NBER WORKING PAPER SERIES

HOUSEHOLD FINANCIAL TRANSACTION DATA

Scott R. Baker
Lorenz Kueng

Working Paper 29027
http://www.nber.org/papers/w29027

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

The authors wish to thank Kwok Yan Chiu and Weijia Zhao for their assistance and Brian Baugh, Anthony J. Cookson, Jonathan Parker, and Scott Schuh for their comments and questions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Household Financial Transaction Data
Scott R. Baker and Lorenz Kueng
NBER Working Paper No. 29027
July 2021
JEL No. C81,D14,G5,H31

ABSTRACT

The growth of the availability and use of detailed household financial transaction microdata has dramatically expanded the ability of researchers to understand both household decision-making as well as aggregate fluctuations across a wide range of fields. This class of transaction data is derived from a myriad of sources including financial institutions, FinTech apps, and payment intermediaries. We review how these detailed data have been utilized in finance and economics research and the benefits they enable beyond more traditional measures of income, spending, and wealth. We discuss the future potential for this flexible class of data in firm-focused research, real-time policy analysis, and macro statistics.

Scott R. Baker
Kellogg School of Management
Northwestern University
2211 Campus Drive
Evanston, IL 60208
and NBER
s-baker@kellogg.northwestern.edu

Lorenz Kueng
Northwestern University
Kellogg School of Management
Suite 4394
2211 Campus Drive
Evanston, IL 60208
and NBER
lorenz.kueng@gmail.com

1 Introduction

Detailed measurement of consumption expenditures is a key ingredient for questions spanning a range of fields across economics and finance. To address such questions, researchers have typically turned to government surveys of households, 'scanner' data with product-level consumer good purchases, or imputing consumption as the residual of the budget constraint using administrative tax records. These approaches each had disadvantages such as high costs, small sample sizes, short panels, being slow to update, covering only portions of the household budget, or being prone to measurement error.

Recent years have seen a surge in research taking advantage of a new class of data that is centered around comprehensive panels of financial transactions from credit, debit, and checking accounts across many thousands or even millions of individuals. Derived from sources such as banks, financial aggregator services (eg. FinTech apps), and credit card companies, the data are often able to be linked to complete views of a households finances across many types of accounts and containing information on both detailed flows as well as high frequency balances. Researchers have employed this data to answer questions about consumption behavior, household balance sheets, and the response to income fluctuations or government policy changes and have greatly enhanced the understanding of heterogeneity in income, balance sheets, beliefs, preferences, and household decision-making in general.

The data has seen such a swift uptake among researchers due to a host of advantages it offers relative to existing survey, scanner, or administrative data. Massive sample sizes, often in the millions of individuals, offer increased statistical precision and the ability to drill down on small sub-samples of interest. Moreover, the granularity of the data and the fact that it is often paired with other financial accounts owned by the individual mean that questions can be asked about the entirety of the household balance sheet or financial flows. Researchers can better control for any number of confounding variables and reject alternative explanations for a given question.

As one example of the rapid proliferation of research in this area, Figure 1 displays the trend in papers utilizing transaction data over time in two of the major repositories for economics and finance research: the National Bureau of Economic Research's (NBER) working paper series and the Social Science Research Network (SSRN) working paper repository. In both cases, there has been a steady growth in papers utilizing this class of data over the past two decades. Notably, 2020 saw the most dramatic growth in papers written using transaction data as researchers took advantage of the granularity of this data and the speed with which it can be accessed to quickly write high impact papers about the COVID-19 pandemic's economic impacts on households and businesses. This new type of data also allowed researchers to quickly evaluate some of the fiscal, monetary, and regulatory interventions enacted all over the world.

As this class of data has increased in availability, researchers have also worked to reduce many of its inherent disadvantages by linking transaction databases with external credit records, conducting surveys of users, or matching to administrative data. Such external linkages increase the ability of researchers to gain qualitative data about a household's characteristics, beliefs, and financial resources.

A first wave of research using direct access to consumer's financial transactions was built around individual brokerage accounts. A range of papers focused on behavioral finance questions about investors, such as Odean (1998), Barber & Odean (2000), Barber & Odean (2001), and Choi et al. (2001). With transaction level brokerage data, researchers could better understand how retail investors made trading decisions, uncovering a number of important stylized facts about heterogeneity across observable types (eg. gender or trading propensities) and how many investors seemed to make systematic mistakes that reduced their overall returns.

Detailed transaction-level data from consumers' bank accounts and credit cards was later to appear, though some work utilized bank and credit card data aggregated to individual-month levels, such as Agarwal et al. (2007). Beginning in the early 2010s, additional research began to appear that utilized transaction data across large sample sizes that was also linked to bank account balances.

Many of these papers aimed to use the new-found level of granularity to provide additional insights into the relationship between income and consumption at a household or individual level. For instance, in one of the first papers to leverage transaction level spending data, Gelman et al. (2014) showed that households responded to the receipt of predictable income by increasing spending in the days immediately after a paycheck, demonstrating that many of the classic models of consumption smoothing may break down at high-frequency levels.

The remainder of the paper is organized as follows. Section 2 describes the sources of transaction data as well as common linkages made to transaction-level data. Section 3 compares transaction-level data to some alternative sources of consumption or household financial flows data. Section 4 details some of the research that has taken advantage of transaction-level data. Section 5 discusses emerging strands of research utilizing transaction-level data across a number of fields and Section 6 concludes.

2 Sources of Payments Data

While earlier work such as Barber & Odean (2000) and Choi et al. (2001) utilized data on financial transactions obtained from brokerage firms, recent transaction data have been derived from a multitude of sources and has expanded to cover all sorts of financial accounts.

One major source of transaction data has been banking institutions. Researchers around the world have been able to form partnerships with both brick and mortar banks (eg. Aydin (2019), Ganong & Noel (2019), Andersen et al. (2020b)) as well as solely online banks (eg. Loos et al. (2020), Baker et al. (2021), D'Acunto et al. (2019)) to gain access to anonymized transaction-level data from their customers. Most banks offer a wide range of services to customers such as savings and checking accounts, debit and credit cards, and mortgages and other lending facilities. Given this span of services, these partnerships offer the possibility of not only observing comprehensive spending and income flows but also various types of assets, borrowing behavior, and other demographic information collected by the banks. Researchers are sometimes limited to a subset of types of financial accounts or to transactions from a limited sample of the bank's customers.

A more limited option has been to obtain access to transaction data directly from credit card firms (eg. Einav & Klenow (2021)) or credit card issuers (eg. Gathergood et al. (2019b)). Such sources will be limited to spending through particular channels and cards and likely have no information regarding income for a given individual or household. Even narrower sets of transactions have been obtained from other proprietary sources. Cookson (2018), for instance, uses transaction-level ATM withdrawal information obtained from casinos throughout the United States. The rise of payment apps such as Apple Pay, Venmo, and AliPay have also allowed for a new avenue to directly observe household spending transactions across large portions of the population in many countries.[1]

While long-standing financial institutions have provided valuable access to transaction level data, much of the pioneering work with this class of data has been enabled by an emerging ecosystem of FinTech products and apps. These products take two extremely broad forms: those that act mainly to aggregate and track a customer's disparate financial accounts for better monitoring and those that are more focused on a particular end-goal such as increasing savings rates or assisting in paying down debt.

For instance, work such as Baugh et al. (2018), Carlin et al. (2017), Baker (2018), and Medina (2020) utilize data obtained from aggregator apps. A significant benefit of these apps is that users typically have a more comprehensive set of financial accounts linked, allowing for both more flexibility in the sorts of questions being asked and also more confidence in the validity of a given set of results. In contrast, Kuchler & Pagel (2018) and Baker et al. (2020c) use data from apps that focus more on encouraging savings or debt repayment and thus both have more potential issues with sample selection but also the completeness of an individual's finances visible within the product.

At least one government entity has also worked to develop a transaction database through a government survey; see Schuh (2018). The Diary of Consumer Payment Choice, run by the Federal Reserve Bank of Atlanta, measures spending behavior through a transaction-level diary among a group of survey respondents. This approach combines much of the utility of surveys with the granularity of transaction data and makes the data widely available to researchers. However, the survey suffers relative to proprietary data obtained from financial institutions in a much smaller sample size and more limited time window (only a few days of transaction data per respondent).

2.1 Ancillary Data Linkages

The basic building blocks of transaction databases shares many similarities across the large variety of sources of such data: generally a complete view of time-stamped financial flows at a transaction level for a set of financial accounts across a large sample of individuals. This provides high quality information regarding household consumption flows. Depending on the source, transaction data

[1]One paper to utilize this data is Xing et al. (2021), wherein the authors use data from Alipay to determine the impact of short-term spending subsidies on consumer spending.

can be accompanied and enriched by other types of financial and demographic characteristics.

Some of the additional data takes the form of account or transaction level metadata. Making use of not only the date of a transaction but also the precise time within a day that a transaction takes place, Gu et al. (2018) identifies 'shirking' at work by examining online shopping purchases during working hours. Aggregators and other financial platforms also generally track how individuals interact with their various online accounts. Olafsson & Pagel (2017) utilize internal data from a FinTech app in order to proxy for attention paid to investment accounts by measuring the number of times a user logs into her account.

Beyond the metadata from the transactions and accounts, many sources link self-reported demographic information to the individual in question, such as age, gender, marital status, levels of education, and number of children (eg. Farrell et al. (2020)). This can also be combined with either self-reported or geo-tagged information on locations of residence or of each individual transaction. While not all sources contain self-reported or administrative data on demographics, some of the firms take the approach of constructing proxies or best guesses for demographic characteristics (eg. gender and age ranges) using observable financial or account characteristics.

A number of researchers are also able to put detailed location data to good use in conjunction with various observed transactions. For instance, Gelman et al. (2019a) utilize locations to identify relevant gasoline prices that individuals face, while Meeuwis et al. (2019) use zipcode level locations to proxy for likely political leanings among users of an investment service. Agarwal et al. (2020b) focus explicitly on location-driven peer effects in Singapore as they examine how bankruptcies affect spending and other financial behavior among co-residents of individual apartment blocks.

Another key element of data linked to transaction databases have been surveys, either administered by the source firm for its own purposes or by the firm on behalf of researchers. Work such as Bachas et al. (2020), Schuh (2018), Baker et al. (2020c), and D'Acunto et al. (2020) all employ surveys of individuals to better contextualize the transactions that are directly observed.[2] Such efforts can help researchers better understand the beliefs, expectations, and reasoning behind the financial decisions that are directly observable, at least partially mitigating one of the issues with context-free transaction data.

[2] Others, such as Giglio et al. (2021), have taken this approach to map brokerage transactions to survey responses regarding beliefs of investors.

For some sources, researchers may also be able to link to various credit reporting agency information like credit scores, numbers of accounts, or negative marks on credit. For instance, Ponce et al. (2017) and Elizondo & Seira (2017) are able to link bank account data to credit bureau data in collaboration with the government in Mexico and Baker (2018) observes credit scores for a subset of users, taking advantage of fact that users of this financial aggregator often choose to connect their credit report.

Transaction databases stemming from financial aggregators are some of the most complete and flexible. This is driven, in part, by the ancillary data that is linked to the financial accounts. Many aggregators contain self-reported data not only on demographics, but also on financial goals and budgets, financial risk tolerances, and financial sophistication. Such data can help researchers to understand how shocks can have heterogeneous responses across agents that vary greatly in key modeling parameters like risk preferences.

Access to the underlying data from FinTech apps has also enabled research into the impacts of this emerging technology, itself, on household financial behavior. Carlin et al. (2017) take a broad overview of the usage of a FinTech app in Iceland, examining differences in patterns of usage across generations. Becker (2017) uses transaction data from a large European bank to demonstrate that access to a money management tool encourages households to start saving and increase savings balances. Levi & Benartzi (2020) utilize the roll-out of a FinTech app across different device types (eg. Android, iOS) to establish that users increase attention paid to their finances and cut discretionary expenditures, with effects concentrated among lower income and high spending-to-income users.

Beyond explicitly linked financial accounts or other data, many researchers have been able to make numerous inferences about individual preferences characteristics using transaction data. For instance, a number of researchers identify a measure of financial sophistication through the propensity for individuals to make financial 'mistakes' like overdrafting or paying credit card penalties when possessing sufficient liquidity to avoid them (eg. Jorring (2019) and Carvalho et al. (2019)). Attention to and salience of financial data and accounts are measured with the frequency by which individuals login to their accounts, and labor market status can be deduced from the receipt of regular direct deposit paychecks or social security income. Account balances may even be able to be partially imputed (when lacking account balance data) using interest income in bank accounts.

These types of inferences are continually refined and augmented across a number of transaction data sources as more researchers begin to work with this class of data.

3 Comparison with Alternate Measures of Consumption

Measuring consumption using household transaction data offers many benefits but also some downsides compared to other sources of consumption data. Table 1 offers a comparison of the attributes of transaction data to those of more traditional sources of consumption data. Importantly, the relative weights placed on these attributes can vary depending on the nature of the question the researcher is asking; each class of data has situations in which it dominates.

For instance, research necessitating product level consumption data will necessitate something akin to scanner data or detailed consumer diary surveys, which sees consumption at a granularity finer than store-level transactions available from credit card statements for example. Questions revolving around durable purchases and financing may be difficult to understand using transaction data, as such purchases often are observed within bank accounts only by a stream of financing charges rather than the actual purchase information. Here, data from tax registries or credit bureau data may dominate. Transaction data will also be inferior when approaching questions regarding aggregate consumption trends over time, since long time horizons are generally unavailable from such sources.

3.1 Advantages of Transaction Data

Some of the longest-standing measures of consumption available to researchers are government-run surveys such as the Consumer Expenditure Survey in the US and the Living Costs and Food Survey in the UK. However, surveys are limited in scope and complexity due to both monetary and time costs. Individuals cannot be asked to answer every question about types of spending or sources of income that might ever be of interest to researchers. Data collection in transaction databases is fully automated, and thus adding an additional transaction incurs no time cost for users and requires minimal monetary costs for the database to log and store it in perpetuity.

Unconstrained by marginal cost concerns on both the part of the user and the entity collecting the data, every individual transaction can be stored along with metadata regarding that transaction.

Therefore, while a survey question cannot easily be asked retroactively, the granular and complete nature of transaction data means that it can be mined far in the future for indicators of behavior or purchases. This is often made possible by the extreme detail offered by the transaction description field, as elaborated on in Section 4.5.

Conditional on having complete views of all of an individual's accounts, values of spending, income, and assets are exact and are not as prone to recall bias as in a survey setting. This can be especially useful when working to measure the value of highly volatile assets like equity portfolios or when investigating atypical spending patterns.

This completeness across financial accounts also drives the main benefits relative to so-called 'scanner data' such as the AC Nielsen Consumer Panel, which tracks subsets of household purchases at an individual product level. Scanner data can grant even finer detail on elements of household spending by decomposing a transaction into its constituent products. However, the set of covered products and stores is limited (generally to nondurable purchases from retail stores) and other financial attributes of the individual are often lacking.[3]

Similarly, researchers employing data from credit bureaus can track some aggregated versions of household spending through the observation of credit card balances. Durables purchased with debt such as vehicles and homes are well tracked in such data, but credit card balances mix revolving debt with current charges (see Fulford & Schuh (2020)), making it difficult to ascertain the flow of household consumption even when all household purchases are made via credit card rather than debit or cash spending.

Beyond any of the advantages stemming the types of data available from these databases, transaction data is often distinguished from other data sources by its sheer size. Table 2 enumerates the sample sizes for some common survey databases as well as some of transaction databases used by researchers in recent years. For many of these databases, sample sizes are often one or two orders of magnitude larger, yielding correspondingly smaller standard errors, especially when examining behavior of a particular subset of the population.

The advantage of size has fueled a substantial expansion in research into the micro-foundations of the macroeconomy and has focused new attention on models that incorporate heterogeneity across individuals and households. The ability to divide a sample jointly by income, liquidity,

[3]For instance, the NCP covers approximately $400 of household spending per month for its panel members, a small fraction of total household spending in the United States.

credit, and age can help to test new tools such as Heterogeneous Agent New Keynesian (HANK) models, which add both heterogeneity and idiosyncratic risk to Representative Agent New Keynesian (RANK) models. More broadly, new questions about inequality of wealth and income can be more thoroughly answered when the top 1% of a sample is composed of 10,000 individuals rather than 500.

Recent years have also seen an increase in researchers employing detailed administrative data on income and wealth to compute 'imputed' consumption as the residual of the household's budget constraint. Countries such as Denmark, Sweden, and Norway have made annualized financial data on all citizens available for research, including things like bank balances, income, and asset holdings. This approach is exemplified by work such as Koijen et al. (2014) in Sweden, De Giorgi et al. (2020) in Denmark, and Fagereng & Halvorsen (2017) in Norway. Relative to transaction-level data, these databases substantially reduce issues of sample selection and incompleteness and also make other links to government databases possible. However, they do not provide any high frequency measures of financial activity and can generally not directly measure spending, meaning that only an aggregate annual level of spending or consumption can be imputed. In addition, calculating consumption as a residual implies that different types of consumption (eg. spending on non-durables, food, services, or durables) cannot be observed.

In developing countries with fewer resources to devote to systematically collecting economic data, transaction data can also fill a gap in coverage. While not all countries collect detailed survey-based microdata, all have substantial penetration of banks, card payment systems, or FinTech products. Just as many countries are leap-frogging wired telephone infrastructure in favor of cell phones, some countries have largely bypassed credit and debit cards, moving to mobile-based payment systems. In such areas, collaborating with FinTech apps can give researchers and policymakers a clearer view of the economy through this type of transaction data.

Finally, transaction data can increasingly provide a tool to understand not only what has happened, but also what is happening. Seen most clearly with the proliferation of research employing transaction data during the COVID-19 pandemic, transaction data can be utilized to give policymakers a new and highly detailed view into the current state of the real economy, beyond financial markets, without having to wait weeks, months, or years for surveys to be administered.

The high frequency nature of the data allows researchers to unpack fast-moving events like a government shutdown (eg. Gelman et al. (2018)) or the onset of the COVID-19 pandemic (eg. Chen et al. (2020)) or even to examine patterns of spending across hours of the day (eg. Gu et al. (2018)). This enables researchers to more directly observe economic trends and shocks rather than imputing effects from data aggregated to monthly, quarterly, or annual level.

3.2 Limitations of Transaction Data

3.2.1 Account Completeness

While transaction data offers a number of substantial advantages over more traditional data sources, they exhibit disadvantages, as well. One basic difficulty is in interpreting the unit of analysis because, for many such databases, the unique identifier is an 'account'. While in a survey, a surveyor can set well-defined boundaries on the unit of analysis, households often use individual financial accounts and aggregator apps in a variety of ways. Take, for instance, a married couple who possess a number of individual financial accounts as well as several joint accounts. Each member of this couple may have their own account with the financial aggregator app that then contains links to their own individual as well as joint accounts. Another possibility is that they may have a single account with the app that contains both sets of individual accounts as well as all the joint accounts.

In general, a researcher can never be certain that they are observing all relevant financial accounts owned by an individual or a household, and must make assumptions or sample restrictions to try to isolate the more 'complete' users. Many researchers take the approach of excluding users with small numbers of transactions per month, no visible sources of income, or financial transfers to accounts that are unlinked. Such ambiguity may make it difficult to interpret the raw levels of income and spending and how they can be compared to other sources of information where the unit of analysis is more clearly documented. More broadly, these databases tend to lack the decades of supporting documentation found in government surveys and other sources that are designed from the ground up to be used by researchers and to make accurate economic inferences.

3.2.2 Sample Selection

Even in the most information-rich and comprehensive transaction databases, some disadvantages remain. One common criticism is that, by their very nature, financial transaction databases are composed of a selected sample of individuals rather than a nationally representative one. Even when dealing with data sourced from large national financial institutions, there exists a sizable minority of the population that do not have bank accounts or credit cards of any type and deal primarily in cash.

The selection issue becomes more acute in instances where the transaction data is obtained from more specialized financial institutions or FinTech apps that are targeted at particular slices of the population. For instance, some apps may be aimed at poorer individuals, who have trouble saving, at individuals who enjoy trading stocks, or at individuals who have a large number of varied financial accounts that need managing. Each of these data sources may exhibit a somewhat different bias relative to the population when investigating a given question. Baker (2018) uses the demographic and financial characteristics linked with transaction data to demonstrate that such sample selection issues can be mostly resolved using population weights given sufficient numbers of users.

3.2.3 Transaction Categorization

Spending done via cash or paper checks can pose problems for researchers using transaction data. Cash transactions cannot be observed directly, though may be able to be inferred by measuring ATM withdrawals. Check transactions can be observed but cannot be automatically categorized into the types of spending that they represent. This concern has declined as cash and checks have come to represent smaller fractions of spending, though the growth in peer-to-peer cash transfer apps (eg. Venmo) often present some of the same concerns.

More broadly, unlike scanner or some types of survey data, categorization of spending is necessarily done at a transaction level rather than a product level. While a transaction at a restaurant is easily interpretable, a transaction at a big box retailer (or at an online retailer like Amazon.com) may contain a range of products from non-durable consumables like groceries to semi-durable and durables like clothing, furniture, or mechanical equipment. Researchers may choose to focus on these more unambiguous categories such as restaurants when, for instance, they discuss responses of 'non-durable' spending to an income shock.

Categorization of spending within household transaction data is generally done by algorithms constructed by the parent organization collecting the transaction data (eg. the bank or account aggregator). Typically, these utilize a combination of Merchant Category Codes (MCC) linked to merchant IDs as well as textual analysis of transaction descriptions to arrive at a category for a given merchant. In some cases, miscategorizations may occur or transactions may be left as 'uncategorized'. However, the textual descriptions and other metadata linked to each transaction can allow researchers to improve on such categorization algorithms or directly observe particular transactions of interest without being limited to the categories assigned by the parent organization.

3.2.4 Replication

The proprietary nature of these databases can also present problems for accessibility for future research as well as for the ability to replicate a given analysis. At an extreme case, some private sources of data such as smaller FinTech apps may go out of business or deprecate older data to minimize storage costs, making it impossible to replicate results. More broadly, research done with some of these databases is the result of personal relationships between researchers and the firms and there is often no established way for other researchers to gain access.

These issues stand in stark contrast to many of the traditional data sources that are primarily sourced from public agencies. Many common government survey data are publicly accessible online or have a standard application process for gaining access to more detailed data. One avenue for expanding the utility of transaction databases is for governments or large research institutions to cement more permanent means of collaboration between researchers and private data providers.

4 Research Avenues Exploiting Benefits of Linked Accounts

4.1 Micro-based Aggregate Consumption Statistics

When it comes to tracking consumer spending in the economy, many economists have turned either to national aggregates such as Personal Consumption Expenditures (PCE) from the National Income and Product Accounts (NIPA) or to national surveys such as the Census Bureau's Retail Trade Survey. However, these data are produced with a substantial lag and provide no detail at a sub-national level. Using transaction data as a building block for better aggregate statistics offers a promising path to a clearer picture of the economy.

Batch et al. (2019) test whether modern machine learning paired with alternative (and faster-to-update) data can provide better advance GDP estimates or lower the size of GDP revisions. While they themselves only have access to an aggregated version of transaction data provided by a financial services vendor, they find that such data is extremely valuable for better predicting components of national accounts. Aladangady et al. (2019) go further, using transaction-level data from a large electronic payments company to provide a disaggregated and higher frequency version of these national consumption statistics. They demonstrate the utility of this measure in tracking spending during fast-moving economic shocks like the 2019 government shutdown and major hurricanes in 2017.

Of these various efforts, those found in Chetty et al. (2020) are perhaps the most ambitious. In the authors' own words: "We build a publicly available platform that tracks economic activity at a granular level in real time using anonymized data from private companies." Many of the statistics presented and tracked are derived from detailed transaction data that allow the authors to provide many dimensions of heterogeneity in high-frequency spending behavior, including disaggregated statistics by geographic area, income, and industry. The ability to publicly provide such detailed statistics in a continuous and up-to-date format is of great importance to both researchers as well as policy-makers when seeking to understand economic events as they occur rather than waiting months for reliable data.

4.2 Linked Financial Accounts

Household transaction data by itself allows for a wide range of new insights into household behavior and the responses to economic shocks. For instance, Einav & Klenow (2021), Baugh & Correia (2020), and Baker et al. (2020a) work solely with transaction databases that do not feature much other information about account balances, assets, or demographics.

However, while transaction data has proven useful even when unaccompanied by a fuller suite of financial information, this class of data has had the largest impact when paired with more complete household financial information. Given the sources of the data are most commonly banks and financial aggregators, it is common to have comprehensive overviews financial information for an individual across a range of important account types. One of the most important components that such comprehensive financial data offer is a time-varying measure of individual liquidity and credit access. This has allowed a large group of researchers to demonstrate the importance of liquidity in driving household responses to both anticipated and unanticipated income shocks. Baker (2018), Agarwal & Qian (2014), Ganong & Noel (2019), Andersen et al. (2020c), and Kueng (2018) all leverage metrics like account balances and credit availability to map out household level heterogeneity in consumption behavior.

For instance, Agarwal & Qian (2014) and Baker (2018) take this data to the classic question about the consumption response to unanticipated income shocks and how the consumption response varies across households with different levels of liquid assets and credit access. The former uses transaction level data from a Singaporean bank to investigate the consumption response to an economic stimulus policy in Singapore, finding that households begin to respond upon the announcement of the upcoming payments and that households with lower levels of liquidity and credit respond more strongly to its receipt. The latter uses shocks to households' employers to instrument for permanent income shocks, finding similar levels of heterogeneity across credit access and liquidity.

Another common linkage is between transaction accounts and home mortgage accounts, enhancing the ability of researchers to examine how home values, mortgage terms, and mortgage performance interact with consumption decisions. Park (2016) takes advantage of information from household transactions to predict mortgage delinquencies, using the propensity of households being charged late fees or overdrafts to proxy for risk. Similarly, Jorring (2019) shows that mistakes in checking or credit accounts (eg. avoidable overdrafts or late fees) can predict the consumption response to predictable changes in mortgage payments. Households who seem to make such financial mistakes often fail to smooth consumption around these theoretically known events.

A number of papers, such as Ganong & Noel (2020) and Bernstein (2020), have been able to directly examine the impact of home values, mortgage policies and mortgage terms like ARM resets on household consumption behavior and labor supply. Given the prominence that homes and mortgages have among household balance sheets around the world, such questions are of first order importance when thinking about the transmission of monetary policy or associated changes to federal home loan policies.

Beyond houses and mortgages, linkages of transaction data to investment and retirement accounts have also been able to provide additional clarity about household balance sheets. For instance, Brauer et al. (2019) and Meyer & Pagel (2019) look to more fully characterize consumption responses to dividends as well as both realized and unrealized capital gains and losses. Again, transaction data provide a way to credibly identify how asset price fluctuations may be transmitted through the economy and can highlight heterogeneity across owners of these assets, as well.

Others rely on transactions across multiple financial accounts to test for offsetting behavior that may net out in aggregate or may be missed if solely observing one type of financial accounts. Medina (2020) highlights the importance of testing for offsetting financial behavior across separate financial accounts. Garcia et al. (2020) examine whether similar offsetting behavior comes about after inducing individuals to save more through a similar intervention, finding little response on the credit utilization. Kukk & Raaij (2020) look at savings distributions across individual and joint accounts within households, finding a wide range of joint allocations.

4.3 Evaluating Heterogeneity and Models of Household Behavior

Models working to understand and rationalize household behavior, especially consumption behavior, are also able to be more thoroughly tested when increased sample sizes are combined with the extreme detail allowed for by transaction data. Key parameters of these models like risk aversion and salience can be more directly captured and important financial characteristics such as precise levels of liquidity and credit access are readily available. For instance, Ponce et al. (2017) note that relative prices, such as interest rates, are poor predictors of how households make decisions about which cards to utilize and which to pay down. In an experiment with randomized interest rate offers, they find that households respond to these salient price changes, suggesting models of limited attention or mental accounting might rationalize the empirical findings.

The sensitivity of consumption to beliefs, news, and income shocks of various types has generated a large array of research using transaction data. Most of these papers address models of intertemporal optimization, often test various aspects of the permanent income hypothesis or buffer stock models. Transaction data at an individual or household level is invaluable to such efforts, as many of the common explanations for consumption fluctuations revolve around income volatility or uncertainty, levels of liquidity and wealth, spending shocks, and durables spending. With detailed transaction data linked with financial account balances and credit information, these explanations can be tested jointly and with great precision across a large sample size. Moreover, identifying precise income shocks, as well as understanding just how anticipated or unanticipated those shocks were, is made much easier when each individual income transaction can be observed and fully characterized according to its source.

As one example, Ganong & Noel (2020) link transaction data with mortgage data to assign mortgage defaults into one of two categories: strategic defaults, where households default when debt is too high relative to home values, and liquidity-driven defaults, where monthly payments are too high given current household resources. They find evidence for the overwhelming majority of defaults being liquidity or 'adverse event' driven rather than strategic, informing a debate over optimal mortgage and debt relief policies.

Looking at responses to macroeconomic news events, Garmaise et al. (2020) use household transaction data to identify persistent excess sensitivity of consumption, even when such events are not necessarily meaningful for the household. For more salient events – those that garner more media attention – responses are higher, as are responses among lower income households.

Other papers address the response of households to personal income shocks. Olafsson & Pagel (2018) highlight excess spending among individuals on their paydays among most of the population of Iceland. They show that even for high income and highly liquid individuals, paydays tend to induce appreciable spending responses across many categories of consumption despite the fact that paydays are known in advance. They explicitly test multiple models of household behavior and reject the possibility that either current or imagined future liquidity constraints are likely the driver of such spending responses. In Olafsson & Pagel (2019), the authors turn towards *unanticipated* income shocks from lotteries, finding MPCs above one – financed by borrowing – for small wins and low MPCs for large wins, again pointing away from liquidity constraints as the main driver of large MPCs out of windfalls.

Households are shown to exhibit excess sensitivity to high profile and perfectly anticipated payments from the Alaska Permanent Fund, as well. Kueng (2018) notes that excess sensitivity in this case is actually increasing in income and is high even for those with high levels of liquid assets. While such behavior stands in contrast to most classical rational models, he shows that the actual welfare loss from such non-smoothing behavior is exceedingly small for these households, suggesting that many consumption deviations identified by other researchers are also likely minimal from a welfare perspective.

Other evidence against liquidity constraints being the primary driver of excess sensitivity of consumption is given by Gelman (2019a). The author tests a range of parameters of the canonical buffer stock model and notes that a quasi-hyperbolic version of the model performs better than the standard exponential formulation in explaining the observed levels of excess sensitivity. In Gelman (2019b), the author takes advantage of the vast sample size afforded by transaction data to decompose the aggregate or average MPCs across individuals with varying observable and unobservable characteristics. He returns to the buffer-stock model and provides perhaps the first clear data-driven breakdown of the importance of these persistent characteristics and time-varying shocks.

Heterogeneity across individuals is also the prime focus of Kuchler & Pagel (2018). They demonstrate that individuals often fail to stick to explicit plans to pay down their debt, consistent with present-biased agents that feature in models with quasi-hyperbolic discounting. They push further with detailed transaction data, showing that heterogeneity across individuals in their debt repayments is strongly correlated with cross-sectional variation in short-run impatience and financial sophistication.

Gathergood et al. (2019a) and Gathergood & Olafsson (2020) are able to use such detailed data to test theories of mental accounting at a household level. In Gathergood et al. (2019a), the authors find that households do not treat all credit card expenditures equally; they tend to pay down debt incurred by certain types of expenses more quickly than others. In particular, households tend to pay down debt more quickly for non-durable expenditures relative to durable expenditures. Agarwal et al. (2020a) examine how household spending changes in the days around credit card statement dates. They suggest that household spending tended to rise substantially following these statement dates but that most of the variation could be rationally accounted for by households taking advantage of the free float that credit cards offer until the following pay period.

Lastly, in Gathergood & Olafsson (2020), the authors turn to the puzzle of households that co-hold liquid balances alongside high-interest debt. Using linked transactions, credit, and bank balances, they calculate these propensities at a daily level, finding that instances of co-holding are usually very short term but that rational explanations are likely subordinate to mental accounting on the part of the household.

4.4 Experimental Interventions

Transaction data has also been utilized in conjunction with experiments on households. With such detailed data, gathering post-treatment data is automatic and comprehensive, allowing for larger samples and more detailed questions to be investigated.

As one example, Aydin (2019) randomizes the extension of credit lines to customers of a large Turkish bank, finding substantial propensities to borrow even among those with sizable assets. He takes advantage of the completeness of the credit, income, spending, and asset data to build support for a buffer-stock model rather than more behavioral explanations. D'Acunto et al. (2020) also investigate the spending responses to the provision of a new credit facility for borrowers and elicit preferences through the FinTech app that they obtain transaction data from. With this data, they are able to fully characterize spending responses across both the income and liquidity distribution as well as across users with various beliefs and risk preferences.

Such linked data can also offer convincing evidence for a role for for behavioral 'mistakes' in household behavior. Medina (2020) runs an experiment on users of a financial account aggregator in Brazil, randomly sending reminders to pay credit card bills to users of the service. She finds that households significantly increase on-time payments but also increase the probability of over-drafting on other accounts they own. This finding highlights one strength of transaction data stemming from financial aggregators: researchers are easily able to observe potentially offsetting behavior occurring in different financial accounts owned by the same individual in response to a shock.

4.5 Utilization of Transaction Text Fields

One of the most flexible components of the common classes of transaction data thus far used by researchers is the 'text' field. That is, the textual description that accompanies the amount and time-stamp for each transaction. This text (or a simplified version of it) is what is observed, for instance, when someone examines their itemized credit card statement at the end of a billing cycle. While this text may feature complex numerical identifiers or abbreviations of firms, it is generally possible to use this text to determine which firm a particular expenditure transaction was paid to or which employer a particular income transaction is from.

This field thereby allows a researcher to pick out transactions that may signal an employment relationship, the timing of government benefit receipt, spending at a particular retailer, or the occurrence of an event of interest. This can all be accomplished without having to explicitly ask a household about any of these linkages or events and it gives researchers the ability to extend their analysis years backwards in time.

In Baker (2018), the author is able to link households to their employers in order to identify unanticipated changes to households' source of income. Given the scarcity of data linking employees and employers, the ability to generate these links is quite valuable in understanding the transmission of shocks through the economy. Baker & Yannelis (2017) and Gelman et al. (2018) also utilize this text field to identify households that contain employees of the federal government in order to analyze differential spending behavior when such households were subjected to missed paychecks as part of the 2013 Federal Government Shutdown.

Others, such as Agarwal et al. (2021), Einav & Klenow (2021) and Baker et al. (2020a), have used a similar approach to map transactions to the numbers of firms that these individuals shop at, mostly retailers or consumer service firms. Linking to a broad swatch of firms enables these researchers to analyze the firms' customer bases and provide new dimensions of information regarding risks to firm growth and an ability to disaggregated their sources of revenue. Baugh et al. (2018) focuses on transactions at a single firm, Amazon.com, and examines responses to changes in sales tax policy. Such linkages to spending at individual firms is not possible using alternative consumption data sources like government surveys, scanner data, or by imputing consumption via tax registries.

Another common tactic when using transaction data has been to identify less 'typical' sources of income or benefits. While administrative data may generally be required to observe precise amounts of unemployment insurance paid out to individuals, Ganong & Noel (2019) and Andersen et al. (2020c) are able to pick out UI payments using the textual descriptions within transaction databases. Similarly, Karger & Rajan (2020) is able to take the same approach with CARES Act Stimulus payments in 2020, Kueng (2018) with payments from the Alaska Permanent Fund, and Olafsson & Pagel (2019) with winnings from national lotteries. Baugh (2016) and Carvalho et al. (2019) pick out loans from and repayments to payday lenders, and Gelman et al. (2019b) can identify the precise timing and amounts of tax returns deposited into accounts via direct deposit.

Overall, the text field found in many types of transaction data offers researchers unparalleled room for creativity in identifying particular sources of income or types of spending.

4.6 Transaction Data During the COVID-19 Pandemic

A unique demonstration of the utility of transaction data came during the COVID-19 pandemic that swept the world in 2020. This pandemic brought about a deep recession in many countries that was notable for not only its size but its speed, as well. The rapidity of the economic shocks was unlike anything seen in modern history, with many countries seeing the largest ever single-month or single-quarter declines in employment, spending, and GDP.

The speed at which this event occurred meant that governments were grappling with policy proposals for household and business support without being able to observe many of the economic indicators that are only produced with a lag of weeks or months. Turning to transaction data, many researchers were able to quickly produce diagnoses of the size and distribution of consumption and income declines. Such analyses were conducted in a wide range of countries; eg. Andersen et al. (2020a) for Denmark and Sweden, Chetty et al. (2020) and Cox et al. (2020) for the United States, Chen et al. (2020) for China, Carvalho et al. (2020) for Spain, Chronopoulos et al. (2020) and Hacioglu et al. (2020) for the United Kingdom, and Bounie et al. (2020) for France.

Governments and central banks often focus on the behavior of financial markets to provide high-frequency views on economic activity, especially in response to recent policy initiatives. The newfound flexibility and utility of transaction-based consumption data during a fast-moving crisis may prompt governments or other institutions to set up persistent indicators of consumption leveraging high frequency transaction data. If these consumption indicators are available in the future, policymakers may be better able to calibrate counter-cyclical policy armed with more accurate data.

5 Future Directions of Research

5.1 Analysis of Firms Using Transaction Data

5.1.1 Household-level Transaction Data

An emerging research agenda has begun to leverage households' financial transactions to study features of and risks facing firms. Researchers are able to map individual households' transactions to particular firms using either unique firm identifiers or the text fields from credit and debit transactions. Novel firm-level panels detailing revenue sources and customer base attributes can thus be constructed from household transactions.

For instance, Agarwal et al. (2021) use credit card data to produce disaggregated firm-level sales data and find that such data is predictive of future stock returns. Moreover, they find that the composition of customer bases is important, with stronger stock returns following increases in spending from customers with higher credit scores, wealth and loyalty.

Baker et al. (2020b) demonstrate that aggregated firm-level statistics derived from household transaction data match well to traditional sources of firm-level data like Compustat. They develop two novel statistics relating to customer bases that cannot be generated using traditional data sources, the rate of churn in a firm's customer base and the pairwise similarity between two firms' customer bases. They then demonstrate that these statistics can be used to inform alpha-generating trading strategies as well as providing new insights into firm investment and markups.

Finally, Einav & Klenow (2021) use debit and credit card transactions spanning Visa's network to demonstrate the importance of measuring customer bases when working to understand firm growth. That is, they demonstrate that extensive margin adjustments (eg. customers leaving or joining a customer base) are responsible for the vast majority of firm-level sales variation over time.

While this window into firms provides substantial additional value to these household transaction datasets, there are limitations to this approach. Since the source of the data is ultimately a database of consumer transactions, firms that obtain sizable amounts of revenue from other businesses or governments will be poorly observed or measured. Moreover, only firm revenue, not costs or employment and production characteristics, can be directly observed. Despite these drawbacks, consumer-facing firms constitute a large portion of most economies and feature prominently in questions spanning fields like finance, marketing, and industrial organization.

5.1.2 Firm-level Transaction Data

Another promising area of research leverages expanded sources of transaction data to directly observe *firm-level transactions*. Much as transaction data has enabled a more thorough analysis of household behavior and heterogeneity, the transactions made by the tens of millions of small businesses can increase visibility within this large sector of the national economy. Especially in countries with less reliable measures of small business activity, the ability to directly observe transactions from these firms by working with banks or other financial intermediaries can greatly empower researchers and policymakers.

Two pioneers in this field can be seen in Weber (2019) and Kim et al. (2020). Weber (2019) links the personal financial transactions from millions of individuals at a large financial institution to the universe of small business transactions at that institution. The author focuses on the founders of small businesses, measuring personal financial characteristics such as cash flow management abilities, which are observable from the founders' personal accounts. He notes that such financial management skills are highly predictive of the financial success of the small businesses subsequently founded by these individuals.

In Kim et al. (2020), the authors work with the JPMCI to gain access to transaction-level data from hundreds of thousands of small businesses. In particular, they are able to match transactions of small businesses to the personal accounts of those businesses' proprietors and map out how the fortunes of the businesses impact proprietors' consumption during the COVID-19 pandemic.

5.2 Integration with Governments

Transaction data may have the potential for substantial future contributions to research and policy through additional integration with governmental agencies and central banks. Current research using transaction data often involves one-off relationships between researchers and data providers. For a range of reasons that encompass both competitive and legal concerns, only a small number of firms have allowed for more regular access to their transaction databases. As a result, obtaining and maintaining access to such data is difficult or impossible for most researchers, and the small scale of most collaborations requires re-learning the idiosyncrasies of each particular transaction database.

Partnerships with government agencies may provide one way to provide a more standardized process for access and also leverage the expertise that many agencies have built up in producing credible and well-documented statistics regarding national product accounts or other economic indicators. Oversight of data access by federal regulators, built on years of facilitating access to IRS or Census records, may also provide the legal assurance firms need to ensure that vetted researchers will not misuse their data.

In addition, linkages of transaction data to other sensitive economic databases, such as tax or credit registries, could be facilitated. For instance, in Denmark, Andersen et al. (2020d) use transaction data matched to government registry data to measure the extent to which informal transfers from family and friends act as insurance for negative income or labor supply shocks. In the United States, Chetty et al. (2020) work with a large number of private sector partners to build a database of economic indicators that is both granular and timely. Overarching sponsorship or organization of such efforts by a neutral party like the Federal Reserve could provide an avenue for the expansion of these efforts.

Governments would be well-served by aiding in these efforts to partner with data providers in systematic way. As has been seen during the COVID-19 pandemic, transaction data can provide some of the most up-to-date statistics and can serve to inform policy decisions at a much finer level than aggregate statistics used in the past. In addition, as governments around the world tackle persistent questions of inequality, transaction level data offers a ready-made window into the economic conditions and decisions of millions of households across the income distribution.

6 Conclusion

The past decade has seen a tremendous amount of growth in research utilizing a new class of data centered on individual financial transactions. This data is often linked to other financial data such as asset balances, credit, and equity holdings as well as self-reported demographic characteristics and geographical information. Whereas many of these detailed aspects of household finance were available for a small slice of the population or could be imputed from less granular survey questions, financial transaction data makes most financial and labor market attributes of a household directly observable.

These databases have enabled researchers to more decisively answer classic economic questions using new identification strategies and to obtain more precise empirical estimates. Such questions range from how households respond to economic shocks, the transmission of government monetary and fiscal policy, and the importance of household balance sheets.

Beyond this revisiting of questions first tested using pre-existing survey and administrative data, this class of data has also proven to be some of the most flexible in terms of its potential for wider research. Researchers have taken advantage of the massive sample sizes to highlight the importance of heterogeneity in both economic means and also economic beliefs and behavior in explaining how the economy moves in response to aggregate shocks. The ability to use the textual transaction descriptions to ex-post identify particular firms, income sources, and detailed types of spending has enabled researchers more room to think creatively about particular shocks to individuals or to define multidimensional characteristics that can be important for structural estimation.

The speed of analysis possible using this data has also been highlighted during the COVID-19 pandemic, with both economic fundamentals as well as the response to government fiscal interventions being tracked nearly in real time and with unprecedented granularity. This powerful demonstration of value seems likely to cement transaction data as an important tool for policy-makers and may aid in the construction of more accessible datasets of this type for researchers.

As the number of sources of transaction data continues to climb around the world, many worries – about representativeness, the proprietary nature of the many of the data sources, and the ability to replicate research – are declining. And more researchers from additional fields are beginning to recognize the value and flexibility of such detailed data, leading to an expansion in use among fields like marketing, IO, development, asset pricing, and corporate finance.

References

Agarwal S, Bubna A, Lipscomb M. 2020a. Timing to the statement: understanding fluctuations in consumer credit use. *Management Science*

Agarwal S, Liu C, Souleles NS. 2007. The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data. *Journal of Political Economy*

Agarwal S, Qian W. 2014. Consumption and debt response to unanticipated income schocks: Evidence from a natural experiment in singapore. *The American Economic Review*

Agarwal S, Qian W, Zou X. 2020b. Thy neighbor's misfortune: peer effects on consumption. *American Economic Journal: Economic Policy*

Agarwal S, Qian W, Zou X. 2021. Disaggregated sales and stock returns. *Management Science*

Aladangady A, Aron-Dine S, Dunn W, Feiveson L, Lengermann P, Sahm C. 2019. From transactions data to economic statistics: constructing real time, high frequency, geographic measures of consumer spending. *Working Paper*

Andersen AL, Hansen ET, Johannesen N, Sheridan A. 2020a. Pandemic, shutdown and consumer spending: lessons from scandinavian policy responses to covid19. *Working Paper*

Andersen AL, Hansen ET, Johannesen N, Sheridan A. 2020b. Social distancing laws cause only small losses of economic activity during the covid19 pandemic in scandinavia. *Proceedings of the National Academy of Sciences*

Andersen AL, Jensen AS, Johannesen N, Kreiner CT, Leth-Petersen S, Sheridan A. 2020c. How do households respond to job loss? lessons from multiple high-frequency data sets. *Working Paper*

Andersen AL, Johannesen N, Sheridan A. 2020d. Bailing out the kids: new evidence on informal insurance from one billion bank transfers. *Working Paper*

Aydin D. 2019. Consumption response to credit expansions: Evidence from experimental assignment of 45307 credit lines. *Working Paper*

Bachas P, Gertler P, Higgins S, Seira E. 2020. How debit cards enable the poor to save more. *Working Paper*

Baker SR. 2018. Debt and the response to household income shocks: validation and application of linked financial account data. *Journal of Political Economy*

Baker SR, Baugh B, Kueng L. 2020a. Income fluctuations and firm choice. *Journal of Financial and Quantitative Analysis*

Baker SR, Baugh B, Sammon M. 2020b. Measuring customer churn and interconnectedness. *Working Paper*

Baker SR, Farrokhnia RA, Meyer S, Pagel M, Yannelis C. 2020c. Income, liquidity and the consumption response to the 2020 economic stimulus payments. *Working Paper*

Baker SR, Kueng L, Meyer S, Pagel M. 2021. Consumption imputation errors in administrative data. *Working Paper*

Baker SR, Yannelis C. 2017. Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics*

Barber BM, Odean T. 2000. Trading is hazardous to your wealth: the common stock investment performance of individual investors. *The Journal of Finance* 55(2):773–806

Barber BM, Odean T. 2001. Boys will be boys: gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*

Batch A, Chen J, Driessen A, Dunn A, Kyle K. 2019. Off to the races: a comparison of machine learning and alternative data for predicting economic indicators. *Working Paper*

Baugh B. 2016. Payday borrowing and household outcomes: evidence from a natural experiment. *Working Paper*

Baugh B, Ben-David I, Park H. 2018. Can taxes shape an industry? evidence from the implementation of the "amazon tax". *The Journal of Finance*

Baugh B, Correia F. 2020. Does paycheck frequency matter? evidence from micro data. *Working Paper*

Becker G. 2017. Does fintech affect household saving behavior: findings from a natural field experiment. *Working Paper*

Bernstein A. 2020. Negative home equity and household labor supply. *The Journal of Finance*

Bounie D, Camara Y, Galbraith JW. 2020. Consumers' mobility, expenditure and online offline substitution response to covid19: evidence from french transaction data. *Working Paper*

Brauer K, Hackethal A, Hanspal T. 2019. Consuming dividends. *Working Paper*

Carlin B, Olafsson A, Pagel M. 2017. Fintech adoption across generations: financial fitness in the information age. *Working Paper*

Carvalho L, Olafsson A, Silverman D. 2019. Misfortune and mistake: the financial conditions and decision making ability of high cost loan borrowers. *Working Paper*

Carvalho VM, Garcia JR, Hansen S, Mora SR, Ortiz A, et al. 2020. Tracking the covid19 crisis with high resolution transaction data. *Working Paper*

Chen H, Qian W, Wen Q. 2020. The impact of the covid19 pandemic on consumption: learning from high frequency transaction data. *Working Paper*

Chetty R, Friedman JN, Hendren N, Stepner M, Team TOI. 2020. How did covid19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data. *Working Paper*

Choi JJ, Laibson D, Metrick A. 2001. How does the internet affect trading? evidence from investor behavior in 401(k) plans. *Journal of Financial Economics*

Chronopoulos DK, Lukas M, Wilson JOS. 2020. Consumer spending responses to the covid19 pandemic: an assessment of great britain. *Centre for Economic Policy Research*

Cookson JA. 2018. When saving is gambling. *Journal of Financial Economics* 129(1):24–45

Cox N, Farrell D, Gannong P, Greig F, Noel P, et al. 2020. Initial impacts of the pandemic on consumer behavior: evidence from linked income, spending and savings data. *Working Paper*

D'Acunto F, Rauter T, Scheuch CK, Weber M. 2020. Perceived precautionary savings motives: evidence from fintech. *Working Paper*

D'Acunto F, Rossi AG, Weber M. 2019. Crowdsourcing financial information to change spending behavior. *Working Paper*

De Giorgi G, Frederiksen A, Pistaferri L. 2020. Consumption network effects. *Review of Economic Studies*

Einav L, Klenow P. 2021. Customers and retail growth. *Working Paper*

Elizondo A, Seira E. 2017. Are information disclosures effective? evidence from the credit card market. *American Economic Journal: Economic Policy*

Fagereng A, Halvorsen E. 2017. Imputing consumption from norwegian income and wealth registry data. *Journal of Economic and Social Measurement*

Farrell D, Ganong P, Greig F, Jones D, Noel P, Wheat C. 2020. Wealth, race and consumption smoothing of typical income shocks. *Working Paper*

Fulford S, Schuh S. 2020. Revolving v olving versus conv ersus convenience use of cr enience use of credit car edit cards: Evidence fr vidence from u.s. credit bureau data. *Working Paper*

Ganong P, Noel P. 2019. Consumer spending during unemployment: positive and normative implications. *The American Economic Review*

Ganong P, Noel P. 2020. Why do borrowers default on mortgages? a new method for causal attribution. *Working Paper*

Garcia I, Medina P, Pagel M. 2020. Does saving cause borrowing. *Working Paper*

Garmaise M, Levi Y, Lustig H. 2020. Spending less after (seemingly) bad news. *Working Paper*

Gathergood J, Loewenstein G, Quispe-Torreblanca EG, Stewart N. 2019a. The red, the black and the plastic: paying down credit card debt for hotels, not sofas. *Management Science*

Gathergood J, Olafsson A. 2020. The co-holding puzzle: new evidence from transaction level data. *Working Paper*

Gathergood J, Sakaguehi H, Stewart N, Weber J. 2019b. How do consumers avoid penalty fees? evidence from credit cards. *Working Paper*

Gelman M. 2019a. The self-constrained hand to mouth. *Working Paper*

Gelman M. 2019b. What drives heterogeneity in the marginal propensity to consume? temporary shocks vs persistent characteristics. *Working Paper*

Gelman M, Gorodnichenko Y, Kariv S, Koustas D, Shapiro MD, et al. 2019a. The response of consumer spending to changes in gasoline prices. *Working Paper*

Gelman M, Kariv S, Shapiro MD, Silverman D. 2019b. Rational illiquidity and consumption: theory and evidence from income tax withholding and refunds. *Working Paper*

Gelman M, Kariv S, Shapiro MD, Silverman D, Tadelis S. 2014. Harnessing naturally occurring data to measure the response of spending to income. *Science*

Gelman M, Kariv S, Shapiro MD, Silverman D, Tadelis S. 2018. How individuals respond to a liquidity shock: evidence from the 2013 government shutdown. *Journal of Public Economics*

Giglio S, Maggiori M, Stroebel J, Utkus S. 2021. Five facts about beliefs and portfolios. *American Economic Review* 111(5):1481–1522

Gu Q, He J, Qian W. 2018. Housing booms and shirking. *Working Paper*

Hacioglu S, Kanzig D, Surico P. 2020. Consumption in the time of covid19: evidence from uk transaction data. *Working Paper*

Jorring AT. 2019. Financial sophistication and consumer spending. *Working Paper*

Karger E, Rajan A. 2020. Heterogeneity in the marginal propensity to consume: evidence from covid19 stimulus payments. *Working Paper*

Kim O, Parker J, Schoar A. 2020. Revenue collapses and the consumption of small business owners in the early stages of the covid-19 pandemic. *Working Paper*

Koijen R, Van Nieuwerburgh S, Vestman R. 2014. Judging the quality of survey data by comparison with truth as measured by administrative records: Evidence from sweden. *Improving the Measurement of Consumer Expenditure*

Kuchler T, Pagel M. 2018. Sticking to your plan: the role of present bias for credit card paydown. *Working Paper*

Kueng L. 2018. Excess sensitivity of high income consumers. *The Quarterly Journal of Economics*

Kukk M, Raaij WFV. 2020. Joint and individual savings within families: evidence from bank accounts. *Working Paper*

Levi Y, Benartzi S. 2020. Mind the app: mobile access to financial information and consumer behavior. *Working Paper*

Loos B, Meyer S, Pagel M. 2020. The consumption effects of the disposition to sell winners and hold losers. *Working Paper*

Medina PC. 2020. Side effects of nudging: Evidence from a randomized intervention in the credit card market. *Working Paper*

Meeuwis M, Parker JA, Schoar A. 2019. Belief disagreement and portfolio choice. *Working Paper*

Meyer S, Pagel M. 2019. Fully closed: individual responses to realized gains and losses. *Working Paper*

Odean T. 1998. Are investors reluctant to realize their losses. *The Journal of Finance* 53(5):1775–1798

Olafsson A, Pagel M. 2017. Family finances: Intra-household bargaining spending, and capital structure. *Working Paper*

Olafsson A, Pagel M. 2018. The liquid hand-to-mouth: Evidence from personal finance management software. *The Review of Financial Studies*

Olafsson A, Pagel M. 2019. Borrowing in response to windfalls. *Working Paper*

Park H. 2016. Financial fragility and mortgage delinquency. *Working Paper*

Ponce A, Seira E, Zamarripa G. 2017. Borrowing on the wrong credit card? evidence from mexico. *The American Economic Review*

Schuh S. 2018. Measuring consumer expenditures with payment diary surveys. *Economic Inquiry* (56):13–49

Weber A. 2019. Financial management skills and entrepreneurial success: evidence from transaction level data. *Working Paper*

Xing J, Zou E, Yin Z, Wang Y, Li Z. 2021. Quick response economic stimulus: The effect of small-value digital coupons on spending. *NBER Working Paper 27596*

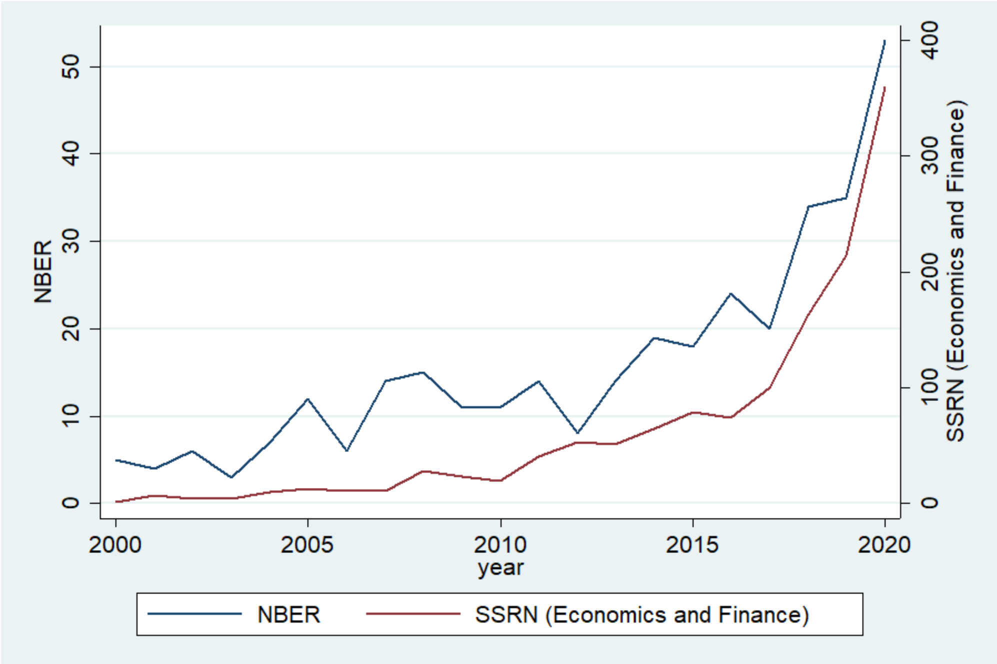

Figure 1: Counts of Papers Employing Transaction Data, by Year; Notes: Yearly counts obtained from searches of papers containing the terms ‘credit card data’, ‘bank account data’, ‘transaction data’ and ‘transaction level data’ from SSRN and the NBER working paper series from 2000 through 2020. Searches were conducted on the full text of articles. Searches limited to Economics and Finance papers in the SSRN database.

Table 1: Attributes of Consumption Data Sources

Properties / Data Source	Transaction	Survey	Admin	Scanner
Examples	JPMCI, Mint.com	PSID, CEX	Norwegian Tax Registry	Nielsen Consumer Panel
Frequency	Very High	Often Quarterly	Often Annual	High
Measurement error ($)	Low	Medium	Low	Low
Measurement error (category)	Medium	Low	High	Low
Sample size	High	Low	Very High	Medium
Comprehensiveness	Medium to Very High	Medium to High	High	Low
Demographics	Medium	High	Low	Medium
Geographic location	High	Low to High (confidential)	High	Medium
Representativeness	Varied	High (falling; tail issues)	Very High	Medium
Historical coverage	Low	High	Medium	Low
Panel length	Medium (growing)	Typically Low	High	Medium
Granularity	High	Medium	Low	Very High
Asset Coverage	Low to Very High	Medium	High	Low

Properties / Data Source	Banks	Credit Cards	Neutral Aggregators	Focused Aggregators	Payment Apps
Example	JPMCI	Visa	Mint.com	ReadyForZero	Alipay
Sample Size	High	High	High	Low	Medium
Account Linkages	High	Low	Very High	Medium	Low
Observable Income	Yes	No	Yes	Yes	No
Representativeness	High	High	Medium	Low	Low
External Links and Surveys	Medium	Low	Varied	Varied	Low
Legal Access Hurdles	High	Very High	Medium	Low	Medium

Table 2: Sample Sizes

Paper	Data Source	Sample Size & Composition
Various	Current Population Survey	60,000 households
Various	Consumer Expenditure Survey	7,000 households
Various	Nielsen Consumer Panel	40-60,000 households
Various	Living Costs and Food Survey	6,000 UK households
Carvalho et al. (2020)	BBVA	6.3m individuals
Baugh et al. (2018)	Large US aggregator app	2.7m individuals
Andersen et al. (2020d)	Danske Bank	1.1m individuals
Bachas et al. (2020)	Bansefi	348,802 individuals
Agarwal et al. (2020b)	Leading bank in Singapore	180,000 individuals
Ganong & Noel (2019)	JP Morgan Chase Institute	182,361 households
Kim et al. (2020)	JP Morgan Chase Institute	380,000 businesses
Baker & Yannelis (2017)	Large US aggregator app	152,810 gov. employees
Olafsson & Pagel (2018)	Meniga (aggregator app)	65,000 Icelandic households

Notes: Sample sizes obtained from summary statistics tables within listed papers. Survey data sample sizes are approximate.