

High Frequency Evidence on the Demand for Gasoline[†]

By LAURENCE LEVIN, MATTHEW S. LEWIS, AND FRANK A. WOLAK*

Daily city-level expenditures and prices are used to estimate the price responsiveness of gasoline demand in the United States. Using a frequency of purchase model that explicitly acknowledges the distinction between gasoline demand and gasoline expenditures, the price elasticity of demand is consistently found to be an order of magnitude larger than estimates from recent studies using more aggregated data. Estimating demand using higher levels of spatial and temporal aggregation is shown to produce increasingly inelastic estimates. A decomposition is then developed and implemented to understand the relative importance of several different factors in explaining this result. (JEL C51, L71, Q35)

The price elasticity of demand for gasoline has been extensively studied over the last 40 years, and for good reason. Understanding gasoline demand responsiveness is critical in determining gasoline tax rates and evaluating alternative policies that target the negative externalities associated with automobile use (pollution, road congestion, etc.). In the United States, continuing pressure to address climate change has prompted a variety of policy proposals at the national level as well as legislative action at the state level. Starting in 2015, California extended its cap-and-trade program for greenhouse gas emissions to cover transportation fuels, and a number of other states have significantly increased their gasoline tax rates. Gasoline prices have also become increasingly volatile as a result of periodic shortages in available refining capacity and increased uncertainty in world oil markets. Understanding consumers' ability to respond to such price fluctuations is crucial for predicting the potential macroeconomic impacts of future petroleum supply disruptions and for evaluating the benefits of policy measures intended to mitigate these effects, such as the maintenance and use of the US Strategic Petroleum Reserve or the use of temporary gasoline tax suspensions.

While empirical studies of gasoline demand have adopted a variety of different estimation strategies, nearly all of them face the challenge of having to work with

*Levin: VISA Decision Sciences, 901 Metro Center Boulevard, Mail Stop: M3-2B Foster City, CA 94404 (email: llevin@visa.com); Lewis: John E. Walker Department of Economics, Clemson University, 228 Sistine Hall, Box 341309, Clemson, SC 29634 (email: mslewis@clemson.edu); Wolak: Program on Energy and Sustainable Development and Department of Economics, Stanford University, Stanford, CA 94305 (email: wolak@stanford.edu). We thank Visa for providing data and seminar participants at the Ohio State University, Clemson University, and DePaul University, and conference participants of the NBER Summer Institute EEE Meeting, the UC Berkeley Energy Institute Summer Camp, and the International Industrial Organization Conference for helpful comments and suggestions.

[†]Go to <https://doi.org/10.1257/pol.20140093> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

imperfect and often highly aggregated data.¹ Many utilize monthly, quarterly, or even annual aggregate proxies of gasoline usage and average prices, often from a single national time series.² Others rely on cross-sectional data in an attempt to identify demand elasticities based on price variation across regions or countries. In reality, individuals make gasoline consumption decisions on a daily basis, responding directly to the gasoline prices observed in their local area on that day. Empirical models relating monthly or annual gasoline volumes to average prices across broad geographic areas necessarily aggregate these different consumption decisions and are likely to mask a significant share of the response by consumers to a local price change. Moreover, the use of highly aggregated data generally requires strong assumptions that restrict the demand relationship from varying across locations or over time. As a result, unobserved heterogeneity in underlying demand has the potential to bias elasticity estimates. Perhaps not surprisingly given such challenges, these aggregate studies have produced a wide range of different estimates of demand elasticity. Academic and government studies evaluating potential policy interventions in gasoline markets frequently rely on estimates from this literature or adopt the same problematic methods to obtain estimates despite the fact that the elasticity values can often substantially impact predicted policy outcomes.

Our study uses daily gasoline prices and citywide gasoline expenditures for 243 US cities from 2006 through 2009 to analyze the impact of daily prices on daily gasoline demand. These high-frequency panel data have several important advantages. First, the expenditure information comes from credit card transactions at the point of sale and, therefore, offers a much more direct measure of consumer purchases. Second, observing sales and prices daily at the city-level allows us to model demand at the level that purchase decisions actually occur, rather than by relating average consumption across cities over some time period to a corresponding average price. Finally, our panel data can be exploited by including extensive fixed effects to better control for persistent differences in gasoline demand over time and across locations.

Our approach yields robust estimates of the own-price elasticity of demand for gasoline by utilizing a more accurate point-of-sale measure of consumption and adopting an empirical strategy that leverages the higher frequency and greater geographic detail of our consumption and price data to avoid the biases potentially impacting the estimates and policy conclusions of previous studies employing aggregate data. Second, we derive a decomposition identifying the different sources of bias that arise in more aggregate models and then examine the relative magnitudes of these different biases by estimating demand models at varying levels of data aggregation. The results of the decomposition help to clarify why the different empirical approaches utilized in the literature tend to generate such different estimates of the demand elasticity.

¹There are a number of survey articles and meta-analyses available (including Dahl and Sterner 1991, Goodwin 1992, Espey 1998, Basso and Oum 2007, and Brons et al. 2008) that summarize and analyze the literature on gasoline demand estimation.

²The most commonly used proxy for gasoline consumption from the US Energy Information Administration measures the volume of gasoline disappearing from refineries or pipelines minus the estimated volume of gasoline exported.

Unlike the existing literature our analysis also directly addresses the fundamental difference between gasoline demand (or usage) and gasoline expenditures—a difference that becomes even more pronounced when using daily data. The price of gasoline on a particular day can influence both how much gasoline a consumer uses in that day as well as whether they decide to make a gasoline purchase. Our approach more accurately models the demand for gasoline in a manner that recognizes the distinction between expenditure and consumption. We specify a two-equation model of the consumer's probability of gasoline purchase and daily gasoline demand that separates the demand decision from the purchase decision in the most flexible manner possible given our city-level expenditure data. Because we observe both the number of gasoline transactions and the total expenditures on gasoline, we are able to separately identify changes in consumers' probability of purchase from changes in consumers' underlying demand for gasoline. Aggregating our two-equation model of individual gasoline purchase and demand over all individuals in a metropolitan area yields a model of daily aggregate gasoline expenditures that we can use to recover a price elasticity of demand for the metropolitan area.

Using this model, we obtain estimates of gasoline demand elasticity ranging from -0.27 to -0.35 . These responses are nearly an order of magnitude more elastic than some recent (and commonly cited) estimates from comparable studies (Hughes, Knittel, and Sperling 2008; Small and Van Dender 2007; Park and Zhao 2010), implying that changes in gasoline prices or taxes may have a much larger impact on gasoline usage or greenhouse gas emissions than one might otherwise predict. Recent events seem to reflect this substantial responsiveness in gasoline demand. Declining gas prices during 2015 and 2016 have reportedly resulted in dramatic reductions in public transit ridership, increases in vehicle miles traveled, and surge in sales of vehicles with lower fuel economy (Morath 2016; Sommer 2015). While studies using aggregate data from the 1970s and 1980s commonly reported gasoline demand elasticity estimates around -0.25 to -0.30 ,³ our aggregation results suggest that many of these estimates were also likely to have been biased and that actual demand response in earlier decades may have been substantially more elastic than previously thought.

Like other studies using similar static models of gasoline demand, our estimates should be interpreted to reflect the short-run demand response that occurs over a period of months rather than the long-run demand responses that occur over many years (potentially capturing expansions in public transit or the development of more efficient cars).⁴ However, our daily data also allow us to investigate whether demand or consumer purchase decisions respond differently in the very short run. We consider a dynamic two-equation frequency of purchase model that incorporates lagged prices, allowing the immediate response of demand to a price change to differ from the overall short-run demand response. We find evidence that gasoline expenditures respond even more strongly in the days immediately following a price change

³Surveys of these earlier results include Dahl and Sterner (1991), Goodwin (1992), Espey (1998), and Basso and Oum (2007).

⁴See Section IA for a more complete discussion of the relevant response time horizons captured by different demand specifications.

(mainly due to a response in the probability of purchase), but this temporary additional response largely dissipates after four to five days. The response that remains is nearly identical to the overall short-run elasticity estimate obtained in our main (static) analysis. This further supports the argument that the estimates from our baseline static model accurately reflect the same short-run response that previous studies attempt to estimate using monthly or annual data, rather than the very-short run response that occurs in the days immediately following the shock.

In the second part of our analysis, we investigate the consequences of estimating demand using more highly aggregated data, as is common in previous studies. As a first step, we use our data to estimate a standard demand model aggregated over time and across cities to varying degrees. The resulting estimates become increasingly less elastic as the level of data aggregation increases. Estimating the model using our data aggregated to a national time series of monthly total expenditures and average prices results in elasticities that are indistinguishable from zero, suggesting that studies using aggregated data may substantially underestimate consumers' price responsiveness.

Next, to better understand the impacts of data aggregation, we derive a decomposition identifying three distinct components of bias that arise in more aggregate demand models. First, demand models that assume a common price coefficient are often estimated even though the underlying coefficients have the potential to differ across cities. While the coefficient obtained from this restricted model does estimate a weighted average of these coefficients, the weighting applied to each individual coefficient will generally change when the model is estimated at a more spatially and temporally aggregated level. A second possible source of bias arises from correlation between the aggregated prices and the aggregated day-of-sample and city fixed effects that appear in the disaggregated model. While this bias will not occur in panel data models that retain a complete set of two-way fixed effects, many of the studies in the existing literature estimating time-series or cross-sectional models or panel models with incomplete fixed effects may be subject to this form of bias. The final source of potential bias arises when correlation between the average prices and error term is created as a result of aggregation. The typical OLS identification assumptions for a disaggregate demand model would require the price in a given city on a given day to be uncorrelated with the unobserved demand shock in that city on that day. However, correlation between prices and demand shocks on other days or in other cities can cause the prices and errors in the aggregated panel data to be correlated, resulting in biased estimates.

Comparing our daily city-level estimates with those obtained using aggregated data, we are able to estimate the magnitude of the bias resulting from each of these three components. The results reveal how the primary source of bias differs depending on the dimension and degree of aggregation. Observed biases are largest in time series models where time-period fixed effects can no longer be used to control for demand differences over time.

The sources of bias identified in our decomposition and the magnitudes suggested in our aggregated regressions help to provide a more systematic explanation of why studies using different methodologies have obtained different elasticity estimates. Consistent with our findings, recent time-series studies including Hughes, Knittel,

and Sperling (2008) and Park and Zhao (2010) tend to produce very inelastic estimates, potentially caused by significant positive bias. Studies estimating panel regressions (though at a more aggregated level than ours) including Davis and Kilian (2011) and Li, Linn, and Muehlegger (2014) are not subject to the most significant source of bias and report more elastic estimates. However, these panel estimates are still lower than those obtained in our daily city-level regressions. Our decomposition provides guidance to researchers in evaluating whether particular market environments or data sources may be more susceptible to the forms of correlation that are likely to result in bias when estimating demand models using aggregate data and is sufficiently general that it can be used to evaluate almost any form of aggregation.

Accurate estimates of gasoline demand elasticity are crucial to making effective public policy decisions. As we discuss in Section VI, the stronger demand response revealed in our analysis suggests, for example, that supply restrictions, such as a refinery or pipeline outages, are likely to have a significantly smaller impact on gasoline prices than might otherwise be concluded based on estimates from earlier studies. Consequently, one might conclude that policies used to reduce fuel price volatility, such as the management of the US Strategic Petroleum Reserve, may be both less necessary and less effective than previously thought.

Like other recent studies of gasoline demand, our empirical approach obtains estimates of the short-run elasticity of gasoline demand and is not ideally suited for evaluating the long-run response to sustained changes in gasoline prices or tax levels. However, the channels through which gasoline usage responds to prices in the short run (reduced miles driven, increased public transit usage, heavier usage of more fuel-efficient vehicles in multi-vehicle households, etc.) remain important components of demand response in the longer run. Moreover, because robust estimates of long-run gasoline demand elasticity are difficult to identify empirically, many of these policy analyses still rely partially or entirely on short-run elasticity estimates. The long-run estimates that do exist may also be subject to many of the same sources of potential aggregation bias we demonstrate here. As a result, we believe that the more responsive gasoline demand revealed in our analysis also has important implications regarding the effectiveness of long-term policies such as carbon taxes and fuel taxes or cap-and-trade programs aimed at reducing greenhouse gas emissions.

I. Approaches to Estimating Gasoline Demand

A. Model Specification

There are a number of survey articles available (including Dahl and Sterner 1991, Goodwin 1992, Espey 1998, and Basso and Oum 2007) that summarize and analyze the literature on gasoline demand estimation. Nevertheless, it is helpful to provide a brief overview of some of the benefits and limitations of a number of commonly used empirical approaches in order to motivate our analysis and highlight the contribution of our study.

Most studies of aggregate gasoline demand estimate a simple log-linear model of quantity as a function of the gasoline price and other variables, such as average

income to control for shifts in demand. This specification is chosen mostly for convenience, so that the coefficient on price represents an estimate of demand elasticity. A number of studies have investigated alternative functional forms, and some find alternative forms to have a better fit (e.g., Hsing 1990), but generally the resulting elasticity estimates are not substantially impacted by specification (see Sterner and Dahl 1992, and Espey 1998).

Such “static” demand models can be appropriate as long as the researcher believes that the price and quantity data in the sample represent observations from a stable relationship. In situations where it is thought that gasoline demand may take multiple periods to adjust to price changes, researchers often adopt a more “dynamic” approach by including a lagged dependent variable (the parsimonious partial adjustment model approach) or specifying a distributed lag model (the more flexible approach) in which lagged values of prices and other control variables are also included in the demand equation. Coefficients on lagged prices capture how demand temporarily deviates from the longer run equilibrium relationship as adjustment occurs, allowing the sum of the coefficients on current and lagged prices to more accurately capture the overall elasticity of demand in the longer run.

When estimated using more aggregate (e.g., quarterly or annual) data, these dynamic models are typically used to separately identify the short-run demand response and long-run demand response (which would include adjustments made to where people live and work or what cars are purchased). In Section IVB, we also adopt a dynamic distributed lag model in order to estimate response over different time horizons. However, given our daily data, the lagged prices serve to separate the response in the days immediately following the price change from the more persistent response that remains weeks or months after the price change. In static models (like our baseline specification), the single price coefficient captures an average demand response over the entire duration of the price change. Estimating the dynamic model allows us to both examine immediate demand responses that are not identified in more highly-aggregated studies and confirm that the elasticity estimates obtained in our static model accurately reflect the same more persistent short-run demand responses that previous studies have attempted to investigate.

B. Identification

Regardless of functional form, demand specifications are most typically estimated using OLS. As a result, if sufficient controls cannot be included in the demand equation, endogeneity of the price variable can arise as a result of supply side responses to unobserved demand shifts. The instrumental variables approaches often adopted in other contexts are rarely used in studies of gasoline demand due to a lack of credible instruments. Time-series studies and even some panel-data studies rely almost exclusively on macroeconomic control variables, like average income and seasonal dummy variables (when relevant), which are certainly jointly determined with aggregate gasoline demand but do not perfectly predict fluctuations in gasoline usage. In this case, it is reasonable to expect this endogeneity bias to make price elasticity estimates more inelastic than they should be. Higher levels of data aggregation tend to exacerbate this problem, making it difficult to include sufficient

control variables and still have enough variation left in the data to identify demand. Our identification strategy leverages the availability of high-frequency panel data to include extensive sets of time and cross-sectional fixed effects in order to eliminate potential sources of endogeneity.

Two main factors determine the severity of bias that will result from price endogeneity in OLS estimates of demand: the prevalence and magnitude of unobserved demand shocks relative to unobserved supply shocks and the elasticity of the supply curve. If demand shifts are relatively small or the relevant supply curve is fairly flat, then observed variation in price will mainly be the result of upward and downward shifts in the supply curve. Correspondingly, our empirical specifications attempt to include fixed effects that will remove most of the unobserved variation in demand while leaving a sufficient level of supply variation to accurately identify the demand elasticity. Moreover, given the nature of the gasoline storage and distribution infrastructure, the daily supply of gasoline is likely to respond very elastically to any remaining daily demand shocks that are not absorbed by the included fixed effects.⁵

Most of the variation in daily city-level gasoline demand is likely to come from persistent differences across cities in per capita gasoline consumption and from economic changes that influence gasoline demand over time nationwide. These demand shifts represent the primary sources of potential endogeneity, and so our baseline demand specification controls for these by including both city and day-of-sample fixed effects. The fixed effects also remove a significant amount of supply variation including that generated by fluctuations in crude oil prices over time or by persistent regional differences in costs, competition, environmental restrictions, or gasoline tax levels. However, temporary regional gasoline supply shocks resulting from refinery outages or pipeline disruptions or regional variation in seasonal gasoline content restrictions can create significant variation in the relative price levels across cities that will not be absorbed by fixed effects. Our identification of demand relies largely on these periodic shifts in local supply. While it is possible that demand shifts remain even with city and day-of-sample fixed effects included in the model, these will mostly come from predictable seasonal patterns or day-of-week purchase patterns, which can be adequately planned for through adjustments to inventories, refinery production, or pipeline delivery schedules. Any additional idiosyncratic day-to-day fluctuations in local purchasing are likely small enough to be easily accommodated using local terminal inventories, so that relative price fluctuations not absorbed by city and day-of-sample fixed effects should entirely reflect localized supply-side cost shocks. In other words, within the context of our model, the relevant daily supply curve is likely to be almost perfectly elastic, minimizing any potential for endogeneity bias in our OLS estimates. For robustness, however, we also estimate and obtain similar results from alternative specifications that include city-specific sets of month-of-year and day-of-week fixed effects to control for potential differences

⁵Based on average inventory levels and daily consumption, there is typically enough gasoline stored at local distribution terminals and refineries to supply four weeks worth of demand (for more information see the US Energy Information Administration's Petroleum Supply Monthly). In addition, gas stations often have several days of inventory in their underground tanks. Day-to-day fluctuations in demand can easily be supplied from a combination of these sources, allowing intertemporal arbitrage that minimizes the likelihood of any substantial price response.

across cities in seasonal or weekly demand patterns, yet still preserve the supply variation resulting from unexpected refinery or pipeline disruptions.

While studies using aggregate data generally must resort to controlling for demand shocks with observable controls and/or a more limited set of fixed effects, several recent papers have attempted to identify and utilize instrumental variables to more accurately model demand. Hughes, Knittel, and Sperling (2008) estimate a specification of nationwide monthly gasoline demand using crude oil production disruptions as instruments for monthly gasoline prices but find the resulting elasticity estimates to be nearly indistinguishable from those obtained in their baseline OLS specifications. Davis and Kilian (2011) utilize monthly state-level aggregate gasoline consumption and average prices to estimate a first-differenced model in which changes in state gasoline tax rates serve as instrumental variables for changes in gasoline price. Their preferred IV estimate suggests a demand elasticity of -0.46 (s.e. = 0.23), while the corresponding OLS monthly state-level panel regression produces a substantially less elastic estimate of -0.19 (s.e. = 0.04), and an estimation using data aggregated to the national monthly time-series level produces an even smaller elasticity of -0.09 (s.e. = 0.04).⁶ Davis and Kilian's IV estimate is much closer to our disaggregated elasticity estimates, suggesting that both their IV approach and our disaggregated panel fixed effects approach may be overcoming the potential simultaneity bias caused by baseline demand differences over time.⁷ While we don't observe enough state tax changes during our sample to consider such an instrument, as another robustness check we do estimate an IV specification (described in online Appendix D) using regional wholesale spot gasoline prices as an instrument for local retail prices and obtain results similar to our OLS regressions.

II. Retail Gasoline Price and Expenditure Data

Our data contains daily gasoline price and expenditure data for 243 metropolitan areas throughout the United States from February 1, 2006 to December 31, 2009. For each city, average daily retail prices of unleaded regular gasoline are obtained from the American Automobile Association's (AAA) Daily Fuel Gauge Report. The prices reported by AAA are provided by the Oil Price Information Service (OPIS), which constructs the city average prices using station-level prices collected from fleet credit card transactions and direct feeds from gas stations.⁸

Our expenditure data were obtained from the financial services company Visa Inc. The data reflect the total dollar amount of purchases by all Visa debit and credit card users at gas stations within a city on a given day. As with the price data, cities are defined based on the geographic definition of the associated Metropolitan

⁶Coglianese et al. (2015) point out that the IV estimate of Davis and Kilian (2011) may be biased as a result of consumers anticipating the tax change and buying more gas in the month before the tax increase than in the month after. When one month leads and lags are included to control for this, the elasticity estimate falls to -0.36 (s.e. = 0.24), nearly identical to the estimate we obtain from our frequency of purchase model.

⁷An important caveat is that their measure differs from ours in that it focuses on the demand response occurring during the month in which the corresponding state gasoline tax rate change occurs, which is likely to reflect a shorter run elasticity than is captured in our baseline model.

⁸The OPIS price survey is the most comprehensive in the industry and is commonly used in research on gasoline pricing.

Statistical Area (MSA). In addition to total citywide expenditures, the Visa data also include the number of gasoline transactions or purchases taking place at gas stations in each city during each day. This allows us to separate the daily probability of purchase from the daily demand for gasoline. We also observe the total number of Visa cards that are actively purchasing (any product) within the current month. We use this as a measure of the total population of cards at risk of recording a gasoline purchase during each day of that month.⁹

Although the data has many advantages, Visa does not directly observe the price paid at the pump or the quantity of gasoline purchased by the customer. Instead, we construct a measure of the total quantity of gasoline purchased in a particular MSA on a particular day by dividing the total expenditures of Visa card customers at gas stations by the average regular unleaded gasoline price in the city on that day. This approximation raises several potential issues which we explore and address below.

First, in dividing total gas station expenditures by the regular gasoline price, we are ignoring the fact that around 15 percent of gas purchases are for mid-grade or premium gasoline, which sell at higher prices. If the fraction of regular-grade purchases were fairly constant over time, we would not expect this unobserved price difference to impact our demand elasticity estimates. However, the results of Hastings and Shapiro (2013) suggest that some consumers may substitute from premium to regular grade gasoline when prices increase. In online Appendix B, we discuss this possibility in more detail, derive an expression for the potential bias, and use the estimates of Hastings and Shapiro (2013) to show that any such bias is very likely to be negligible in our context.

A second and potentially more important source of bias arises from the fact that total gas station expenditures also include some non-gasoline purchases. Simply dividing total expenditure by the gasoline price to produce a measure of quantity ignores this possibility, causing expenditures to appear more elastic to gasoline price changes. Even if the prices and demand for non-gasoline items are not correlated with the price of gasoline, dividing these expenditures by the price of gasoline will mechanically generate an elasticity of -1 for the non-gasoline portion of the transaction.¹⁰ In general, the share of revenues generated by non-gasoline items is small but nontrivial. According to the 2007 US Economic Census, gasoline stations in the United States receive just over 21 percent of their total revenues from nonfuel sales.¹¹ Fortunately, biases from nonfuel purchases are only a concern for in-store transactions, and our data includes the daily city-level expenditures and number of transactions separately for pay-at-pump and in-store purchases. Pay-at-pump purchases represent over 76 percent of total expenditures and over 64 percent of all transactions in our data.¹² Estimating gasoline demand using only pay-at-pump

⁹At two points during the sample period (August 1, 2006 and August 1, 2007), Visa made small adjustments to their merchant classification methodology, which produce discrete jumps in our measure of gas station expenditures in some cities. To correct for this, we estimate all models with additional city-specific data-period fixed effects allowing the average expenditures and transaction counts to differ before, between, and after these two dates.

¹⁰See online Appendix C for a more complete discussion of the resulting bias.

¹¹Most of these nonfuel revenues come from food, cigarettes, and alcohol. Fuel sales often generate less than half of a station's profits, but given the high volume sold, it still represents the vast majority of station revenues.

¹²On average, pay-at-pump transactions are larger (in dollar value) than in-store transactions. The most likely explanation is that some in-store transactions include only nonfuel items, which tend to be less expensive than the

transactions gives an alternative estimate of elasticity that is not subject to this bias and may give some indication of the magnitude of the bias for in-store transactions. In addition, in online Appendix C, we derive the magnitude of the bias that would be expected when non-gasoline expenditures are present in the data, and use outside estimates of non-gasoline expenditures to construct bias-corrected elasticity estimates. The results are consistent with the notion that estimating elasticities using pay-at-pump purchases fully eliminates (or perhaps even overcompensates for) any bias resulting from non-gasoline expenditures.

A third and related concern involves the possibility that the gasoline price level may influence consumers' decision of whether to purchase with cash or credit, which, consequently, may also influence whether they purchase in the store or at the pump. Industry experts and consumer surveys (NACS 2007) regularly reveal that consumers are somewhat more likely to purchase with credit cards when gas prices increase. Although the magnitude of such an effect is unclear, it implies that our observed credit card purchases may appear less elastic than overall gasoline demand in response to a price change. In addition, if a higher gasoline price causes consumers to substitute from paying with credit in the store to paying with credit at the pump, perhaps because they are less likely to be also purchasing food, then pay-at-pump purchases will be even less elastic than overall credit card gasoline purchases. We do not have the necessary information or data to attempt to correct for these effects, and so we consider the estimates obtained from our pay-at-pump data to represent a conservative estimate of the true elasticity of gasoline demand. In the interest of completeness our analysis will present estimates utilizing all purchases as well as estimates using only pay-at-pump purchases. However, we will generally take the more conservative approach of using pay-at-pump purchase data in our alternative specifications and supplementary analysis.

Although our data only contain gasoline purchases made on Visa credit and debit cards, we believe them to be reasonably representative of all gasoline buyers using credit or debit cards. If cash buyers exhibit different purchase behavior, this will not be captured in our elasticity estimates. However, as roughly 70 percent of all gasoline purchases are made using a credit or debit cards (NACS 2007), our findings reflect the demand responsiveness of the vast majority of gasoline buyers.

A. Descriptive Statistics

Before we begin our empirical analysis, it is helpful to highlight some important features of the data. First, the price data reveal significant idiosyncratic fluctuation across cities. Though prices in all cities are impacted by common factors like world oil prices, there are many other factors that influence prices locally. Persistent price differences across states arise as a result of differences in gasoline tax rates or in the blends of gasoline that are required. More importantly, significant transitory differences in daily prices across the MSAs arise frequently during our sample period. Figure 1 compares retail price fluctuations in Los Angeles, Chicago, and New York

typical gasoline purchase. Unfortunately, our data do not allow us to examine the distribution of individual transaction amounts since we only observe the total expenditure for the day in each city.

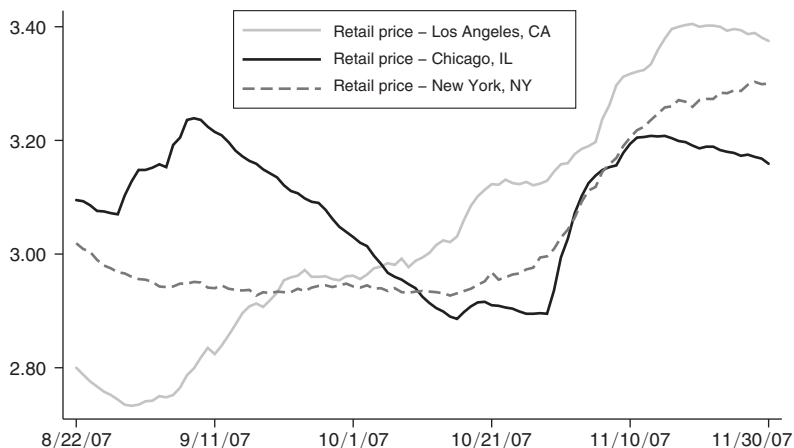


FIGURE 1. DAILY AVERAGE RETAIL GASOLINE PRICES FOR SELECTED CITIES

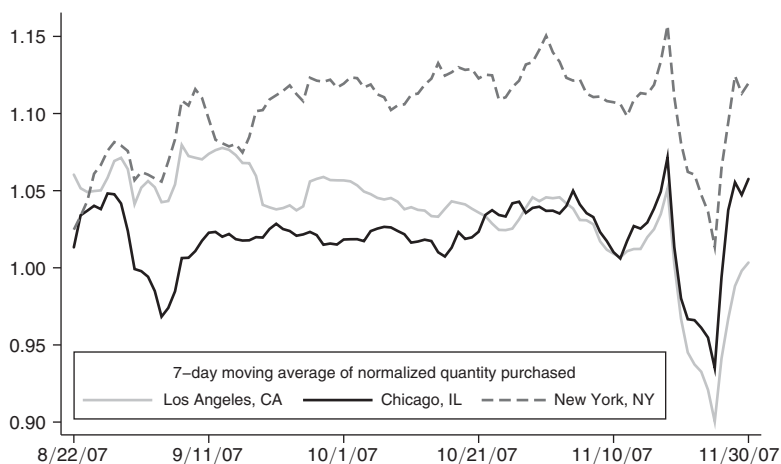


FIGURE 2. SEVEN-DAY MOVING AVERAGE OF TOTAL QUANTITY PURCHASED FOR SELECTED CITIES NORMALIZED BY THE CITY AVERAGE

over a typical 100 day period in 2007. It is clear that daily city-level prices provide much richer price variation than monthly data with which to study demand response.

The daily quantity of gasoline purchased also exhibits different patterns across MSAs, presumably due to both independent retail price movements and other city-specific events. Figure 2 presents a normalized seven-day moving average of this measure of the daily quantity purchased for the same three cities depicted in Figure 1 over the same period.¹³ The daily quantities for each MSA are normalized by the average quantity purchased in that MSA over the sample period. Just as with the prices, the quantities move together at times but also exhibit significant differences.

¹³ A moving average of daily quantity is used here to eliminate the strong within-week purchase patterns that are described below.

As an example of the type of variation that will be important for our empirical identification of demand, note that as the daily average price in California increases from being the lowest to the highest of the three cities by the end of the time period in Figure 1, the normalized quantity demanded falls from being the highest to the lowest over the same time period in Figure 2. The fact that these relative price and quantity differences between cities often persist over many months illustrates why the elasticity estimates from our baseline (static) demand model should be interpreted to reflect demand responses occurring over weeks or months even when estimated using daily data.

Daily expenditure data also allow us to examine high frequency features of gasoline purchase patterns. Gasoline demand is known to exhibit strongly seasonal variation, and our data reflects this general pattern. We are also able to document a very strong within-week pattern in gasoline purchasing behavior. Our data show that consumers buy roughly 17 percent more gasoline on Fridays than the daily average and buy 15 percent less on Sundays than the daily average. Within each city, this pattern varies to some extent, but Friday is always the highest expenditure day, and Sunday is always the lowest expenditure day.

Variation in the total expenditures across days can result either from fluctuations in the number of transactions that occur or from fluctuations in the amount people purchase per transaction. The within-week pattern in expenditure per transaction is notably different from that of overall expenditures, with Mondays and Fridays exhibiting the largest purchase sizes. However, expenditure per transaction exhibits much less day-to-day variation overall, revealing that the within-week pattern observed in total expenditures results largely from fluctuations in the number of transactions occurring in each day. Figures reporting average expenditures by day of week and average expenditures per transaction by day of week are presented in online Appendix F.

III. Model of Consumer Demand and Purchase Behavior

Because we are working with daily data, the effect of price on the amount of gasoline purchased may be very different from the effect on the amount of gasoline people are actually demanding at any given time. Consumers can buy and store gasoline in their car, which implies that a consumer's daily demand for gasoline can differ from its expenditures on gasoline. This section presents a theoretical model that recovers an estimate of the daily price elasticity of the unobserved demand for gasoline from data on the daily number of purchases and expenditures on gasoline for each MSA. A latent customer-level daily demand for gasoline and daily purchase probability give rise to an econometric model for customer-level daily gasoline expenditures that we then aggregate to the MSA level.

Suppose the daily demand for each customer in a city j on a day d takes the form

$$(1) \quad d_{jd} = \exp(\alpha_j + \lambda_d + \beta \ln(p_{jd}) + \epsilon_{jd}),$$

where α_j is a fixed-effect for MSA j , λ_d is the fixed-effect for day-of-sample d , p_{jd} is the price of gasoline for day d in region j , and β is the price elasticity of demand. For

each j , the ϵ_{jd} are a sequence of unobserved mean-zero random variables that may be heteroscedastic and correlated over time within each MSA, but are distributed independently across MSAs, and are independent of p_{jd} . Consumers must periodically purchase gasoline to satisfy this level of daily usage. The probability that a consumer in MSA j purchases gasoline on a day d is assumed to equal

$$(2) \quad \rho_{jd} = \gamma_j + \delta_d,$$

where γ_j is a fixed-effect for MSA j , and δ_d is the day-of-sample fixed effect for day d . We assume that the expenditure on gasoline during day d by each customer in MSA j , e_{jd} , is related to the consumer's daily purchase probability and daily gasoline demand through the following relationship:

$$(3) \quad e_{jd} = \frac{p_{jd}d_{jd}}{\rho_{jd}}.$$

This model implies that the actual quantity of gasoline purchased (if purchase occurs) times the daily probability of purchase is equal to the daily quantity demanded by that customer. Because our data is at the MSA level, we aggregate the customer-level model of daily gasoline expenditures over the total number of customers in MSA j during day d , N_{jd} . The number of customers in MSA j during day d making a gasoline purchase is equal to n_{jd} . Therefore, E_{jd} , total gasoline expenditures during day d for MSA j can be expressed as

$$(4) \quad E_{jd} = e_{jd}n_{jd} = \frac{p_{jd}d(p_{jd}, \epsilon_{jd})n_{jd}}{\rho_{jd}}.$$

Because we observe the total number of active Visa cards (N_{jd}) in MSA j during day d , and the total number of gasoline transactions (n_{jd}), n_{jd}/N_{jd} is an unbiased estimate of ρ_{jd} , the probability of purchase for MSA j during day d . Accordingly, we can estimate the parameters of equation (2) using OLS applied to

$$(5) \quad \frac{n_{jd}}{N_{jd}} = \gamma_j + \delta_d + \nu_{jd},$$

where the ν_{jd} are a sequence of mean-zero random variables that may be heteroscedastic and correlated with ϵ_{jd} , and over time within each MSA, but are distributed independently across MSAs. We can use the fitted values $\hat{\rho}_{jd} = \hat{\gamma}_j + \hat{\delta}_d$ to obtain a consistent estimates of the ρ_{jd} . Substituting the estimated purchase probability into equation (4) and taking logs generates our econometric model of gasoline expenditures:

$$(6) \quad \ln(E_{jd}) = \alpha_j + \lambda_d + (\beta + 1) \ln(p_{jd}) + \ln(n_{jd}) - \ln(\hat{\rho}_{jd}) + \epsilon_{jd}.$$

This model can alternatively be expressed in terms of the quantity purchased:

$$(7) \quad \ln(Q_{jd}) = \alpha_j + \lambda_d + \beta \ln(p_{jd}) + \ln(n_{jd}) - \ln(\hat{\rho}_{jd}) + \epsilon_{jd}.$$

The empirical model in equation (7) makes it possible to identify the underlying MSA-level elasticity of demand for gasoline (β) using data on prices, the quantity purchased, and the number of transactions. In equations (1) and (2), the demand and probability of purchase are both assumed to vary by city and day of sample, but different combinations of fixed effects can easily be used to generate alternative specifications for each of these functions. We will also consider a specification that includes both lagged and current prices in the demand and purchase probability equations.

IV. Estimation and Results

A. Frequency of Purchase Model

We begin by estimating city-level demand for gasoline using the model expressed in equation (7). Results are reported in Table 1. Following the discussion of data concerns in Section II, our main specifications are estimated using pay-at-pump purchases only, but we also report the results when all purchases are used. The model implies that coefficients on $\ln(n_{jd})$ and $\ln(\hat{\rho}_{jd})$ should be 1 and -1 , respectively. We estimate the model with this restriction imposed and without it. To facilitate comparisons with earlier studies of gasoline demand, we also include (in columns 1 and 4) estimates from a basic log-linear aggregate demand model. For the basic model, we report heteroskedasticity consistent standard error estimates that allow for arbitrary serial correlation within each city. Standard error estimates for the purchase model are generated using a nonparametric bootstrap to account for the fact that the predicted probability of purchase is estimated in a first-stage regression.¹⁴

When estimated using pay-at-pump transactions only, the model with all restrictions yields an elasticity estimate of -0.36 , while the unrestricted model produces a similar elasticity estimate of -0.30 . The unrestricted coefficient on $\ln(n_{jd})$ is very close to one, but the coefficient on $\ln(\hat{\rho}_{jd})$ is very close to zero; far from the -1 implied by the theory. This may be because the fixed effects absorb most of the variation in the probability of purchase (given the functional form specified in equation (7)), and any variation left may be measured with error. The estimated price elasticity is still similar to that from the model with restrictions imposed. Estimates of both the restricted and unrestricted models are somewhat more elastic when using data from all purchases as would be expected if the inclusion of in-store purchases

¹⁴ Given that there are C cities and D days in the sample, the procedure first resamples with replacement C pairs of OLS residual vectors of length D from the purchase probability equation and the expenditure equation for the same city. These are combined with the OLS estimates of the fitted values for both equations to compute resampled purchase frequencies and expenditure levels vectors of length CD . The two-step estimation procedure is then repeated for the resampled purchase frequency equation and then for the expenditure equation with the logarithm of the fitted value from the estimated purchase frequency equation used as a regressor in the expenditure equation to obtain estimates for the parameter values for both equations for that resample. The sample variance of these re-samples is then used to compute the estimated standard errors for both parameter vectors.

TABLE 1—ESTIMATES OF BASELINE EMPIRICAL MODEL OF DEMAND

	Dependent variable = $\ln(\text{quantity}_{jd})$					
	Pay-at-pump only			All purchases		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\text{price}_{jd})$	-0.295 (0.031)	-0.364 (0.022)	-0.301 (0.006)	-0.378 (0.030)	-0.465 (0.022)	-0.408 (0.008)
$\ln(\text{number of transactions}_{jd})$		1	0.996 (0.002)		1	1.021 (0.004)
$\ln(\text{predicted probability of purchase}_{jd})$		-1	-0.003 (0.001)		-1	0.006 (0.002)

Notes: Day-of-sample fixed effects and city fixed effects are included in all specifications. Standard errors in columns 1 and 4 are robust and clustered to allow serial correlation within city. Standard errors for the remaining specifications are generated using a nonparametric bootstrap that allows errors to be serially correlated within a city and jointly distributed with the error term in the first-stage regression.

TABLE 2—ESTIMATES FROM ALTERNATIVE SPECIFICATIONS

	Dependent variable = $\ln(\text{quantity}_{jd})$			
	Pay-at-pump transactions			
	(1)	(2)	(3)	(4)
$\ln(\text{price}_{jd})$	-0.278 (0.011)	-0.271 (0.012)	-0.301 (0.006)	-0.351 (0.005)
$\ln(\text{number of transactions}_{jd})$	1.007 (0.006)	1.010 (0.006)	0.989 (0.002)	0.994 (0.002)
$\ln(\text{predicted probability of purchase}_{jd})$	0.009 (0.006)	0.011 (0.006)	-0.003 (0.001)	-0.002 (0.001)
Fixed effects:				
Day of week	X	X		
Month of sample	X			
Day of sample			X	X
City	X			
Month of year \times city			X	
Month of sample \times city		X		X

Note: Standard errors are generated using a nonparametric bootstrap that allows errors to be arbitrary serial correlated within a city and jointly distributed with the error term in the first-stage regression.

results in some bias from non-gasoline transactions. Nevertheless, the pattern of estimates across specifications is similar for the pay-at-pump and all-transaction samples. The coefficient estimates from the basic demand model in columns 1 and 4 are quite close to those generated by our frequency of purchase model.

Including time fixed effects helps to control for shifts in demand, but it can also mask supply shifts that could help to better identify demand elasticity. Hence, as an alternative specification, we estimate the demand model using month-of-sample fixed effects rather than day-of-sample. Day-of-week fixed effects are also included to control for the weekly pattern in demand suggested by online Appendix Figure F1. The resulting coefficient estimates, reported in Table 2, column 1, are very similar to those of the benchmark pay-at-pump specification, with the price elasticity

estimated to be -0.28 , suggesting that month-of-sample and day-of-sample fixed effects appear to work similarly in controlling for macroeconomic or gasoline market specific fluctuations that might impact gasoline demand at the national level.

In all the specifications estimated to this point, demand elasticities are identified off of city-specific variation in price and quantity purchased. However, large city-specific or regional demand shifts occurring over time have the potential to bias these elasticity estimates. Therefore, we estimate an additional set of specifications that include additional controls for city-specific demand patterns. Table 2, column 2 reports the results of a regression including month-of-sample effects for each city as well as day-of-week fixed effects. Column 3 contains the estimates from a model including a full set of national day-of-sample fixed effects in addition to the city-specific month-of-year effects to allow seasonal patterns in demand to differ across cities. Column 4 adds further flexibility by including both day-of-sample and city-specific month-of-sample fixed effects. The estimated elasticities vary somewhat across these models but are all of a similar magnitude to the baseline estimates in Table 1. It is important to note, however, that once separate month-of-sample fixed effects are included for each city they absorb any deviations from the national average that last longer than one month. As a result, the elasticity estimate in column 4 only reflects the response in demand that occurs in the four weeks following a change in price. This estimate is slightly more elastic than in the other specifications suggesting that very short run demand response is even more elastic than that which persists over a longer time horizon.¹⁵ In general, however, the results reveal that relatively elastic estimates of demand can be obtained using a variety of different sources of variation within the data and are not sensitive to a particular functional form.

Though the ability to include extensive fixed effects should minimize any potential endogeneity bias that might result from unobserved demand shocks, it is difficult to conclude with certainty that these demand shocks have been entirely eliminated. Unobserved city-specific demand shocks could still bias our elasticity estimates downward if local distribution terminals are not able to plan ahead or use inventories to costlessly absorb these differences (i.e., if the local supply curve is not highly elastic with respect to daily adjustments). As a final robustness check, we have estimated an instrumental variables specification that utilizes spot market (wholesale) gasoline prices from large regional refining centers (New York Harbor, the Gulf Coast, or Los Angeles) as instruments for local retail prices. Market-wide fluctuations in wholesale gasoline prices resulting from a combination of changes in demand and changes in crude oil and refining costs are captured by day-of-sample fixed effects, but differences in wholesale gasoline prices between regions still exhibit significant variation largely driven by unexpected refinery shocks. Under the assumption that this regional variation in spot prices is relatively unaffected by temporary city-specific demand fluctuations, using this IV approach could help to eliminate any remaining endogeneity generated by correlation between city-level prices and local demand shocks.

¹⁵This result is confirmed when we explore the differences in longer run versus shorter run demand response in more detail in Section IVB.

The results of these IV specifications (described in more detail and reported in online Appendix D) largely confirm that the robustness of the main OLS findings. The IV estimates of both the basic aggregate demand model and the frequency of purchase model are slightly more elastic but of a relatively similar magnitude to those reported in Tables 1 and 2.

What is most striking about these findings, in general, is that the elasticity estimates from both our frequency of purchase model and the basic log-linear aggregate demand model are consistently several times more elastic than those from comparable recent studies including Hughes, Knittel, and Sperling (2008), whose estimates range from -0.034 to -0.077 for the period 2000–2006, and Park and Zhao (2010), whose time-varying estimates range from around -0.05 in 2000 to around -0.15 in 2008. Our daily city-level purchase data clearly reveal a much greater degree of demand response than has been suggested by much of the literature.

It is important to highlight that the use of daily data rather than, say, monthly data does not imply that our elasticity estimates describe a “shorter run” demand response. The relevant response horizon of any elasticity estimate depends on the variation in prices used to identify the response parameter and when these movements occur relative to when demand is observed. Prices change in this market on a daily basis and consumers make purchase decisions on a daily basis, so behavior is likely to be more accurately represented using a model of daily demand for gasoline. However, as in most studies, our baseline model is static and does not allow the degree of demand responsiveness to change depending on how long it has been since a price change occurred, so price changes occurring several months ago are just as important as price changes occurring days or weeks ago in terms of identifying demand elasticity.¹⁶ As a result, we believe that our main elasticity estimates reflect the same type of consumer response that other studies attempt to measure using more aggregated monthly or quarterly data. In the next section, we leverage our daily data by relaxing these assumptions in the model to investigate if demand responds differently in the very short run.

B. Immediate versus Short-Run Demand Response

To this point, our models have assumed that prices influence gasoline demand entirely through the current gasoline price, implying that all demand response occurs immediately. In practice, however, it is not unusual for the elasticity of demand to change when considering different response horizons. Perhaps the most common of these situations occurs when consumers can hold inventories and choose to add to or withdraw from inventories in response to price changes even when their underlying consumption does not change. Gasoline consumers obviously hold small inventories of gasoline in their vehicle’s tank, so this behavior is feasible on a limited scale. Similarly, consumers may have the ability to postpone (or expedite) some necessary trips or utilize public transportation in response to a temporary increase (or

¹⁶Specifically, our model of daily demand implies that if the daily price rose by 10 percent and remained at this higher level, daily gasoline demand would remain lower by an amount equal to 10 percent times our demand elasticity for as long as the daily price remained 10 percent higher.

decrease) in price, regardless of how they change their overall driving habits. These types of behavior imply that, for a given gasoline price today, the amount of gasoline purchased today might depend on whether the price has been at or near its current level for a while or whether it was significantly higher or lower a few days or a few weeks ago.

The elasticity estimate from most gasoline demand models (including our baseline model) represents some average of these shorter run and longer run responses. However, with daily data, it is possible to separately identify these different responses by allowing demand to depend on past prices as well as current price levels. Moreover, by using the structure of our consumer purchase model and including past prices along with current prices in both the individual demand and purchase equations, we are able to decompose any potential short-run responses to examine whether consumers appear to be significantly altering gasoline usage or simply shifting when they make purchases in the days following a price change. If consumers are substituting away from driving in response to price increases then their daily demand may be influenced by past prices. If consumers are using their inventories of gasoline strategically, both current and past prices may influence a consumer's probability of purchase. We alter equations (1) and (2) to allow for these types of behavior. The demand for each customer in a city j on a day d can be specified as

$$(8) \quad d_{jd} = \exp\left(\alpha_j + \lambda_d + \beta \ln(p_{jd}) + \sum_{l \in L} \zeta_l \ln(p_{j,d-l}) + \epsilon_{jd}\right),$$

where $p_{j,d-l}$ represents the price l days prior to the current period and L represents the set of lags lengths included in the specification. Similarly, the probability of purchase can be expressed as

$$(9) \quad \rho_{jd} = \gamma_j + \delta_d + \psi \ln(p_{jd}) + \sum_{l \in L} \eta_l \ln(p_{j,d-l}).$$

Leaving the consumer purchase model from Section III otherwise unchanged results in the following final representation of the aggregate quantity purchased in city j on day d :

$$(10) \quad \ln(Q_{jd}) = \alpha_j + \lambda_d + \beta \ln(p_{jd}) + \sum_{l \in L} \zeta_l \ln(p_{j,d-l}) + \ln(n_{jd}) - \ln(\hat{\rho}_{jd}) + \epsilon_{jd},$$

where the predicted purchase probability can be estimated from an OLS regression of

$$(11) \quad \frac{n_{jd}}{N_{jd}} = \gamma_j + \delta_d + \psi \ln(p_{jd}) + \sum_{l \in L} \eta_l \ln(p_{j,d-l}) + \nu_{jd}.$$

In both the demand equation and the purchase probability equation we include log of the current price and the lagged log prices from each of the previous 5 days as well as longer lags of 10 and 20 days. Lags longer than 20 days are omitted as

TABLE 3—PURCHASE MODEL WITH LAGGED PRICES

	Traditional model	Purchase frequency model	
	Demand equation (1)	Purchase equation (2)	Demand equation (3)
$\ln(\text{price}_{j,d})$	-0.775 (0.081)	-0.009 (0.002)	-0.458 (0.009)
$\ln(\text{price}_{j,d-1})$	-0.607 (0.098)	-0.028 (0.002)	0.098 (0.008)
$\ln(\text{price}_{j,d-2})$	0.637 (0.091)	0.023 (0.001)	-0.002 (0.007)
$\ln(\text{price}_{j,d-3})$	0.393 (0.044)	0.014 (0.001)	0.052 (0.006)
$\ln(\text{price}_{j,d-4})$	0.184 (0.046)	0.003 (0.002)	0.020 (0.007)
$\ln(\text{price}_{j,d-5})$	0.055 (0.039)	0.003 (0.002)	-0.011 (0.007)
$\ln(\text{price}_{j,d-10})$	-0.076 (0.021)	-0.003 (0.001)	0.0002 (0.003)
$\ln(\text{price}_{j,d-20})$	-0.119 (0.021)	-0.001 (0.001)	0.004 (0.004)
$\ln(\text{number of transactions}_{j,d})$			0.995 (0.002)
$\ln(\text{predicted probability of purchase}_{j,d})$			-0.003 (0.001)
Total implied elasticity 20 days after a price change	-0.308	0.061	-0.296

Notes: City and day of sample fixed effects are present in all specifications. The dependent variable in column 1 is the log of the average quantity purchased at the pump per capita by Visa customers in city j on day d . Standard errors in column 1 are robust and clustered to allow arbitrary serial correlation within a city. The dependent variables columns 2 and 3 are the share of Visa customers purchasing at the pump and the log of the quantity purchased at the pump by Visa customers in city j on day d . Standard errors in columns 2 and 3 are generated using a nonparametric bootstrap that allows errors to be arbitrary serial correlated within a city and jointly distributed with the error term in the first-stage regression.

their inclusion requires the use of a shorter sample for estimation, but when price lags of 40 and 60 days are included, their coefficients are small in magnitude and do not substantially affect the estimates of the existing coefficients. For comparison we also estimate a similar version of the more traditional single-equation demand specification that includes lagged prices. Coefficient estimates from the traditional model are reported in column 1 of Table 3, and estimates from the demand and purchase probability equations from the frequency of purchase model are reported in columns 2 and 3. All specifications are estimated using only pay-at-pump purchases. The final row of the table includes the total implied elasticity of the probability of purchase or of demand response after 20 days.¹⁷ Estimates from the traditional demand model and the demand equation in the purchase model are directly comparable to the corresponding specifications without lags in Table 1.

¹⁷For the demand equations, this is simply the sum of all the log-price coefficients. For the probability of purchase equation this is the sum of all the log-price coefficients divided by the mean probability of purchase.

In the traditional demand model (column 1), the coefficients on the current and previous day's log price are negative and much larger in magnitude than the corresponding elasticity estimated without lags. The sum of the first two coefficient estimates in column 1 imply that the amount of gasoline purchased one day after a 1 percent price increase will be 1.38 percent lower than it would have been without the price increase. During the following three to four days, however, the amount purchased tends to increase sharply, back toward its original level, canceling out much of the very strong initial response in purchasing. The 10- and 20-day lags reveal that the price response becomes slightly stronger once again, several weeks after the price change. Adding together the coefficients of all the price lags in the regression gives the response of demand 20 days after a permanent price change. This sum of coefficients is reported in the last row of Table 3 and implies that the elasticity of demand response after 20 days is -0.31 , nearly identical to the elasticities of -0.30 identified in our baseline model with no lags.

Coefficient estimates from the frequency of purchase model in columns 2 and 3 of Table 3 reveal that the large response in the amount of gasoline purchased in the days following a price change results almost entirely from a temporary change in the probability of making a purchase rather than a change in gasoline demand or usage. The probability of purchase falls (rises) significantly on the day of and particularly on the day following a price increase (decrease). The coefficients on $\ln(p_{j,d})$ and $\ln(p_{j,d-1})$ imply that the purchase probability one day after a price change exhibits an elasticity with respect to price of around -1.12 all else equal.¹⁸ However, this response in the probability of purchase in the day of and the day after a price change is entirely counteracted over the following few days to leave the elasticity of the overall response of purchase probability after 20 days to be small and slightly positive at 0.06.

In the demand equation (column 3), the inclusion of lagged prices causes the coefficient on the current value of $\ln(p_{jd})$ to increase in magnitude, suggesting an even larger immediate demand response to price changes. As in the basic demand specification of column 1, however, the sum of the coefficients on the current and lagged values of $\ln(p_{jd})$ in the demand equation are very similar to the coefficient estimates for $\ln(p_{jd})$ when no lagged prices are included (Table 1, column 3). In other words, the total demand response to a price shock that lasts longer than a few days exhibits a demand elasticity of around -0.30 , nearly identical to the estimates in our baseline purchase model. The results also reveal a small additional response within the first few days of a price shock, consistent with the idea that consumers delay/expedite gasoline usage by a few days in response to price fluctuations, but this effect does not appear to be as large as those resulting from changes in the probability of purchase.

The results of both the traditional demand model and the purchase frequency model provide evidence of a stronger responsiveness to price changes in the very short run, but also confirm that these responses are *in addition to* a more persistent response, which very closely resembles that implied by our earlier specifications.

¹⁸The average probability of purchase by a cardholder on a given day is 0.033, so the elasticity of the probability of purchase is computed as $(-0.009 - 0.028)/0.033 = -1.12$.

We conclude from this that previous studies utilizing monthly or annual data did not obtain lower elasticity estimates a result of consumers being less responsive to price changes that persist over longer time periods. Instead, as we show in the next section, the use of such temporally aggregated data is likely to bias elasticity estimates causing demand to appear less responsive than can be revealed when using daily data.

V. Examining the Divergence from Previous Findings

Given the rather large discrepancy between our elasticity estimates and those found in other recent studies, we discuss in this section a number of differences in our analyses that could potentially explain this disparity.

A. Sources of Gasoline Consumption Data

Perhaps the biggest challenge in studying gasoline demand is finding an accurate measure of consumption. Nearly all available measures are recorded at a highly aggregated level and are likely to measure actual gasoline usage within the specified time interval with a considerable amount of error. The most common source used in recent time series or panel studies (e.g., Hughes, Knittel, and Sperling 2008; Park and Zhao 2010; Lin and Prince 2013) is the US Energy Information Administration's (EIA's) data on finished motor gasoline "product supplied." These data are constructed from surveys of refineries, import/export terminals, and pipeline operators, and the volumes reported reflect the disappearance of refined product from these primary suppliers into the secondary distribution system (local distributors and storage facilities). Each month the EIA reports amounts disappearing as the product supplied in each of the nation's five Petroleum Area Defense Districts (PADDs). Unfortunately, given distribution lags and storage capabilities, the amount of product flowing from secondary distributors to retailers and ultimately to consumers could differ substantially from the amount received by these suppliers in any given time period. In addition, to generate a measure that represents domestic gasoline usage, the EIA must net out from total production the estimated quantity of gasoline exported for use in other countries. This step provides yet another dimension for potential error, and created serious measurement issues in 2011 during a period of rapidly growing refined product exports.¹⁹

Another potential data source for gasoline consumption is the Federal Highway Administration (FHWA), which collects information from each state on the number of gallons of motor gasoline for which state excise taxes have been collected each month. This measure would appear to be more closely linked to consumption, and it is available at the state level rather than the PADD level. However, a significant amount of measurement error is generated by the fact that each state has its own procedures and systems for collecting this information. The point in the supply chain at which the fuel is taxed also varies across states. Some require taxes to be paid when

¹⁹See Cui (2012).

the distributor first receives the fuel, while others tax the volume of gasoline sold by the distributor. In fact, the FHWA includes in its publications the disclaimer that the reported volumes “may reflect time lags of six weeks or more between wholesale and retail levels.”²⁰

In contrast to the EIA and FHWA data, our measure of gasoline expenditures from Visa is recorded at the final step of the distribution process—when the consumer purchases the product from the retailer. This eliminates the possibility that changes in consumer purchase volumes are masked by additions or withdrawals from local storage. Moreover, it allows daily city-specific volumes to be observed and more accurately linked with contemporaneous local prices, providing for a more direct identification of demand response.

B. Estimating Demand Elasticity Using Aggregated Data

In general, using highly aggregated data can mask much of the temporal and geographic co-movements in prices and quantities that result from consumer demand response and also make it more difficult to empirically identify consistent estimates of such response. To illustrate this point, suppose the per capita daily demand for gasoline in MSA c during day d can be represented as

$$(12) \quad q_{cd} = D_{cd}(p_{cd}, X_{cd}) + \epsilon_{cd},$$

where p_{cd} is the price of gasoline in region c on day d and X_{cd} is the vector of characteristics of region c and day d that enter the demand function for that region and day. These daily demand functions for each MSA imply that Q_m , the national average daily per capita demand for gasoline during month m , is equal to the sum of total consumption ($q_{cd}N_{cd}$) across cities and days divided by the sum of the total population (N_{cd}) across cities and days:

$$(13) \quad Q_m = \frac{\sum_{d \in \mathcal{A}_m} \sum_{c=1}^C N_{cd} D_{cd}(p_{cd}, X_{cd})}{\sum_{d \in \mathcal{A}_m} \sum_{c=1}^C N_{cd}} + \frac{\sum_{d \in \mathcal{A}_m} \sum_{c=1}^C N_{cd} \epsilon_{cd}}{\sum_{d \in \mathcal{A}_m} \sum_{c=1}^C N_{cd}},$$

where C is the total number of MSAs in the sample and \mathcal{A}_m denotes the set of days in month m . This aggregation process implies that the national monthly demand for gasoline depends on the daily prices for all days during that month for all MSAs, rather than simply a single monthly national average price. Similarly, a monthly state-level average demand would depend on the daily prices for all days during that month for all MSAs in that state. However, this is not the model estimated in most empirical studies of aggregate demand. Typically, these studies only have access to aggregated price measures as well and, as a result, estimate something like

$$(14) \quad Q_m = \tilde{D}(\bar{p}_m, \tilde{X}_m) + \tilde{\epsilon}_m \quad \text{where} \quad \bar{p}_m = \frac{1}{A_m} \frac{1}{C} \sum_{d \in \mathcal{A}_m} \sum_{c=1}^C p_{cd}.$$

²⁰FHWA Highway Statistics (2010, table MF-33GA, footnote 1).

Given the data generating process in (12), the model in (14) will only hold under very specific assumptions, such as the case in which the $D_{cd}(p_{cd}, X_{cd})$ function is linear and identical across cities, and each ϵ_{cd} is uncorrelated not just with p_{cd} and X_{cd} but also with the values of these variables in other cities and periods over which aggregation occurs.²¹ If these conditions are not collectively satisfied in the markets being examined, then biases arising from aggregation have the potential to explain the large differences observed between our elasticity estimates and those from previous studies. Fortunately, we are able to investigate the impact of aggregation by using our data to create new datasets with varying levels of temporal and geographic aggregation. We construct daily datasets of state-level and nationwide total quantity purchased, and average price, as well as three monthly datasets at the city, state, and national levels. To facilitate a more direct comparison with other studies, we estimate basic log-log aggregate demand models and use aggregate per capita quantities calculated as the corresponding sum of the daily quantity purchased divided by the total number of Visa customers in the combined area. For consistency, average prices are also computed as a per capita weighted average across cities and days.

As in our main analysis, a complete set of time period and cross-sectional fixed effects are used whenever possible to control for shifts in demand. When using daily national time-series data, we include day-of-week and month-of-sample fixed effects. For the monthly national time series we are restricted to using month of year (i.e., seasonal) fixed effects, so per capita real personal disposable income is included as an additional control for demand shifts.²² This final specification is identical to that of Hughes, Knittel, and Sperling (2008).

The demand estimates for each level of aggregation appear in panel A of Table 4. The top rows report the results when estimated using pay-at-pump purchases only while the bottom rows report the results for all purchases. The first column reports the daily city-level results again for comparison. The next three columns contain panel regressions with varying levels of temporal and/or geographic aggregation. Price elasticity estimates from these specifications are all very similar to each other (between -0.22 and -0.25 for pay-at-pump purchases and between -0.30 and -0.34 for all purchases) and are less elastic than corresponding estimates from the disaggregated regression in column 1. The elasticity estimates from the two time series regressions in columns 5 and 6 are even smaller in magnitude, being indistinguishable from zero for pay-at-pump purchases and ranging from -0.12 to -0.14 for all purchases; much closer to the elasticities reported by Hughes, Knittel, and Sperling (2008) in their national time-series study.²³ Clearly, increasing levels of aggregation lead to less elastic estimates of demand, particularly when moving from panel to time series data.

²¹The various issues necessitating such strong assumptions will be detailed in Section VC.

²²Data on per capita personal disposable income comes from the Bureau of Economic Analysis.

²³In online Appendix E, we replicate the exact model of Hughes, Knittel, and Sperling (2008) using their data source, but from our 2006–2009 time period, and show the results to be very similar to our pay-at-pump national time-series estimates.

TABLE 4—REGRESSIONS USING AGGREGATED DATA

Geography:	City	City	State	State	National	National
Periodicity:	Daily	Monthly	Daily	Monthly	Daily	Monthly
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. log-log model: Dependent variable = ln(quantity per capita)</i>						
Pay-at-pump only:						
ln(price _{it})	-0.323 (0.025)	-0.246 (0.028)	-0.245 (0.059)	-0.215 (0.066)	-0.014 (0.067)	-0.002 (0.021)
ln(income _{it})						0.339 (0.171)
All purchases:						
ln(price _{it})	-0.405 (0.024)	-0.338 (0.028)	-0.325 (0.065)	-0.295 (0.073)	-0.137 (0.064)	-0.130 (0.016)
ln(income _{it})						0.458 (0.146)
<i>Panel B. Linear model: Dependent variable = quantity per capita</i>						
Pay-at-pump only:						
price _{it}	-0.046 (0.004)	-0.037 (0.004)	-0.030 (0.009)	-0.025 (0.010)	-0.009 (0.008)	-0.001 (0.003)
income _{it}						0.005 (0.002)
Implied elasticities:						
price	-0.312	-0.247	-0.201	-0.164	-0.062	-0.008
income						0.379
All purchases:						
price _{it}	-0.077 (0.006)	-0.064 (0.007)	-0.058 (0.016)	-0.050 (0.018)	-0.035 (0.009)	-0.027 (0.003)
income _{it}						0.010 (0.002)
Implied elasticities:						
price	-0.379	-0.313	-0.285	-0.245	-0.174	-0.133
income						0.613
Fixed effects:						
Day of sample	X		X			
Day of week					X	
Month of sample		X		X	X	
Month of year						X
City	X	X				
State			X	X		

Notes: Standard errors for panel specifications are robust and clustered at the level of the cross-sectional unit to allow for arbitrary serial correlation. Standard errors for time-series specifications are estimated using the Newey and West (1987) procedure and are robust to the forms of serial correlation they consider. Implied elasticities are calculated at the sample-wide mean values of quantity per capita, price, and income.

C. Sources of Aggregation Bias

In this section, we decompose the impact of using different forms of aggregate price and quantity data on the resulting elasticity estimates. Rather than assess the impact of aggregating data with a log-log demand specification and have to deal with the fact that the sum across days or cities of the log price or quantity is not the

same as the log of the sum, we consider a linear demand specification that can be aggregated using linear operations. To evaluate the importance of this functional form choice, we begin by estimating these linear demand models at different levels of aggregation, just as was done with the log-log model. In panel B of Table 4, we report both coefficient estimates and the implied elasticities evaluated at mean values of quantity per capita, price, and income per capita. Although the linear and log-log specifications have the potential to yield different results, for our data, the levels of the elasticity estimates evaluated at the sample mean of the data and the degree to which they change across increasing levels of aggregation are quite similar. We conclude from this comparison that the differences in estimates resulting from data aggregation apparent in our setting is relatively pervasive and is not the result of a particular functional form assumption or a particular mode of aggregation.

Although most studies of gasoline demand focus on a model in which all locations are assumed to have the same price coefficient, our decomposition analysis considers a more general data generating process in which the daily city-level per capita demand for gasoline depends on a city-specific fixed effect, a day-of-sample fixed effect, and the price of gasoline with a coefficient that may differ across cities:

$$(15) \quad Q_{cd} = \lambda_c + \alpha_d + \beta_c p_{cd} + \epsilon_{cd},$$

where $c = 1, \dots, C$; $d = 1, \dots, D$; $C = 241$ is the number of cities, and $D = 1,430$ is the number of days in our sample.

To write this model in matrix form we utilize the following standard notation: $\mathbf{1}_N$ represents an $N \times 1$ vector on 1s, \mathbf{I}_N is an $N \times N$ identity matrix, and $\mathbf{M}_X = \mathbf{I}_N - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ represents the residual maker matrix of dimension N , where N is the number of rows in \mathbf{X} . Let \mathbf{P} equal a $CD \times C$ matrix with the c th column having zeros everywhere but the c th block of length D , which is replaced with the $D \times 1$ vector $\mathbf{p}_c = (p_{c1}, p_{c2}, \dots, p_{cD})'$, $\mathbf{Q} = (\mathbf{Q}'_1, \mathbf{Q}'_2, \dots, \mathbf{Q}'_C)'$, $\mathbf{Q}_c = (Q_{c1}, Q_{c2}, \dots, Q_{cD})'$, $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_C)'$, $\boldsymbol{\epsilon} = (\epsilon'_1, \epsilon'_2, \dots, \epsilon'_C)'$, and $\boldsymbol{\epsilon}_c = (\epsilon_{c1}, \epsilon_{c2}, \dots, \epsilon_{cD})'$. In addition, $\boldsymbol{\lambda}$ is a $C \times 1$ vector of the city fixed effects, $\boldsymbol{\alpha}$ is a $D \times 1$ vector of the day-of-sample fixed effects, $\boldsymbol{\delta} = (\boldsymbol{\lambda}', \boldsymbol{\alpha}')$, and $\mathbf{Z} = [\mathbf{I}_C \otimes \mathbf{1}_D | \mathbf{1}_C \otimes \mathbf{I}_D]$ a $CD \times (C + D)$ matrix. In terms of this notation (15) becomes

$$(16) \quad \mathbf{Q} = \mathbf{Z}\boldsymbol{\delta} + \mathbf{P}\boldsymbol{\beta} + \boldsymbol{\epsilon}.$$

OLS can be applied to equation (16) to recover an estimate of each city's β_c . Let \mathbf{b} , \mathbf{d} , and \mathbf{e} equal the OLS estimates of $\boldsymbol{\beta}$, $\boldsymbol{\delta}$, and $\boldsymbol{\epsilon}$, respectively.

It is more common to estimate a restricted model of gasoline demand in which all locations are assumed to have the same price coefficient, implying the following true model:

$$(17) \quad \mathbf{Q} = \mathbf{Z}\boldsymbol{\delta}^* + (\mathbf{P}\mathbf{1}_C)\boldsymbol{\beta}^* + \boldsymbol{\epsilon}^*,$$

where $\mathbf{P}\boldsymbol{\nu}_C$ is a $CD \times 1$ vector that equals $\mathbf{p} = (\mathbf{p}'_1, \mathbf{p}'_2, \dots, \mathbf{p}'_C)'$ and β^* is a scalar. The Appendix derives the following expression for the OLS estimate of β^* in terms of \mathbf{b} , the OLS estimates of the vector $\boldsymbol{\beta}$,

$$(18) \quad b^* = (\boldsymbol{\nu}'_C \mathbf{P}' \mathbf{M}_Z \mathbf{P} \boldsymbol{\nu}_C)^{-1} \boldsymbol{\nu}'_C \mathbf{P}' \mathbf{M}_Z \mathbf{P} \mathbf{b} = (\mathbf{p}' \mathbf{M}_Z \mathbf{p})^{-1} \mathbf{p}' \mathbf{M}_Z \mathbf{P} \mathbf{b}.$$

Note that if (16) is the true data generation process, then the OLS estimate b^* is a weighted average of the OLS estimates of β_c in which the relative weighting of each city's β_c is based on the degree of within-city variation in the residual prices (after controlling for city and day-of-sample fixed effects). As we demonstrate below, different data aggregations will yield different weighted sums of the OLS estimates of the β_c as a component of the estimated slope coefficient.

Using this notation, we can also derive expressions for the OLS estimate of the slope coefficient with different quantity and price aggregates. Consider a regression model aggregated to the month and state level. Suppose there are S states and M months. This can be accomplished by pre-multiplying \mathbf{Q} and \mathbf{p} by a $SM \times CD$ aggregation matrix \mathbf{A} defined in the Appendix. Let $\mathbf{p}^+ = \mathbf{A}\mathbf{p}$ and $\mathbf{Q}^+ = \mathbf{A}\mathbf{Q}$. These are the $SM \times 1$ vectors of monthly per capita prices and quantities for each state. Let $\mathbf{Z}^+ = [\mathbf{I}_S \otimes \boldsymbol{\nu}_M \mid \boldsymbol{\nu}_S \otimes \mathbf{I}_M]$ be the $SM \times (S + M)$ matrix of indicator variables for the state and month-of-sample fixed effects. The OLS estimate of the slope coefficient using state and month per capital variables is

$$(19) \quad b^+ = (\mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{p})^{-1} \mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{Q}.$$

The Appendix derives the following expression for the difference between b^+ and b^* :

$$(20) \quad b^+ - b^* = [(\mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{p})^{-1} \mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} - (\mathbf{p}' \mathbf{M}_Z \mathbf{p})^{-1} \mathbf{p}' \mathbf{M}_Z] \mathbf{P} \mathbf{b} \\ + (\mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{p})^{-1} \mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{Z} \mathbf{d} + (\mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{p})^{-1} \mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+} \mathbf{A} \mathbf{e}.$$

This difference can be decomposed into three components (represented by the three terms of (20)). The first results from the fact that aggregation changes how the city-specific price response parameters in \mathbf{b} are weighted within b^+ relative to how they are weighted in b^* . Specifically, the b^+ estimate from the aggregate model weights elements of \mathbf{b} from each state according to the variance in the state's monthly average prices over time relative to other states, while the b^* estimate from the disaggregate model weights each city's coefficient based on the variance in the city's daily prices over time relative to other cities.²⁴ The second component of the difference arises when the residuals of the state-level, monthly weighted average

²⁴Note that the product $\mathbf{P}\mathbf{b}$ represents a vector of the elements $\beta_c p_{cd}$. Thus, in the second term of the β^+ estimate in (A5), premultiplying $\mathbf{P}\mathbf{b}$ by \mathbf{A} produces a population-weighted average of these region-specific predicted price effects across days of the month and cities within each state, while premultiplying these averages by $(\mathbf{p}' \mathbf{M}_Z \mathbf{p})^{-1} \mathbf{p}' \mathbf{A}' \mathbf{M}_{Z^+}$ produces a variance-weighted average across states and months of these population-weighted state-month average price effects.

prices regressed on state and month-of-sample fixed effects are correlated with the state-level, monthly weighted average of the city, and day-of-sample fixed effects. If the models using aggregate panel data retain a complete set of time-series and cross-section fixed effects, this term will be zero. However, this will not be true when aggregating to a national time series. The third term is nonzero if the residual from regressing state and month-of-sample weighted-average prices on state and month-of-sample fixed effects is correlated with residual from regressing the state and month-of-sample weighted-average values of the OLS residuals from estimating equation (18), \mathbf{e} , on state and month-of-sample fixed effects. Although a necessary condition for the consistency of the OLS estimates of δ and β is that p_{cd} and ϵ_{cd} are contemporaneously uncorrelated, if there is correlation between any prices in the same state and month, and any values of demand shock (ϵ) in the same state and month, this third term will be nonzero.

When comparing estimates obtained using time-series aggregated data, the specific terms of the decomposition change but maintain a similar structure. For example, aggregating to the daily national average level requires replacing the aggregation matrix \mathbf{A} with a new matrix \mathbf{B} defined in the Appendix. Using this notation, $\mathbf{B}\mathbf{p}$ is the daily national per capita weighted average price, and $\mathbf{B}\mathbf{Q}$ is the daily national per capita weighted average quantity. Although it is no longer possible to include day-of-sample fixed effects, more time-aggregated fixed-effects can be included. Let $\mathbf{Z}^\#$ equal the $D \times M$ matrix that is composed of M blocks each of length equal to the number of days in month m for $m = 1, 2, \dots, M$. Suppose that the daily weighted-average quantity is regressed on the daily weighted-average price and month-of-sample fixed effects. The slope coefficient from this regression is

$$(21) \quad b^\# = (\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{Q}.$$

The Appendix derives the following expression for the difference between $b^\#$ and b^* :

$$(22) \quad b^\# - b^* = [(\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B} - (\mathbf{p}'\mathbf{M}_{\mathbf{Z}}\mathbf{p})^{-1}\mathbf{p}'\mathbf{M}_{\mathbf{Z}}]\mathbf{P}\mathbf{b} \\ + (\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{Z}\mathbf{d} + (\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{\mathbf{Z}^\#}\mathbf{B}\mathbf{e}.$$

The three components of (22) differ in several important ways from those in (20). Because the matrix \mathbf{B} generates a population-weighted national average, the first term in (22) now represents the difference between a simple population-weighted average of the elements of \mathbf{b} and the residual-price-variation weighted average reflected in b^* . The second term is due to correlation between the residual of daily weighted average prices regressed on month-of-sample fixed effects and the residual from regressing the fitted values from the city and day-of-sample fixed effects on month-of-sample fixed effects. In other words, within-month fluctuations in nationwide demand that are correlated with national average prices but do not represent demand response will now result in endogeneity bias due to the inability to include

TABLE 5—DIFFERENCE BETWEEN AGGREGATE AND DISAGGREGATE ESTIMATES OF RESPONSIVENESS TO PRICE

(Pay-at-pump transactions only)					
Geography: Periodicity:	Aggregated panel models			Time-series models	
	City Monthly (1)	State Daily (2)	State Monthly (3)	National Daily (4)	National Monthly (5)
Disaggregated:					
b^* (from equation (17))	-0.047 (0.004)	-0.047 (0.004)	-0.047 (0.004)	-0.047 (0.004)	-0.047 (0.004)
Implied elasticity	-0.312	-0.312	-0.312	-0.312	-0.312
Aggregated:					
b^+ (from equation (19))	-0.037 (0.004)	-0.030 (0.009)	-0.025 (0.010)	-0.009 (0.007)	-0.001 (0.004)
Implied elasticity	-0.247	-0.201	-0.164	-0.062	-0.004
Decomposition of the difference ($b^+ - b^*$):					
Equation (20), term 1	0.0002	0.004	0.005	-0.0001	-0.0001
Implied impact on elasticity	0.001	0.028	0.035	-0.001	-0.001
Equation (20), term 2	0	0	0	0.042	0.046
Implied impact on elasticity	0	0	0	0.282	0.305
Equation (20), term 3	0.010	0.012	0.017	-0.005	0.0005
Implied impact on elasticity	0.064	0.083	0.112	0.031	0.003

Notes: Included fixed effects in each specification are the same as those indicated in Table 4. Standard errors for panel specifications are robust and clustered at the level of the cross-sectional unit to allow for arbitrary serial correlation. Standard errors for time-series specifications are estimated using a Newey and West (1987) procedure and are robust to first-order serial correlation. Implied elasticity terms represent the associated coefficient estimate multiplied by the sample mean price and divided by the sample mean of the quantity per capita measure.

daily fixed effects. The third term is nonzero if the residual from regressing daily weighted-average prices on month-of-sample fixed effects is correlated with the residual from regressing the month-of-sample weighted average values of the OLS residuals from estimating (16), \mathbf{e} , on state and month-of-sample fixed effects. If there is correlation between average prices on a day and the value of the demand shock on a different day in the same month, then this term will be nonzero.

While we allow the true price responsiveness of demand to vary across cities, our decomposition reflects a comparison to the restricted OLS model, which assumes a common slope coefficient, in large part because nearly all studies of gasoline demand estimate such models. In fact, for our daily city-level data, OLS estimation of the unrestricted model (15) does reveal some degree of heterogeneity, though the city-specific price coefficients are reasonably tightly centered around the mean of -0.049 with a standard deviation of 0.010 . This mean also turns out to be quite close to the OLS estimate from the restricted model (17), which produces $b^* = -0.047$. In other words, in the disaggregate data, the estimate of the weighted average of the β_c reflected in b^* is very close to the unweighted average of the estimates of the β_c . We will see that with aggregation this will no longer be the case.

In Table 5, we examine the extent to which the three different sources of bias discussed above arise in our data. Each column presents the estimation results for a particular degree of aggregation. The top panel presents the price coefficient from

the aggregate model and compares it to the price coefficient obtained using disaggregate data. Because a linear demand model is used in all specifications, we also report the elasticity implied by the coefficient when evaluated at the mean of the data. In the lower panel of Table 5, we decompose the difference in estimates from aggregate and disaggregate models by calculating the values of the three different terms in equation (20). In the aggregated panel regressions (first three columns), the first term in equation (20) contributes a moderate positive impact, making demand appear up to 11 percent less elastic in aggregate models. This suggests that cities in which demand is more responsive are being relatively underweighted in the aggregated models, either because weighting within state is now based on population or because weighting relative to cities in other states now depends on the degree variation in state weighted average prices. In addition, the third term in equation (20) consistently contributes an even larger positive impact on the aggregated panel estimates, making demand appear of up to 36 percent less elastic as a result of positive correlation between prices and unexplained demand shifts occurring on other days within the month and in other cities within the state.

In the aggregated time-series models (columns 4 and 5 of Table 5) the source of bias is entirely different, resulting instead from the inability to include the time-period fixed effects that are present in the panel regressions. The incomplete set of time dummy variables are only partially able to control for macroeconomic fluctuations that may influence both gasoline demand and gasoline prices, and as a result, these models exhibit very inelastic estimates (due to the large positive impact from the second term in equation (22)). Similar issues will arise if data is aggregated over time necessitating the use of cross-sectional or short-panel models that are no longer able to include a complete set of location fixed effects. In this case, the inability to control for unobserved locational characteristics that influence gasoline demand and are correlated with price local levels has the potential to substantially bias elasticity estimates (most likely toward higher magnitudes).²⁵ While not reported here, models estimated after aggregating our data to city- or state-level cross-sections result in estimates that are several times more elastic than our disaggregated results suggest.

We conclude from our decomposition that several important sources of bias may result from estimating gasoline demand using aggregate data, and that the source and degree of bias is likely to differ across the various levels of potential aggregation. Moreover, the generalized decomposition provided above can be used to evaluate the effects of almost any form of aggregation. As long as the aggregation process can be written as a matrix times the original vector of daily city-level data, a decomposition of the sources of bias can be computed as shown.

D. Implications for Studies Utilizing Aggregate Data

The findings from our bias decomposition analysis provide a framework for researchers designing or implementing new studies of gasoline demand response to

²⁵For example, people in highly populated areas might consume less gasoline per capita due to higher density and/or public transportation, but might also face higher gasoline prices that reflect higher costs of distribution and retailing.

evaluate the scope of potential biases that might arise from their data and empirical setting. In addition, by revealing how estimates from different empirical designs are likely to be biased and in what direction they may be biased, this will enable researchers and policymakers to utilize more appropriately and effectively demand elasticity estimates from existing studies. The wide ranging estimates of demand elasticity observed in the literature can be rationalized as reflecting, at least in part, the varying biases resulting from different levels and dimensions of data aggregation.

As an example, consider recent studies by Li, Linn, and Muehlegger (2014) and Rivers and Schaufele (2015) that examine how gasoline demand responds to changes in gasoline-related taxes. Li, Linn, and Muehlegger (2014) estimate a panel two-way fixed effects demand model using state-level annual averages of gasoline deliveries and prices (very similar to the type of log-log models we estimate in Table 4), obtaining estimates that are relatively consistent with our similarly aggregated panel results. Rivers and Schaufele (2015) take a similar approach using monthly province-level data from Canada. The focus of both studies, however, is on distinguishing price changes generated by gasoline tax rate adjustments from other non-tax related price changes and separately estimating the elasticity of demand response to these two types of price changes. They consistently find demand to be more elastic with respect to a change in gasoline tax. While both studies conclude that this result may reflect differences in persistence or salience, our aggregation model illustrates how different types of price changes (i.e., tax versus non-tax related) could result in different degrees of aggregation bias. For a given level of aggregation, the magnitude of the bias depends on the degree of heterogeneity and cross-correlation between consumption and prices in different cities or on different days over which aggregation is occurring. Unlike gas price fluctuations, gasoline taxes (and changes) are uniform across cities in the state and exhibit only occasional (and coordinated) discrete changes across the days of each year. As a result, tax-related price changes are less likely to be correlated with unobserved demand fluctuations in other cities. In this case, an estimate of the elasticity of demand to state tax changes could be more elastic simply because it exhibits less bias than an estimate based on nontax-related price changes. This mechanism may also explain why Davis and Kilian (2011) and Coglianesi et al. (2015) obtain more elastic estimates when using gasoline tax changes as an instrumental variable for gasoline price in order to identify demand response.

VI. Summary and Conclusions

In this study, we use high frequency panel data on gasoline prices and expenditures to reexamine the nature of gasoline demand in the United States. We specify a model of gasoline purchase behavior that allows us to identify a measure of the short-run elasticity of gasoline demand from data on gasoline expenditures. Our demand estimates are significantly more elastic than those of other recent studies. To investigate this discrepancy, we aggregate our data, estimate demand models similar to those used in previous studies, and then perform a decomposition that identifies the degree of aggregation bias resulting from each of several possible sources. The results suggest that the sources of bias can differ depending on the degree and

dimension of aggregation, but clearly show that the strongest bias occurs in time series models where a complete set of time-period fixed effects can not be included.

We also take advantage of the high frequency of our data to more carefully study how consumers respond immediately following a change in gasoline prices. Using our purchase frequency model, we are able to separately identify the elasticity of gasoline usage or demand from the change in consumers' probability of purchase. Our findings reveal a temporary response in the probability of purchase in the days following a price change as well as an immediate response in usage that does not dissipate over time. The estimated demand response over durations longer than a few days is consistent with the estimate from our static model.

Gasoline demand elasticity estimates are frequently used in policy evaluation and in broader economic research. Our results provide robust evidence that gasoline demand may be considerably more responsive to price fluctuations than one might conclude based on the recent literature, and the estimates differ by magnitudes large enough to substantially impact subsequent policy evaluation or market analysis.

Hughes, Knittel, and Sperling (2008), for example, estimate that gasoline demand has become extremely inelastic in recent years and conclude that emission reduction goals that rely on reductions in gasoline consumption will be difficult to achieve using gasoline taxes alone. They suggest that alternative measures such as CAFE standards may be needed to meet desired reductions while avoiding politically infeasible tax levels. Based on our more elastic estimates, however, the required gasoline tax increase needed to achieve a particular reduction in consumption would be at least five times smaller than what would be predicted using the estimates of Hughes, Knittel, and Sperling (2008).

Studies evaluating the impacts of cap-and-trade policies, like Borenstein et al. (2015), who analyze expected permit prices under California's greenhouse gas emissions cap-and-trade program, also rely directly on existing parameter estimates such as the price responsiveness of gasoline demand. Partly in response to recent estimates like those of Hughes, Knittel, and Sperling (2008), they adopt a rather inelastic value for gasoline demand elasticity which may contribute to their overall prediction that the supply of emissions abatement will be relatively inelastic to permit prices. Acknowledging a greater elasticity in the demand for gasoline could reduce predicted GHG permit price levels and volatility.

More accurate elasticity estimates can also substantially impact the inferences one draws when evaluating the macroeconomic costs of gasoline and oil market disruptions and the benefits of policies responses like maintaining a Strategic Petroleum Reserve (SPR) that are intended to reduce these costs. If gasoline demand were significantly more elastic than previously thought, prices will likely increase substantially less than would otherwise be predicted in response to a oil supply disruption, and the quantity consumers will purchase at these prices will be substantially smaller. As a result, the overall macroeconomic displacement effect is likely to be much smaller than would have previously been predicted. In addition, when consumers have more elastic demand, the release of a certain volume of fuel from the SPR during a market disruption will not be as effective as a policy lever aimed at reducing price levels. By demonstrating a significantly larger price-responsiveness of gasoline demand, our elasticity results strengthen any argument in favor of

eliminating or reducing the size of the SPR or in support of mechanisms imposing prices for greenhouse gas emissions.

In addition to providing a more robust and precise estimate of gasoline demand response, our aggregation analysis and decomposition of the sources of aggregation bias in this setting should help researchers and policy analysts to more successfully evaluate the reliability of existing estimates and to improve empirical design and identification in future studies.

APPENDIX

A. Aggregated Model Slope Estimates in Terms of Disaggregate Model Estimates

This Appendix derives the estimates of the slope coefficients and the relationships between them presented in Section VC. Applying OLS to equation (16) yields the following sample analogue:

$$(A1) \quad \mathbf{Q} = \mathbf{Zd} + \mathbf{Pb} + \mathbf{e},$$

where \mathbf{d} , \mathbf{b} , and \mathbf{e} are the OLS estimates of δ , β , and ϵ , respectively. Pre-multiplying both sides of (A1) by $\mathbf{M}_Z = \mathbf{I}_{DC} - \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'$ yields

$$(A2) \quad \mathbf{M}_Z\mathbf{Q} = \mathbf{M}_Z\mathbf{Pb} + \mathbf{e},$$

where $\mathbf{b} = (\mathbf{P}'\mathbf{M}_Z\mathbf{P})^{-1}\mathbf{P}'\mathbf{M}_Z\mathbf{Q}$. In terms of this notation, the OLS estimate of β^* is $b^* = (\iota_C'\mathbf{P}'\mathbf{M}_Z\mathbf{P}\iota_C)^{-1}\iota_C'\mathbf{P}'\mathbf{M}_Z\mathbf{Q}$. Substituting the expression for \mathbf{Q} from (A1) and recognizing that $\mathbf{M}_Z\mathbf{Z} = 0$ and $\mathbf{M}_Z\mathbf{e} = \mathbf{e}$ yields

$$(A3) \quad b^* = (\iota_C'\mathbf{P}'\mathbf{M}_Z\mathbf{P}\iota_C)^{-1}\iota_C'\mathbf{P}'\mathbf{M}_Z\mathbf{Pb} = (\mathbf{p}'\mathbf{M}_Z\mathbf{p})^{-1}\mathbf{p}'\mathbf{M}_Z\mathbf{Pb}.$$

We can also derive an expression for the OLS estimate of the slope coefficient from the model aggregated to the month and state level in terms of \mathbf{d} , \mathbf{b} , and \mathbf{e} . The first step in this process requires defining a $SM \times CD$ matrix \mathbf{A} that creates month and state-level data from day and city-level data. Assume that the rows of \mathbf{A} are ordered in blocks by states and within each block by months. Assume the columns of \mathbf{A} are ordered in blocks by cities and then by days within each block. The m th row of the s th block is composed of zeros for all elements except those where city c is contained in state s and day d is contained in month m . In those cases, the element of the matrix is equal to N_{cd} , the number of active purchasers in city c in day d divided by the sum of all of the N_{cd} values in the m th row of block s . The analog to equation (A1) when using monthly state-level aggregate data to estimate a single price coefficient is

$$(A4) \quad \mathbf{AQ} = \mathbf{Z}^+\mathbf{d}^+ + \mathbf{APb}^+ + \mathbf{e}^+,$$

where $\mathbf{Z}^+ = [\mathbf{I}_S \otimes \iota_M | \iota_S \otimes \mathbf{I}_M]$ the $SM \times (S + M)$ matrix of indicator variables for the state and month-of-sample fixed effects, and \mathbf{d}^+ , b^+ , and \mathbf{e}^+ are the resulting

OLS estimates. Following the logic above, b^+ can be expressed as in equation (19), where $M_{Z^+} = \mathbf{I}_{SM} - \mathbf{Z}^+(\mathbf{Z}^+\mathbf{Z}^+)^{-1}\mathbf{Z}^+$. Substituting equation (A1) into this expression for b^+ yields

$$(A5) \quad b^+ = (\mathbf{p}'\mathbf{A}'\mathbf{M}_{Z^+}\mathbf{A}\mathbf{p})^{-1}\mathbf{p}'\mathbf{A}'\mathbf{M}_{Z^+}\mathbf{A}\mathbf{Z}\mathbf{d} + (\mathbf{p}'\mathbf{A}'\mathbf{M}_{Z^+}\mathbf{A}\mathbf{p})^{-1}\mathbf{p}'\mathbf{A}'\mathbf{M}_{Z^+}\mathbf{A}\mathbf{P}\mathbf{b} \\ + (\mathbf{p}'\mathbf{A}\mathbf{M}_{Z^+}\mathbf{A}\mathbf{p})^{-1}\mathbf{p}'\mathbf{A}\mathbf{M}_{Z^+}\mathbf{A}\mathbf{e}.$$

This produces the expression for the difference between b^+ and b^* in equation (20) of Section VC.

Estimating a daily national weighted average quantity on the daily weighted-average price regression requires a $D \times CD$ matrix \mathbf{B} that transforms daily city-level data into daily national data. Each row of \mathbf{B} is composed of C blocks of length D . All of the elements of the d th row are zero except for the d th element of each of the C blocks, each of which is equal to N_{cd} divided by $N_d = \sum_{c=1}^C N_{cd}$. In terms of this notation, the slope coefficient from the regression of the daily national weighted-average quantity on the daily weighted-average price and month-of-sample fixed effects is

$$(A6) \quad b^\# = (\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{Z}\mathbf{d} + (\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{P}\mathbf{b} \\ + (\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{p})^{-1}\mathbf{p}'\mathbf{B}'\mathbf{M}_{Z^\#}\mathbf{B}\mathbf{e},$$

where $M_{Z^\#} = \mathbf{I}_D - \mathbf{Z}^\#(\mathbf{Z}^\#\mathbf{Z}^\#)^{-1}\mathbf{Z}^\#$. This yields the following expression for the difference between $b^\#$ and b^* given in equation (22) in Section VC.

REFERENCES

- Basso, Leonardo J., and Tae Hoon Oum.** 2007. "Automobile Fuel Demand: A Critical Assessment of Empirical Methodologies." *Transport Reviews* 27 (4): 449–84.
- Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins.** 2015. "Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design." National Bureau of Economic Research (NBER) Working Paper 20999.
- Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld.** 2008. "A meta-analysis of the price elasticity of gasoline demand: A SUR approach." *Energy Economics* 30 (5): 2105–22.
- Coglianesi, John, Lucas W. Davis, Lutz Kilian, and James H. Stock.** 2015. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." National Bureau of Economic Research (NBER) Working Paper 20980.
- Cui, Carolyn.** 2012. "Questions Arise on Gasoline Data." *Wall Street Journal*, March 26, C1.
- Dahl, Carol, and Thomas Sterner.** 1991. "Analysing gasoline demand elasticities: A survey." *Energy Economics* 13 (3): 203–10.
- Davis, Lucas W., and Lutz Kilian.** 2011. "Estimating the Effects of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 26 (7): 1187–1214.
- Espey, Molly.** 1998. "Gasoline demand revisited: An international meta-analysis of elasticities." *Energy Economics* 20 (3): 273–95.
- Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas.** 2007. "Revisiting the Income Effect: Gasoline Prices and Grocery Purchases." National Bureau of Economic Research (NBER) Working Paper 13614.
- Goodwin, P. B.** 1992. "A Review of New Demand Elasticities with Special Reference to Short and Long Run Effects of Price Changes." *Journal of Transport Economics and Policy* 26 (2): 155–69.
- Hastings, Justine S., and Jesse M. Shapiro.** 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics* 128 (4): 1449–98.

- Hsing, Yu.** 1990. "On the variable elasticity of the demand for gasoline: The case of the USA." *Energy Economics* 12 (2): 132–36.
- Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling.** 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29 (1): 113–34.
- Levin, Laurence, Matthew S. Lewis, and Frank A. Wolak.** "High Frequency Evidence on the Demand for Gasoline: Dataset." *American Economic Journal: Economic Policy*. <https://doi.org/10.1257/pol.20140093>.
- Li, Shanjun, Joshua Linn, and Erich Muehlegger.** 2014. "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy* 6 (4): 302–42.
- Lin, C.-Y. Cynthia, and Lea Prince.** 2013. "Gasoline price volatility and the elasticity of demand for gasoline." *Energy Economics* 38 (1): 111–17.
- Morath, Eric.** 2016. "Lower Gas Prices Drove Down Transit Use, So Why Can't You Find a Seat on the Train?" *Wall Street Journal*, April 6. <https://blogs.wsj.com/economics/2016/04/06/low-gas-prices-drove-down-transit-use-so-why-cant-you-find-a-seat-on-the-train/>.
- National Association of Convenience Stores (NACS).** 2007. "2007 Gas Price Kit." <http://www.nacsonline.com/YourBusiness/FuelsReports/Documents/GasPriceKit2007.pdf>.
- National Association of Convenience Stores (NACS).** 2014. "NACS Consumer Fuels Report." http://www.nacsonline.com/YourBusiness/FuelsReports/2014/Consumer-Research/Documents/ConsumerFuelsReport_Apr2014.pdf.
- Newey, Whitney K., and Kenneth D. West.** 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–08.
- Park, Sung Y., and Guochang Zhao.** 2010. "An estimation of U.S. gasoline demand: A smooth time-varying cointegration approach." *Energy Economics* 32 (1): 110–20.
- Rivers, Nicholas, and Brandon Schaufele.** 2015. "Salience of carbon taxes in the gasoline market." *Journal of Environmental Economics and Management* 74: 23–36.
- Small, Kenneth A., and Kurt Van Dender.** 2007. "Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect." *Energy Journal* 28 (1): 25–51.
- Sommer, Jeff.** 2015. "Cheap Gas Is a Thrill, but a Costly One." *New York Times*, December 13, BU8.
- Sterner, Thomas, and Carol A. Dahl.** 1992. "Modeling transport fuel demand." In *International Energy Economics*, edited by Thomas Sterner, 65–79. London: Chapman-Hall.