THE RESPONSE OF CONSUMER SPENDING TO CHANGES IN GASOLINE PRICES*

Michael Gelman[a], Yuriy Gorodnichenko[b,c], Shachar Kariv[b], Dmitri Koustas[b],

Matthew D. Shapiro[a,c], Dan Silverman[d,c], and Steven Tadelis[b,c]

August 8, 2016

*This research is carried out in cooperation with a financial aggregation and bill-paying computer and smartphone application (the "app"). We are grateful to the executives and employees who have made this research possible. This research is supported by a grant from the Alfred P. Sloan Foundation. Shapiro acknowledges additional support from the Michigan node of the NSF-Census Research Network (NSF SES 1131500). Gorodnichenko thanks the NSF for financial support.

[a]University of Michigan
[b]University of California, Berkeley
[c]National Bureau of Economic Research
[d]Arizona State University

THE RESPONSE OF CONSUMER SPENDING TO CHANGES IN GASOLINE PRICES

This paper estimates how overall consumer spending responds to changes in gasoline prices. It uses the differential impact across consumers of the sudden, large drop in gasoline prices in 2014 for identification. This estimation strategy is implemented using comprehensive, daily transaction-level data for a large panel of individuals. The estimated marginal propensity to consume (MPC) is approximately one, a higher estimate than estimates found in less comprehensive or well-measured data. This estimate takes into account the elasticity of demand for gasoline and potential slow adjustment to changes in prices. The high MPC implies that changes in gasoline prices have large aggregate effects.

Michael Gelman
Department of Economics
University of Michigan
Ann Arbor MI 48109-1220
mgelman@umich.edu

Yuriy Gorodnichenko
Department of Economics
University of California, Berkeley
Berkeley, CA
and NBER
ygorodni@econ.berkeley.edu

Shachar Kariv
Department of Economics
University of California, Berkeley
Berkeley, CA
kariv@berkeley.edu

Dmitri Koustas
Department of Economics
University of California, Berkeley
Berkeley, CA
dkoustas@econ.berkeley.edu

Matthew D. Shapiro
Department of Economics
and Survey Research Center
University of Michigan
Ann Arbor, MI 48109-1248
and NBER
shapiro@umich.edu

Dan Silverman
Department of Economics
Arizona State University
Tempe AZ
and NBER
Daniel.Silverman.1@asu.edu

Steven Tadelis
Haas School of Business
University of California, Berkeley
Berkeley, CA
and NBER
stadelis@berkeley.edu

I. Introduction

Few macroeconomic variables grab headlines as often and dramatically as do oil prices. In 2014, policymakers, professional forecasters, consumers and businesses all wondered how the decline of oil prices from over $100 per barrel in mid-2014 to less than $50 per barrel in January 2015 would influence disposable incomes, employment, and inflation. A key component for understanding macroeconomic implications of this shock is consumers' spending from resources freed up by lower gasoline prices. Estimating the quantitative impact of this element is central to policy decisions. Yet, because of data limitations a definitive estimate has proved elusive. Fortunately, "big data" has opened unprecedented opportunities to shed new light on the matter. This paper uses extraordinarily detailed transaction-level data provided by a personal financial management service to assess the spending response of consumers to changes in gasoline prices over the 2013-2016 period.

Specifically, we use this information to construct high-frequency measures of spending on gasoline and non-gasoline items for a panel of more than one million consumers. We use cross-consumer variation in the intensity of spending on gasoline interacted with the large, exogenous, and permanent decline in gasoline prices to identify and estimate the marginal propensity to consume (MPC) out of savings generated by reduced gasoline prices. Our baseline estimate of the MPC is approximately one. That is, consumers on average spend *all* of their gasoline savings on non-gasoline items. There are lags in adjustment, so the strength of the response builds over a period of weeks and months.

Our results are useful and informative in several dimensions. First, our estimate of the MPC is largely consistent with the permanent income hypothesis (PIH), which is a theoretical framework that became a workhorse for analyses of consumption, and that has been challenged in previous studies. Second, our findings suggest that, ceteris paribus, falling oil prices can give a considerable boost to the U.S. economy via increased consumer spending (although other factors can offset output growth). Third, and also consistent with the PIH, we show that consumers' liquidity was not important for the strength of the consumer spending response to gasoline price shocks. Fourth, our analysis highlights the importance of having high-frequency transaction data at the household level for estimating consumer reactions to income and price shocks.

This paper is related to several strands of research. The first strand, surveyed in Jappelli and Pistaferi (2010), is focused on estimating consumption responses to income changes.

Typically, studies in this area examine if and how consumers react to anticipated income shocks, which in many cases are transitory. A common finding in this strand is that, in contrast to predictions of the PIH, consumers often spend only upon the realization of an income shock, rather than upon its announcement, although the size of this "excess sensitivity" depends on household characteristics. For example, Kaplan and Violante (2014) and many others document that consumers with low liquidity holdings respond more strongly to anticipated income shocks. At the same time, estimating spending responses to unanticipated, highly persistent income shocks has been challenging, because identifying such shocks is particularly difficult.

We contribute to this literature in several ways. First, we exploit a particularly clear-cut source of variation in household budgets (spending on gasoline) with a number of desirable properties. Specifically, we use a large, salient, unanticipated, permanent (or perceived to be permanent) shock. Second, we examine spending responses at the weekly frequency while, due to data limitations, the vast majority of previous studies estimate responses at much lower frequencies. As we discuss below, the high-frequency dimension allows us to obtain crisp estimates of the MPC and thus provide a more informative input for policy making.

The second strand to which we contribute studies the effects of oil prices on the economy. In surveys of this literature, Hamilton (2008) and Kilian (2008) emphasize that oil price shocks can influence aggregate outcomes via multiple channels (e.g., consumer spending, changes in expectations) but disruption of consumers' (and firms') spending on goods other than energy is likely to be a key mechanism for amplification and propagation of the shocks. Indeed, given the low elasticity of demand for gasoline, changes in gasoline prices can materially affect non-gasoline spending budgets for a broad array of consumers. As a result, a decrease in gasoline prices can generate considerable savings for consumers which could be put aside (e.g., to pay down debt or save) or used to spend on items such as food, clothing, furniture, etc.

Strikingly, despite the importance of the MPC out of gasoline savings, research on the sensitivity of consumer non-gasoline spending to changes in the gasoline price has been scarce. One reason for the dearth of research on the matter has been data limitations. Available household consumption data tend to be low frequency, whereas consumer spending, gasoline prices, and consumer expectations can change rapidly. For example, the interview segment of the U.S. Consumer Expenditure Survey (CEX) asks households to recall their spending over the

previous month. These data likely suffer from recall bias and other measurement errors that could attenuate estimates of households' sensitivity to changes in gasoline prices (see Committee on National Statistics 2013). The diary segment of the CEX has less recall error, but the panel dimension of the segment is short (14 days), making it difficult to estimate the consumer response to a change in prices. Because the CEX is widely used to study consumption, we do a detailed comparison of our approach using the app data with what can be learned from using the CEX. We find that analysis based on the CEX produces noisy and attenuated estimates.

Grocery store barcode data, such as from AC Nielsen, have become a popular alternative to measure higher-frequency spending. These data cover, however, only a limited category of goods. For example, gasoline spending by households is not collected in AC Nielsen, making it impossible to exploit heterogeneity in gasoline consumption across households. As a result, most estimates of MPC tend to be based on time series variation in aggregate series (see e.g. Edelstein and Kilian 2009).

There are a few notable exceptions, nevertheless. Using loyalty cards, Hastings and Shapiro (2013) are able to match grocery barcode data to gasoline sold at a select number of grocery stores with a gasoline station on site. Perhaps not surprisingly, we show that households typically visit multiple gasoline station retailers in a month, suggesting limitations to focusing on consumer purchases at just one gasoline retailer. There is also some recent work using household data to identify a direct channel between gasoline prices and non-gasoline spending. Gicheva, Hastings and Villas-Boas (2010) use weekly grocery store data to examine the substitution to sale items as well as the response of total spending. They find that households are more likely to substitute towards sale items when gasoline prices are higher, yet they focus only on a subset of goods bought in grocery stores (cereal, yogurt, chicken and orange juice), making it difficult to extrapolate.

Perhaps the closest work to ours is a policy report produced by the J.P. Morgan Chase Institute (2015), which also uses "big data" to examine the response of consumers to the 2014 fall in gasoline prices, and finds an average MPC of approximately 0.6. This report differs from our study in both its research design and its data. Most important, our data include a comprehensive view of spending, across many credit cards and banks. In contrast, the Chase report covers a vast number of consumers, but information on their spending is limited to Chase

accounts. If, for example, consumers use a non-Chase credit card or checking account, any spending on that account would be missed in the J.P. Morgan Chase Institute analysis, and measurement of household responses may therefore be critically incomplete. In this paper, we show that an analysis based on accounts in one financial institution leads to a significantly attenuated estimate of the response of spending to changes in gasoline prices.

This paper proceeds as follows. Section II describes trends in gasoline prices, putting the recent experience into historical context. In Section III, we discuss the data, Section IV describes our empirical strategy, and Section V presents our results. Specifically, we report baseline estimates of the MPC and the elasticity of demand for gasoline. We contrast these estimates with the comparable estimates one can obtain from alternative data. In Section V we also explore robustness of the baseline estimates and potential heterogeneity of responses across consumers. Section VI concludes.

II. Recent Changes in Gasoline Prices: Unanticipated, Permanent and Exogenous

In this section, we briefly review recent dynamics in the prices of oil and gasoline and corresponding expectations of future prices. We document that the collapse of oil and gasoline prices in 2014-2015 was highly persistent, unanticipated, and exogenous to demand conditions in the United States. These properties of the shock are important components of our identification strategy.

A. Unanticipated and Permanent

In Panel A of Figure 1, the solid black line shows the spot price of gasoline at New York Harbor, an important import and export destination for gasoline. The New York Harbor price is on average 70 cents lower than average retail prices, although the two series track each other very closely. The dashed line shows the one-year-ahead futures price for that date. The futures price tracks the spot price closely, suggesting the market largely treats gasoline price as a random walk—i.e., the best prediction for one-year-ahead price is simply the current price.

Panel B shows the difference between the realized and predicted spot price. The behavior of one-year-ahead forecast errors indicates that financial markets anticipated neither the run-up

nor the collapse of gasoline prices in 2007-2009. Likewise, the dramatic decline in gasoline prices in 2014-2015 was not anticipated.

The Michigan Survey of Consumers has asked households about their expectations for changes in gasoline prices over the next one-year and five-year horizons. Panel C of Figure 1 plots the mean and median consumer expectations along with the actual price and the one-year-ahead future price from Panel A. While consumers expect a slightly higher price relative to the present price than the futures market, the basic pattern is the same as in Panel A: the current price appears to be a good summary of expected future prices. Consistent with this observation, Anderson, Kellogg and Sallee (2012) fail to reject the null of a random walk in consumer expectations for gasoline prices. Thus, consumers perceive changes in gasoline prices as permanent. Also similar to the financial markets, consumers were not anticipating large price changes in 2007-2009 or 2014-2015 (Panel D).

Figure 1 shows large movements in prices during the Great Recession (shaded). Unlike the recent episode that is the subject of this paper, we would *not* use it to identify the MPC because this fluctuation in commodity prices in the Great Recession surely represents an endogenous response to aggregate economic conditions.

When put into historical context, the recent volatility in gasoline prices is large. Table 1 ranks the largest one-month percent changes in oil prices since 1947. When available, the change in gasoline prices over the same period is also shown.[1] The price drops in 2014-2015 are some of the largest changes in oil and gasoline prices in the last 60 years. Note that in 1986, gasoline prices and oil prices actually moved in opposite directions, indicating that the process generating gasoline prices can sometimes differ from oil.

B. Exogenous

Why did prices of oil and oil products such as gasoline fall so much in 2014-2015? While, obviously, many factors could have contributed to the dramatic decline in the prices, the consensus view, summarized in Baffes et al. (2015), attributes a bulk of the decline to supply-

[1] Oil spot prices exist back to 1947, while the BLS maintains a gasoline price series for urban areas back to 1976. In our analysis, we use AAA daily gasoline prices retrieved from Bloomberg (3AGSREG). The series comes from a daily survey of 120,000 gasoline stations. These data almost perfectly track another series from the EIA which are point in-time estimates from a survey of 900 retail outlets as of 8am Monday.

side factors. Specifically, this view emphasizes that key forces behind the decline were first, OPEC's decision to abandon price support, and second, rapid expansion of oil supply from alternative sources (shale oil in the U.S., Canadian oil sands, etc.). Consistent with this view, other commodity prices had modest declines during this period, which would not have happened if the decline in oil prices was driven by global demand factors. Observers note that the collapse of oil prices in 2014-2015 is similar in many ways to the collapse in 1985-1986, when more non-OPEC oil supply came from Mexico, the North Sea and other sources, and OPEC also decided to abandon price support. In short, available evidence suggests that the 2014-2015 decline in oil prices is a shock that was supply-driven and exogenous to U.S. demand conditions. In contrast, for the 2007-2009 episode, Hamilton (2009) and others observe that the run up in oil and gasoline prices around 2007-2009 can be largely attributed to booming demand, stagnant production, and speculators, and the consequent decline of the prices during this period, to collapsed global demand (e.g. the Great Recession and Global Financial Crisis).

III. Data

Our analysis uses high-frequency data on spending from a financial aggregation and bill-paying computer and smartphone application (henceforth, the "app").[2] The app had approximately 1.4 million active users in the U.S. in 2013.[3] Users can link almost any financial account to the app, including bank accounts, credit card accounts, utility bills, and more. Each day, the app logs into the web portals for these accounts and obtains central elements of the user's financial data including balances, transaction records and descriptions, the price of credit and the fraction of available credit used.

We draw on the entire de-identified population of active users and data derived from their records from January 2013 until February 2016. The app does not collect demographic information directly and instead uses a third party business that gathers both public and private sources of demographics, anonymizes them, and matches them back to the de-identified dataset.

[2] These data have previously been used to study the high-frequency responses of households to shocks such as the government shutdown (Gelman et al. 2016) and anticipated income, stratified by spending, income and liquidity (Gelman et al. 2014).

[3] All data are de-identified prior to being made available to the project researchers. Analysis is carried out on data aggregated and normalized at the individual level. Only aggregated results are reported.

Table 1 in Gelman et al. (2014) (replicated in Appendix Table A1) compares the gender, age, education, and geographic distributions in a subset of the sample to the distributions in the U.S. Census American Community Survey (ACS), representative of the U.S. population in 2012. The app's user population is not representative of the U.S. population, but it is heterogeneous, including large numbers of users of different ages, education levels, and geographic location.

A. Identifying Spending Transactions

Not every transaction reported by the app is spending. For example, a transfer of funds from one account to another is not spending. To avoid double counting, we exclude transfers across accounts, as well as credit card payments from checking accounts that are linked within the app. If an account is not linked, but we still observe a payment, we count this as spending when the payment is made. We identify transfers in several ways. First, we search if a payment from one account is matched to a receipt in another account within several days. Second, we examine transaction description strings to identify common flags like "transfer", "tfr", etc. To reduce the chance of double counting, we exclude the largest single transaction that exceeds $1,000 in a given week, as this kind of transaction is very heavily populated by transfers, credit card payments, and other non-spending payments (e.g., payments to the U.S. Internal Revenue Service). We include cash withdrawals from the counter and ATM in our measure of spending. To ensure that accounts in the app data are reasonably linked and active, we keep all users who were in the data for at least 8 weeks in 2013 and who did not have breaks in their transactions for more than two weeks. More details are provided in Appendix A.

B. Using Machine Learning to Classify Type of Spending

Our analysis requires classification of spending by type of goods. To do so, we address several challenges in using transactional data from bank accounts and credit cards. First, transactional data are at the level of a purchase at an outlet. For many purchases, a transaction will include many different goods. In the case of gasoline, purchases are carried out mainly at outlets that exclusively or mainly sell gasoline. Hence, gasoline purchases are relatively easy to identify in transactional data. Second, for the bulk of transactions in our data, we must classify the outlet from the text of the transaction description, rather than classifications provided by financial

institutions. We therefore use a machine learning (ML) algorithm to classify spending based on transaction descriptions. In this section, we provide an outline of the classification routine, and compare our ML predictions in the data provided by the app with external data.

We use an ML algorithm to construct a set of rules for classifying the data as gasoline or non-gasoline. This requires a training data set to build a classification model, and a testing data set not used in the training step to validate the model predictions. Two of the account providers in the data classify spending directly in the transaction description strings, using merchant category codes (MCCs). MCCs are four digit codes used by credit card companies to classify spending and are also recognized by the U.S. Internal Revenue Service for tax reporting purposes. Our main MCC of interest is 5541, "Automated Fuel Dispensers." Purchases of gasoline could also fall into MCC code 5542, "Service Stations," which in practice covers gasoline stations with convenience stores.[4] Because distinguishing gasoline purchases classified as 5542 or 5541 is nearly impossible with the information in transaction descriptions,[5] we group transactions with these two codes together.

A downside of this approach is that transactions at the convenience store associated with a gasoline station can be classified as a purchase of gasoline; that is, buying a food item in a gasoline station's convenience store can be classified as a purchase of gasoline. According to the National Association of Convenience Stores (NACS), which covers gasoline stations, purchases of non-gasoline items at gasoline stations with convenience stores (i.e. "Service Stations") account for about 30 percent of sales at "Service Stations." Although the app data do not permit us to differentiate gasoline and non-gasoline items at "Service Stations," we can use transaction data from "Automated Fuel Dispensers" (which do not have an associated convenience store), as well as external survey evidence to separate purchases of non-gasoline items from purchases of gasoline. Specifically, according to the 2015 NACS Retail Fuels Report (NACS 2015), only 35 percent of gasoline purchases are associated with going inside a gasoline station's store. Conditional on going inside the store, the most popular activities are to "pay for gasoline at the register" (42%), "buy a drink" (36%), "buy a snack" (33%), "buy cigarettes" (24%), and "buy lottery tickets" (22%). The last four items are likely to be associated with small amounts of

[4] To be clear, "Service Stations" do not include services such as auto repairs, motor oil change, etc.

[5] E.g., a transaction string with word "Chevron" or "Exxon" could be classified as either MCC 5541 or MCC 5542.

spending. This conjecture is consistent with the distribution of transactions for "Service Stations" and "Automated Fuel Dispensers." In particular, approximately 60 percent of transactions at "Service Stations" are less than $10 while the corresponding share for "Automated Fuel Dispensers" is less than 10 percent. As we discuss below, the infrequent incidence of gasoline purchases totaling less than $10 is also consistent with other data sources. Thus, we exclude transactions less than $10 to filter out purchases of non-gasoline items at "Service Stations."

Using one of the two providers with MCC information (the one with more data), we train a Random Forest ML model to create binary classifications of transactions into those made at a gasoline station/service station and those that were made elsewhere. We then use the second providers to validate the quality of our ML model.[6] The ML model is able to classify spending with approximately 90% accuracy in the second provider not used to train the model, which is a high level of precision. More details on the procedure can be found in Appendix B.

We can also use the app data to investigate which gasoline stations consumers typically visit. The top ten chains of gasoline stations in the app data account for most of gasoline spending. On average, the app data suggest that the typical consumer does 66 percent of his or her gasoline spending in one chain and the rest of gasoline spending is spread over other chains. Thus, while for a given consumer there is a certain degree of concentration of gasoline purchases within a chain, an analysis focusing on just one gasoline retailer only, such as in Gicheva, Hastings and Villas-Boas (2010) or Hastings and Shapiro (2013), particularly one not in the top ten chains, would still miss a substantial amount of gasoline spending.

C. Comparison with the Consumer Expenditure Survey

We compare our measures of gasoline and non-gasoline spending with similar measures from the Consumer Expenditure Survey (CEX). We use both the CEX Diary Survey and Interview Survey. In the diary survey, households record all spending in written diaries for 14 days. Therefore, this survey provides an estimate of daily gasoline spending that should be comparable

[6] Card providers use slightly different transaction strings, and one may be concerned that training the model on a random subsample of data from both card providers, and testing it on another random subsample, can provide a distorted sense of how our ML model performs on data from other card providers. Thus, using a card from one account provider to train, and testing on an entirely different account provider, helps to assure that the ML model is valid outside of the estimation sample. Classification of transactions based on ML applied to both card providers yields very similar results.

to the daily totals we observe in the app. In Figure 2, we compare the distribution of spending in our data (solid lines) and in the diary survey (dash line). We find that the distributions are very similar, with one notable exception: the distribution of gasoline purchases in the app data has more mass below $10 (solid gray line) than the CEX Diary data. As we discussed above, this difference is likely to be due to our inability to differentiate gasoline purchases and non-gasoline purchases at "Service Stations." In what follows, we restrict our ML predictions to be greater than $10 (solid black line).

The CEX Diary Survey provides only a limited snapshot of households' gasoline and other spending. In particular, since a household on average only makes 1 gasoline purchase per week in the diary, we expect only to observe 2 gasoline purchases per household, which can be a noisy estimate of gasoline spending at the household level. Idiosyncratic factors in gasoline consumption that might push or pull a purchase from one week to the next could drastically influence the measure of a household's gasoline purchases by 50% or more. In addition, because the survey period in the diary is so short, household fixed effects cannot be used to control for time-invariant household heterogeneity. Hence, while a diary survey could be a substitute for the app data in principle, the short sample of the CEX diary survey makes it not a good substitute in practice.[7]

The CEX Interview Survey provides a more complete measure of total spending, as well as a longer panel (4 quarters), from which we can make a comparison with estimates based on spending reported by the app at longer horizons. Panel A of Figure 3 reports the histogram (bin size is set to $1 intervals) of monthly spending on gasoline in the CEX Interview data for 2013-2014.[8] The distribution has clear spikes at multiples of $50 and $100 with the largest spikes at $0 and $200. In contrast, the distribution of gasoline purchases in the app data has a spike at $0 but the rest of the distribution exhibits considerably less bunching, particularly at large values like $200 or $400 that correspond with reporting $50 or $100 per week, respectively. In addition, the

[7] We have done a comparison of the CEX diary spending for January 2013 through December 2014 (the last time period that the CEX is available at this point). In a regression of log daily spending for days with positive spending on month time effects and day of week dummies, the month effects estimated in the CEX and app have a correlation of 0.77. (Finer than monthly comparison of the app and CEX is not possible because the CEX provides only the month and day of week, but not the date, of the diary entry.)

[8] The survey question in the CEX Interview Survey asks households to report their "Average monthly expense for gasoline."

distribution of gasoline spending has a larger mass at smaller amounts in the app data than in the CEX Interview data. These differences are consistent with recall bias in the CEX Interview Survey data. As argued by Binder (2015), rounding in household surveys can reflect a natural uncertainty of households about how much they spent in this category.

Table 2 compares moments for gasoline and non-gasoline spending across the CEX and the app data. We find that the means are similar across data sources. For example, mean biweekly gasoline spending in the CEX Diary Survey is $84.72, while the app counterpart is $81.68. Similarly, non-gasoline spending is $1,283 in the CEX Diary Survey and $1,469 in the app data. The standard deviation, however, tends to be larger in the app data than in the CEX, which reflects a thick right tail of spending in the app data. This pattern is consistent with under-representation of higher-income households in the CEX, a well-documented phenomenon (Sabelhaus et al. 2015). Medians and inter-quartile ranges, measures of central tendency and variations that are less sensitive to the tails, are smaller in the app data than the CEX Diary Survey's counterparts. The moments in the CEX Interview Survey (quarterly frequency) are generally closer to the moments in the app data. For example, mean (median) spending on gasoline is $647 ($540) in the CEX Interview Survey data and $628 ($475) in the app data, while the standard deviations (interquartile ranges) are $531 ($630) and $588 ($660) respectively. In each panel of Table 2, we also compare the distribution of the ratio of gasoline spending to non-gasoline spending, a central ingredient in our analysis. The moments for the ratio in the CEX and the app data are very similar. For instance, the mean ratio is 0.08 for the CEX Interview Survey and 0.07 for the app data, while the standard deviation of the ratio is 0.07 for both the CEX Interview Survey and the app data.[9]

In summary, spending in the app data is similar to spending in the CEX data. Some of the differences could reflect survey recall bias, consumers buying gasoline on cards that are not linked to the app (such as credit cards specific to gasoline station chains), the ML procedure missing some gasoline stations, or gasoline spending done in cash that we could not identify.[10] We will address these potential issues in our robustness tests.

[9] Appendix Figure 3 shows the density of the gasoline to non-gasoline spending ratio for the CEX and app data.

[10] According to the 2015 NACS Retail Fuels Report (NACS 2015), approximately 80 percent of gasoline purchases are made with debit or credit card.

IV. Empirical Strategy

The discourse on potential macroeconomic effects of a fall in gasoline prices centers on the question of how savings from the fall in gasoline prices are used by consumers. Specifically, policymakers and academics are interested in the marginal propensity to consume (MPC) from savings generated by reduced gasoline prices. Define *MPC* as

$$dC_{it} = -MPC * d(P_t Q_{it}) \quad (1)$$

where i and t index consumers and time, C is spending of non-gasoline items, P is the price of gasoline, and Q is the quantity of consumed gasoline. Note that we define the *MPC* as an *increase* in spending in response to a *decrease* in the price of gasoline.[11]

Equation (1) is a definition, not a behavioral relationship. Of course, Q_{it}, the quantity of gasoline purchased, and overall non-gasoline spending, C_{it}, are simultaneously determined, with simultaneity being an issue at the individual as well as aggregate level. In this section, we develop an econometric relationship that yields identification of the *MPC* based on the specific sources of variation of gasoline prices discussed in the previous sections.

At the macroeconomic level, one important determinant of gasoline spending is aggregate economic conditions. As discussed in Section II, the 2007-2008 collapse in gasoline prices has been linked to the collapse in global demand due to the financial crisis; demand for gasoline fell driving down the price, at the same time demand was falling for other goods. Individual-level shocks are another important source of simultaneity bias and threat to identification. Consider a family going on a road trip to Disneyland; this family will have higher gasoline spending (long road trip) and higher total consumption in that week due to spending at the park. Yet another example is a person who suffers an unemployment spell; this worker will have lower gasoline spending (not driving to work) and lower other spending (a large negative income shock).

This discussion highlights that gasoline purchases and non-gasoline spending are affected by a variety of shocks. Explicitly modelling all possible shocks, some of which are expected in advance by households (unobservable to the econometrician), would be nearly impossible. Fortunately, this is not required to properly identify the policy-relevant parameter–the sensitivity

[11] The MPC is likely to be different across groups of people, but our notation and estimation refers to the average MPC.

of non-gasoline spending to changes in gasoline spending induced by exogenous changes in the price of gasoline. This parameter may be interpreted as a partial derivative of non-gasoline spending with respect to the price of gasoline and thus could be mapped to a coefficient estimated in a regression. For this, we only need to satisfy a weaker set of conditions. First, we need exogenous, unanticipated shocks to gasoline prices. These shocks should be unrelated to the regression residual absorbing determinants of non-gasoline consumption unrelated to changes in gasoline prices. Second, we need to link non-gasoline spending to the price of gasoline (i.e., P_t), rather than purchases of gasoline ($P_t Q_{it}$).

As we established in Section II, shocks to gasoline prices in the period of our analysis were unanticipated, exogenous, and permanent so that we have an exogenous source of variation, an essential ingredient for consistent estimation with regression analysis. To link the partial derivative of interest to a regression coefficient and to link it with cross-sectional variation in pre-determined propensity to spend on gasoline, we manipulate equation (1) as follows:

$$
\begin{aligned}
\frac{dC_{it}}{\bar{C}_i} &= d\log C_{it} = -MPC \times \frac{d(P_t Q_{it})}{\bar{C}_i} = -MPC \times \frac{d(P_t Q_{it})}{(\overline{PQ})_i} \times \frac{(\overline{PQ})_i}{\bar{C}_i} \\
&= -MPC \times \frac{d(P_t Q_{it})}{(\overline{PQ})_i} \times s_i \\
&= -MPC \times \frac{\bar{Q}_i dP_t + \bar{P} dQ_{it}}{(\overline{PQ})_i} \times s_i \\
&= -(MPC \times s_i \times \frac{dP_t}{\bar{P}} + MPC \times s_i \times \frac{dQ_{it}}{\bar{Q}_i}) \\
&= -MPC \times s_i \times d\log P_t - MPC \times s_i \times \left(\frac{dQ_{it}}{\bar{Q}_i} \times \frac{\bar{P}}{dP_t}\right) \times \frac{dP_t}{\bar{P}} \\
&= -MPC \times s_i \times d\log P_t - MPC \times s_i \times \epsilon \times d\log P_t \\
&= -MPC \times (1+\epsilon) \times s_i \times d\log P_t \qquad (2)
\end{aligned}
$$

where bars denote steady-state values, $s_i \equiv \frac{(\overline{PQ})_i}{\bar{C}_i}$ is the ratio of gasoline spending to non-gasoline spending,[12] and ϵ is the price elasticity of demand for gasoline (a negative number). Now the only source of time variation in the right-hand side of the equation is the price of gasoline. The

[12] We calculate s_i as the ratio of consumer *i*'s annual spending on non-gasoline items to his/her annual spending on gasoline in 2013. Using annual frequency in this instance helps to address seasonal variation in gasoline spending as well as considerable high frequency variation in the intensity of gasoline spending (e.g., trips to gasoline stations, spending per trip). Additionally, the use of 2013 data to calculate the share makes it pre-determined with respect to the shock to gasoline prices in the estimation period.

identifying variation in equation (2) comes from the interaction of the aggregate gasoline price shock, with consumers' exposure to the shock as measured by s_i. Thus, when we regress log non-gasoline spending on the log of gasoline price multiplied by the ratio of gasoline spending to non-gasoline spending, we get $-MPC(1+\epsilon)$.

Note that we have an estimate of −MPC scaled by $1+\epsilon$, but the scaling should be small if demand is inelastic. As discussed below, there is some variation in the literature on ϵ's estimated using household versus aggregate data. To ensure that a measure of ϵ is appropriate for our sample, we note:

$$\begin{aligned} d\log P_t Q_{it} &= d\log P_t + d\log Q_{it} \\ &= d\log P_t + d\log P_t \frac{d\log Q_{it}}{d\log P_t} = \left(1+\frac{d\log Q_{it}}{d\log P_t}\right) \times d\log P_t \\ &= (1+\epsilon) \times d\log P_t. \end{aligned} \tag{3}$$

Similar to equation (2), the only source of time variation in the right-hand side of equation (3) is the price of gasoline. Thus, a regression of $d\log P_t Q_{it}$ on $d\log P_t$ yields $\widehat{(1+\epsilon)}$, which is the partial derivative of gasoline spending with respect to the price of gasoline, and the residual in this regression absorbs determinants of gasoline purchases unrelated to the changes in the price of gasoline.[13] The estimated $(1+\epsilon)$ and $-MPC(1+\epsilon)$ can be combined to obtain the MPC.

In the derivation of equations (2) and (3) we deliberately did not specify the time horizon over which sensitivities are computed, as these may vary with the horizon. For example, with lower prices, individuals may use their existing cars more intensively or may purchase less fuel-efficient cars. There may be delays in adjustment to changes in prices (e.g., search for a product). It might take time to notice the price change (Coibion and Gorodnichenko 2015). The very-short-run effects may also depend on whether a driver's tank is full or empty when the shock hits.

To obtain behavioral responses over different horizons, we build on the basic derivation above and estimate a multi-period "long-differences" model, where both the MPC and the price elasticity are allowed to vary with the horizon. Additionally, we introduce aggregate and idiosyncratic shocks to overall spending, as well as idiosyncratic shocks to gasoline spending. Hence,

[13] Because the dependent variable is spending on gasoline rather than volume of gasoline, elasticity ϵ estimated by this approach also includes substitution across types of gasoline (Hastings and Shapiro 2013).

$$\Delta_{\mathrm{k}} \log C_{it} = \beta_k \times s_i \times \Delta_{\mathrm{k}} \log P_t + \psi_t + \varepsilon_{it} \quad (4)$$

$$\Delta_{\mathrm{k}} \log PQ_{it} = \delta_k \Delta_{\mathrm{k}} \log P_t + u_{it} \quad (5)$$

where $\beta = -MPC(1+\epsilon)$, $\delta = (1+\epsilon)$, $\Delta_k x_t = x_t - x_{t-k}$ is a *k*-period-difference operator, ψ_t is the time fixed effect, and ϵ_{it} and u_{it} are individual-level shocks to spending.[14] By varying k, we can recover the average impulse response over k-periods so that we can remain agnostic about how quickly consumers respond to a change in gasoline prices.[15]

Note that gasoline and oil prices are approximately random walks and thus $\Delta_{\mathrm{k}} \log P_t$ can be treated as an unanticipated, permanent shock. To the extent oil prices and, hence, gasoline prices are largely determined by global factors or domestic supply shocks, rather than domestic demand—which is our maintained assumption for our sample period—OLS yields consistent estimates of MPC and ϵ. Formally, we assume that the idiosyncratic shocks to spending are orthogonal to these movements in gasoline prices. Given the properties of the shock to gasoline prices in 2014-2015, the PIH model predicts that the response of spending from the resulting change in resources should be approximately equal to the change in resources ($MPC \approx 1$) and take place quickly.

The approach taken in specifications (4) and (5) has several additional advantages econometrically. First, as discussed in Griliches and Hausman (1986), using "long differences" helps to enhance signal-to-noise ratio in panel data settings. Second, specifications (4) and (5) allow straightforward statistical inference. Because our shock ($\Delta_k \log P_t$) is national and we expect serial within-user correlation in spending, we cluster standard errors on two dimensions: time and person. This simplicity is particularly convenient in our case because we estimate equations (4) and (5) as a system to recover MPC from estimates of $-MPC(1+\epsilon)$ and $(1+\epsilon)$.

[14] Note that there are time effects only in equation (4). Since we have argued that changes in gasoline prices are exogenous over the time period, time effects are not needed for consistency of estimation of either (4) or (5). In (4), they may improve efficiency by absorbing aggregate shocks to overall spending. It is not possible to include time effects in (5) because they would completely absorb the variation in gasoline prices. But again note that the presence of an aggregate component in *u* does not make the estimates of δ biased under our maintained assumption that gasoline prices are exogenous to the U.S. economy in the estimation period. (The standard errors account for residual aggregate shocks.)

[15] For example, if $\log C_{it} = \sum_{s=0}^{\infty} \psi_s shock_{t-s} + u_t$ and u_t summarizes variation orthogonal to the shock series of interest shock, then the impulse response is $\{\psi_s\}_{s=0}^{\infty}$ and the long-difference regression recovers $\beta_k = k^{-1} \sum_{s=0}^{k-1} \psi_s$.

To summarize, our econometric framework identifies the MPC from changes in gasoline prices by interacting two sources of variation: a large, exogenous, and permanent change in gasoline prices, with the pre-determined share of spending on gasoline. The econometric specification also accounts for the response of spending on gasoline to lower prices by allowing a non-zero elasticity of demand for gasoline and allowing for lagged adjustment of gasoline spending to changes in gasoline prices.

V. Results

In this section, we report estimates of MPC and ϵ for different horizons, frequencies, and populations. We also compare estimates based on our app data to the estimates based on spending data from the CEX.

A. Sensitivity of Expenditure to Gasoline Prices

We start our analysis with the estimates of MPC and ϵ at weekly frequency for different response horizons. Panel A of Figure 4 shows $\hat{\epsilon}$ for $k = 0, \dots, 26$ weeks. Table 3, Row 1, gives the point estimates for selected horizons. The point estimates indicate that elasticity of demand for gasoline is increasing in the horizon (i.e., over time, consumers have greater elasticity of demand): estimated elasticity changes from -0.20 at the horizon of 15 weeks to -0.24 at the horizon of 25 weeks. Confidence intervals are very wide at short horizons; estimates become quite precise at horizons of 12 weeks and longer.

This estimate is broadly in line with previously reported estimates. Using aggregate data, the results in Hughes, Knittel and Sperling (2008) suggest that U.S. gasoline demand is significantly more inelastic today compared with the 1970s. Regressing monthly data on aggregate per capita consumption of gasoline on changes in gasoline prices, they estimate a short-run (monthly) price elasticity of -0.034 to -0.077 for the 2001 to 2006 period, compared with -0.21 to -0.34 for the 1975-1980 period. The Environmental Energy Institute (EIA 2014) also points to an elasticity close to zero, and also argues this elasticity has been trending

downward over time.[16] In contrast to Hughes, Knittel and Sperling (2008), our findings suggest that gasoline spending could still be quite responsive to gasoline price changes. In general, our results lie in between the Hughes, Knittel and Sperling's estimates and previous estimates using household expenditure data to measure gasoline price elasticities. Puller and Greening (1999) and Nicol (2003) both use the CEX interview survey waves from the 1980s to the early 1990s to estimate the elasticity of demand. The approaches taken across these papers are very different. Nicol's (2003) approach is to estimate a structural demand system. Puller and Greening (1999), on the other hand, take advantage of the CEX modules about miles traveled that were only available in the 1980s, as well as vehicle information. Both of these papers find higher price elasticities of demand at the quarterly level, with estimates in Nicol (2003) ranging from -0.185 for a married couple with a mortgage and 1 child, to -0.85 for a renter with two children, suggesting substantial heterogeneity across households. Paul and Greening's (1999) baseline estimates are -0.34 and -0.47, depending on the specification. A more recent paper by Levin, Lewis and Wolak (2016) uses city level price data and city level expenditure data obtained from Visa credit card expenditures. They estimate the elasticity of demand for gasoline to be closer to ours, but still higher, ranging from −0.27 to −0.35. Their data are less aggregate than the other studies, but more aggregate than ours because we observe individual level data. Also, we observe expenditures from all linked credit and debit cards and are not restricted only to Visa.

Panel B of Figure 4 shows the dynamics of $\widehat{MPC}$ over the same horizons with point estimates at selected horizons in the first row of Table 3. At short time horizons (contemporaneous and up to 3 weeks), the estimates vary considerably from nearly 2 to 0.5 but the estimates are very imprecise. Starting with the four-week horizon, we observe that $\widehat{MPC}$ steadily rises over time and becomes increasingly precise. After approximately 12 weeks, $\widehat{MPC}$ stabilizes at just below 1.0 with a standard error of 0.3. This estimate suggests that consumers spend nearly *all* their gasoline savings on non-gasoline items. The standard errors are somewhat smaller at monthly horizons (4-5 weeks) since the shock. We suspect this is because the residual variance in consumption tends to be lower at monthly frequency due to factors like recurring

[16] EIA (2014) reports, "The price elasticity of motor gasoline is currently estimated to be in the range of -0.02 to -0.04 in the short term, meaning it takes a 25% to 50% decrease in the price of gasoline to raise automobile travel 1%. In the mid 1990's, the price elasticity for gasoline was higher, around -0.08."

spending, bill pay, and rules of thumb/behavioral reasons (shopping once per month), while in other weeks, the consumption process has considerably more randomness.

There are not many estimates of the MPC derived from changes in gasoline prices. The JPMorgan Institute (2015) report examines the same time period that we do using similar data. It finds an MPC of 0.6, much lower than our estimate. This finding likely arises from the use of data from a single financial institution rather than our more comprehensive data. This is an important advantage of the app data because many consumers have multiple accounts across financial institutions. The app's users have accounts on average in 2.6 different account providers (the median is 2). As a result, we have a more complete record of consumer spending. In contrast, previous "big data" attempts to estimate the MPC from gasoline savings had information from a single financial institution. To illustrate the importance of this point, we rerun our specification focusing on a subgroup of consumers with accounts at the top three largest providers.[17] Specifically, we restrict the sample to accounts only at a specific provider so that we can mimic the data observed by a single provider. In rows (2), (4) and (6) of Table 3 we report estimates of ϵ and the MPC at horizons 5, 15 and 25 weeks for the case when we use *any* account at the provider. The MPC estimates based on data observed by a single provider are lower and have larger standard errors than the baseline, full-data MPC estimates reported in row (1). For example, the $\widehat{MPC}$ for Provider 1 (row 2) at the 25-week horizon is 0.515, which is approximately half of the baseline $\widehat{MPC}$ at 0.963, but the standard error for the former estimate is 0.387, so that we cannot reject equality of the estimates as well as equality of the former estimate to zero.

One may be concerned that having only one account with a provider may signal incomplete information because the user did not link all accounts with the app. To address this concern, we restrict the sample further to consider users that have at least one checking and one credit-card account with a given provider. In this case, one may hope that the provider is servicing "core" activities of the user. In rows (3), (5) and (7), we re-estimate our baseline specification with this restriction. We find estimates largely similar to the case of any account,

[17] These providers cover 49.6 percent of accounts in the data and 55.0 percent of total spending.

that is, the estimated sensitivity to changes in gasoline prices is attenuated and more imprecise relative to the baseline where we have accounts linked across multiple providers.

These results for the single-provider data are consistent with the view that consumers can specialize their card use. For example, one card (account) may be used for gasoline purchases while another card (account) may be used for other purchases. In these cases, because single-provider information systematically misses spending on other accounts, MPCs estimated on single-provider data could be attenuated severely.

B. Robustness

While our specification has several important advantages, there are nevertheless several potential concerns. First, if s_i in specification (4) is systematically underestimated because a part of gasoline spending is missing from our data, for instance, due to gasoline retailer cards that are not linked to the app, then our estimate of the MPC will be mechanically higher. Second, suppose instead that we are misclassifying some spending, or that consumers buy a large portion of their gasoline in cash, so that this spending shows up in our dependent variable. Misclassifying gasoline spending as non-gasoline spending will generate a positive correlation between gasoline spending and the gasoline price. Third, while a random walk may be a good approximation for the dynamics of gasoline prices, one may be concerned that gasoline prices have a predictable component, so that estimated reaction mixes up responses to unanticipated and predictable elements of gasoline prices. Indeed, some changes in gasoline prices are anticipated due to seasonal factors.[18]

A practical implication of the first concern (i.e., cases where consumers use gasoline retailer cards that are not linked to the app) is that consumers with poorly linked accounts should have zero spending on gasoline. To evaluate if these cases could be quantitatively important for our estimates of MPC and ϵ, we estimate specifications (4) and (5) on the sample that excludes households with zero gasoline spending in 2013 (recall that the app data have a larger spike at zero than the counterpart in the CEX Interview Survey). Row (2) of Table 4 reports MPC

[18] In the summer time, many states require using a summer blend of gasoline which is more expensive than a typical winter blend.

estimates for this restricted sample at horizons $k = \{5,15,25\}$. We find that these estimates are very close to the baseline reported in row (1).

To address the second concern about cash spending, we note that cash spending only shows up in the dependent variable, generating a positive correlation that will cause us to *underestimate* the MPC. In a robustness check, we exclude ATM and other cash withdrawals from the dependent variable. We find (row 3) that both the MPC and elasticity of demand estimated on these modified data are nearly identical to the baseline estimates. This finding is consistent with the intensity of using cash as means of payment being similar for gasoline and non-gasoline spending.

For the third concern relating to expected changes in gasoline prices, we turn to data from the futures market. In particular, we use changes in one-month-ahead futures for spot prices at New York Harbor (relative to last week's prediction for the month ahead) instead of the change in gasoline prices since last week. Specifically, let F_t^h denote the futures price at time t for month $t+h$. Then, in lieu of $\Delta_k \log P_t$ in our baseline specification (4), we instead use $\Delta_k \log \mathcal{F}_t \equiv \log F_t^1 - \log F_{t-k}^1$ for $k \in \{1, \dots, 25\}$. While the focus on one-month change is arguably justified given approximate random walk in gasoline prices, we also try the average change in the yield curves for gasoline prices over longer horizons (two years) to have a measure of changes in gasoline prices that are perceived as persistent: $\Delta_k \widetilde{\log \mathcal{F}}_t \equiv \frac{1}{24}\sum_{h=1}^{24}\left(\log F_t^h - \log F_{t-k}^h\right)$. In either one-month change (row 4 of Table 4) or average change over two years (row 5), the results are very similar to our baseline.

C. Comparison with MPC using CEX

To appreciate the significance of using high-quality transaction-level data for estimating the sensitivity of consumers to income and price shocks, we also estimate the sensitivity using conventional, survey-based data sources such as the Consumer Expenditure Survey (CEX). This survey provides comprehensive estimates of household consumption across all goods in the household's consumption basket and is the most commonly used household consumption survey. In this exercise, we focus on the interview component of the survey which allows us to mimic the econometric analysis of the app data.

In this survey, households are interviewed for 5 consecutive quarters and asked about their spending over the previous quarter. Note that the quarters are not calendar quarters; instead, households enter the survey in different months and are asked about their spending over the previous three months. The BLS only makes available the data from the last 4 interviews; therefore we have a one-year panel of consumption data for a household. Given the panel design of the CEX Interview Survey, we can replicate aspects of our research design described above. Specifically, we calculate the ratio of gasoline spending to non-gasoline spending in the first interview. We then estimate the MPC in a similar regression over the next three quarters for households in the panel.[19] For this specification, we use BLS urban gasoline prices which provide a consistent series over this time period (see note for Table 1).

In the first row of Table 5, we estimate our baseline specification for the app data at the quarterly frequency: the estimates are slightly different from the estimates based on the weekly frequency, though much less precise. The standard errors are so large that we cannot reject the null of equality of the estimates over time or across frequencies.

Panel B of Table 5 presents estimates based on the CEX. To maximize the precision of CEX estimates, we apply our approach to the CEX data covering 1980-2014. The point estimates indicate that non-gasoline spending *declines* in response to decreased gasoline prices. Standard errors are so large that we cannot reject the null of no response. The estimated elasticity of demand for gasoline is approximately -0.4, which is a double of the estimates based on the app data and is similar to some of the previous CEX-based estimates (e.g., Nicol, 2003).

One should be concerned that the underlying variation of gasoline prices is potentially different across datasets. The dramatic decline in gasoline prices in 2014-2015 was largely determined by supply-side and foreign-demand factors, but it is less clear that one may be equally confident about the dominance of this source of variation over a longer sample period. Indeed, Barsky and Kilian (2004) and others argue that oil prices have often been demand-driven in the past. In this case, one may find wrong-signed or a non-existent relationship between gasoline prices and non-gasoline spending. To address this identification challenge, we focus on instances when changes in oil prices were arguably determined by supply-side factors.

[19] Our build of the CEX data follows Coibion et al. (2012).

Specifically, we follow Hamilton (2009, 2011) and consider several episodes with large declines in oil prices: (i) the 1986 decline in oil prices (1985-1987 period); and (ii) the 1990-1991 rise and fall in oil prices (1989-1992 period). Estimated MPCs and elasticities for each episode are reported in rows (3) and (4). Only the 1986 episode generates positive MPCs but the standard errors continue to be too high to reject the null of no response. The other episode continues to yield negative MPCs with large standard errors.

Note that in estimates from the app in row 1 we continue to use complete histories of consumer spending over 2013-2016 while the CEX tracks households only for four quarters. To assess the importance of having a long spending series at the consumer level, we "modify" the app data to bring it even closer to the CEX data. Specifically, for every month of our sample, we randomly draw a cohort of app users and track this cohort for only four consecutive quarters, thus mimicking the data structure of the CEX. Then, for a given cohort, we use the first quarter of the data to calculate s_i and use the remainder of the data to estimate ϵ and MPC. Results are reported in row 2 of Table 5. Generally, patterns observed in row 1 are amplified in row 2. In particular, the elasticity of demand for gasoline is even lower at shorter horizons and even greater at the longer horizons. In a similar spirit, the estimated MPC increases more strongly in the horizon when we track consumers for only four quarters relative to the complete 2013-2016 coverage.

In summary, the CEX-based point estimates are extremely imprecise. The data are inherently noisy. Moreover, when limited to sample periods that have credibly exogenous variation in gasoline prices, the sample sizes are far too small to make precise inferences. Furthermore, these estimates do not appear to be particularly robust. These results are consistent with a variety of limitations of the CEX data such as small sample size, recall bias, and under-representation of high-income households. These results also illustrate advantages of using high-frequency (weekly) data relative to low-frequency (quarterly) data for estimating sensitivity of consumer spending to gasoline price shocks. The app's comprehensive, high frequency data, combined with a natural experiment—the collapse of oil and gasoline prices in 2014—help us resolve these issues and obtain precise, stable estimates of MPC and elasticity of demand for gasoline.

D. Heterogeneity in Responses

Macroeconomic theory predicts that the responses of consumers to changes in income (or prices) could be heterogeneous with important implications for macroeconomic dynamics and policy. For example, Kaplan and Violante (2014) present a theoretical framework where consumers with liquidity constraints ("hand-to-mouth" consumers, HtM) should exhibit a larger MPC to *transitory*, *anticipated* income shocks than consumers for whom these constraints are not binding (non-HtM consumers). Kaplan and Violante (2014) document empirical evidence consistent with these predictions and quantify the contribution of consumer heterogeneity in terms of liquidity holdings for the 2001 Bush tax rebate. In a similar spirit, Mian and Sufi (2014), McKay, Nakamura and Steinsson (2016), and many others document that consumers' liquidity and balance sheets can play a key role for aggregate outcomes.

The conventional focus in this literature is the consumption response to transitory, anticipated income shocks because the behavior of HtM and non-HtM consumers should be particularly different in this case. First, HtM consumers spend an income shock when it is realized rather than when it is announced, while non-HtM consumers respond to the announcement and exhibit no change in spending at the time the shock is realized. Second, the MPC of non-HtM consumers should be small (this group smooths consumption by saving a big fraction of the income shock), while the MPC of HtM consumers should be large (the income shock relaxes a spending constraint for these consumers).

This sharp difference in the responses hinges on the temporary, anticipated nature of the shock. For other shocks, the responses may be alike across HtM and non-HtM consumers. For example, when the shock is permanent and unanticipated, HtM and non-HtM consumers should behave in the same way (Mankiw and Shapiro 1985): both groups should have $MPC = 1$ at the time of the shock. Intuitively, non-HtM consumers have $MPC = 1$ because their lifetime resources change permanently and, accordingly, these consumers adjust their consumption by the size of the shock when the shock happens. HtM consumers have $MPC = 1$ because they are in a "corner solution" and would like to spend away every dollar they receive in additional income the moment they receive it. Thus, macroeconomic theory predicts that, in this case, the MPC should be similar across HtM and non-HtM consumers and that the MPC should be close to one.

Given that the shock to the price of gasoline in our analysis was unanticipated and perceived as permanent, we focus this section on testing these two predictions.

For these tests one needs to identify HtM and non-HtM consumers. This seemingly straightforward exercise has proved to be a challenge in applied work due to a number of data limitations, which have made researchers use proxies for liquidity constraints. As a result, estimated MPCs should be interpreted with caution and important caveats. For example, Kaplan and Violante (2014) argue that identification of "hand-to-mouth" consumers requires information on consumers' liquidity holdings *just before* they receive pay checks.[20] Because the Survey of Consumer Finances (SCF), the dataset used in Kaplan and Violante (2014), reports *average* balances for a household as well as *average* monthly income, Kaplan and Violante are forced to make assumptions about payroll frequency (also not reported in the SCF) and behavior of account balances (e.g., constant flow of spending). Given heterogeneity in payment cycles (i.e., weekly, biweekly, monthly) and spending patterns across consumers, this procedure can mix hand-to-mouth (HtM) and non-hand-to-mouth (non-HtM) consumers. As a result, MPC estimated in Kaplan and Violante (2014) could provide a lower bound for the true MPC.

In contrast, the app data allow us to take Kaplan and Violante (2014)'s definition literally. We identify the exact day of a consumer's payroll income (if any), and examine bank account and credit card balances of the consumer the day before this payment arrives. A consumer is classified as HtM in a given month if, in the previous month, the consumer has virtually no liquid assets (less than $100 in the consumer's checking or savings accounts net of credit card debt), or the consumer is in debt (the sum of the consumers' liquid assets and available balance on credit cards is negative) and is within $100 of the consumer's credit card limits. Denote the dummy variable identifying hand-to-mouth consumers at this frequency with D_{it}^{*}. We find that, in the app data, roughly 20% of consumers are HtM, which is in the lower end of the range reported in Kaplan and Violante (2014).[21]

[20] Intuitively, hand-to-mouth consumers do not carry liquid assets from period to period. Hence, just before receiving a pay check (an injection of liquidity), a hand-to-mouth consumer should have zero liquid wealth.

[21] While the app data are close to ideal for identification of hand-to-month (i.e., low liquidity) consumers, the app data are not suitable for further disaggregation of consumers into wealthy hand-to-mouth and poor hand-to-mouth because the app does not collect information on consumer durables (e.g., vehicles), housing and other illiquid assets which are not backed by corresponding loans and mortgages.

To allow for heterogeneity in the MPC by liquidity, we add interaction terms to the baseline specification (4)-(5):

$$\Delta_{\mathrm{k}} \log C_{it} = \beta_1 \times s_i \times \Delta_{\mathrm{k}} \log P_t + \beta_2 \times s_i^{gas} \times \Delta_{\mathrm{k}} \log P_t \times D_{it} + \mu_0 \times D_{it} + \mu_1 \times s_i \times D_{it} + \psi_t + \omega_t \times D_{it} + \varepsilon_{it} \quad (6)$$

$$\Delta_{\mathrm{k}} \log PQ_{it} = \delta_1 \times \Delta_{\mathrm{k}} \log P_t + \delta_2 \times \Delta_{\mathrm{k}} \log P_t \times D_{it} + \xi \times D_{it} + u_{it} \quad (7)$$

where D_{it} is a variable measuring the presence/intensity of liquidity constraints identifying HtM consumers, and $\omega_t \times D_{it}$ is the time fixed effect specific to HtM consumers.

We have several options for D_{it}. One could use a dummy variable equal to one if a consumer is liquidity constrained in period $t-k-1$ (recall that Δ_{k} operator calculates the growth rate between periods $t-k$ and t). We denote this "lagged" measure of HtM with $\widetilde{D}_{it} \equiv D^*_{i,t-k-1}$. Alternatively, because liquidity constraints may be short-lived, one may want to use measures that are calculated over a longer horizon to identify "serial" HtM consumers. To this end, we construct three measures on the 2013 sample which are not used in the estimation of MPC and ϵ. Specifically, for each month of data available for consumer i in 2013, we use three metrics to classify consumers as hand-to-mouth or not. We consider the average value of D^*_{it} (this continuous variable provides a sense of frequency of liquidity constraints; we denote this measure with $\overline{D}_{i,2013}$), the modal value of $D^*_{i,t}$ (most frequent value;[22] we denote this measure with $\breve{D}_{i,2013}$), or the minimum value of $D^*_{i,2013}$ during the 2013 part of the sample. The latter measure, which we denote with $\widehat{D}_{i,2013}$, is equal to one only if a consumer is identified as HtM in *every* month in 2013.

Irrespective of which measure we use, we find in results reported in Table 6 that estimated MPCs are very similar for HtM and non-HtM consumers. Although the point estimates for HtM consumers tend to be larger at short horizons (e.g., 5 weeks), we cannot reject the null of equal MPCs across the groups for any horizon or definition of HtM status. Furthermore, we cannot reject that estimated MPCs are equal to one. These results are consistent with theoretical predictions and thus lend more credibility to our baseline estimate of MPC equal to one.

[22] We classify a household as HtM if there is a tie.

VI. Conclusion

How consumers respond to changes in gasoline prices is a central question for policymakers and researchers. We use big data from a personal financial management service to examine the dynamics of consumer spending during the 2014-2015 period when gasoline prices crashed by 50 percent. Given low elasticity of demand for gasoline, this major price reduction generated large windfall savings for consumers equal to approximately 2 percent of total consumer spending. (Average total household spending in 2014 was $53,495 in total, while the average household spending on gasoline was $2,468.)

We document that the marginal propensity to consume (*MPC*) out of these savings is approximately one, which is consistent with the predictions of the permanent income hypothesis, given that the change in gasoline prices was unexpected and permanent. The effects on the aggregate economy depend on a number of factors that are beyond the scope of this paper. Yet it is clear that a persistent increase of spending of 2 percent would have noticeable effects on aggregate economic performance.

We also show why previous attempts to estimate the MPC out of gasoline savings led to lower and/or more imprecise estimates due to data limitations (e.g., low frequency of data, incomplete coverage of consumer spending, short panel) in earlier studies. Our analysis highlights enormous potential of big data for enhancing national economic statistics, as well as estimates of key, policy-relevant macroeconomic parameters.

References

Anderson, Soren T., Ryan Kellogg and James M. Sallee. 2012. “What do consumers believe about future gasoline prices?” *Journal of Environmental Economics and Management.* 66(3): 383-403.

Baffes, John, M. Ayhan Kose, Franziska Ohnsorge, and Marc Stocker, 2015. "The Great Plunge in Oil Prices: Causes, Consequences, and Policy Responses," World Bank Group Policy Research Note #15/01.

Barsky, Robert B., and Lutz Kilian. 2004. "Oil and the Macroeconomy Since The 1970s," *Journal of Economic Perspectives* 18(4), 115-134.

Binder, Carola. 2015. “Measuring Uncertainty Based on Rounding: New Method and Application to Inflation Expectations.” Available at SSRN: http://ssrn.com/abstract=2576439.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2012. “Innocent Bystanders? Monetary Policy and Inequality in the U.S.” NBER Working Paper No. 18170.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review* 105, 2644–2678

Committee on National Statistics. 2013. *Measuring What We Spend: Toward a New Consumer Expenditure Survey*, Don A. Dillman and Carol C. House, Editors. Washington, D.C.: National Academies Press.

Edelstein, P., and Kilian, L. (2009). “How sensitive are consumer expenditures to retail energy prices?” *Journal of Monetary Economics*, 56(6), 766–779.

Energy Information Administration (EIA). 2014. “Gasoline prices tend to have little effect on demand for car travel.” Last accessed June 2, 2016. Available at: http://www.eia.gov/todayinenergy/detail.cfm?id=19191

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman and Steven Tadelis. 2016. "How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data." Working Paper.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman and Steven Tadelis. 2014. "Harnessing Naturally Occurring Data to Measure the Response of Spending to Income," *Science*, 345(6193): 212-215.

Gicheva, Dora, Justine Hastings and Sofia Villas-Boas. 2010. "Investigating Income Effects in Scanner Data: Do Gasoline Prices Affect Grocery Purchases?" *American Economic Review: Papers & Proceedings* 100: 480-484.

Griliches, Zvi, and Jerry A. Hausman. 1986. "Errors in variables in panel data," *Journal of Econometrics* 31(1), 93-118.

Hamilton, James D. 1983. "Oil and the Macroeconomy since World War II." *Journal of Political Economy*, 91(2): 228-248.

Hamilton, James D. 2009. "Causes and Consequences of the Oil Shock of 2007-08." *Brookings Papers on Economic Activity*. Spring 2009. 215-283.

Hamilton, James D. 2011. "Historical Oil Shocks," in *Routledge Handbook of Major Events in Economic History*, pp. 239-265, edited by Randall E. Parker and Robert Whaples, New York: Routledge Taylor and Francis Group, 2013.

Hamilton, James. 2008. "Oil and the Macroeconomy," in *New Palgrave Dictionary of Economics*, 2nd edition, edited by Steven Durlauf and Lawrence Blume, Palgrave McMillan Ltd.

Hastings, Justine S. and Jesse M. Shapiro. 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *The Quarterly Hournal of Economics*. 128(4): 1449-1498.

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): 93-114.

J.P. Morgan Chase Institute. 2015. *How Falling Gas Prices Fuel the Consumer: Evidence from 25 Million People.* Available at https://www.jpmorganchase.com/corporate/institute/document/jpmc-institute-gas-report.pdf.

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes," *Annual Review of Economics* 2(1), 479-506.

Johnson, David S., Jonathan A. Parker and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review,* 96(5): 1589-1610.

Kaplan, Greg and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica.* 82(4): 1199-1239.

Kilian, Lutz. 2008. "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature,* 46(4): 871-909.

Kilian, Lutz. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review* 99(3), 1053-69.

Levin, Laurence, Matthew S. Lewis and Frank A. Wolak. 2016. "High Frequency Evidence On The Demand For Gasoline," NBER Working Paper no. 22345

Mankiw, N. Gregory, and Matthew D. Shapiro. 1985. "Trends, Random Walks, and Tests of the Permanent Income Hypothesis," *Journal of Monetary Economics* 16, 165-174.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, 2016. "The Power of Forward Guidance Revisited," *American Economic Review*. Forthcoming.

Mian, Atif, and Amir Suri. 2014. *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening*. Princeton University Press.

National Association of Convenience Stores (NACS). 2015. *Retail Fuels Report*. Available at http://www.nacsonline.com/YourBusiness/FuelsReports/2015/Documents/2015-NACS-Fuels-Report_full.pdf .

Nicol, C.J. 2003. "Elasticities of demand for gasoline in Canada and the United States." *Energy Economics*. 25: 201-214.

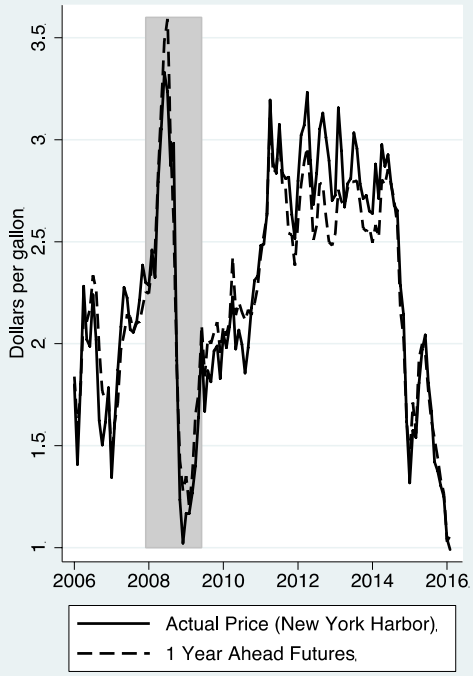
Puller, Steven and Lorna A. Greening. 1999. “Household adjustment to gasoline price change: an analysis using 9 years of US survey data.” *Energy Economics* 21: 37-52.

Sabelhaus, John, David Johnson, Stephen Ash, David Swanson, Thesia I. Garner, John Greenlees, and Steve Henderson. 2015. “Is the Consumer Expenditure Survey Representative by Income?” in Carroll, Crossley, and Sabelhaus, eds., *Improving the Measurement of Consumer Expenditures*, University of Chicago Press.

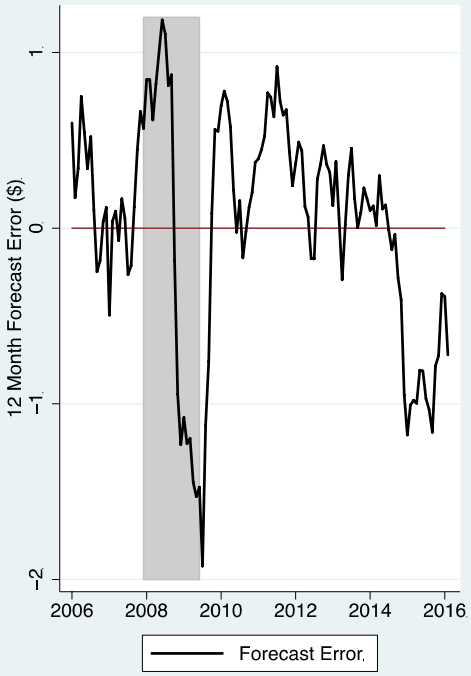
Shapiro, Matthew D. and Joel Slemrod. 2003. "Consumer Response to Tax Rebates." *American Economic Review*, 93(1): 381-396.

Figure 1. Gasoline prices and expectations

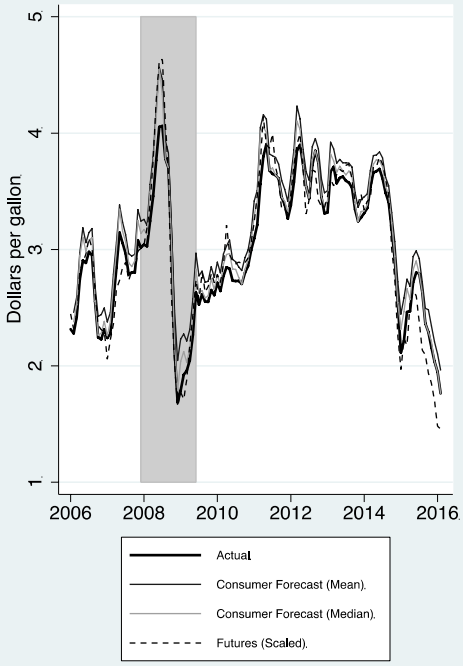
Panel A. Current and Futures Price



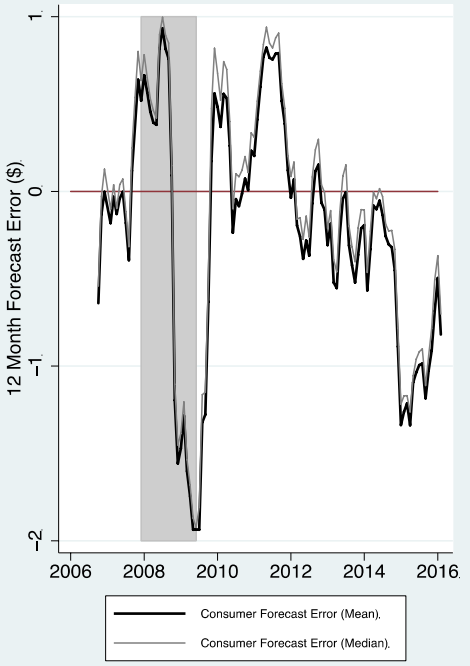
Panel B. Futures' Forecast Error



Panel C. Consumer/SPF Expectations

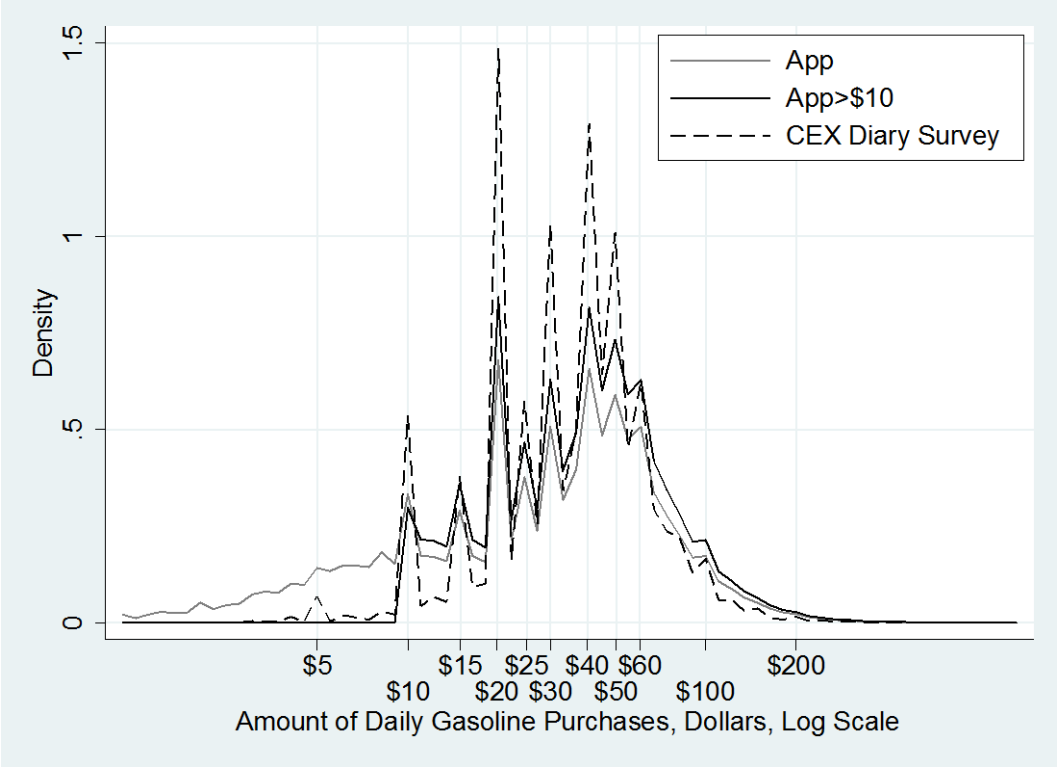


Panel D. Consumers' Forecast Error



Notes: Panel A shows the New York Harbor spot price and the 1 year ahead future price. Panel B shows the 1 year ahead forecast error, defined as the difference between the realization of the spot price and the forecast 1 year earlier. Panel C shows the gasoline price, and the weighted mean and median expectations from the Michigan Surveys of Consumers. See https://data.sca.isr.umich.edu/sda-public/cgi-bin/hsda?harcsda+sca. In the survey, households are asked, "About how many cents per gallon do you think gasoline prices will (increase/decrease) during the next twelve months compared to now?" We add the household response to this question to the current gasoline price. We also plot the futures used in panel A, scaled by the average difference between the spot price and retail price over the period (dotted line). Panel D shows retail gasoline prices and the consumer forecast made 12 months earlier. Shaded area is the Great Recession.

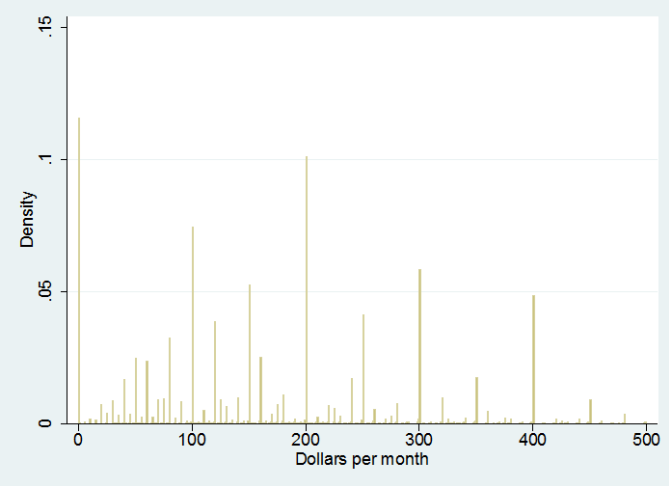
Figure 2. Distribution of log gasoline spending: CEX Diary versus App



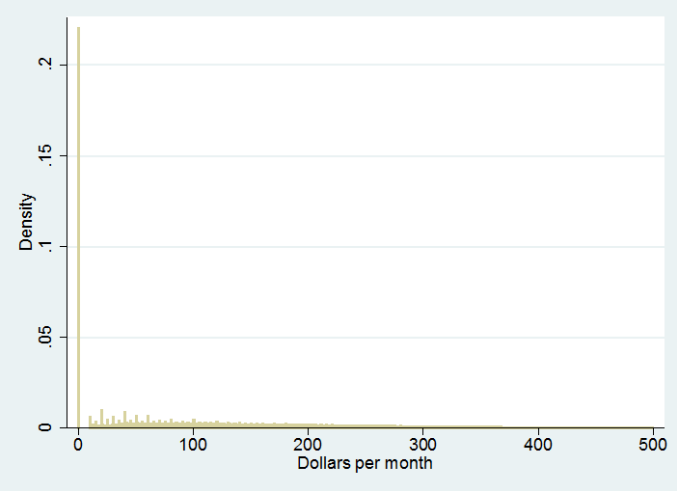
Notes: the figure shows the distribution of daily log spending on gasoline in the Diary segment of the Consumer Expenditure Survey (CEX) and in the app data. Gasoline spending in the app data is identified using machine learning (ML). *App* includes all transactions that ML identifies as purchases of gasoline. *App>$10* includes transactions that ML identifies as purchases of gasoline and that are greater than $10. See text for further details.

Figure 3. Reported gasoline spending (monthly)

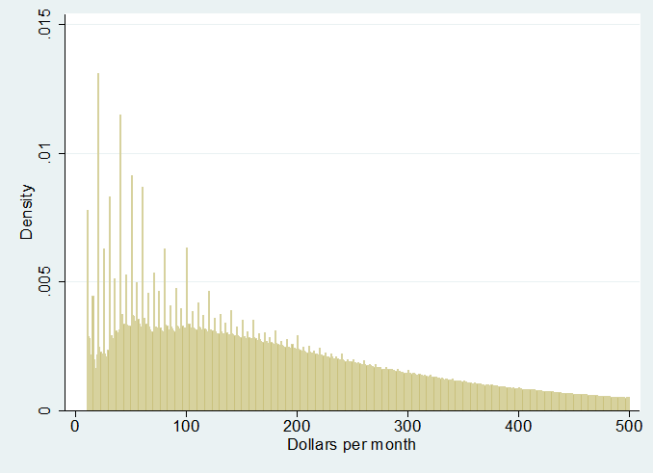
Panel A. CEX Interview Survey



Panel B. App data



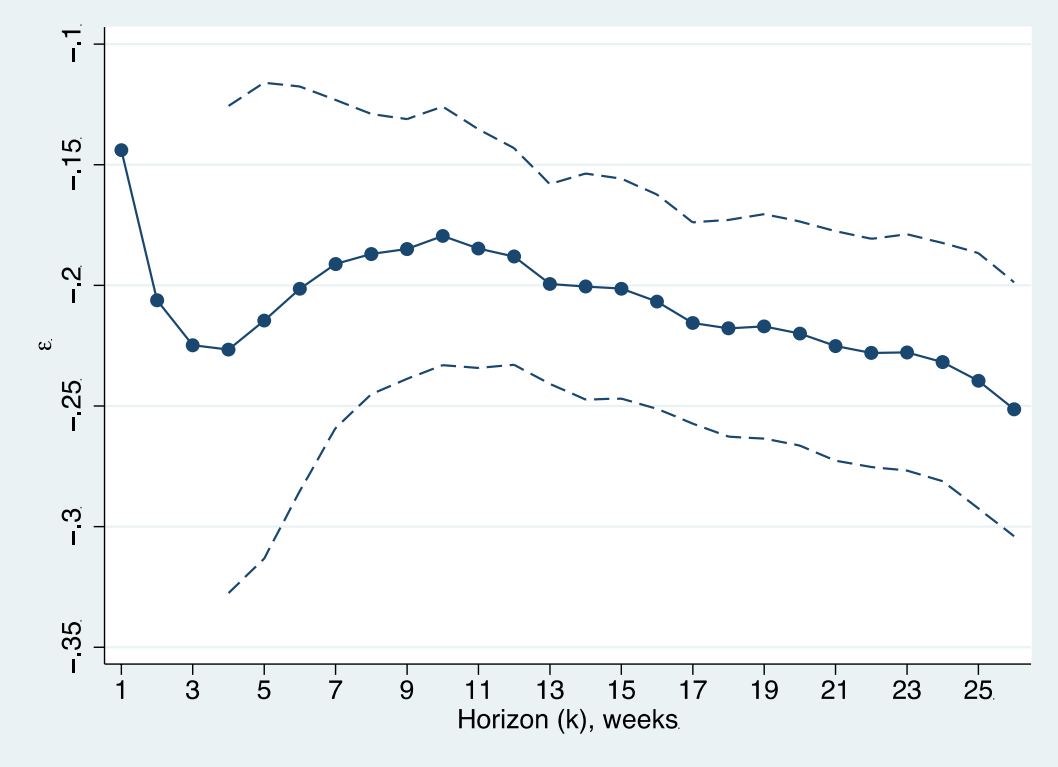
Panel C. App data excluding zero spending



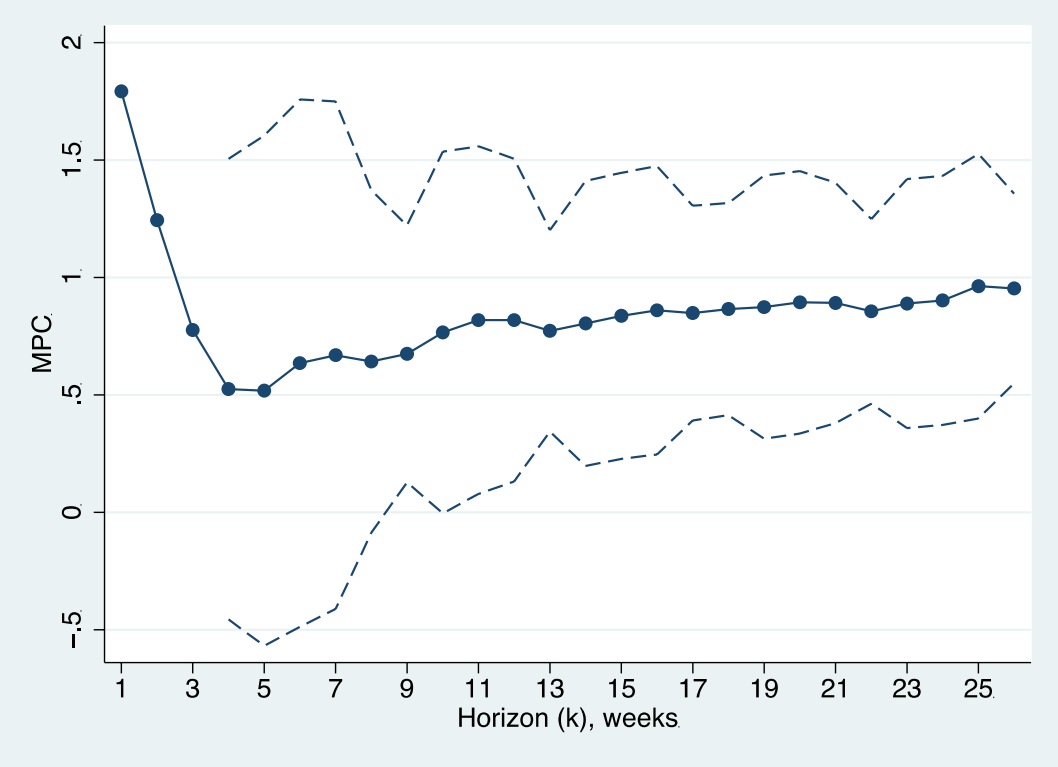
Notes: the figure reports monthly spending on gasoline in the Interview segment of the Consumer Expenditure Survey (CEX) and in the app data. The horizontal axis is in dollars. The size of the bin in is set to $1 in all panels.

Figure 4. Dynamic response to a change in gasoline price

Panel A. Elasticity of demand for gasoline, ϵ



Panel B. MPC



Notes: the figure reports estimates of elasticity of demand for gasoline ϵ (Panel A) and marginal propensity to consume (MPC) (Panel B) based on specifications (4) and (5). Standard errors clustered at both the week and user level. First three standard errors are omitted for readibility of the graph. See text for further details.

Table 1. Largest monthly changes in oil and gasoline prices

Largest Decreases			Largest Increases		
Date	Percent Change		Date	Percent Change	
	Oil	Gas		Oil	Gas
1986:2	-33	-6	1974:1	135	
2008:12	-28	-21	1990:8	47	10
2008:10	-26	-14	1986:8	30	-5
2008:11	-25	-32	1948:1	24	
2014:12	-22	-11	1990:9	23	9
2015:1	-20	-18	2009:3	23	1

Notes: Table shows the month-to-month percent change in West Texas Intermediate spot oil prices (FRED series OILPRICE and MCOILWTICO) and the corresponding change in average monthly regular gasoline prices, when available, from January 1946 – February 2016. For gasoline prices, the table use the BLS U.S. city average (BLS series APU000074714), since it is available further back in time than other available gasoline price data.

Table 2. Comparison of spending in the CEX and app data, 2013

Frequency and type of spending	Moment			
	Mean	St. Dev.	Median	Interquartile range
Panel A. Biweekly				
Spending on gas, dollars				
CEX Diary Survey	84.72	84.72	83.42	65.00
App	81.68	203.75	30.99	79.66
Spending on non-gasoline items, dollars				
CEX Diary Survey	1,283.36	1,470.93	790.56	1,380.66
App	1,468.65	3,617.73	557.38	975.32
Ratio of gasoline to non-gasoline spending				
CEX Diary Survey	0.15	0.25	0.06	0.14
App	0.12	0.25	0.04	0.12
Panel B. Quarterly				
Spending on gasoline, dollars				
CEX Interview Survey	646.63	530.87	540.00	630.00
App	627.94	588.24	475.33	660.18
Spending on non-gasoline items, dollars				
CEX Interview Survey	10,143.78	8,141.67	7,728.70	7,406.49
App	11,264.85	11,391.42	8,392.24	8,605.46
Ratio of gasoline to non-gasoline spending				
CEX Interview Survey	0.08	0.07	0.06	0.08
App	0.07	0.07	0.05	0.07

Notes: Means and standard deviation are from the distribution winsorized at the 1% level. The variables from the CEX use population sample weights. For Panel A, the ratio for a consumer/household is calculated as average value of the sum all gasoline spending during a biweekly period in 2013 divided by total non-gasoline spending in the corresponding biweekly period in 2013. For Panel B, the ratio for a consumer/household is calculated as the sum of all gasoline spending in a quarter, divided by total non-gasoline spending in that quarter.

Table 3. Estimated elasticity of demand and MPC: Baseline and estimates for single financial providers

Accounts	Sample	Row	Elasticity of demand for gasoline, ϵ			MPC		
			Horizon (weeks)			Horizon (weeks)		
			5	15	25	5	15	25
			(4)	(5)	(6)	(7)	(8)	(9)
Baseline	All	1	-0.215	-0.201	-0.240	0.518	0.837	0.963
			(0.050)	(0.023)	(0.027)	(0.554)	(0.311)	(0.287)
Large Provider #1	Any Account	2	-0.256	-0.265	-0.314	0.125	0.586	0.515
			(0.048)	(0.024)	(0.029)	(0.835)	(0.482)	(0.387)
	Core	3	-0.290	-0.261	-0.293	0.199	0.666	0.661
			(0.053)	(0.025)	(0.027)	(0.928)	(0.526)	(0.427)
Large Provider #2	Any Account	4	-0.242	-0.260	-0.367	0.064	0.321	0.642
			(0.069)	(0.038)	(0.054)	(0.425)	(0.218)	(0.266)
	Core	5	-0.223	-0.237	-0.326	-0.057	0.429	0.616
			(0.065)	(0.034)	(0.051)	(0.487)	(0.306)	(0.321)
Large Provider #3	Any Account	6	-0.183	-0.179	-0.246	0.281	0.561	0.639
			(0.051)	(0.034)	(0.045)	(0.736)	(0.431)	(0.371)
	Core	7	-0.161	-0.170	-0.241	0.204	0.575	0.734
			(0.052)	(0.035)	(0.046)	(0.595)	(0.366)	(0.349)

Notes: The table reports estimates of elasticity of demand for gasoline ϵ and marginal propensity to consume (MPC) based on specifications (4) and (5) for horizons 5, 15, and 25 weeks. Row 1 presents the baseline estimates based on the full sample. In the rest of the table, the sample is restricted to a single provider indicated in the left column. In other words, we restrict the sample to accounts only at a specific provider so that we can mimic the data observed by a single provider. In rows (2), (4) and (6), the table report estimates for the case when we use any account of a provider. In rows (3), (5) and (7), the table report estimates based on "core accounts"; that is, to be part of the estimation sample, a user has to have at least one checking and one credit-card account with a given provider and have at least one transaction per month on each account. In all specifications, robust standard errors are clustered at both the consumer and week level. See text for further details.

Table 4. Robustness of MPC estimate

Sample	Row	Elasticity of demand for gasoline, ϵ			MPC		
		Horizon (weeks)			Horizon (weeks)		
		5	15	25	5	15	25
		(1)	(2)	(3)	(4)	(5)	(6)
Baseline	1	-0.215	-0.201	-0.240	0.518	0.837	0.963
		(0.050)	(0.023)	(0.027)	(0.554)	(0.311)	(0.287)
Exclude zero gasoline spending in 2013	2	-0.214	-0.201	-0.239	0.541	0.865	1.021
		(0.050)	(0.023)	(0.027)	(0.563)	(0.312)	(0.293)
Exclude ATM withdrawals	3	-0.213	-0.198	-0.235	0.401	0.870	1.062
		(0.050)	(0.023)	(0.027)	(0.754)	(0.443)	(0.391)
Change in one-month-ahead gasoline futures	4	-	-	-	0.204	0.855	0.907
					(0.440)	(0.234)	(0.216)
Average change in the yield curve of gasoline futures	5	-	-	-	0.388	1.033	1.110
					(0.531)	(0.281)	(0.259)

Notes: the table reports estimates of elasticity of demand for gasoline ϵ and marginal propensity to consume (MPC) based on specifications (4) and (5) for horizons 5, 15, and 25 weeks. Row 1 presents the baseline estimates based on the full sample. The estimation sample in row 2 excludes consumers with zero spending on gasoline in 2013. In row 3, we exclude ATM withdrawals and other cash withdrawals in calculation of the growth rate of non-gasoline spending. In rows 4 and 5, we replace actual changes in gasoline prices with changes in futures prices of gasoline in specification (4); specification (5) is estimated as in the baseline, so ϵ is the same as in row 1. Specifically, let F_t^h denote the futures price made at time t for period $t+h$. Then, in lieu of $\Delta_k \log P_t$ in our baseline specifications (4), we instead use $\Delta_k \log \mathcal{F}_t \equiv \log F_t^1 - \log F_{t-k}^1$ for $k \in \{1, \ldots, 25\}$ in row 4 and the average change in the yield curves for gasoline prices over longer horizons (two years) $\Delta_k \widetilde{\log \mathcal{F}}_t \equiv \frac{1}{24}\sum_{h=1}^{24}\left(\log F_t^h - \log F_{t-k}^h\right)$ in row 5. In all specifications, robust standard errors are clustered at both the consumer and week level. See text for further details.

Table 5. Elasticity of demand for gasoline and MPC: Consumer Expenditure Survey (CEX) versus App

Data and Sample	Row	Elasticity of demand for gasoline, ϵ			MPC		
		Horizon (quarters)			Horizon (quarters)		
		1	2	3	1	2	3
		(1)	(2)	(3)	(4)	(5)	(6)
Panel A: App data (quarterly)							
Baseline	1	-0.084	-0.126	-0.383	1.000	1.050	2.320
		(0.118)	(0.110)	(0.157)	(0.478)	(0.627)	(0.797)
CEX sample design	2	0.005	-0.111	-0.432	0.732	1.574	4.942
		(0.073)	(0.066)	(0.057)	(0.548)	(0.287)	(1.017)
Panel B: CEX							
1980-2014	3	-0.429	-0.330	-0.352	-0.917	-0.526	-1.929
		(0.041)	(0.027)	(0.027)	(0.659)	(0.638)	(1.265)
1985-1987	4	-0.478	-0.396	-0.449	9.624	5.074	4.116
		(0.168)	(0.145)	(0.090)	(5.959)	(4.057)	(4.031)
1990-1992	5	-0.636	-0.562	-0.512	-6.420	-3.005	-4.917
		(0.136)	(0.142)	(0.106)	(4.775)	(4.516)	(6.450)

Notes: the table reports estimates of elasticity of demand for gasoline ϵ and marginal propensity to consume (MPC) based on specifications (4) and (5) for horizons 1, 2, and 3 quarters. The CEX estimates use the ratio of gasoline spending to non-gasoline spending calculated in the first interview, and exclude this period from estimation. For the baseline estimates in Row 1, we use the same 2013 ratio of gasoline spending to non-gasoline spending as in the baseline estimates, and aggregate the spending and gasoline prices to the quarterly level. In row 2, we replicate the CEX sampling scheme, randomly selecting a start month for a user and keeping only the data for the 12 month period that follows it (if a full 12 months of data follow). We similarly use the non-gasoline consumption calculated in the first quarter, and exclude this period from the estimation. In all specifications, robust standard errors are clustered at both the consumer and week level. See text for further details.

Table 6. MPC by liquidity status.

Measure of Hand-to-mouth consumers (HtM)	Elasticity of demand for gasoline, ϵ			MPC		
	Horizon (weeks)			Horizon (weeks)		
	5	15	25	5	15	25
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Lagged HtM						
Non-HtM	-0.191	-0.155	-0.179	0.569	0.894	0.660
	(0.057)	(0.025)	(0.026)	(0.749)	(0.438)	(0.356)
HtM	-0.239	-0.240	-0.293	0.547	1.044	0.896
	(0.058)	(0.030)	(0.031)	(0.829)	(0.472)	(0.471)
P-value (Non-HtM=HtM)	0.052	0.000	0.000	0.915	0.883	0.658
Panel B. Average HtM in 2013						
Non-HtM	-0.186	-0.164	-0.199	0.618	0.776	0.783
	(0.054)	(0.024)	(0.027)	(0.619)	(0.311)	(0.299)
HtM	-0.274	-0.317	-0.373	0.742	1.336	1.383
	(0.055)	(0.030)	(0.035)	(0.838)	(0.537)	(0.428)
P-value (Non-HtM=HtM)	0.007	0.000	0.000	0.932	0.073	0.077
Panel C. Modal HtM in 2013						
Non-HtM	-0.191	-0.172	-0.209	0.597	0.899	0.893
	(0.053)	(0.024)	(0.027)	(0.651)	(0.358)	(0.320)
HtM	-0.255	-0.287	-0.340	0.896	0.851	0.965
	(0.053)	(0.028)	(0.032)	(0.546)	(0.270)	(0.291)
P-value (Non-HtM=HtM)	0.012	0.000	0.000	0.579	0.410	0.641
Panel D. Extreme HtM in 2013						
Non-HtM	-0.195	-0.182	-0.221	0.613	0.924	0.933
	(0.053)	(0.024)	(0.028)	(0.639)	(0.355)	(0.319)
HtM	-0.284	-0.318	-0.365	1.029	0.912	1.064
	(0.056)	(0.029)	(0.033)	(0.571)	(0.297)	(0.310)
P-value (Non-HtM=HtM)	0.002	0.000	0.000	0.470	0.472	0.753

Notes: the table reports estimates of MPC and ϵ based on equations (6)-(7) over k periods, where k is shown in the top row of the table. s_i^{gas} is the ratio of gasoline spending to non-gasoline spending for 2013 for consumer *i*. The title of each panel indicates how the presence/intensity of liquidity constraints is measured. Denote the dummy variable identifying hand-to-mouth consumers for a given month with D_{it}^*. Panel A uses a dummy variable equal to one if a consumer is liquidity constrained in period $t-k-1$ (recall that Δ_k operator calculates the growth rate between periods $t-k$ and t), i.e. $\widetilde{D}_{it} \equiv D_{i,t-k-1}^*$. For other panels, we construct three measures on the 2013 sample which is not used in the estimation of MPC and ϵ: the average value of D_{it}^* (this continuous variable provides a sense of frequency of liquidity constraints; we denote this measure with $\overline{D}_{i,2013}$), the modal value of D_{it}^* (most frequent value; we denote this measure with $\breve{D}_{i,2013}$), or the minimum value of $D_{i,2013}^*$ during the 2013 part of the sample. The latter measure, which we denote with $\widehat{D}_{i,2013}$ and refer to as "extreme," is equal to one only if a consumer is identified as hand-to-mouth in *every* month in 2013. Robust standard errors are clustered by week and consumer and are reported in parentheses. *P-value (Non-HtM=HtM)* is the p-value for the test of HtM and non-HtM responses being equal. See text for further details.