plausible that informing the per-liter price ratio at the point of sale was more effective at raising salience to consumers in these markets than displaying a table that stated the distances a R$ 50 bill was expected to fetch alternatively on gasoline and ethanol. Salvo and Huse (2013) found from telephone surveys that consumers were quite aware of the 70% threshold, on which the price-ratio flyer was based.

Controlling for station-day fixed effects. The average marginal effect from either treatment on the choice of the favorably priced fuel, for the different price points, is shown in Table 5, using first a linear probability model (upper part of table) and then a probit model (lower part). For each model I include station-visit fixed effects; along with treatment, these help to explain 6-8% of the variation in choices. Estimated station visit-clustered standard errors are reported in parentheses.

Results, including standard errors, are similar to the simple differences in means reported earlier in Tables 3 and 4. (Estimates are also very similar if we additionally control for individual demographics, as in the next section.) While it can be statistically significantly different from zero, the effect of nudging consumers to choose the favorably priced fuel, by making true effective prices more salient at the point of sale, was modest. This finding is, of course, based on each subject's single observed choice, since subsequent transactions, i.e., the consumer's choices in future shopping trips, are unobservable.

To summarize the experiments in which effects were significant, in markets with very favorably priced ethanol [**EE**], treatment with the price-ratio flyer or the km-per-R$ 50 flyer similarly raised the proportion of consumers who chose ethanol over gasoline by 5 to 6 percentage points compared to the control, with estimated standard errors on these estimates of just over 2 percentage points. In the control group, 57% of subjects chose ethanol over gasoline. Thus, increasing price salience raised the adoption of the very favorably priced fuel by *one-tenth*. Compared to the other experiments, significant treatment effects may have been obtained in these markets in part because relative prices had recently shifted.

In contrast, in experiments with very favorably priced gasoline [**GG**] and a recent history of price stability, treatment effects were not significant, despite being precisely estimated, with standard errors of just over 1 percentage point. In such markets, for every ten owners of bi-fuel vehicles, the one consumer who holds on to ethanol despite gasoline being so competitively priced at the pump likely has a strong taste (resp., distaste) for ethanol (resp., gasoline).

4 Salience as a fuel demand shifter

I specify two alternative discrete-choice models of demand for gasoline and ethanol that incorporate the effect of raising the salience of relative prices. In the first model, salience enters the indirect utility function and can shift a consumer's sensitivity to price differences and tastes for the different fuels (akin to individual demand models in which exposure to advertising and in-store promotion are controls in the utility specification). In the second model, salience can shift the likelihood that a consumer pulling up at the pump chooses from a bi-fuel set, containing gasoline and ethanol, rather than a choice set limited to a single fuel, i.e., gasoline-only or ethanol-only. Whereas all consumers drive vehicles equipped with bi-fuel capability, I model the consumer's mindset as possibly being single-fuel: the car may be "flexible fuel" but the consumer is not. I estimate these models using the individual choice data—repeated cross-sections—collected over the course of the field visits. The exercise quantifies how raising price salience at the pump impacts fuel substitution over a wide range of price variation, and compares the effect of raising salience to observable variation in consumer characteristics, such as education and wealth. I use the second model to perform a tentative welfare calculation, conditional on this being the true model, acknowledging that welfare analysis under "nonstandard" decision making may be ambiguous (Gillingham and Palmer 2014).

4.1 Model 1: Salience shifts a consumer's relative utility

Consider a bi-fuel vehicle driver choosing between regular-grade gasoline, g, midgrade gasoline, $\bar{g}$, and regular-grade ethanol, e, at the pump. Having pulled up at station l on date t, consumer i chooses the fuel $j \in \{g, \bar{g}, e\}$ that maximizes utility:

$$u_{ijlt} = \alpha_i p_{ijlt} + x_l'\beta_{1j} + x_{jlt}'\beta_{2j} + x_i'\beta_{3j} + \xi_{jl} + \varepsilon_{ijlt}. \tag{1}$$

In the first term, p_{ijlt} is the effective price consumer i faces for fuel j in market lt, namely R$ per km traveled should his vehicle operate on that fuel purchased at that location on that date. Price sensitivity α_i can vary across consumers according to:

$$\alpha_i = \alpha_1 + x_i'\alpha_2 + T_i'\alpha_3. \tag{2}$$

Vector x_i contains demographic variables such as the consumer's vehicle price—a proxy for his wealth or income—and his schooling. Also shifting price sensitivity, for $\alpha_3 \neq 0$, dummy variables T_i indicate whether the consumer is in either treatment group, $T_i^{ratio} = 1$ for the price-ratio flyer and $T_i^{km} = 1$ for the km-per-R$50 flyer, or in the control, $T_i^{ratio} = T_i^{km} = 0$. Treatments can shift the sensitivity to price differences to capture the notion that a treated consumer may make different tradeoffs between price and non-price characteristics compared to a control.

The second to fourth terms in (1) allow choice probabilities for each fuel to shift differently according to the following sets of observables: (i) x_l, retailer specific characteristics, e.g., the average income (proxied by observed vehicle value) of a station's customer pool; (ii) x_{jlt}, fuel by market characteristics, e.g., the observed number of nozzles dispensing each fuel across all the station's pumps on the date of the visit; and (iii) x_i, individual characteristics, e.g., the consumer's gender, age, education, wealth proxy. Treatment can also directly shift the taste for each fuel, irrespective of prices, via the term $x_i'\beta_{3j}$, by adding T_i to vector x_i.

The fifth term ξ_{jl} captures unobserved product-retailer effects, e.g., the perceived quality of fuel j at station l, assumed to be constant over a station's consumers in the 11-month sample. Finally, ε_{ijlt} is an unobserved consumer-specific idiosyncratic taste for fuel j, distributed—for analytical convenience—extreme-value type I, i.i.d. across consumers, fuels, and markets.

To estimate this model, we can collect terms that do not vary across consumers and time, denoting as mean utility:

$$\delta_{jl} = x_l'\beta_{1j} + \xi_{jl}.$$

The model can be estimated via an inner fixed point algorithm. In the inner loop, conditional on every guess of $\theta_2 = (\alpha_1, \alpha_2, \alpha_3, \beta_{2j}, \beta_{3j})$, a contraction mapping yields mean utilities, $\delta_{jl}(\theta_2)$. This contraction is based on equating model-predicted fuel shares, aggregated across consumers observed at each retailer, to shares observed in the data (Berry 1994). In the outer loop, θ_2 is estimated by maximum likelihood. Parameters $\theta_1 = (\beta_{1j})$ and product-retailer effects, ξ, to the extent they are of interest, may be estimated from estimated mean utilities, $\hat{\delta}$.

Specifically, the outer loop solves:

$$\arg\max_{\theta_2} \log L\left(data|\, \theta_2\right),$$

where

$$L\left(data|\, \theta_2\right) = \prod\nolimits_{i(,l,t)} \prod\nolimits_{j\in\{g,\bar{g},e\}} \left(s_{ijlt}(\theta_2)\right)^{1(i \text{ chooses } j)},$$

indicator $1\left(i \text{ chooses } j\right)$ equals 1 if consumer i chooses fuel j, and 0 otherwise, and model-predicted individual-level fuel shares are given by the expression:

$$s_{ijlt}(\theta_2) = \frac{\delta_{jl}(\theta_2) + \mu_{ijlt}(\theta_2)}{\sum_{k\in\{g,\bar{g},e\}} \left(\delta_{kl}(\theta_2) + \mu_{iklt}(\theta_2)\right)} \tag{3}$$

with $\mu_{ijlt}(\theta_2) = \alpha_i p_{ijlt} + x'_{jlt}\beta_{2j} + x'_i\beta_{3j}$.

Key to identification of this demand model is the design feature that each station was visited on multiple dates as relative prices varied for supply side reasons, namely, variation in the world price of sugar driving sugarcane ethanol prices and government control of wholesale gasoline prices (Salvo and Huse 2010, 2013). To the extent that unobserved fuel-station specific attributes influence prices, by conditioning on the part of the error correlated with price, ξ_{jl}, the remaining individual error ε_{ijlt} is uncorrelated with price (e.g., as in Goolsbee and Petrin 2004).

The above discussion assumed that consumers at the pump choose between three fuels: regular-grade gasoline, midgrade gasoline, and regular-grade ethanol, i.e., their choice set is $\{g, \bar{g}, e\}$. In the data, whereas regular-grade gasoline and ethanol were always available during station visits, midgrade gasoline was not available at (or not carried by) the retailer on 3% of observed purchases. In these occasions where $\bar{g}$ was not available, the consumer's choice set reduces to $\{g, e\}$.[18] By design, there is no outside good and, similar to Houde (2012) and Hastings and Shapiro (2013), I take fuel quantities in km purchased as exogenous. Salvo and Geiger (2014) find no evidence that road usage, as measured by recorded traffic congestion and travel times in the city of São Paulo, varied between 2008 and 2011 as relative fuel prices varied.

[18] In contrast to midgrade gasoline, the availability of midgrade ethanol—denote this fuel by $\bar{e}$—was rare during the sample period. In fact, in the sample of 10,422 consumers, I observe a mere 93 consumers, or less than 1% of the sample, purchase $\bar{e}$. I drop these 93 observations from the structural analysis. Thus there are 10,329 observations.

Results. Table 6 reports estimates of the salience-shifts-relative-utility model, under different specifications for (1), using the full sample of choices (10,329 purchases). One strength of this more structural approach is that we can naturally pool the estimation sample across all relative price points, observed across the different cities over both waves of field activity (compare the analysis to that of the previous section).

The specification reported in column I includes city-fuel fixed effects, whereas columns II to IV specify more granular station-fuel fixed effects, ξ_{jl}. Across the 52 stations visited repeatedly, the mean number of observations is 199 fuel purchases per station. In columns I and II, price sensitivity can vary with wealth, as proxied by the value of the consumer's vehicle. Compared to column II, price sensitivity can additionally shift with age in column III, or with education in column IV.

Estimates are quite stable across specifications. Treated subjects, particularly in the price-ratio treatment, are on average more sensitive to price differences compared to the control. For example, in column IV, $\hat{\alpha}_3^{ratio} = -4.5$ with a standard error (s.e.) of 2.0 in row 3, whereas in the km-per-R\$50 treatment, $\hat{\alpha}_3^{km} = -2.4$ (s.e. 2.0, row 4), on top of a base price sensitivity of $\hat{\alpha}_1 = -29.6$ (s.e. 9.6, row 1), for an individual with no more than primary schooling. Estimates also suggest that treatment does not affect fuel preferences other than through its effect on price sensitivity—see rows 11-12 and 21-22.

In column IV, price sensitivity decreases in wealth, with $\hat{\alpha}_2^{\text{VehiclePrice}} = +4.7$ (s.e. 2.7, row 2). Estimates further suggest that sensitivity to price differences grows with schooling. In particular, consumers who reached secondary school or are college educated respond to price differences more than those who did not complete primary school, i.e., $\hat{\alpha}_2^{\text{HighSchool}}$ and $\hat{\alpha}_2^{\text{College}}$ are, respectively, -8.4 (s.e. 4.3, row 8) and -9.3 (s.e. 4.2, row 9). One interpretation is that more educated consumers have a higher level of financial literacy (Lusardi and Mitchell 2008), or they place relatively more weight on the price characteristic across fuels they perceive as otherwise similar (Bronnenberg et al. 2015).

The price of the vehicle is also associated with a stronger taste for gasoline over ethanol, particularly midgrade gasoline (rows 10 and 20). Other demographic tastes come out strongly significant, such as female consumers favoring regular-grade—but not midgrade—gasoline over ethanol (rows 13 and 23), and a preference for midgrade gasoline monotonically increasing with the consumer's age (rows 24-26).

In an additional specification, which the table does not report for brevity, I control for the (logarithm of the) number of nozzles dispensing each fuel across all pumps at the station on the date of the visit—recall x_{jlt} in (1). One can interpret this covariate as "shelf space," capturing potential shifts in promotional activity, though one should note that promotions are considerably less common a feature here compared to other consumer markets such as soda or detergent. Accordingly, x_{jlt} varies mostly across stations—variation that is already picked up in the ξ_{jl}—than over time. Estimates are very similar to the baseline estimates reported in column IV, and I find the number of nozzles to be positively and significantly associated with fuel choice: $(\hat{\beta}_{2g}, \hat{\beta}_{2\bar{g}}, \hat{\beta}_{2e}) = (0.36, 1.29, 1.47)$, with s.e. of $(0.07, 0.11, 0.06)$.[19]

Counterfactual experiments. I conduct counterfactual exercises to evaluate whether the estimated price salience effects are economically large, at the different price points, compared to other in-sample variation. Figure 3 shows that this is not the case. The left panel shows that treating all consumers with salient price information at the pump, even with the price-ratio flyer, is predicted to have a small effect on the choice of the favorably priced fuel: aggregate demand—with ethanol demand plotted—is *not* substantially more elastic. By contrast, in the right panel of Figure 3, consumers with college education exhibit substantially more elastic fuel demand compared to a "control" group with no more than primary education.

For example, when ethanol is 10% more expensive than gasoline in \$/km, with $p_e/p_g = 0.7 \times 1.1$, the probability that gasoline is chosen in the price-ratio treatment is 76% compared to 74% in the control—only 2 percentage points apart. By contrast, the probability that a college-educated consumer chooses the very favorably priced fuel is 75% compared to 70% for a consumer with no more than primary education—this difference of 5 percentage points between the groups with different schooling is more than twice the price salience treatment effect.

Consider another example, this time when gasoline is 10% more expensive than ethanol, with $p_e/p_g = 0.7/1.1$. Informing consumers that ethanol is favorably priced and showing them the price-ratio flyer increases demand for ethanol by the same amount as a 2.9 percent ethanol price cut. By contrast, the share of ethanol among college educated consumers at this very favorable price point is equivalent to the share of ethanol among the less schooled only when the ethanol price is further cut by 5.9 percent!

Figure 4, left panel further compares consumer choices as one shifts from wealth at the 10th percentile to the 90th percentile of the distribution of vehicle prices. The panel shows just how less elastic (and tilted toward gasoline) is demand at the high end of the wealth distribution compared to the low end.

In the right panel of Figure 4, I illustrate just how local tastes are for the non-price characteristics of the fuel. I re-estimate the specification shown in Table 6, column IV *separately* using: (i) the subsample of choices observed in São Paulo, and (ii) the Belo Horizonte subsample. São Paulo and Belo Horizonte are the two cities with the largest numbers of sampled consumers, respectively with 2878 and 2970 observations. Using

[19] Other specifications I have estimated, but similarly do not report, allow other demographic characteristics to shift price sensitivity, such as gender, or allow interactions of demographics x_i and treatment T_i to differentially shift price sensitivity. Instead of denominating p_{ijlt} in R\$/km based on each vehicle's specific fuel economy, another robustness check specifies effective prices using the media-reported 70% conversion rate. This assumes that consumers form a price heuristic for ethanol around the comparatively stable gasoline series, taking gasoline prices at 70% of their per-liter posted prices (and ethanol prices at the per-liter price). Results for these specifications are available from the author.

each separate set of estimates, I predict demand at each counterfactual price point.

For example, when ethanol is 10% more expensive than gasoline in \$/km, ethanol choice probabilities in each city's bi-fuel vehicle population are a very different 35% in São Paulo and 17% in Belo Horizonte—18 percentage points apart. Compared to Belo Horizonte, São Paulo is clearly "pro-ethanol."[20] This follows decades of local advertising by the local sugar industry, with ethanol positioned as "the green fuel" and as "promoting local jobs" (Salvo and Huse 2013). Importantly, gasoline and ethanol were both universally available at the pump in São Paulo and in Belo Horizonte, so differences in ethanol demand are not due to differences in distribution. Specifically, in the large representative external sample of Figure 1, gasoline and ethanol were both available for over 99% of the 16,789 station-week pairs in São Paulo and over 99% of the 4,557 station-week pairs in Belo Horizonte, surveyed between May 2011 and March 2012.[21]

In sum, I conclude from the analysis that the effect of raising price salience, while statistically significant, is economically limited when compared to variation in observable individual characteristics, including education, wealth, and "home bias."

4.2 Model 2: Salience shifts a consumer's choice set

Driving a bi-fuel vehicle and sampled in station l, consumer i now chooses only among fuels that are in a random choice set $\mathbf{C}_{il} \subseteq \{g, \bar{g}, e\}$. This latent variable is the set of fuels that the consumer actually chooses from when at the pump. (See Goeree 2008, Santos et al. 2012, Dinerstein et al. 2014 and Gaynor et al. 2016 for recent modeling of "consideration sets" in the economics literature.) There are three possible sets: gasoline only, $\{g, \bar{g}\}$; ethanol only, $\{e\}$; and gasoline and ethanol, $\{g, \bar{g}, e\}$. The consumer then picks the fuel $j \in \mathbf{C}_{il}$ that yields maximal utility. The indirect utility function follows (1), with the station-fuel intercept ξ_{jl} restricted to a simple fuel fixed effect ξ_j,[22] as it is now the choice-set density that shifts by station.[23] Moreover, price sensitivity now shifts only with demographic variables x_i such as wealth, but not treatment:

$$\alpha_i = \alpha_1 + x_i' \alpha_2 \tag{4}$$

[20] Similar to the comparison between São Paulo and Belo Horizonte, Curitiba is "pro-ethanol" relative to Recife. At a price point of $p_e/p_g = 0.7 \times 1.1$, ethanol choice probabilities are 36% in Curitiba and 20% in Recife—16 percentage points apart. To draw an analogy with the United States, midwestern states such as Iowa or Minnesota are plausibly "pro-ethanol" relative to other states (Anderson 2012).

[21] Also important, beginning in 2003, adoption of bi-fuel vehicles over single-fuel ones was equally fast in São Paulo and in Belo Horizonte (Salvo and Huse 2010). This means that we are comparing the wider population of drivers in both cities, *not* a subpopulation of early and keen adopters of a product in one city to the wider population in the other.

[22] Alternatively, taste shocks can vary by fuel-city pair, ξ_{jl}. This yields qualitatively similar estimates to the baseline estimates reported in Table 7, column II. In particular, I estimate coefficients (s.e.) of 0.45 (0.00), 0.15 (0.00), 0.73 (0.11), and -61.75 (29.64) for rows CS5, CS10, CS11 and U1, respectively.

[23] Since (1) does not include a station-fuel shifter, as it did in model 1, I allow utility to shift with the average income (proxied by observed vehicle value) of the station's customer pool, i.e., the term in x_l.

I assume that consumers in the control group observed at station l (i.e., $T_{il} = 0$) are distributed as follows:

$$\mathbf{C}_{il}|\, T_{il} = 0 = \begin{cases} \{g, \bar{g}\} \text{ with probability } \lambda_l^g \\ \{e\} \text{ with probability } \lambda_l^e \\ \{g, \bar{g}, e\} \text{ with probability } 1 - (\lambda_l^g + \lambda_l^e) \end{cases}$$

where parameters λ_l, $\mathbf{0} \leq \lambda_l = (\lambda_l^g, \lambda_l^e) \leq \mathbf{1}$, are allowed to shift with location. As do the unobserved taste shocks ξ_{jl} in model 1, single-fuel choice-set probabilities λ_l capture perceived fuel quality or home bias for gasoline or ethanol among residents (or commuters) in the micro-region around a station or the wider city—now to be interpreted by way of a single-fuel versus bi-fuel mindset rather than a shock entering the utility specification. These λ_l parameters to be estimated are assumed to be constant over a station's consumers in the 11-month sample. The identifying restriction is that the composition of a station's consumers, namely those who consider purchasing only a single fuel despite their bi-fuel vehicle capability, does not vary over time. In particular, densities λ_l are longer-term attributes and do not vary with prices (within a relevant range). For comparison, Goeree (2008) writes probabilities on alternative products entering the choice set, not directly on alternative choice sets.

In contrast, random choice sets for consumers at station l who are treated with salient relative price information at the pump ($T_{il} = 1$) are distributed as follows:

$$\mathbf{C}_{il}|\, T_{il} = 1 = \begin{cases} \{g, \bar{g}\} \text{ with probability } \phi_l \lambda_l^g \\ \{e\} \text{ with probability } \phi_l \lambda_l^e \\ \{g, \bar{g}, e\} \text{ with probability } 1 - \phi_l\, (\lambda_l^g + \lambda_l^e) \end{cases}$$

Thus, price salience treatment—whether with the price-ratio or the km-per-R$50 flyers—can reduce the proportion of "single-fuel minded" consumers, in proportion to λ_l^g and λ_l^e, by $1 - \phi_l$, where $0 \leq \phi_l \leq 1$. In other words, the treatment parameter ϕ_l captures the proportionate shift in density from limited gasoline or ethanol choice sets to a bi-fuel choice set upon treatment. To reduce the number of parameters to be estimated, I restrict ϕ_l to vary by city rather than by station and to be equal across treatments. (Alternatively, I could allow λ_l to shift only by city and ϕ_l to shift by station.)

Intuitively, the average choices over time at a station help pin down the choice-set densities λ_l, whereas the variation in within-station choices as relative prices vary over time pin down both the λ_l and the price sensitivity parameters (α_1, α_2). Any increased substitution across fuels in the price salience treatments relative to the control is picked up by the salience (or information) parameter ϕ_l. Thus, in model 2, intervention can expand the consumer's choice set but it does not impact utility directly, as it did in model 1, in which intervention entered (1) both in level and interacted with price. In model 2, salience shifts a consumer's choice set in a stage prior to the choice stage; in particular, it does not shift a consumer's sensitivity to price differences.

I also estimate model 2 by maximum likelihood. In terms of notation, collect the parameters to be estimated in vectors $\theta_2 = (\alpha_1, \alpha_2, \beta_{1j}, \beta_{2j}, \beta_{3j}, \xi_j)$ and $\theta = (\lambda_l, \phi_l, \theta_2)$, and solve:

$$\begin{aligned} &\arg\max_{\theta} \log L\left(data|\,\theta\right), \\ \text{subject to } \lambda_l^g &\geq 0, \quad \lambda_l^e \geq 0, \quad \lambda_l^g + \lambda_l^e \leq 1, \quad 0 \leq \phi_l \leq 1 \quad \forall l \end{aligned}$$

with

$$L\left(data|\,\theta\right) = L\left(control\ group\ choices|\,\theta\right) L\left(treatment\ group\ choices|\,\theta\right).$$

The first factor of the likelihood, pertaining to observed choices in the control group, is:

$$\prod\nolimits_{i\in\{T_{il}=0\}(,l,t)} \left(\begin{array}{c} \lambda_l^g 1\left(i \text{ chooses } g \text{ or } \bar{g}\right) \prod\nolimits_{j\in\{g,\bar{g}\}} \left(s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g}\}\right)\right)^{1(i \text{ chooses } j)} \\ +\lambda_l^e 1\left(i \text{ chooses } e\right) \\ +\left(1-\left(\lambda_l^g+\lambda_l^e\right)\right) \prod\nolimits_{j\in\{g,\bar{g},e\}} \left(s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g},e\}\right)\right)^{1(i \text{ chooses } j)} \end{array} \right)$$

noting, in order, that

$$s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g}\}\right) = \frac{\delta_{jl}(\theta_2) + \mu_{ijlt}(\theta_2)}{\sum_{k\in\{g,\bar{g}\}}\left(\delta_{kl}(\theta_2) + \mu_{iklt}(\theta_2)\right)},$$

that a consumer with a single-fuel ethanol choice set necessarily chooses ethanol at the pump, and that $s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g},e\}\right)$ is given by (3). Similarly, the second factor in the likelihood pertains to observed choices in the treatment groups:

$$\prod\nolimits_{i\in\{T_{il}=1\}(,l,t)} \left(\begin{array}{c} \phi_l\lambda_l^g 1\left(i \text{ chooses } g \text{ or } \bar{g}\right) \prod\nolimits_{j\in\{g,\bar{g}\}} \left(s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g}\}\right)\right)^{1(i \text{ chooses } j)} \\ +\phi_l\lambda_l^e 1\left(i \text{ chooses } e\right) \\ +\left(1-\phi_l\left(\lambda_l^g+\lambda_l^e\right)\right) \prod\nolimits_{j\in\{g,\bar{g},e\}} \left(s_{ijlt}\left(\theta_2|\,\mathbf{C}_{il} = \{g,\bar{g},e\}\right)\right)^{1(i \text{ chooses } j)} \end{array} \right)$$

Results. Table 7 reports estimates of the salience-shifts-choice-set model, using the full sample of choices. Columns I and II report estimates for alternative specifications in which choice-set densities $(\lambda_l^g, \lambda_l^e)$ shift, respectively, (i) by city ($2 \times 4 = 8$ parameters), and (ii) by station ($2 \times 52 = 104$ parameters).

The estimated utility function parameters (rows U1 to U23) are similar to those estimated under model 1 in terms of sign and statistical significance. For example, I again estimate price sensitivity to increase in schooling (rows U3-4) and decrease in wealth (row U2).[24] Moreover, price sensitivity is estimated to be higher in model

[24] Further, as in model 1, a vehicle's price is associated with a stronger taste for gasoline, female consumers favor regular-grade gasoline, and tastes for midgrade gasoline grow with age.

2 compared to model 1 (row U1, also see below). This is intuitive, since the gradual aggregate substitution across gasoline and ethanol that we observe in the data as relative prices vary is now also explained by an inclination for consumers to be single-fuel minded at the point of purchase. This frees the consumer's price sensitivity α_i, on choosing from a bi-fuel choice set, to grow in magnitude.

The estimated choice-set parameters are informative. In column I, the proportion of single-fuel gasoline-minded consumers exceeds the proportion of single-fuel ethanol-minded consumers in all cities, but the difference in favor of gasoline is lowest in comparatively pro-ethanol São Paulo and Curitiba, where $\lambda_l^g - \lambda_l^e$ is estimated at 0.11 and 0.10, respectively, compared to 0.64 and 0.37 in Belo Horizonte and Recife. Among the population of flexible fuel vehicle drivers, the density of bi-fuel minded—"flexible"—consumers, i.e., $1 - \hat{\lambda}_l^g - \hat{\lambda}_l^e$, is lowest in Belo Horizonte, at 0.29, compared to 0.69 in Curitiba, 0.54 in Recife and 0.41 in São Paulo. Roughly speaking, model 2 suggests that only about *one-half* the population of consumers equipped with bi-fuel vehicle technology actually considers bi-fuel choice sets when making fuel purchases!

In addition, the estimated treatment parameter comes out statistically significantly lower than 1 (the no-effect baseline) in the case of São Paulo, at $\hat{\phi}_{\text{SaoPaulo}} = 0.78$, in both columns I (s.e. 0.11) and II (s.e. 0.09) (see row CS11). This suggests that, viewed through the lens of this salience-shifts-choice-set model, treatment lowers the mass of single-fuel minded São Paulo consumers by about one-fifth, "freeing" them to purchase fuel from a bi-fuel choice set. Across columns I and II, I do not obtain robust statistically significant evidence of a shift in mindset in other locations ($\hat{\phi}_{\text{BeloHorizonte}}$ is statistically significantly less than 1 in column II but not in column I).

Counterfactual experiments. To illustrate the salience-shifts-choice-set model estimates, Figure 5 reports, for the different counterfactual price points, predicted ethanol shares among bi-fuel vehicle drivers fueling at two different stations in the sample. (I use the more flexible specification reported in column II of Table 7.) The left panel considers a given station in São Paulo, with estimated proportions of single-fuel minded gasoline and ethanol consumers, respectively, of $\hat{\lambda}_l^g = 0.41$ (s.e. 0.08) and $\hat{\lambda}_l^e = 0.38$ (s.e. 0.15). The anticlockwise rotation from the solid (blue) curve to the dashed (blue) curve indicates the estimated treatment effect, $\hat{\phi}_{\text{SaoPaulo}} = 0.78$ (s.e. 0.09), by which treatment shifts the mass of single-fuel to bi-fuel choice sets by 22 percentage points. The elastic curve indicated by the dotted (green) curve indicates the predicted ethanol shares in the counterfactual scenario that each consumer is endowed with a bi-fuel choice set with probability 1.[25]

By contrast, the right panel illustrates the model's prediction for ethanol demand at

[25] Again by analogy (see note 8), this thought experiment is similar to that in Dinerstein et al. (2014) in which all information frictions are eliminated and "consumers are shown the entire set of (relevant) listings available on the platform" (p.18).

a station sampled in Belo Horizonte. Here, proportions of single-fuel minded gasoline and ethanol consumers are estimated at very different $\hat{\lambda}_l^g = 0.66$ (s.e. 0.24) and $\hat{\lambda}_l^e = 0.15$ (s.e. 0.09), respectively, recalling that the point estimate for the treatment parameter is $\hat{\phi}_{\text{BeloHorizonte}} = 0.75$ (s.e. 0.10). Again, the model suggests that the extensive margin of fuel demand—consumers choosing gasoline or ethanol—would be much more elastic were all bi-fuel vehicle motorists to effectively *choose* from a bi-fuel choice set.

Consumer welfare. The information intervention in this second model impacts expected welfare insofar as it shifts the probability distribution over the possible choice sets, and thus the distribution over the possible choice realizations. It does not impact expected utility conditional on a choice set. For this reason, I use the model to perform a welfare calculation of the estimated treatment effect. The distributional assumption on the idiosyncratic taste shock ε_{ijlt} allows one to compute expected utility analytically. For consumer i in market lt subject to treatment, this is:

$$\begin{cases} \ln\left(\sum_{j\in\{g,\bar{g}\}} \exp\left(\delta_{jl}(\theta_2) + \mu_{ijlt}(\theta_2)\right)\right) \text{ with probability } \phi_l\lambda_l^g \\ \ln\left(\exp\left(\delta_{el}(\theta_2) + \mu_{ielt}(\theta_2)\right)\right) \text{ with probability } \phi_l\lambda_l^e \\ \ln\left(\sum_{j\in\{g,\bar{g},e\}} \exp\left(\delta_{jl}(\theta_2) + \mu_{ijlt}(\theta_2)\right)\right) \text{ with probability } 1-\phi_l\left(\lambda_l^g+\lambda_l^e\right) \end{cases}$$

A similar expression follows for a consumer's expected utility in the absence of treatment, replacing the treatment parameter ϕ_l by 1.

For each one of the bi-fuel vehicle drivers sampled in São Paulo, I use the estimated model to compute the expected gain in welfare from raising the salience of prices—treatment versus no treatment, with $\hat{\phi}_{\text{SaoPaulo}} = 0.78$—at a given price point. I then average across consumers. Rather than express the gain in utils, I compute the "equivalent variation" of the aggregate welfare gain as a welfare-equivalent proportionate price reduction for all fuels, i.e., a drop in the general fuel price level.[26] I repeat this exercise for each different price point, ranging from ethanol very favorably priced to gasoline very favorably priced. Results are reported in the left panel of Figure 6. For example, p_e/p_g in São Paulo hovered around 0.55 throughout 2009 (Salvo and Geiger 2014). At this price point, raising price salience would enhance consumer welfare by an equivalent magnitude to cutting fuel prices—both gasoline *and* ethanol—by 2.2%.

[26] Formally, fixing a price point by the vector $p = (p_g, p_{\bar{g}}, p_e)$ and denoting the welfare-equivalent proportionate fuel price discount by ρ, I search for the scalar ρ that solves:

$$E_i\left[\max_j u_{ij}\,|T_i = 1, p\right] = E_i\left[\max_j u_{ij}\,|T_i = 0, (1-\rho)\,p\right]$$

I keep $(p_g, p_{\bar{g}})$ at their sample values and vary the price point by picking r such that $p_e = rp_g, r \in [0.5, 0.9]$. The exercise ignores welfare effects from consumers responding along the intensive margin (e.g., driving more) and firms responding to demand shifts by changing prices. Such responses may magnify or attenuate consumer welfare gains.

The right panel reports results averaged across Belo Horizonte's sampled consumers, with $\hat{\phi}_{\text{BeloHorizonte}} = 0.75$. Throughout the 11-month sample period, p_e/p_g in Belo Horizonte hovered around 0.78. At this price point, raising price salience or cutting all fuel prices by about 0.5% would have had an equivalent effect on consumer welfare. When gasoline is favorably priced, the model predicts lower welfare gains from raising salience in Belo Horizonte because this is where the density of consumers without gasoline in their choice sets is estimated to be lowest, e.g., median $\hat{\lambda}_l^e$ across Belo Horizonte's 14 stations is 0.04, compared to 0.31 across São Paulo's 14 stations. By the same token, predicted welfare gains from salience would be high were Belo Horizonte bi-fuel vehicle drivers faced with favorably priced ethanol, given the relative absence of ethanol from their choice sets.

5 Conclusion and policy discussion

This paper addresses a puzzle: why do many of Brazil's energy consumers driving flexible fuel vehicles not realize energy savings at the pump, by choosing the fuel that yields the highest distance traveled by dollar of spending? The savings can be significant. Consider a typical example. Noting that most drivers do not fill up at the pump, one-half of a tank of gasoline is worth US$ 34 and 240 km of urban driving. Faced with ethanol priced at a 20% discount relative to gasoline in \$/km, 240 km can be purchased for \$7 less, yet about 20% of bi-fuel vehicle owners choose gasoline.[27] That is, such consumers lack the "flexibility" that their vehicles come equipped with. It is worth noting that the energy savings are available *today*, not in the future as in most of the housing and durables stock settings examined in the "energy efficiency gap" literature. The question is relevant in the context of a literature that pushes for policies such as a carbon tax (or against policies such as efficiency standards) under the assumption that drivers will make optimal fuel purchase decisions if given the correct incentives.

One possibility is that price differences across gasoline and ethanol are not salient at the pump. In phone interviews with bi-fuel vehicle owners, Salvo and Huse (2013) found that the vast majority of consumers were aware that gasoline and ethanol exhibit different energy contents per volume unit in which prices are posted at the pump. However, it may be that consumers underweigh the price attribute when pulling up to refuel (Bordalo et al. 2013). A related possibility is that, since price differences are not salient, some consumers only consider one type of fuel when purchasing energy for their bi-fuel vehicles.

This paper studies the effect of truthfully informing consumers, both verbally and by way of a flyer, of the fuel that is favorably priced at the point of sale as soon as

[27] The example considers gasoline priced at R$ 2.80/liter (US$ 1.40/liter), $k_g = 9.8$ km/liter (E25) and a tank capacity of 48 liters, such as the Fiat Palio ELX 2010 with a 1.0-liter engine.

they pull up at the pump. Among different experiments, the largest treatment effect I obtained was to shift one-tenth of consumers who would have counterfactually chosen expensive gasoline, to choose very favorably priced ethanol instead. This shift, while statistically significant, is small relative to a comparison of choices between observably different consumer types, such as consumers with no more than primary schooling versus their college educated counterparts—I show that the latter are more likely to pick the favorably priced fuel, consistent with more weight placed on price differences than on perceived or real differences in non-price attributes.

One of the models that I estimate suggests that a similar policy to raise the salience of effective prices could raise consumer welfare by the equivalent of a 1 to 3% drop in (all) motor fuel prices. For example, policymakers can mandate that more comparable gasoline and ethanol prices *per km* be posted at the pump for representative consumer vehicles, such as 1, 1.4 and 1.8 liter engines. The effect of such a long-run intervention might exceed the short-run effect that I am able to estimate under the design I used. Moreover, such a policy may have larger effects in markets, such as the US, where price salience is possibly lower than among consumers who have been exposed to ethanol for decades.

While the effect of raising price salience is economically meaningful to some extent, the considerable mass of consumers who still purchase the less favorably priced fuel when reminded of effective prices indicates a significant role for heterogeneous preferences for fuels' non-price characteristics, including perceived or real drivability, maintenance and external impacts, as well as long-run habit formation. This finding matters to the design of policies to displace consumption of oil derivatives as the motor fuels of choice. Besides promoting alternative technologies on the supply side, preferences matter even where one might think consumers care only about prices.

References

[1] Allcott, H. and M. Greenstone (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives* 26(1), 3-28

[2] Allcott, H. and S. Mullainathan (2010). Behavior and energy policy. *Science* 327(5970), 1204-1205

[3] Allcott, H. and R. Sweeney (2015). Can retailers inform consumers about energy costs? Evidence from a field experiment. E2e working paper 8

[4] Anderson, S. T. (2012). The demand for ethanol as a gasoline substitute. *Journal of Environmental Economics and Management* 63(2), 151-168

[5] Anderson, S. T. and J. M. Sallee (2011). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review* 101(4), 1375-1409

[6] Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics* 25(2), 242-262

[7] Bordalo, P., N. Gennaioli and A. Shleifer (2013). Salience and consumer choice. *Journal of Political Economy* 121(5), 803-843

[8] Bronnenberg, B. J., J.-P. Dubé, M. Gentzkow and J. M. Shapiro (2015). Do pharmacists buy Bayer? Informed shoppers and the brand premium. *Quarterly Journal of Economics* 130(4), 1669-1726

[9] Busse, M. R., N. Lacetera, D. G. Pope, J. Silva-Risso and J. R. Sydnor (2013). Estimating the effect of salience in wholesale and retail car markets. *American Economic Review* 103(3), 575-579

[10] Chetty, R., A. Looney and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review* 99(4), 1145-1177

[11] Clerides, S. and P. Courty (2016). Sales, quantity surcharge and consumer inattention. *Review of Economics and Statistics*, forthcoming

[12] Choi, J. J., D. Laibson and B. C. Madrian (2011). $100 bills on the sidewalk: Suboptimal investment in 401(k) plans. *Review of Economics and Statistics* 93(3), 748-763

[13] DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315-372

[14] Dinerstein, M., L. Einav, J. Levin and N. Sundaresan (2014). Consumer price search and platform design in internet commerce. Stanford University working paper

[15] Dranove, D. and G. Z. Jin (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature* 48(4), 935-963

[16] Fowlie, M., M. Greenstone, and C. Wolfram (2015). Are the non-monetary costs of energy efficiency investments large? Understanding low take-up of a free energy efficiency program. *American Economic Review* 105(5), 201-204

[17] Gaynor, M., C. Propper, and S. Seiler (2016). Free to choose? Reform, choice, and consideration sets in the English National Health Service. *American Economic Review* 106(11), 3521-3557

[18] Gillingham, K. and K. Palmer (2014). Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy* 8(1), 18-38

[19] Goeree, M. S. (2008). Limited information and advertising in the US personal computer industry. *Econometrica* 76(5), 1017-1074

[20] Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103(7), 2643-2682

[21] Hastings, J. S. and J. M. Shapiro (2013). Fungibility and consumer choice: Evidence from commodity price shocks. *Quarterly Journal of Economics* 128(4), 1449-1498

[22] Holland, S. P., J. E. Hughes and C. R. Knittel (2009). Greenhouse gas reductions under low carbon fuel standards? *American Economic Journal: Economic Policy* 1(1), 106-146

[23] Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review* 102(5), 2147-2182

[24] Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review* 104(4), 1417-1438

[25] Kahneman, D., J. L. Knetsch and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 76(4), 728-741

[26] Kling, J. R., S. Mullainathan, E. Shafir, L. C. Vermeulen and M. V. Wrobel (2012). Comparison friction: Experimental evidence from Medicare drug plans. *Quarterly Journal of Economics* 127(1), 199-235

[27] Lacetera, N., D. G. Pope and J. R. Sydnor (2012). Heuristic thinking and limited attention in the car market. *American Economic Review* 102(5), 2206-2236

[28] Lusardi, A. and O. S. Mitchell (2008). Planning and financial literacy: How do women fare? *American Economic Review* 98(2), 413-417

[29] Pires, T. (2016). Costly search and consideration sets in storable goods markets. *Quantitative Marketing and Economics*, forthcoming

[30] Salvo, A. and F. M. Geiger (2014). Reduction in local ozone levels in urban São Paulo due to a shift from ethanol to gasoline use. *Nature Geoscience* 7(6), 450-458

[31] Salvo, A. and C. Huse (2010). Is arbitrage tying the price of ethanol to that of gasoline? Evidence from the uptake of flexible-fuel technology. *Energy Journal* 32(3), 119-148

[32] Salvo, A. and C. Huse (2013). Build it, but will they come? Evidence from consumer choice between gasoline and sugarcane ethanol. *Journal of Environmental Economics and Management* 66, 251-279

[33] Santos, B. D. L., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review* 102(6), 2955-2980

[34] Valor Econômico (2012). Crise de confiança afeta consumo de etanol. (Crisis of confidence affects ethanol consumption. In Portuguese).

A Appendix

A.1 Data

Survey. For details on CNPBrasil, the market research firm hired to interact with consumers by the pump, I refer the reader to Salvo and Huse (2013, Appendix A), which includes a discussion of CNPBrasil's "internal control" procedures. Gaining access to fueling stations to carry out the current project turned out to be substantially more complex, due to its experimental nature (see note 12). Interacting with a seller's buyers before they place their order and the sale is consummated is bound to be more sensitive commercially compared with passively observing consumers make their purchases, as in Salvo and Huse (2013). In addition, the current project required the accurate handling of price information. It involved a significantly larger number of subjects per station visit—54 consumers compared to 12 in Salvo and Huse (2013). The duration of a visit (including breaks) averaged 9.4 hours versus 2.5 hours in Salvo and Huse (2013, Table A1). On the other hand, bi-fuel vehicles accounted for an even greater share of the light vehicle stock during the current campaign's 2011/2012 survey period—likely over 50% of vehicle-miles traveled—compared to the 2010 survey in Salvo and Huse (2013). This meant that enumerators would wait less time for each qualifying consumer to arrive. As in Salvo and Huse (2013), a consumer would qualify for the subject pool if he drove an originally manufactured bi-fuel vehicle on his private, not company nor taxi, business.

Besides obtaining authorization from management at Shell-Cosan, all the way to its chief executive officer, securing consumer access at stations required my constant intervention directly with Shell-Cosan sales executives, as well as station managers and franchisees, during the entire execution of fieldwork.[28] This, in turn, ensured my continued monitoring of fieldwork, with communication between myself and CNPBrasil almost on a daily basis. This communication included the scanning and emailing of material which was used, such as flyers, and completed, such as survey forms, during fieldwork. Examples of such back-and-forth communication are available on demand. On a few occasions, CNPBrasil reported that station visits had been interrupted, and observations discarded, because, for example, the station manager had called off the visit or one of the fuels had run out. Further, in the final sample I chose to drop eight station visits for which the pattern of fuel choices by consumers seems unusual yet the enumerator failed to report any irregularity such as the station running out of one fuel. For example, consider one such visit, stationid 212053 in São Paulo, visited on June 3, 2011 when ethanol was very favorably priced relative to gasoline, namely, $p_e/p_g \simeq 0.57$. Auditing flagged that while as many as 15 out of 18 consumers in the first group (control) had chosen ethanol, only one consumer in the subsequent treatment group did so, and no consumer in the final treatment group of the visit chose ethanol, strongly suggesting that the station had run out of ethanol soon after the control stage was completed, yet the enumerator failed to record the occurrence and abort the visit. The results reported here are robust to adding these eight station visits, comprising 4% of the original sample, back into the analysis.

Vehicle-specific fuel economy. I use the National Institute for Meteorology's (Inmetro) laboratory measurements to predict the fuel economy of all bi-fuel vehicles sampled at stations. For 230 model-engine-year combinations from 2009 to 2012, and following U.S. EPA guidelines, Inmetro reports the kilometerage per liter, k, under separate "urban" and "highway" driving cycles, when alternatively burning gasoline E22 and ethanol E100, denoted by subscripts g and e, respectively. Since lab measures are available for only a subset of all vehicle model-engine-year combinations (e.g., GM Celta 1.0-liter 2009) sampled in the field, I first use the available data and project k on other vehicle characteristics; I then use the fitted regression and predict k for each vehicle

[28] I am deeply grateful to the numerous managers, sales personnel and franchisees who worked to grant me access to consumers, free of charge and with the sole purpose of enabling research into energy consumption.

sampled at stations based on observables. Table 9 reports regressions of three vehicle-specific fuel economy measures, k_e^{urban}, k_g^{urban}, and k_e^{urban}/k_g^{urban}, in columns I to III, respectively (Salvo and Huse 2013 provide further details, including regressions of fuel economy under highway driving). Each dependent variable is projected on two sets of independent variables: specification (a) in the upper part of the table or specification (b) in the lower part of the table.

When subsequently predicting the fuel economy of vehicles sampled in the field, fitted regression (a), which includes model fixed effects, is used for vehicle models for which there are lab measurements and thus a model fixed effect. This is the case for 80% of bi-fuel vehicles sampled in the field. Alternatively, fitted regression (b), which includes fixed effects for vehicle make and for vehicle class, is used to predict the fuel economy for the remaining vehicle models that were not tested in the lab.

For perspective, the interdecile range (90th percentile minus 10th percentile) for the fitted value of the ratio k_e^{urban}/k_g^{urban} (under specification (a) and for the lab observations) is 0.035, i.e., $0.691-0.656$, suggesting that the media-reported parity threshold of 70% is slightly on the high side, perhaps due to rounding. The interdecile range for specification (b) is similar, at $0.688 - 0.664 = 0.024$.

Expected distances to be traveled on R$ 50 of fuel for different engine sizes. Relative price information conveyed in the second price-salience treatment's flyers was based not only on per-liter fuel prices at the pump on the day, but also on fuel economy for three standard vehicle engine sizes, namely 1.0, 1.4 and 1.8 liter engines (Figure 2B). These engine sizes were by far the mostly commonly observed in circulation. I took kilometerage per liter, k, under the urban cycle for the 67 bi-fuel vehicles (model-engine-year combinations), powered either by gasoline E22 or ethanol E100, as included in Inmetro's 2009 and 2010 reports. I first adjusted k_g^{urban} from gasoline E22, as tested in the lab, to either gasoline E25 or gasoline E20, depending on the gasoline blend that was being sold during the course of the experiments. For every one of the 67 tested vehicles and each of the two retailed fuels, gasoline (E25 or E20) and ethanol E100, dividing R$ 50 by the corresponding price per liter, then multiplying by k, provides a measure of the distance traveled by the vehicle on a R$50 purchase of the fuel. The distances shown on the flyer were then the means across vehicles by engine category. The price-ratio flyer printed a "thumbs up" for gasoline when the ratio of gasoline km to ethanol km exceeded 1.011 for each of the 1.0, 1.4 and 1.8 liter engines. Similarly, a thumbs-up for ethanol was printed when the ratio of gasoline km to ethanol km was lower than 0.989 for each of the three engine sizes. Cases in which (at least one) ratio approached 1 stated that both fuels offered similar yields. Such calculations and messaging, specific to each station visit, were automated on a print-ready calculation-protected spreadsheet.

Other vehicle characteristics, including values. Characteristics for sampled vehicles were obtained, based on the recorded make, model, engine size and model-year, from the following online sources: (i) (secondary market prices) *Tabela de Preços Fipe/Quatro Rodas*, and (ii) (tank capacities) *Carrosnaweb* as well as manufacturer websites.[29] I collected used vehicle prices for a cross-section, the August 2012 secondary market. Since experiments took place over almost one year, I adjusted for the fact that a vehicle of a given make, model, engine and model-year sampled on an early date was worth more

[29] Some websites are http://quatrorodas.abril.com.br/tabela-de-precos, and http://www.carrosnaweb.com.br/.

than the same vehicle make, model, engine and model-year sampled on a later date. Take, for example, a GM Celta Life 1.0 liter model-year 2009: the value of this vehicle was higher if sampled on May 30, 2011 than if sampled (say) on May 30, 2012, after an additional year of depreciation. Used vehicle prices in the 2012 cross-section suggested a mean annual depreciation factor of 7% within make-model-engine (so, for the GM Celta Life 1.0 liter, a model-year 2010 would, on average, trade at 93% the price of a model-year 2011). To stay with the example, I then used this empirically determined depreciation factor of 7% to upward adjust the value of the GM Celta model-year 2009 if observed on May 30, 2011 compared to when observed on May 30, 2012.

The recorded vehicle make, model, engine size and model-year serve as a consistency check that enumerators did indeed sample only from the population of bi-fuel vehicles, and not from the pool of single-fuel vehicles. Single-fuel vehicles—gasoline-only and, to a lesser extent, ethanol-only—were sold primarily prior to 2005, but they include some premium imported vehicles with low market penetration (Salvo and Huse 2010). Reassuringly, very few vehicles as recorded in the data sample had not been commercialized in the flex version, according to *Carrosnaweb* and manufacturer websites. It is also reassuring that the mean value of vehicles sampled at stations that carried midgrade ethanol (a plausible proxy for the affluence of a station's customers, as in Salvo and Huse 2013) is statistically significantly higher than the mean value of vehicles recorded at stations where midgrade ethanol was not offered (respectively R$ 30,408 against R$ 28,751; the p-value for a test of equality of means is 0.000).

Similarly, it is reassuring that for very few observations did the recorded purchase size in liters exceed the nominal tank capacity corresponding to the recorded vehicle. In fact, in line with note 14, observed consumers tended to purchase a volume equivalent to only a fraction of their tank capacity. The median "fraction of tank purchased" across sampled consumers is 35%, and the 75th and 90th percentiles are, respectively, 57% and 79% of tank capacity. The convenience of station attendants who are readily available to fuel vehicles, with queues rarely forming at the pump, and the preference for paying in cash, help explain this market feature.

A.2 Experiments where both fuels were priced similarly

Ethanol and gasoline were priced similarly in km equivalent units mostly during some first-wave visits in São Paulo, Curitiba and Belo Horizonte, predominantly with station attendants making effective prices salient to treated subjects (Table 1, last column). Namely, in these experiments, $0.7/1.05 \leq p_e/p_g < 0.7 \times 1.05$. As Table 8 indicates, a low 33% of subjects in the control group chose ethanol over gasoline, with the proportion of ethanol adopters rising to 38% in the treatment groups, a difference that is marginally statistically significant (in part due to the test's low power). It is important to recall that in such markets, ethanol prices had just fallen sharply back to parity with gasoline (Figures 1A and 1B), and some consumers may have been slow to substitute out of gasoline and into ethanol.

A.3 Descriptive plots of fuel choices

Fuel choices are depicted in the four panels of Figure A1, at all the different price points (vertical axes) and the distinct recent price histories (left versus right panels) in the sample, separately by control and treatment groups (top versus bottom panels). Each panel also indicates the best linear predictor for the relative ethanol price, p_e/p_g, against the proportion of subjects who chose ethanol in each control or treatment group of 18. An observation in these panels is a station visit and group pair. Demand slopes downward, yet there is substantial heterogeneity across station visits (within treatment type). The structural models I estimate account for this by specifying unobservable product-retailer taste (model 1) or choice set (model 2) shifters.

Any effect of treating subjects with price information—in the format I experiment with in this study—is not visible in Figure A1: demand in the bottom panels (treatment) does not seem more elastic around parity compared to demand in the top panels (control). In the left panels, depicting first-wave experiments conducted in São Paulo, Curitiba and Belo Horizonte in the wake of a sharp drop in ethanol prices, ethanol choice probabilities tend to be less than 0.5 for p_e/p_g close to 70%. Compared to choices observed months later (right panels) as well as a year earlier (Salvo and Huse 2013, Figure 4), this may reflect short-run dynamics.

Anecdotal evidence both from the press and from my own interviews with consumers at the pump suggest that some consumers may have been slow to "return" to ethanol in 2011, following a second and more pronounced price spike in two years, with the aim to "punish" the industry (in the spirit of Kahneman et al. 1986). For example, Valor Econômico (2012) spoke of "a lack of 'patience' by the consumer (with ethanol)" and cited the head of the Brazilian fuel retail trade association, Paulo Miranda Soares, who argued that "instability in the ethanol market is perceived by the consumer, who appears to have developed a 'stubbornness' (*birra*) with respect to the biofuel."

Figure A2 plots average treatment effects by station visit, for each of the three price informant-wave combinations implemented in the field (in the columns), separately for the price-ratio and km-per-R$ 50 treatments (in the rows). If treatment effects were large, points would concentrate in the second and fourth quadrants of each panel, i.e., less (resp., more) choice of ethanol over gasoline for treated relative to control for markets with high (resp., low) ethanol prices. There is a somewhat greater mass of points in the second and fourth quadrants compared to the first and third quadrants, but this is not a strong effect.

p_e/p_g at the pump:	$< 0.7/1.1$	$0.7/1.1$ to $0.7/1.05$	$\geq 0.7 \times 1.1$	0.7×1.05 to 0.7×1.1	$0.7/1.05$ to 0.7×1.05
Favorably priced fuel (rel. to substitute)	Ethanol	Ethanol	Gasoline	Gasoline	$\sim$ Similar
How favorably priced?	"Very"	"Somewhat"	"Very"	"Somewhat"	
	[**EE**]	[**E**]	[**GG**]	[**G**]	
Attendant informing prices, Wave 1					
São Paulo	1296	702	0	54	270
Curitiba	162	702	0	0	594
Belo Horizonte	0	0	540	972	270
Recife	0	0	1026	162	0
Attendant informing prices, Wave 2					
São Paulo	–	–	–	–	–
Curitiba	0	0	270	54	54
Belo Horizonte	0	0	594	108	54
Recife	0	0	324	0	0
Enumerator informing prices, Wave 2					
São Paulo	54	108	0	0	486
Curitiba	0	0	108	108	0
Belo Horizonte	0	0	432	0	0
Recife	0	0	756	108	54
All price informant-wave combinations					
São Paulo	1350	810	0	54	756
Curitiba	162	702	378	162	648
Belo Horizonte	0	0	1566	1080	324
Recife	0	0	2106	270	54
Total	1512	1512	4050	1566	1782

Table 1: Relative price variation at the pump for each of three price informant by wave field combinations, by city. Entries in the table are the number of subjects at the different relative price points (per-liter ethanol-to-gasoline, regular grade). The price informant is the person (a station attendant or an enumerator) responsible for treating the subject, namely greeting the subject as he pulled up at the station and offering verbal and printed relative price information. Wave 1 lasted between May and July 2011. Wave 2, following preliminary analysis of wave 1 and repeated instruction of enumerators, lasted between September 2011 and March 2012.

	Control		Price-ratio flyer			Km-per-R$50 flyer		
	Obs	Mean (S.D.)	Obs	Mean (S.D.)	Diff Mean (S.E.)	Obs	Mean (S.D.)	Diff Mean (S.E.)
Female (1=yes)	3474	0.340 (0.474)	3474	0.345 (0.475)	0.004 (0.011)	3474	0.347 (0.476)	0.007 (0.011)
Aged less than 25 years (1=yes)	3474	0.124 (0.330)	3474	0.128 (0.335)	0.004 (0.008)	3474	0.125 (0.331)	0.001 (0.008)
Aged 25 to 40 years (1=yes)	3474	0.458 (0.498)	3474	0.469 (0.499)	0.011 (0.012)	3474	0.479 (0.500)	0.021* (0.012)
Aged 40 to 65 years (1=yes)	3474	0.372 (0.484)	3474	0.368 (0.482)	-0.005 (0.012)	3474	0.359 (0.480)	-0.013 (0.012)
Aged more than 65 years (1=yes)	3474	0.045 (0.208)	3474	0.035 (0.185)	-0.010** (0.005)	3474	0.037 (0.188)	-0.008* (0.005)
At most secondary school, incomplete (1=yes)	3474	0.070 (0.255)	3474	0.071 (0.257)	0.001 (0.006)	3474	0.063 (0.244)	-0.006 (0.006)
Secondary school, completed (1=yes)	3474	0.241 (0.428)	3474	0.244 (0.429)	0.003 (0.010)	3474	0.246 (0.430)	0.004 (0.010)
College, incomplete (1=yes)	3474	0.144 (0.351)	3474	0.133 (0.340)	-0.011 (0.008)	3474	0.144 (0.351)	0.000 (0.008)
College, completed (1=yes)	3474	0.545 (0.498)	3474	0.552 (0.497)	0.007 (0.012)	3474	0.547 (0.498)	0.002 (0.012)
Vehicle price (R$ $\times$ 1000)	3474	29.048 (9.502)	3474	28.975 (9.236)	-0.073 (0.225)	3474	29.106 (9.901)	0.058 (0.233)
Vehicle engine size (liters)	3474	1.355 (0.346)	3474	1.358 (0.347)	0.003 (0.008)	3474	1.356 (0.365)	0.001 (0.009)
Vehicle make is Fiat (1=yes)	3474	0.292 (0.455)	3474	0.298 (0.457)	0.006 (0.011)	3474	0.298 (0.458)	0.007 (0.011)
Vehicle make is GM (1=yes)	3474	0.202 (0.401)	3474	0.203 (0.402)	0.001 (0.010)	3474	0.197 (0.397)	-0.006 (0.010)
Vehicle make is Volkswagen (1=yes)	3474	0.215 (0.411)	3474	0.221 (0.415)	0.006 (0.010)	3474	0.199 (0.400)	-0.016 (0.010)
Vehicle make is Ford (1=yes)	3474	0.102 (0.302)	3474	0.096 (0.295)	-0.005 (0.007)	3474	0.104 (0.306)	0.003 (0.007)
Number of subjects declining to participate	193	3.332 (3.731)	193	2.674 (3.154)	-0.658* (0.352)	193	2.482 (3.315)	-0.850** (0.359)

Table 2: Summary statistics by control and treatment group. Number of observations, mean and standard deviation (S.D.) by control and treatment group. Difference in means and its standard error (S.E.) between each treatment group and control. Educational attainment is stated rather than observed. Vehicle price is based on vehicle characteristics (see the appendix). In all but the last row, an observation is a subject. In the last row, an observation is a group (control or treatment) within a station visit. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Ethanol very/somewhat favorably priced	Control		Price-ratio flyer			Km-per-R$50 flyer		
	Obs	Mean (S.D.)	Obs	Mean (S.D.)	Diff Mean (S.E.)	Obs	Mean (S.D.)	Diff Mean (S.E.)
Choice environment: $p_e/p_g < 0.7/1.1$								
[**EE**] e very favorably priced relative to g								
Chose very favorably priced fuel e (1=yes)	504	0.565 (0.496)	504	0.627 (0.484)	0.062** (0.031)	504	0.619 (0.486)	0.054* (0.031)
					[Either treatment:	1008	0.623 (0.485)	0.058** (0.027)]
Value of fuel purchase (R$)	504	50.7 (30.1)	504	50.6 (29.5)	-0.2 (1.9)	504	49.4 (25.5)	-1.3 (1.8)
					[Either treatment:	1008	50.0 (27.8)	-0.7 (1.6)]
Stated vehicle usage (km/week)	411	345 (344)	409	333 (352)	-11 (24)	400	312 (270)	-33 (22)
					[Either treatment:	809	323 (314)	-22 (20)]
Choice environment: $0.7/1.1 \leq p_e/p_g < 0.7/1.05$								
[**E**] e somewhat favorably priced relative to g								
Chose somewhat favorably priced fuel e (1=yes)	504	0.593 (0.492)	504	0.563 (0.496)	-0.030 (0.031)	504	0.560 (0.497)	-0.034 (0.031)
					[Either treatment:	1008	0.562 (0.496)	-0.032 (0.027)]
Value of fuel purchase (R$)	504	49.9 (27.8)	504	50.8 (30.1)	0.9 (1.8)	504	51.8 (30.5)	1.9 (1.8)
					[Either treatment:	1008	51.3 (30.3)	1.4 (1.6)]
Stated vehicle usage (km/week)	415	328 (397)	412	311 (326)	-17 (25)	410	290 (272)	-37 (24)
					[Either treatment:	822	301 (300)	-27 (22)]

Table 3: Proportion choosing ethanol, value of fuel purchase and stated weekly usage of vehicle, when ethanol was favorably priced, by control and treatment group. Number of observations, mean and standard deviation (S.D.) for each response variable. Difference in means and its standard error (S.E.) between each treatment group and control. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Gasoline very/somewhat favorably priced	Control		Price-ratio flyer			Km-per-R$50 flyer		
	Obs	Mean (S.D.)	Obs	Mean (S.D.)	Diff Mean (S.E.)	Obs	Mean (S.D.)	Diff Mean (S.E.)
Choice environment: $p_e/p_g \geq 0.7 \times 1.1$								
[**GG**] g very favorably priced relative to e								
Chose very favorably priced fuel g (1=yes)	1350	0.884 (0.320)	1350	0.899 (0.301)	0.015 (0.012)	1350	0.886 (0.318)	0.001 (0.012)
					[Either treatment:	2700	0.893 (0.310)	0.008 (0.011)]
Restrict to choice of regular-grade g (1=yes)	1350	0.828 (0.377)	1350	0.844 (0.363)	0.016 (0.014)	1350	0.829 (0.377)	0.001 (0.015)
					[Either treatment:	2700	0.836 (0.370)	0.008 (0.012)]
Value of fuel purchase (R$)	1350	51.1 (34.9)	1350	53.1 (35.3)	1.9 (1.4)	1350	50.8 (34.6)	-0.3 (1.3)
					[Either treatment:	2700	52.0 (35.0)	0.8 (1.2)]
Stated vehicle usage (km/week)	1170	323 (366)	1180	312 (340)	-11 (15)	1171	308 (308)	-15 (14)
					[Either treatment:	2351	310 (324)	-13 (13)]
Choice environment: $0.7 \times 1.05 \leq p_e/p_g < 0.7 \times 1.1$								
[**G**] g somewhat favorably priced relative to e								
Chose somewhat favorably priced fuel g (1=yes)	522	0.839 (0.368)	522	0.856 (0.351)	0.017 (0.022)	522	0.837 (0.370)	-0.002 (0.023)
					[Either treatment:	1044	0.847 (0.360)	0.009 (0.020)]
Restrict to choice of regular-grade g (1=yes)	522	0.743 (0.437)	522	0.782 (0.414)	0.038 (0.026)	522	0.774 (0.419)	0.031 (0.026)
					[Either treatment:	1044	0.778 (0.416)	0.034 (0.023)]
Value of fuel purchase (R$)	522	53.2 (35.4)	522	55.4 (37.8)	2.2 (2.3)	522	51.4 (35.2)	-1.8 (2.2)
					[Either treatment:	1044	53.4 (36.5)	0.2 (1.9)]
Stated vehicle usage (km/week)	439	359 (464)	447	315 (330)	-45* (27)	426	348 (362)	-11 (28)
					[Either treatment:	873	331 (346)	-28 (25)]

Table 4: Proportion choosing gasoline, value of fuel purchase, and stated weekly usage of vehicle, when gasoline was favorably priced, by control and treatment group. Number of observations, mean and standard deviation (S.D.) for each response variable. Difference in means and its standard error (S.E.) between each treatment group and control. * p<0.1, ** p<0.05, *** p<0.01.

Experiment	$p_e/p_g < 0.7/1.1$ [**EE**] e very favorably priced	$0.7/1.1 \leq p_e/p_g < 0.7/1.05$ [**E**] e somewhat favorably priced	$p_e/p_g \geq 0.7 \times 1.1$ [**GG**] g very favorably priced	$0.7 \times 1.05 \leq p_e/p_g < 0.7 \times 1.1$ [**G**] g somewhat favorably priced
Dependent variable	Chose e (1=yes)	Chose e (1=yes)	Chose g (1=yes)	Chose g (1=yes)
Model	Linear probability	Linear probability	Linear probability	Linear probability
	Coeff. (S.E.)	Coeff. (S.E.)	Coeff. (S.E.)	Coeff. (S.E.)
Price-ratio flyer (1=yes)	0.062*** (0.021)	−0.030 (0.027)	0.015 (0.013)	0.017 (0.023)
Km-per-R$50 flyer (1=yes)	0.054** (0.026)	−0.034 (0.037)	0.001 (0.013)	−0.002 (0.024)
R^2	0.076	0.057	0.068	0.070
Model	Probit	Probit	Probit	Probit
	Marg.Eff. (S.E.)	Marg.Eff. (S.E.)	Marg.Eff. (S.E.)	Marg.Eff. (S.E.)
Price-ratio flyer (1=yes)	0.060*** (0.021)	−0.030 (0.027)	0.015 (0.013)	0.016 (0.023)
Km-per-R$50 flyer (1=yes)	0.053** (0.026)	−0.034 (0.036)	0.002 (0.014)	−0.003 (0.024)
Log likelihood	−956	−988	−1272	−630
Station-visit fixed effects	Yes	Yes	Yes	Yes
Mean of dependent variable	0.604	0.572	0.900	0.844
Number of observations	1512	1512	4050	1566

Table 5: Choosing the favorably priced fuel when a fuel is very or somewhat favorably priced relative to the other fuel: Linear probability model and Probit model, each with station-visit fixed effects. For the probit model, estimated marginal effects at the mean for each price subsample are reported. An observation is a subject, treated or control. Station visit-clustered standard errors (S.E.). * p<.1, ** p<.05, *** p<.01

		I		II		III		IV	
		Coeff.	(S.E.)	Coeff.	(S.E.)	Coeff.	(S.E.)	Coeff.	(S.E.)
1	Fuel price (R$/km)	−43.00***	(9.34)	−36.42***	(9.02)	−35.61***	(9.16)	−29.58***	(9.57)
2	Fuel price*Vehicle price (1000 R$, log)	6.42**	(2.71)	4.11	(2.66)	3.36	(2.69)	4.65*	(2.70)
3	Fuel price*Treatment Price Ratio (1=yes)	−4.30**	(1.99)	−4.48**	(2.03)	−4.46**	(2.03)	−4.47**	(2.03)
4	Fuel price*Treatment Km-per-R$50 (1=yes)	−2.20	(1.96)	−2.39	(2.00)	−2.30	(2.00)	−2.40	(2.00)
5	Fuel price*Aged 25 to 40 years (1=yes)					0.44	(2.77)		
6	Fuel price*Aged 40 to 65 years (1=yes)					3.57	(2.85)		
7	Fuel price*Aged more than 65 years (1=yes)					3.83	(5.19)		
8	Fuel price*Secondary school (1=yes)							−8.39*	(4.34)
9	Fuel price*College educated (1=yes)							−9.26**	(4.18)
10	g (1=yes)*Vehicle price (1000 R$, log)	0.41***	(0.09)	0.31***	(0.04)	0.30***	(0.04)	0.31***	(0.04)
11	g (1=yes)*Treatment Price Ratio (1=yes)	−0.02	(0.06)	−0.01	(0.06)	−0.01	(0.06)	−0.01	(0.06)
12	g (1=yes)*Treatment Km-per-R$50 (1=yes)	−0.07	(0.06)	−0.08	(0.06)	−0.07	(0.06)	−0.07	(0.06)
13	g (1=yes)*Female consumer (1=yes)	0.27***	(0.05)	0.24***	(0.05)	0.24***	(0.05)	0.24***	(0.05)
14	g (1=yes)*Aged 25 to 40 years (1=yes)	−0.13*	(0.08)	−0.12	(0.08)	−0.12	(0.08)	−0.12	(0.08)
15	g (1=yes)*Aged 40 to 65 years (1=yes)	−0.15*	(0.08)	−0.13	(0.08)	−0.11	(0.08)	−0.13	(0.08)
16	g (1=yes)*Aged more than 65 years (1=yes)	0.21	(0.15)	0.17	(0.15)	0.19	(0.15)	0.17	(0.15)
17	g (1=yes)*Secondary school (1=yes)	0.22*	(0.13)	0.11	(0.12)	0.10	(0.12)	0.03	(0.12)
18	g (1=yes)*College educated (1=yes)	0.30**	(0.12)	0.18	(0.12)	0.17	(0.12)	0.09	(0.12)
19	g (1=yes)*Uses vehicle extensively (1=yes)	0.01	(0.06)	0.07	(0.06)	0.07	(0.06)	0.07	(0.06)
20	$\bar{g}$ (1=yes)*Vehicle price (1000 R$, log)	0.73***	(0.16)	0.82***	(0.08)	0.82***	(0.08)	0.81***	(0.08)
21	$\bar{g}$ (1=yes)*Treatment Price Ratio (1=yes)	0.02	(0.12)	0.02	(0.12)	0.02	(0.12)	0.03	(0.12)
22	$\bar{g}$ (1=yes)*Treatment Km-per-R$50 (1=yes)	−0.05	(0.11)	−0.05	(0.12)	−0.05	(0.12)	−0.05	(0.12)
23	$\bar{g}$ (1=yes)*Female consumer (1=yes)	−0.08	(0.11)	−0.12	(0.11)	−0.12	(0.11)	−0.12	(0.11)
24	$\bar{g}$ (1=yes)*Aged 25 to 40 years (1=yes)	0.17	(0.18)	0.12	(0.17)	0.12	(0.18)	0.13	(0.17)
25	$\bar{g}$ (1=yes)*Aged 40 to 65 years (1=yes)	0.46***	(0.18)	0.43**	(0.18)	0.40**	(0.18)	0.44**	(0.18)
26	$\bar{g}$ (1=yes)*Aged more than 65 years (1=yes)	1.17***	(0.25)	1.30***	(0.25)	1.25***	(0.25)	1.29***	(0.25)
27	$\bar{g}$ (1=yes)*Secondary school (1=yes)	0.14	(0.24)	0.05	(0.23)	0.04	(0.23)	0.10	(0.23)
28	$\bar{g}$ (1=yes)*College educated (1=yes)	0.39*	(0.23)	0.31	(0.22)	0.30	(0.22)	0.37*	(0.22)
29	$\bar{g}$ (1=yes)*Uses vehicle extensively (1=yes)	0.21**	(0.11)	0.21*	(0.11)	0.21*	(0.11)	0.21**	(0.11)
30	City-fuel fixed effects	Yes		-		-		-	
31	Station-fuel fixed effects, ξ_{jl}	No		Yes		Yes		Yes	
	Log likelihood	−0.6790		−0.6577		−0.6576		−0.6575	

Table 6: Estimated coefficients and standard errors (based on the inverse of the analytically derived Hessian) for the salience-shifts-relative-utility model. The sample consists of 10,329 observations (purchases) and 30,633 alternatives in total. * p<0.1, ** p<0.05, *** p<0.01.

		I		II	
		Coeff.	(S.E.)	Coeff.	(S.E.)
	Choice Set (CS) parameters:				
CS1	$\lambda^g_{\text{SaoPaulo}}$ (**Single-fuel gasoline mindset**)	0.35***	(0.04)		
CS2	$\lambda^g_{\text{Curitiba}}$	0.21***	(0.06)		
CS3	$\lambda^g_{\text{BeloHorizonte}}$	0.67***	(0.07)		
CS4	$\lambda^g_{\text{Recife}}$	0.41***	(0.11)		
CS5	Mean λ^g_l across 52 stations			0.44***	(0.00)
CS6	$\lambda^e_{\text{SaoPaulo}}$ (**Single-fuel ethanol mindset**)	0.24***	(0.08)		
CS7	$\lambda^e_{\text{Curitiba}}$	0.10	(0.07)		
CS8	$\lambda^e_{\text{BeloHorizonte}}$	0.03*	(0.02)		
CS9	$\lambda^e_{\text{Recife}}$	0.04	(0.04)		
CS10	Mean λ^e_l across 52 stations			0.17***	(0.00)
CS11	ϕ_{SaoPaulo} (**Treatment effect: Bi-fuel mindset**)	0.78**	(0.11)	0.78**	(0.09)
CS12	ϕ_{Curitiba}	1.00	(0.22)	1.00	(0.19)
CS13	$\phi_{\text{BeloHorizonte}}$	0.95	(0.09)	0.75***	(0.10)
CS14	ϕ_{Recife}	1.00	(0.67)	0.89	(0.20)
	Utility (U) function parameters:				
U1	Fuel price (R$/km)	−77.77***	(28.18)	−75.77***	(26.59)
U2	Fuel price*Vehicle price (1000 R$, log)	14.28**	(7.20)	13.01	(8.07)
U3	Fuel price*Secondary school (1=yes)	−17.55*	(9.61)	−19.24*	(10.14)
U4	Fuel price*College educated (1=yes)	−18.99**	(9.48)	−19.01**	(9.22)
U5	g (1=yes)*Mean vehicle price at station (1000 R$, log)	5.90***	(1.46)	−1.09	(4.67)
U6	$\bar{g}$ (1=yes)*Mean vehicle price at station (1000 R$, log)	6.05***	(1.53)	−1.01	(4.77)
U7	g (1=yes)*Vehicle price (1000 R$, log)	0.63***	(0.21)	0.65***	(0.22)
U8	g (1=yes)*Female consumer (1=yes)	0.46***	(0.16)	0.66***	(0.16)
U9	g (1=yes)*Aged 25 to 40 years (1=yes)	−0.24	(0.16)	−0.27	(0.20)
U10	g (1=yes)*Aged 40 to 65 years (1=yes)	−0.26	(0.16)	−0.21	(0.19)
U11	g (1=yes)*Aged more than 65 years (1=yes)	0.63*	(0.37)	0.44	(0.52)
U12	g (1=yes)*Secondary school (1=yes)	0.14	(0.28)	−0.15	(0.28)
U13	g (1=yes)*College educated (1=yes)	0.24	(0.27)	−0.07	(0.26)
U14	g (1=yes)*Uses vehicle extensively (1=yes)	0.11	(0.12)	0.09	(0.15)
U15	$\bar{g}$ (1=yes)*Vehicle price (1000 R$, log)	1.07***	(0.23)	1.11***	(0.25)
U16	$\bar{g}$ (1=yes)*Female consumer (1=yes)	0.16	(0.19)	0.36*	(0.18)
U17	$\bar{g}$ (1=yes)*Aged 25 to 40 years (1=yes)	0.09	(0.23)	0.06	(0.26)
U18	$\bar{g}$ (1=yes)*Aged 40 to 65 years (1=yes)	0.36	(0.23)	0.41*	(0.25)
U19	$\bar{g}$ (1=yes)*Aged more than 65 years (1=yes)	1.75***	(0.43)	1.57***	(0.55)
U20	$\bar{g}$ (1=yes)*Secondary school (1=yes)	0.28	(0.32)	0.02	(0.33)
U21	$\bar{g}$ (1=yes)*College educated (1=yes)	0.51	(0.32)	0.21	(0.31)
U22	$\bar{g}$ (1=yes)*Uses vehicle extensively (1=yes)	0.28*	(0.15)	0.27	(0.18)
U23	Fuel fixed effects, ξ_j	Yes		Yes	
	Log likelihood	−0.6800		−0.6683	

Table 7: Estimated coefficients and standard errors (based on the inverse of the analytically derived Hessian) for the salience-shifts-choice-set model. The sample consists of 10,329 observations (purchases) and 30,633 alternatives in total. Constrained optimization using the solver Knitro, with estimates robust to initial values. * p<0.1, ** p<0.05, *** p<0.01. In rows CS11 to CS14, asterisks correspond to the p-value for a one-tailed test with a null of 1 against the alternative that the parameter is less than 1.

Similarly priced fuels	Control		Price-ratio flyer			Km-per-R$50 flyer		
	Obs	Mean (S.D.)	Obs	Mean (S.D.)	Diff Mean (S.E.)	Obs	Mean (S.D.)	Diff Mean (S.E.)
$0.7/1.05 \leq p_e/p_g < 0.7 \times 1.05$								
Chose similarly priced fuel e (1=yes)	594	0.333 (0.472)	594	0.369 (0.483)	0.035 (0.028)	594	0.389 (0.488)	0.056** (0.028)
					[Either treatment:	1188	0.379 (0.485)	0.045* (0.024)]

Table 8: (Appendix) Proportion choosing ethanol when ethanol is similarly priced relative to gasoline, by control and treatment group. Difference in means and its standard error (S.E.) between each treatment group and control. * p<0.1, ** p<0.05, *** p<0.01.

Dependent variable	k_e^{urban}	k_g^{urban}	k_e^{urban}/k_g^{urban}
Specification (a)	**I(a)**	**II(a)**	**III(a)**
Size of the engine (liters)	-1.17^{***} (0.14)	-1.70^{***} (0.22)	-0.00 (0.01)
Cohort age, 2012 − model-year (years)	-0.11^{***} (0.03)	-0.16^{***} (0.04)	-0.00 (0.00)
Non-manual transmission (1=yes)	-0.14^{**} (0.07)	0.01 (0.08)	-0.02^{***} (0.00)
Intercept	8.79^{***} (0.22)	13.00^{***} (0.39)	0.68^{***} (0.01)
Vehicle model fixed effects	Yes	Yes	Yes
R^2	0.80	0.82	0.48
Fitted value at lab means (km/l)	7.12 (0.03)	10.47 (0.03)	0.68 (0.00)
Specification (b)	**I(b)**	**I(b)**	**III(b)**
Size of the engine (liters)	-1.38^{***} (0.13)	-1.99^{***} (0.18)	-0.00 (0.01)
Cohort age, 2012 − model-year (years)	-0.11^{***} (0.03)	-0.15^{***} (0.04)	-0.00 (0.00)
Non-manual transmission (1=yes)	-0.11 (0.07)	-0.00 (0.09)	-0.01^{**} (0.00)
Vehicle class is compact (1=yes)	-0.50^{***} (0.11)	-0.60^{***} (0.15)	-0.01 (0.01)
Vehicle class is mid-size (1=yes)	-0.49^{***} (0.12)	-0.45^{***} (0.16)	-0.02^{*} (0.01)
Vehicle class is full-size (1=yes)	-0.72^{***} (0.17)	-0.90^{***} (0.21)	-0.01 (0.01)
Vehicle class is small truck (1=yes)	-0.50^{***} (0.12)	-0.71^{***} (0.16)	-0.00 (0.00)
Vehicle class is SUV (1=yes)	-1.02^{***} (0.19)	-1.17^{***} (0.26)	-0.02^{**} (0.01)
Vehicle class is minivan (1=yes)	-2.08^{***} (0.13)	-2.88^{***} (0.16)	-0.01 (0.01)
Intercept	9.50^{***} (0.18)	13.80^{***} (0.25)	0.69^{***} (0.01)
Make fixed effects	Yes	Yes	Yes
R^2	0.70	0.72	0.22
Fitted value at lab means (km/l)	7.12 (0.03)	10.47 (0.04)	0.68 (0.00)
Number of observations	230	230	230

Table 9: (Appendix) Predicting vehicle-specific fuel economy in the urban driving cycle using laboratory test data. An observation is a model-engine-year combination in the Inmetro test sample. Specification (a) omits the Fiat Palio dummy variable. Specification (b) omits the Fiat and the subcompact dummy variables. Heteroscedasticity-robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01

1A

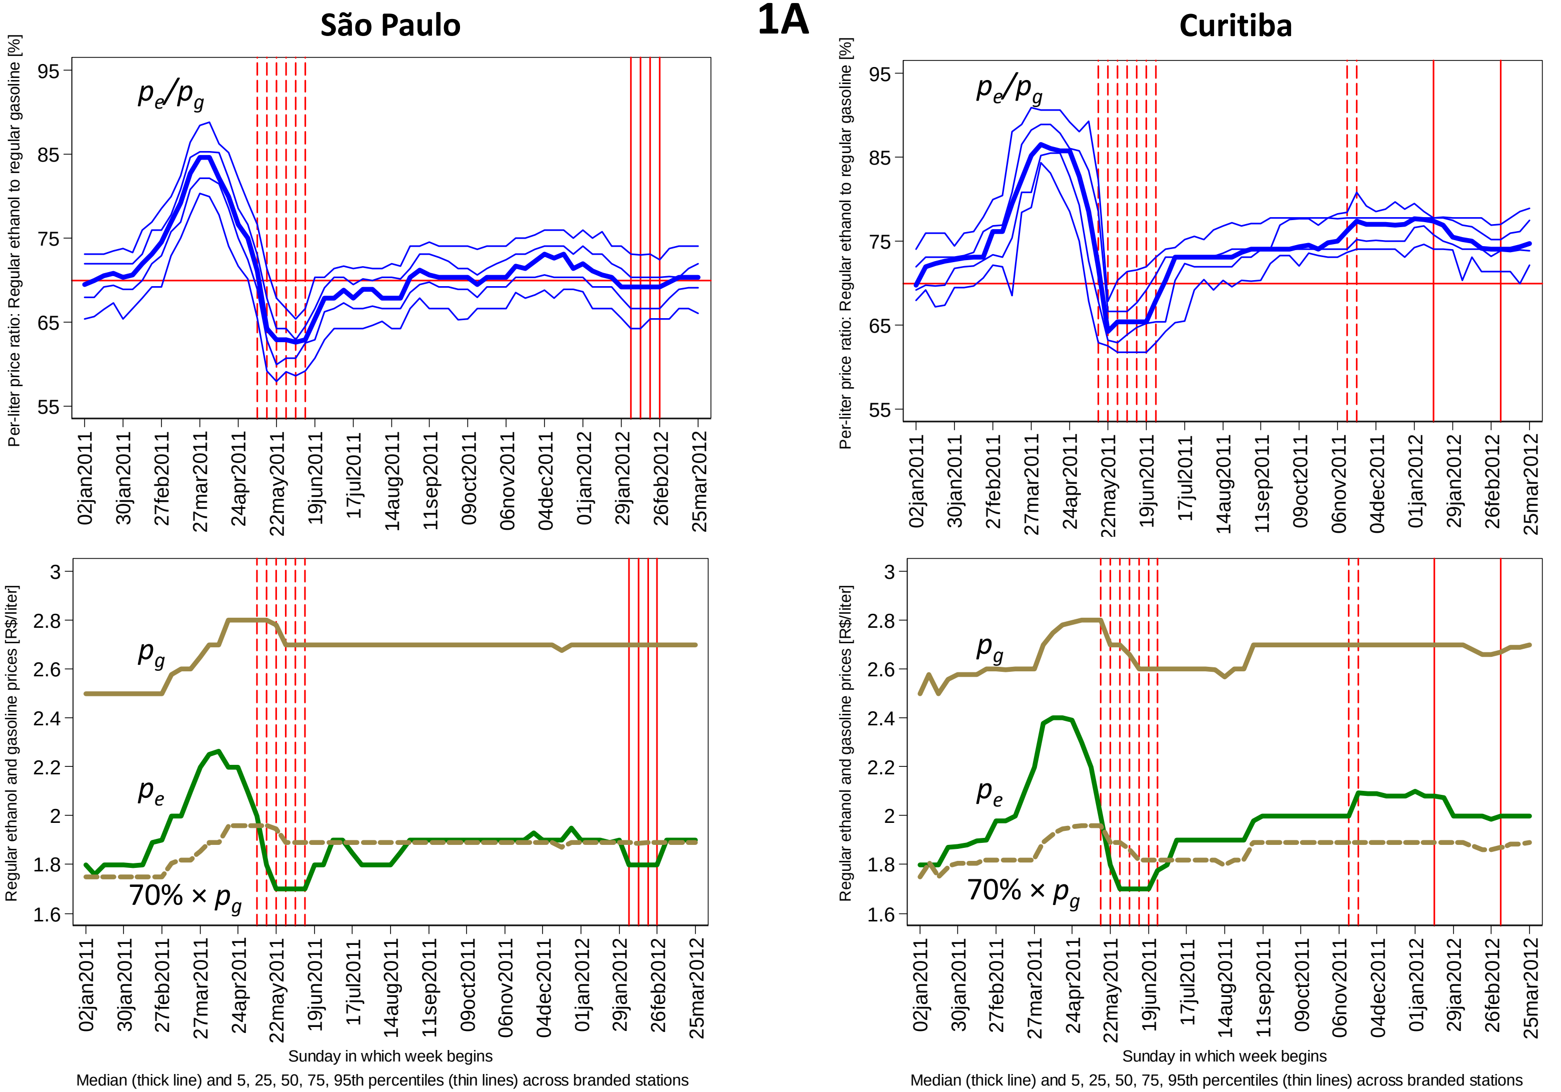

Figure 1: Price paths from weekly external samples of fueling stations in (**1A**) **São Paulo and Curitiba**, and (1B) Belo Horizonte and Recife. The **top panels** report the per-liter regular-grade **ethanol-to-gasoline price ratio**, and the **bottom panels** report **price levels**, in (nominal) R$/liter. Thick curves indicate medians and the thin curves indicate the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles (top panels only) in the cross-sections of branded stations. Horizontal red lines mark the 70% media-reported "parity" ratio. Vertical red lines indicate field activity, with treatment either by the station attendant (dashed lines) or the enumerator (solid lines). Source: ANP's retail price database.

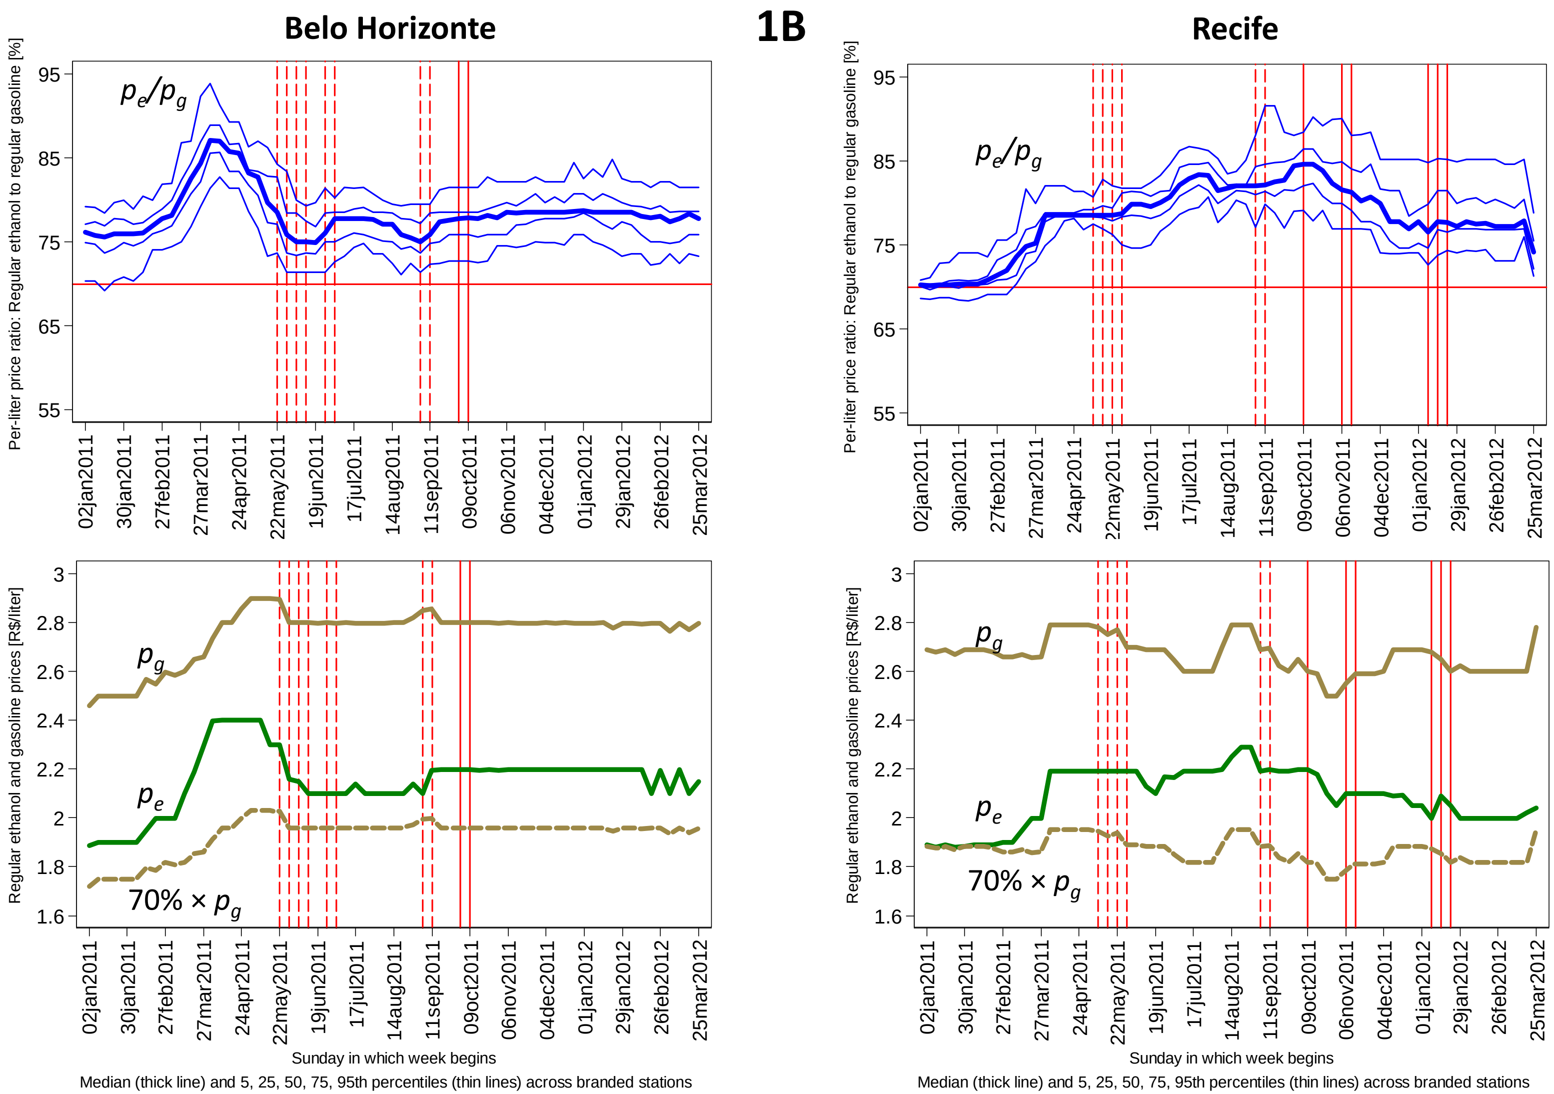

Figure 1: Price paths from weekly external samples of fueling stations in (1A) São Paulo and Curitiba, and (**1B**) **Belo Horizonte and Recife**. The **top panels** report the per-liter regular-grade **ethanol-to-gasoline price ratio**, and the **bottom panels** report **price levels**, in (nominal) R$/liter. Thick curves indicate medians and the thin curves indicate the 5th, 25th, 75th, and 95th percentiles (top panels only) in the cross-sections of branded stations. Horizontal red lines mark the 70% media-reported "parity" ratio. Vertical red lines indicate field activity, with treatment either by the station attendant (dashed lines) or the enumerator (solid lines). Source: ANP's retail price database.

CAMPANHA:

Oferecendo Informação a Você!

Nesse posto hoje: **Preço Álcool Comum / Preço Gasolina Comum = 66%**

Mais vantagem: Álcool

Prezado Sr./Sra. Motorista de Veículo "Flex":

Nesse posto hoje, o litro do Álcool Comum representa **66%** do preço do litro da Gasolina Comum.

Especialistas avisam que quando essa relação de preços (entre Álcool e Gasolina) estiver:

→ abaixo de 70%, o Álcool está mais vantajoso.*

→ acima de 70%, a Gasolina está mais vantajosa.*

* Ou seja, oferece o menor preço por Km rodado. A relação de paridade exata depende do modelo de seu veículo, de seu estilo de dirigir, e da composição dos combustíveis. Considera-se gasolina com adição de 25% de etanol anidro combustível, conforme a lei. Fonte: Tabelas de Eficiência Energética, INMETRO.
Para mais informações, escreva para: a-salvo@kellogg.northwestern.edu

Realização: CNPBrasil NORTHWESTERN UNIVERSITY

Informações: INMETRO

Expto 2
31020

Figure 2A: Example of the “price ratio relative to 70% threshold” flyer presented to subjects in one treatment group, used at a station in São Paulo visited on June 13, 2011 in which (p_e,p_g)=(1.649,2.499).

CAMPANHA:

Oferecendo Informação a Você!

		Álcool	**Gasolina**	**Mais vantagem: Álcool**
				
Motor	Exemplo	R$ 50 de Álcool Comum comprado aqui rende:	R$ 50 de Gasolina Comum comprada aqui rende:	
Motor 1.0	VW Gol	**227 KM**	**218 KM**	
Motor 1.4	GM Corsa	**212 KM**	**204 KM**	
Motor 1.8	Fiat Stilo	**195 KM**	**188 KM**	

Prezado Sr./Sra. Motorista de Veículo "Flex":

A tabela acima mostra a quantidade estimada de KM (quilômetros) que cada R$ 50 de Álcool Comum ou de Gasolina Comum comprados nesse posto hoje permitem rodar na cidade.*

* Média calculada para 49 modelos 2009 ou 2010 com classificação no Programa Brasileiro de Etiquetagem (PBE), rodando na cidade. Considera-se gasolina com adição de 25% de etanol anidro combustível, conforme a lei. Fonte: Tabelas de Eficiência Energética, INMETRO.
Para mais informações, escreva para: a-salvo@kellogg.northwestern.edu

Realização: CNPBrasil NORTHWESTERN UNIVERSITY

Informações: INMETRO

Expto 3
31020

Figure 2B: Example of the “km per R$ 50” flyer presented to subjects in another treatment group, used at a station in São Paulo visited on June 13, 2011 in which (p_e,p_g)=(1.649,2.499).

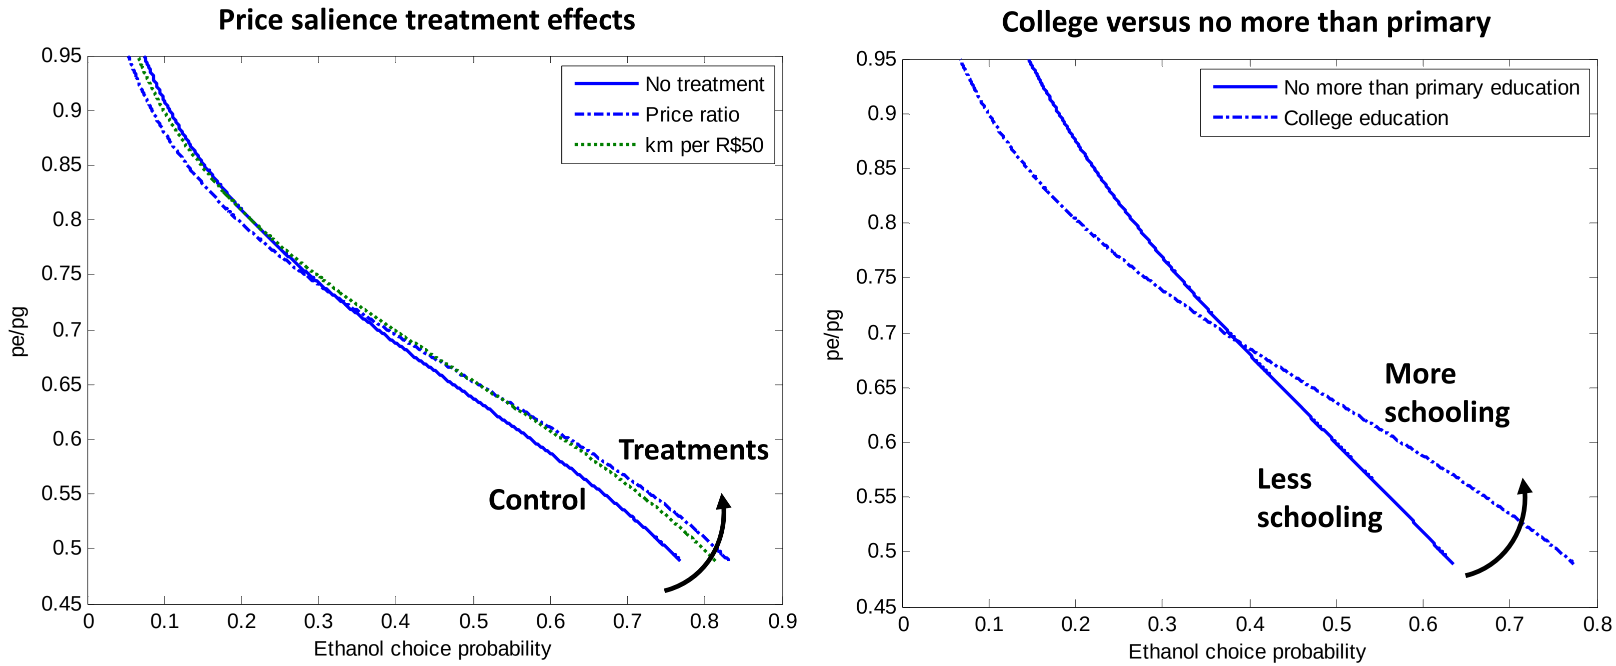

Figure 3: (Model 1, Salience shifts a consumer's relative utility) **Predicted ethanol share** in the full sample of choices (all cities, both waves) under different counterfactual scenarios, as the relative price point varies, for p_e/p_g varying from 0.50 to 0.95. The **left panel** compares different **price salience treatments**, with all consumers counterfactually exposed to: (i) the price-ratio treatment (dashed, blue), (ii) the km-per-R$ 50 treatment (dotted, green), or (iii) no treatment (solid, blue). The **right panel** reports, for comparison, different levels of education (and no price salience treatment), with the following alternative characteristics counterfactually "turned on" for all consumers: (i) college education (dashed), or (ii) no more than primary education (solid). Source: Specification in column IV, Table 6. The price of ethanol is varied while holding gasoline (E25) prices constant at the sample values. Other individual and market characteristics are at their observed values.

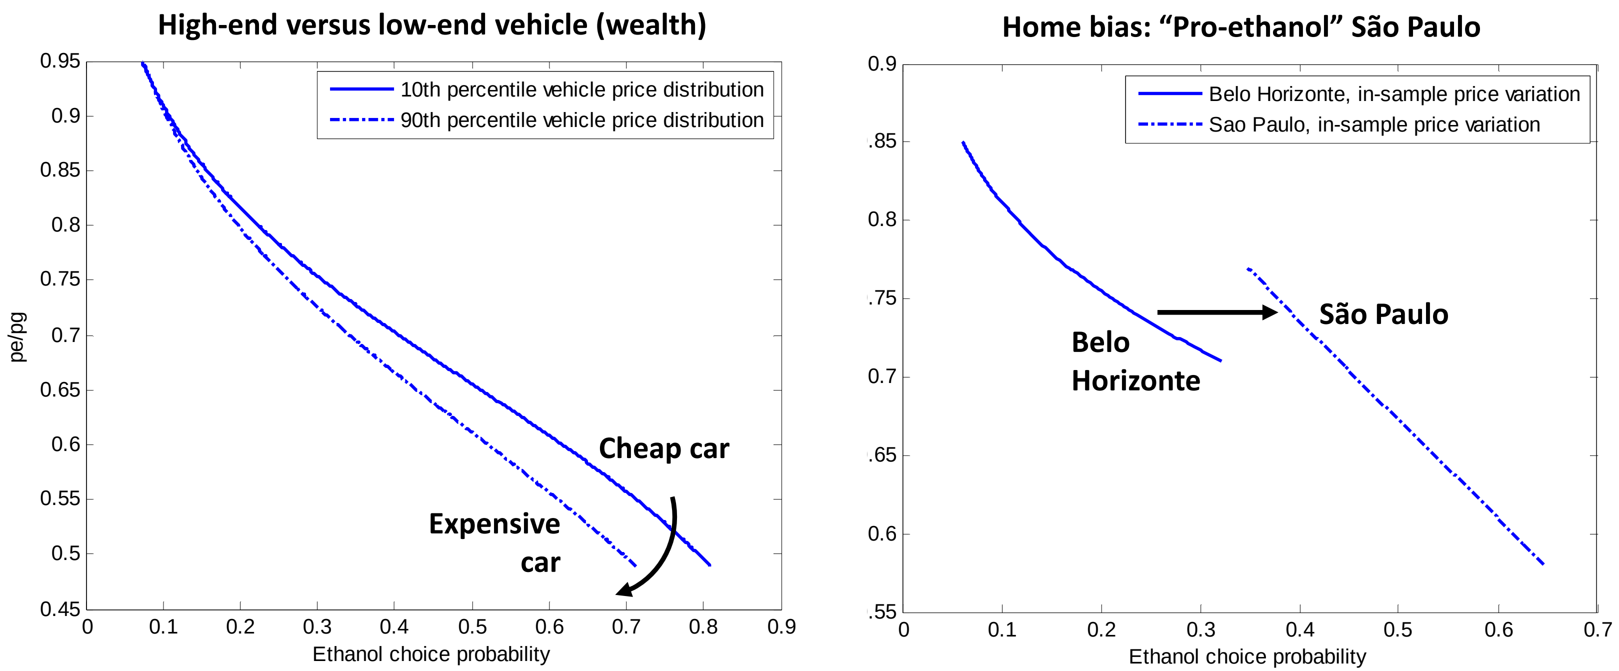

Figure 4: (Model 1, Salience shifts a consumer's relative utility) **Predicted ethanol share** under further counterfactual scenarios, as the relative price point varies, for p_e/p_g varying from 0.50 to 0.95. The **left panel** compares demand in the full sample at different proxied wealth levels: vehicle price at (i) the 90th percentile (dashed), or (ii) the 10th percentile (solid) of the empirical distribution of vehicle prices. The **right panel** reports demand estimated separately from the subsample of choices observed in São Paulo (dashed) and the subsample of choices in Belo Horizonte (solid). Source: Specification in column IV, Table 6. The price of ethanol is varied while holding gasoline (E25) prices constant at the sample values. Other individual and market characteristics are at their observed values.

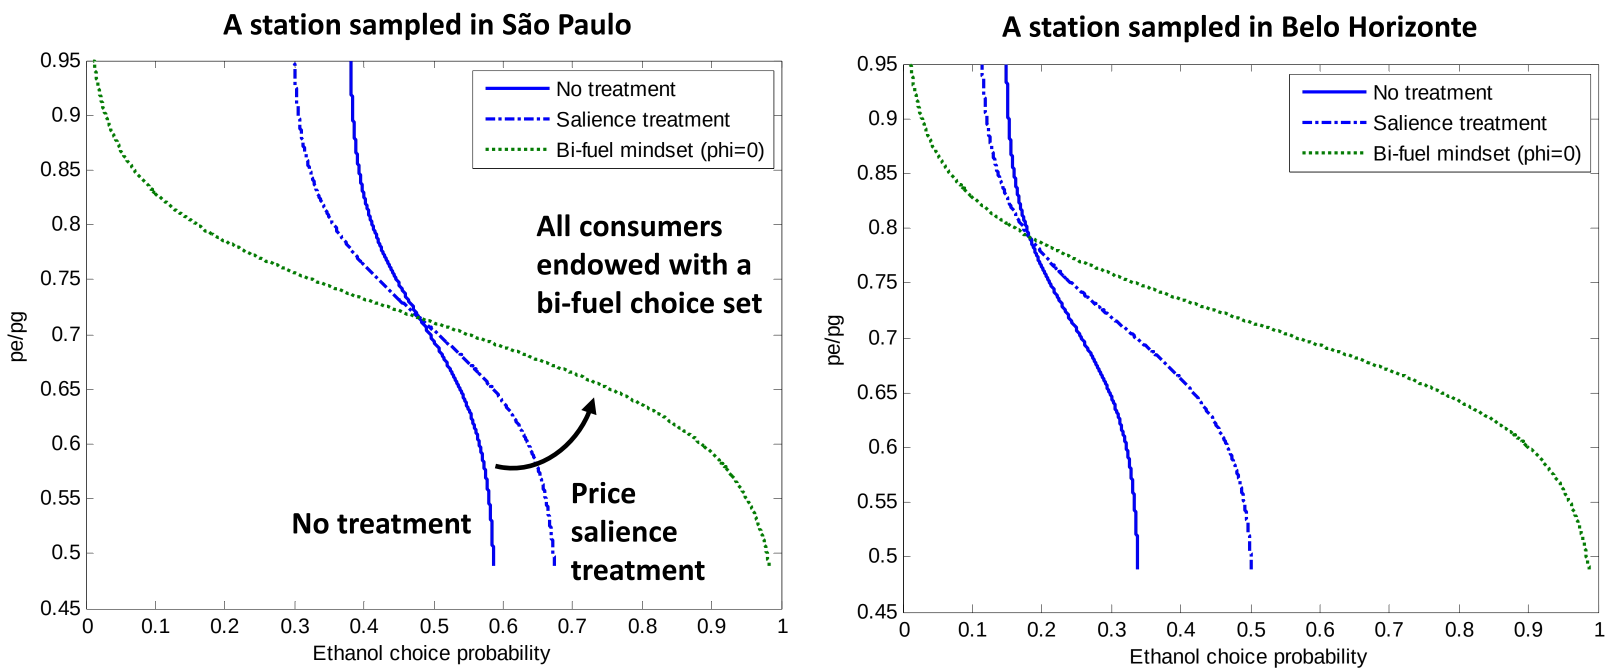

Figure 5: (Model 2, Salience shifts a consumer's random choice set) **Predicted ethanol share** at two different stations in the sample, as the relative price point varies, under alternative counterfactual scenarios: (i) no treatment (solid, blue), (ii) all sampled consumers undergo price salience treatment, with either the price-ratio or the km-per-R$ 50 flyers (dashed, blue), or (iii) all sampled consumers are endowed with bi-fuel choice sets (dotted, green). **Left panel**: A specific station sampled in São Paulo, with estimated proportions of single-fuel minded gasoline and ethanol consumers, respectively, of $\hat{\lambda}_l^g$=0.41 and $\hat{\lambda}_l^e$=0.38. **Right panel**: A specific station sampled in Belo Horizonte, with $\hat{\lambda}_l^g$=0.66 and $\hat{\lambda}_l^e$=0.15. Source: Specification of column II, Table 7. The price of ethanol is varied while holding gasoline (E25) prices constant at the sample values. Other individual and market characteristics are at their observed values.

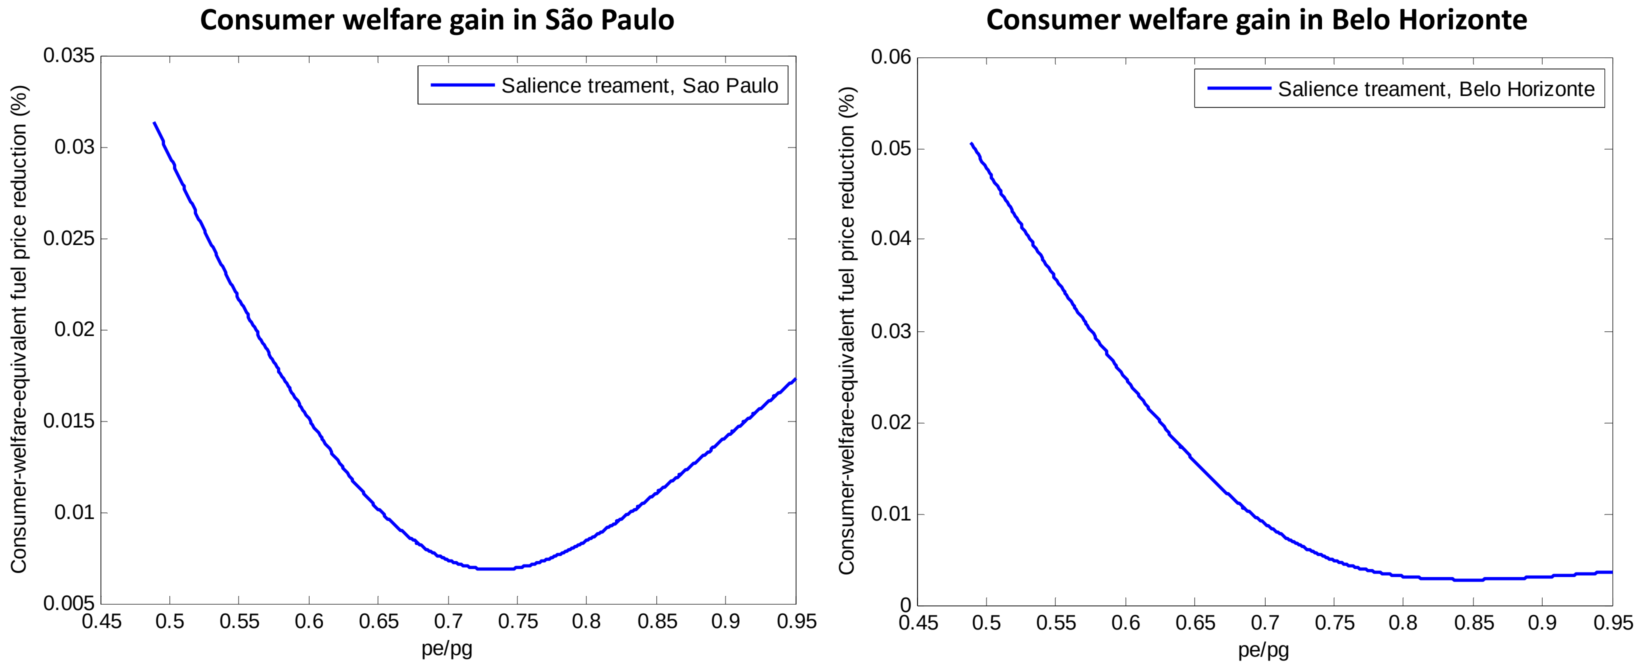

Figure 6: (Model 2, Salience shifts a consumer's random choice set) **Aggregate consumer welfare gains from raising the salience of prices** in São Paulo (**left panel**) and in Belo Horizonte (**right panel**), as the relative price point varies, for p_e/p_g varying from 0.50 to 0.95. Source: Specification of column II, Table 7. The price of ethanol is varied while holding gasoline (E25) prices constant at the sample values. Other individual and market characteristics are at their observed values.

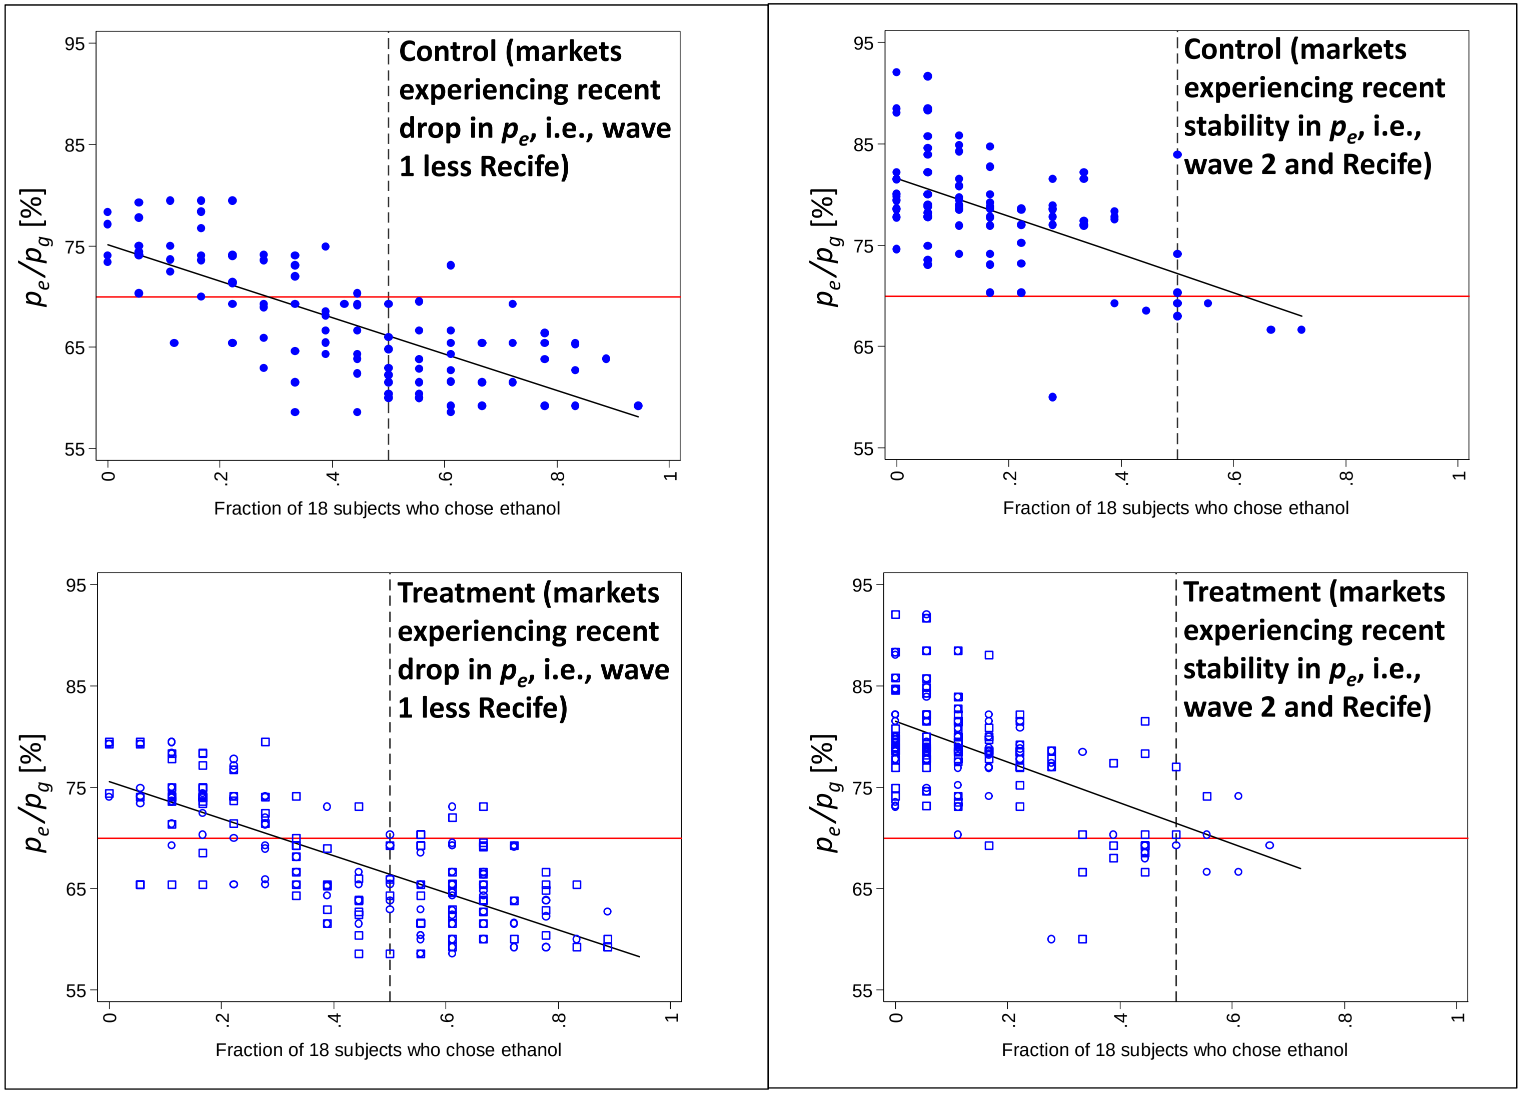

Figure A1: **Empirical demand** at the station-visit level, separately for the **control** group (**top panels**) and the **treatment** groups (**bottom panels**). Per-liter regular-grade ethanol-to-gasoline (p_e/p_g) plotted against ethanol's share in the 18 choices observed in each control or treatment group in each station visit. The control and treatment panels on the **left** depict fuel choices during wave 1 experiments in São Paulo, Curitiba and Belo Horizonte, in which **ethanol prices had fallen sharply**. The control and treatment panels on the **right** depict fuel choices in all other experiments, including wave 2, in which **ethanol prices were quite stable**. In the bottom panels, price-ratio treatment choices are marked with circles and those for km-per-R$50 treatments with squares.

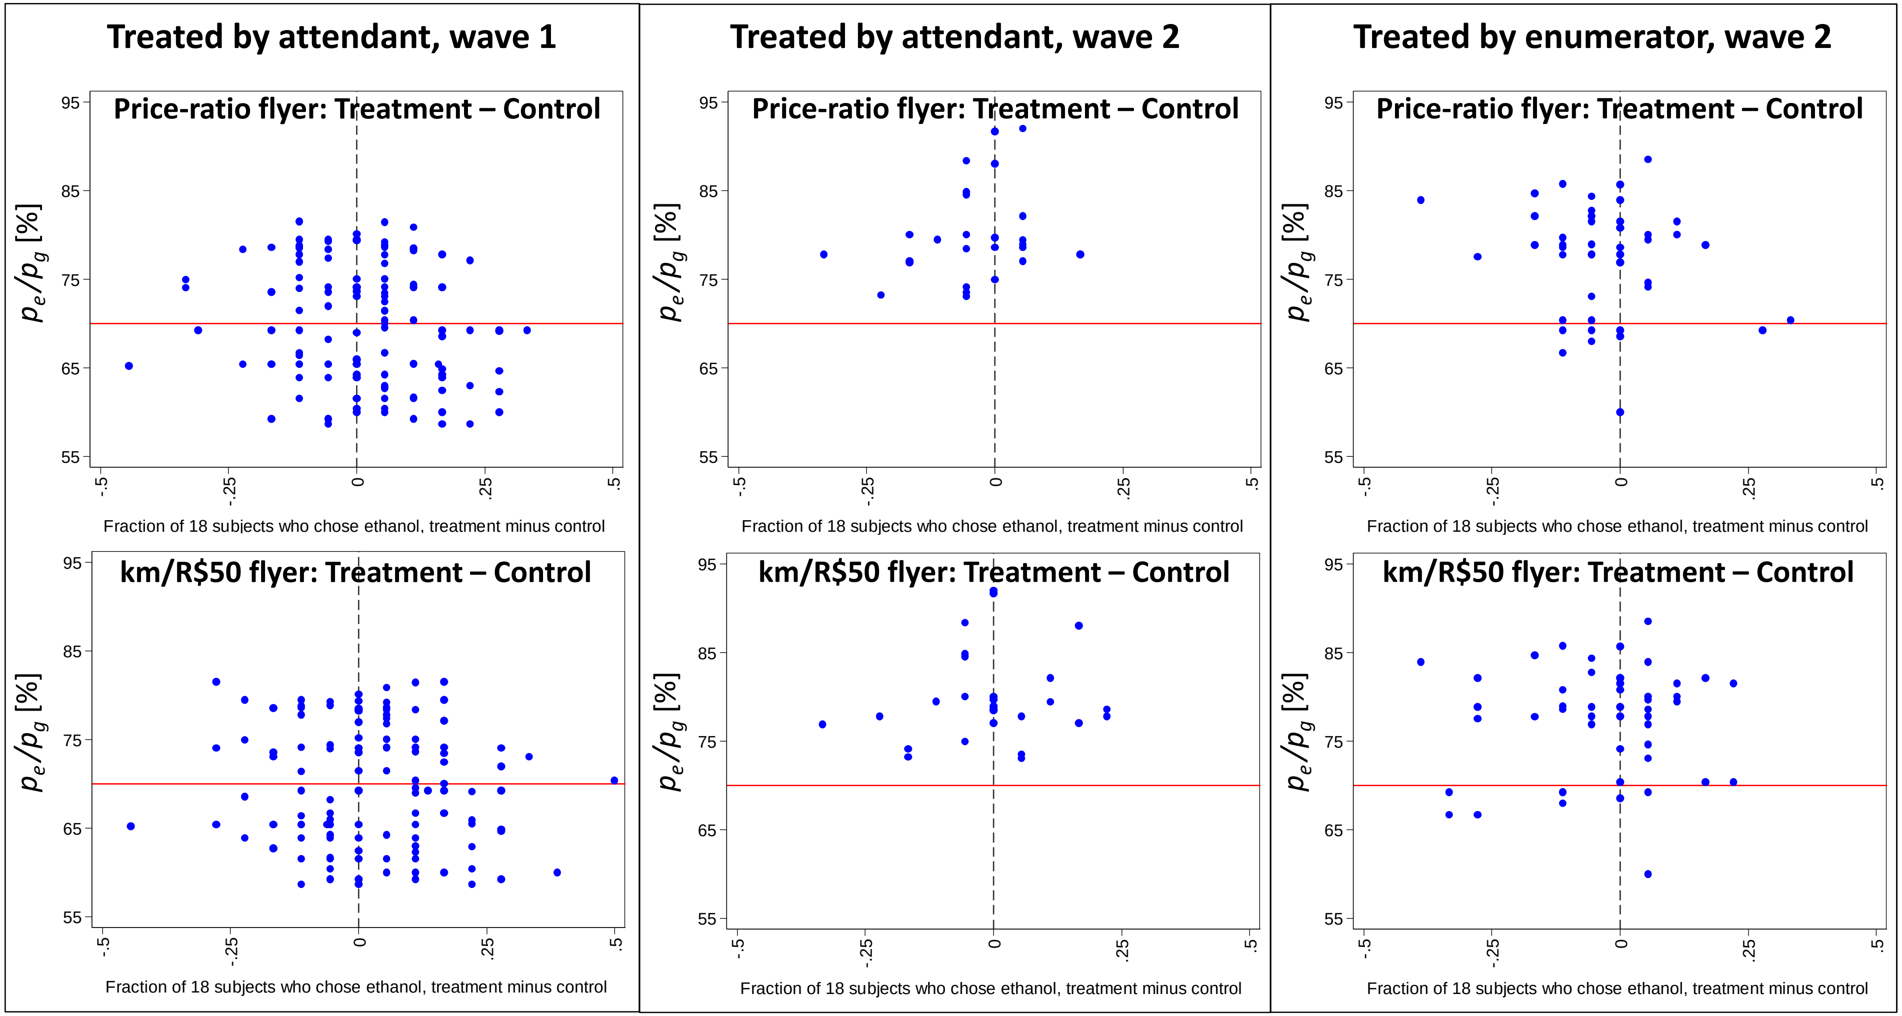

Figure A2: **Average treatment effects by station visit**, for the **price-ratio treatment** (**top panels**) and the **km-per-R$50 treatment** (**bottom panels**). Per-liter regular-grade ethanol-to-gasoline (p_e/p_g) plotted against ethanol's share in the 18 choices observed in a treatment group minus ethanol's share in the 18 choices observed in the control. The **left** panels depict treatment effects in **wave 1 experiments with subjects treated by the station attendant**, the **middle** panels depict **wave 2 experiments with treatment by the attendant**, the **right** panels depict **wave 2 experiments with treatment by the enumerator** (see Table 1).