Real Time Measurement of Household Electronic Financial Transactions in a Population Representative Panel

Marco Angrisani,[a] Arie Kapteyn,[a,b] and Swaroop Samek[a]

[a] *University of Southern California, Center for Economic and Social Research*

[b] *National Bureau of Economic Research*

December 2020

Abstract

From September 2016 to September 2018, we piloted the collection of financial transaction and account balance data in the Understanding America Study (UAS), an Internet panel representative of the U.S. adult population. Unlike previous studies relying on electronic transaction information from financial aggregators, our data allow us to explore heterogeneity in consumer behavior as driven by demographics, health, cognitive ability, and financial literacy, among others. In this paper, we describe the results of this pilot project, documenting selectivity issues and the major barriers associated with participation in the study. We carry out illustrative exercises to highlight how the combination of surveys and electronic financial records can open new research avenues to better understand individual financial decision making and well-being.

1 Introduction

The measurement of financial and consumer behavior mainly relies on self-reports. Such reports are burdensome for respondents and tend to suffer from serious measurement and selectivity issues. Since most individual financial and spending behaviors (use of payment instruments, investments, assets and liabilities) are recorded in electronic form, direct access to electronic records may both improve accuracy of measurement and greatly reduce respondent burden. Yet, the mere availability of transaction data does not typically suffice to examine heterogeneity in consumer behavior and to better understand financial decisions across different segments of the population. For this purpose, the combination of rich survey information and transaction data would be ideal. To date, such combination of data has not been available.

In this paper, we describe a pilot project in which we have explored the collection of financial electronic records of respondents of the Understanding America Study (UAS), an Internet panel representative of the adult population in the United States, by accessing their electronic records directly through their financial institutions. We have designed and implemented a data collection system, using a contract with a major financial information aggregation firm, and asked the panel members' consent to share with us their financial information through the aggregator platform. The outcomes of this project provide important insights for future data collection efforts of this kind. Relying on both survey and transaction data, we carry out illustrative exercises, which highlight the potential of combining rich individual-level background information with financial electronic records for examining heterogeneity in consumer behavior and for future research avenues more generally.

A number of ongoing U.S. studies collect financial and expenditure information from households, including the Survey of Consumer Finances (SCF), the Health and Retirement Study (HRS), the Panel Study of Income Dynamics (PSID), the Consumer Expenditure Survey (CE), and the Survey of Income and Program Participation (SIPP). These studies are used in a wide variety of research projects. Typically, respondents are asked a large number of questions about their expenditures, financial holdings, and/or financial behavior. Typically, questions about finances and spending are burdensome and among the least popular for respondents to answer. For example, the median interview length in the 2010 SCF was about 90 minutes, but some interviews took more than three hours (Bricker et al., 2012). While the value of such data cannot be understated, several studies have assessed the quality of wealth information in major U.S. surveys with generally less than satisfactory conclusions (Bound et al., 2001).

Economic outcomes are mostly measured by eliciting factual concepts, involving exact amounts of money, instead of opinions, attitudes, or other subjective concepts. This poses several important challenges, which are bound to affect data quality from surveys. First, the respondent may not be certain which concept the investigator is interested in (e.g., should income be reported before or after taxes?). Defining concepts precisely is easier with narrow categories, but this implies asking more questions, thereby increasing respondent burden. Second, during the interview, a respondent may not think about rarely used accounts or other unusual assets, fail to report incidental income components, or not remember certain expenses. Studies that ask income, expenses, or wealth in multiple categories have higher averages than studies that ask only broad totals (Kapteyn et al., 1988; Browning et al., 2014). This is often interpreted as evidence that multiple categories are “better" (Winter, 2014) and should be used even at the cost of increasing survey length. However, while decomposition into narrow categories can make reports of irregular frequent behavior more accurate, it often implies less accurate reports of regular, frequent behavior (Menon, 1997). Third, when asking about income or expenses in a certain reference period, the respondent may inadvertently include income that was received or expenses that were incurred shortly before or after the reference period. The potential for errors is generally larger with longer reference periods or reference periods in a more distant past. For the purpose of accurately recording information, reference periods should ideally be short and recent. However, such short reference periods lead to substantial random variation (Hurd and Rohwedder, 2009) and greatly increase volatility of household expenditures (Angrisani et al., 2014). Fourth, respondents may misremember amounts, make rough guesses about amounts that they do not precisely know, or provide rounded amounts. Furthermore, respondents and interviewers may make typographical errors, or interviewers may mishear what respondents say.

Due to these and other problems, the measurement of income, asset values and expenditures in surveys is far from perfect (Moore et al., 1999; Hurd, 1999; Browning et al., 2003; Carroll et al., 2014). Inaccurate measures may lead to biased estimates of key economic parameters, which, in turn, may be used to poorly inform policy interventions. For instance, survey data make it difficult to disentangle transitory variations in labor income from changes attributable to reporting error (Meghir and Pistaferri, 2004). This may prevent a precise assessment of the degree of income risk faced by individuals and hamper the design of programs aimed at reducing it. Similarly, measurement errors in recall consumption data may greatly affect the estimation of consumption models and the testing of theoretical predictions (Attanasio, 2000).

An alternative to survey data on income and expenditure is represented by electronic records of financial transactions. Gelman et al. (2014) illustrate the potential of financial aggregator data for studying key economic questions. They use data on income and spending of more than 75,000 customers of Check (https://check.me) to study the responsiveness of spending to the arrival of anticipated income. Previous studies that have looked at this question using data from the CE lack precision (Souleles, 1999; Parker, 1999), while other contributions relying on administrative data are not as comprehensive (e.g., Agarwal et al., 2007 exploiting data from a single credit card company only).

Although the work by Gelman et al. (2014) indicates the enormous potential of such data, it also helps us understand the limitations. One of the challenges with the data is selectivity. For instance, Check's users are more likely to be male, young (25-44 years of age), and are less likely to have a graduate degree than the average person living in the U.S. But, they may also differ in other characteristics, such as financial literacy or cognitive and non-cognitive traits, which are relevant for financial decisions. Another issue is that these data provide a partial picture of the financial transactions and balances of a household. The individual links accounts that s/he has with different financial institutions, including joint accounts s/he may share with a partner. However, the data do not include the individual accounts of one's partner and other household members. Finally, financial aggregators do not collect or share demographic information. This greatly limits the range of research questions that can be addressed with such data. Our approach, exploiting rich background information and self-reports at both individual and household level, aims at addressing all three of these limitations.

After the pioneering study by Gelman et al. (2014), several other authors have relied on data from financial aggregators to study a range of topics. These include spending responses to both regular and irregular income (Olafsson and Pagel, 2018) or predetermined payments (Kueng, 2015), interaction between household balance sheets, income, and consumption during the Great Recession (Baker, 2014), changes in household spending in the midst of the COVID-19 pandemic (Baker et al., 2020), and debt repayment behavior (Kuchler, 2015).[1] The data limitations described above, largely stemming from the lack of detailed background information on the individuals for whom financial transactions are observed, apply to these studies. In describing our pilot project, we aim to provide relevant insights on how these limitations could potentially be overcome, to identify the major hurdles associated with the collection of transaction data in a representative sample, and to indicate how the combination of surveys and electronic financial records can open new research avenues to better understand individual financial decision making and well-being.

The remainder of the paper proceeds as follows. In Section 2, we describe the main features of our data collection effort and document respondents' recruitment and participation in the study. In Section 3, we provide descriptive statistics of collected transaction data. In Section 4, we compare self-reported and actual monthly expenditure obtained from transaction data across spending categories. In Section 5, we examine individuals' spending behavior over the pay cycle. Section 6 concludes.

[1] Recently, transaction data have also been used to construct daily estimates of retail spending at the metropolitan statistical area (Aladangady et al., 2019).

2 Combining Survey and Transaction Data

The context for our proposed data collection effort and research project on financial behavior is the Understanding America Study (UAS), a household panel managed by the Center for Economic and Social Research at the University of Southern California. Currently, the UAS has approximately 9,000 members representative of the U.S. adult population. At the time we started this pilot project, the number of UAS members was about 4,500. The UAS is an Internet Panel, which means that respondents answer surveys on a computer, tablet, or smartphone, wherever they are and whenever it is convenient for them to participate. Anyone willing to participate, who does not have a computer or Internet access, is provided with a tablet and broadband Internet.

Panel members answer surveys about once or twice a month. Surveys are restricted to about 30 minutes per interview. Since all data can be linked across surveys, a large amount of information is available about panel members, including demographics, health, financial behavior and financial literacy, labor force status, cognitive capability, and personality.[2]. Respondents receive compensation for their time spent answering questions at a rate of $20 per 30 minutes of interview time. Annual attrition rates are modest, on the order of 6-7% per year.[3]

From the viewpoint of representativeness, it is important to note that the UAS is a probability-based Internet panel. That is, respondents are drawn from a well-defined sampling frame (U.S. Postal addresses) with known inclusion probabilities. Probability Internet panels have to be distinguished from convenience Internet panels, where respondents are recruited from among existing Internet users by placing banners on web-sites to invite respondents,

[2]It is worth noting that all panel members answer the full survey instrument of the Health and Retirement Study (HRS) every 2 years. This contains some 130 minutes worth of interview time (administered over a number of separate sessions to stay within a 30 minute limit per survey) with information on health, income, assets, labor market outcomes, retirement and expectations. Angrisani et al. (2019) examine sample characteristics and elicited survey measures of HRS and UAS.

[3]The UAS webpage (https://uasdata.usc.edu/index.php) provides full details about recruitment, response rates and attrition.

and inclusion probabilities and sampling properties are unknown. Several studies have shown that probability Internet panels and convenience panels differ fundamentally in the quality of information they provide about the U.S. population.[4]

2.1 Aggregator data

Firms such as Mint (https://www.mint.com) and Yodlee (http://www.yodlee.com) provide financial aggregation services to consumers. Individuals signing up for any of these aggregators list their various accounts with financial institutions and share passwords with the aggregator. The software of the aggregator (which reflects agreements with financial institutions) then combines the information from the various accounts and provides overviews of spending in broad categories, use of payment instruments and balances. Users receive weekly overviews and alerts if spending or changes in balances exceed some pre-specified trigger levels. Most major banks provide similar services to their own customers, but by necessity that information is only based on the data contained in the bank's own records. The backend of banks' services is often provided by Yodlee, whose platform is used by 7 of the 10 top U.S. banks, more than 700 global financial institutions, and over 50 million consumers.

Based on Yodlee's capabilities and discussions about the services it can provide, we invited UAS panel members to join Yodlee, enabling the UAS to gain access to the financial information that Yodlee collects. We asked consent from panel members to sign up to Yodlee and for their permission for Yodlee to share their financial information with us. Panel members were incentivized by receiving monetary compensation for their effort. We purposefully decided not to provide participants with any type of feedback about their account balances and expenditure, so to reduce influence on their financial behavior

[4] Chang and Krosnick (2009) administered the same questionnaire to a telephone sample, an Internet probability sample, and a non-probability sample of volunteers who do Internet surveys for money. They found that the telephone sample has the most measurement error, while the non-probability (convenience) sample exhibits most bias. On balance, the probability Internet sample produced the most accurate results. Yeager et al. (2011) document similar findings.

stemming from their participation in the study.

We obtained information on every electronically recorded financial transaction of UAS panel members who signed up for our project directly from Yodlee. For the recruitment (from among UAS members) and maintenance of the sample that signs up with Yodlee, we varied certain parameters experimentally to gauge their effect on participation and data quality. The Yodlee data are supplemented with data from other UAS surveys, both to assess and improve the quality of the collected data and to retrieve information that is particularly powerful for answering research questions in combination with actual transaction data.

2.2 Recruiting respondents

Before this pilot project, which we will refer to as the **UASFin study**, we fielded a consent survey among members of the UAS to gauge willingness to participate in a study involving sharing financial information.[5] Results indicated that about 60% of the respondents were interested in participating. We also varied the level of promised incentives for participation, but found no significant effects. Given the budget available for the pilot project and the anticipated 60% response rate, we decided to invite 1,110 panel members to join the study, expecting to eventually have a sample of about 600-650 respondents. It should be noted that participation involves a number of steps to be taken by the panel members; at each step there is potential for attrition. These steps are:

1. Provide consent to participate.

2. Create an account on the financial management web-site (the web-site was created by us, but connected to the financial aggregator system).

3. Once an account has been created, the respondent needs to add financial institutions to the account, which includes the sharing of passwords with the financial aggregator.

[5] UASFin is also the name we used on the UAS platform for participating members.

4. The accounts need to be kept up-to-date (e.g, if the account password was changed with the bank, the password needs to be changed with the aggregator for it to continue accessing account data).

The 1,100 selected UAS respondents received the following invitation (the amounts varied experimentally across respondents, as will be explained below):

We are interested in how Americans spend their money and how they are doing financially. We would ask you to sign up with a custom made financial management web-site. The web-site has been developed in collaboration with one of the biggest financial management service companies in the world: Yodlee. For instance, Yodlee provides services to 12 of the 20 largest banks in the United States.

We will NOT have access to your passwords or any other identifying information*; this information will be safeguarded by Yodlee. We will use the data in the same way we use surveys you participate in: to make summary tables or graphs to better understand how Americans are doing. Just like the information you provide through surveys, you will be compensated for the information that you share with us.*

If you agree to participate, we will pay you $25 just for signing up with the financial management web-site, plus $5 for every one of your financial institutions that you add on the web-site.

For example, imagine you have a checking and a savings account with one financial institution, a credit card with another, a brokerage account with another financial institution, and a retirement account (such as a 401(K) or IRA) with yet another one. That means you have a total of 5 accounts at 4 financial institutions. You will earn $20 (4 × $5) if you sign all 4 of your financial institutions up.

Every month after that, we will pay you $2 per institution that you signed up to the web-site. That means the earlier you sign up all of your financial institutions, the sooner you can start earning money, just for letting Yodlee summarize information about your accounts for us. You'll get the monthly amount as long as you keep information about each institution current in the *system.*

Of the 1,110 invitees 509 stated they would be willing to take part in the study for a 46% consent rate, which is lower than the anticipated 60% obtained from an earlier hypothetical question. Possibly, some respondents, who initially were interested in potentially participating, changed their mind when they were actually asked to commit. Table 1 compares the characteristics of consenters and non-consenters. As can be seen, consenters tend to be younger and better educated than non-consenters and are less likely to be married. The two groups do not differ significantly in terms of gender, race, work status, or household income.

Once respondents consented, they were asked to create an account on the financial management web-site. Out of the 509 respondents who consented, 355 eventually created an account ("signed-up"). Hence, the unconditional sign-up rate was about 32%. There may be several reasons why respondents who initially consent to participate do not create an account. One reason is that respondents find this more cumbersome than anticipated; another possibility is that the process of creating an account is another opportunity for reflection on perceived risks and, hence, for the decision not to proceed further. While we did not explicitly ask UAS members to state the reasons why they did not consent to take part in the UASFin study or did not link an institution after consenting, we have anecdotal evidence from respondents' survey comments and contacts with the UAS help desk that privacy/security concerns were the main barrier to participation.

The first two columns in Table 2 show marginal effects from Logit regressions of the sign-up indicator on demographic variables. The first column refers to the unconditional sign-up and uses the entire sample. The results confirm the earlier observation that better educated respondents are more likely to create an account with the financial aggregator, while older and married individuals are less likely to sign-up. In the second column, we examine the likelihood of signing-up conditional on consent. The results show that, even among consenters, older individuals are less likely to participate, while individuals with higher education are more likely to participate. In addition to

Table 1: Individual Characteristics (sample proportions)

	Consenters	Non-Consenters	Δ sig
Male			
	41.26	44.42	
Non-White			
	24.16	24.96	
Age			
18-34	21.81	16.47	**
35-44	24.36	16.64	***
45-54	17.68	17.64	
55-64	21.81	26.12	*
65+	14.34	23.13	***
Education			
High School or Less	23.97	29.62	**
Some College	40.86	36.94	
Bachelor or More	35.17	33.44	
Married			
	54.03	64.23	***
Working			
	40.67	44.09	***
Income			
$< \$30k$	33.01	27.79	*
$\$30k - \$60k$	24.75	28.95	
$\$60k - \$99k$	22.40	25.12	
$\$100k+$	19.84	18.14	
N	509	601	

Note: ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.

demographics, one would expect study participation to be related with online banking experience. The regression results in Table 2 reveal that to be true. Respondents who are active in internet banking, either by paying bills, checking account balances or making transfers online, are more likely to create an account with the financial aggregator.

Once respondents have signed up, they need to add financial institutions to their dashboard. This is a critical step: without adding institutions, which implies sharing username and password with the financial aggregator, no transactional data can be retrieved. Out of the 355 respondents who created an

Table 2: Determinants of Participation: Demographics

	Unconditional Sign-up	Sign-up Conditional on Consent	Linking Institutions
Male	0.003	-0.029	-0.028
	(0.029)	(0.043)	(0.053)
Non-White	-0.058*	-0.053	-0.111*
	(0.031)	(0.050)	(0.063)
Age 35-44	-0.012	-0.045	-0.029
	(0.046)	(0.056)	(0.072)
Age 45-54	-0.063	-0.038	-0.030
	(0.048)	(0.060)	(0.078)
Age 55-64	-0.152**	-0.144**	0.009
	(0.044)	(0.061)	(0.083)
Age 65+	-0.192**	-0.192**	-0.114
	(0.048)	(0.080)	(0.092)
Some College	0.081**	0.139**	0.081
	(0.034)	(0.057)	(0.071)
Bachelor or More	0.094**	0.174**	0.175**
	(0.039)	(0.062)	(0.079)
Married	-0.083**	0.006	0.003
	(0.031)	(0.044)	(0.055)
Working	-0.035	-0.045	-0.024
	(0.033)	(0.046)	(0.060)
HH Income 30-60k	-0.041	0.006	0.099
	(0.038)	(0.057)	(0.074)
HH Income 60-100k	-0.018	0.044	0.110
	(0.042)	(0.063)	(0.078)
HH Income 100k+	0.036	0.104	0.143
	(0.051)	(0.067)	(0.087)
Internet Banking Activity	0.187**	0.214**	
	(0.042)	(0.104)	
N	1110	509	355

Note: ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.

account, only 135 actually linked at least one financial institution. The results in the third column of Table 2 show that non-White respondents are less likely to add a financial institution. The likelihood of linking a financial institution is about 18 percentage points higher for individuals with at least a Bachelor's degree. Among the 135 participants who linked a financial institution, everybody has had experience with online banking activity by either paying bills, or checking account balances, or transferring money over the internet. When adding financial literacy and cognitive ability to the set of regressors, we note that they are positively correlated with the conditional and unconditional probabilities of signing-up, but are uncorrelated with the likelihood of linking institutions.

2.3 Incentive experiments

Various factors need to be considered when designing an incentive scheme for UAS panel members to join a study that involves sharing information with a financial aggregator. The incentives should be large enough to make it attractive for panel members to participate, with minimal distortion of reporting behavior. In order for respondents to sign-up with Yodlee, a one-time incentive is appropriate. However, since participants should provide as complete information as possible, the incentive should be higher as more financial institutions (and therefore accounts) are linked. A third consideration has to do with attrition versus change. If respondents change a password for any of their accounts, if they move an account to another, new institution, or if they open a new account with a new institution, those changes need to be communicated to the financial aggregator. Hence, an incentive to keep information and linkages up-to-date is important. We tested two incentive schemes, each characterized by different incentive combinations as summarized in Table 3. We adopted Scheme I for the first batch of 110 invited members. Under this scheme, the sign-up incentive is relatively generous and can be \$10, \$25 or \$50, while adding an institution was rewarded with either \$2 or \$5. The incentive to maintain the information up-to-date was either \$1 or \$2. Given the low fraction of individuals who had linked a financial institution after creating an account within the first batch, we decided to move to Scheme II for the other 900 invitees, by decreasing the sign-up incentive to \$5 and increasing the reward for adding an institution up to \$15.

Table 3: Incentive Schemes

Scheme I			
	Sign-Up	Add Institution	Monthly Payment
Treatment 1	$10	$2	$1
Treatment 2	$25	$2	$1
Treatment 3	$50	$2	$1
Treatment 4	$10	$5	$2
Treatment 5	$25	$5	$2
Treatment 6	$50	$5	$2
Scheme II			
	Sign-Up	Add Institution	Monthly Payment
Treatment 1	$5	$3	$1
Treatment 2	$5	$5	$1
Treatment 3	$5	$10	$1
Treatment 4	$5	$5	$2
Treatment 5	$5	$10	$2
Treatment 6	$5	$15	$2

We have found no clear pattern suggesting that different monetary incentives may produce different participation behavior. Specifically, among the 110 respondents subject to incentive Scheme I, the probability of creating an account with Yodlee does not vary significantly with whether the sign-up monetary reward is $10, $25, or $50, being 36%, 36% and 33%, respectively. For the 35 respondents in Scheme I who created an account with Yodlee, the likelihood of linking a financial institution is 53% when the monetary incentive to add an institution is $5 and 39% when the monetary incentive to add an institution is $2. This difference, however, is not statistically significant, which by itself may just reflect the small sample sizes. Among the 320 individuals subject to incentive Scheme II who decided to participate in the UASFin study, the probability of adding a financial institution is 37% when the incentive to add an institution is $3, 39% when the incentive is $5, 33% when the incentive is $10, and 42% when the incentive is $15. We have tested all possible combinations pair-wise and found no evidence of statistically significant differences across different monetary rewards. Overall, the outcomes of these analyses suggest the existence of barriers to participation, such as privacy and security concerns, which monetary incentives are unlikely to help overcome.

3 Descriptive Statistics of Transaction Data

Before proceeding with the analysis of the transaction data, we take the following sample selection steps. We start with the 135 respondents who have added financial institutions. We drop all account entries with missing transaction date or transaction amount (0.25% of the entire data set). In doing this, we lose one respondent who linked accounts for which the data were never retrieved. Finally, we eliminate accounts with less than 10 observations or with data covering less than a month and 1 checking account with persistently large negative balance (more than $100,000) over the observation period. After this, we are left with 130 individuals whose accounts are observed from September 1, 2016 through September 1, 2018. All analyses in this paper are based on this sample. Not all these 130 participants linked the three types of account we will be focusing on, namely checking, savings, and credit card accounts. Thus, the samples for analyses by account type may not include all UASFin participants. Figure 1 illustrates what individual data look like. It shows the balance and the transaction amount of a checking account of a study participant over the observation period.

We begin by comparing the features of key outcome variables in our sample with those reported by Gelman et al. (2014). In Table 4, the total number of accounts per respondent is similar across the two studies, albeit slightly lower in our study. The breakdown by account type reveals that the number of checking accounts is similar. The main differences are observed in the number of savings accounts, which is larger in our sample than in the Gelman et al. (2014) data set, and in the number of credit cards, which is substantially larger in Gelman et al. (2014) than in our sample. As far as the number of daily transactions is concerned, it is apparent that individuals in the Gelman et al. (2014) data set are more active. These observed differences are likely due to selection. The Gelman et al. (2014) study is based on data from customers

Figure 1: A Respondent's Checking Account Over Time

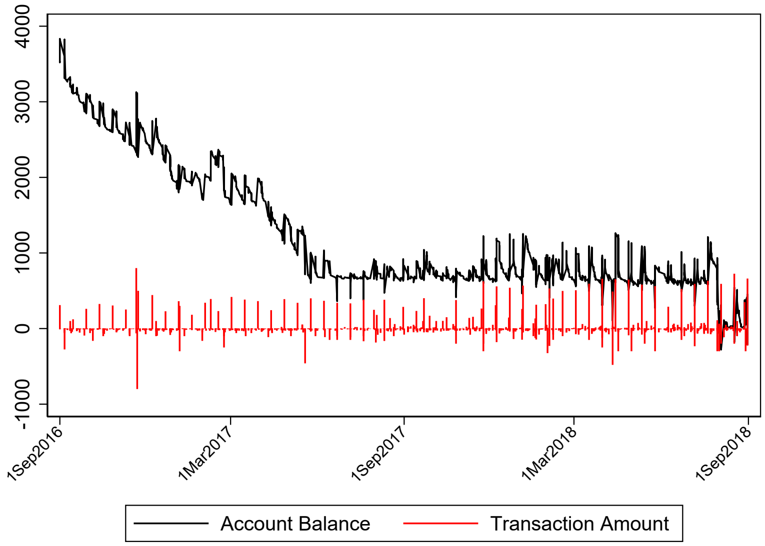

of Check (https://check.me).[6] Individuals who decide to sign-up with a financial aggregator are likely more financially savvy and probably feel a greater need for the services of an aggregator, due to the complexity of their financial portfolio and number of transactions. The comparison of account balances in the bottom panel of Table 4 reflects the similarities and differences described above. Specifically, checking account balances appear comparable in the two studies, while individuals in our sample keep twice as much in savings accounts compared to those in Gelman et al. (2014) data set. Both credit card balances and limits are substantially larger in Gelman et al. (2014).

Unlike Gelman et al. (2014), and other studies relying exclusively on transaction data from financial aggregators, we have very rich background information about study participants that has been collected across several UAS surveys. Among these, the 2016 and 2017 Survey of Consumer Payment Choice (SCPC), designed by the Federal Reserve Bank of Boston, are ideal for assessing coverage of one's financial accounts through transaction data. The SCPC asks about ownership and number of individually owned checking, savings,

[6] In 2014, Mint owner, Intuit, purchased Check. Since Spring 2017, Mint and Check are integrated in the Mint platform.

Table 4: Number of Accounts, Daily Transactions, and Balances

	Gelman et al. (2014)		UASFin Study	
	Mean	Median	Mean	Median
Number Accounts				
Total	5.84	5	5.18	4
Checking	1.35	1	1.72	1
Savings	0.79	1	1.78	1
Credit card	3.58	3	2.50	2
Number of Daily Transactions				
Total	4.54	3	2.14	2
Checking	3.03	2	0.97	0
Savings	0.22	0	0.08	0
Credit card	1.23	1	0.77	0
Account Balances				
Checking	6,969	1,400	6,981	1,296
Savings	4,476	400	8,634	1,060
Credit card	7,228	3,600	3,146	1,587
Credit limit	23,019	11,900	10,023	7,500

Note: Number of individuals is 130 for Total, 80 for Checking, 50 for Savings, and 108 for Credit Card accounts. Number of observed transactions is 72,006. Number of individual-day observations over which balances are computed is 52,221.

and credit card accounts. The 2016 and 2017 SCPCs were administered in the UAS in October 2016 and October 2017 respectively, thereby providing information about financial accounts that is contemporaneous to the time when transaction data were collected. Hence, we can compare the extent to which self-reported information about financial accounts matches what is available in the transaction data.

We have self-reported account information for 119 out of the 130 UASFin active participants. Of these, 118 reported owning a checking account in the SCPC, with 71 (60%) linking at least one checking account. As shown in Table 5, among the 49 participants who reported owning just one checking account in the SCPC, 30 (61%) linked one checking account and 19 (39%) linked two or more checking accounts. Among the 22 participants who reported owning two or more checking accounts in the SCPC, 10 (45%) linked only one checking account and 12 (55%) linked two or more checking accounts. The number of individuals who reported owning a savings account in the SCPC is 100, with 42 (42%) adding at least one savings account in the financial aggregator. Three participants linked one savings account while reporting not owning one in the SCPC. The number of individuals who reported owning one savings accounts in the SCPC is 21, with 14 (67%) linking one savings account and 7 (33%) linking two or more savings accounts. Similarly, 21 participants reported owning 2 or more savings accounts in the SCPC. Among them, 14 (67%) linked one savings account and 7 (33%) linked two or more savings accounts. As far as credit card information is concerned, 115 reported having a credit card in the SCPC, with 100 (87%) linking at least one. Among the 12 participants who reported owning just one credit card in the SCPC, 11 (92%) linked one account and 1 (8%) linked two or more accounts in the UASFin study. Among the 88 participants reporting having more than one credit card, 25 (28%) linked just one credit card account and 63 (72%) linked two or more credit cards with the financial aggregator. Overall, this comparison shows very good coverage of credit card accounts, while coverage of checking accounts is less satisfactory. Individuals' reluctance to link checking accounts compared to credit cards may be partially explained by better fraud protection offered by the latter.

In Table 6, we take further advantage of individual-level background information at our disposal and present a breakdown of median account balances by education (less than a college degree/college degree or more), financial literacy (low/high as defined by below/above median financial literacy score in the sample), and cognitive ability (low/high as defined by below/above median cognition score in the sample). Median checking account balances increase monotonically with education, financial literacy, and cognition. Individuals with lower levels of education, financial literacy and cognitive ability tend to maintain a higher balance in their savings accounts compared to their counterparts. Credit card limits are twice as large for individuals with a college degree compared to those without a college degree and for individuals with financial literacy and cognitive ability above the median compared to those with scores

Table 5

Checking			
Accounts in the SCPC	Accounts linked 1	Accounts linked 2+	Total
1	30 (61%)	19 (39%)	49 (100%)
2+	10 (45%)	12 (55%)	22 (100%)
Savings			
Accounts in the SCPC	Accounts linked 1	Accounts linked 2+	Total
0	3 (100%)	0 (0%)	3 (100%)
1	14 (67%)	7 (33%)	21 (100%)
2+	14 (67%)	7 (33%)	21 (100%)
Credit Cards			
Accounts in the SCPC	Accounts linked 1	Accounts linked 2+	Total
1	11 (92%)	1 (8%)	12 (100%)
2+	25 (28%)	63 (72%)	88 (100%)

below the median. Interestingly, median credit card balances are virtually the same for college and non-college graduates. In contrast, they are between \$300 and \$800 higher for participants with below median financial literacy and cognitive ability scores compared with those with above median scores. This suggests differences in financial decision-making and use of credit related to an individual's level of financial knowledge and cognitive functioning.

Table 7 shows medians of monthly income, monthly total expenditure and a few expenditure components.[7] We calculate a measure of regular income, including salary, earnings, and income from pensions and annuities. We add to this measure investment and interest income as well as other income to

[7]For each study participant, we sum all amounts classified as income/expenditure within a month to obtain monthly income/expenditure measures. Thus, if at least one income/expenditure transaction is observed each month, each study participant would have 24 data points, corresponding to the 24 months throughout the observation period. In Table 7, we report median values across all available data points obtained in this fashion.

Table 6: Median Account Balances by
Education, Financial Literacy and Cognitive Ability

	Education		Fin. Literacy		Cog. Ability	
	NoCol	Col	Low	High	Low	High
Checking	775	1,726	997	1,854	1,149	1,515
Savings	1,420	800	1,215	756	1,234	650
Credit card	1,531	1,630	1,888	1,274	1,970	1,213
Credit limit	5,500	10,000	5,300	10,000	6,000	10,000

Note: Number of individual-day observations over which balances are computed is 20,480 for Checking, 4,943 for Savings, and 26,798 for Credit Card accounts. Individuals without and with a college degree are 34% and 66% of the sample, respectively.

obtain total income. When computing total expenditure, we exclude outlays associated with service fees for retirement and investment accounts, mortgages and loans. Since we are mainly interested in a measure of recurrent monthly expenses to be compared with monthly income, we also calculate total expenditures without check payments. The reason for doing this is that, within our sample, the adoption of checks as a method of payment is by an order of magnitude lower than the adoption of debit, cash and credit card, and typically observed for large, one-off purchases (Greene and Stavis, 2018). Given that the number of individuals who linked credit cards is more than double the number of individuals who linked checking accounts, we have many more data points for expenditure than for income. This implies that the composition of the samples with expenditure and income information is likely very different, preventing a meaningful comparison of income and expenditure to infer saving rates. We, therefore, restrict attention to the sub-sample of study participants with both income and expenditure data. This leaves us with 569 observations for both income and expenditure.

As can be seen from the first column of Table 7, median monthly expenditure is about $700 greater than median total monthly income ($5,145 vs $4,411), but when we exclude check payments it becomes $300 lower ($4,126 vs $4,411). Median values of expenditures on sub-categories appear to be in line with official statistics, although differences in how expenditures are categorized and in the consumer unit expenditures refer to prevent a precise comparison. For example, median monthly expenditure on grocery in the UASFin study is $320, while the mean value of food at home obtained from the 2018 CE is about $380 per month.[8] The median value of monthly restaurant expenditure in our study is $259, while the 2018 CE reports an average value of food away from home of $290 per month. UASFin participants spend $180 at the median in automotive-related expenditures, which mainly include gasoline and maintenance. In the 2018 CE, average spending on gasoline and vehicle maintenance is $250 per month. Given how the merchandise category is defined in our study, a meaningful comparison with the 2018 CE is not possible in this case.

Table 7 also shows median monthly income and expenditures by education (no-college/college), financial literacy (low/high defined as below/above median score) and cognitive ability (low/high defined as below/above median score). Income and expenditures are higher at the median for respondents with more education and higher levels of financial literacy and cognitive ability. Differences are more sizeable between individuals with and without a college degree than between individuals below and above median financial literacy and cognition scores. Total expenditures remain larger than total income across groups, although the difference between these two measures tends to be smaller for respondents with higher education and cognitive ability. This may partly reflect a more comprehensive capturing of income sources through transaction data for these groups. Total expenditures excluding check payments are below total monthly income to a much greater extent for better educated individuals and those with higher levels of financial ability and cognition, suggesting relatively higher saving rates for these groups.

In the last column of Table 7, we further restrict the sample to those individuals with most or all responsibility for shopping and paying bills within the household, which we identify using self-reported information from the SCPC. In this case, we have 260 income and expenditure data points at our

[8] Average expenditures from the 2018 CE can be retrieved at `https://www.bls.gov/cex/csxcombined.htm#2018`.

Table 7: Median Monthly Income and Expenditures

	All	Education		Fin. Literacy		Cog. Ability		FinResp
		NoCol	Col	Low	High	Low	High	
Reg Inc	4,024	1,773	4,641	3,322	4,160	3,200	4,328	3,779
Tot Inc	4,411	2,481	5,420	3,825	4,745	3,696	4,816	4,328
Tot Exp	5,145	3,076	5,803	4,426	5,486	5,024	5,227	4,321
$(\text{TotExp})_{nocheck}$	4,126	2,297	4,718	3,749	4,329	3,788	4,260	3,524
Grocery	320	229	333	300	343	372	282	280
Restaurant	259	146	313	276	235	242	266	203
Merchandise	545	506	554	549	533	504	573	482
Automotive	180	98	196	181	179	235	153	157

Note: Study participants with both salary/regular income and expenditure data only (Obs = 569). In the last column (FinResp), we consider the sub-sample of those who have most or all responsibility for shopping and paying bills within the household (Obs = 260). Reg Inc is salary and regular income; Tot Inc also includes investment and retirement income, interest income, and other income. Tot Exp includes all expenditures except service fees associated with retirement and investment accounts, mortgages and loans. $(\text{Tot Exp})_{nocheck}$ is total expenditure without check payments. Merchandise comprises purchases of different goods typically at large retailers such as Amazon, Walmart, Target, etc.

disposal. For these individuals, we would expect a larger share of household expenses to be paid using linked accounts and, therefore, the comparison monthly income/monthly expenditure to be more reliable. Within this group, total monthly income and expenditure are virtually the same. Excluding check payments, total monthly expenditure is about \$800 lower than total income and about \$250 lower than regular income. This would imply saving rates of 18% and 7%, respectively.

4 Comparison of Self-Reported and Transaction Data Expenditures

A comprehensive expenditure module was administered in the UAS between August and November 2016 to all active members at that time. The same respondents were invited to answer the module one more time throughout 2017. In the same year, the expenditure module was administered to all newly recruited UAS members. As a result of this, we have at our disposal self-reported expenditure measures for all 130 individuals in our sample that are roughly contemporaneous to the time transaction data were obtained.

The UAS expenditure module elicits expenditure amounts on a wide range of spending categories at the household level with reference to the previous calendar month. We do observe the date when the expenditure module was answered by each respondent. Hence, we match the self-reported expenditure measures with the transaction data of the calendar month prior to the month when the expenditure module was answered (provided that transaction data for that month exist). Out of the 130 respondents in our original sample, we are able to match self-reported and transaction expenditure measures for 115.[9]

We compare self-reported and electronically recorded expenditures and evaluate distance between these two measures across expenditure categories. This type of exercise, which we merely use to illustrate the potential use of transaction data within the UAS, may be particularly valuable to determine the reliability of self-reports in different spending domains. Before proceeding, though, it is worth pointing out the limitations of this analysis. First, we cannot treat the transaction data expenditures as the “true" level of expenditures. Since we do not know how many accounts we are missing for each individual, we cannot determine what part of spending is missing in the transaction data. Second, self-reported and transaction expenditures are conceptually different. The former are expenditures at the household level, the latter are electronic transactions observed in an individually owned account. Even if we had complete coverage of an individual’s accounts, observed transaction expenditures would coincide with household level expenditure only if all household purchases were made by the study participant. Because of this, self-reported measures may systematically exceed transaction data measures across spending categories, with the possible exception of one-person households.[10] Third, the

[9]Some individuals answered the UAS expenditure module prior to joining the UASFin study. For a few respondents, no expenditure information could be retrieved from transaction data (e.g., they only linked retirement or brokerage accounts). About one-third of these 115 respondents are observed twice (i.e. we are able to match two self-reports). The results that follow are unaffected by whether we use weights that account for this or not.

[10]About 32% of UASFin study participants do not cohabit with other household members.

transaction categorization extracted from electronic records is coarser than the one available in the UAS expenditure module. For some spending categories, such as groceries, restaurant, utilities, and mortgage payments, there exists an immediate and exact correspondence between the two sources of data. For others, like automotive, entertainment, and insurance payments, it is harder to achieve a good match. Thus, observed differences between self-reports and transaction data may reflect differences in the composition of the categories that are being compared. To reduce this issue, we will focus on grocery and restaurant expenditures, which are the two most common types of expenditures in the transaction data (together with merchandise expenditures) and the two most directly comparable to their self-reported counterparts.

For these two spending categories, we compare the self-reported expenditure measure for a given month with the same-month expenditure measure taken from transaction data, as well as with the average monthly expenditures over the entire observation period, again calculated using transaction data. The idea is to shed light on the cognitive process behind self-reports, whether it is an episodic memory process, by which individuals try to recall what they actually spent in the previous calendar month, or a rate-based estimation process, by which they try to compute an average over a few months and apply it to the previous calendar month.

Figures 2 and 3 show scatter plots of transaction data expenditures (y-axis) against self-reported expenditures (x-axis), using same-month expenditures from transaction data (upper panel) and average monthly expenditure from transaction data (bottom panel). Self-reported household-level grocery expenditures are systematically larger than transaction data outcomes. In view of our discussion above, this should not be surprising. Median/mean self-reported grocery expenditure is \$350/\$414 versus \$160/\$277 for same-month grocery expenditure from transaction data. The correlation between these two measures is 0.4. The distance between self-reports and transaction data is slightly reduced when we consider average monthly expenditure over the transaction data observation period. Specifically, median/mean average monthly grocery expenditure from transaction data is \$223/\$308 and the cor-

Figure 2: Grocery Expenditure Comparison

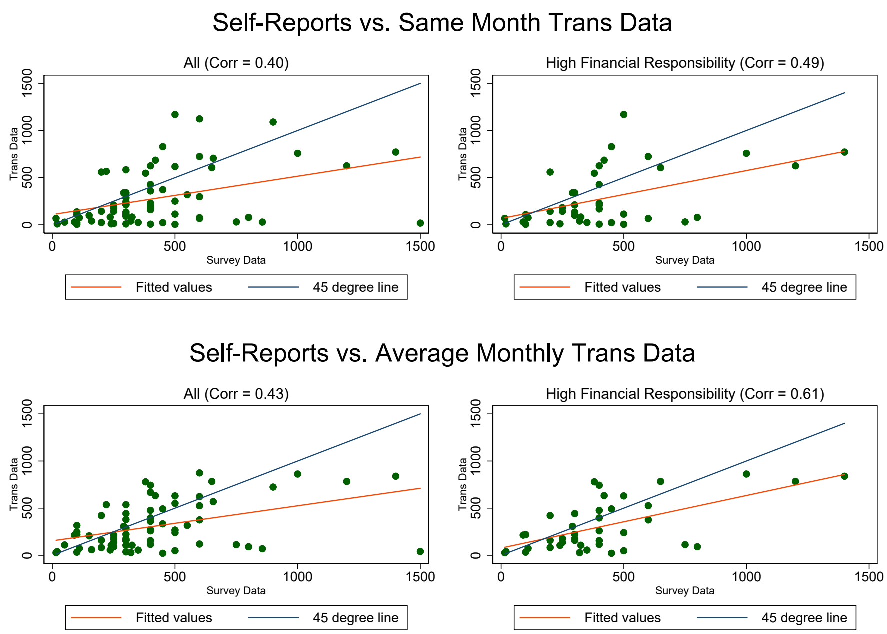

relation between self-reported and average transaction data expenditures is 0.43. When we restrict attention to the sub-sample of individuals who have most or all responsibility for shopping and paying bills within the household, the alignment between self-reported and transaction data measures improves significantly. The correlation coefficient is 0.49, when using same-month expenditure, and 0.61, when using average monthly expenditure.

Self-reported and transaction data restaurant expenditures are relatively close, with a tendency of the latter to be slightly larger. Median/mean self-reported restaurant expenditure is $150/$187 versus $134/$207 and $184/$256 for same-month and average monthly restaurant expenditure from transaction data, respectively. Figure 3 shows a rather close alignment between self-reports and transaction data in the lower half of the distribution, while in the upper half of the distribution, values from transactional data systematically exceed self-reported ones. In the entire sample, the correlation between self-reports and same-month expenditures is 0.42, whereas the correlation between self-reports and average monthly expenditure is 0.55. The latter correlation increases to 0.7 when we only consider individuals with high financial responsibility within the household who are expected to be relatively more knowledgeable about expenses.

Figure 3: Restaurant Expenditure Comparison

Overall, self-reported expenditures show less variation and less skewness than transaction data measures. A plausible explanation for this is that self-reports over longer periods, like a month, partly reflect "normal" or "typical" expenditures, rather than precise fluctuations from month to month. This conclusion is backed by the observed closer alignment between self-reports and average monthly expenditures over the period covered by our transaction data. The closer correspondence between self-reports and transaction data measures when the sample is restricted to individuals who have most or all responsibility for shopping and paying bills within the household is reassuring. With a more comprehensive coverage of an individual's accounts through transaction data and a specifically designed questionnaire asking about spending categories that match those retrievable from transaction data, this exercise, which we have performed mainly for illustrative purposes, can be very valuable to assess the reliability of expenditure measures via surveys.

5 Expenditure Smoothing

Of great interest in economics is the extent to which households smooth consumption over time, especially over pay cycles, as predicted by a standard model with rational, forward-looking agents. This issue has been investigated using transaction data (Gelman et al., 2014; Olafsson and Pagel, 2018). We revisit it exploring heterogeneity by education and financial literacy, an exercise that once again exploits the rich background information at our disposal alongside recorded electronic transactions. For this purpose, we consider the subset of respondents with clearly identifiable paydays and investigate to which extent expenditures are different before and after payday. We do this separately for individuals with no college, with college, and with more than college degree, as well as by tercile of the financial literacy score in our sample. For now, this is just an illustrative exercise, given the limited size of the sample underlying this analysis. Also, expenditures do not need to coincide with consumption. Moreover, for larger items (e.g. rent), payment patterns may be synchronized with income receipt, without affecting consumption (housing services, in this example).

The subsequent analysis requires a number of assumptions and data manipulations. We focus on transactions classified as “salary and regular income." First, for each individual we identify a “payment date." That is, a date when a transaction classified as salary/regular income is recorded. If multiple payment dates exist within a two-week period, we ignore those for which the amount of salary/regular income received is less than 25% of the maximum amount received within the two-week period. Second, we compute the distance between two subsequent payment dates. Third, we define the individual-specific frequency of payment as the within-individual median of distances between two successive payment dates.

Out of 41 respondents for whom we observe salary and regular income payments over time, 3 are paid every 15 days, 24 every 14 days, and 5 every 13 days. In our analysis, we will only use these 32 individuals, assuming they are paid bi-weekly. As far as the others are concerned, 4 are paid every 7 days, 2 every 8 days, 1 every 10 days, and 2 every 30 days.

We observe these 32 individuals over the period September 1, 2016 – September 1, 2018. We consider three types of expenditures: grocery, restaurant, and merchandise. The latter may comprise a variety of goods. The transaction description for this category typically features big retailers such as Amazon, Walmart, Target, etc. For each individual we compute average expenditures over the observation period and then take the ratio of each date's expenditure to the average expenditure. This ratio, defined for each expenditure category, is the dependent variable in our analysis.

We run regressions of expenditure ratios on a set of dummies for the 7 days before the payment date, the payment date itself, and the 6 days after the payment date (and before the next pay cycle starts). The omitted dummy is the one for the day before payment day. We also include day-of-the-week dummies. Standard errors are clustered at the individual level. The figures in this section show the estimated coefficients from these regressions alongside with their 90% confidence intervals.

Figure 4 reports the estimates for grocery expenditures by education and financial literacy. Given the way the coefficients are estimated, the dependent variable is always exactly zero on day -1. Because of the very limited sample size, coefficients are not precisely estimated. Yet, it is apparent that expenditure variability is substantially higher for individuals with no college degree and for those in the bottom third of the financial literacy score distribution. Except for individuals with more than a college degree and with the highest level of financial literacy, the day before payday shows the lowest spending level and spikes in grocery expenditure, often statistically significant, are observed in the subsequent days.

As shown in Figure 5, restaurant expenditures exhibit similar patterns and heterogeneity across groups. Specifically, individuals with the lowest level

Figure 4: Grocery Expenditures - Fraction of Daily Average Spending

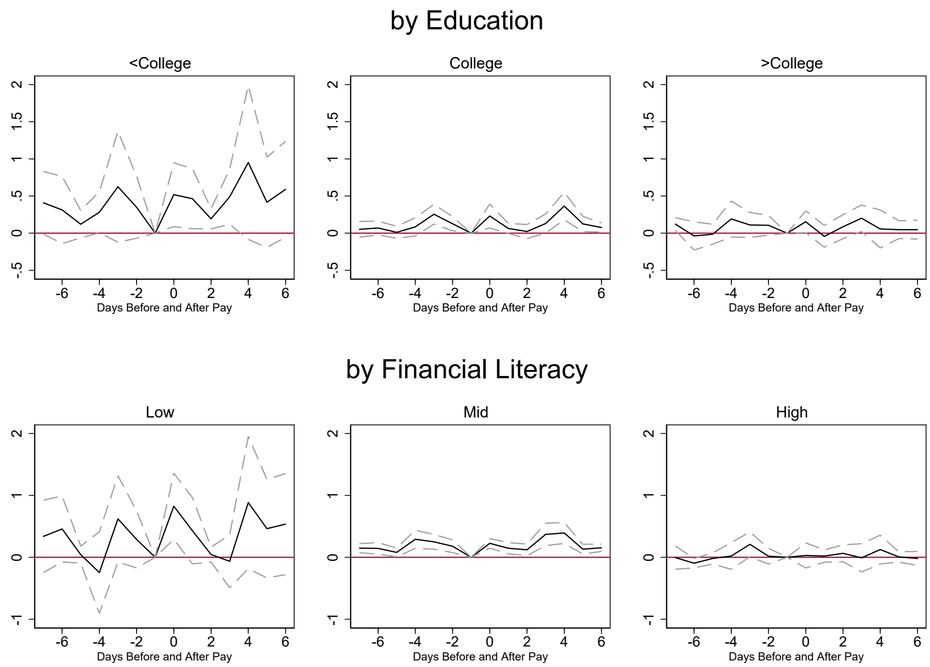

of education seem to have the highest expenditure variability, but also those in the middle of the education and financial literacy score distributions tend to decrease outlays right before payday and to increase them right after it. For merchandise (Figure 6), there seems to be a clear tendency to increase expenditures after payday, but again those who show the clearest cyclical pattern over the pay cycle are individuals without a college degree and in the bottom tercile of the financial literacy score distribution.

6 Conclusions

This paper describes a first attempt at collecting electronic transaction data in a population representative household panel and use them in combination with survey data for research purposes. This pilot project reveals a number of challenges and provides valuable insights for future data collection efforts of

Figure 5: Restaurant Expenditures - Fraction of Daily Average Spending

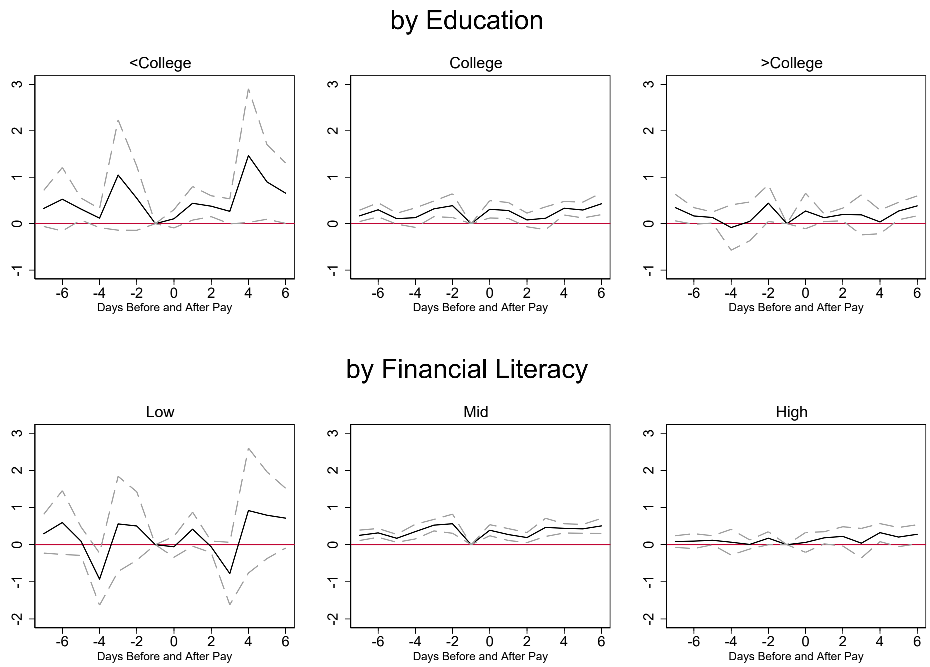

this kind.

The participation rate is relatively low at about 12%. Important determinants of participation are respondents' willingness and/or skills to register their accounts. The biggest hurdle is represented by privacy and security concerns, while monetary incentives have little or no effect. Even among active study participants, coverage of all accounts is not necessarily guaranteed and possible omission of some accounts may hinder certain types of analysis (e.g., a study of individual earnings requires observing the checking account where labor income is received). We can retrieve data from individually owned accounts. The extent to which this information is representative of the household's financial situation and behavior depends on household composition as well as on the financial responsibility and involvement in financial decisions of different household members.

Our approach of recruiting study participants from a representative panel

Figure 6: Merchandise Expenditures - Fraction of Daily Average Spending

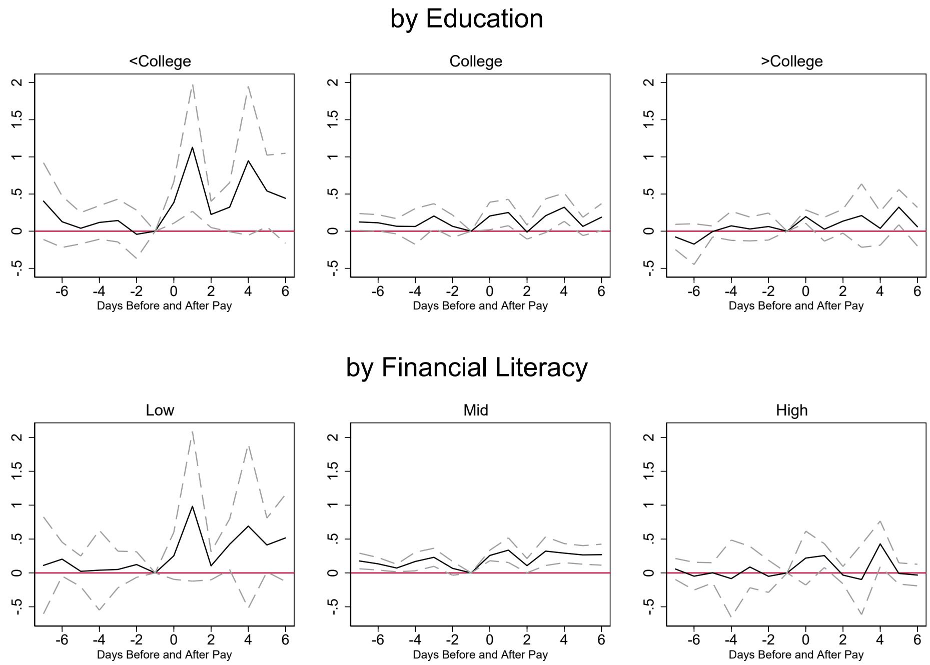

like the UAS has several advantages and offers ways to reduce the importance of the aforementioned limitations. Given the wealth of individual-level information at our disposal, we can determine selectivity of the resulting sample in great detail. This is not possible in typical data sets from financial aggregators comprising only transaction data. By eliciting appropriate information through survey questions, we can evaluate coverage of a household's accounts and complement transaction with survey data to obtain a more comprehensive picture of a household's financial situation than what can be obtained from transaction data alone. Also, knowing the role played by each study participant within the household in terms of expenditure, saving, and investment decisions, we can establish the extent to which the available transaction data tell us about a household's overall financial behavior.

As demonstrated by the exercises carried out in the paper, the combination of transaction and survey data opens the possibility to pursue interesting and novel research questions. For example, the comparison of electronically recorded expenditures with self-reports can provide valuable insights about the quality and validity of individual reports. This is an important methodological issue with relevant implications for the design of consumer spending surveys and their use in policy analysis. Also, the ability to investigate spending behavior separately for different demographic groups provides a unique opportunity to gauge what shapes financial decisions across households and what barriers individuals face when it comes to saving and planning for the future. This knowledge is crucial for devising and implementing interventions that can effectively improve financial decision-making and financial well-being of households, particularly among more disadvantaged and vulnerable segments of the population.

Based on the results of this pilot study and the lessons we have learned from it, we have recently developed a new platform for collecting a new wave of transaction data in the UAS. We have invited the entire pool of UAS respondents (about 8,000) to share their electronic financial records and should soon have at our disposal a larger and more complete data set with both transaction and survey information.

References

Agarwal, S., C. Liu, and N. Souleles (2007). The reaction of consumer spending and debt to tax rebates–evidence from consumer credit data. *Journal of Political Economy 115*, 986–1019.

Aladangady, L., S. Aron-Dine, W. Dunn, L. Feiveson, P. Lengermann, and C. Sahm (2019). From transactions data to economic statistics: constructing real-time, high-frequency, geographic measures of consumer spending. *NBER Working Paper No. 26253.*

Angrisani, M., B. Finley, and A. Kapteyn (2019). Can internet match high quality traditional surveys? Comparing the Health and Retirement Study and its online version. In K. Huynh, D. Jacho-Chavez, and G. Tripathi (Eds.), *Advances in Econometrics, Vol. 39. The Econometrics of Complex Survery Data: Theory and Applications.* Emerald Publishing Limited.

Angrisani, M., A. Kapteyn, and S. Schuh (2014). Measuring household spending and payment habits: The role of “typical" and “specific" time frames in survey questions. In C. Carroll, T. Crossley, and J. Sabelhaus (Eds.), *Imrpoving the Measurement of Consumer Expenditures. Chapter 15.* Chicago: University of Chicago Press.

Attanasio, O. P. (2000). Consumption. In J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, pp. 741–812. Amsterdam: Elsevier Science.

Baker, S. R. (2014). Debt and the consumption response to household income shocks. *mimeo*, Stanford University.

Baker, S. R., R. A. Farrokhnia, S. Meyer, M. Pagel, C. Yannelis, and J. Pontiff (2020). How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic. *The Review of Asset Pricing Studies 14*.

Bound, J., C. Brown, and N. Mathiowetz (2001). Measurement error in survey data. *Hanbook of Econometrics 5*, 3705–3843.

Bricker, J., A. B. Kennickell, K. B. Moore, and S. J. (2012). Changes in U.S. family finances from 2007 to 2010: Evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin 98*, 1–80.

Browning, M., T. F. Crossley, and G. Weber (2003). Asking consumption questions in general purpose surveys. *The Economic Journal 113*, 540–567.

Browning, M., T. F. Crossley, and J. Winter (2014). The measurement of household consumption expenditure. *Annual Review of Economics 6*, 475–501.

Carroll, C., T. Crossley, and J. Sabelhaus (2014). *Imrpoving the Measurement of Consumer Expenditures*. Chicago: University of Chicago Press.

Chang, L. and J. A. Krosnick (2009). National surveys via RDD telephone interviewing versus the internet: Comparing sample representativeness and response quality. *Public Opinion Quarterly 73*, 641–678.

Gelman, M., S. Kariv, M. Shapiro, D. Silverman, and S. Tadelis (2014). Harnessing naturally occurring data to measure the response of spending to income. *Science 345*, 212–215.

Greene, C. and J. Stavis (2018). The 2016 and 2017 surveys of consumer payment choice: Summary results. *Federal Reserve Bank of Boston, Research Data Reports 18*.

Hurd, M. (1999). Anchoring and acquiescence bias in measuring assets in household surveys. *Journal of Risk and Uncertainty 19*, 111–136.

Hurd, M. and S. Rohwedder (2009). Methodological innovations in collecting spending data: The HRS Consumption and Activities Mail Survey. *Fiscal Studies 30*, 435–449.

Kapteyn, A., P. Kooreman, and R. J. Willemse (1988). Some methodological issues in the implementation of subjective poverty definitions. *Journal of Human Resources 23*, 222–242.

Kuchler, T. (2015). Sticking to your plan: Empirical evidence on the role of present bias for credit card paydown. *mimeo*, New York University.

Kueng, L. (2015). Explaining consumption excess sensitivity with near-rationality: Evidence from large predetermined payments. *NBER Working Paper No. 21772*.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*, 1–32.

Menon, G. (1997). Are the parts better than the whole? The effects of decompositional questions on judgments of frequent behaviors. *Journal of Marketing Research 34*, 335–346.

Moore, J. C., L. L. Stinson, and E. J. Welniak (1999). Income reporting in surveys: Cognitive issues and measurement error. In M. Sirken, D. Herrmann, S. Schechter, N. Schwarz, J. Tanur, and R. Tourangeau (Eds.), *Cognition and Survey Research*. New York: Wiley.

Olafsson, A. and M. Pagel (2018). The liquid hand-to-mouth: Evidence from personal finance management software. *Review of Financial Studies hhy055, https://doi.org/10.1093/rfs/hhy055.*

Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic Review 89*, 959–973.

Souleles, N. (1999). The response of household consumption to income tax refunds. *American Economic Review 89*, 947–958.

Winter, J. (2014). Response bias in survey-based measures of household consumption. *Economics Bulletin 3*, 1–12.

Yeager, D. S., J. A. Krosnick, L. Chang, H. S. Javitz, M. S. Levendusky, A. Simpser, and R. Wang (2011). Comparing the accuracy of rdd telephone and internet surveys conducted with probability and non-probability samples. *Public Opinion Quarterly 75*, 709–747.