FEDERAL RESERVE BANK of ATLANTA
RESEARCH DATA REPORTS

The 2016 and 2017 Surveys of Consumer Payment Choice: Technical Appendix

2018 ● No. 18-04

Marco Angrisani, Kevin Foster, and Marcin Hitczenko

Abstract: This document serves as the technical appendix to the 2016 and 2017 Surveys of Consumer Payment Choice administered by the Dornsife Center for Economic and Social Research (CESR). The Survey of Consumer Payment Choice (SCPC) is an annual study designed primarily to collect data on attitudes toward and use of various payment instruments by consumers over the age of 18 in the United States. The main report, which introduces the survey and discusses the principal economic results, can be found at frbatlanta.org/banking-and-payments/consumer-payments/survey-of-consumer-payment-choice. In this data report, we detail the technical aspects of the survey design, implementation, and analysis.

JEL classification: D12, D14, E4

Key words: survey design, sample selection, raking, survey cleaning, poststratification estimates

https://doi.org/10.29339/rdr2018-04

This version: August 2018

Marco Angrisani is an associate economist at the University of Southern California's Dornsife Center for Economic and Social Research. Kevin Foster is a survey methodologist and Marcin Hitczenko is a statistician, both in the Research Department of the Federal Reserve Bank of Atlanta. The coauthors' email addresses are marco.angrisani@usc.edu, kevin.foster@atl.frb.org, and marcin.hitczenko@atl.frb.org. This paper, which may be revised, is available on the website of the Federal Reserve Bank of Atlanta at frbatlanta.org/banking-and-payments/consumer-payments. To receive email notifications about new data reports, go to frbatlanta.org/forms/subscribe.

The authors thank their colleagues and management in the Consumer Payments Research Center and the Boston and Atlanta Fed research departments. In addition, we thank the management and staff at the CESR. We'd like to thank these people from the Boston Fed: Scott Schuh, Joanna Stavins, and Bob Triest. From the Atlanta Fed, we'd like to thank Claire Greene. And we'd like to thank these people from the CESR: Tania Gursche, Arie Kapteyn, Bart Orriens, and Bas Weerman. Special thanks go to Erik Meijer from the CESR, who contributed to earlier SCPC appendices, which formed the basis for this paper. Finally, the authors acknowledge John Sabelhaus and the staff of the Survey of Consumer Finances at the Federal Reserve Board of Governors for their advice and mentorship. Geoff Gerdes and May Liu from the Board also shared advice and knowledge.

1 Introduction

This document serves as the technical appendix for the 2016 and 2017 Surveys of Consumer Payment Choice (SCPC), an annual survey sponsored by the Consumer Payment Research Center (CPRC) at the Federal Reserve Bank of Boston (Boston Fed) since 2008. The CPRC is responsible for writing the questionnaire, analyzing the collected data, and publishing results. The programming of the survey instrument for online use, sample selection, and data collection is outsourced to an external survey vendor. From the initial version of the SCPC in 2008 until 2013, the CPRC worked exclusively with the RAND Corporation (RAND) and their American Life Panel (ALP). In 2014, in addition to RAND, the CPRC contracted with the Center for Economic and Social Research (CESR) at the University of Southern California to use their Understanding America Study (UAS) panel for additional observations. Since 2015, the CPRC has worked exclusively with the CESR and the UAS.

This document is designed to be the primary reference for the 2016 and 2017 SCPCs. The organization of this work is identical to that of previous years, based on the natural, chronological progression of considerations involved in conducting and analyzing a survey. As a result, comparisons of strategies, methodologies, and results across years can be easily done by referencing corresponding sections in earlier versions of the technical appendix.

We begin by establishing the context and goals of the survey in Section 2 and follow that by highlighting changes in the survey questionnaire from the 2015 version to the 2016 version and from the 2016 version to the 2017 version in Section 3. In Section 4, we detail the sample selection strategy in the context of that used in previous years and present statistics relating to survey response and completion. Section 5 delineates the methodology used to generate the sample weights, which are used to make inferences about the entire population of U.S. consumers. Section 6 discusses our general philosophy toward data preprocessing of categorical and quantitative variables and provides detailed algorithms for key data-editing procedures. Finally, in Section 7, we present the statistical methodology used for estimating and comparing population estimates.

2 Survey Objective, Goals, and Approach

In this section we describe the SCPC survey program's overall objectives, goals, and approach, and explain the choices made in selecting the observation unit and the interview mode of the SCPC. In both cases, the choice was influenced by best survey practices along with constraints imposed by the SCPC budget.

2.1 Survey Objective and Goals

The main objective of the CPRC is to measure U.S. consumer payment behavior, as outlined in greater detail in Foster, Schuh, and Zhang (2013). To this end, the primary purpose of the SCPC data program is to provide an annual consumer-level dataset to support research on consumer payments and to provide aggregate data on trends in U.S. consumer payments. Key attributes of the data are the longitudinality, in which responses are collected from many of the same individuals in subsequent years, and an increasingly strong link with the Diary of Consumer Payment Choice, a second source of payments information. The change in primary survey vendor necessarily implies a discontinuation of the longitudinality built with RAND from 2008 to 2014. Differences between the RAND panel and the UAS panel are discussed further in Section 4, and the challenges of such a change in panel is discussed in Section 7.1.

2.2 Unit of Observation

The SCPC uses the individual consumer as both the sampling unit and the observation unit. This choice stands in contrast to those of many other finance surveys, most notably the Survey of Consumer Finances, which is organized by primary economic units in the household, and the Consumer Expenditure Survey, which uses the household as the sampling unit and observation unit.

One practical reason that the SCPC focuses on the consumer is that it is less expensive to collect data about an individual than an entire household. Household surveys require either thorough interviews with all adult household members, which is logistically difficult, or having one selected household member enter data for the entire household. The latter strategy imposes a considerable burden on the household representative. Since SCPC incentives are based on the average length of time it takes respondents to complete the survey, the cost of each survey would increase if the household were the unit of observation. This, in turn, would limit the total number of responses within a fixed budget.

In addition, there is a methodological reason for the choice of observational unit. Namely, for many economic concepts on which the SCPC focuses, it seems that asking each respondent about his or her behavior rather than the entire household's is likely to yield more accurate data. Prime examples include information about getting, carrying, and using cash and the number of non-bill payments made in a month. It may be difficult for one household member to accurately report the behavior of other household members, and, even if asked, household members may not feel comfortable sharing such information with one another at such a level of detail. Therefore, it is most appropriate to ask the individual consumer about his or her own behavior and not about the habits of other household members.

Admittedly, the use of the consumer as the unit of observation may not be ideal for other variables, most notably the payment of bills or other expenses more closely associated with the household than an individual. For consistency within the survey instrument, the SCPC asks respondents to estimate only the number of bills that they physically pay themselves, either by mail, by phone, online, in person, or by having set up an automatic payment. In theory, this should produce accurate measurements on bill payment behavior when averaged across the entire sample. The difficulty comes from the fact that these kinds of payments may be paid out of joint accounts or pooled resources, or may have been set on an automatic schedule so long ago that no household member can recall who set up the automatic payments. As a result, it can be difficult to attribute responsibility for such payments, often leading to under-counting, if they are not reported at all, or double-counting, if several household members each claim responsibility for the same payment. In addition, research on SCPC data suggests that survey respondents are more likely to have a higher share of financial responsibility within the household than would be expected if household members were selected at random, and thus tend to be more likely to make certain types of payments than an average sample of the population (Hitczenko 2015). Treating such a sample as representative of all consumers may lead to overestimation of the number of bills paid. Therefore, to accurately measure bills, it might be preferable to ask about the entire household's bill payment behavior.

2.3 Interview Mode

The SCPC is a computer-assisted web interview (CAWI). This mode of interview fits best with our sampling frame, which is an internet-based panel. To maximize coverage, all ALP members are given internet access upon recruitment into the panel. The survey instrument is the NubiS survey system, developed by team members at the CESR.[1]

The CAWI mode is beneficial to the SCPC because of the length and complexity of the survey. The projected median length in minutes for the SCPC survey in each year is around

[1] More information on NubiS is available at https://cesr.usc.edu/nubis/node/2.

30 minutes. In 2016, the median time spent on survey screens was 37.6 minutes and the middle 50 percent spent between 28.5 and 49.1 minutes on the survey screens. Changes to the survey from 2016 to 2017, described in detail in Section 3, resulted in a significantly shorter survey. Specifically, in 2017, the median SCPC took 30.6 minutes, while the middle 50 percent took between 23.1 and 39.6 minutes. Using a CAWI allows the respondent to log off and come back to the survey later if interrupted. This is cheaper than using face-to-face interviews or telephone because there are no interviewers who need to be trained and paid. Finally, respondents may be more willing to answer some sensitive questions, like the amount of cash stored in their home, if the survey is conducted via the web (De Leeuw 2005).

2.4 Public Use Datasets

Individuals who are interested in downloading the original, unprocessed datasets can obtain these from the UAS website.[2] The SCPC website also contains a link to the UAS data download site.[3] Interested users must create a username and password to download data from the UAS website. These data contain only the survey variables found directly in the survey instrument itself. These survey variables have not been edited or processed. For example, survey items that allow the respondent to choose a frequency (week, month, or year) have not been converted to a common frequency. For those interested in using these data, we recommend identifying survey variables by directly finding them in the SCPC questionnaire, which can be downloaded from the SCPC website.

An extension of the data, which includes edited variables and new variables created by the CPRC, can be downloaded at the SCPC website as well. The data are available in Stata, SAS, and CSV formats. Information about the definitions and naming schemes for all new variables not found in the original dataset are described in the companion document, "SCPC Data User's Guide: 2016-17" (Foster 2018), which is also available at the SCPC website. Before using the data, it is useful to read the warning against using consumer-level estimates to aggregate up to U.S. total population estimates, in Section 7.2.1 of this paper.

The variable `prim_key` is the unique identifier for each respondent. This variable is used as the primary key for both SCPC and the Diary of Consumer Payment Choice (DCPC) datasets, and is based on the terminology established by RAND and used in SCPCs and DCPCs prior to 2015. We choose to continue with this nomenclature for internal consistency,

[2] 2016 SCPC https://uasdata.usc.edu/UAS-62, 2017 SCPC https://uasdata.usc.edu/UAS-105

[3] Because of a change in ownership of the SCPC and DCPC from the Federal Reserve Bank of Boston to Atlanta, all relevant information can now be found at:https://www.frbatlanta.org/banking-and-payments/consumer-payments.aspx

even though the UAS dubs the unique individual identifier as uasid. As a result, to merge any UAS dataset, including the the raw, uncleaned dataset, with any processed SCPC and DCPC dataset, the user must merge the variables prim_key and uasid. In some statistical software programs, this requires renaming one of the variables to match the other's name.

3 Questionnaire Changes

The SCPC questionnaire is written by the CPRC and is available to download at the SCPC website. For the most part, the survey questions for the 2016 and 2017 SCPC are the same as or similar to those in prior versions, although changes are introduced every year either to collect new information or to improve the data collection process for the same information. This section describes the changes in two parts, first, the changes to the questionnaire from 2015 to 2016, then the changes from 2016 to 2017.

3.1 Changes from 2015 SCPC to 2016 SCPC

Table 1: New questions in the 2016 SCPC, Assessment of Characteristics section

Variable ID	Question description
as004	How do you rate the security of the following means of making a payment?
as005_g	How would you rate the security of each type of debit card transaction? Online without providing a security code (CVV)
as005_i	How would you rate the security of each type of debit card transaction? During a voice telephone call, without security code (CVV)
as012	Please tell us which payment characteristic is most important when you decide which payment method to use.

Table 2: New questions in the 2016 SCPC, Virtual Currency section

Variable ID	Question description
pa133	In the past 12 months, have you exchanged virtual currency for US dollars or exchanged US dollars for virtual currency?
pa135	In the past 12 months, how many times did you exchange virtual currency for US dollars or vice versa?
pa128	In the past 12 months, have you used virtual currency to make a payment for goods or services to another person?

Table 3: New questions in the 2016 SCPC, other sections

Variable ID	Question description
pa045_c	In the past 12 months, have you authorized a text message payment using one of the following methods? Authorize your mobile phone company to pay for you.
ph025_g	Do you use any of the following online personal financial management services or apps to budget and monitor your spending, saving, or account balances? MoneyWiz
ph025_h	Do you use any of the following online personal financial management services or apps to budget and monitor your spending, saving, or account balances? GoodBudget
de012	Please tell us the total combined income of all members of your family living here during the past 12 months. (only asked if item de010 = 18, or "$500,000 or more")

Table 4: Questions that were edited from 2015 to 2016, all sections

Variable ID	Question description	Description of change
as005	How would you rate the security of each type of debit card transaction	Changed some of the text for existing questions, added two new questions.
pa001	Screen for reporting number of checking and savings accounts	Edited instruction text to be shorter and more concise.
pa001	Screen for reporting number of checking and savings accounts	Added definition of savings accounts
pa055	Table of questions to determine underbanked status	Edited to present items as two tables on the same screen instead of one
pa120	List of virtual currencies	Dropped Stellar and added Ethereum
pa129	Who did you pay using virtual currency?	Edited list of response options
pu011	How would you compare your unpaid balance on your credit card to your unpaid balance 12 months ago?	Added new response option – I did not have a balance 12 months ago
pa189_c	In the past 12 months, have you used a mobile phone to make any of these kinds of payments?	Edited the response for using a mobile app
pa001_e	List of mobile apps	Added Samsung Pay and Square Cash
pa001_f	List of mobile apps	Added Samsung Pay and Square Cash
ph004	In the past 12 months, have you, or anyone you know well, been a victim of identity theft?	Added definition of identity theft

Table 5: Dropped questions from 2015 to 2016, all sections

Variable ID	Question description
pa007_b	Is your savings account linked to your primary checking account?
pa015_c	About how much cash stored elsewhere are you holding for cash payments?
pa015_d	About how much cash stored elsewhere do you have set aside for long-term savings?
pa198_n	Prepaid card type: Other types of passes or membership cards (museum, gym, parking, recreation)
pa202_n	Prepaid card logos: Other types of passes or membership cards (museum, gym, parking, recreation)
pa203_n	Prepaid cards that can pay anywhere: Other types of passes or membership cards (museum, gym, parking, recreation)

3.2 Changes from 2016 SCPC to 2017 SCPC

Table 6: Dropped questions from 2016 to 2017, adoption section (Part One)

Variable ID	Question description
pa007	At what type of financial institution is your primary savings account?
pa006_a, b	At what type of financial institution is your [primary, secondary] checking account?
pa073_a, b	About how much money do you have in your [primary, secondary] checking account?
pa085_a, b	Average balance of [primary, secondary] checking account
pa086_a, b	Drop-down list with dollar ranges for balance of [primary, secondary] checking account
pa080_a, b	With whom do you share your jointly owned [primary, secondary] checking account?
pa022	Please choose the most important reason why you don't have an ATM card.
pa021	Please choose the most important reason why you don't have a debit card.
pa034	If you are given a choice while completing a debit card purchase, do you prefer to enter your PIN or give your signature?

Table 7: Dropped questions from 2016 to 2017, adoption section (Part Two)

Variable ID	Question description
pa055	In the past 12 months, did you use any of the following financial services? (Underbanked screen)
pa027	Please choose the most important reason why you don't have a credit card.
pa051_a-e	How many of each kind of these kinds of credit cards that are also branded with a company logo?
pa195	Please choose the most important reason why you don't have a general purpose prepaid card.
pa192	Do you use any phone apps that are funded by buying a prepaid card and entering the number on the card into your app?
pa109	Please choose the most important reason why you don't have any automatic bill payments set up.
pa001_d2, d3	Do you have an account with any of the following payment services? [Google Wallet, Amazon Payments]
pa044_b, c	In the past 12 months, have you used [Google Wallet, Amazon Payments] to make a purchase or pay another person?
pa047_a, b, c	About how much money do you have in your [PayPal, Google Wallet, Amazon Payments] account?
pa048_a2-e2, a3-e3	In the past 12 months, have you used any of the following methods to make payments with your [Google Wallet, Amazon Payments] account?

Table 8: Dropped questions from 2016 to 2017, cash section

Variable ID	Question description
Cash Q's	The entire cash section has been moved to Day 2 of the Diary
pa015_a	About how much cash do you have in your wallet, purse, and/or pocket?
pa015_b	About how much cash do you have stored elsewhere in your home, car, office, etc.
pa016	When you get cash, where do you get it most often?
pa017_a	When you get cash from [FILL WITH ANSWER FROM PA016], what amount do you get most often?
pa018_1	In a typical period (week, month, or year), how often do you get cash from [FILL WITH ANSWER FROM PA016]?
pa017_b	When you get cash from all other sources besides [fill from answer PA016], what amount do you get most often?
pa018_2	In a typical period (week, month, or year), how often do you get cash from all other sources besides [fill from answer PA016]?

Table 9: Dropped questions from 2016 to 2017, all other sections

Variable ID	Question description
as004	How do you rate the security of the following means of making a payment?
as005	How would you rate the security of each type of debit card transaction?
as012	Please tell us which payment characteristic is most important when you decide which payment method to use.
ph025_a-h	Do you use any of the following online personal financial management (PFM) service or app to budget and monitor your spending, saving, or account balances?

Table 10: New questions in the 2017 SCPC, adoption section

Variable ID	Question description
pa909	Have you ever had a store-branded card linked to your bank account? Yes/No
pa031_b	Have you ever had blank paper checks for any of your checking accounts? Yes/No
pa008_c	Number of store-branded cards linked to your bank account
pa042_a	Did you purchase any of the money orders you used in the past 12 months from a non-bank source?
pa042_e	Did you send any of the remittances you used in the past 12 months from a non-bank source?
pa041_e	Have you ever sent a remittance, even once?
pa120_b7	Have you heard of Bitcoin Cash (BCH)?
pa052	Do you own any kinds of credit cards that also are branded with a company logo?
pa001g1, 2, 3	In the past 12 months, have you used any of the following features of your bank's mobile banking app?

Table 11: New questions in the 2017 SCPC, all other sections

Variable ID	Question description
pu020	On your last bill(s), about how much were the new charges made to all of your credit cards and/or charge cards?
pu004_f1-3	Number of bill payments by mail, in person, or by phone paid using your bank account and routing numbers
pu021_g	Person-to-person payments made by using account-to-account payments using a nonbank service such as PayPal or Venmo

Table 12: Questions that were edited from 2016 to 2017, all sections

Variable ID	Question description	Description of change
text screens	Any screen with all text—instructions, definitions, etc.	Reduced the number of words on each of these screens
pa075_a, b	Is your primary [secondary] checking account jointly owned with someone else?	Changed from Yes/No to four different response options from pa080
pa050	In the past 12 months, have you used any of the following payment methods, even once? Cash	Moved to a table with other payment instruments
pa040_e	In the past 12 months, have you used any of the following payment methods, even once? Remittance	Moved to a table with other payment instruments
pa055_a2, b1-b5	Underbanked questions	Pa055_a2 is now on its own screen, Pa055_b1-b5 are on a separate screen in a table
VC section	All questions	Added ticker symbols for currencies
VC section	List of currencies	Removed Dash and Dogecoin, added IOTA and NEM
pu010, pu013	Credit card balance (pu010) and limit (pu013)	Moved these questions to be inside a skip condition
pa201	Do you have any of the following types of prepaid cards?	Changed from asking Yes/No to asking Number of Cards
pa001_e, f	Do you have any of the following mobile apps or online accounts?	Added Amazon Payments, Zelle, and "Your bank's mobile banking app"
pu005_b, pu006a_c, pu006c_c	Payment use section questions about debit cards	Added text about store branded linked debit cards
pu021_c, d	Person-to-person payments using debit cards and credit cards	Added text about Venmo
pu021_e	Person-to-person payments using account-to-account payments	Added text "using a service provided by your bank"

4 Data Collection

This section describes various aspects of the data collection process for the 2016 and 2017 SCPC. Once the survey instrument is finalized, the goal is to recruit a sample of respondents that can be effectively used to make inferences about the U.S. population of consumers and, then, effectively administer the survey to those individuals. The methodologies and their underlying philosophies adopted by the CPRC in this process are outlined below. In addition, outcome statistics related to the fielding of the survey are detailed. Similar expositions focusing on the previous editions of the SCPC can be found in the official releases by the CPRC (Angrisani, Foster, and Hitczenko 2013; 2014; 2015; 2016; 2017).

As in 2015, the SCPC in 2016 and 2017 relied exclusively on the Understanding America Study (UAS), a panel of respondents created and managed by the Center for Economic and Social Research's (CESR) at the University of Southern California, after a change in 2014 from RAND's American Life Panel. In the long term, the CPRC believes that the best practices used to construct and maintain the UAS panel, most notably address-based sampling for recruiting panelists, will provide for better population estimates in the future. The motivations for the change in survey vendor are discussed in greater detail in the 2015 SCPC Technical Appendix (Angrisani, Foster, and Hitczenko 2017). However, in the short term, panel effects, biases related to the particularities of respondent selection and survey administration associated with each panel vendor, introduce possible challenges to time series analysis (Kennedy et al. 2016). In Section 7 of this document, we suggest a statistical methodology for inference based on data collected from multiple panels.

4.1 Understanding America Study

The UAS originated in 2014 as a collection of individuals recruited to participate in a wide variety of surveys. The vast majority of panelists are recruited as part of the "Nationally Representative" subsample, intended to well-represent unincarcerated individuals aged 18 years and older who live in the United States, the target population for the CPRC. In addition, the UAS features two smaller subsamples constructed for specific research projects hosted by CESR: individuals of Native American origin and families with young children who live in Los Angeles county. The Native American cohort was used in 2015, when the small size of the UAS panel required it to ensure a large enough sample. Since then, no individuals from this cohort were recruited for the SCPC. Individuals from the Los Angeles county cohort have never been included in the SCPC sample, because effectively incorporating data from such a specialized subpopulation into general population estimates is difficult.[4] Therefore, the description below relates exclusively to the Nationally Representative subset of the UAS panel.

The UAS panel uses address-based sampling, in which zip codes and then addresses within those zip codes are drawn at random, to generate a probability panel. A detailed description of the recruitment process can be found at the UAS website https://uasdata.usc.edu/recruitmentoverview/1, but a brief summary is given below. After an initial introductory postcard, an invitation featuring a $5 prepaid card and a 10-minute paper survey is mailed to the selected households. An additional incentive of $15 is provided for the return of the survey. The use of mail in the initial outreach ensures that all households, even those without internet or an online presence, are included. Those who join the panel but do not have internet access are eventually provided with internet access and a tablet with which to take surveys. Multiple stages of correspondence, featuring reminders, follow-up surveys, and additional financial incentives, are conducted through mail, phone, and the internet to build a relationship with the individual and encourage enrollment in the panel. A key feature of the recruitment process is that individuals are encouraged and incentivized to enroll fellow household members. As a result, about 18 percent of enrolled households featured multiple members in the panel.

Recruitment into the UAS panel has been conducted in waves by the CESR. The initial recruitment wave, commenced in February 2014, was also the largest with 9,284 households selected for contact. An additional wave of 1,799 households were invited in September 2015, followed by seven waves of recruitment from January 2016 to June 2016 that invited from around 1,800 to around 3,500 households each. The waves conducted in 2016 were used to improve demographic representation of the panel by favoring ZIP codes more heavily populated with demographic groups lacking in the panel. In addition, the CESR conducted experiments on the details of the recruitment process itself, such as the incentives and types and order of correspondence. In general, the recruitment waves had about 15 percent success rate, defined as the percent of individuals contacted who eventually became panel members. The recruiting effort in 2016 resulted in a sizeable increase in panel size, up to 4,776 at the time of sampling in 2016 from 2,140 at the same time in 2015. The panel size decreased by 18 people to 4,759 at the time of sampling in 2017, a net result of attrition, which is less than 10 percent annually, and a lack of additional recruitment efforts other than allowing household members to join. Thus, between August 2016 and August 2017, the only new

[4] To see how the Native American cohort was incorporated into the sample via poststratification weights, see the 2015 Technical Appendix.

panelists were additions from households with a member already featured in the UAS.

Once in the panel, UAS members are eligible to be recruited for surveys, unless they formally ask to be removed or stop participating in surveys over a prolonged period of time. At the beginning of each year, the CESR contacts all members who did not take any survey for at least a year and removes them from the panel, unless they explicitly declare continued interest in participating. Since inactive members are removed only once a year, the pool of those invited in August to answer the survey at a given point in time may include inactive members.

4.2 SCPC Sample Selection

Sample selection for the SCPC is done jointly with the Diary of Consumer Payment Choice (DCPC), a second data collection instrument developed by the CPRC and co-sponsored by the Federal Reserve Banks of San Francisco and Richmond. The DCPC is a diary that asks individuals to track details of all payments and financial transactions over the course of three consecutive days in October. As discussed in Section 2, the DCPC has increasingly grown to be a complement to the SCPC. In fact, the 2016 and 2017 version of the DCPC pulls in information collected in the SCPC. Because of this strengthening connection, since 2015, the CPRC has introduced a strict requirement that respondents must take the SCPC prior to their assigned diary period. Selection of respondents is primarily based on trying to maximize the number of respondents that participate in both surveys.

Traditionally, the fielding of a survey such as the SCPC by the CESR simply involves making it accessible to each potential respondent online on some specified release date, with notification of its availability coming in the form of an email and link. However, the unique nature of the DCPC, which distributes responses throughout an entire month and requires prior commitment from individuals in order to mail diary-related materials, requires an additional consent survey. This consent survey is treated like a typical survey fielded by the CESR, released at the end of August with a $5 incentive for individuals. The incentive was introduced in 2016 in an effort to increase the percentage of individuals who answer the consent survey. Although fielding this consent survey with an incentive reduces the budget available for the SCPC and DCPC, experience has shown that without the incentive a much lower percentage of panelists respond, thus making planning for the DCPC much more difficult. Based on the improved results from 2015, a year in which no incentive was used for the consent survey, to 2016, the incentive was continued in 2017. The consent survey describes both surveys, including the $20 incentive for the SCPC and the $70 incentive for the DCPC, and asks respondents if they are willing to take both surveys. Those who decline are asked for a reason why and are presented with a second request that aims to assuage their concerns. In 2016, individuals who did not agree to participate in both surveys were asked if they would be willing to take the SCPC only. In 2017, because of the high consent rate and the smaller target sample size, this option was removed, so that only those willing to take both surveys were able to take the SCPC.

The selection of the individuals to whom the consent survey is sent is governed by two basic principles. First, it is our desire to have the demographic composition of the sample match that of the U.S. population of adults as closely as possible. To ascertain this, the population is divided into 30 disjoint strata defined by household income (3 groups), age (3 groups), gender (2 groups), and race (2 groups), and sample proportions in each stratum are compared to those implied by the March supplement of the Current Population Study of the relevant year. The second principle is a desire to maintain a longitudinal panel across years. The benefits of such a longitudinal panel, in the form of added power associated with tracking trends at the individual level, are well known (Baltagi 2008; Duncan and Kalton 1987; Frees 2004; Lynn 2009). For many research agendas, it is advantageous to base results on a longitudinal panel, rather than on a sequence of cross-sectional studies. The CPRC does not construct a formal longitudinal panel, governed by rules for replacing members who drop out in order to maintain similar panel composition. Instead, we simply put a heavy emphasis on inviting respondents who participated in previous years. Because the participation rates in the UAS panel are quite high, this somewhat informal strategy effectively yields a longitudinal panel.

The number of desired respondents for the joint SCPC and DCPC is determined by the available budget and the cost of management, programming, and data collection in a given year. In 2016, the desired number of respondents was 3,300, while in 2017 it was 3,150. These numbers along with the predicted participation rates, estimated to be between 60 and 70 percent, inform the number of individuals necessary to invite to achieve the target number of respondents. In 2016, getting an adequate sample size meant inviting the entire set of nationally representative individuals, so no selection according to demographics was done. In 2017, initial calculations revealed that perhaps not all panelists would need to be invited. So, stratified quota sampling was used to determine the number of respondents in each of the 30 demographic strata that would need to be invited to yield a representative sample of the desired size. To maximize longitudinality, respondents were classified into two tiers according to whether they had any prior experience with the SCPC or DCPC or not. The invitation survey was first sent out to the first tier, which consisted of 3,677 individuals, so that spots that could be filled with individuals with more experience were not filled by new respondents simply because they responded faster. Based on observed participation rates among the first tier, it was determined worthwhile to send the second wave of invitations to all of the remaining 1,082 panelists who had no prior experience with the survey. This second wave was sent two weeks after the first wave. Although the strategy was effectively the same in 2017 as in 2016, with all panelists sent an invitation, the tiering used in 2017 shows how longitudinality can be maintained in future years as the UAS panel increases in size.

Table 13 shows the number of individuals invited and the number who consented to participate in 2016 and 2017. In both years, the pattern of consent is similar. The major barrier to consent was getting panelists to take the invitation survey, as only 74 percent of individuals did so in 2016 and 68 percent did so in 2017. However, willingness to participate among those who took the consent survey was very high. In both years, about 94 percent of respondents agreed to participate in the SCPC and DCPC, yielding an overall consent rate of close to 65 percent. Additionally, over 98 percent of those who consented to both surveys did so with the initial invitation, without requiring assurances about their concerns. In 2016, of the 211 individuals who did not agree to the joint survey, 105 or nearly 50 percent agreed to take the SCPC only.

Table 13: Number of survey invitations, consents, and participations for the 2016 and 2017 SCPC.

	Year	
	2016	2017
Panel Size	4,776	4,759
Sent Invitation	4,776	4,759
Took Invitation Survey	3,572	3,293
Consent to Both		
Consented to Initial Invite	3,291	3,110
Consented to Secondary Invite	70	48
Consented to SCPC Only	105	NA
Total	3,466	3,158
Started the SCPC	3,404	3,099
Completed the SCPC	3,386	3,075

Source: Authors' calculations

Figure 1 shows the sizes of various panels within the UAS framework dating back to the 2014 SCPC survey. The most notable feature of the figure is the significant jump in sample size from 2015 to 2016, an increase of around 150 percent. As a result, while the four-year panel is composed of around 600 individuals, the two-year panel starting in 2016 is much larger, with over 2,600 respondents. The expectation is that the longitudinal panel beginning in

2016 and stretching into the future year will continue to be over 2,000 individuals.

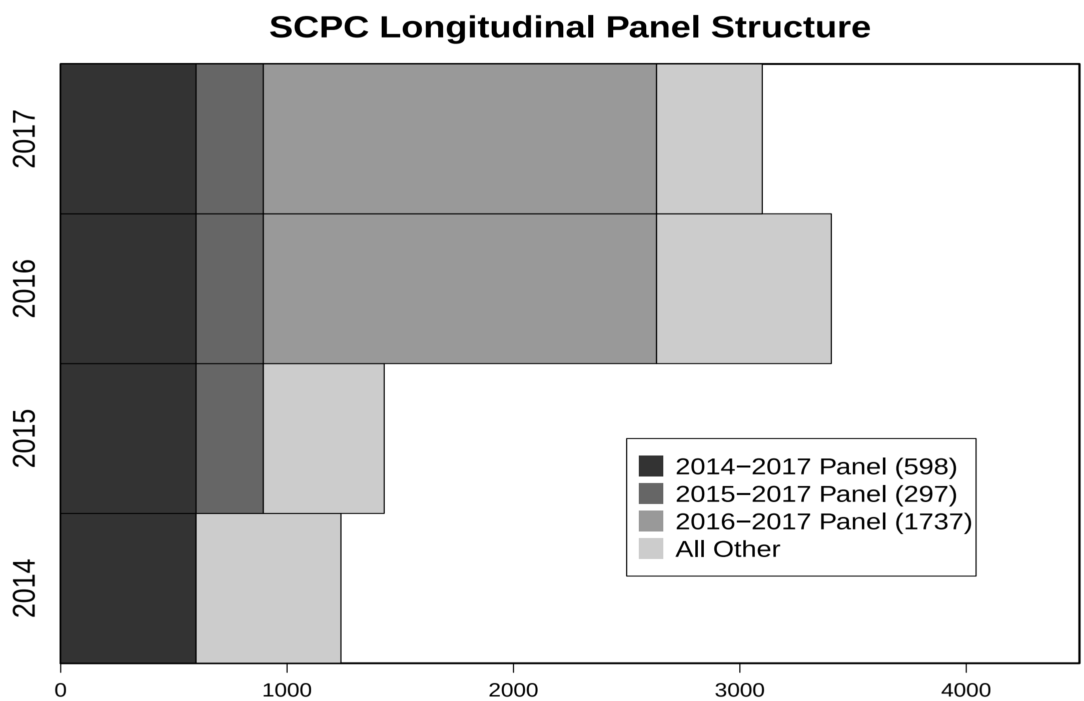

Figure 1: The longitudinal panel structure from 2014 to 2017.
Source: Authors' calculations

4.3 Survey Participation

The CPRC has generally limited the survey release to a period in the fall, ranging from the end of September through October. From an economic point of view, this time of year is a reasonably representative period of time with respect to certain economic variables such as employment or sale volumes; it includes no major holidays and falls between summer and winter. For the same reason, all previous versions of the DCPC, with the exception of the 2015 version, were also administered in October, with the added incentive that responses from both surveys could be linked more easily if they correspond to the same period of economic activity.[5] The increasing link between the SCPC and the DCPC, in which the former is a prerequisite to the latter, effectively forces the release of the SCPC into the middle of September.

[5] The 2015 DCPC was delayed in administration due to complications of coordinating with multiple vendors. More information is provided in the 2015 Technical Appendices.

The official release of the SCPC involves emailing a notification and a survey link to all respondents who have consented to participate. Each respondent can begin the survey at any point after receiving the notification of the SCPC release. In both 2016 and 2017, the SCPC was released on September 19th in order to give the earliest diarists 10 days to complete the SCPC before the first module of the DCPC, which can occur as early as September 29th. Participants do not have to finish the survey in one sitting and have the option to log off and continue later. Respondents were given reminders every so often at the discretion of the CESR if they had not logged on to the SCPC. In particular, reminders were sent out a few days before the assigned DCPC dates to each respondent to ensure that individuals have completed the SCPC by the time their DCPC commences.

Starting the survey, which is defined as logging on to the survey, is an important threshold because everyone who does so is considered a participating respondent and is assigned a survey weight. The lower half of Table 13 shows that, in both 2016 and 2017, 98 percent of respondents who consented to participate in the SCPC started the survey. A second metric used to gauge participation is that of completion, which we define as logging off after the final survey screen. It is important to note that logging off may not accurately reflect total completion of the survey, as it is possible to finish the survey without logging out. Other standards to define survey completion can be used. For example, one such standard would be answering all of the SCPC questions and reaching the last screen, which asks individuals for feedback on the survey questionnaire itself, but not logging out. Indeed, reaching the last question is the minimum requirement for the respondent to receive the financial incentive. As can be seen in Table 13, fewer than 1 percent of those who started did not complete the survey, which marks an improvement from previous years when 2 to 3 percent did not complete the survey.

As expected due to the joint recruitment effort, the sample overlap between the SCPC and DCPC is high in both years. In 2016, 90 percent (all but 356) of individuals who participated in the SCPC also took the DCPC. In 2017, this number was 93 percent (all but 228). However, if one excludes the 105 individuals who only committed to the SCPC in the first place in 2016, the overlap rate between both samples is very similar in both years. More information about how DCPC diary periods relate to the SCPC will be provided in the technical appendices for the DCPCs for each year.

Figure 2 shows the proportion of surveys completed as a function of the number of days since the survey was distributed for the 2012 – 2017 versions.[6] Throughout the life of the

[6]Similar plots that include data from 2008 to 2011 can be found in earlier versions of the SCPC Technical Appendix.

survey, the pace of completion for the 2016 and 2017 SCPC was generally much faster than in previous years and very similar within the two years. Whereas about 50 percent of respondents had started the SCPC within the first three days in 2016 and 2017, only 30 to 40 percent had done so in previous years. Within a month of release, the completion rate is around 95 percent in the two most recent years, as compared to below 90 percent in previous years. A higher completion rate is also reached, as noted above. A major part of this trend might be the necessity of taking the SCPC before the DCPC, which results in a greater number of reminders and a greater incentive to do so.

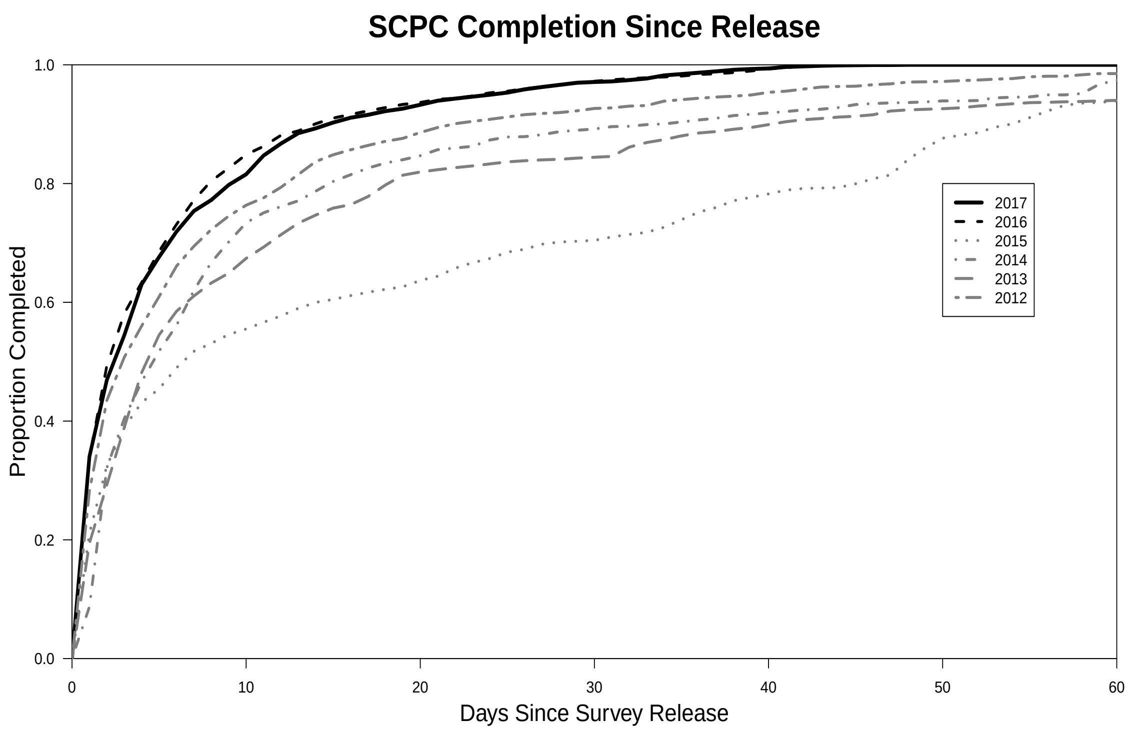

Figure 2: The proportion of respondents who completed the survey as a function of the number of days since the survey was received.

Source: Authors' calculations

Figure 3 shows the proportion of surveys completed by each calendar day within each of the years from 2012 to 2017. This depiction reveals the relatively wide range of time frames in which the bulk of annual responses are collected. In earlier versions, a majority of data responses were collected at the end of September, while in 2014 and 2015, most data were gathered near the middle of October. In 2016 and 2017, responses were collected even earlier, with most coming in the middle part of September.

An ideal survey design would specify that responses should be collected in a way that standardizes the response period across years. From an analytical point of view, trends from year to year are more easily identified if differences in behavior are not attributable to seasonal behavioral variation. Although the SCPC asks respondents about behavior in a "typical" month to reduce seasonal effects, it is possible that recent activity may influence responses. For example, to the extent that taking the survey near the beginning or end of a month influences responses, one should be aware of this when comparing results across years. In addition, if typical behavior changes in November due to the ensuing holiday season, payment use responses for a larger fragment of the 2015 SCPC sample than in other years will reflect this. Optimal survey fielding may involve accounting not only for the time of year and the time within a month, but also the time within a week. This is difficult to do, as respondents must be given a reasonable window of time in which to take the survey. The observed temporal gaps are even more extreme at the individual level, where a particular respondent might respond in October of one year and as late as January in a different year. Again, this raises potential issues of comparability.

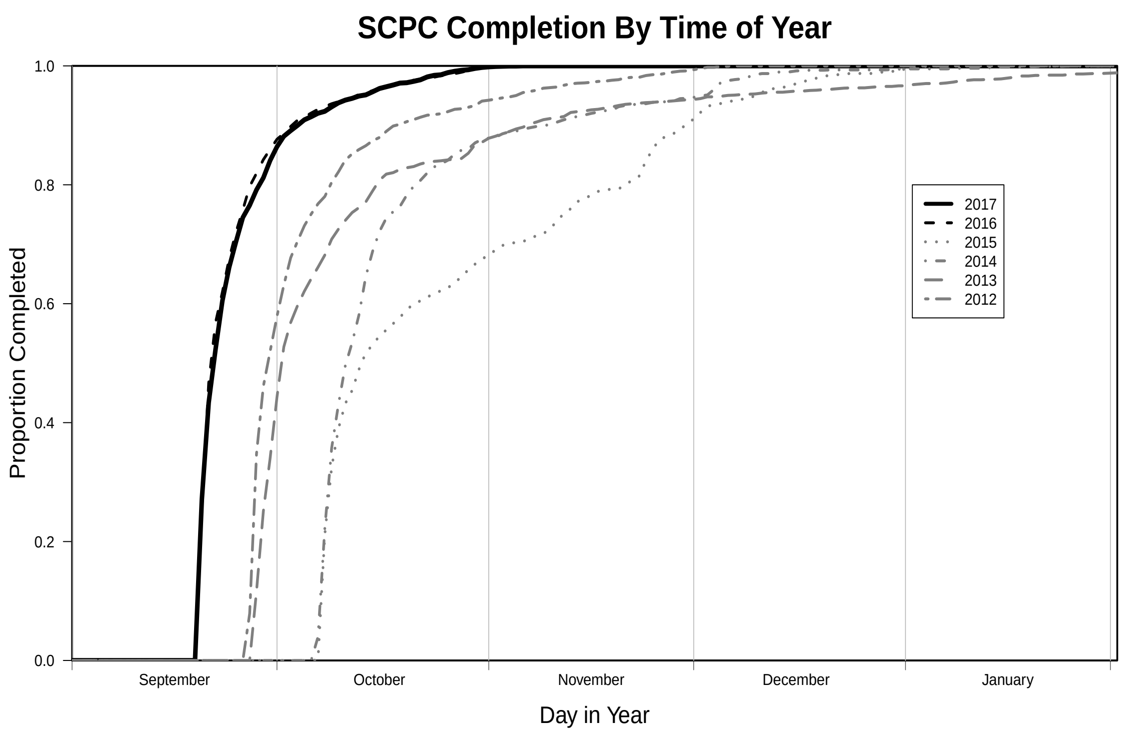

Figure 3: The proportion of respondents who completed the survey as a function of the date within the year.

Source: Authors' calculations

Figure 4 compares the distributions of the number of minutes it took respondents to complete the survey for the past five years of the SCPC, defined here as the difference in minutes between the time of first log-in to the survey and the last log-out. Individuals who take breaks while taking the survey will thus have long completion times, yielding the skewed-right nature of the observed timing curves. Changes in median time of completion generally correspond to changes in the length or complexity of the questionnaire. Although the 2013 survey has a median completion time of 32 minutes, almost every respondent provided additional information in a follow-up survey, dubbed Module B, which had a median time of 15 minutes (see Angrisani, Foster, and Hitczenko (2015) for details). The 2014 survey was designed to be considerably shorter than all previous versions (and was not paired with a follow-up survey) and has a median completion time of 29.5 minutes. A detailed description of the changes, mostly the removal of questions, from 2013 to 2014 made can be found in the 2014 Technical Appendix (Angrisani, Foster, and Hitczenko 2016). The 2015 and 2016 versions reincorporated a lot of information from Module B, and the 2016 SCPC began asking more detailed questions that were fed to the DCPC, such as checking account balances. The 2017 version was again made shorter, as documented in more detail in the following section, with a result of shortening the median time of completion. Of course, other factors, such as greater familiarity with the survey questions by the repeat respondents, may also contribute to changes in completion time distributions. Finally, in 2015, there was evidence that completion times depended on the speed of survey loading and the CESR servers, an effect that could influence observed distributions in any given year.

4.4 Item Response

For a survey to provide a valid picture of the overall population, it is very important that the item response rates for each question be high. High nonresponse rates not only mean there is less information on which to base estimates but also raise concerns about potential bias in the estimates. If the fact that an observation is missing is independent of the value of the observation, a condition referred to as "missing at random" (Little and Rubin 2002), imputation procedures can be used to generate estimates of sample statistics. However, if there is a confounding variable that relates to both the value of a variable and the likelihood of nonresponse, it is impossible to adjust for the effects on sample statistics. Certain economic variables, such as net worth or personal cash holdings, are potentially sensitive topics, and it is possible that there is a correlation between the true values and the willingness of respondents to provide these values. Naturally, variables with low nonresponse rates are less susceptible to this type of bias.

The SCPC has over 200 survey variables, although the survey itself is administered with a

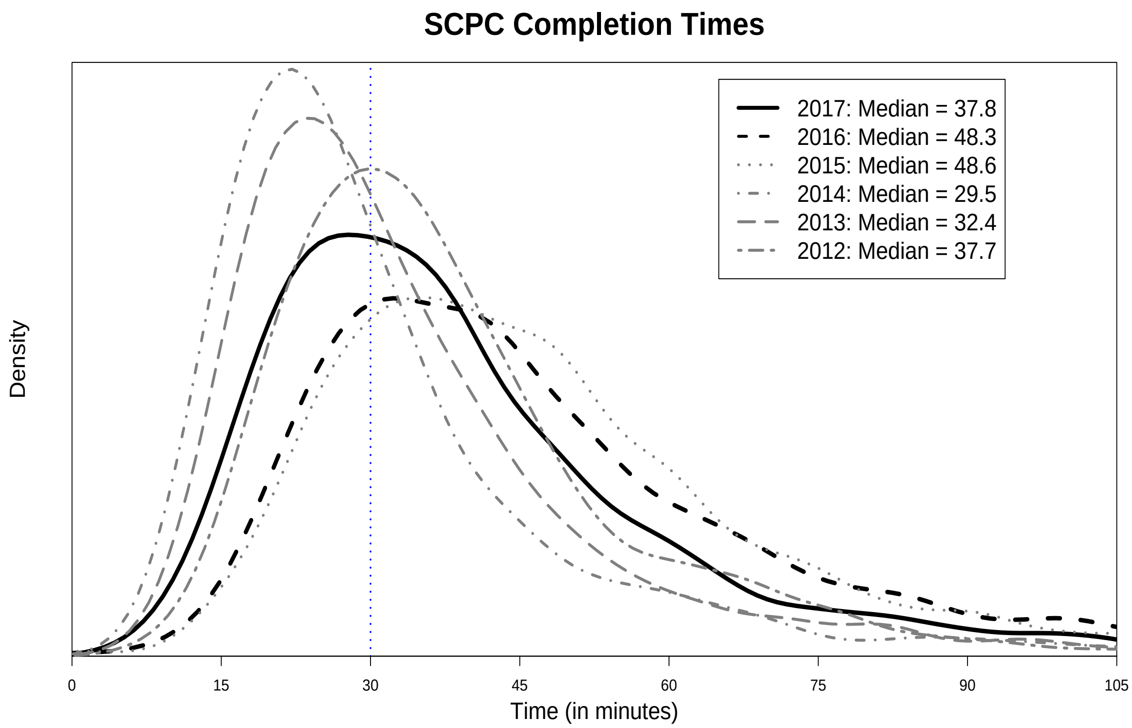

Figure 4: The proportion of respondents who completed the survey as a function of time. The vertical line at 30 minutes represents the intended average length of completion.

Source: Authors' calculations

relatively complicated skip logic so not everyone answers the same set of questions. However, taking a set of eight questions asked of everyone, dispersed throughout the survey, we found item nonresponse rates were low in both years, as shown in Table 14. Seven out of the eight nonresponse rates were under 1 percent. The exception is a question about the number of checking accounts (`pa001_a`), with around 3 percent of respondents not answering in both years. One possible explanation is that individuals who are nonadopters leave the answer box blank rather than writing in zero. Overall, the response rate is very high within the SCPC, which may be partly attributable to the fact that respondents have volunteered to take surveys and are being paid to do so. In 2016 and 2017, slightly less than 97 and 96 percent of respondents, respectively, answered all eight of the selected questions.

Table 14: Nonresponse rates (%) for eight questions in the 2016 and 2017 SCPC. The exact text of the corresponding questions can be found in the 2015 SCPC Questionnaire.

Question	fr001_a	as003a4	pa001_a	pa050	pa053	pa024	ph006	de011
Section in SCPC	II	III	IV	IV	IV	IV	VI	VII
2016 SCPC	0.23	0.32	2.82	0.47	0.47	0.53	0.56	0.73
2017 SCPC	0.23	0.45	3.10	0.94	0.35	0.48	0.71	0.87

Source: Authors' calculations.

5 Sampling Weights

5.1 Post-Stratification

An important goal of the SCPC is to provide accurate estimates of payment statistics for the entire population of U.S. consumers over the age of 18. Although the UAS panel uses address-based sampling, a form of probability sampling, relatively low response rates may mean that the final set of respondents taking the SCPC may not be a good representation of the target population. As a rough estimate, consider that only around 15 percent of individuals invited joined the UAS panel, and, of those, around 65 percent took the SCPC, yielding an overall response rate of 9.8 percent. If there are systematic differences in the likelihoods that a randomly selected individual ends up taking the SCPC across demographics, this could manifest itself in a final sample that looks different from the original population of invitees. Even relatively minor shifts in demographic composition can lead to bias, if not accounted for, for economic variables that show a lot of heterogeneity across demographic strata, as is the case with many payments variables (Stavins 2016).

Table 15 shows the unweighted and weighted marginal proportions of various demographic groups in the 2016 and 2017 SCPCs. Overall, the demographic composition of the samples in the two years is very similar, as should be expected from the fact that most respondents feature in both samples. Based on the weighted results, the raw sample overrepresents females, older individuals, people who identify as white, and well-educated individuals. This reflects imbalances in the UAS panel itself as there is no evidence that there are strong demographic effects relating to propensity of consenting once invited. Slight deviations in the weighted demographics are partly due to changes in the true population values and partly due to the differences in the unweighted sample compositions in the two years.

Optimal allocation of respondents across strata depends on the degree of variation of responses across demographics, with strata that have greater variability in responses requiring more observations. If the variance of responses is fixed across demographic strata, the most efficient estimate will be based on a sample in which the number of responses for each stratum is proportional to its overall frequency in the population. For this reason, without a priori knowledge about demographic patterns, proportional representation of strata in the sample is most often the goal of sample selection. Nevertheless, work by Wang et al. (2009) suggests that nonrepresentative polling can provide relatively accurate estimates with appropriate statistical adjustments.

Table 15: Unweighted percentages for various marginal demographics in the 2014 and 2015 SCPC sample, as well as weighted percentages for the 2015 SCPC. The weighted values are based on CPS data.

Demographics		Unweighted 2016 SCPC	Unweighted 2017 SCPC	Weighted 2016 SCPC	Weighted 2017 SCPC
Gender	Male	43.2	44.0	48.3	48.2
	Female	56.8	57.0	51.7	51.8
Age	18−24	4.1	3.2	6.8	5.1
	25−34	15.8	14.9	23.2	24.6
	35−44	19.6	19.8	16.3	16.2
	45−54	20.6	18.2	17.4	17.1
	55−64	22.9	23.5	16.8	17.0
	65 and older	17.0	20.4	19.5	20.1
Race	White	84.2	83.5	74.3	74.2
	Black	8.5	9.1	13.3	13.7
	Asian	1.8	1.8	3.1	3.6
	Other	5.5	5.6	2.2	1.4
Ethnicity	Hispanic	6.8	6.6	12.1	11.9
Education	No HS diploma	4.6	4.8	7.3	7.0
	High School	19.5	19.8	33.4	33.0
	Some College	38.9	38.0	28.5	28.4
	College	21.5	21.4	17.3	18.0
	Post-graduate	15.5	15.9	13.4	13.5
Income	$< \$25K$	21.3	20.0	21.1	18.3
	$\$25K - \$49K$	23.4	23.7	23.8	23.4
	$\$50K - \$74K$	20.3	20.1	17.8	19.4
	$\$75K - \$99K$	13.1	13.8	12.0	13.2
	$\$100K - \$124K$	8.9	9.0	10.5	10.2
	$\$125K - \$199K$	9.7	9.9	11.1	11.7
	$\geq \$200K$	3.4	3.4	3.7	3.8

Source: Authors' calculations

To enable better inference of the entire population of U.S. consumers, SCPC respondents are assigned post-stratified survey weights designed to align as much as possible the composition of the SCPC sample with that of a reference population. Specifically, each year the benchmark distributions against which SCPC surveys are weighted are derived from the CPS. This follows common practice in other social science surveys, such as the Consumer Expenditure Survey (CES).

5.2 Raking Algorithm

Sampling weights are generated by the CESR, using a raking algorithm (Deming and Stephan 1940; Gelman and Lu 2003) that is very similar to that used by RAND in previous years. This iterative process assigns a weight to each respondent so that the weighted distributions of specific socio-demographic variables in the SCPC sample match their population counterparts (benchmark or target distributions). The weighting procedure consists of two main steps. In the first part, demographic variables from the CPS are chosen and mapped onto those available in the SCPC.

Table 16 shows the 31 variables used in weighting as well as the levels within each variable. The number of levels for each variable should be small enough to capture homogeneity within each level, but large enough to prevent strata containing a very small fraction of the sample, which could cause weights to exhibit considerable variability. The socio-economic variables chosen for the raking procedure result from recent research conducted by RAND regarding the sampling properties of weights based on different demographic factors. Sample weights produced by different combinations of variables were evaluated on the basis of how well they matched the distributions of demographic variables not used as raking factors (test variables). To assess the robustness and accuracy of different combinations of weighting variables, Monte Carlo samples were drawn and demographic distributions of the test variables were generated based on the weights for that particular sample. Mean deviation from the CPS-defined levels for test variables was estimated by averaging over the samples. The combination of variables in Table 16 consistently matched the target distributions of the CPS for a variety of different sample sizes.

The pairing of gender with other socio-demographic variables allows one to correct better for discrepancies between distributions within each gender, while avoiding the problem of small cell counts. In other words, implementing the raking algorithm on the set of pairs shown in Table 16 ensures that the distributions of age, ethnicity, and education in the SCPC are matched separately for men and women to their population counterparts in the

CPS. Moreover, since bivariate distributions imply marginal distributions for each of the two variables, this approach also guarantees that the distributions of gender, age, ethnicity, and education for the entire SCPC sample are aligned with the corresponding benchmarks in the CPS. The same is true for household size and household income.

Table 16: The set of weighting variables. "M" stands for male, and "F" stands for female. The highest income brackets for single households were combined to avoid small cell sizes.

Gender × Age				
M, $18-32$	M, $33-43$	M, $44-54$	M, $55-64$	M, 65+
F, $18-32$	F, $33-43$	F, $44-54$	F, $55-64$	F, 65+

Gender × Ethnicity	
M, White	M, Other
F, White	F, Other

Gender × Education		
M, High School or Less	M, Some College	M, Bachelor's Degree or More
F, High School or Less	F, Some College	F, Bachelor's Degree or More

Household Size × Household Income			
Single, $< \$30K$	Single, $\$30K - \$59K$	Single, $\geq 60K$	
Couple, $< \$30K$	Couple, $\$30K - \$59K$	Couple,$\$60K - \$99K$	Couple, $\geq \$100K$
≥ 3 , $< \$30K$	≥ 3 , $\$30K - \$59K$	≥ 3 ,$\$60K - \$99K$	≥ 3 , $\geq \$100K$

In the second step, the raking algorithm is implemented and sample weights are generated by matching the proportions of predefined demographic groups in the SCPC to those in the CPS. Missing information about education, household size, and income are imputed, if necessary, using ordered logistic regression with gender and age as predictors. Race is imputed using logistic regression. The procedure is sequential, so that variables with the least number of missing values are imputed first and, in turn, used as inputs to impute the variables with the most missing values. Imputations are performed by ordered logistic regression for ordered categorical variables, and by multinomial logistic regression for categorical variables. The order of imputation moves across variables with the least number of missing entries to those with the most, and each level of imputation uses all previous demographic variables. In the case of SCPC, income is the most commonly missing variable. In 2016, 54 individuals had some demographic characteristic imputed, while 36 respondents did so in 2017.

We describe the raking algorithm in greater detail below. Let $d = 1, \ldots, 4$ represent the four bivariate demographic categories presented in Table 16. For each demographic, we index the possible groups with the variable g, so that d_g represents group g for demographic category d. For example, $d = 1$ corresponds to the intersection of gender and age and group $g = 1$ might correspond to males between 18 and 32 years of age. We use the shorthand notation $i \in d_g$ to indicate that individual i belongs to group g in demographic d. Then, we define $w_i^{(k)}$ to be the weight assigned to individual i after iteration k of the raking algorithm. Within each such iteration, the raking algorithm iterates across demographic groups d, so we let $w_i^{(k,d)}$ define the assigned weights after iterating over d.

In 2016, as in previous years, the base weights serve only to distinguish the nationally representative subset of the UAS panel from those recruited through targeted efforts. In practice, because of the sampling design, this means that every base weight in 2016 has a value of 1.0. In 2017, however, CESR began using more advanced processes for base weights. Sampling of households is based on an initial sampling of U.S. ZIP codes, which is an iterative process that depends on demographic compositions of the U.S. ZIP codes (provided by American Community Survey and the 2010 Urban Area to ZIP Code Tabulation Area). The result is that ZIP codes are not equally likely to be selected, thus even within strata, there is a lack of uniformity. The idea behind the base weights is to adjust for the fact that certain individuals are more likely to be sampled due to residence in more highly sought ZIP codes. Thus, the base weight for each individual corresponds to the inverse of the estimated likelihood of that individual's household being selected. This is captured by the product of the likelihood that an individual's ZIP code was selected and the likelihood of that individual's household being conditionally selected from the ZIP code. The former is modeled as a function of characteristics such as census region, population density, population size, sex, race, age, marital status, and education composition of the ZIP code. The latter is simply a ratio of the number of households selected to the number of households in the ZIP code, because household selection within selected ZIP codes is done at random.

Within each iteration, $k = 1, 2, 3, \ldots$, we do the following, mirroring the algorithm found in Valliant, Dever, and Kreuter (2013):

1. $w_i^{(k,0)} = w_i^{(k-1,4)}$
2. Otherwise, for $d = 1, \ldots, 4$, we let $w_i^{(k,d)} = w_i^{(k,d-1)} m_{k,s_d[i]}$, where

$$m_{k,s} = \frac{\sum w_i^{(k,d-1)} 1\left[i \text{ in stratum } s\right]}{\sum w_i^{(k,d-1)}} \times f_{d,g},$$

   where $f_{d,g}$ represents the proportion of the U.S. population that belong in group g of demographic d and where $1\left[i \text{ in stratum } s\right]$ is 1 if individual i belongs in stratum s and is 0 otherwise. This ensures that, after iteration d, the weighted marginal frequencies in the sample for demographic d will match perfectly those in the population.

3. Trim weights by letting $\bar{w}^{(k)}$ represent the average weight within the sample and then assign weight values according to:

$$w_i^{(k,4)} = \begin{cases} 0.25\bar{w}^{(k,4)}, & \text{if } w_i^{(k)} < 0.25\bar{w}^k \\ 4\bar{w}^{(k,4)}, & \text{if } w_i^{(k,4)} > 4\bar{w}^k \\ w_i^{(k,4)} & \text{else.} \end{cases} \tag{1}$$

Therefore, within each iteration, weights that are less than a quarter of the average or more than four times the average are trimmed. Trimming is performed to decrease large weights and increase small weights, thereby decreasing the variation in the weights. While this may sacrifice unbiasedness of estimators it does so by reducing the mean-squared error, which is adversely affected by high variations in weights.

The CESR runs 50 iterations of this algorithm or until all marginal weight matching and trimming specifications are achieved. The algorithm for the SCPC data converges in both years. Upon convergence, we let w_i represent the weight given to individual i. Weights are standardized to have a mean of 1.0, so the maximum weight is 4.0 and the minimum weight is 0.25. Overall, there are 508 unique weights, meaning 508 out of a possible 660 strata implied by the demographics in Table 16 are represented in the 2016 sample. The change in base weights in 2017 leads to 1,786 unique weights, though around the same number of demographic strata are represented. The standard deviation of weights is 0.8 in 2016 and 1.05 in 2017. The increase represents one greater than what one would expect due to the decrease in sample size alone, which would see an increase of around 5 percent rather than the observed 31 percent. The greater than expected variation is likely due to the use of base weights, which additionally account for a different type of variation relating to the demographics of each individual's ZIP codes.

Because the UAS sample itself is not representative of the U.S. population, post-stratification is an important step in inference for the population. The fact that not all strata of interest are represented in the sample makes raking the natural method for assigning weights. However, doing so introduces a few complications related to the statistical framework and analysis of the data. The first relates to the increased difficulty in calculating standard errors of population estimates, which are weighted averages of the sample values. In all tables and publications, the standard errors have been calculated by taking the weights as fixed values, thereby reducing the standard errors. The sampling weights, which are a function of the strata representation in the sample, are random variables, and their variation should be factored into the calculation of standard errors (Gelman and Lu 2003). However, high response rates and targeted sampling (as described in Section 3.2) mean that the variability in the observed sample composition is small, which in turn implies that the variability in the raked weights is small. Therefore, conditional on a chosen weighting scheme, the variance of our estimators can be attributed largely to the variation in the observed responses themselves and not in the sample composition.

The second area of concern regards the effects of the sampling scheme on the weights and on the estimates they produce. In order for the raking algorithm to be appropriate in the sense that the expected weights for each stratum equal those of the population, the sampling procedure must be such that, in expectation, each stratum is proportionally represented in the sample. To be precise, the expected proportion of the sample belonging to a specific stratum is directly proportional to the relative proportion of that stratum within the population. A sampling procedure that does not have this property is likely to consistently produce weights for certain strata that do not reflect the true representation in the entire population. If strata properties correlate with payment behavior, this could lead to biased population-wide estimates. In the case of a sampling procedure in which some strata tend to be overrepresented and others underrepresented, the raking algorithm, which strives to match marginal proportions rather than those of the cross-sections of all the variables, may generate sample weights that fail to align the sample composition with the reference population. Although the sample from the UAS does not perfectly reflect the U.S. population (for example, it tends to have more females than males), the differences between the panel and the broader population are relatively small for the demographics used in weighting. In addition, for many SCPC variables, there is little evidence of strong correlations with these variables used in weighting, so any bias is likely to be small.

Overall, comparisons of changes in the estimates based on the SCPC data from year to year are likely to be meaningful. While the estimate levels themselves naturally vary with different weighting schemes, estimates of trends are likely to be more robust. A discussion of using the post-stratification weights to generate per-consumer as well as aggregate U.S. population estimates appears in Section 7.2.1.

6 Data Preprocessing

Prior to further statistical analysis, it is important to examine the data carefully and develop a consistent methodology for dealing with potentially invalid and influential data points. As a survey that gathers a large range of information from each respondent, much of it about a rather technical aspect of life that people may not be used to thinking about in such detail or many know little about, the SCPC, like any consumer survey, is susceptible to erroneous input or missing values. This section describes the general types of data preprocessing issues encountered in the SCPC and outlines the general philosophy used in studying the reliability of data at the respondent level.

Section 6.1 describes the methodology of imputing missing data, while Section 6.2 describes procedures used to identify and edit data entries that are likely to be erroneous (commonly referred to as "cleaning the data"). Overall, there were no changes in the statistical methodologies used to edit the data from those used in 2015. Nevertheless, the methodologies are described in detail for all edited variables featured in the 2016 and 2017 SCPC.

The CPRC uses the edited variables in its analysis, most notably to generate population estimates provided in the SCPC tables, as the edited distributions hopefully better represent the truth. However, in 2015, even after editing, the data corresponding to two variables measuring the frequency of cash withdrawals were so puzzling that the CPRC decided not to publish any estimates of these variables. We continued to collect this data in 2016 and 2017, but we also continue the approach of not publishing estimates based on these variables. A detailed discussion of this decision is provided in Section 6.2.2. Nevertheless, as for all edited variables, both edited and unedited data are released to the public. A guide on which variables were edited and how to access the pre- and post-processed versions of the variables is given in Section 6.3.

6.1 Data Imputation

As the post-stratification weights depend on certain demographic variables, the CESR imputes the necessary variables for respondents for whom the information is missing. In the case of many demographic variables, such as age group, gender, or race, missing information can be verified from other surveys taken within the context of the UAS. For household income and household size, both attributes that could easily change within a year, values are imputed by the CESR through logistic regression models for the purpose of creating post-stratification weights. The imputations are only used to generate post-stratification weights and are left as missing in the dataset.

The CPRC also relies on imputation to edit certain created categorical variables. The types of categorical variables in the SCPC are diverse, ranging from demographic variables, to binary variables (answers to Yes/No questions), to polytomous response variables (multiple choice questions with more than two possible answers). Currently, the data imputation performed on SCPC data relates to identifying missing values as negations of statements within the question or as implying an answer of 0 for numerical responses. This often relates to questions in which respondents are asked binary questions, such as "Do you have an ATM card?" or questions that ask respondents to enter numerical values for a set of related items, such as the number of credit cards owned for several credit card categories or the dollar value stored on different types of prepaid cards. In either of these cases, if at least one of the items features a non-missing response, we impute the values of all missing responses in the same sequence. Specifically, in the case of binary questions, missing variables are coded as "No," while in the case of numerical values, they are coded as 0.

At the moment, no other types of imputations are done, although multiple imputation procedures are being considered for future editions of the survey results. It is very difficult, without making strong assumptions, to identify irregular or erroneous data inputs, especially for multiple choice questions. Research conducted by the CPRC suggests that response bias in sequences of Likert scale questions introduced by a form of anchoring effects is present, but not of economic significance (Hitczenko 2013). See Daamen and de Bie (1992); Friedman, Herskovitz, and Pollack (1994) for general discussion on anchoring effects. Because the item response rates are high, the effect of missing values is not a major concern for the SCPC. Nevertheless, the CPRC is considering developing multiple imputation techniques for missing numerical data entries.

6.2 Data Editing

The greatest challenge in data preprocessing for the SCPC comes in the form of quantitative variables, especially those that represent the number of monthly payments or dollar value of cash holdings or withdrawals. Measurement errors in such a context, defined as any incongruity between the data entry and the true response, can be attributed to a variety of sources ranging from recall error to rounding errors to data entry errors or even to misinterpretation of the question. A data entry subject to measurement error can take many forms, but practically the only identifiable forms are those that lie outside the realm of possible values and those that fall in the realm of possibility, but take extreme values.

Data entries that defy logic are easily identified by range checks and logical reasoning. The first line of data inspection consists of a basic range and consistency check of the demographic variables to ensure that reported values are logical and that they correspond to established categorical codes. Any response item that fails this check is edited to a missing value. One example is the entry of a negative monthly payment count. A second example of a question in which data entries are potentially changed to missing values is one that first asks respondents whether or not they own various types of credit cards and then asks for the number owned for only the categories that were declared as owned. In such a case, it is technically possible for someone to claim that he or she is an adopter of a card, but, when prompted, say that he or she owns zero of such cards, a clear inconsistency. The CPRC treats responses to questions in any potential sequence as correct until an inconsistency occurs. Then, at all subsequent levels, all responses inconsistent with those to earlier questions are marked as missing. Thus, in the case of credit card adoption, the hypothetical respondent would be recorded as an adopter, but with the number of credit cards owned missing.

Identifying data that are possible, but very unlikely, is much more difficult, as it requires assessing the heterogeneity of behavior within the population. This is especially true for economic variables such as cash holdings and value of assets, which are characterized by highly right-skewed distributions. In other words, it is possible that data entries that by some numerical evaluations are statistical outliers are actually accurate and valid. This issue is not unique to the SCPC. Many consumer surveys, such as the Survey of Consumer Finances (SCF) and the Consumer Expenditure Survey (CES) must also tackle the cleaning of such fat-tailed variables. While the details of the preprocessing of outliers are not provided in either survey, the general strategy is similar to that adopted in the SCPC (Bricker et al. 2012; Bureau of Labor Statistics 2013). First, all relevant information in the data particular to each variable is used to identify statistical outliers and inconsistent responses. Then, values that cannot be confirmed or reconciled are imputed. It should be noted that the SCPC does not benefit from in-person interviews (as does the SCF) or multiple phases and modes of interview for each respondent (as does the CES), making it more difficult to identify inconsistent responses. It is important to distinguish conceptually between influential and potentially invalid data points. An influential point is one whose inclusion or exclusion in any inferential analysis causes a significant difference in estimates (Bollen and Jackman 1990; Cook and Weisberg 1982), and thus the influence of a point depends on the statistical procedure being performed. An invalid data entry is, technically, any entry that does not represent the truth. As mentioned above, data cleaning procedures focus predominantly on identifying invalid entries in the tails of the distribution (Chambers and Ren 2004). An invalid data point need not be influential and an influential point is not necessarily invalid. To the degree possible, the procedures adopted by the CPRC rely on economic intuition to identify potentially invalid data entries. Thus, the cleaning procedures for variables for which we have a higher degree of economic understanding seek to identify invalid entries and edit their value. For variables for which there is less economic intuition available, we rely more on raw statistical procedures such as matching known parametric distributions to the data or Cook's distance to identify influential points in the context of estimating weighted sample means (Cook 1977; Cook and Weisberg 1982).

Below, we outline the considerations and economic motivations in cleaning several different variables and provide adopted algorithms for each. The variables relate to the typical number of monthly uses of payment instruments, reported dollar amounts in various contexts, and the number of payment instruments or accounts owned. In the case of cash withdrawals, we argue that new data patterns observed for the reported frequency of cash withdrawal would require much more aggressive editing methodologies in order to be in line with even vague economic priors. As a result, we conclude that this set of questions is yielding implausible responses from a sufficiently large percentage of respondents to justify discarding estimates based on these variables.

6.2.1 Preprocessing: Typical Monthly Payment Use

The number of typical payments in a month is an aggregate from data entries for 41 different combinations of payment method and transaction type. The SCPC delineates 10 payment methods, nine payment instruments plus income deduction, and seven transaction types. For example, the use of cash is reported in a series of questions about cash use in the context of paying for a service, for a bill, for a product, or as a payment to a specific person. All combinations of payment method and transaction type are listed in the SCPC User's Guide (Foster 2014). In addition, for each of the 41 variables, the SCPC allows the respondent to answer on either a weekly, monthly, or annual frequency, so that recall periods better match frequencies of use that are natural to the respondents. Since only "adopters," defined as those people who claim to possess the payment method, are asked to provide information on use, missing entries for this question are assumed to be zero (for example, a person who has a credit card need not make use of it). Before preprocessing, all 41 payment number variables are standardized to a monthly frequency (multiplied by $\frac{52}{12}$ if reported by week and divided by 12 if reported by year).

The 10 payment methods are indexed by $j = 1, 2, \ldots, 10$. For each payment method, there is a variety of potential transaction types, $k = 1, \ldots, K_j$. In addition, each data entry is associated with an individual, labeled $i = 1, \ldots, N$, and a year, labeled $t = 2014, \ldots, 2017$. Therefore, Y_{ijkt} is the recorded number of typical monthly payments by individual i via payment method j of the k^{th} transaction type for that particular method in year t. Then, $Y_{ijt} = \sum_{k=1}^{K_j} Y_{ijkt}$ is the number of reported monthly payments by payment method j in year

t and $Y_{it} = \sum_{j=1}^{10} Y_{ijt}$ is the number of total number of monthly payments reported in year t.

More economic intuition exists about the total number of monthly payments than about which instruments and in what contexts those payments are made. In addition, economic theories dictate that the number of payments made with a particular payment method depends on the payment methods adopted by the individual. The collection of adopted payment methods is called a "bundle." The general cleaning procedure first identifies a hard threshold for the total number of monthly payments and then, in turn, a bundle-dependent threshold for each payment method. For each payment method, if the reported value exceeds this threshold, the lower-level components are imputed. If an individual component stands out as an outlier, it is winsorized. Otherwise, all components are scaled down to bring the resulting number of payments with the method in question to the threshold, while preserving the relative shares within the payment method. The economic idea behind this latter adjustment is that the individual is likely consistently overestimating use of the payment method.

Although the fundamental idea behind the adopted procedure is based on the common approach of using known distributions to identify potential invalid data points, the unique characteristics of payment choice require some additional assumptions. As a result, many aspects of the procedure are based on original ideas developed at the CPRC. This process is described in more detail below and is fully delineated in Algorithm 1.

An initial threshold for the total number of monthly payments was assumed to be 300, representing 10 payments per day for 30 days. The bottom panel in Figure 5 shows that this roughly corresponds to the 98^{th} percentile of the raw SCPC data for each year, and is also where the yearly distributions seem to start diverging from one another somewhat more. From a statistical point of view, the ability to pool data to estimate empirical distributions is a great advantage, as pooling enables one to base estimates on more information. In the future, other sources, such as the Diary of Consumer Payment Choice (DCPC), could also be used to inform this threshold.

Given a number of monthly payments, the distribution of the number of payments reported for each payment method quite naturally depends on which payment methods are adopted by the individual. A simple model assumes that the number of payments made with each instrument follows a multinomial distribution, conditional on the total number of payment instruments adopted. Thus, the model assumes that with each incoming payment there is some set of probabilities $\{p_j\}$ that correspond to the probability of using payment j. The decision is assumed to be independent for each individual and for each of the necessary

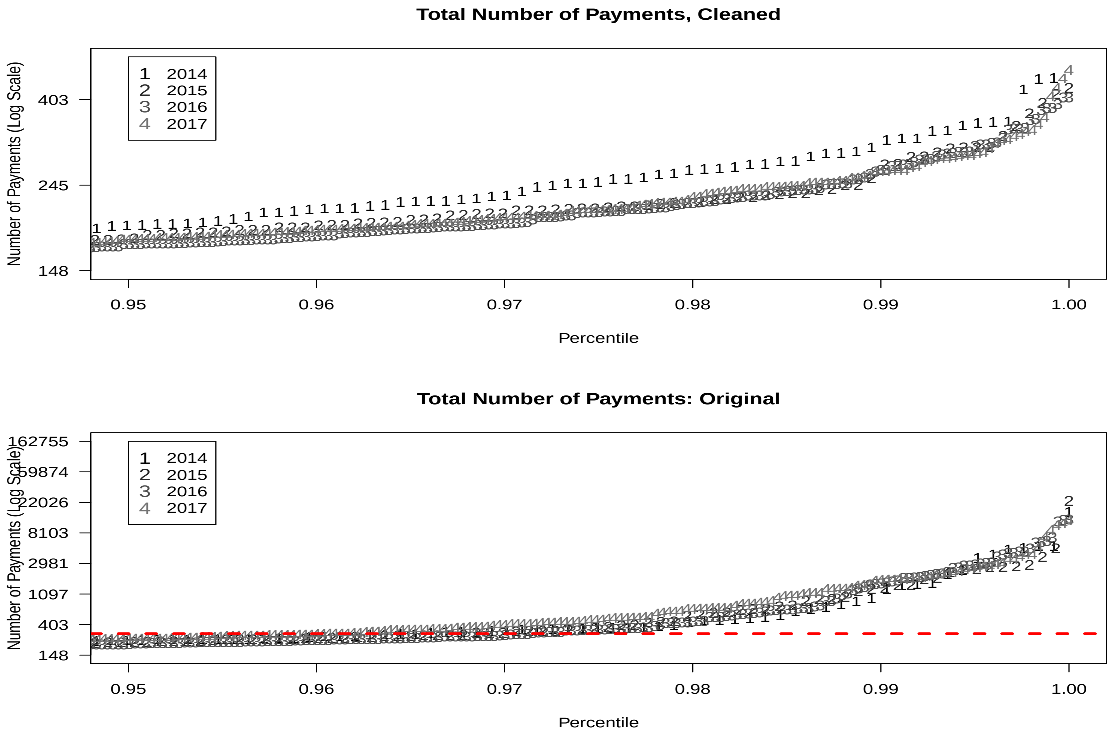

Figure 5: The log values of the largest 5 percent of the total monthly payments data before and after processing for the past four years.

Source: Authors' calculations

payments and to depend only on the individual's adoption choices. While this assumption may not hold completely (for example, the choice of payment method might depend on the dollar value of the transaction), it is a suitable approximation for the purposes of identifying likely invalid data points. To make this more concrete, for individual i in year t, let $\mathcal{B}_{it}$ be the bundle adopted by individual i; for example, $\mathcal{B}_{it} = \{1, 2\}$ for an individual who adopts only cash and checks.

In order to account for the fact that certain payment methods are used much more often than others yet to keep the calculations simple, the probabilities, $\{p_j\}$, are assumed to be proportional to the relative prevalence of the adopted payment methods to one another. Thus, for $j = 1, \ldots, 10$, r_j is defined as the weighted mean of the bottom 95 percent of the number of monthly payments made by method j in the raw data. The 95^{th} percentile is used to prevent undue influence of outliers, and changing this percentile does very little to change the relative prevalence. The intuition then is that r_j represents a prior sense of the typical monthly rate of use of payment method j among the population.

Based on the chosen r_j, the approximated proportion of payments made by individual i with payment method j in year t, defined as p_{ijt} will be

$$p_{ijt} \quad = \quad \frac{r_j}{\sum_{j' \in \mathcal{B}_{it}} r_{j'}} 1_{\{j \in \mathcal{B}_{it}\}}.$$

The value p_{ijt} is a probability and the distribution of these values will be the same for every individual with the same bundle of payment methods. It should be noted that calculations of p_{ijt} are dependent not only on the prior assumptions but also on the assumption that using one payment method does not influence the relative use rates of the other methods. As an example, this means that the relative use ratio of cash to check does not depend on whether or not the individual uses credit cards. While this might be a strong assumption, it is one that avoids the need to make many assumptions about joint use rates for various bundles of payment methods.

The cutoffs for each payment method are then defined as the 98^{th} percentile of the number of monthly payments, with 300 total payments and probability of use p_{ijt}. Therefore, if $Y_{ijt} \sim \text{Binomial}(300, p_{ijt})$, the cutoff c_{ijt} is defined to be such that

$$\text{Prob}(Y_{ijt} \leq c_{ijt}) \quad = \quad 0.98.$$

Based on this, y_{ijt} is flagged whenever $y_{ijt} > c_{ijt}$. This flag indicates that the reported value is unusually high when taking into account the payment methods adopted. It is only at this point that the lowest level of data entry, y_{ijkt}, is studied. Because little intuition exists about the distributions of the y_{ijkt}, comparisons of flagged values are made to the 98^{th} percentile of the empirical distribution estimated by pooling data from the past three years. Specifically, let q_{jk} be the 98^{th} percentile of the pooled set of data comprised of the y_{ijkt} for $t = 2008, \ldots, 2014$ among people for all (i, t) for which $j \in \mathcal{B}_{it}$. Then, for each flagged payment method, the flagged entry is imputed with the minimum of the calculated quantile and the entered value: $y^*_{ijkt} = \min(y_{ijkt}, q_{jk})$. This form of winsorizing means that extremely high reported numbers are brought down to still high, but reasonable levels. If none of the data entries at the lowest level are changed, all y_{ijkt} for the payment method j are scaled down proportionally in order to bring the total for the payment method down to the cutoff value c_{ijt}.

Once data at the lowest level of input are cleaned, aggregated values can naturally be reconstructed. Figure 6 shows the implied number of total monthly payments before and after preprocessing (on the log scale), and Figure 5 also shows the top 5 percent of edited payment totals. The fatter right tail observed in 2014 may reflect an underlying truth, but is also consistent with natural sampling variation. A feature of this algorithm is that, although it uses 300 as a threshold to flag the total number of reported payments, it does allow individuals to have more payments if reported numbers at the lowest level are consistent with others' responses. In each year, there are individuals with as many as 400 monthly payments. Figure 6 also indicates that the smallest number of payments to be edited is around 55, although the changes to the number of payments made are relatively small. Changes on this scale are due to a majority of the reported number of payments being reported for a payment instrument that has very low typical use, such as money orders.

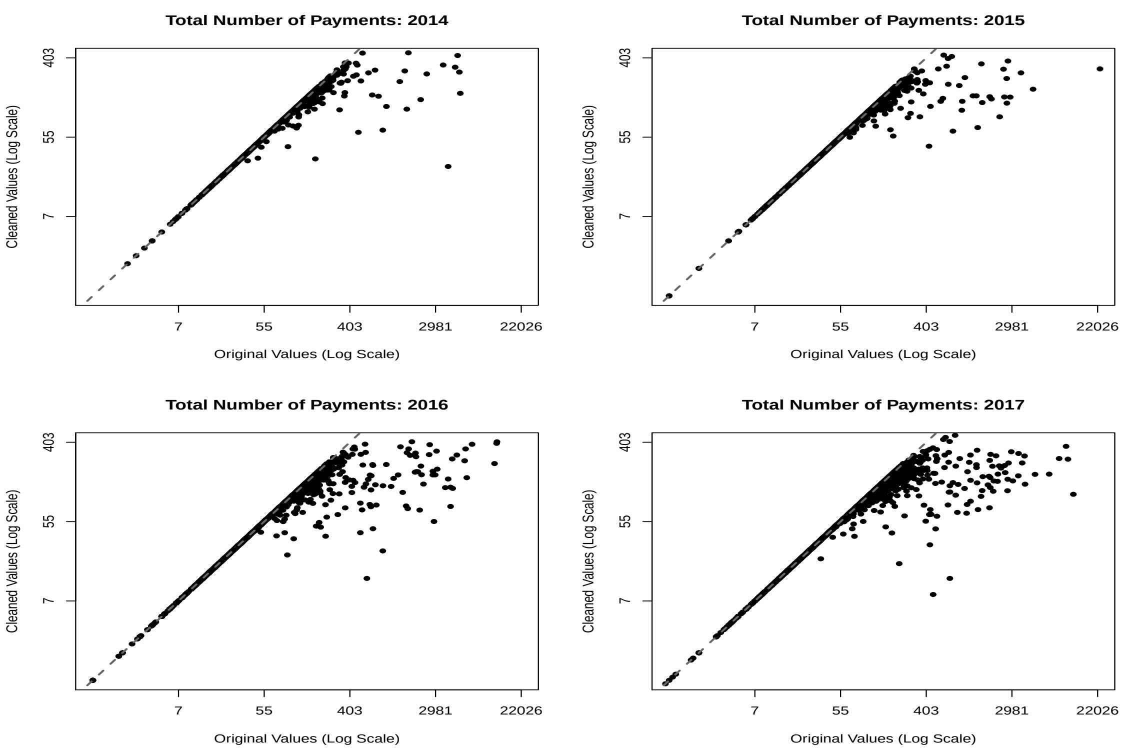

Figure 6: The log values of the cleaned total monthly payments data plotted against the log values of the original values.

Source: Authors' calculations

6.2.2 Preprocessing: Cash Withdrawal

A second concept that requires a fair amount of attention in terms of preprocessing is that of cash withdrawal. We begin by describing the editing algorithm. We then argue that the distribution of data related to the frequency of cash withdrawal is so different compared with other years that we conclude the survey question is not doing an adequate job of capturing the truth for a substantial subset of the population. For this reason, we suppress population

Algorithm 1 Preprocessing: Number of Monthly Payments

```
for i = 1 : N do
  Determine B_it
  for j ∈ B_it do
    Calculate p_ijt and then c_ijt
    if y_ijt > c_ijt then
      Set change.subtotal = 0 {used to keep track if y_ijkt are changed}
      for k = 1 : K_j do
        if y_ijkt > q_jk then
          Set y_ijkt = q_jk
          Set change.subtotal = 1
        end if
      end for
      if change.subtotal = 0 then
        for k = 1 : K_j do
          Set y_ijkt = y_ijkt × c_ijt / y_ijt
        end for
      end if
    end if
  end for
end for
```

estimates for all economic concepts based on cash withdrawal frequency from 2015 through 2017.

Preprocessing Algorithm Cash withdrawal since the 2009 SCPC is reported as a combination of four separate variables: frequency of withdrawal at primary and all other locations and typical dollar amount per withdrawal at primary and all other locations. Because reported dollar amounts correspond to typical values, which could represent the mean, the median, or the mode, the value determined by multiplying the reported frequency and the dollar amount does not necessarily correspond to the average total cash withdrawal either for the primary or for all other locations. In preprocessing, data for the primary and for all other locations are treated separately. The editing process is described below.

Assuming that N independent individuals report positive cash withdrawal in a typical month, let $C_{it} = A_{it}F_{it}$ be the typical monthly value of all cash withdrawals, where A_{it} is the reported amount per visit in year t and F_{it} is the reported frequency of monthly visits in year t. In the case of cash withdrawals, information about the tails comes from distributional assumptions, so empirical estimates that rely on pooling data across years for more statistical power are not necessary. As a result, the subscript corresponding to year t is dropped for simplicity.

If $C_i \sim \text{Log-Normal}(\mu_W, \sigma_W)$ with independence across individuals, then it follows that

$$\log(C_i) = \log(A_i) + \log(F_i)$$

has a normal distribution, which in turn means that $\log(A_i)$ and $\log(F_i)$ are also normally distributed. The fact that individuals who withdraw a larger value of cash will likely need to do so fewer times than those who take out smaller values suggests a negative correlation between the two variables. Thus, the joint distribution will take the form

$$\begin{bmatrix} \log(A_i) \\ \log(F_i) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_A \\ \mu_F \end{bmatrix}, \begin{bmatrix} \sigma_A^2 & \rho_{AF} \\ \rho_{AF} & \sigma_F^2 \end{bmatrix} \right),$$

with ρ_{AF} likely to be negative. For simplicity of notation, let $W_i = [\log(A_i) \;\; \log(F_i)]^T$, where the superscript T refers to a matrix transpose, and let μ and Σ represent the respective mean and covariance of W_i.

In order to determine distributional outliers, consider that if Λ is such that $\Lambda^T \Lambda \Sigma = \mathbf{I}_2$, the 2×2 identity matrix, (in other words, Λ is the cholesky decomposition of Σ^{-1}), then the set of $Z_i = \Lambda^T (W_i - \mu)$ will be independent draws from a two-dimensional standard normal distribution. For the bivariate standard normal, $D_i = \|Z_i\|$ is the Euclidean distance of the i^{th} draw, Z_i, to the point $(0, 0)$. Also, if $f(\cdot \mid \mathbf{0}, \mathbf{I})$ is the density function of the bivariate standard normal distribution, then $D_i^2 > D_{i'}^2$ implies $f(Z_i \mid \mathbf{0}, \mathbf{I}) < f(Z_{i'} \mid \mathbf{0}, \mathbf{I})$. This implies that if $D_i^2 = D_{i'}^2$, then the density at Z_i is equal to that at $Z_{i'}$, which is why the bivariate standard normal curve has circular contour lines. The contour lines of a bivariate normal distribution with mean μ and variance Σ will be an ellipse centered at μ with points W_i and $W_{i'}$ having the same densities if and only if

$$(W_i - \mu)^T \Sigma^{-1} (W_i - \mu) = (W_{i'} - \mu)^T \Sigma^{-1} (W_{i'} - \mu).$$

Transforming the N independent draws from the true distribution to N independent draws of the bivariate distribution makes it easier to work with the data. This transformation preserves the sense of distance from the mean with respect to the assumed density (which is lower for less likely points and decreases as one moves away from the mean). Therefore, if W_i and $W_{i'}$ are such that $D_i^2 > D_{i'}^2$, then $f(W_i \mid \mu, \Sigma) < f(W_{i'} \mid \mu, \Sigma)$. So, the extremity of each of the N points can be measured by comparing the distances D_i^2.

It is known that D_i^2 are independent and identically distributed random variables from the Exp(0.5) or equivalently a Chi-Square(2) distribution. Therefore, we can easily determine the 98^{th} percentile for D_i^2, which we call $q_{.98}$.

Algorithm 2 Preprocessing: Monthly Cash Withdrawal

Let $w_i = (\log(a_i), \log(f_i))$ for all $i = 1, \ldots, N$
Estimate $\hat{\mu} = \text{mean}(w_i)$ and $\hat{\Sigma} = \text{var}(w_i)$ from sample statistics of the w_i
Calculate $\hat{\Lambda}$ such that $\hat{\Lambda}^T \hat{\Lambda} = \hat{\Sigma}^{-1}$
Calculate $q_{.98}$ based on $\hat{\mu}$ and $\hat{\Sigma}$
for $i = 1, \ldots, N$ **do**
  Calculate $z_i = \hat{\Lambda}^T (w_i - \hat{\mu})$
  Calculate $d_i^2 = \|z_i\|^2$
  if $d_i^2 \leq q_{.98}$ **then**
    Calculate z_k^{new}
    Calculate $w_k^{new} = \hat{\mu} + \hat{\Lambda}^{-T} z_k^{new}$
    Replace w_k with w_k^{new}
  end if
end for
Keep changes to w_i only if $\log(a_i) < \hat{\mu}_A$ and $\log(f_i) < \hat{\mu}_F$.

For all observation pairs for which $D_i^2 > q_{.98}$, the procedure reassigns the data entry to a point more consistent with the fitted distribution but a minimum distance from the original value. Specifically, the data point is reassigned so that its new distance is exactly $\sqrt{q_{.98}}$. The imputation procedure is exactly the same as in previous years. First, Z_i is reassigned to Z_i^{new}, which corresponds to a well-known constrained optimization problem. Namely, Z_i^{new} is such that $\|Z_i^{new} - Z_i\|$ (the distance between the old and new points) is minimized, subject to the condition $\|Z_i^{new}\|^2 = q_{.98}$. Optimization programs for this paradigm are available for most computational packages (Press et al. 2007). The new value, Z_i^{new}, is then converted from the standard normal distribution to a corresponding value on the bivariate normal distribution defined by μ and Σ by letting

$$W_i^{new} = \mu + \Lambda^{-T} Z_i^{new}.$$

In practice, μ and Σ are not known and must be estimated from the data. We use lower-case notation, such as $w_i = (\log(a_i), \log(f_i))$, to represent the actual values observed in any given survey year, and estimate the bivariate mean and covariance with $\hat{\mu}$, the sample mean, and $\hat{\Sigma}$, the sample covariance. The entire procedure is outlined in Algorithm 2.

This procedure results in the editing of observations that are extreme with respect to the general mass of the sample data, even if the total monthly dollar value is reasonable. For example, if a person reports an amount of \$1 per withdrawal and a frequency of 0.25 withdrawals per month, the corresponding pair on the log-scale will be $(0, -1.38)$, which could be determined to be extreme given the much higher average values of frequency and amount. Thus, additional rules to exclude points from the editing procedure above may be desired. One option is not to edit any pairs for which the implied monthly dollar total is below some threshold. A second option is to consider outliers by the quadrant they lie in. For the SCPC data, a rule is imposed so that no changes are made to data for which $\log(a_i) < \hat{\mu}_A$ and $\log(f_i) < \hat{\mu}_F$.

Discussion of Cash Withdrawal Frequencies Even after the editing has been done, the cleaned variable for cash withdrawal frequency has shown vastly different distributions especially in the tail over the past few years. In particular, it is characterized by a much fatter right tail in 2014 and 2015, meaning there are many more instances of a very high number of monthly cash withdrawals. Figure 7 shows the 98th percentile of the bivariate normal distribution estimated to fit the log of the dollar amount per withdrawal and the log of the number of monthly withdrawals for the past four years of data. This contour is significant because it is used as the cutoff for trimming in the algorithm described above. Clearly, while the distributions of the log dollar amounts are fairly similar across years, the distribution of frequencies is significantly different in 2015 from those in other years. The jump from 2014 to 2015 is of a much greater magnitude. The results of a further investigation are shown in the technical appendix for the 2015 SCPC, but no explanation was found for the change.

Rather than impose a very aggressive editing algorithm and thus affect a large portion of the data, the CPRC has concluded that this set of questions simply did not collect reliable data. Although the tails of the distribution of withdrawal frequency are much more reasonable in 2016 and 2017, as can be seen in Figure 7, the history of this variable suggests it has flaws as a way of collecting data. Until more research is done and a better understanding of the nature of the responses is reached, these variables are simply not used to generalize to the U.S. population. Therefore, the official table of results does not provide estimates of the frequency of cash withdrawals or of the total dollar value withdrawn per month. Nevertheless, the raw and cleaned data are released in the official dataset. Information about cash withdrawals can be found in the DCPC, which asks respondents to report all cash activity.

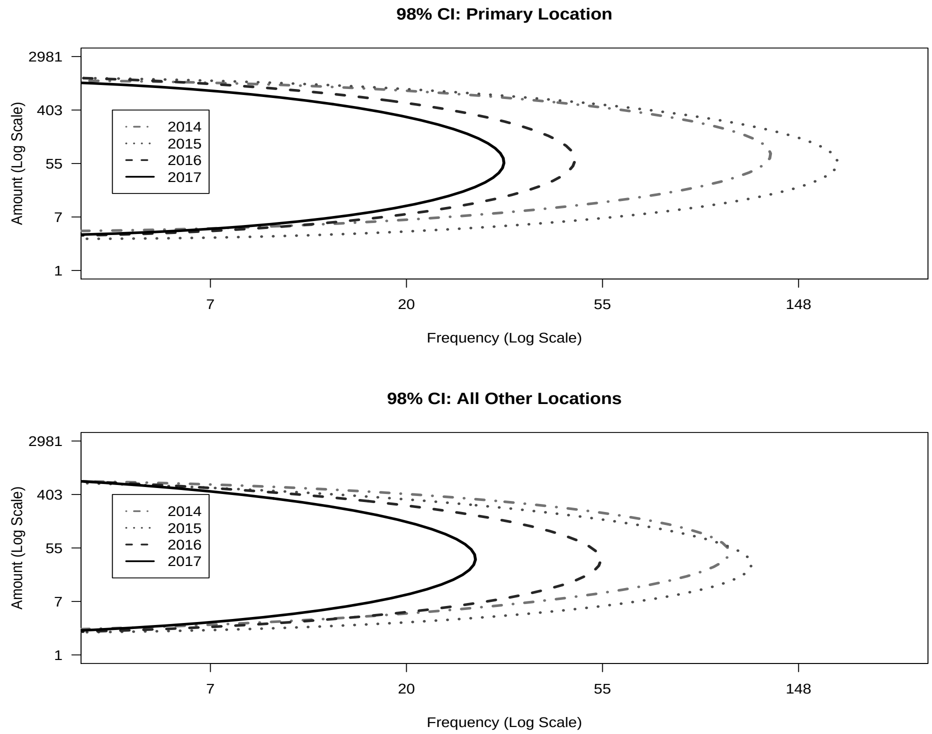

Figure 7: 98th percentiles of estimated bivariate Normal distribution for cash withdrawal data (on the log scale) from 2010 to 2015.

Source: Authors' calculations

6.2.3 Preprocessing: Cash Holdings

The SCPC also collects the dollar value of cash holdings. This concept is collected as two variables: the value of cash holdings on person and the the value of cash holdings stored at home (or other locations). We treat each variable separately, as there is no obvious relationship that one would expect to exist between the two. For the dollar values, we adopt the one-dimensional version of Algorithm 2 used to clean the cash withdrawal variables. Because other than in dimension the algorithms are identical, we do not provide more information about the procedure or delve into any details.

Figure 8 shows the distribution of the right tails of cash holdings for each of the two variables. As indicated, this cleaning procedure results in no edits to the cash holdings on person. The maximum reported values for the four years range from $1,600 to $6,700. According to our cleaning algorithm, the presence of other observations of this magnitude suggests that there is not enough evidence to edit these values. These values are large, and it is certainly possible, maybe even likely, that an input error caused $67.00 to be coded as $6,700. At the same time, the reported values are plausible; cash transactions approaching a value of

$10,000 have been observed in the DCPC.

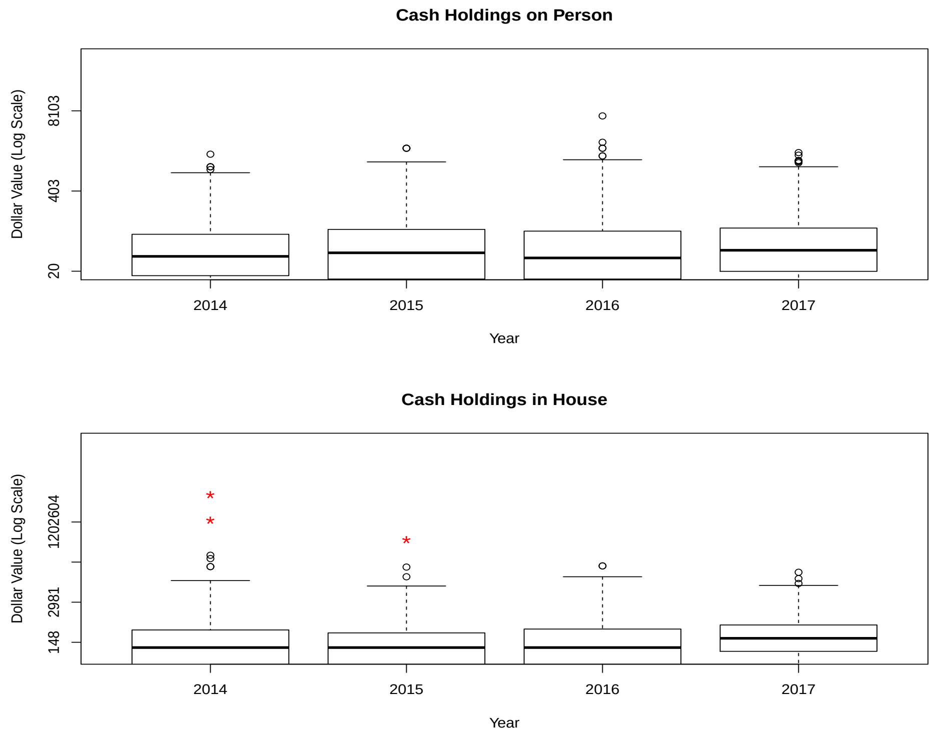

Figure 8: Boxplots of right tails of cash holdings. The asterisk represents the only edited value. Source: Authors' calculations

For cash holdings at home, no edits were made in either year. The highest values reported were $45,000 in 2016 and $28,000 in 2017. There were two values over $20,000 in the former and one in the latter. The nature of the right tail of this distribution is less intuitive, perhaps because many individuals do not disclose the amount of cash stored at home to others. As a result, it is more difficult to say whether the observed values are reasonably likely or not. As we do not have much economic intuition, we adopt a more conservative approach and let the data and the results of the algorithm stand without imposing stricter standards.

6.3 Summary of Edited Variables

In this section, we summarize the variables that are edited by the CPRC. In most cases, the edited variables are created by the CPRC as a function of various survey variables, which are any variables directly measured in the SCPC. In such cases, the underlying survey variables and any other underlying created variables that define the concept of interest are left unedited. The exceptions are the payment use variables, where the frequency-converted survey variables are edited. The original payment use survey variables remain unedited and are still reported in weekly, monthly, or yearly frequencies.

Any created variables that are defined by survey variables that are potentially edited have values determined by the edited version of those survey variables. For example, all variables relating to payment use, such as "csh_typ," which defines the number of cash payments, are aggregates of the lowest-level entries for payment use defined by a combination of payment method and transaction type. All statistics for payment use variables are created using the cleaned versions of data for each combination of payment method and transaction type. Thus, researchers who are interested in comparing the unedited variables must reconstruct any created variables themselves. All unedited variables are available and are denoted by "_unedited" or "_unedit" (in order to keep variable names below a certain number of characters) at the end of the variable name. For example, "csh_wallet_1st" holds all edited entries for the dollar value of cash holdings on person, while "csh_wallet_1st_unedited" defines the unedited version of the data. Table 17 lists all variables that are edited by the CPRC.

Table 17: Summary of edited variables. "Underlying variables" are any survey or created variables that are used to define some created variable.

Variables Cleaned (Description of Algorithm)	**Notes**
Payment Instrument Use (Section 6.2.1) pu002_a, pu002_b, pu002_c, pu002_d, pu002_e, pu003_a, pu003_b, pu003_c, pu003_d, pu004_a, pu004_b, pu004_bmo, pu004_c, pu004_d, pu004_e, pu005_a, pu005_amo, pu005_b, pu005_c, pu005_d, pu005_e, pu006a_a, pu006a_b, pu006a_bmo, pu006a_c, pu006a_d, pu006a_e, pu006c_a, pu006c_b, pu006_bmo, pu006c_c, pu006c_d, pu006c_e, pu021_a, pu021_b, pu021_bmo, pu021_c, pu021_d, pu021_e, pu021_f, pu008_c	Variables based on these variables use edited data.
Cash Withdrawal Value (Section 6.2.2) csh_amnt_1st, csh_freq_1st, csh_amnt_2nd, csh_freq_2nd	Underlying variables remain unedited. Population estimates based on csh_freq_1st and csh_freq_2nd are not generated.
Cash Holdings Value (Section 6.2.3) csh_wallet, csh_house	Underlying variables remain unedited.

7 Population Parameter Estimation

The main goals of data collection in the SCPC are to produce estimates of consumer payment behavior for the entire population of U.S. consumers and to monitor changes from one year to the next. This section presents the model that provides a framework for achieving both of these goals. This framework will work within the assumption of a longitudinal data structure, both looking forward to the future UAS panel and matching previous expositions relating to the ALP data. The model is presented in a general way so that it can easily be applied to a variety of measured variables, ranging from binary measurements of payment instrument adoption to count data such as the typical number of monthly payments. Let Y_{ijt} be the measurement for person i, for variable $j = 1, \ldots, J$ in year $t = 1, \ldots, T$. In the context of the number of monthly payments, for example, j could correspond to the number of payments made with payment method j.

Within all observed data, the respondent identifier i ranges from 1 to N, where N represents the total number of unique respondents in all T years. In fact, for certain variables, N can be lower due to a paucity of observations. As discussed, the rate of item non-response is low in the SCPC, so estimates are simply based on the observed data, with the weights of non-responders distributed evenly across those who did respond.

A natural representation for the population mean with respect to some stratification of the population into disjoint strata, indexed by s, is

$$\mu_{jt} = \sum_{s} f_s \mu_{jt}[s], \tag{2}$$

where f_s refers to the relative proportion of stratum s in the population (so that $\sum_s f_s = 1$), and $\mu_{jt}[s]$ is the average value observed in stratum s for variable j in year t.

We are most generally interested in estimating $\mu_j = [\mu_{j1}\ \mu_{j2}\ \ldots \mu_{jT}]^T$. To this end, we use the following specifications:

$$\begin{aligned} \mathrm{E}\left[Y_{ijt}\right] &= \mu_{jt}[s_i] \\ \mathrm{Var}\left[Y_{ijt}\right] &= \sigma^2_{jt} \\ \mathrm{Cov}\left[Y_{ijt}, Y_{ijt'}\right] &= \rho_{jtt'}, \end{aligned} \tag{3}$$

where s_i represents the stratum of individual i. Note that we do not specify a distribution for responses, so the approach can be generally applied to continuous variables, count data, and binary variables. We also make assumptions about the data dependence, most notably that the variance of data within strata is fixed across all strata. We also assume independence in responses across individuals (even if they are in the same household) and within individuals, but for different variables. Such assumptions are standard for most surveys and should not affect the expected values of estimates, just the associated standard errors.

In order to provide the formulas for estimating the population parameters as a function of the observed sample, we introduce the following variables. Let N_{jt} represent the number of responses obtained for variable j in year t, and let $N_{jtt'}$ represent the number of respondents who gave responses for variable j in both year t and year t'. Defining $N_j = \sum_{t=1}^{T} N_{jt}$, let $\mathbf{Y}_j$ be the $N_j \times 1$ vector with all of the responses relating to variable j over all T years. In addition, let $\mathbf{X}_j$ be a $N_j \times T$ matrix defined as follows. The (k,t)th element of the matrix, $\mathbf{X}_j[k,t]$, will be 1 if the k^{th} element of $\mathbf{Y}_j$ was observed in year t, and 0 otherwise. Finally, $\mathbf{W}_j$ is an $N_j \times N_j$ diagonal matrix such that the k^{th} element of the diagonal corresponds to the weight of the individual corresponding to the k^{th} element in $\mathbf{Y}_j$ in the year when that observation was made. Then, according to established theory (Lohr 1999), the estimates of the population vector μ_j will be

$$\hat{\mu}_j \quad = \quad \left(\mathbf{X}_j^T \mathbf{W}_j \mathbf{X}_j\right)^{-1} \mathbf{X}_j^T \mathbf{W}_j \mathbf{Y}_j. \tag{4}$$

Before we proceed, note that the population estimates calculated from the model, given in (4), correspond to the natural, design-based estimates given by the SURVEYMEANS procedure in SAS (SAS Institute Inc. 1999). Namely, if we define $\mathcal{S}_{jtt'}$ to be the index of all respondents who provided a valid data entry for variable j in year t and t', then

$$\hat{\mu}_{jt} \quad = \quad \frac{\sum_{i \in \mathcal{S}_{jtt}} w_{it} y_{ijt}}{\sum_{i \in \mathcal{S}_{jtt}} w_{it}}. \tag{5}$$

To see how the estimate in (5) mimics the model form in (2), we rewrite (5). Let m_{st} represent the number of respondents who belong to stratum s, and thus have identical weights. Then, let w_{st} represent the weight associated with each individual in stratum s. In other words, $w_{st} = w_{it}$ for all i such that $s_i = s$. Then, we let

$$\bar{y}_{sjt} = \frac{1}{m_{st}} \sum_{i \in \mathcal{S}_{jtt}} y_{ijt} 1\left[s_i = s\right]$$

be the observed sample mean for all respondents in stratum s. Then, grouping individuals by strata leads to

$$\hat{\mu}_{jt} \;=\; \sum_{s} \frac{w_{st} m_{st}}{\sum_s w_{st} m_{st}} \bar{y}_{sjt}. \tag{6}$$

Naturally, the sample mean, $\bar{y}_{stj}$, serves as an estimate of the true stratum mean $\mu_{jt}[s]$ for each s, and the relative weights assigned to each stratum are designed to have expectation equal to f_s, the true frequency of stratum s in the population. If one assumes independence between the weights and the sample observations, the implication is that $\mathrm{E}\left[\hat{\mu}_{jt}\right] = \mu_{jt}$.

It should also be noted that although the point estimates of the μ_j are the same as those in a weighted least squares analysis, we are conceptually fitting a regression model with weights designed to scale the sample data to generate estimates for a finite population (see Lohr 1999, section 11.2.3). Therefore, unlike in the weighted-least squares case, the covariance of the estimates, $\mathbf{\Lambda}_j = \mathrm{Cov}(\mu_j)$ will be estimated by

$$\hat{\mathbf{\Lambda}}_j \;=\; \left(\mathbf{X}_j^T \mathbf{W}_j \mathbf{X}_j\right)^{-1} \mathbf{X}_j^T \mathbf{W}_j \hat{\mathbf{\Sigma}}_j \mathbf{W}_j \mathbf{X}_j \left(\mathbf{X}_j^T \mathbf{W}_j \mathbf{X}_j\right)^{-1},$$

where $\hat{\mathbf{\Sigma}}_j$ is the Huber-White sandwich estimator of the error variances, $\mathrm{Var}(\mathbf{Y}_j)$ (Eicker 1967; Huber 1967; White 1980). In this context, this means that

$$\hat{\sigma}^2_{jt} \;=\; \frac{1}{N_{jt} - T} \sum_{k \in \mathcal{S}_{jtt}} \left(y_{kjt} - \hat{\mu}_{jt}\right)^2$$

and

$$\hat{\rho}_{jtt'} \;=\; \frac{1}{N_{jtt'} - T} \sum_{k \in \mathcal{S}_{jtt'}} \left(y_{kjt} - \hat{\mu}_{jt}\right)\left(y_{kjt'} - \hat{\mu}_{jt'}\right).$$

In addition to the important population means $\hat{\mu}_j$, the analysis above gives the estimates' covariances $\hat{\mathbf{\Lambda}}_j$. The square roots of the diagonal entries of $\hat{\mathbf{\Lambda}}_j$ correspond to the standard errors of the yearly mean estimates. The standard errors for the population estimates corresponding to the 2008 − 2017 SCPC are available at `https://www.frbatlanta.org/banking-and-payments/consumer-payments.aspx`.

7.1 Panel Effects

The model in (4) easily allows for the introduction of panel effects. To do so, we introduce a new variable, p_i, that indicates which panel individual i is from. For example, $p = 1$ might correspond to the ALP, and $p = 2$ might represent the UAS. The most general manifestation of panel effects on first moment estimates is by extending (4) to:

$$\mathrm{E}\left[Y_{ijt}\right] \;\; = \;\; \mu_{jt}[s_i, p_i],$$

so that the expected response depends on the panel itself. The way in which the panel selection affects the expectation can vary, but a relatively simple model is one in which it has an additive effect that is fixed across strata, but not time:

$$\mathrm{E}\left[Y_{ijt}\right] \;\; = \;\; \mu_{jt}[s_i] + \lambda_{jt}[p_i],$$

where $\lambda_{jt}[p]$ represents an additive bias associated with panel p. Under such a model, the weighted estimate given in (5) applied to data from panel p will be such that

$$\mathrm{E}\left[\hat{\mu}_{jt}[p]\right] \;\; = \;\; \mu_{jt} + \lambda_{jt}[p].$$

This unknown panel effect can make it difficult to compare estimates from different panels. A simple example is seen in Figure 9, which shows a time-series of estimates for payment instrument adoption rates based on data from the ALP and UAS from 2013 to 2015. Likely panel effects are most obvious in 2014, since temporal changes cannot be used to explain differences between the two panels. While adoption of certain instruments, such as credit cards (*cc_adopt*), shows a fair amount of consistency across panels, most yield non-overlapping confidence intervals for 2014 estimates. Developing a better understanding of panel effects and developing a comprehensive way to generate realistic trend estimates that incorporates all years of data is a high priority for the CPRC. Although certain economic measures, such as share of payments made by payment instrument, show much more consistency across panels, until a broad methodology for assimilating estimates from different panels is adopted, the CPRC refrains from making comparisons of estimates across different panels. However, a general approach to such an endeavor would be to assign a distribution to the panel effects, likely on a log-scale so that $\lambda_{jt}[p]$ or perhaps its log are draws from a Normal$(0, \sigma_p^2)$, and simultaneously estimate σ_p with other population means from the data. Greater discrepancies between panel results will correspond to greater uncertainty about the true mean, which manifests itself in larger standard errors.

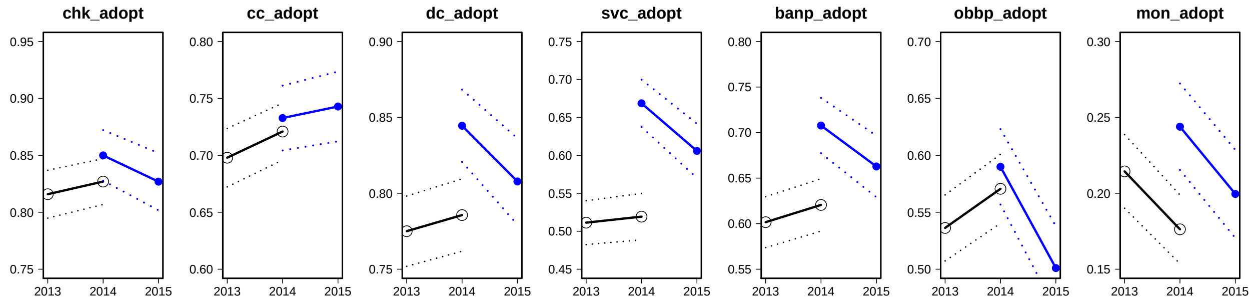

Figure 9: Estimates of payment instrument adoption rates and 95 percent confidence intervals based on ALP data in 2013 – 2014 and UAS data in 2014 – 2015.

Source: Authors' calculations.

7.2 Functions of Population Means

While the most interesting population parameters are the per capita population means, μ_{jt} in (2), we are also interested in some variables that are functions of these population parameters. Perhaps the two most illuminating functions from an economic standpoint are the growth rates and the shares. In this work, we choose to work with the macroeconomic definition of each, meaning that we consider the growth rate of averages rather than the average of the individual growth rates. We thus let

$$g_{jt} = \frac{\mu_{j,t+1} - \mu_{jt}}{\mu_{jt}} \tag{7}$$

be the growth rate of variable j from year t to $t+1$, and

$$s_{jt} = \frac{\mu_{jt}}{\sum_{k=1}^{J} \mu_{kt}} \tag{8}$$

be the share of variable j in year t.

The macroeconomic definitions used in (7) and (8) should be contrasted with their microeconomic alternatives. The former involve defining individual shares for each variable, $s_{ijt} = \frac{y_{ijt}}{\sum_{k=1}^{J} y_{ikt}}$ and estimating s_{jt} by applying (5) to this individual variable. The macroeconomic approach is statistically sounder, as, under most models that treat individuals as independent, it will give the maximum likelihood estimates of the parameters in question. For example, if the total number of payments for person i at time t is Y_{it} modeled as a Poisson random variable and the number assigned to variable j, Y_{ijt} is a binomial distribution conditional on Y_{it} with probability p_{jt}, then the maximum likelihood estimates for the p_{jt} will be given by $\frac{\sum_i Y_{ijt}}{\sum_i Y_{it}}$ rather than $\sum_i \frac{Y_{ijt}}{NY_{it}}$ (in this example, we have made all weights equal to simplify the equations). Thus, throughout this analysis, we generally use the macroeconomic definitions.

7.2.1 Generating U.S. Aggregate Estimates

The term μ_{jt} in (2) represents a per capita average in year t. For example, if the variable of interest is the number of payments made in a typical month with cash, then μ_{jt} represents the average of this value with respect to all U.S. adult consumers. In theory, if $\hat{\mu}_{jt}$ is an estimate of this mean, then a corresponding estimate for the aggregate number among the entire population would be $\hat{\mu}_{jt}$ multiplied by the size of the population. However, such calculations must be taken with caution. The estimates of μ_{jt} from the SCPC are likely to be fairly variable due to the relatively small sample size and variation in the post-stratification weights. Thus, while the estimates might be unbiased, any one estimate based on a particular sample is potentially a relatively poor estimate of μ_{jt}. Any difference between $\hat{\mu}_{jt}$ and μ_{jt} is magnified when multiplied by the U.S. population, making the resulting estimate a potentially poor estimate of the population aggregate. The high degree of error in these aggregate estimates is the reason we recommend that such methodologies be employed with caution. Issues of bias in the estimates could arise as a result of the sampling instrument and potential measurement errors. For example, the SCPC asks respondents for their personal rather than household payment choices. Inability to clearly delineate all payments related to the household, such as bills, could lead to systematically inaccurate responses.

7.2.2 Data Suppression

Many population estimates in the SCPC are based on a subset of the sample. For example, estimates for adopters of payment instruments are naturally based only on respondents who claimed to be adopters of the payment instrument in question. In some cases, the set of eligible respondents can be quite small, resulting in an unreliable estimate. As a result, in the data tables found in the SCPC report, estimates that are based on a small number of responses are suppressed.

The CPRC uses two thresholds: one for categorical data and one for numerical data. The threshold for categorical data is 20 while that for numerical data is 50. That is, if the number of respondents is lower than the corresponding threshold, the estimated population average is not reported in the tables. Numerical data are given a higher threshold because many of the variables, such as those relating to dollar amounts or number of uses, are heavy-tailed and therefore highly variable. Thus, a larger number of responses is required to produce reasonably reliable estimates. As can be seen in Klein et al. (2002), which details rules for suppression in various surveys, the thresholds adopted by the CPRC are comparable to those adopted by other U.S. government agencies.

References

Angrisani, Marco, Kevin Foster, and Marcin Hitczenko. 2013. "The 2010 Survey of Consumer Payment Choice: Technical Appendix." Federal Reserve Bank of Boston Research Data Report 13-3.

Angrisani, Marco, Kevin Foster, and Marcin Hitczenko. 2014. "The 2011 – 2012 Survey of Consumer Payment Choice: Technical Appendix." Federal Reserve Bank of Boston Research Data Report 14-2.

Angrisani, Marco, Kevin Foster, and Marcin Hitczenko. 2015. "The 2013 Survey of Consumer Payment Choice: Technical Appendix." Federal Reserve Bank of Boston Research Data Report 15-5.

Angrisani, Marco, Kevin Foster, and Marcin Hitczenko. 2016. "The 2014 Survey of Consumer Payment Choice: Technical Appendix." Federal Reserve Bank of Boston Research Data Report 16-4.

Angrisani, Marco, Kevin Foster, and Marcin Hitczenko. 2017. "The 2015 Survey of Consumer Payment Choice: Technical Appendix." Federal Reserve Bank of Boston Research Data Report 17-4.

Baltagi, Badi H. 2008. *Econometric Analysis of Panel Data*. Hoboken, New Jersey: John Wiley and Sons.

Bollen, Kenneth A., and Robert W. Jackman. 1990. "Regression Diagnostics: An Expository Treatment of Outliers and Influential Cases." In *Modern Methods of Data Analysis*, eds. John Fox and J. Scott Long, 257–291. Newbury Park, CA: Sage.

Bricker, Jesse, Arthur B. Kennickell, Kevin B. Moore, and John Sabelhaus. 2012. "Changes in U.S. Family Finances from 2007 to 2010: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 98(2).

Bureau of Labor Statistics. 2013. "Consumer Expenditures and Income." In *BLS Handbook of Methods*. BLS Publishing.

CES. Various Years. "Consumer Expenditure Survey." http://www.bls.gov/cex/home.htm.

Chambers, Raymond L., and Ruilin Ren. 2004. "Outlier Robust Imputation of Survey Data." *The Proceedings of the American Statistical Association*.

Cook, R. Dennis. 1977. "Detection of Influential Observations in Linear Regression." *Technometrics* 19(1): 15–18.

Cook, R. Dennis, and Sanford Weisberg. 1982. *Residuals and Influence in Regression.* New York, New York: Chapman and Hall.

Daamen, Dancker D. L., and Steven E. de Bie. 1992. "Serial Context Effects in Survey Interviews." In *Context Effects in Social and Psychological Research*, eds. Norbert Schwarz and Seymour Sudman, 97–113. Springer-Verlag.

DCPC. Various Years. "Diary of Consumer Payment Choice." Federal Reserve Bank of Boston.

De Leeuw, Edith D. 2005. "To Mix or Not to Mix Data Collection Modes in Surveys." *Journal of Official Statistics* 21(5): 233–255.

Deming, W. Edwards, and Frederick F. Stephan. 1940. "On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Tables are Known." *The Annuals of Mathematical Statistics* 11: 427–444.

Duncan, Greg J., and Graham Kalton. 1987. "Issues of Design and Analysis of Surveys Across Time." *International Statistical Review* 55: 97–117.

Eicker, F. 1967. "Limit Theorems for Regression with Unequal and Dependent Errors." *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 59–82.

Foster, Kevin. 2014. "SCPC User's Guide: 2011 – 2012." Technical report. Consumer Payment Research Center, Federal Reserve Bank of Boston.

Foster, Kevin. 2018. "SCPC User's Guide: 2016 – 2017." Technical report. Consumer Payment Research Center, Federal Reserve Bank of Boston.

Foster, Kevin, Scott Schuh, and Hanbing Zhang. 2013. "The 2010 Survey of Consumer Payment Choice." Federal Reserve Bank of Boston Research Data Report 13-2.

Frees, Edward W. 2004. *Longitudinal and Panel Data: Analysis and Applications in the Social Sciences.* Cambridge, UK: Cambridge University Press.

Friedman, Hershey, Paul Herskovitz, and Simcha Pollack. 1994. "Biasing Effects of Scale-Checking Style in Response to a Likert Scale." *Proceedings of the American Statistical Association Annual Conference: Survey Research Methods* 792–795.

Gelman, Andrew, and Hao Lu. 2003. "Sampling Variances for Surveys with Weighting, Post-stratification, and Raking." *Journal of Official Statistics* 19(2): 133–151.

Hitczenko, Marcin. 2013. "Modeling Anchoring Effects in Sequential Likert Scale Questions." Federal Reserve Bank of Boston Working Paper 13-15.

Hitczenko, Marcin. 2015. "Identifying and Evaluating Selection Bias in Consumer Payment Surveys." Federal Reserve Bank of Boston Research Data Report 15-7.

Huber, Peter J. 1967. "The Behavior of Maximum Likelihood Estimates Under Nonstandard Conditions." *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 221–233.

Kennedy, Courtney, Andrew Mercer, Scott Keeter, Nick Hatley, Kyley McGeeney, and Alejandra Gimenez. 2016. "Evaluating Online Nonprobability Surveys." Technical report.

Klein, Richard J., Suzanne E. Proctor, Manon A. Boudreault, and Kathleen M. Turczyn. 2002. "Healthy People 2010 Criteria for Data Supression." Technical report. Centers for Disease Control and Prevention.

Little, Roderick J. A., and Donald B. Rubin. 2002. *Statistical Analysis with Missing Data*. New York, New York: Wiley.

Lohr, Sharon L. 1999. *Sampling: Design and Analysis*. California: Brooks/Cole Publishing.

Lynn, Peter. 2009. *Methodology of Longitudinal Surveys*. Hoboken, New Jersey: John Wiley and Sons.

Press, William H., Saul A. Teukolsky, William T. Vetterlin, and Brian P. Flannery. 2007. *Numerical Recipes: The Art of Scientific Computing*. New York, New York: Cambridge Univeristy Press, 3rd ed.

SAS Institute Inc. 1999. *SAS/STAT User's Guide, Version 8*. SAS Institue Inc., Cary, NC.

SCF. Various Years. "Survey of Consumer Finances." http://www.federalreserve.gov/econresdata/scf/scfindex.htm.

Stavins, Joanna. 2016. "The Effect of Demographics on Payment Behavior: Panel Data with Sample Selection." Technical report.

Valliant, Richard, Jill A Dever, and Famke Kreuter. 2013. *Practical Tools for Designing and Weighting Survey Samples*. New York, NY: Springer.

Wang, Wei, David Rothschild, Sharad Goel, and Andrew Gelman. 2009. "Forecasting Elections with Non-Representative Polls." *Public Opinion Quarterly* 73(5): 895–916.

White, Halbert. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* 48(4): 817–838.