High Frequency Evidence on the Demand for Gasoline

Laurence Levin*
Matthew S. Lewis†
Frank A. Wolak‡

April 12, 2013

Abstract

This paper uses daily expenditure data and prices from February 1, 2006 to December 31, 2009 for 243 United States cities to estimate the price responsiveness of the daily demand for gasoline to changes in daily gasoline prices. Across a wide range of econometric model specifications, we obtain price elasticity estimates ranging from -0.29 to -0.61 or roughly an order of magnitude larger than estimates from recent studies using data that is more aggregated over both time and locations. We investigate the divergence between these results and find that a virtually all of the difference can be explained by the higher level of temporal and spatial aggregation. Since our baseline estimates rely on a traditional demand estimation approach that does not properly account for the distinction between expenditures on gasoline and demand for gasoline, we also specify a frequency of purchase model that recovers the demand for gasoline from expenditures and the number of purchases. In addition to confirming the robustness of our demand elasticity estimates we find evidence that consumers are substantially less (more) likely to purchase gasoline in the days immediately following a price increase (decrease). Finally, we allow for differences in price elasticities across metropolitan areas and show that they appear to be associated with various demographic characteristics as well as the availability of viable alternative modes of transportation in the region.

*VISA Decision Sciences, 901 Metro Center Boulevard, Mail Stop: M3-2B Foster City , CA 94404-2775, llevin@visa.com

†Department of Economics, The Ohio State University, 410 Arps Hall, 1945 N. High St., Columbus, OH 43210, lewis.1015@osu.edu

‡Department of Economics, Stanford University, Stanford, CA 94305-6072, wolak@stanford.edu.

1 Introduction

Over the last decade gasoline prices in the United States have become increasingly volatile, attracting the ire of consumers and the attention of policymakers armed with various proposals to cushion the impact of price shocks on consumers and the broader economy. This price volatility has largely resulted from an increased uncertainty in the supply of oil and available refining capacity—conditions that are likely to persist for the foreseeable future. Understanding consumers' ability to respond to such price fluctuations is crucial for predicting the potential impact of future supply disruptions and for estimating the value of proposed policy measures intended to limit the associated price volatility.[1] Estimates of gasoline demand responsiveness are also frequently used in macroeconomic analysis and in studies by public finance and environmental economists.

As a result, there has been a renewed interest in investigating the elasticity of gasoline demand. The conventional wisdom has been that gasoline demand is fairly inelastic and perhaps even more inelastic in the short run. Surveys of this literature by Dahl and Sterner (1991) and Espey (1998) both find that estimates of short-run price elasticity averaging around $-.26$. However, these studies largely analyze data from the 1970s and 80s. A more recent study by Hughes, Knittel and Sperling (2008) suggests that gasoline demand has become even less elastic in recent decades. They estimate short run gasoline demand elasticities of $-.24$ to $-.34$ for the period from 1975–1980, but find elasticities of $-.034$ to $-.077$ for the period 2001–2006. Taking a different approach Small and Van Dender (2007) estimate a structural model of demand for vehicle fuel efficiency and miles traveled from 1966–2001. They also conclude that demand for gasoline is highly inelastic in the short run: calculating an implied short run elasticity of $-.087$ over the entire sample period and $-.066$ for the period from 1997–2001.

Most of these studies identify demand elasticity estimates using highly aggregated data; often relying on monthly or annual data at a national level. For example, Hughes et al. (2008) use a time-series of monthly U.S. gasoline consumption and nationwide average retail prices. Small and Van Dender (2007) study annual state-level data. Other studies including Puller and Greening (1999) and Kayser (2000) investigate gasoline demand using household level surveys, but these data are typically observed on a quarterly or annual basis.[2] Monthly or annual regressions using

[1] Notable examples of policy proposals include the "gas tax holiday" suggested by several presidential candidates during the 2008 elections and the release of oil from the strategic petroleum reserve which was utilized by President Obama during the summer of 2011 amid oil supply disruptions in Libya and other areas of the Middle East.

[2] Puller and Greening (1999) uses quarterly data from the U.S. Consumer Expenditure Survey during the 1980s while

national or state-level average prices and gasoline volumes could potentially mask a significant share of the response by consumers in a given location to changes in the local gasoline price.

In addition to being highly aggregated, the quantity measures used in most of these studies are also observed fairly high up the supply chain at a point far removed from the final point of sale. This can be problematic if producers respond with a lag to changes in consumer purchases and/or demand fluctuations are buffered by additions or withdraws from gasoline inventories at various points in the distribution system. Moreover, these data often must be adjusted in an attempt to more accurately reflect domestic gasoline usage. For example, the most closely watched and heavily relied upon measure of U.S. gasoline consumption is the "product supplied" measure reported by the Department of Energy's Energy Information Administration (EIA). This measure is constructed from data on refinery output and includes a correction to account for the share of that production estimated to have been exported rather than destined for domestic consumption. Some controversy recently arose over its accuracy when it was revealed that from late 2010 through August of 2011 the EIA had underestimated gasoline export volumes by as much as 50% causing domestic product supplied measures to be inflated by over 3%.[3][4] Clearly the use of highly-aggregated production-based measures of gasoline volumes are not ideal for studying consumer demand.

Our study uses daily gasoline prices and citywide gasoline expenditures from 243 U.S. cities to provide an analysis of 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 demand. Second, the daily price variation that is masked by averaging in monthly, quarterly, or annual demand models can now be used to identify high frequency demand responses. 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.

In the first part of our analysis we adopt a more traditional demand model similar to that used in many of the previous studies. We obtain estimates of short-run demand elasticity ranging from $-.29$ to $-.61$, roughly an order of magnitude more elastic than other recent estimates. To investigate the source of this discrepancy we then aggregate our data to create a nationwide time

Kayser (2000) uses a cross-section from the 1981 wave of the Panel Study of Income Dynamics.

[3]See Cui (2012).

[4]A 3% error is quite large for this market. For example, in August 2011 the reported EIA numbers implied that U.S. gasoline demand had fallen by 3.5 percent from a year earlier when in fact it had actually fallen by an impressive 6.5%.

series of monthly consumption and monthly average price, similar to the data used by Hughes et al. (2008). Using our aggregated data we estimate a demand elasticities similar to those in Hughes et al. (2008), suggesting that studies using aggregated data underestimate consumers' responsiveness to gas price fluctuations.

We also consider the possibility that estimating elasticities using daily data may simply identify a shorter-run demand response that is, in fact, more elastic than the longer-run response captured using monthly aggregated data. To investigate this we incorporate lagged prices into the demand specification and show that the elasticities estimated in our earlier specifications accurately capture the persistent response of demand to changes in price. We do find evidence that gasoline expenditures respond even more strongly in the days immediately following a price change, but this temporary additional response largely dissipates after 4 to 5 days.

In the second part of our analysis we directly address the fundamental difference between gasoline demand (or usage) and gasoline expenditures—a difference that becomes even more pronounced when using daily data. Because consumers can use their car's gas tank for short-term storage, a consumer's demand for gasoline will be very different from their pattern of expenditures. Moreover, consumers may respond to price changes by altering both how much gasoline they use and and when they decide to purchase. Traditional models that simply relate prices to gasoline expenditures cannot separate these two potential effects.

In order to more accurately model the demand elasticity, 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. For example, consumers may be more likely to buy on certain days and the amount they purchase when they buy gasoline might fluctuate over time. Because we observe both the number of gasoline transactions occurring in a city on a given day as well as the total expenditures on gasoline we are able to separately identify changes in consumers' probability of purchase from changes in consumers' underlying demand. 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 of elasticity demand for the metropolitan area.

In general, our purchase model generates estimates of gasoline demand elasticity ranging from $-.28$ to $-.48$. These are similar in magnitude to the elasticity estimates obtained using the more traditional model and again are nearly an order of magnitude larger that estimates from other recent studies. In addition, by extending the model to include prices of previous days in addition to the current price we identify a significant response in the probability of purchase in the days following a price change. The probability of purchase on the day following a price increase tends to fall by .75% for every 1% change in price. However, this effect vanishes after a few days and the remaining impact of a price change on expenditures results entirely from changes in consumer's gasoline usage.

Finally, as an illustration of the questions that can be addressed using our model and our highly disaggregated daily city-level data, we then extend our model to allow the metropolitan area demand elasticities to depend of observable characteristics of the metropolitan area including various demographics and the frequencies with which alternative modes of transportation are used. We find evidence that population density, the size of low income population, and the shares of people commuting to work by subway, carpool, and walking or biking all are important predictors of a metropolitan area's demand elasticity.

2 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 prices collected from fleet credit card transactions and direct feeds from gas stations.[5]

Our expenditure data were obtained from Visa. 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 geographic definition of the associated Metropolitan Statistical Area (MSA). One caveat of the data is that Visa does not directly observe the price paid at the pump or the quantity purchased by the customer. This means that total expenditures at gas stations may include other items purchased along with gasoline. Fortunately, any potential bias arising from the inclusion of such purchases is likely to be small given that non-gasoline items

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

represent only a small portion of total expenditures at gas stations, and our analysis includes fixed effects that control for any variation across cities or over time in the amount of non-gasoline purchases being made. In addition, as a robustness check we utilize an alternative version of the data that includes only purchases paid for at the pump which will not contain non-gasoline items. We revisit this issue in more detail in Section 3.1.

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.

2.1 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 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.

Daily gasoline expenditures also follow different patterns across MSAs, presumably due to both independent retail price movements and other city-specific events. Note that daily expenditures necessarily change with retail prices because they represent the total quantity purchased multiplied by the price paid. We can create a measure of the total quantity of gasoline purchased each day using total daily expenditures divided by the daily average retail price for each MSA . 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.[6] The daily quantities for each MSA are normalized by the average quantity purchased in that MSA over the sample

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

Figure 1: Daily Average Retail Gasoline Prices for Selected Cities

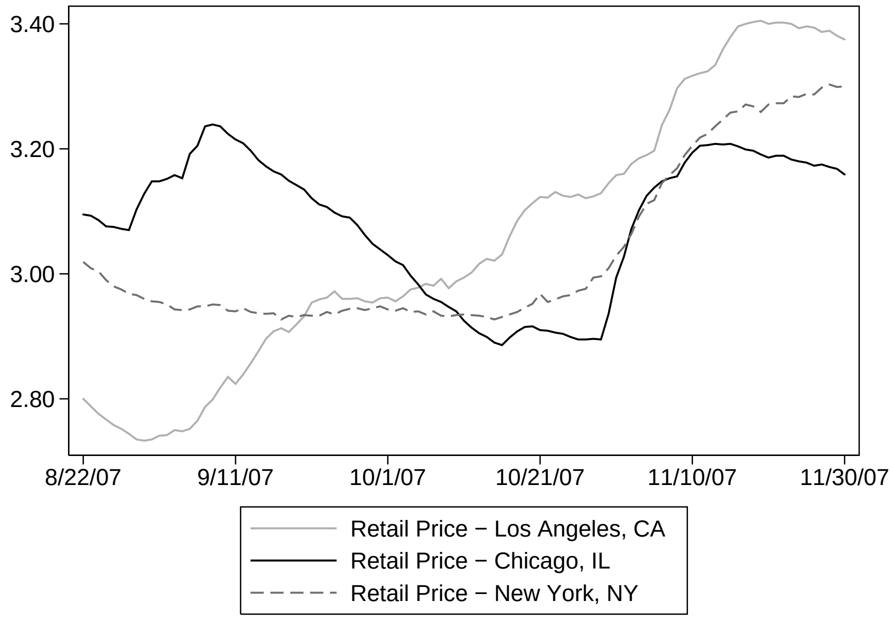

Figure 2: Seven Day Moving Average of Total Quantity Purchased for Selected Cities Normalized by the City Average

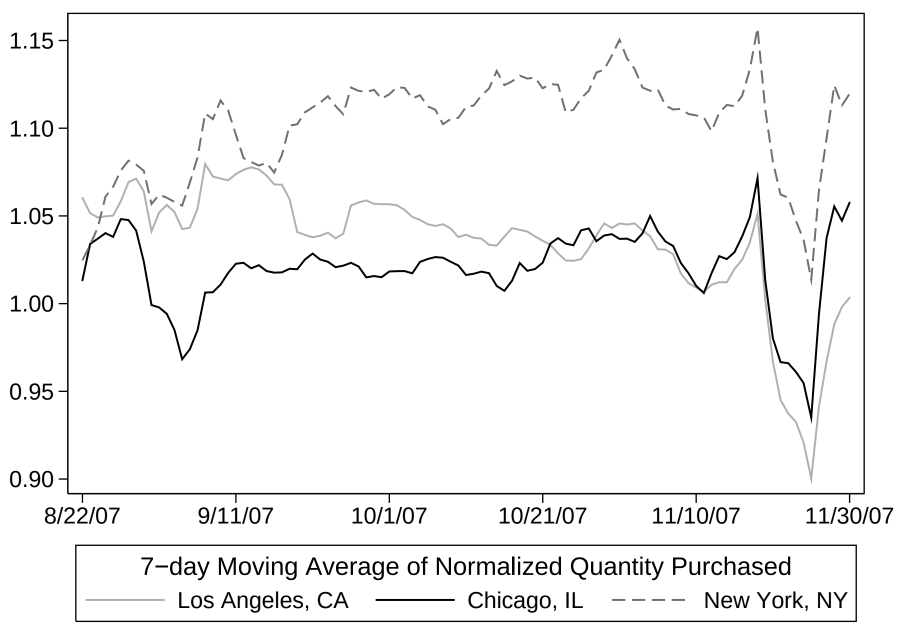

period. Just as with the prices, the quantities move together at times but also exhibit significant differences.

The daily expenditure data allows us to examine high frequency features of gasoline purchase patterns. Gasoline demand is know 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% more gasoline on Fridays than the daily average and buy 15% less on Sundays than the daily average. Figure 3 shows the averages by day of week of the total daily gasoline expenditures of Visa card customers across all 243 cities in our sample.[7] Within each city this pattern varies to some extent but Friday is always the highest demand day and Sunday is always the lowest demand 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. Figure 4 reports the average expenditure per transaction by day of week across all 243 cities in our sample.[8] The within-week pattern in expenditure per transaction is notably different from that of overall expenditures, and it exhibits much less day-to-day variation overall. This reveals that the within-week pattern observed in total expenditures results largely from fluctuations in the number of transactions occurring in each day.

3 Traditional Estimation of Gasoline Demand

In order to facilitate comparisons with earlier studies we begin our analysis of gasoline demand using the traditional and more descriptive approach. This involves regressing a measure of gasoline consumption on gasoline prices (often using a constant-elasticity or log-log form) while including other variables that help control for shifts in demand so that the resulting price coefficient can be interpreted as an elasticity of gasoline demand. While previous studies (particularly time-series studies) generally rely on observable proxies such as income to control for demand shifts, we take advantage of having highly disaggregated panel data to utilize and extensive set of city and time fixed effects.

Following Hughes et al. (2008) we estimate the price responsiveness of per-capita gasoline

[7] Day-of-week averages of the total quantity sold in the sample exhibit the exact same pattern.

[8] This is a quantity weighted average which is equivalent to calculating (for each day of the week) the total expenditures divided by the total number of transactions over the sample period in all cities.

Figure 3: Day-of-Week Averages of Daily Nationwide Gasoline Expenditures by Visa Customers in our 243 Cities

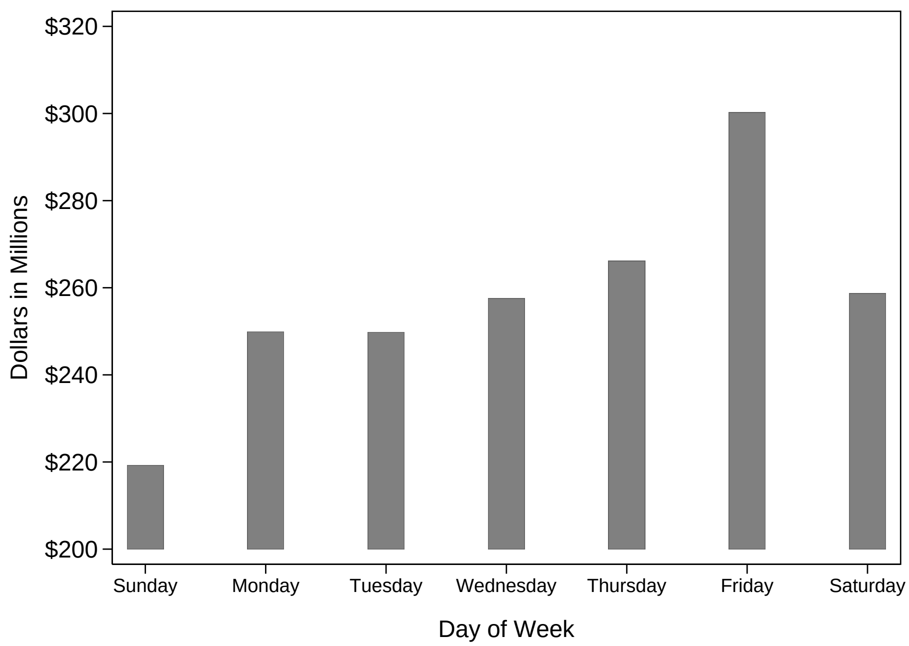

Figure 4: Day-of-Week Averages of Nationwide Gasoline Expenditures per Transaction by Visa Customers in our 243 Cities

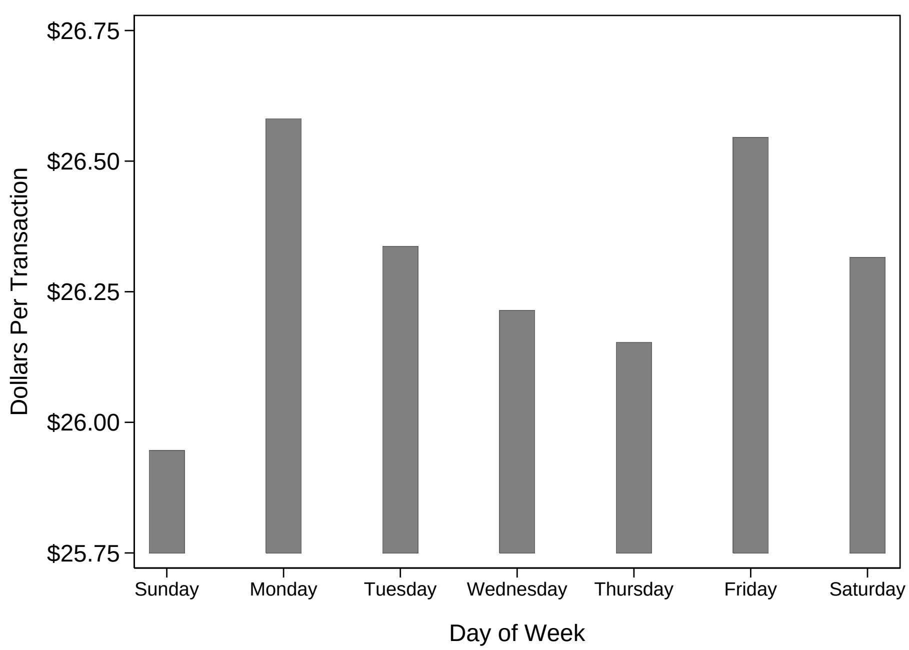

Table 1: Traditional Model of Gasoline Demand

Dependent Variable:	$\ln(q_{jd})$ (1)	q_{jd} (2)	$\ln(q_{jd})$ (3)	q_{jd} (4)
$\ln(\text{price}_{jd})$	-0.358		-0.614	
	(0.027)		(0.034)	
price_{jd}		-0.069		-0.114
		(0.005)		(0.007)
Fixed Effects:				
Day of Sample	X	X	X	X
City	X	X	X	X
Month of Sample $\times$ City			X	X
Implied Elasticity of Demand	-0.358	-0.373	-0.614	-0.623

Note: q_{jt} represents the per-capita amount of gasoline purchased in city j on day d. Standard errors are robust and clustered to allow arbitrary serial correlation within a city. The implied elasticity of demand for linear specifications is calculated at mean levels of price and per-capita consumption.

purchases so that the quantity variable is more comparable across cities. The quantity of gasoline purchased in a city on a given day is measured as the total observed expenditure from gas stations in the city on that day divided by the gasoline price in the city that day. Since we only observe gasoline consumption by Visa card users, we define the population of potential gasoline consumers as the number of consumers who used their Visa card for any type of transaction within that city in that particular month. As a result our baseline specification can be expressed as:

$$ln(q_{jd}) = \alpha_j + \lambda_d + \beta ln(p_{jd}) + \epsilon_{jd}, \tag{1}$$

where q_{jd} and p_{jd} represent the quantity of gasoline purchased per-capita and the average gasoline price in city j on day d, α_j and λ_d are fixed-effects for MSA j and day-of-sample d, and β represents the price elasticity of demand. For robustness we also estimate a linear model. The results of these specifications are reported in Column 1 & 2 of Table 1. Standard error estimates are clustered at the city level to allow for arbitrary serial correlation within each city. The degree of demand elasticity impled by the estimates for mean levels of price and quantity per capita are reported in the last row of the table.

The day-of-sample fixed effects present in the first two specifications effectively control for macroeconomic or gasoline market specific fluctuations that might impact demand at the national level. However, large city specific or regional demand shifts occurring over time could bias these elasticity estimates. For example, a city that is a popular warm-weather tourist destination may have a higher demand in the winter (opposite from the rest of the country) and may also tend to have slightly higher gasoline prices during the tourist season than the off-season. To allow for this possibility we also estimate an additional set of specifications that include separate sets of month-of-sample fixed effects for each city as well as a full set of national day-of-sample fixed effects. These are reported in Columns 3 & 4 of Table 1.

Implied demand elasticity estimates from the linear and log-log specifications are fairly similar. The specifications with city and day-of-sample fixed effects produce elasticities of around $-.36$. When month-of-sample fixed effects are also included for each city the elasticity estimates increase in magnitude to around $-.62$, suggesting that it may be important to control for city specific fluctuations in gasoline demand.

These elasticity estimates are roughly an order of magnitude more elastic than the estimates from the recent study by Hughes et al. (2008) which range from $-.034$ to $-.077$. The daily city-level purchase data appears to reveal a much greater degree of short run demand response than has been previously identified. In the remainder of this section and the following two sections we focus our analysis on establishing the robustness of this result and investigating the causes of the large discrepancy relative to other recent estimates.

3.1 Potential Biases from Non-gasoline Purchases

One potential limitation of our expenditure data is that we only observe total number and dollar amount of all transactions at gasoline stations in a given day, and some of these transactions will include non-gasoline purchases. To this point our analysis has effectively treated all purchases as gasoline only, allowing us to divide the total expenditure by the gasoline price to produce a measure of quantity. If non-gasoline purchases are present this can cause expenditures to appear more elastic to gasoline price changes. For example, if the prices and demand for non-gasoline items are not correlated with the price of gasoline, then dividing these expenditures by the price of gasoline will mechanically generate an elasticity of -1 for the non-gasoline portion of the transaction. In general, the share of revenues generated by non-gasoline items is small but non-trivial. According to the 2007 U.S. Economic Census, the average gasoline station in the U.S. receives just over 21% of its total revenues from non-fuel sales.[9] These non-fuel items may represent a sub-

[9] Most of these non-fuel 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.

stantially smaller share of credit card purchases at gas stations since consumers who purchase in the store and who make small purchases are more likely to pay with cash. Nevertheless, we want to examine the robustness of our empirical analysis to potential biases from non-fuel purchases.

Fortunately, these biases 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. Estimating our previous specification using only pay-at-pump transactions gives an alternative estimate of demand elasticity that is free from this bias and may give some indication the magnitude of the bias for in-store transactions.

Pay-at-pump purchases represent over 76% of total expenditures and over 64% of all transactions in our data.[10] Table 2 reports the coefficient estimates from the traditional demand specifications estimated using only gasoline purchased at the pump. These are directly comparable to the specifications in Table 1 for all gasoline purchases. The estimated price coefficients from the specifications with city and day-of-sample fixed effects in Columns 1 & 2 both imply elasticities of demand for purchasing at the pump of around $-.30$. The specifications that also include month-of-sample fixed effects for each city (Columns 3 & 4) result in elasticity estimates around $-.56$. These are somewhat less elastic than the estimates from the previous section that used both in-store and pay-at-pump purchases.

While estimates of demand elasticity derived from pay-at-pump purchases have the advantage of not being biased away from zero by the presence of non-fuel items, there are other reasons why restricting our analysis exclusively to pay-at-pump expenditures may not be optimal. First, individuals paying at the pump may differ from those paying in the store in terms of their sensitivity to gasoline price changes. To the extent possible with our data, we want to identify an overall measure of the elasticity of demand for gasoline in the population, and focusing on particular groups of consumers may be misleading. In addition, the degree to which the gasoline demand elasticity estimated from in-store purchases will be biased away from zero is not clear. As noted above, we do not know the fraction of expenditures in our data that come from non-fuel items. There is also reason to believe that non-fuel purchases may be negatively correlated with fuel prices. Gicheva, Hastings and Villas-Boas (2007) show that consumers' overall food purchasing behavior does respond to gasoline price fluctuations. When gasoline prices increase consumers

[10]On 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 non-fuel items which tend to be less expensive than the 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.

Table 2: Traditional Model of Gasoline Demand Using only Pay-at-Pump Transactions

Dependent Variable:	$\ln(q_{jd})$ (1)	q_{jd} (2)	$\ln(q_{jd})$ (3)	q_{jd} (4)
$\ln(\text{price}_{jd})$	-0.288		-0.561	
	(0.026)		(0.039)	
price_{jd}		-0.043		-0.076
		(0.003)		(0.005)
Fixed Effects:				
Day of Sample	X	X	X	X
City	X	X	X	X
Month of Sample $\times$ City			X	X
Implied Elasticity of Demand	-0.288	-0.316	-0.561	-0.559

Note: q_{jt} represents the per-capita amount of gasoline purchased at the pump in city j on day d. Standard errors are robust and clustered to allow arbitrary serial correlation within a city. The implied elasticity of demand for linear specifications is calculated at mean levels of price and per-capita consumption.

tend to substitute spending away from "food away from home" (which is relatively more expensive) and buy more at the grocery store. In addition, within the grocery store, consumers tend to substitute from more expensive to less expensive items. Within the context of convenience stores, the National Association of Convenience Stores' *2009 Consumer Fuels Report* indicates that, out of 1,100 consumers surveyed, 24% indicated that they would purchase fewer items in the store when gas prices rise. If non-fuel purchases fall when gas prices increase, this will mitigate the impact that non-fuel sales have on the estimated elasticity for gasoline.

In the interest of completeness most of the remaining analysis will present both estimates utilizing all purchases and estimates using only pay-at-pump purchases. Following the logic above, the estimates using all purchases are likely to represent an overestimate of the true overall demand elasticity while estimates using only pay-at-pump purchases may represent an underestimate.

4 Examining the Divergence from Previous Findings

Our empirical analysis consistently identifies elasticities that are nearly an order of magnitude more elastic than those obtained in recent studies. In this section we will discuss, in turn, a number of differences in our analyses that could potentially explain this disparity.

4.1 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 with a considerable amount of error. The most common source is that used by Hughes et al. (2008)—the U.S. 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 distributers and storage facilities). Each month the EIA reports the product supplied in each of the nation's 5 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.[11]

Another potential data source 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 the distributer 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 6 weeks or more between wholesale and retail levels."[12]

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

[11] See Cui (2012).

[12] FHWA Highway Statistics 2010, Table MF-33GA, Footnote 1.

masked by additions or withdraws from local storage. Moreover, the volumes are observed at a much more geographically and temporally disaggregated level, allowing changes in consumption to be linked much more directly to contemporaneous local price fluctuations.

4.2 Estimating Demand Elasticity Using Aggregated Data

In general, using highly aggregated data has the potential to mask much of the temporal and geographic co-variation in prices and quantities that result from consumer demand response. The following simple theoretical model illustrates this point. Let $D_{jd}(p_{jd}, X_{jd})$ equal the daily demand for gasoline in MSA j during day d, where p_{jd} is price of gasoline in region j on day d and X_{jd} is the vector of characteristics of region j 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 aggregate demand for gasoline during month m, is equal to

$$Q_m = \sum_{d \,\in\, S(m)} \sum_{j=1}^{J} D_{jd}(p_{jd}, X_{jd}), \tag{2}$$

where $S(m)$ is the set of days in month m and J is the total number of MSAs in the sample. 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, the monthly state-level demand depends on the daily prices for all days during that month for all MSAs in that state. Put another way, specifying an empirical model in which total monthly state-level consumption is a function of only the average price for the month in that state is equivalent to assuming that consumers' daily demand for gasoline in each city responds only to the average gasoline price for that month in the state rather than the actual price of gasoline in that city on that particular day. Models using monthly state-level average prices are also unable to utilize any within-month or within-state variation in prices around the average level to help identify the demand elasticity. Consequently, it is not surprising that regressing the monthly aggregate quantity of gasoline consumed at the state or national level on a single monthly average price at the state or national level would yield a substantially smaller (in absolute) estimated monthly demand response to changes in the monthly average price.

We are able to investigate the degree to which such data aggregation might impact elasticity estimates by using our data to create new data sets with varying levels of temporal and geographic aggregation. We construct daily data sets of state level and nationwide total quantity purchased and average price, as well three monthly data sets at the city, state, and national levels of aggregation. Aggregate per-capita quantities are simply the corresponding sum of the daily quantity purchased divided by the total number of Visa customers in the combined area. Since prices may be averaged across cities (or days) with very different quantities purchased, we compute a quantity weighted average price. To facilitate a more direct comparison with Hughes et al. (2008), average prices in the national monthly specification are converted to constant 2005 dollars using the GDP implicit price deflator. Using real prices would not impact the other specifications because the day- or month-of-sample fixed effects would absorb any price adjustment.

As in our main analysis, a complete set of time period and cross-sectional fixed effects are used whenever possible to control for unobserved 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.[13] This final specification is closest to that of Hughes et al. (2008).

The demand estimates for each level of aggregation appear in Table 3. The top rows report the results when estimated using all purchases while the bottom rows report the results for pay-at-pump purchases only. 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 $-.27$ and $-.29$ for all purchases and between $-.21$ and $-.18$ for pay-at-pump purchases) and are noticeably less elastic than corresponding estimates from the disaggregated regression in Column 1. The elasticity estimates from the two time series regressions in Columns 5 & 6 are even smaller in magnitude at around $-.14$ for all purchases and near zero for pay-at-pump purchases. Clearly, increasing levels of aggregation lead to less elastic estimates of demand, particularly when moving from panel to time series data.

Once aggregated to a national time series our elasticity estimates are fairly close to those of Hughes et al. (2008). However, their study examined a slightly earlier time period than ours. To generate a more accurate comparison we continue our analysis by replicating the Hughes et al. (2008) specification using their data sources but for our later time period. As in their study, gasoline consumption is measured using the EIA's monthly nationwide estimate of motor gasoline

[13]Data on per capita personal disposable income comes from the Bureau of Economic Analysis and are converted to constant 2005 dollars using the GDP implicit price deflator.

Table 3: Regressions Using Aggregated Data

Dependent Variable = ln(quantity per capita)

Geography:	city	city	state	state	national	national
Periodicity:	daily	monthly	daily	monthly	daily	monthly
	(1)	(2)	(3)	(4)	(5)	(6)
All Purchases:						
$\ln(\text{price}_{it})$	-0.358	-0.290	-0.297	-0.270	-0.143	-0.127
	(0.027)	(0.015)	(0.007)	(0.025)	(0.065)	(0.024)
$\ln(\text{income}_{it})$						-0.244
						(0.440)
Pay-at-Pump Purchases Only:						
$\ln(\text{price}_{it})$	-0.288	-0.214	-0.206	-0.176	-0.020	0.002
	(0.026)	(0.017)	(0.008)	(0.028)	(0.069)	(0.024)
$\ln(\text{income}_{it})$						-0.734
						(0.333)
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		

Note: 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-West procedure and are robust to first-order serial correlation.

"product supplied". Price is measured using the U.S. city average price for unleaded regular gasoline as reported in the U.S. Bureau of Labor Statistics' CPI-Average Price Data and is converted to constant 2005 dollars using the GDP implicit price deflator.

To check our ability to replicate the Hughes et al. (2008) analysis we first estimate their baseline double-log specification (equivalent to Column 5 of Table 3 above) for the period 2001–2006. Results are reported in table 4, column 1. The estimate of price elasticity ($-.042$) is identical to that of Hughes et al. (2008).[14] Estimating the same specification using data from 2006–2009 yields a price elasticity that is slightly positive and not significantly different from zero. This

[14]Our estimate of the income elasticity is .32 as opposed to their estimate of .53. This discrepancy appears to have been caused by the fact that Hughes et al. (2008) use previously published estimates of disposable personal income that have since been revised by the BEA. We utilize the updated estimates so that our income measures are consistent with those available for the 2006–2009 sample period.

Table 4: Elasticity Estimates from Replication of Hughes et al. (2008)

Dependent Variable = ln(quantity per capita)				
Date Range:	2000–2006	2006–2009	2006–2009	2006–2009
Data Source:	EIA/BLS	EIA/BLS	Visa/AAA (all purchases)	Visa/AAA (pay-at-pump)
	(1)	(2)	(3)	(4)
ln(price$_t$)	-0.042	0.026	-0.127	0.002
	(0.010)	(0.024)	(0.024)	(0.024)
ln(income$_t$)	0.321	-1.272	-0.244	-0.734
	(0.066)	(0.405)	(0.440)	(0.333)
Month-of-Year Fixed Effects	X	X	X	X

Note: Columns (1) and (2) use EIA data on "product supplied" to measure quantity and the CPI average price data for unleaded regular to measure price. Column (3) uses our Visa expenditure data to measure quantity and a weighted average of our AAA average price data to measure price. Standard errors are estimated using a Newey-West procedure and are robust to first-order serial correlation.

is very similar to our estimate from the same specification using aggregated Visa pay-at-pump purchase data (reported in Column 4) but is much less elastic than our estimate using all Visa purchases (Column 3). The use of an alternative data source may be partially responsible for differences between our elasticity estimates and those of previous studies, but overall the results above suggest that most of the difference likely a result of the level of data aggregation.

4.3 Controlling for Demand Shocks

In addition to avoiding measurement error and aggregation bias the other important advantage to using disaggregated panel data when estimating demand response is the ability to utilize fixed effects to better control for demand differences across days and cities. Studies using time-series data must either identify other observable variables (such as per-capita income) that account for changes in demand or rely on instrumental variables to isolate gasoline price changes associated with supply shocks. With panel data we observe many different locations experiencing the same macroeconomic demand shocks so it is possible to identify demand response by observing how idiosyncratic price deviations between cities result in corresponding quantity changes.

The results in Table 3 reveal that aggregating our panel data to a daily or monthly time-series results in demand estimates that are substantially less elastic. In part this may be caused by the fact that the full set of day-of-sample fixed effects can no longer be included. Elasticity

Table 5: Regressions Using Different Time Fixed Effects

Dependent Variable = ln(quantity per capita)

Geography:	city	city	city
Periodicity:	daily	daily	daily
	(1)	(2)	(3)
All Purchases:			
ln($price_{it}$)	-0.358	-0.252	-0.158
	(0.027)	(0.014)	(0.005)
ln($income_{it}$)			-0.323
			(0.136)
Pay-at-Pump Purchases Only:			
ln($price_{it}$)	-0.288	-0.159	-0.039
	(0.026)	(0.014)	(0.004)
ln($income_{it}$)			-0.444
			(0.165)
Fixed Effects:			
Day of Sample	X		
Month of Sample		X	
Month of Year			X
Day of Week		X	X
City	X	X	X

Note: Standard errors are robust and clustered by city to allow for arbitrary serial correlation.

estimates are likely to be biased downward if there are demand shocks that are not controlled for by the day-of-week and month-of-sample fixed effects in the daily specification or by the month-of-year and per-capita income variables in the monthly specification. In fact, the negative income elasticities that are estimated in our 2006–2009 time series specifications are most likely biased because observed income and month-of-year fixed effects are unable to sufficiently control for the large shifts in gasoline demand that occurred before and after the recession of 2008.

In order to examine the extent to which the lower elasticity estimates in our aggregated specifications are a result of the lack of control variables or fixed effects we estimate similar specifications (i.e., with incomplete time fixed effects) using our fully disaggregated data. The specifications reported in Table 5 are all estimated using daily city-level data, and all include city fixed effects to account for differences across cities in average consumption and day-of-week fixed effects (where identified) to control for weekly purchase patterns. The first column once again reports the results from our baseline model with a full set of day-of-sample fixed effects. The second specification includes only month-of-sample fixed effects. Here elasticity estimates drop significantly to $-.25$ for all purchases and $-.16$ for pay-at-pump. The final specification includes only month-of-year fixed effects and a measure of the national per-capita disposable income, just as in the monthly time series specification in Table 3, Column 6. The elasticity estimates (of $-.16$ and $-.04$) are very similar to those of the corresponding aggregated specifications. Together the results reveal that the ability to effectively control for changes in demand across days and locations is a major reason why our estimates from our aggregated analysis are substantially different from our disaggregated analysis.

We conclude that the typical consumer has a much more elastic short-run demand for gasoline than national time series estimates would imply. One reason for this is that we are able to utilize a measure of quantity consumed that is recorded at the final consumer purchase stage rather than further upstream. The second, and perhaps more important, reason is that the use of more highly disaggregated data on both price and quantity avoids problems of aggregation bias and makes it possible to better control for demand shifts.

5 Short Run vs Longer Run Demand Elasticity

The findings of the previous section suggest that earlier studies may have produced relatively inelastic estimates of gasoline demand due to their reliance on more aggregated data. We now want to consider the alternative possibility that both estimates are, in fact, correct but that they characterize distinctly different relationships—that our demand estimate (using daily data) is more elastic because it captures consumers' initial response following a price change while estimates using more aggregated monthly data capture a longer run response. It is not unusual for the demand curves to be more elastic in the short run than in the long run. Perhaps the most common of these situations occurs when consumers can hold inventories and in the short run choose to add to or withdraw from inventories in response to price changes even when they do not significantly change their consumption in the long run. 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 in response to a temporary increase (or 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

Table 6: Traditional Demand Model with Lagged Prices

Dependent Variable = ln(quantity per capita)

	All Purchases (1)	Pay at Pump (2)
$\ln(\text{price}_{jd})$	-0.824	-0.853
	(0.073)	(0.079)
$\ln(\text{price}_{j,d-1})$	-0.513	-0.610
	(0.084)	(0.100)
$\ln(\text{price}_{j,d-2})$	0.545	0.661
	(0.078)	(0.092)
$\ln(\text{price}_{j,d-3})$	0.323	0.403
	(0.041)	(0.046)
$\ln(\text{price}_{j,d-4})$	0.146	0.183
	(0.042)	(0.047)
$\ln(\text{price}_{j,d-5})$	0.080	0.067
	(0.036)	(0.041)
$\ln(\text{price}_{j,d-10})$	-0.068	-0.085
	(0.018)	(0.022)
$\ln(\text{price}_{j,d-20})$	-0.025	-0.034
	(0.018)	(0.018)
Fixed Effects:		
Day of Sample	X	X
City	X	X
Total Implied Elasticity 20 Days After a Price Change	-0.338	-0.267

Note: Standard errors are robust and clustered to allow arbitrary serial correlation within a city. The implied elasticity of demand for linear specifications is calculated at mean levels of price and per-capita consumption.

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.

It is important to note that although we are using daily data the traditional demand specification doesn't allow for this behavior. It assumes that the amount purchased today depends on the current price and is not allowed to differ depending on prices in the recent past. This mitigates the extent to which such behavior could be responsible for the more elastic estimates found in the daily city-level analysis above. We can, however, alter our specification by incorporating past prices along with current prices in order to examine whether such behavior has an important impact on gasoline purchases.

Table 6 reports the results of a specification that includes the log of the current price and the lagged prices from each of the previous 5 days as well as longer lags of 10 and 20 days. The coefficients in Column 1 are estimated using all purchases while those in Column 2 are estimated using pay-at-pump purchases only. When additional price lags of 40 and 60 days are included they are jointly insignificant and do not substantially affect the estimates of the existing coefficients.

For both the all-purchase and pay-at-pump data, 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. For example, the results imply that the amount of gasoline purchased at the pump one day after a 1% price increase will will be 1.45% lower than it would have been without the price increase. During the following 3 to 4 days, however, the amount purchased tends to increase sharply, back towards 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 6 and implies that the elasticity of demand response after 20 days is $-.34$ for all purchases and $-.27$ for pay-at-pump purchases. These responses are roughly equivalent to the elasticities of $-.36$ and $-.29$ identified in our baseline model with no lags. We conclude from this that the lower elasticity estimates in previous studies and in our monthly aggregated regressions do not appear to have resulted from consumers being less responsive to price changes that persist over longer time periods. Instead it appears that the types of temporal aggregation bias and measurement error discussed in the previous section are likely to be responsible for reducing estimates of demand response below those reflected in the daily data. This explanation is also consistent with the fact that additional aggregation of the analysis geographically produces a further reduction in estimated demand response despite the fact that there is little reason to suspect that consumers shift demand from state to state in response to relative price differences.

Nonetheless, the very-short-run responses we are able to identify using daily lagged prices are interesting in their own right. Consumers appear to substantially alter their purchases in the days following a price change. They purchase more gasoline sooner when prices fall and they reduce their purchases for several days after prices rise, perhaps waiting to see if prices will fall again before they have to buy. Of course, the more descriptive demand model we have been using to this point makes it impossible to determine whether consumers actually alter their driving intensity following a price change or whether they simply delay or expedite purchases in the days following the change. This is one of the main goals of the consumer purchase model described in the next section—attempting to separate consumers' demand (or usage) decision from their purchase decision to gain a better understanding of how consumers respond to price fluctuations.

6 Model of Consumer Demand and Purchase Behavior

To this point we have shown that the gasoline purchases are much more responsive to price fluctuations than previous studies have suggested. However, since 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. Because consumers can buy and store gasoline in their car, a consumer's daily demand for gasoline can differ from the consumer's 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:

$$d_{jd} = exp(\alpha_j + \lambda_d + \beta ln(p_{jd}) + \epsilon_{jd}), \tag{3}$$

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:

$$\rho_{jd} = \gamma_j + \delta_d. \tag{4}$$

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:

$$e_{jd} = \frac{p_{jd}d_{jd}}{\rho_{jd}}. \tag{5}$$

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. Since 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:

$$E_{jd} = e_{jd}n_{jd} = \frac{p_{jd}d(p_{jd}, \epsilon_{jd})n_{jd}}{\rho_{jd}}. \tag{6}$$

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:

$$\frac{n_{jd}}{N_{jd}} = \gamma_j + \delta_d + \nu_{jd}, \tag{7}$$

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 6 and taking logs generates our econometric model of gasoline expenditures:

$$ln(E_{jd}) = \alpha_j + \lambda_d + (\beta + 1)ln(p_{jd}) + ln(n_{jd}) - ln(\hat{\rho}_{jd}) + \epsilon_{jd}. \tag{8}$$

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

$$ln(Q_{jd}) = \alpha_j + \lambda_d + \beta ln(p_{jd}) + ln(n_{jd}) - ln(\hat{\rho}_{jd}) + \epsilon_{jd}, \tag{9}$$

The empirical model in Equation 9 makes it possible to identify the underlying MSA-level elasticity of demand for gasoline (β) using only data on prices, the quantity purchased, and the number of transactions. In Equations 3 & 4 the demand and probability of purchase are both assumed to vary by city and day of sample, but, as in our earlier analysis, different combinations of fixed effects can easily be used to generate alternative specifications for each of these conditional mean functions. We will also consider a specification that includes both lagged and current prices in the demand and purchase probability equations.

Table 7: Estimates of Demand Using Purchase Model

Dependent Variable = ln(quantity$_{jd}$)

	All Purchases			Pay-at-Pump Only		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(price$_{jd}$)	-0.448	-0.396	-0.480	-0.341	-0.288	-0.351
	(0.019)	(0.009)	(0.005)	(0.020)	(0.006)	(0.002)
ln(# of transactions$_{jd}$)	1	0.999	1.025	1	0.997	0.993
		(0.005)	(0.003)		(0.002)	(0.001)
ln(predicted probability of purchase$_{jd}$)	-1	0.007	0.008	-1	-0.003	-0.007
		(0.002)	(0.002)		(0.001)	(0.001)
Fixed Effects:						
Day of Sample	X	X	X	X	X	X
City	X	X	X	X	X	X
Month of Sample × 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.

7 Estimation of the Frequency of Purchase Model

Coefficient estimates from the model in Equation 9 are reported in Columns 1 & 2 of Table 7. Standard error estimates are generated using a nonparametric bootstrap to account for the fact that the predicted probability of purchase is estimated in a first-stage regression. 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 and without for robustness. The model with all restrictions yields an elasticity estimate of $-.45$. The unrestricted model in column 2 produces a similar elasticity estimate of $-.40$. The unrestricted coefficient on $ln(n_{jd})$ is very close to 1, but the coefficient on $ln(\hat{\rho}_{jd})$ is around .03; 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. As in our traditional demand analysis we also estimate a model including city-specific month-of-sample fixed effects in addition to the day-of-sample fixed effects. This is reported in Column 3. The elasticity estimate is larger in magnitude, at $-.48$, but not as large as was seen in the traditional demand analysis.

Column 4 through Column 6 of Table 7 report the results of the same three specifications when estimated using only pay-at-pump transactions. These elasticity estimates exhibit a very similar pattern to those in Column 1 through Column 3, though each is about .1 less elastic than the estimates using all transactions. In general, the demand elasticity estimates generated using the purchase model (ranging from $-.29$ to .48) are similar in magnitude to the corresponding elasticity estimates found using the traditional demand model described above. We conclude from these findings that the relatively strong short-run response of expenditures to gas price changes that was revealed using the traditional model is largely a result of changes in consumers' gasoline usage rather than a response in the timing of purchase.

7.1 Purchase Model with Lagged Prices

Aside from identifying elasticities that were substantially larger than previous estimates, our traditional demand analysis also identified a sizable additional response in expenditures in the days immediately following a price change once lagged prices were included in the analysis. This response could result from consumers temporarily postponing (or expediting) some of their driving or from consumers delaying purchases (or purchasing sooner) to try to take advantage of temporary price swings. Using the structure of our consumer purchase model we have the ability to decompose this very-short-run response 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.

We can allow for these very-short-run responses by incorporating past prices along with current prices into our model of individual demand and purchase. If consumers are substituting driving intertemporally in response to price changes 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 3 & 4 to allow for these types of behavior. The demand for each customer in a city j on a day d can be specified as:

$$d_{jd} = exp(\alpha_j + \lambda_d + \beta ln(p_{jd}) + \sum_{l \in L} \zeta ln(p_{j,d-l}) + \epsilon_{jd}), \tag{10}$$

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:

$$\rho_{jd} = \gamma_j + \delta_d + \psi ln(p_{jd}) + \sum_{l \in L} \eta ln(p_{j,d-l}). \tag{11}$$

Leaving the consumer purchase model from Section 3 otherwise unchanged results in the following final representation of the aggregate quantity purchased in city j on day d:

$$ln(Q_{jd}) = \alpha_j + \lambda_d + \beta ln(p_{jd}) + \sum_{l \in L} \zeta ln(p_{j,d-l}) + ln(n_{jd}) - ln(\hat{\rho}_{jd}) + \epsilon_{jd}, \tag{12}$$

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

$$\frac{n_{jd}}{N_{jd}} = \gamma_j + \delta_d + \psi ln(p_{jd}) + \sum_{l \in L} \eta ln(p_{j,d-l}) + \nu_{jd}. \tag{13}$$

As in our traditional demand analysis with lagged prices we include in the demand equation lags of the log of price for each of the previous 5 days, as well as from 10 and 20 days previous. These lags are also included in the purchase probability equation. Columns 1 & 2 of Table 8 report the results when using all purchases and Columns 3 & 4 report results for pay-at-pump purchases only. The final row of the table includes the total implied elasticity of the probability of purchase or of demand response after 20 days.[15] The demand estimates are directly comparable to the specifications without lags in Columns 2 & 5 of Table 7.

In general, the presence of a lagged price in the demand specification 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. The coefficients on lagged prices are smaller and generally positive, indicating that the degree of response is reduced after the first few days. In all specifications, sum of the coefficients on the current and lagged values of $ln(p_{jd})$ in the demand equations are very similar to the coefficient estimates for $ln(p_{jd})$ when no lagged prices are included (Table 7, Columns 2 & 5). In other words, the total demand response to a price shock that lasts longer than a few days exhibits a demand elasticity of around $-.42$ for all purchases or $-.28$ for pay-at-pump purchases—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.

The response in purchase probability to a price change is somewhat different. 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 $-.76$ (or -1.12 for pay-at-pump purchases) all else equal. However, this response in the probability of

[15]For 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.

Table 8: Purchase Model with Lagged Prices

	All Purchases		Pay at Pump	
	Purchase Equation	Demand Equation	Purchase Equation	Demand Equation
	(1)	(2)	(3)	(4)
$\ln(\text{price}_{jd})$	-0.007	-0.579	-0.009	-0.449
	(0.004)	(0.022)	(0.003)	(0.015)
$\ln(\text{price}_{j,d-1})$	-0.032	0.001	-0.028	0.092
	(0.004)	(0.020)	(0.004)	(0.012)
$\ln(\text{price}_{j,d-2})$	0.024	0.090	0.023	0.0002
	(0.004)	(0.018)	(0.003)	(0.001)
$\ln(\text{price}_{j,d-3})$	0.016	0.058	0.014	0.054
	(0.002)	(0.013)	(0.001)	(0.008)
$\ln(\text{price}_{j,d-4})$	0.001	0.039	0.003	0.021
	(0.002)	(0.013)	(0.002)	(0.008)
$\ln(\text{price}_{j,d-5})$	0.007	0.001	0.003	-0.012
	(0.002)	(0.011)	(0.002)	(0.008)
$\ln(\text{price}_{j,d-10})$	-0.004	-0.005	-0.003	0.0003
	(0.001)	(0.006)	(0.001)	(0.003)
$\ln(\text{price}_{j,d-20})$	-0.0001	-0.025	-0.001	0.013
	(0.001)	(0.007)	(0.001)	(0.005)
$\ln(\text{\# of transactions}_{jd})$		0.998		0.996
		(0.007)		(0.004)
ln(predicted probability of purchase$_{jd}$)		0.025		-0.008
		(0.007)		(0.003)
Fixed Effects:				
Day of Sample	X	X	X	X
City	X	X	X	X
Total Implied Elasticity 20 Days After a Price Change	0.096	-0.420	0.061	-0.281

Note: The dependent variable in Column 1 and Column 2 are (respectively) the share of Visa customers purchasing and the log of the average quantity purchased per capita by Visa customers in city j on day d, and in Column 3 & 4 are the share of Visa customers purchasing at the pump and the log of the average quantity purchased at the pump per capita by Visa customers in city j on day d. 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. The implied elasticity of demand for linear specifications is calculated at mean levels of price and per-capita consumption.

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 to be small and slightly positive at .09 for all purchases and .06 for pay-at-pump. As a result, it appears that both gasoline usage and the probability of purchasing on a given day fall (rise) immediately after prices increase (decrease), but the effects of consumers shifting when they purchase are very short lived, while most of the impact of a price change on usage remains permanently.

These findings help to explain the estimated response of overall expenditure levels that were identified when we included lagged prices in our traditional demand analysis. The results from our consumer purchase model imply that much of the temporary portion (i.e., lasting several days) of the very large response in expenditures to a price change is due to consumers delaying or expediting purchases while the fraction of the response in expenditures that persists is largely due to changes in underlying gasoline usage.

8 Geographic Variation in Demand Elasticity

In order to both test the plausibility of our estimates and illustrate the types of questions that can be addressed using our high-frequency city-level panel data, we present one final extension the purchase model that allows the elasticity of demand to vary across cities. One simple way to illustrate the degree of variation in demand responsiveness is to include interactions between the city fixed effects and the $ln(p_{jd})$ term in Equation 9. The results of this model reveal significant variation across cities, though nearly all remain within a fairly tight range between $-.35$ and $-.45$. Individual city elasticity estimates are fairly precise, with most having standard errors of .005 or less. Figure 5 shows a histogram of elasticities for the cities in the sample.

Variation in demand elasticity across cities most likely results from differences in the way consumers use gasoline or differences in resources consumers have to purchase gasoline. When studying city level measures of demand elasticity it is difficult to identify exactly how differences in consumer behavior or budget translate into price sensitivity. Nevertheless, we can examine how city level demographic measures relate to our overall elasticity estimates.

MSA-level demographic characteristics from the 2000 U.S. Decennial Census are collected for each of the cities in our sample. These include measures of population density in the MSA and the share of the population in the MSA with income greater than twice the poverty level, as well as

Figure 5: Histogram of Demand Elasticity Estimates Across Sample Cities

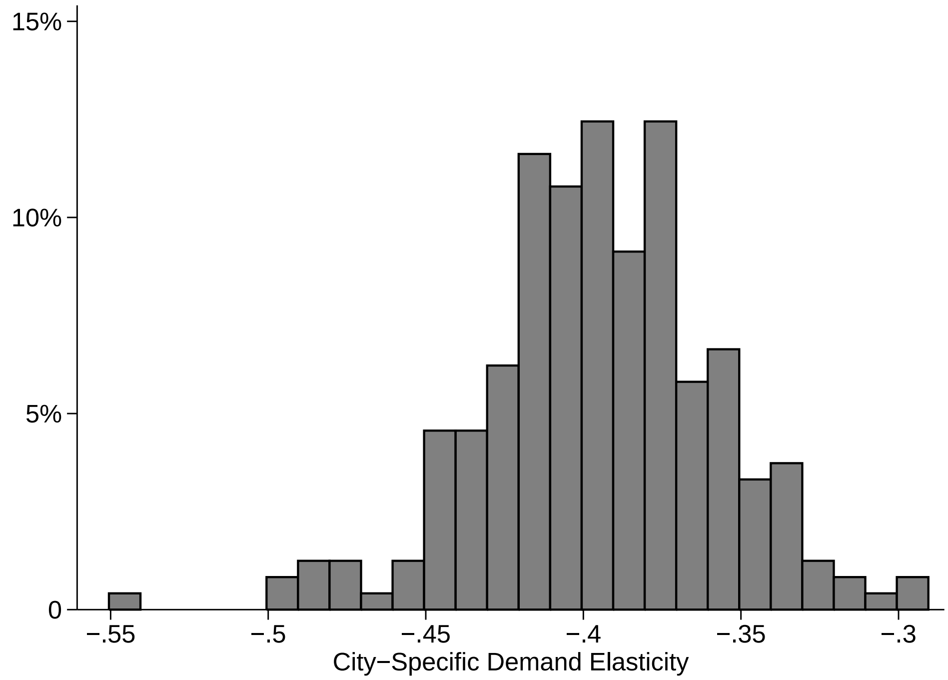

the shares of workers in the MSA that use each of the following modes of transportation to work: drive alone, carpool, bus, subway or rail, walk or bike. Summary statistics for these variables are reported in Table 9.

In theory, gasoline demand is likely to be more elastic in areas where there are one or more alternative modes of transportation available to consumers that are similar in cost and convenience to driving a car. Unfortunately, it is difficult to construct measures of the relative cost and convenience of various modes of transit. Using measures of consumers' actual transit choices (from several years prior to our sample period) hopefully reveals information about the relative attractiveness of the different modes, capturing differences across cities in geographic layout, the quality of the various transit networks, prices of various modes, highway congestion, etc.

We estimate a region-varying elasticity model using the following estimation procedure. Our demand model takes the same form as given in equation 7 except that we hypothesize the following functional form for β_{jd}, the region-varying elasticity of demand: $\beta_{jd} = \gamma_0 + Z'_{jd}\gamma$, where γ_0 ia an unknown parameter, Z_{jd} is the vector MSA-level variables described above, and γ

Table 9: Summary Statistics for City Characteristics
(Number of Cities = 224)

	Mean	S.D.	Median	Min	Max
persons per square mile	1771	932	932	535	6313
ln(persons per square mile)	7.35	0.494	7.32	6.28	8.75
population share over twice the poverty level	0.704	0.071	0.712	0.483	0.898
share commuting by:					
car (alone)	0.823	0.044	0.831	0.518	0.889
subway or rail	0.003	0.012	0.0002	0	0.088
bus	0.016	0.018	0.010	0.001	0.178
carpool	0.125	0.025	0.124	0.082	0.203
walking or bicycle	0.034	0.019	0.029	0.010	0.103

is unknown vector of parameters. Replacing β in equation 7 by this expression for β_{jd} yields:

$$ln(Q_{jd}) = \alpha_j + \lambda_d + \gamma_0 ln(p_{jd}) + ln(p_{jd}) * Z'_{jd}\gamma + ln(n_{jd}) - ln(\hat{\rho}_{jd}) + \epsilon_{jd}, \tag{14}$$

where $ln(p_{jd}) * Z'_{jd}$ means that each element of Z_{jd} is multiplied by $ln(p_{jd})$. We estimate the model using all purchases as well as using only pay-at-pump purchases. The results are reported in Table 10.

The signs and approximate magnitudes of the coefficient estimates are generally consistent across the all-purchases and pay-at-pump specifications, with the exception being that the coefficient estimate on share of workers commuting in a carpool is smaller when using only pay-at-pump purchases. The estimates reveal that, all else equal, more densely populated MSAs and those with more low income households tend to have have more elastic demand for gasoline. MSAs with more workers commuting by carpool or by walking or biking tend to have slightly less elastic demand, while those in which more workers commute by subway or rail have (if anything) slightly more elastic demand.

The magnitudes of the estimated coefficients are fairly small, but this is not surprising given the relatively narrow range of city-specific demand elasticity estimates. For example, cities where the share of commuters walking or biking to work is two standard deviation larger tend to have a demand elasticity that is 0.02 larger in absolute value, all else equal. Cities in which the share of the population with income below twice the poverty level is two standard deviations larger are predicted to have demand elasticities that are around 0.03 lower. While these results are only suggestive, they show that the cross-city differences in demand elasticities estimated by our

Table 10: Estimated Demand Elasticities as a Function of City Characteristics

Dependent Variable = ln(quantity$_{jd}$)

	(all purchases) (1)	(pay-at-pump) (2)
ln(p) = logarithm of price of gasoline	-0.551	-0.420
	(0.047)	(0.033)
ln(p)*ln(persons per square mile)	-0.010	-0.007
	(0.005)	(0.004)
ln(p)*population share over twice the poverty level	0.214	0.207
	(0.040)	(0.029)
ln(p)*share commuting by subway or rail	-0.302	-0.382
	(0.237)	(0.153)
ln(p)*share commuting by bus	0.034	0.009
	(0.171)	(0.113)
ln(p)*share commuting by carpool	0.492	0.178
	(0.120)	(0.076)
ln(p)*share commuting by walking or bicycle	0.475	0.472
	(0.152)	(0.103)
ln(number of transactions)	0.997	0.996
	(0.005)	(0.002)
ln(predicted probability of purchase)	0.006	-0.003
	(0.003)	(0.001)
MSAs Fixed Effects	X	X
Day-of-Sample Fixed Effects	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.

empirical model appear to be correlated with differences in consumers' characteristics and with the existence and attractiveness of other viable modes of transportation.

9 Conclusions

In this study we use high frequency panel data on gasoline prices and expenditures to re-examine the nature of gasoline demand in the U.S. Our demand estimates are significantly more elastic than those of other recent studies. To investigate this discrepancy we aggregate our data and estimate demand models similar to those used in previous studies. The results suggest that using more aggregated data can lead to more inelastic estimates of gasoline demand. In addition, since our data is recorded directly from consumer purchases it may also provide a more accurate measure of actual demand response, generating a more elastic estimate than found in previous studies that use refinery level sales data. 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. We specify a model of gasoline purchase behavior that allows us to separately identify the short run elasticity of gasoline usage or demand from 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. Given that gasoline demand elasticity estimates are commonly used in policy evaluation and in broader economic research, our results provide valuable new evidence that gasoline demand may be more responsive to short term price fluctuations than was previously believed.

Given that gasoline demand elasticity estimates are commonly used in policy evaluation and in broader economic research, our results provide valuable new evidence that gasoline demand may be more responsive to short term price fluctuations than was previously believed. Moreover, the results can substantially impact the inferences one draws when evaluating market disruptions. For example, in early October of 2012 an Exxon refinery near Los Angeles experienced a power outage and was shut down for several weeks. This refinery represents 15% of total gasoline production in the state and 25% of production in Southern California. Several refineries in the state were already out of operation or operating under full capacity and inventories were fairly low, so the unexpected outage led prices in the Los Angeles area to increase by roughly 50 cents (an increase of around 13%) within a matter of days. According to our demand elasticity estimates of approximately -0.4, such a price increase might have caused the quantity demand to fall by as much as 5%, substantially contributing to the alleviation of the temporary supply shortfall. In contrast, using demand elasticity estimates closer to those generated by other recent studies (say for example -.05) would imply a negligible demand response of only 0.05%, suggesting that almost the entire shortfall must have been made up through further withdraws from storage and expediting additional supplies from other markets. Having an accurate estimate of demand response is crucial for understanding how the market adjusts for such disruptions, and our results reveal that demand may play a much more important stabilizing role in the market than has recently been suggested.

References

Cui, Carolyn, "Questions Arise on Gasoline Data," *The Wall Street Journal*, March 21, 2012, p. C1.

Dahl, Carol and Thomas Sterner, "Analysing Gasoline Demand Elasticities: A Survey," *Energy Economics*, July 1991, *13* (3), 203–210.

Espey, Molly, "Gasoline Demand Revisited: An International Meta-analysis of Elasticities," *Energy Economics*, 1998, *20* (3), 273–295.

Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas, "Revisiting the Income Effect: Gasoline Prices and Grocery Purchases," November 2007. NBER Working Paper No. 13614.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling, "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *The Energy Journal*, 2008, *29* (1), 93–114.

Kayser, Hilke A., "Gasoline Demand and Car Choice: Estimating Gasoline Demand Using Household Information," *Energy Economics*, June 2000, *22* (3), 331–348.

Puller, Steven L. and Lorna A. Greening, "Household Adjustment to Gasoline Price Change: an Analysis Using 9 Years of US Survey Data," *Energy Economics*, February 1999, *21* (1), 37–52.

Small, Kenneth A. and Kurt Van Dender, "Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect," *Energy Journal*, 2007, *28* (1), 25–51.