The Self-Constrained Hand-to-Mouth

Michael Gelman*

December 11, 2018

Abstract

This paper examines the response of food expenditures to the receipt of regular paychecks using financial account data from a personal finance app. Like previous studies, I find that food expenditures increase during the pay-week. While the standard explanation for this result is temporary liquidity constraints, I argue individuals are unlikely to be constrained during the pay-week. Both the data and a buffer stock model of consumption show that consumption behavior is not affected by liquidity during the pay-week. This novel evidence implies that excess sensitivity is not explained by temporary liquidity constraints but rather by individual preferences.

JEL Classification: D12, E21

*Claremont McKenna College (mgelman@cmc.edu). This research project is carried out in cooperation with a financial aggregation and bill-paying computer and smartphone application (the app). The project is grateful to the executives and employees who have made this research possible. This project is supported by a grant from the Alfred P. Sloan Foundation with additional support from the Michigan node of the NSF-Census Research Network (NSF SES 1131500). I would like to thank Miles Kimball, John Leahy, Matthew Shapiro, and Melvin Stephens for valuable comments, suggestions, and support. I also thank Daphne Chang, Michael Gideon, Yuriy Gorodnichenko, Shachar Kariv, Minjoon Lee, Dhiren Patki, Dan Silverman, Steven Tadelis, and Fudong Zhang for helpful conversations and suggestions.

1 Introduction

Since Hall's seminal work on testing the Life-cycle/permanent-income hypothesis (LC-PIH) (Hall, 1978), many studies (see Jappelli and Pistaferri (2010) for an excellent survey) have documented the fact that consumption responds to the arrival of predictable income (excess sensitivity). Many of these studies show that the strength of the consumption response varies by some measure of liquidity constraints such as income, liquid wealth, age, or occupation.[1] These empirical results have led researchers to conclude that excess sensitivity is caused by temporary liquidity constraints.

This paper challenges this interpretation of excess sensitivity by arguing that individuals who receive regular paychecks are unlikely to be liquidity constrained during the week in which they receive their paycheck. This intuition is formalized by specifying a parsimonious buffer stock model of consumption with realistic paycheck dynamics. Model simulations show that in the week the paycheck is received, consumption behavior is unlikely to be affected by liquidity levels and so behavior is driven purely by preferences. By using a novel dataset on high frequency expenditure and liquid savings behavior, I show that indeed expenditure behavior on pay weeks is not affected by how much liquidity an individual holds. This simple buffer stock model can explain both patterns in the level of expenditures as well as the joint behavior of expenditure growth and liquidity levels. The main contribution of the paper is to show empirically that the relationship between expenditure growth and liquidity is consistent with a buffer stock model that includes realistic paycheck dynamics.

The idea that excess sensitivity is caused by preferences and not temporary liquidity constraints is not new. There are a few papers such as Laibson (1997) and Shapiro (2005) which argue that quasi-hyperbolic discounting can explain the high frequency responses to changes in income. While not focusing specifically on high frequency excess sensitivity, Parker (2017) argues that excess sensitivity must be driven by persistent rather than temporary factors. Similarly, Carroll et al. (2017) show that heterogeneity in preferences can explain heterogeneity in consumption behavior without having to appeal to liquidity constraints.[2] Testing whether persistent characteristics like preferences explain excess sensitivity is hard without high frequency data on consumption, income, and liquidity. Gelman et al. (2014), Kuchler (2015), and Olafsson and Pagel (2018) were the first studies to use high frequency personal finance app data to document excess sensitivity to regular paychecks.[3] Olafsson and Pagel (2018) is closely related to this paper in that it also claims that excess sensitivity cannot be explained

[1] For example Zeldes (1989a) and Stephens (2006).

[2] Although less directly related, Sahm, Shapiro and Slemrod (2010) also show that individual "types" rather than constraints can explain consumption behavior.

[3] Kueng (2018) uses similar personal finance website data to document excess sensitivity to a lower frequency income source (Alaska Permandent Fund).

by liquidity constraints. The main difference is that Olafsson and Pagel (2018) use theories from the corporate finance literature to measure liquidity constraints while this paper defines liquidity constraints in the context of a buffer stock model with realistic paycheck dynamics.[4] Furthermore, Olafsson and Pagel (2018) remain agnostic as to what can explain excess sensitivity while this paper shows that the empirical results are consistent with an appropriately modified buffer stock model. The upside of this result is that we can interpret excess sensitivity not as a failure of the LC-PIH, but as optimal behavior that reflects preferences. This paper is also related to Gelman (2017) which uses the same data set and also attempts to disentangle preferences and constraints. The main difference is that this paper uses high-frequency weekly data and focuses on the response to paychecks while Gelman (2017) uses the response to receiving a tax refund to estimate preference parameters.

The rest of the paper is organized as follows. Section 2 discusses the dataset and defines the main variables used in the analysis. Section 3 establishes facts regarding the high frequency expenditure response to receiving a paycheck. Section 4 lays out the theoretical framework I use to separately identify the effects of preferences and constraints at high frequency. Section 5 compares model simulations to the data to confirm that individuals do not exhibit signs of being constrained during the week the paycheck is received. Section 6 uses the receipt of tax refunds to further test the implications of the model and section 7 concludes.

2 High frequency account data

This section describes the data source, sample filters, variable definitions, and descriptive statistics.

2.1 Data source

This paper utilizes a novel dataset derived from de-identified transactions and account data, aggregated and normalized at the individual level. The data are captured in the course of business by a personal finance app.[5] More specifically, the app offers financial

[4] Both this paper and Olafsson and Pagel (2018) argue that liquidity constraints are hard to measure because constrained behavior is a function of preferences and income uncertainty. For example, for an individual with little uncertainty, holding few liquid assets may not be a sign of liquidity constraints. Similarly, an individual who faces high levels of uncertainty may still feel constrained even though they hold a large buffer of liquidity. Therefore, instead of using levels of liquidity like many previous studies, I test for liquidity constraints by estimating the relationship between consumption growth and liquidity inspired by the consumption euler equation.

[5] These data have previously been used to study the high-frequency responses of households to shocks such as the government shutdown (Gelman et al., 2015) and anticipated income, stratified by spending, income and liquidity (Gelman et al., 2014). Similar account data from other apps have been used in Baugh,

Table 1: App user demographics

Education	Not Completed College	Completed College	Completed Graduate School
ACS	66.62	24.02	9.36
App	70.42	23.76	5.83

Ages 25 and over. Sample size - ACS: 2,176,103 App: 28,057

Age	18-20	21-24	25-34	35-44	45-54	55-64	65+
ACS	5.85	7.28	17.44	17.24	18.78	16.00	17.41
App	0.59	5.26	37.85	30.06	15.00	7.76	3.48

Sample size - ACS: 2,436,714 App: 35,417

Gender	Male	Female
ACS	48.56	51.44
App	59.93	40.07

Sample size - ACS: 2,436,714 App: 59,072

Region	Northeast	Midwest	South	West
ACS	17.77	21.45	37.36	23.43
App	20.61	14.62	36.66	28.11

Sample size - ACS: 2,441,532 App: 63,745

Source: Gelman et al. (2014).

aggregation and bill-paying services. 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. Prior to analysis, the data are stripped of personally identifying information such as name, address, or account number. The data have scrambled identifiers to allow observations to be linked across time and accounts.

The paper's sample draws on the entire de-identified population of active users and data derived from their records from December 2012 until July 2016. For a subset of the data, we have made use of demographic information provided to the app by a third party. Table 1 compares the age, education, gender, and geographic distributions in the sample that matched with an email address to the distributions in the U.S. Census American Community Survey (ACS), representative of the U.S. population in 2012.

Figure 1 compares the income distribution in the app to total family income in the ACS. Users who use the app are on average higher income than individuals surveys in

Ben-David and Park (2014), Baker (2017), Baker and Yannelis (2017), Kuchler (2015), Ganong and Noel (2016), and Kueng (2018).

Figure 1: Income comparison

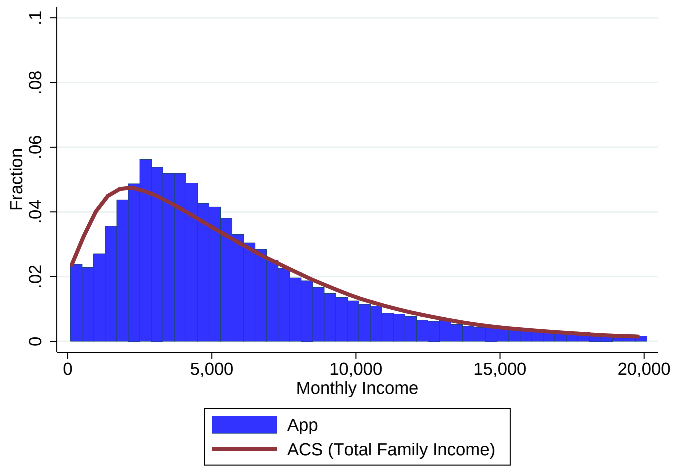

Source: Gelman et al. (2014).

the ACS.

In summary, the app is not perfectly representative of the US population, but it is heterogeneous, including large numbers of users of different ages, education, income, and geographic location.

2.2 Defining the sample

The sample is filtered on various characteristics to mitigate measurement error. I filter users based on length of panel, number of accounts, connectedness of accounts, regular paycheck status, and completeness of income data.

2.2.1 Defining account linkage

If all accounts that are used for receiving income and making expenditures are not observed, we may mistake mismeasurement for excess sensitivity. For example, an individual may have a checking account that is used to pay most bills and a credit card that it used when income is low. If credit card expenditures are not properly observed, it may look like expenditures are lower the week after a paycheck is received relative to the week in which the paycheck is received.

In order to identify linked accounts, I use a method that calculates how many credit card balance payments are also observed in a checking account. I define the variable *linked* as the ratio of the number of credit card balance payments observed in all checking accounts that matches a particular payment that originated from all credit card accounts. For example, a typical individual will pay their credit card bill once a month. If they existed in the data for the whole year, they will have 12 credit card balance payments. If 10 of those credit card payments can be linked to a checking account the variable $linked = \frac{10}{12} \approx 0.83$.

One drawback to this approach is that it requires individuals to have a credit card account. To ensure that those without credit cards are still likely to have linked accounts, I also condition on individuals who have three or more accounts.

2.2.2 Identifying regular paychecks

In order to identify regular paychecks, I start by using keywords that are commonly associated with these transactions.[6] I condition on five statistics to ensure that these transactions represent regular paychecks.

1. Number of paychecks ≥ 5
2. Median paycheck amount $>$ \$200
3. Median absolute deviation of days between paychecks is ≤ 5
4. Coefficient of variation of the paycheck amount ≤ 1
5. Weekly or bi-weekly payroll schedule

For bi-weekly paychecks there are two possible payment schedules. I define these bi-weekly payroll patterns by "odd" or "even." Although this is an arbitrary definition, the main role of this variable is to create two mutually exclusive groups. My definition of week starts on Thursday and week 0 is December 6, 2012. Therefore "even" weeks are the weeks starting Dec 20, 2012, Jan 3, 2012, etc. I define a payroll schedule for a particular individual as "even" if 90% of paychecks are received on an even week. The odd week schedule is defined similarly.

2.3 Variable definitions

Most survey data sets such as the consumer expenditure survey (CE), panel study of income dynamics (PSID), and survey of consumer finances (SCF) are created with the explicit goal of facilitating academic research. The data set used in this study is naturally occurring and was not explicitly designed for use in academic studies. Constructing variables in this data set to match our models is not necessarily a trivial exercise. In order to study the expenditure response to receiving a paycheck, the main variables I utilize are expenditure, paycheck income, and liquid assets.

[6] Keywords used to identify paychecks are "dir dep","dirde p","salary","treas xxx fed","fed sal","payroll","ayroll","payrll","payrl","payrol","pr payment","adp","dfas-cleveland","dfas-in" and DON'T include the keywords "ing direct","refund","direct deposit advance","dir dep adv."

2.3.1 Expenditures

The empirical analysis will focus on food expenditures away from home because they can be considered non-durable and non-storable at high frequency. These features allow my expenditure measure to closely track the concept of consumption used in the theoretical model that I will later introduce.

The raw data consists of individual transactions with characteristics such as amount, transaction type (debit or credit), and transaction description. While the type of expenditure (food, non-food) is not directly observed, I use a machine learning (ML) algorithm (see Appendix A.1 for more details) to aid in categorization. The goal of the ML algorithm is to provide a mapping from transaction descriptions to expenditure categories. For example, any transaction with the keyword "McDonalds" should map into "Fast Food." A subset of these categories are then combined to create the expenditure variable.

The finest level of categorization is derived from merchant category codes (MCCs) which are directly observable in two of the account providers in the data. MCCs are four digit codes used by credit card companies to classify expenditures and are also recognized by the U.S. Internal Revenue Service for tax reporting purposes. The ML algorithm works by using a subset of the data where the truth is known in order to create a mapping from transaction description to MCCs.

After training the ML algorithm on the data where the truth is known, the algorithm is then applied to the rest of the data set. I then define expenditure as expenditures on fast food and restaurants.

2.3.2 Cash on hand and liquid assets

Cash on hand is defined as $X_{it} = A_{it-1} + Y_{it}$ where A_{it-1} represents liquid balances for individual i in the previous period and Y_{it} represents income received in the current period.

Liquid balances (A) are defined as the sum of checking and saving account balances observed in the app. These balances are captured daily as the app takes a snapshot of the balance from each provider.

It is crucial to the analysis that I am able to observe high frequency (daily) observations of liquid balances. There are few datasets that contain high frequency measures of liquidity and as a result there are few studies that analyze the interaction between high frequency liquidity and expenditures.

3 The expenditure response to paycheck arrival

This section documents the expenditure response to the arrival of a bi-weekly paycheck. By using two different bi-weekly schedules, I show that the expenditure response seen in the data is due to the receipt of a paycheck and not confounded with other events such as first of the month effects.

3.1 Time series figures

When analyzing high frequency excess sensitivity, it's important to focus on non-durable expenditures to make sure expenditures line up with consumption as much as possible. As discussed in the previous section, I use fast food and restaurant expenditures to test excess sensitivity of expenditure. Figure 2 compares this expenditure measure to a comparable expenditures series from the Census Bureau.[7] Because the app data and the Census data are in different units, I plot the log difference relative to Jan 2013 on the y-axis. While the app data is more volatile than the Census data, they both exhibit a similar upward trend over the time period.

Figure 2: Monthly food expenditures

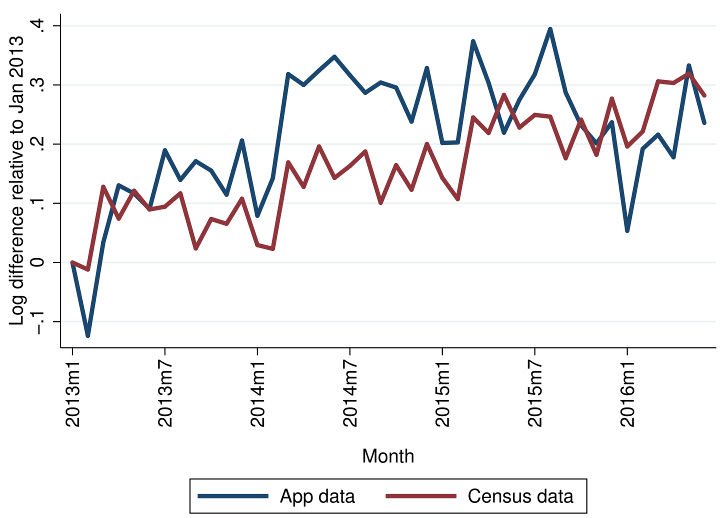

Using the high frequency nature of the data, Figure 3 plots weekly food expenditures for bi-weekly and weekly paycheck receivers. For bi-weekly paycheck receivers, I further distinguish between "odd" and "even" pay schedules. It's clear from the figure that there is a strong bi-weekly pattern in food expenditures. Furthermore, the opposing bi-weekly pay schedules make it clear that the spikes are associated with paycheck receipt

[7] I combine the series "7221: Full service restaurants" and "7222: Limited service eating places" from the U.S. Census Bureau Monthly Retail Trade and Food Services report.

and not other recurring events like the first of the month. The weekly paycheck series is much smoother but still follows the overall trend seen in the bi-weekly paycheck schedules.

Figure 3: Weekly food expenditures

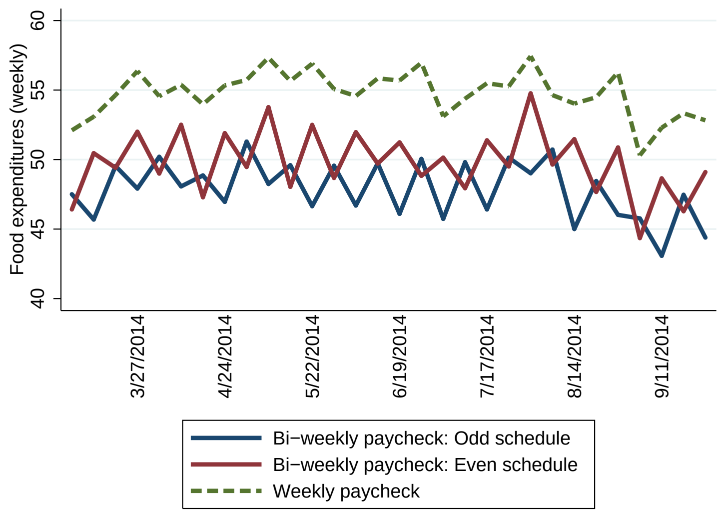

3.2 Excess sensitivity of food expenditures

The time series for bi-weekly paycheck receivers in Figure 3 indicate that expenditures rise in sync with weeks in which individuals are paid. The time series for weekly paycheck receivers reveal that expenditures rise in some weeks even for those that receive a paycheck every week. In order to estimate the rise in expenditures from receiving a paycheck while controlling for seasonal expenditure fluctuation, I estimate the following specification.

$$log(Food_{it}) = \alpha_i + \beta_1 Even_t + \beta_2 Payweek_{it}^{Even} + \beta_3 Payweek_{it}^{Odd} + \varepsilon_{it} \qquad (1)$$

where $Even_t$ is an indicator variable for whether week t is an even week, $Payweek_{it}^{Even}$ and $Payweek_{it}^{Odd}$ are indicator variables for whether individual i receives bi-weekly paychecks on week t on the even and odd schedule respectively, and α_i represents an individual fixed effect. β_2 and β_3 capture the growth rate of food expenditures on payweeks for those on the bi-weekly even and odd schedule respectively. β_1 captures the growth rate of food expenditures on even weeks. Including the weekly paid individuals helps to control for these seasonal trends that aren't necessarily associated with receiving a paycheck like first of the month effects or holidays that tend to fall on even weeks.

Column (1) of table 2 shows the coefficient estimates from estimating specification

(1). The estimate of 0.012 on $Even_{it}$ represents the fact that food expenditures grow by 1% on average during even weeks. The coefficients on $Payweek_{it}^{Even}$ and $Payweek_{it}^{Odd}$ are nearly identical and we cannot reject the null hypothesis that the magnitudes are the same. These estimates imply that food expenditures grow by an additional 5.5% on weeks in which bi-weekly individuals are paid after controlling for general seasonal trends. The magnitude of these estimates are in line with Stephens (2003), Shapiro (2005), Stephens (2006), and Kuchler (2015). The granularity of the data allow for much more accurate measurement of receipt of paychecks which results in more precise estimates relative the the previous studies.

Breaking out the bi-weekly paycheck schedule into odd and even help to ensure that the effect we are estimating is the arrival of the paycheck rather than other seasonal effects that may be correlated with paycheck arrival such as beginning of the month effects. Column (2) of of table 2 estimates the expenditure response by pooling together both even and odd schedules. As expected, the pooled effect is very similar to the estimates in column (1). The analysis in the later part of the paper will focus on the pooled effect.

Table 2: Excess sensitivity estimates

	(1)	(2)
VARIABLES	$ln(Food_{it})$	$ln(Food_{it})$
$Even_t$	0.012***	0.012***
	(0.002)	(0.001)
$Payweek_{it}^{Even}$	0.055***	
	(0.003)	
$Payweek_{it}^{Odd}$	0.054***	
	(0.003)	
$Payweek_{it}$		0.055***
		(0.001)
Observations	3,193,752	3,193,752
R-squared	0.276	0.276

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The standard explanation for the excess sensitivity seen in table 2 is that individuals are temporary liquidity constrained. Following the literature, table 3 re-estimates equation (1) by terciles of 2013 average liquidity.[8] The estimation only uses data from 2014 and onward to ensure that there is no mechanical correlation with the measure used to split the sample. In line with the previous literature, individuals that have lower levels of liquidity tend to react more strongly to the receipt of a paycheck relative to those with higher levels of liquidity. For example, food expenditures increase by 10% on average during weeks in which a paycheck is received for individual with low average levels of liquidity relative to 2% for individuals with high levels of liquidity. The coefficient on $Even_t$ is fairly similar across liquidity terciles. This is consistent with the view that that $Even_t$ captures aggregate trends that are common to all individuals.

Table 3: Excess sensitivity estimates by liquidity tercile

	(1)	(2)	(3)
VARIABLES	Low avg liquidity	Medium avg liquidity	High avg liquidity
$Even_t$	0.009***	0.010***	0.014***
	(0.003)	(0.003)	(0.003)
$Payweek_{it}^{Even}$	0.100***	0.043***	0.021***
	(0.005)	(0.005)	(0.005)
$Payweek_{it}^{Odd}$	0.099***	0.039***	0.017***
	(0.006)	(0.005)	(0.005)
Observations	748,692	754,908	701,221
R-squared	0.292	0.305	0.306

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This section has documented the presence of excess sensitivity of food expenditures to the receipt of a paycheck using financial account data from a personal finance app. The estimates are in line with the previous literature and provide more precise estimates than previous studies. The main goal of this section is to set the stage to further investigate whether the standard explanation that liquidity constraints explain excess sensitivity of expenditure to paychecks is correct. The next section introduces a theoretical model of consumption which will allow us to more formally test the standard explanation.

[8] Pooling the estimates across payweek schedules results in weighted means of the coefficients on $Payweek_{it}^{Even}$ and $Payweek_{it}^{Even}$. I do not report the coefficients here but can supply them upon request.

4 Buffer stock model of consumption

This section describes the model used to analyze consumption decisions. Individuals behave according to the standard "buffer-stock" saver model in the spirit of Zeldes (1989b), Deaton (1991), and Carroll (1997).

Optimization problem An individual solves the following utility maximization problem

$$\max_{\{C_j\}_{j=t}^{\infty}} \mathbb{E}_t \left[\sum_{j=t}^{\infty} \beta^{j-t} \frac{C_j^{1-\theta}}{1-\theta} \right] \tag{2}$$

subject to

$$A_{t+1} = (1+r)(A_t + Y_t - C_t) \tag{3}$$

$$A_{t+1} \geq b \tag{4}$$

$$Y_t = \bar{Y} + \varepsilon_t \tag{5}$$

$$\varepsilon_t \overset{iid}{\sim} N(\mu_y, \sigma_y^2) \tag{6}$$

where β, r, C_t, A_t and Y_t represent the time discount factor, the interest rate, consumption, liquid assets, and income respectively. Each period t represents a week. Y_t is further broken down into a constant term $\bar{Y}$ which represents a recurring paycheck and a stochastic term ε_t that represents non-paycheck income.

Income process I model the income process to match individuals who receive bi-weekly paychecks. Therefore, individuals receive a paycheck every other period. Overall, paycheck income comprises 70% of total income.

Solution The consumption problem specified above does not admit a closed form solution and is therefore solved computationally. I reformulate the individual's problem in terms of a functional equation and define cash on hand $x_t = a_t + y_t$ to simplify the state space. This variable represents the amount of resources available to the individual in the beginning of the period.

The individual then solves the optimization problem

$$V(x_t) = \max_{a_{t+1}} \{u(c_t) + \beta \mathbb{E}[V(x_{t+1})]\} \tag{7}$$

subject to

$$x_{t+1} = (1+r)(x_t - c_t) + y_{t+1} \tag{8}$$

and the previous constraints (4), (5), and (6).

Substituting in for c_t and x_{t+1} results in an equation in terms of x_t, a_{t+1}, and y_{t+1}

$$V(x_t) = \max_{a_{t+1}} \left\{ u\left(x_t - \frac{a_{t+1}}{1+r}\right) + \beta \mathbb{E}[V(a_{t+1} + y_{t+1})] \right\} \tag{9}$$

The individual maximizes utility by choosing next period saving (a_{t+1}) conditional on cash on hand (x_t). The model is solved using value function iteration which results in the value function $V(x_t)$ and the policy function $a_{t+1}(x_t)$ which maps the state variable x_t into the optimal control variable a_{t+1}. The consumption function is calculated using constraint (4) so that $c_t(x_t) = x_t - \frac{a_{t+1}}{1+r}$.

5 Model analysis

The buffer stock model introduced in the previous section can help us understand the cause of the excess sensitivity observed in section 3.2. In this section, I test whether the model can generate similar patterns as seen in the data. Furthermore, I explore which parameters are important for explaining the observed data. The parameter values used to calibrate the model are listed in Table 4 and represent weekly time periods. The utility function is specified as constant relative risk aversion (CRRA) with $\theta = 1$.

Table 4: Parameter values

Parameter	Value	Description	Notes
$u(x)$	$\frac{x^{1-\theta}}{1-\theta}$	utility function	CRRA utility
θ	1	coefficient of relative risk aversion	standard
β	0.95	time preference	calibrated to match the data
μ_y	0.30	non-paycheck income share	estimated from data
σ_y	0.10	S.D. of temporary shocks	from Gelman (2017)
$\bar{y}$	1.4	paycheck income share	estimated from data
r	0.01 / 52	interest rate	weekly r on checking/saving
b	0	borrowing limit	no borrowing condition

Notes: The parameters correspond to a weekly frequency.

5.1 Understanding excess sensitivity

As seen in figure 3, one important feature of the data when observed at a weekly aggregation is the consistent spike in expenditures during the paycheck week with a subsequent drop in the non-paycheck week. Figure 4 panel (a) below plots weekly log deviations of food expenditures to their average from March to October of 2014.[9] Panel (b) plots a random subsample of simulated time series in the buffer-stock model. By incorporating the receipt of a regular bi-weekly paycheck, the model can easily explain the spikes in expenditures upon paycheck receipt.

Figure 4: Weekly time series for bi-weekly paycheck receivers

(a) Data (b) Model

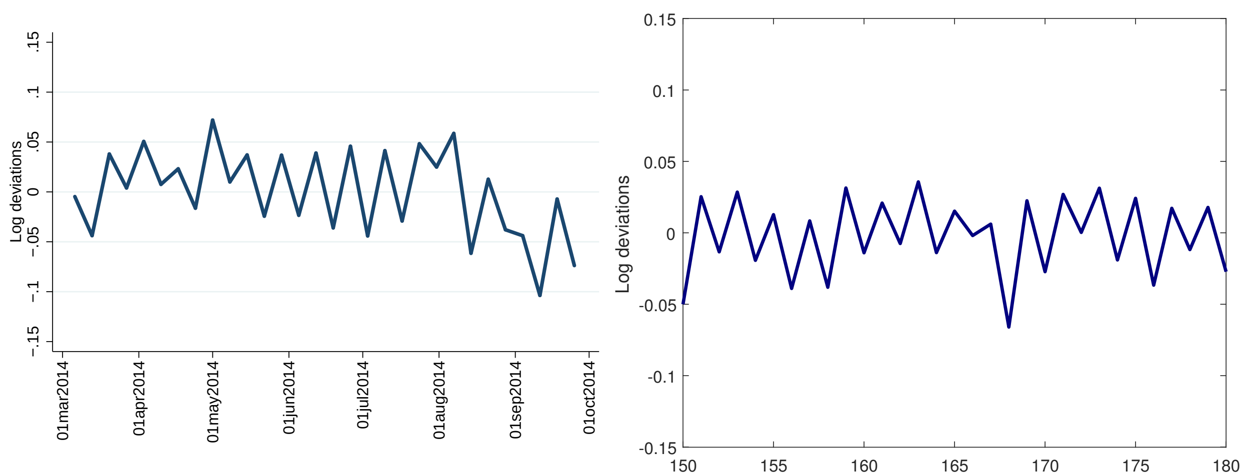

The model further allows us to investigate what causes these spikes in expenditures. In this particular model, the time discount factor is the most important parameter that influences the spike in expenditures. This is seen in figure 5 panel (b) where I simulate the model for different time preference parameters. For patient individuals with high time preference, the time series is relatively smooth. Conversely, impatient individuals with low time preference exhibit much larger spikes. In the data, splitting up individuals into average liquidity terciles as in panel (a) leads to differences in the peaks and troughs of log deviations. Individuals with low average liquidity tend to react more strongly to the receipt of a paycheck relative to individuals with high average liquidity. Most studies see this evidence and conclude that temporary liquidity constraints explains excess sensitivity. However, in the model, temporary liquidity constraints cannot explain excess sensitivity because individuals are rarely constrained during the week in which they receive their paycheck. It is during the week in which they are paid that individuals make the decision on how to allocate expenditures between this week and next week. The week after the paycheck is received is simply a reaction to the decisions made during the paycheck week. The next section will make this more clear by formally exploring how expenditure growth is determined in the paycheck week and the non-paycheck week.

[9] I only plot individuals on the bi-weekly odd schedule to simplify the output. The results would look similar but opposite for the even schedule.

Figure 5: Weekly time series for bi-weekly paycheck receivers (heterogeneity)

(a) Data (b) Model

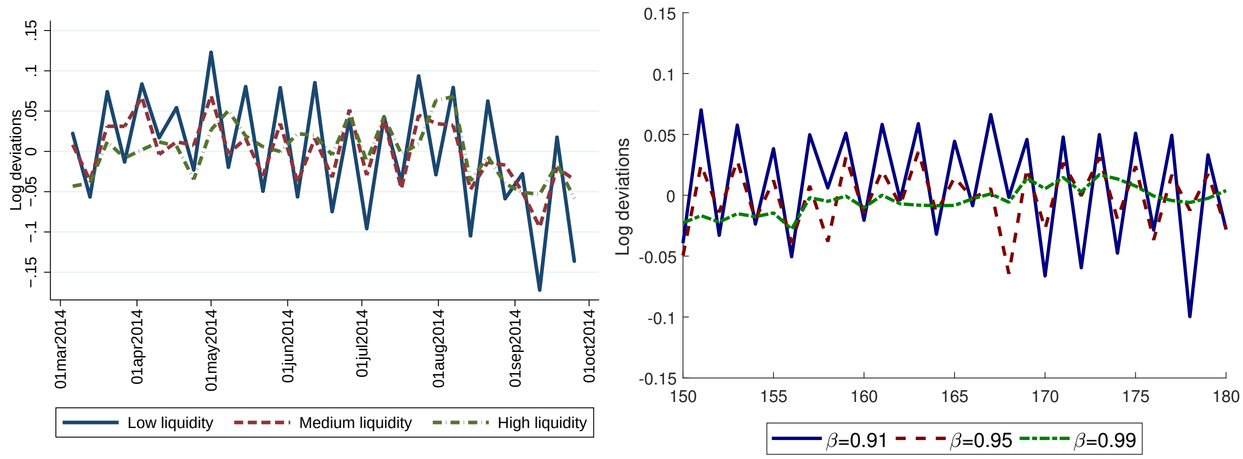

5.2 Excess sensitivity and liquidity constraints

The excess sensitivity documented in the previous sections can be interpreted as positive consumption growth in weeks in which a paycheck is not received and negative consumption growth in weeks in which a paycheck is received. In order to understand excess sensitivity, it's important to understand what influences consumption growth. The model provides a key equation that can provide much insight. The key equation can be derived from the optimality conditions of the consumption optimization problem specified in section 4. The second order approximation of the optimality condition is commonly known as the consumption euler equation and is written below as

$$\underbrace{\Delta ln(c_{t+1})}_{\text{consumption growth}} \approx \overbrace{\frac{r-\delta}{\theta}}^{\text{impatience}} + \underbrace{\frac{\theta}{2}\sigma^2_{t+1}(x_t)}_{\text{precautionary saving}} + \underbrace{\lambda_t(x_t)}_{\text{liquidity constraint}} + \varepsilon_{t+1} \tag{10}$$

where c_t is consumption, $\delta = \frac{1}{\beta} - 1$ is the discount rate, θ is the coefficient of relative risk aversion, σ^2_t is a measure of consumption growth volatility, r is the interest rate, and ε_t is a mean zero rational expectations error.

The equation shows that consumption growth is influenced by three terms. The first term is constant and represents desired consumption growth in the absence of any precautionary saving or liquidity constraints. It is driven by the difference between the interest rate and the time discount rate scaled by the intertemporal elasticity of substitution.

The second term represents precautionary saving motives. As explained in Kimball

(1990), a positive third derivative of the utility function induces a precautionary saving motive which will tend to cause individuals to save for tomorrow in favor of consuming today. This term will tend to increase consumption growth by lowering consumption today.

Lastly, the third term represents liquidity constraints. If the constraint is binding, this term will also increase consumption growth because individuals cannot increase consumption today relative to their desired amount.

In general, it is difficult to derive analytical expressions for the precautionary saving and liquidity constraint terms. However, we do know that they are functions of cash on hand x_t. Variation in x_t is driven by both uncertain income as well as predictable changes that arise from different consumption levels in paycheck and non-paycheck weeks. For the liquidity constraint term, there is a value of x_t for which the constraint will begin to bind and so it is an increasing function of x_t. Similarly, the precautionary saving motive is an increasing function of x_t. The intuition is that when x_t is small, an individual is not able to smooth shocks very well leading to a wide range of possible consumption values in the next period depending on the realization of the labor income shock. This translates into high variability in consumption growth. Conversely, when x_t is high, an individual is easily able to smooth consumption in the face of income shocks so there will be little variation in consumption growth. In the limit, as $x_t \to \infty$, liquidity constraints will be unlikely to bind and precautionary fears become irrelevant. In that case, consumption growth will be dominated by the impatience term.

In order to better understand these mechanisms, panel (a) of figure 6 plots expected consumption growth from the model on the y-axis against relative liquidity for weeks in which the paycheck is received on the x-axis. Relative liquidity is defined as the log difference of liquidity in time t from it's average. In general, expected consumption growth is negative because consumption tends to be higher in the paycheck week relative to the non-paycheck week. The main result here is that expected consumption growth does not vary much by relative liquidity. This is not surprising if we assume the precautionary saving and liquidity constraints terms in equation (10) are small during the payweek.

Typically, the theoretical relationship plotted in panel (a) is hard to estimate empirically. There are few datasets for which liquidity is observed at such a high frequency jointly with expenditure growth and the timing of paycheck arrival. Utilizing these unique features of the financial app data, panel (b) of figure 6 estimates the empirical analogue to panel (a) by using realized food expenditure growth. More specifically, panel (b) plots a smoothed local linear relationship between food expenditure growth and log deviations from average liquidity in the week in which the paycheck is received. This relationship is estimated for each tercile of average liquidity. Similar to the theoretical predictions, food expenditure growth does not vary much with relative liquidity. During weeks in which individuals do receive their bi-weekly paycheck, individuals are likely to be flush with liquidity and so their decisions should not depend on liquidity. To the best of my knowledge, this is the first study to show that food expenditure growth is not a function of liquidity in the week that the paycheck is received. Under this interpretation, for individuals who receive regular paychecks, persistent non-zero expenditure growth in the payweek or "excess sensitivity" is not due to temporary liquidity constraints but must be due to persistent characteristics such as preferences. Lastly, individuals with low average liquidity tend to have lower levels of food expenditure growth in the pay week relative to high average liquidity individuals. The interpretation under the buffer stock model is that low levels of time preference will jointly generate low expenditure growth and low levels of average liquidity.

Figure 6: Consumption/expenditure growth and relative liquidity (pay week)

(b) Data

(a) Model

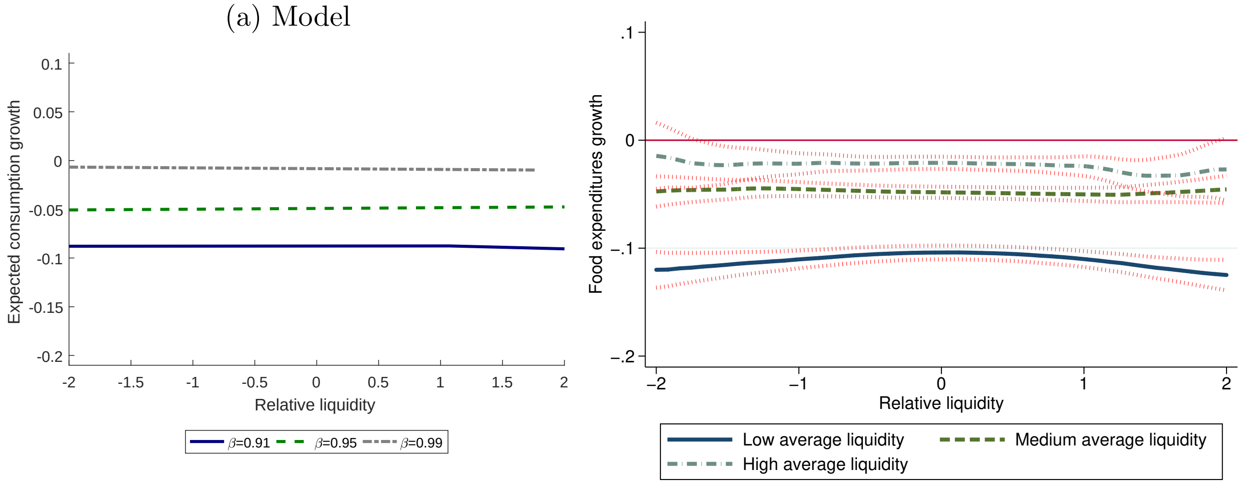

While explaining consumption behavior during the payweek is the main focus of the analysis, it is also helpful to analyze consumption behavior in the non-payweek. It may be the case that in our data, food expenditure growth is not sensitive to relative liquidity and so just happens to fit the theory laid out by the model. Panel (b) of figure 7 plots a smoothed local linear relationship between food expenditure growth and log deviations from average liquidity in the week in which the paycheck is not not received. There are several important differences relative to figure 6. First and most importantly, food expenditure growth does appear to vary with relative liquidity. This is important because it shows that the result in figure 6 is not due to the fact that a relationship between food expenditures growth and relative liquidity cannot be estimated in the data. Second, food expenditure tends to be positive in the non-pay week because expenditures increase when the paycheck is received. Third, the negative relationship between food expenditure growth and relative liquidity appears to be stronger for lower liquidity individuals. Panel (a) of 7 plots the theoretical relationship in the model between expected consumption growth and relative liquidity in weeks in which the paycheck is not received. All three of the empirical facts seen in the data are also present in the theoretical relationship. For example, the model shows that (1) there is a negative relationship between expected consumption growth and relative liquidity, (2) expected consumption growth tends to be positive, and (3) the relationship between expected consumption growth and relative liquidity is stronger for more impatient individuals.

Figure 7: Consumption/expenditure growth and relative liquidity (non-pay week)

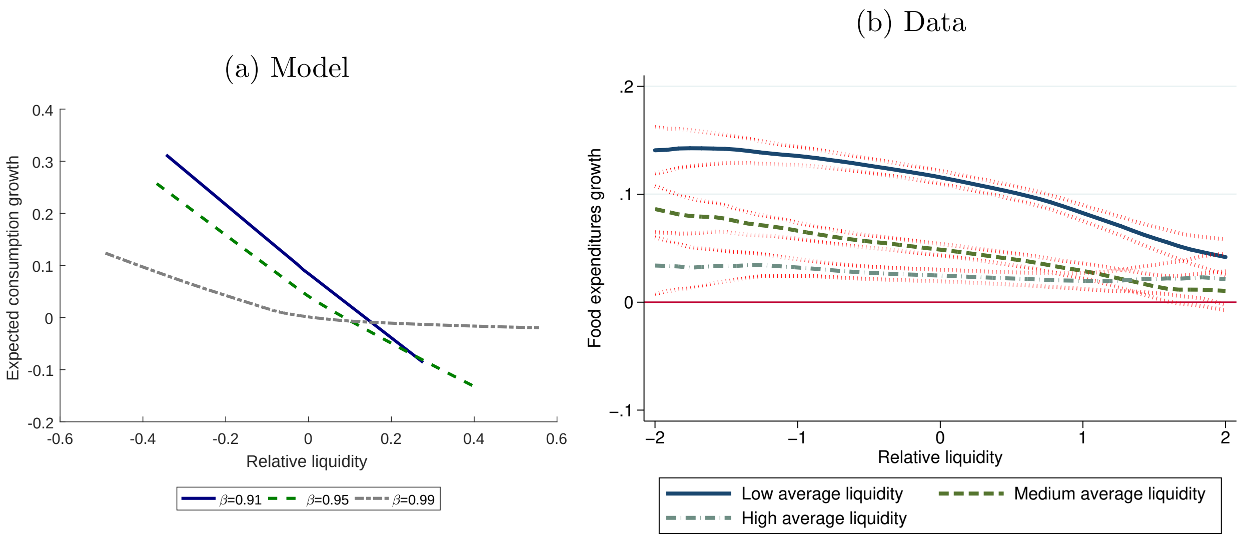

Although these three facts are qualitatively present in both the model and the data, there are some important differences. First, the range of relative liquidity is much smaller in the model. This is due to the fact that in the pay period, the certain paycheck keeps liquidity from deviating too much from its mean value. In the actual data, there appear to be more extreme shocks to liquidity that may arise from one-off transfers in and out of accounts that are not incorporated into the model. Second, in the model, consumption growth doesn't appear to flatten out around zero for those that are more impatient. This is due to the fact that when faced with a very positive shock, impatient individuals might actually consume more in the non-pay period relative to the pay period if they receive a large negative shock in the pay period. In the data, this rarely seems to be the case and on average food expenditure growth does not become negative even when relative liquidity in the non-pay period is quite high. Lastly, the relationship between expected consumption growth and relative liquidity is much steeper in the model relative to the data. This likely reflects the fact that individuals in our data may use credit cards and the delay of recurring expenses such as bills (as seen in Gelman et al. (2015)) to smooth consumption better than the model predicts.

While the model can be modified to better match the features of the data seen in panel (b) of figure 7, I believe it would needlessly complicate the model without changing the main results that obtain in figure 6. Namely, the main conceptual insight and novel empirical result is to show that in weeks in which individuals are paid, their consumption growth is not affected greatly by liquidity.

Table 5 summarizes the results from this section. It lists the estimated coefficients from the specification

$$\Delta ln(food_{it+1}) = \alpha_i \times payweek_{it} + \beta_2 \times liq_{it-1}^{pay} + \beta_3 \times liq_{it-1}^{nopay} + \varepsilon_{it+1} \qquad (11)$$

where $\alpha_i \times payweek_{it}$ represents individual fixed effects for both paycheck and non-paycheck weeks, and liq_{it-1}^{pay} and liq_{it-1}^{nopay} represent $t-1$ log liquidity in the payweek and non-pay week respectively for individual i. I use $t-1$ liquidity because I want to measure the resources individual have when they enter period t. The individual fixed effects for both paycheck and non-paycheck weeks allow us to interpret liquidity as the percent change in the previous week relative to the pay and non-pay week. This relative measure is important because the liquidity levels are different in pay and non-pay weeks. Equation 11 is then estimated for each liquidity tercile. Intuitively, the coefficients from the econometric specification estimate the slope of the linear relationship captured in panel (b) of figures 6 and 7.

Table 5: Relationship between expenditure growth and relative liquidity

	(1)	(2)	(3)
VARIABLES	Low avg liquidity	Medium avg liquidity	High avg liquidity
Pay week	0.001	-0.003	-0.002
	(0.002)	(0.003)	(0.004)
Non pay week	-0.037***	-0.026***	-0.009**
	(0.003)	(0.003)	(0.004)
Observations	363,714	416,502	383,654
R-squared	0.056	0.036	0.033

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The table confirms the results discussed in the figures above. They are:

1. There is no statistically significant relationship between expenditure growth and relative liquidity in the pay week.
2. There is a statistically significant relationship between expenditure growth and relative liquidity in the non-pay week.
3. The relationship between expenditure growth and relative liquidity is stronger for individuals with lower average liquidity.

These results are in line with the model as well as the intuition that individuals are very unlikely to be constrained in the week in which they receive their paycheck and so should not respond much to their liquidity levels at the beginning of the week.

In previous studies, researchers have often observed that average liquidity levels are strong predictors of how individuals respond to paychecks. The analysis in this section makes it clear that we should be cautious about interpreting these results as evidence that temporary liquidity constraints explain excess sensitivity. Instead, the results imply that preferences jointly generate excess sensitivity as well as lower levels of average liquidity. In this simple buffer stock model, excess sensitivity reflects preferences and not constraints.

Another way to interpret the results is that liquidity constraints are endogenous. The model makes it clear that individuals who are impatient will tend to consume more of their paycheck in the pay week leaving them vulnerable to shocks in the non pay week. It is in this sense that individuals are constraining themselves by not transferring sufficient resources across weeks to buffer high frequency shocks. Only by looking at behavior during the pay week when individuals are unlikely to be constrained, can we observe the true preference of individuals.

5.3 Excess sensitivity and income

If the explanation in the previous section is true, average liquidity can be thought of as a proxy for preferences. Conversely, paycheck income in the model is exogenous and so does not reflect preferences. To test this assumption, figure 8 estimates the relationship between food expenditure growth and relative liquidity for different terciles of paycheck income. The results show that paycheck income terciles do not differentiate between different levels of food expenditure growth as well as liquidity terciles. Furthermore, the ordering of the relationships by tercile doesn't generally match the model predictions.

Figure 8: Expenditure growth and relative liquidity (by income heterogeneity)

(a) Non-pay week

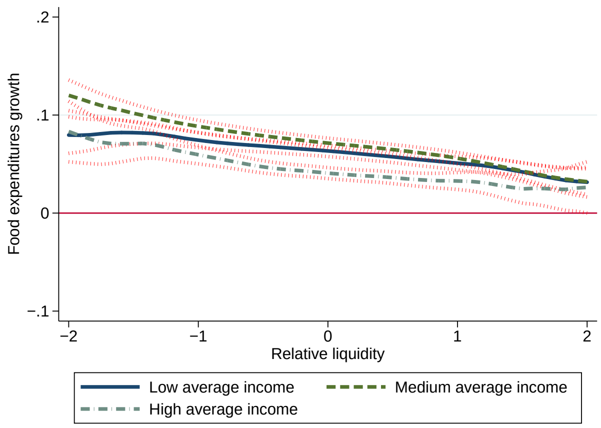

(b) Pay week

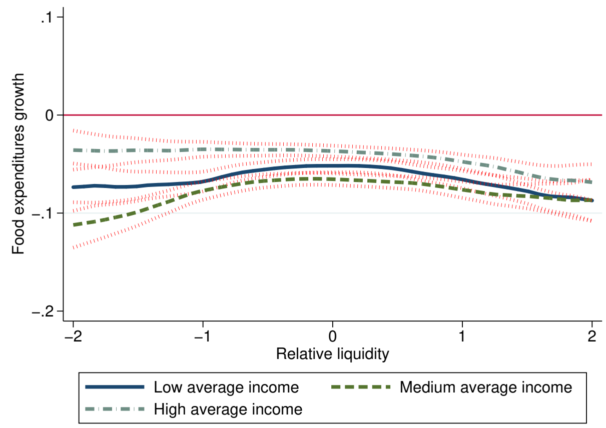

The results from figure 8 show that not all variables used to split the sample will generate the results seen in figure 6 and 7. While the assumption that paycheck income is not correlated with preferences may not be entirely true, it seems reasonable to assume that average levels of buffer stock saving are more strongly correlated with preferences than paycheck income.

6 Testing the implications of liquidity constraints using tax refunds

The main result in this study is that expenditure growth is relatively unaffected by liquidity in the pay week while it is significantly affected by liquidity during the non-pay week. One way to see this mechanism in action is by looking at the response to receiving a tax refund. More specifically, if individuals are liquidity constrained during the non-pay week, we should observe a stronger reaction to the refund if it is received during a non pay week relative to a pay week.

6.1 Impulse response to tax refund

This section analyzes the spending response to the tax refund in order to validate the use of receiving a tax refund as an increase in liquidity.

I estimate the distributed lag of receiving a tax refund using the following specification.

$$ln(food_{it}) = \alpha_i + \sum_{j=-10}^{10} \beta_j \times refund_{it-j} + \delta_t + \varepsilon_{it} \tag{12}$$

where $refund_{it-j}$ represents an indicator variable for whether person i received a refund in week $t-j$, δ_t represents week fixed effects, and ε_{it} is the error term. Figure 9 below plots the β_j from equation (12). The estimates show that there is little anticipatory response of food expenditures to receiving a tax refund. On the week the refund is received, there is a large spike in food expenditures that slowly decays over time. This response is consistent with Souleles (1999) which examines the consumption response to income tax refunds in the CEX and Baugh, Ben-David and Park (2014) which studies the weekly response of spending to the arrival of tax refunds using similar account data.

Figure 9: Consumption response to the receipt of a tax refund

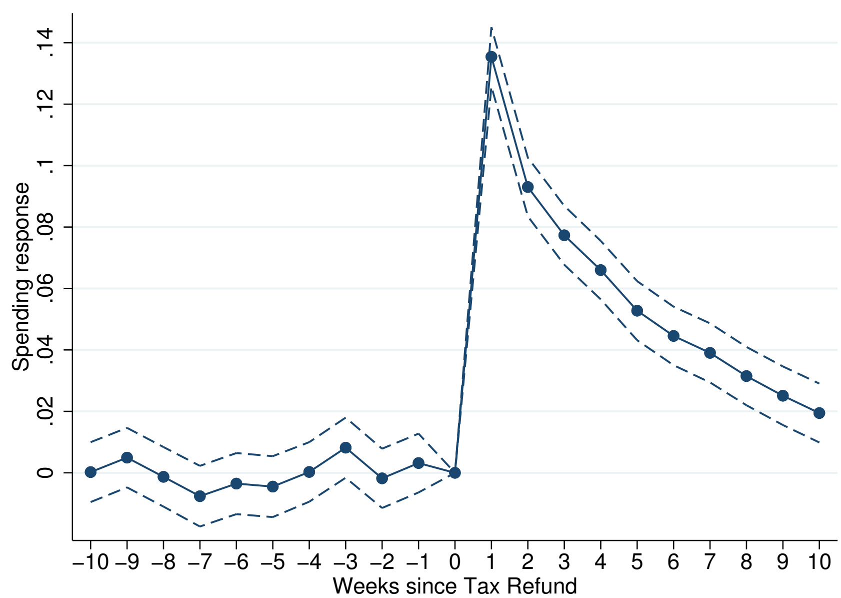

In summary, the estimated response of food expenditures to the receipt of a tax refund is similar in dynamics to previous studies. This fact helps to confirm that the spending response to these refunds is a valid instrument to test the main results of this study.

6.2 The effect of receiving a tax refund on expenditure growth

This section estimates the effect of receiving a tax refund on expenditure growth. It also tests whether the effect is different depending on whether the week in which the tax refund is received is a pay period or a non-pay period. The econometric specification is

$$\Delta ln(food_{it+1}) = \alpha_i + \beta_1 \times refund_{it} + \beta_2 \times payweek_{it} + \beta_3 \times refund_{it} \times payweek_{it} + \varepsilon_{it+1} \tag{13}$$

where $refund_{it}$ and $payweek_{it}$ are indicator variables for whether a refund or a paycheck was received for person i in week t and α_i is an individual-level fixed effect.

Table 6 shows the results from estimating equation (13) for each liquidity tercile. The coefficient on $payweek_{it}$ shows that expenditure growth is negative in weeks in which a paycheck is received. This is in line with the excess sensitivity captured in earlier results. Similarly, the coefficient on $refund_{it}$ shows expenditure growth is negative in weeks in which a tax refund is received. This indicates that individuals increase expenditures when they receive a tax refund. The positive coefficients on $refund_{it} \times payweek_{it}$ show that expenditure growth is less negative when the refund is received during weeks in which the paycheck is also received. This is consistent with the notion that individuals are more liquidity constrained in weeks in which they don't receive a paycheck. For example, for individuals with low average liquidity, expenditure growth is 12% lower during weeks in which a refund is received and a paycheck is not received. This reflects the fact that expenditures exhibit a large spike upon arrival of a tax refund on non-pay weeks. If the refund is received in the same week that the paycheck is received, expenditure growth is only 3% lower relative to weeks in which the refund is not received. The results show that whether individuals react strongly to a tax refund depends mostly on whether it is a pay week or not (and hence how liquidity constrained they are) and less so on what average liquidity group they belong to.

Table 6: Coefficient estimates

VARIABLES	(1) Low avg liquidity	(2) Medium avg liquidity	(3) High avg liquidity
$refund_{it}$	-0.118***	-0.056***	-0.010
	(0.023)	(0.021)	(0.023)
$payweek_{it}$	-0.221***	-0.094***	-0.047***
	(0.004)	(0.003)	(0.004)
$refund_{it} \times payweek_{it}$	0.084***	0.012	0.006
	(0.032)	(0.029)	(0.032)
Observations	375,965	419,802	385,760
R-squared	0.013	0.005	0.004

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

6.3 The effect of receiving a tax refund on the relationship between expenditure growth and relative liquidity

This section takes a closer look at how receiving a tax refund affects the relationship between expenditure growth and relative liquidity. Figure 10 shows the relationship between expenditure growth and relative liquidity during weeks in which the paycheck is not received. As seen earlier, in weeks in which the paycheck is not received, expenditure growth has a strong negative relationship with relative liquidity. However, on weeks in which the tax refund is received, that strong negative relationship no longer holds. One way to interpret this is that individuals are usually very cash starved during weeks in which they don't receive their paycheck because they choose to consume more during weeks in which they receive their paychecks. Receiving a tax refund relaxes the liquidity constraints that usually bind. Due to the constraints being relaxed, expenditure growth is no longer affected by the amount of liquidity individuals carry over from the previous period.

Figure 10: Expenditure growth and relative liquidity (non-pay week)

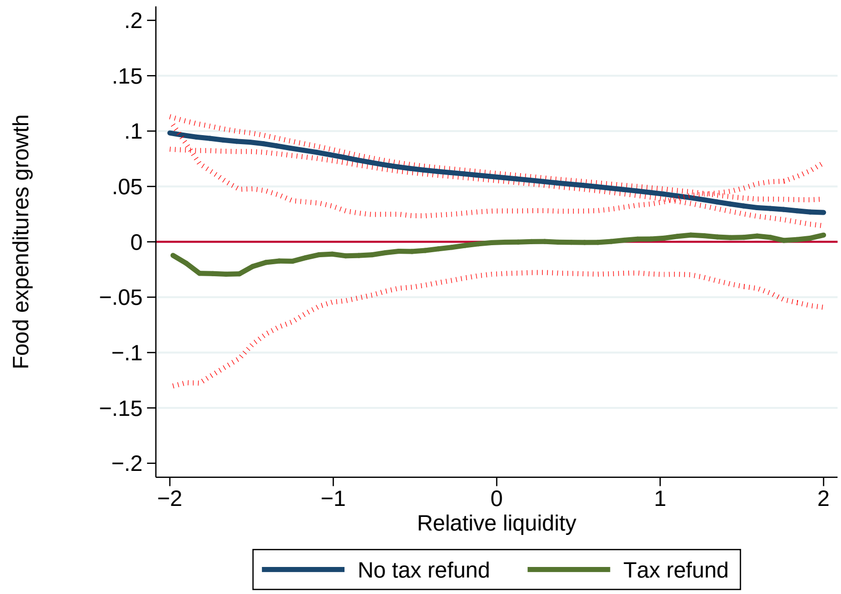

Figure 11 performs the same analysis but for weeks in which a paycheck is received. As seen in the previous sections, the relationship between expenditure growth and relative liquidity is much weaker during weeks in which the paycheck is received. Furthermore, because an individual typically has so much liquidity during pay weeks, the relationship does not appear to be very different in weeks in which a tax refund is also received.

Figure 11: Expenditure growth and relative liquidity (pay week)

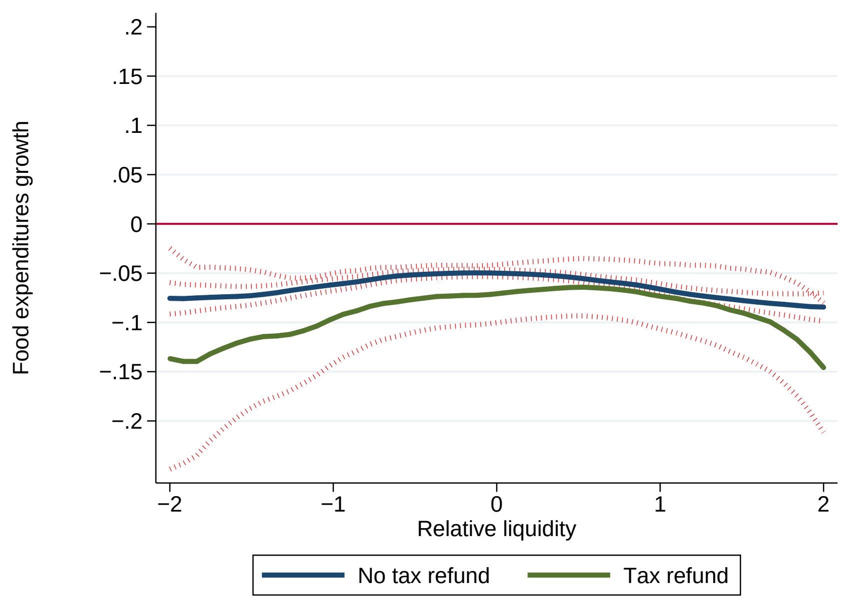

Because tax refunds are only received once a year, the results conditioning on weeks in which a tax refund is received are much less precise. In order to more formally analyze how receiving a refund affects the relationship between expenditure growth rate and relative liquidity, I estimate the following econometric specification

$$\begin{aligned}\Delta ln(food_{it+1}) = \alpha_i + \alpha_i \times payweek_{it} + \beta_1 \times liq^{pay}_{it-1} + \beta_2 \times liq^{pay}_{it-1} \times refund_{it} + \\ \beta_3 \times liq^{nopay}_{it-1} + \beta_4 \times liq^{nopay}_{it-1} \times refund_{it} + \\ \beta_5 \times refund_{it} + \beta_6 \times refund_{it} \times payweek_{it} + \varepsilon_{it+1} \quad (14)\end{aligned}$$

where liq^{pay}_{it-1} and liq^{nopay}_{it-1} capture the log of liquidity in the previous period when the current period is a pay week or non-pay week respectively. The specification aims to capture the differential marginal effect of relative liquidity on expenditure growth in weeks in which a tax refund is received. The negative coefficient on liq^{nopay}_{it-1} replicates the earlier result that relative liquidity is an important determinant of expenditure growth in non-pay weeks. Furthermore, the small and statistically insignificant result on liq^{pay}_{it-1} replicates the earlier result that relative liquidity is *not* an important determinant of expenditure growth in *pay* weeks. The new results of interest are the coefficients on $liq^{pay}_{it-1} \times refund_{it}$ and $liq^{nopay}_{it-1} \times refund_{it}$. They represent the additional effect of liquidity on expenditure growth during weeks in which the tax refund is received. The small and statistically insignificant coefficient on $liq^{pay}_{it-1} \times refund_{it}$ confirms that since liquidity is already high on pay weeks, receiving additional liquidity in the form of a tax refund does not have much of an effect. The positive and statistically significant coefficient on $liq^{nopay}_{it-1} \times refund_{it}$ confirms that since individuals tend to be liquidity constrained during non-pay weeks, receiving extra liquidity cancels out the negative relationship between relative liquidity and expenditure growth during non-pay weeks. To test this idea more formally, I calculate $\beta_3 + \beta_4 = 0.0116$ with a p-value of 0.201. Therefore, the econometric specification confirms the results in figure 10 that liquidity no longer affects expenditure growth in non-pay weeks after the tax refund relieves liquidity constraints.

Table 7: Relationship between expenditure growth and relative liquidity

	(1)
VARIABLES	$\Delta ln(food_{it+1})$
liq^{pay}_{it-1}	0.000
	(0.002)
$liq^{pay}_{it-1} \times refund_{it}$	0.011
	(0.008)
liq^{nopay}_{it-1}	-0.026***
	(0.002)
$liq^{nopay}_{it-1} \times refund_{it}$	0.037***
	(0.009)
$refund_{it}$	-0.349***
	(0.073)
$refund_{it} \times payweek_{it}$	0.234**
	(0.096)
Observations	1,394,974
R-squared	0.037

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

To summarize, this section tests the implications of the effects of liquidity on expenditure growth during pay and non-pay weeks. The main analysis suggests that liquidity only affects expenditure growth in non-pay weeks because this is when liquidity is low. It tests this implication by studying a case in which liquidity is increased in the form of a tax refund. Similarly to what the theory and empirics suggest, receiving a tax refund has different effects depending on whether it is received on a pay week or non-pay week. In general, expenditure growth is negative in weeks in which a tax refund is received as individuals increase expenditure relative to weeks in which a tax refund is not received. However, the analysis in this section shows that the impact of receiving a tax refund is greater in non-pay weeks. The analysis then proceeds by estimating the effect of receiving a tax refund on the relationship between expenditure growth and relative liquidity. The analysis shows that in weeks in which the tax refund is received, liquidity no longer affects expenditure growth in the non-pay week. These results are consistent with the interpretation that individuals are liquidity constrained during the non-pay week. The receipt of the tax refund allows us to test this assumption and confirms that indeed when liquidity constraints are relaxed, relative liquidity no longer affects expenditure growth.

7 Conclusion

This paper provides an alternative explanation for the excess sensitivity of expenditure to the receipt of a regular paycheck. In particular, it establishes that in a standard buffer stock model of consumption, excess sensitivity is best understood as a result of preferences rather than temporary liquidity constraints. The logic is that since the receipt of the paycheck relieves any liquidity constraints, behavior during the pay week can be seen as an individual choice. The key empirical fact that confirms this logic is that during the week individuals are paid, expenditure growth is not a function of liquidity. Because some individuals are impatient, they endogenously constrain themselves in the non paycheck period. This implies that excess sensitivity can be interpreted as a fall in expenditures due to the lack of liquidity in the non-pay period rather than an increase in expenditures upon arrival of a paycheck.

To formalize this intuition, I specify a parsimonious buffer stock model of consumption with realistic paycheck dynamics. Model simulations show that during the week in which a paycheck is received, consumption growth is not affected by changes in liquidity. I then turn to the data and show that indeed liquidity does not affect expenditure growth in the week in which the paycheck is received.

Under the specified model, the spike up in expenditures during the pay week is driven by the fact that some individuals are impatient and prefer to consume more when they have resources. Indeed, in the data, excess sensitivity is strongest for those with low average liquidity. This is consistent with the model as impatient individuals react more strongly to paychecks while at the same time holding less liquidity on average.

In the model, impatient individuals intentionally leave less liquidity for themselves next period, thus making them vulnerable to shocks in weeks in which a paycheck is not received. I further test this assumption by showing how an influx of liquidity affects expenditure behavior. In pay weeks, individuals are already awash with liquidity so they should not react much to extra liquidity. Conversely, in non pay weeks, individuals that have left themselves with fewer resources should react strongly to liquidity. Using the extra liquidity provided by the receipt of a tax refund, I find that expenditure behavior once again matches the predictions of the model.

Both the model and the empirical results imply that excess sensitivity is not caused by temporary liquidity constraints. Instead, excess sensitivity is an optimal outcome for impatient individuals that face high frequency fluctuations in income.

References

Baker, Scott (2017) "Debt and the response to household income shocks: Validation and application of linked financial account data," *Journal of Political Economy.*

Baker, Scott R and Constantine Yannelis (2017) "Income changes and consumption: Evidence from the 2013 federal government shutdown," *Review of Economic Dynamics*, Vol. 23, pp. 99–124.

Baugh, Brian, Itzhak Ben-David, and Hoonsuk Park (2014) "Disentangling Financial Constraints, Precautionary Savings, and Myopia: Household Behavior Surrounding Federal Tax Returns,"Technical report, National Bureau of Economic Research.

Carroll, Christopher D. (1997) "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, Vol. 112, No. 1, pp. 1–55.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White (2017) "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, Vol. 8, No. 3, pp. 977–1020.

Deaton, Angus (1991) "Saving and Liquidity Constraints," *Econometrica*, Vol. 59, No. 5, pp. 1221–1248.

Ganong, Peter and Pascal Noel (2016) "How Does Unemployment Affect Consumer Spending?"

Gelman, Michael (2017) "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics."

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*, Vol. 345, No. 6193, pp. 212–215.

——— (2015) "How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data," Working Paper 21025, National Bureau of Economic Research.

Hall, Robert E (1978) "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, Vol. 86, No. 6, pp. 971–87.

Jappelli, Tullio and Luigi Pistaferri (2010) "The Consumption Response to Income Changes," *Annual Review of Economics*, Vol. 2, No. 1, pp. 479–506.

Kimball, Miles S. (1990) "Precautionary Saving in the Small and in the Large," *Econometrica*, Vol. 58, No. 1, pp. 53–73.

Kuchler, Theresa (2015) "Sticking to your plan: Hyperbolic discounting and credit card debt paydown," *Available at SSRN 2629158*.

Kueng, Lorenz (2018) "Excess Sensitivity of High-Income Consumers," SSRN Scholarly Paper ID 2627893, Social Science Research Network, Rochester, NY.

Laibson, David (1997) "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, Vol. 112, No. 2, pp. 443–478.

Olafsson, Arna and Michaela Pagel (2018) "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *The Review of Financial Studies*.

Parker, Jonathan A. (2017) "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment," *American Economic Journal: Macroeconomics*, Vol. 9, No. 4, pp. 153–183.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011) "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, Vol. 12, pp. 2825–2830.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod (2010) "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," in *Tax Policy and the Economy, Volume 24*. National Bureau of Economic Research, Inc, pp. 69–110.

Shapiro, Jesse M. (2005) "Is there a daily discount rate? Evidence from the food stamp nutrition cycle," *Journal of Public Economics*, Vol. 89, No. 23, pp. 303–325.

Souleles, Nicholas S. (1999) "The Response of Household Consumption to Income Tax Refunds," *The American Economic Review*, Vol. 89, No. 4, pp. 947–958.

Stephens, Melvin (2003) ""3rd of tha Month": Do Social Security Recipients Smooth Consumption between Checks?" *The American Economic Review*, Vol. 93, No. 1, pp. 406–422.

——— (2006) "Paycheque Receipt and the Timing of Consumption," *The Economic Journal*, Vol. 116, No. 513, pp. 680–701.

Zeldes, Stephen P (1989a) "Consumption and liquidity constraints: An empirical investigation," *The Journal of Political Economy*, Vol. 97, No. 2, pp. 305–346.

Zeldes, Stephen P. (1989b) "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, Vol. 104, No. 2, pp. 275–298.

A Appendix

A.1 Machine learning algorithm

Most transactions in the data do not contain direct information on expenditure category types. However, category types can be inferred from existing transaction data. In general, the mapping is not easy to construct. If a transaction is made at "McDonalds," it's easy to surmise that the category is "Fast Food Restaurants." However, it is much harder to identify smaller establishments such as "Bob's store." "Bob's store" may not uniquely identify an establishment in the data and it would take many hours of work to look up exactly what types of goods these smaller establishments sell. Luckily, the merchant category code (MCC) is observed for two account providers in the data. 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. If an individual uses an account provider that provides MCC information "Bob's store" will map into a expenditure category type.

The mapping from transaction data to MCC can be represented as $Y = f(X)$ where Y represents a vector of MCC codes and X represents a vector of transactions data. The data is partitioned into two sets based on whether Y is known or not.[10] The sets are also commonly referred to as training and prediction sets. The strategy is to then estimate the mapping $\hat{f}(\cdot)$ from (Y_1, X_1) and predict $\hat{Y}_0 = \hat{f}(X_0)$.

One option for the mapping is to use the multinomial logit model since the dependent variable is a categorical variable with no cardinal meaning. However, this approach is not well suited to textual data because each word would need its own dummy variable. Furthermore, interactions may be important for classifying expenditure categories. For example "jack in the box" refers to a fast food chain while "jack s surf shop" refers to a retail store. Including a dummy for each word can lead to about 300,000 variables. Including interaction terms will cause the number of variables to grow exponentially and will typically be unfeasible to estimate.

In order to handle the textual nature of the data I use a machine learning algorithm called random forest. A random forest model is composed of many decision trees that map transaction data to MCCs. This mapping is created by splitting the

[10] Y_0 represents the set where Y is not known and Y_1 represents the set where Y is known.

sample up into nodes depending on the features of the data. For example, for transactions that have the keyword "McDonalds" and transaction amounts less that \$20, the majority of the transactions are associated with a MCC that represents fast food. To better understand how the decision tree works, Figure 12 shows an example. The top node represents the state of the data before any splits have been made. The first row "transaction_amount $\leq$ 19.935" represents the splitting criteria of the first node. The second row is the Gini measure which is explained below. The third row show that there are 866,424 total transactions to be classified in the sample. The fourth row "value=[4202,34817,...,27158,720]" shows the number of transactions in each expenditure category. The last row represents the majority class in this node. Because "Restaurants" has the highest number of transactions, assigning a random transaction to this category minimizes the categorization error without knowing any information about the transaction. At each node in the tree, the sample is split based on a feature. For example, the first split will be based on whether the transaction amount is $\leq$ 19.935. The left node represents all the transactions for which the statement is true and vice versa. Transactions $\leq$ 19.935 are more likely to be "Restaurants" expenditure while transactions $>$ 19.934 are more likely to be "Gas and Grocery." In our example, the sample is split further to the left of the tree. Transactions with the string "mcdonalds" are virtually guaranteed to be "Restaurant" expenditure. A further split shows that the string "amazon" is almost perfectly correlated with the category "Retail Shopping." How does the algorithm decide which features to split the sample on? The basic intuition is that the algorithm should split the sample based on features that lead to the largest disparities in the different groups. For example, transactions that have the word "mcdonalds" will tend to split the sample into fast food and non-fast food transactions so it is a good feature to split on. Conversely, "bob" is not a very good feature to split on because it can represent a multitude of different types of expenditure depending on what the other features are.

Figure 12: Decision tree example

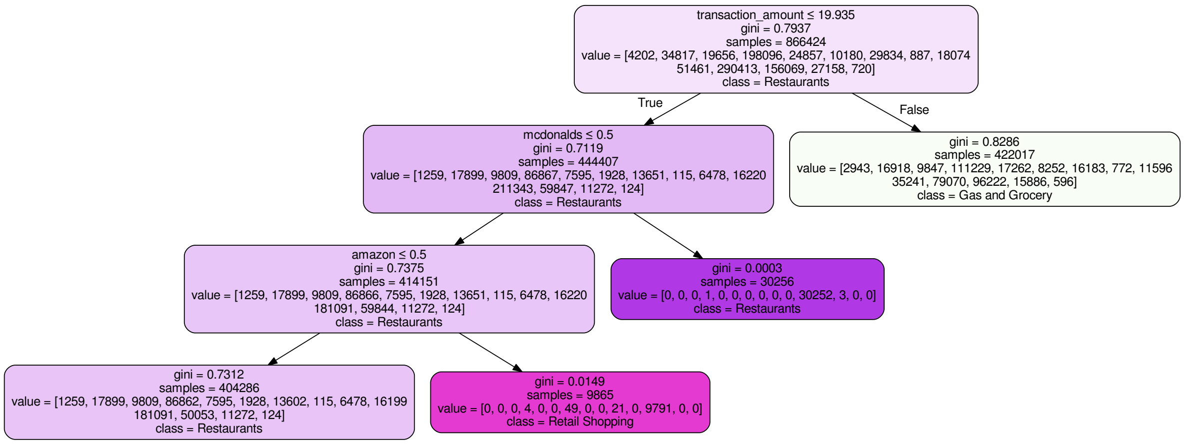

I state the procedure more formally by adapting the notation used in (Pedregosa et al., 2011). Define the possible features as vectors $X_i \in R^n$ and the expenditure categories as vector $y \in R^l$. Let the data at node m be presented by Q. For each candidate split $\theta = (j, t_m)$ consisting of a feature j and threshold t_m, partition the data into $Q_{left}(\theta)$ and $Q_{right}(\theta)$ subsets so that

$$Q_{left}(\theta) = (X, y)|x_j \leq t_m \tag{15}$$

$$Q_{right}(\theta) = Q \setminus Q_{left}(\theta) \tag{16}$$

The goal is then to split the data at each node in the starkest way possible. A popular quantitative measure of this idea is called the Gini criteria and is represented by

$$H(X_m) = \sum_k p_{mk}(1 - p_{mk}) \tag{17}$$

where $p_{mk} = 1/N_m \sum_{x_i \in R_m} \mathbb{I}(y_i = k)$ represents the proportion of category k observations in node m.

If there are only two categories, the function is is minimized at 0 when the transactions are perfectly split into the two categories[11] and maximized when the transactions are evenly split between the two categories.[12]

Therefore, the algorithm should choose the feature to split on that minimizes the Gini measure at node m

[11] because 0*1 + 1*0 = 0.

[12] because 0.5*0.5 + 0.5*0.5 = 0.5.

$$\theta^* = argmin_\theta \frac{n_{left}}{N_m} H(Q_{left}(\theta)) + \frac{n_{right}}{N_m} H(Q_{right}(\theta)) \tag{18}$$

The algorithm acts recursively so the same procedure is performed on $Q_{left}(\theta^*)$ and $Q_{right}(\theta^*)$ until a user-provided stopping criteria is reached. The final outcome is a decision rule $\hat{f}(\cdot)$ that maps features in the transaction data to expenditure categories.

This example shows that decision trees are much more effective in mapping high dimensional data that includes text to expenditure categories. However, fitting just one tree might lead to over-fitting. Therefore, a random forest fits many trees by bootstrapping the samples of the original data and also randomly selecting the features used in the decision tree. With the proliferation of processing power, each tree can be fit in parallel and the final decision rule is based on all the decision trees. The most common rule is take the majority decision of all the trees that are fit.