Research Data Reports

FEDERAL RESERVE BANK OF BOSTON™

No. 17-4

The 2015 Survey of Consumer Payment Choice: Technical Appendix

Marco Angrisani, Kevin Foster, and Marcin Hitczenko

Abstract:
This document serves as the technical appendix to the 2015 Survey 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 to 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 http://www.bostonfed.org/economic/cprc/SCPC. In this data report, we detail the technical aspects of the survey design, implementation, and analysis.

Keywords: survey design, sample selection, raking, survey cleaning, poststratification estimates

JEL codes: D12, D14, E4

Marco Angrisani is an associate economist at the University of Southern California Dornsife Center for Economic and Social Research. Kevin Foster is a survey methodologist and Marcin Hitczenko is a statistician; both are members of the Consumer Payments Research Center in the research department of the Federal Reserve Bank of Boston. Their e-mail addresses are marco.angrisani@usc.edu, kevin.foster@bos.frb.org, and marcin.hitczenko@bos.frb.org, respectively. This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at http://www.bostonfed.org/economic/rdr/index.htm.

The Survey of Consumer Payment Choice is a product of the Consumer Payments Research Center (CPRC) in the research department at the Federal Reserve Bank of Boston.

The authors thank their colleagues and management in the CPRC and the Boston Fed research department. In addition, we thank the management and staff at CESR. From the Boston Fed: Tamás Briglevics, Sean Connolly, Claire Greene, Vikram Jambulapati, Adam Karabatakis, Suzanne Lorant, William Murdock, Scott Schuh, Oz Shy, Joanna Stavins, and Bob Triest. From CESR: Tania Gursche, Arie Kapteyn, Bart Orriens, and Bas Weerman. Special thanks go to Erik Meijer from 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.

This version: July 2017

CONSUMER PAYMENTS Research Center

Contents

1 Introduction

This document serves as the technical appendix for the 2015 Survey of Consumer Payment Choice (SCPC) administered by the Understanding America Study (UAS). The SCPC is an annual survey created and sponsored by the Consumer Payment Research Center (CPRC) at the Federal Reserve Bank of Boston (Boston Fed). Each year, 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). In 2014, in addition to RAND, the CPRC contracted with the UAS for additional observations. Since the 2015 SCPC, the CPRC has worked exclusively with UAS. The UAS is a panel at the University of Southern California (USC) and is a project of the USC Dornsife Center for Economic and Social Research (CESR).

While this document builds on SCPC technical appendices corresponding to data from earlier years, it is designed to be the only necessary reference for the 2015 SCPC data. The organization of this work follows the natural, chronological progression of considerations involved in conducting and analyzing a survey. As a result, the structure of this technical appendix is identical to that of previous years, so 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 2014 version to the 2015 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 made to use best survey practices, within the constraints of the SCPC budget.

2.1 Survey Objective and Goals

As noted in Foster, Schuh, and Zhang (2013), the main objective of the SCPC program is to measure U.S. consumer payment behavior. The main goals of the program are to provide a consumer-level longitudinal dataset to support research on consumer payments and to provide aggregate data on trends in U.S. consumer payments. The change in primary survey vendor necessarily implies a discontinuation of the longitudinality built with RAND from 2008 to 2014. Nevertheless, the CPRC plans to establish a new longitudinal panel with more optimal properties within the CESR panel going forward. Differences between the RAND panel and the CESR panel are discussed further in Section 4.

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 the choices of 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 reason the SCPC focuses on the consumer is that it is less expensive to collect data about an individual rather 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 collect data for the entire household. Both strategies impose a considerable burden on the respondent. 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.

In addition, 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 nonbill 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 effective to ask the individual consumer about his or her own behavior and not about the habits of other household members.

However, 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 associated more closely with the household than an individual. 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. To measure bills accurately, it would be ideal to ask about the entire household's bill payment behavior. Nevertheless, 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. This should produce accurate measurements on bill payment behavior when averaged across the entire sample.

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 panel members are given internet access upon recruitment into the panel. The survey instrument is the NubiS survey system, developed by team members at the UAS panel.[1]

The CAWI mode is beneficial to the SCPC because of the length of the survey. The projected median length in minutes for the SCPC survey in each year is around 30 minutes. The 2015 SCPC median time spent on survey screens was 38.4 minutes and the middle 50 percent spent between 28.8 and 49.3 minutes on the survey screens. Using a CAWI allows the respondent to log off and come back to the survey later if interrupted. In addition, it 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).

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

2.4 Public Use Datasets

Users who are interested in downloading the original, unprocessed datasets can obtain these from the UAS website.[2] The Boston Fed SCPC website also contains a link to the UAS data download site. 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 finding them directly in the SCPC questionnaire, which can be downloaded as a Microsoft Word document from the Boston Fed's SCPC website.[3]

An extension of this dataset, which includes edited variables and new variables created by the CPRC, which are functions of the original survey variables, 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: 2015" (Foster 2017), 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. The name `prim_key` was chosen for internal consistency with prior years' SCPC and DCPC datasets. In the raw, uncleaned dataset from UAS, this unique identifier is named `uasid`. To merge a UAS dataset with the Boston Fed's processed dataset, the user must merge on the variables `prim_key` to `uasid`. In most 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 Boston Fed's SCPC website. For the most part, the survey questions for the 2015 SCPC are the same as or similar to those in the 2014 version, although changes are introduced every year either to collect

[2] https://uasdata.usc.edu/UAS-28
[3] http://www.bostonfed.org/economic/cprc/SCPC

new information or to improve the data collection process for the same information. This section describes the changes to the questionnaire from 2014 to 2015.

Tables 1a through 2b identify changes to the questionnaire according to the following types of changes:

1. Tables 1a–1i: Questions that are new to the SCPC in 2015.
2. Tables 2a–2b: Questions that were edited from 2014 to 2015. If a question was changed in the 2014 SCPC from a previous version, then that change remains in effect in the 2015 SCPC, unless stated otherwise.

No questions were deleted from the 2014 SCPC for the 2015 SCPC.

Table 1a: New questions in the 2015 SCPC, Assessment of characteristics section

Variable ID	Question description
as005_e	How would you rate the security of each type of debit card transaction? Voice telephone call
as005_f	How would you rate the security of each type of debit card transaction? Mobile app

Table 1b: New questions in the 2015 SCPC, Bank accounts and instruments section

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa007	At what type of financial institution is your savings account?
pa007_b	Is your savings account linked to your primary checking account?
pa006_a	At what type of financial institution is your primary checking account? (This question is equivalent to pa006 from previous years)
pa006_b	At what type of financial institution is your 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	*Please tell us a range for the balance of your primary/secondary checking account*
pa075_a, b	*Is your primary/secondary checking account jointly owned with someone else?*
pa076_a, b	*Does your primary/secondary checking account pay interest?*
pa079_a, b	*Does your primary/secondary checking account have overdraft protection?*
pa080_a, b	*With whom do you share your jointly owned primary/secondary checking account?*
pa004_a, b	*What interest rate do you earn on the balance of your primary/secondary checking account?*
pa008_b1, 2, 3	*Number of ATM cards for primary/secondary/all other checking accounts.* (This question is equivalent to pa008_b from previous years)

Table 1c: New questions in the 2015 SCPC, Bank accounts and instruments section, continued

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa008_a1, 2, 3	*Number of debit cards for primary/secondary/all other checking accounts.* (This question is equivalent to pa008_a from previous years)
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.
pa108_a, b	*Error checking questions for respondents reporting more than one card per checking account.*
pa011_a, b, c	*Does your primary/secondary/other debit card give rewards?* (This question is equivalent to pa011 from previous years)
pa026	Have you set up any of the following methods of accessing your checking accounts? Mobile banking
pa059	Which of these non-bank or other financial services did you use? (Underbanked follow-up question.)

Table 1d: New questions in the 2015 SCPC, Cash accounts section

Variable ID	Question description
pa015_c	About how much cash are you holding for cash payments (either for planned spending or emergencies)?
pa015_d	About how much cash have you set aside for long-term savings?

Table 1e: New questions in the 2015 SCPC, Virtual currency section

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa120_b1–b5	Have you heard of any of these other virtual currencies?
pa120_c	Have you heard of any other virtual currencies not listed above?
pa131_a	How familiar are you with bitcoin and how it works?
pa121_b–f	Do you have or own any of these virtual currencies? (for secondary virtual currencies that are not bitcoin)
pa122_b–f	Have you ever had or owned any of these virtual currencies? (for secondary virtual currencies that are not bitcoin)
pa125	What is the main reason that you do not own any virtual currency?
pa126_a, b	*Please tell us your primary/secondary reason for owning virtual currency.*
pa123_b–f	How much virtual currency do you have or own? (number of coins and USD amounts)
pa123_other	*Error checking screen for "Number of coins" and "USD amounts"*
pa139	Did you have to refer to records or websites to know the number of coins or the equivalent value in U.S. dollars?
pa124_b–f	In the past 30 days, have you used virtual currency to make a payment or transaction?
pa128_a–f	In the past 12 months, have you used virtual currency to make a payment or transaction?
pa127_a–f	In the past 30 days, how many payments did you make using a virtual currency?
pa140_a–f	In the past 12 months, how many payments did you make using a virtual currency?
pa129_a, b	*Have you ever used virtual currency to pay a merchant/person?*
pa129_c	Please list up to three merchants you have paid using virtual currency.
pa130_a, b, c	*How do you expect the value of one bitcoin to change over the following time periods? (Next week/month/year)*

Table 1f: New questions in the 2015 SCPC, Credit/charge accounts section

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa027	Please choose the most important reason why you don't have a credit card.
pa052_a–e	Do you own any of these kinds of cards which are branded with a logo of a company, store, or gas station?
pa054_z	*Error checking question for number of credit cards*
pa051_a–e	*You told us you have N Visa/MC/Disc/Amex charge/Amex credit cards. How many of these are branded with a logo of a company, store, or gas station?*
pu012	Last month, did you carry an unpaid balance on any credit card from one month to the next (that is, you did not pay the balance in full at the monthly due date)?
pu013	Today, about how much is the total credit limit of all your credit cards?
pu013_a	*Error checking question for credit limits and credit card balances*

Table 1g: New questions in the 2015 SCPC, Prepaid accounts section

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa045	Have you made a text message payment in the past 12 months?
pa045_a, b	*In the past 12 months, have you authorized a text message payment using one of the following methods? Via your bank/via non-bank service such as PayPal*
pa201_a–f, h, i	*Do you have any of the following types of (general purpose reloadable) prepaid cards?*
pa198_n	Please tell us how many of each type of prepaid card you have. Other types of passes or membership cards
pa202_a–n	Do any of these cards have a logo from Visa, MasterCard, Discover or American Express?
pa203_a–n	Can any of these cards be used to make purchases anywhere credit or debit cards are accepted?
pa195	Please choose the most important reason why you don't have a general purpose prepaid card.
pa194	In the past 12 months, have you used one of these electronic toll payment devices to pay a toll?
pa193	How is the electronic toll payment device that you use most often funded?
pa041_b, c	*Have you ever used a traveler's check/cashiers check, even once?*
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?
pa189_a, b, c	*In the past 12 months, have you used a mobile phone to make any of these kinds of payments? (tap and pay, scanned QR code, mobile app)*
pa188	When you pay with your mobile phone, what payment method do you use most often?

Table 1h: New questions in the 2015 SCPC, Non-bank payment accounts section

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa048_a1–e1	In the past 12 months, have you used any of the following methods to make payments with your PayPal account? (credit card, debit card, bank account, money stored with PayPal, some other method)
pa048_a2–e2	In the past 12 months, have you used any of the following methods to make payments with your Google Wallet account? (credit card, debit card, bank account, money stored with Google Wallet, some other method)
pa048_a3–e3	In the past 12 months, have you used any of the following methods to make payments with your Amazon Payments account? (credit card, debit card, bank account, money stored with Amazon Payments, some other method)
pa044_a–c	*In the past 12 months, have you used PayPal/Google Wallet/Amazon Payments to make a purchase or pay another person?*
pa047_a–c	*About how much money do you have in your PayPal/Google Wallet/Amazon Payments account?*
pa001_e	*Do you have any of the following mobile apps or online accounts? (list of 10 apps)*
pa001_f	*Which one of these mobile apps or online accounts do you have? (check-all-that-apply list of those 10 apps)*

Table 1i: New questions in the 2015 SCPC, Automatic bill payments, payment use, payment history, and demographics sections

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)
pa109	Please choose the most important reason why you don't have any automatic bill payments set up.
pu002_f	*In a typical period (week, month, or year), how many automatic bill payments do you make? (using prepaid cards)*
pu003_e	*In a typical period (week, month, or year), how many online bill payments do you make? (using prepaid cards)*
ph025_a–f	*Do you use any of the following online personal financial management (PFM) services or apps to budget and monitor your spending, saving, or account balances?*
de010	Which category represents the total combined income of all members of your family living here during the past 12 months?
de018	While answering this survey, did you look at your financial statements, receipts, old checks, websites, or other records at any time?

Table 2a: Questions that were edited from 2014 to 2015, all sections

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)	Description of change
definitions	Screen containing definitions for different payment instruments	Added definition for "virtual currency."
newobbp	*"Have you ever…" question for online banking bill payment*	Made small change in question wording.
pa016	When you get cash, where do you get it most often?	Added new response option "Payday lender."
pu009	*During the last 12 months did you have an unpaid (credit card) balance?*	Added new skip pattern in response to new item pu012.
pu010, pu011	*Two questions about unpaid credit card balances*	Edited skip pattern in response to new item pu012.
pu003	*Typical number of online bill payments*	Edited skip pattern for screen to account for new item pu003_e.
pu005	*Typical number of nonbill internet payments*	Changed title of screen to read "Nonbill internet payments for goods and services."

Table 2b: Questions that were edited from 2014 to 2015, all sections, continued.

Variable ID	Question description (questions in *italics* are paraphrases of the original question. See the questionnaire for exact wording.)	Description of change
ph004	*Question about being a victim of identity theft*	Added question text "In the past 12 months."
de010	*Question about household income*	Added several new categories at the upper end of the income distribution, up to "$500,000 or more."
de014	*Approximate value of primary home*	This question is now answered in dollars, not thousands of dollars.
de015	*Approximate amount owed on loans for primary home*	This question is now answered in dollars, not thousands of dollars.
de016	*Approximate value all non-home assets*	This question is now answered in dollars, not thousands of dollars.
de017	*Approximate amount owed on all non-home debts*	This question is now answered in dollars, not thousands of dollars.

4 Data Collection

This section describes various aspects of the data collection process for the 2015 SCPC. Once the survey instrument is finalized, the goal is to recruit a sample of respondents that can be used to make inferences about the U.S. population of consumers and, then, effectively administer the survey to those individuals. The methodologies and 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 of the CPRC (Angrisani, Foster, and Hitczenko 2013; 2014; 2015; 2016).

Perhaps the most important change to the SCPC in 2015 is that a new survey vendor was used to provide respondents and manage the survey. Whereas in prior years, official population estimates were based on data collected from RAND's American Life Panel (ALP), the 2015 estimates are based on data from 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. The main reasons for the switch were to take advantage of UAS improvements in panel recruitment (hence representativeness) and to avoid some limitations in the ALP. In 2014 and 2015, data were collected from multiple panels. In particular, the CPRC fielded the SCPC within the UAS as a pilot study of that panel in 2014. In 2015, the third-party vendor, GfK, provided around 500 respondents in addition to those from UAS. All such data will be released for public use, but will not be used in official CPRC estimates of population means. Virtually all discussion below will focus on the 2015 UAS data, although important contrasts between the ALP and the UAS will be highlighted.

A change in panel vendor naturally disrupts continuity and makes certain types of analyses, especially those studying temporal trends, more difficult. For one, the structure of the ALP allowed the CPRC to largely maintain a longitudinal panel from 2009 to 2014, with many of the same respondents taking the SCPC in subsequent 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. This advantage is naturally lost in changing survey vendors. However, the planned expansion of the UAS will allow for the construction of a longitudinal panel in the future. The CPRC believes that the best practices used to construct and maintain the UAS panel will provide for better population estimates in the future, not only in reflecting temporal changes, but also for absolute levels themselves.

A second complication comes in the form of "panel effects," biases related to the particularities of respondent selection and survey administration associated with each panel vendor (Kennedy et al. 2016). The structure of these panel effects can be quite complicated; they can vary across time or across subpopulations, for example. Their practical effect, however, is to introduce systematic shifts in expectation for point estimates of affected variables across panels. Therefore, it can be difficult to determine whether observed differences in two estimates based on data from different panels reflect real economic changes or are merely the result of panel effects. This concept is discussed in more detail in Section 7. Evidence of such effects in comparing the ALP and the UAS results are the motivation for not making temporal comparisons from 2014 to 2015.

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 SCPC is not administered to this latter subsample, because effectively incorporating data from such a specialized subpopulation into general population estimates is difficult. Therefore, the description below will focus on the other types of panelists.

Because the CESR was developed and is managed by many of the same individuals who did the same for the ALP, many aspects of the two panels are similar, and the UAS benefits from experience gained while working with the ALP. A key difference, however, is in the method used to recruit panelists. The ALP relied on a variety of methodologies, including some that were not probability-based, and it was initially intended to represent a population of retirees. By contrast, the recruitment of all individuals into the UAS panel occurs through address-based sampling, in which zip codes and then addresses within those zip codes are drawn at random. A detailed description of this 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 of the Fall of 2015, about 18 percent of enrolled households featured multiple members in the panel.

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 to answer the survey at a given point in time may include inactive members. Nevertheless, the annual attrition rate is roughly 10 percent, so the proportion of such cases is likely to be relatively small.

Since its inception, the UAS panel has been growing in size through subsequent waves of recruitment. At the time of sample recruitment for the 2015 SCPC, there were 2,140 eligible respondents from the Nationally Representative subsample. Most of these individuals were added during the initial wave of recruitment in February 2014, when 9,284 households were selected and invited to join the panel. An additional wave of 1,799 households were invited in September 2015. Around 4,300 households with Native American origin were selected for invitation via two waves conducted in June 2014 and January 2015. Of these, only around 3 percent, or 124 individuals, would be UAS panelists in the Fall of 2015, with the lower percentage due to the fact that some invitees were eventually excluded because they did not meet ethnic requirements for the subsample.

4.2 SCPC Sample Selection

A key aspect of the SCPC data collection process in 2015 is the inextricable link between the SCPC and the Diary of Consumer Payment Choice (DCPC). The DCPC, a diary developed by the CPRC and co-sponsored by the Federal Reserve Banks of San Francisco and Richmond, asks individuals to track details of all payments and financial transactions over the course of three consecutive days. The DCPC had not been fielded since 2012, and, unlike that version, the 2015 version introduced a strict requirement that respondents take the SCPC prior to their assigned diary period. Therefore, initial recruitment of respondents was done jointly for the SCPC and the DCPC.

The allocated budget of the sponsoring banks allowed for the fielding of the SCPC and DCPC to a greater number of individuals than the number of UAS panelists at the time of sample selection. Therefore, the opportunity to participate was extended to all UAS panelists except those who were part of the Los Angeles County subset.[4] The effective census of the UAS panel mirrors the approach taken with UAS for the 2014 pilot, when the entire panel of around 1,500 individuals were invited and yielded 1,238 respondents. This simple approach contrasts with the methodology used in previous years for sample selection within the ALP (except 2008), where stratified quota sampling was used to ensure an expected level of representativeness of the sample with respect to certain demographics. However, because the ALP was not very representative of the U.S. population in its earlier years, sample selection often involved balancing longitudinality with representation, as the latter could be improved by adding recently added members instead of those with more experience.

[4] An additional 504 individuals, selected to meet demographic criteria matching the U.S. population of adults, were provided by GfK for the SCPC and DCPC

Because the DCPC is difficult to manage, in order to ensure adequate participation throughout the period of interest, a short "consent" survey is used to determine dates when respondents may not be able to participate. The fielding of such a consent survey is atypical, as most surveys are simply made accessible to each potential respondent online, with notification coming in the form of an email and link. In 2015, the consent survey was released in September to all UAS panelists. Individuals were asked whether they were willing to participate in the DCPC for a three-day diary period sometime between October 13 and December 17. The SCPC was later released in the normal fashion to all individuals who agreed to participate in the DCPC, offering a $20 incentive.

The consent survey was answered by 1,482 individuals, representing only 65.5 percent of panelists. However, of those who responded, 1,291 agreed to participate in the surveys, yielding an overall consent rate of 57 percent, but an 87 percent consent rate conditional on viewing the SCPC/DCPC invitation. This overall consent rate is significantly lower than the 75 to 85 percent observed in previous years, including the 2014 UAS. A major reason for this may be the additional recruitment step presented by the consent survey, which was offered without an incentive, thus not making it an enticing survey for panelists.

In order to increase the number of responses, on November 23, 2015, an additional 701 UAS members, chosen from the subset who had not responded to the consent survey, were invited to take the SCPC only (upon completion, they were offered the opportunity to take the DCPC). Of these, 165 agreed to take the SCPC, improving the overall consent rate to 64 percent, still much lower than in previous years. Table 3 summarizes the recruitment process and acceptance rates.

Table 3: Number of survey invitations, consents, and participations for the 2015 SCPC.

	Subsample	
	Nationally Representative	Native Americans
Invitation		
Initial Invite (SCPC and DCPC)	2,140	124
Secondary Invite (SCPC only)	701	0
Consent		
Consented to Initial Invite	1,211	80
Consented to Secondary Invite	165	0
Total	1,376	80
Started the SCPC	1,349	80
Completed the SCPC	1,332	75

Source: Authors' calculations.

4.3 Suvey Participation

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. Optimal survey fielding may involve accounting not only for the time of year, but also the time within a month, and even 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 CPRC has broadly 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 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.

In 2015, however, an unexpected delay in fielding the DCPC due to negotiations with RAND meant that responses could only be collected for the second half of October. As a result, the CPRC opted to field the DCPC for approximately two months, from the middle of October to the middle of December. This choice not only tested the feasibility of managing such a complicated survey for a longer period of time, but also allowed for the collection of more observations. Most notably, around 500 respondents were selected to participate in two diary periods, separated by two to four weeks to minimize potential diary fatigue. The additional respondents from GfK could only be recruited in time for diary periods in the later half of November. An added benefit of the longer diary period is that data from the 2015 DCPC can be used to ascertain whether the diary is able to capture any changes in payment behavior due to to the holiday period that coincides with the second half of the diary period in 2015. Despite this wide range of diary dates, the 2015 SCPC was released on October 6th, the same day as in 2014 and no more than two weeks later than the 2010 – 2013 SCPCs.

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. Participants do not have to finish the survey in one sitting and have the option to log off and continue later. Respondents are given reminders every so often at the discretion of the CESR if they have not logged on to the SCPC. In particular, in 2015, reminders are sent out a few days before the assigned DCPC dates to each respondent, in order to ensure that individuals have completed the SCPC by the time their DCPC commences.

Of the 1,456 individuals who consented to participate, only 27, or 1.8 percent, never started the survey, which is defined as logging on. Starting the survey is an important threshold as everyone who does so is considered a participating respondent and is assigned a survey weight. 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 individuals who answered all of the SCPC questions and reached the last screen, which asks individuals for feedback on the survey questionnaire itself, but did not log out. Indeed, reaching the last question is the minimum requirement for the respondent to receive the financial incentive. Only 22 of those who started did not complete the survey, representing 1.5 percent of participants. This is comparable to noncompletion rates in previous years: 1.8 percent of those who started the SCPC did not complete it among the 2014 ALP sample. Table 3 provides a further breakdown of participation among the UAS subsample. Figure 1 shows the sample sizes for the official SCPC and DCPC, those respondents who were used to estimate population parameters in CPRC releases, for all prior years of surveys.

Of those who participated in the SCPC, 1,155, or 85.6 percent, also participated in the DCPC, meaning they logged on for at least one diary day. In fact, 516 SCPC participants took the diary twice. More information about how DCPC diary periods relate to the SCPC will be provided in a comparable document for the 2015 DCPC. SCPC responses from 2015 can also be compared to those from the previous year, as 917 individuals also participated in the 2014 SCPC (512 did not).

Figure 2 shows the proportion of surveys completed as a function of the number of days since the survey was distributed for the 2010 – 2015 versions.[5] The initial pace of completion was comparable for the 2015 UAS sample and those from the ALP in previous years. In fact, the 2015 version took less time to reach a completion rate of 33 percent than the 2014 version. However, beyond the first four days, the marginal rate of completion for 2015 was much slower than in previous years. In all five surveys from the ALP panel, roughly 80 percent of respondents had completed the survey within two weeks of release and around 95 percent had done so within a month. The comparable numbers for the 2015 SCPC are 55 and 68 percent, respectively. Almost all of this is explained by the late recruitment of 165 SCPC respondents near the end of November.

Figure 3 shows the proportion of surveys completed by each calendar day within each of the years from 2010 to 2015. This depiction reveals the relatively wide range of timeframes in which the bulk

[5] Similar plots that include data from 2008 and 2009 can be found in earlier versions of the SCPC Technical Appendix.

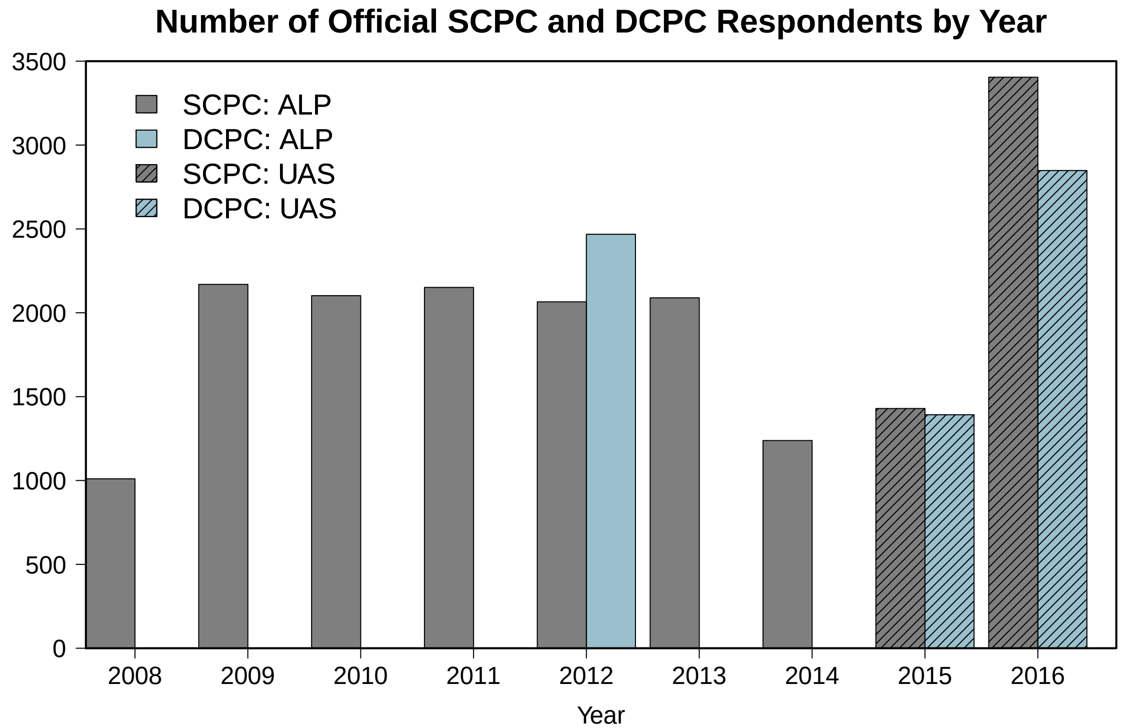

Figure 1: The sample sizes and sources for the SCPC and DCPC from 2008 to 2015. Sample sizes correspond to the subset of observations that the CPRC uses to make population inferences.

Source: Authors' calculations.

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 beginning of October. 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 the 2015 results to those of previous years.

More general seasonal effects may affect comparisons across years as well. For example, 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. 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 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. Figure 4 indicates that from 2010 to 2012 the survey was getting longer, with the

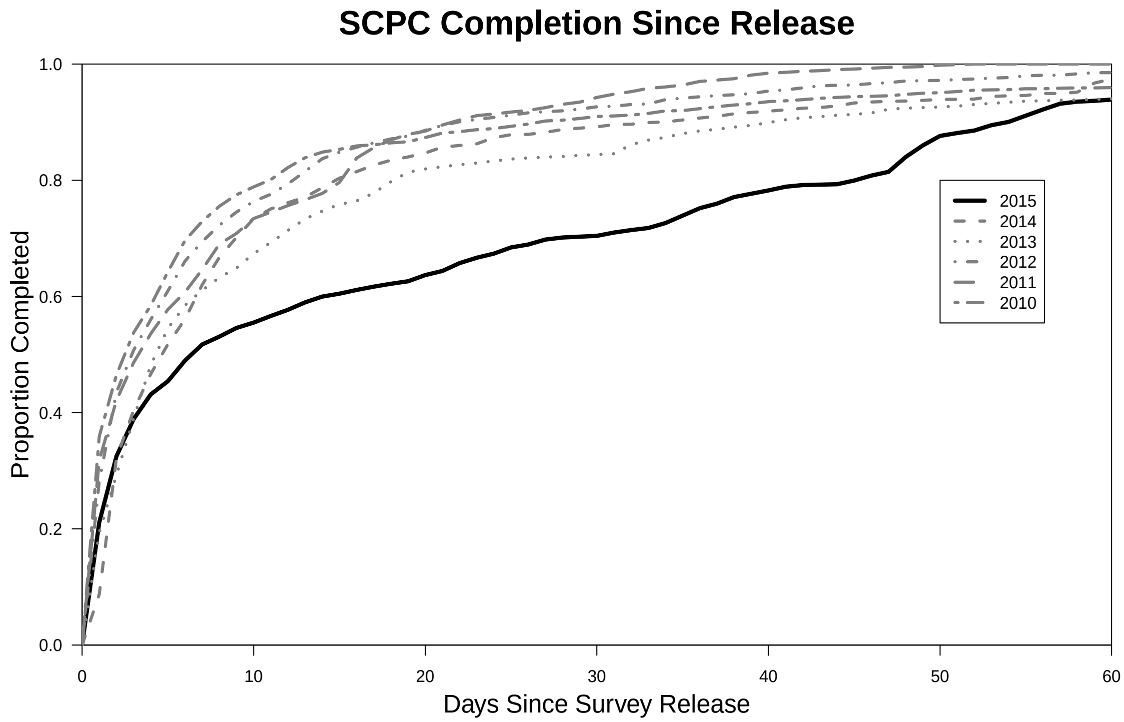

Figure 2: The proportion of respondents who completed the survey as a function of the number of days since the survey was received.

Source: Authors' calculations.

median completion time going from 33 minutes in 2010 to almost 38 minutes in 2012. However, the 2013 survey has a median completion time of 32 minutes, although 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, which was not paired with a follow-up survey, has a median completion time of 29.5 minutes, making it a considerably shorter than all previous versions.

There are two distributions for the 2015 data: one based on all data and one based on surveys taken after October 20, 2015. The former shows a significant jump in the median completion time, with a median time of 49 minutes. Part of this results from the fact that the survey added questions from 2014 to 2015, as documented in Section 3. In addition, perhaps less familiarity with the survey among UAS panelists led to longer survey times (a majority of RAND participants in 2014 had taken the survey for at least three years prior to 2014). Finally, it was discovered that the UAS servers were inconsistent in how long survey screens were taking to update, with surveys taken much later in the survey period displaying faster times. Table 4 shows the median completion times according to the date range in which the SCPC was taken and reveals a drop in median time

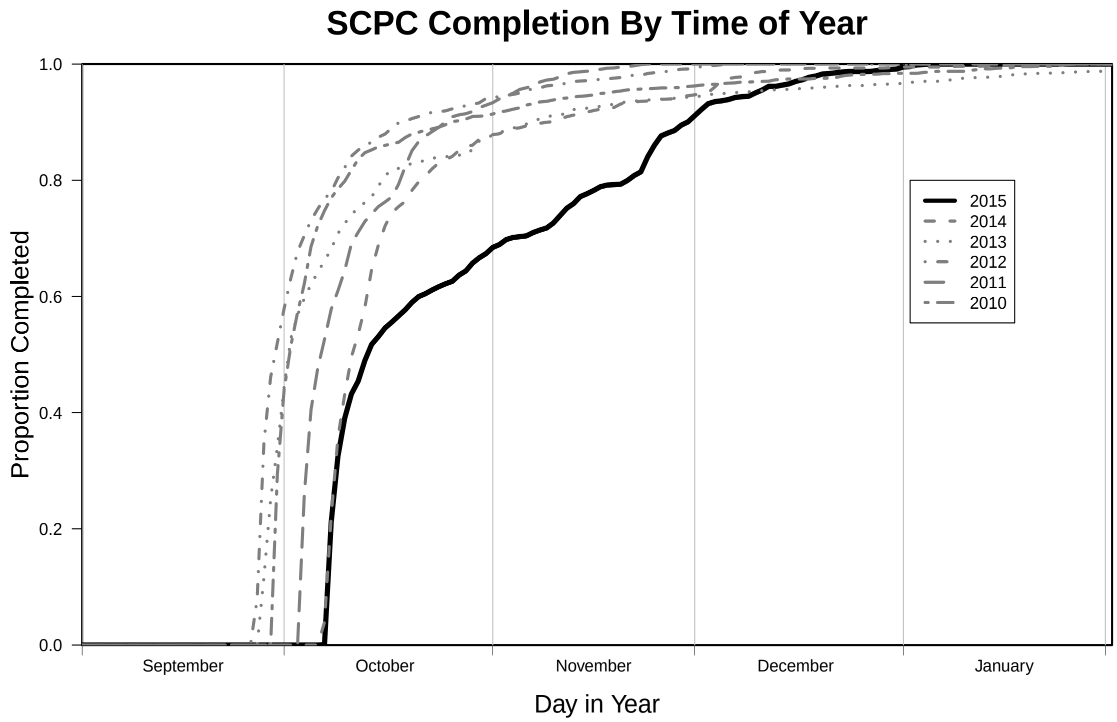

Figure 3: The proportion of respondents who completed the survey as a function of the date within the year.

Source: Authors' calculations.

of up to 15 minutes between the first week since of release and the later weeks. As shown in Figure 4, the elapsed time to take the survey based on those taken with faster updates looks much more similar to that in previous years and has a median of 39 minutes.

Table 4: Median time of completion by date on which survey was started.

Date Range	Number of Observations	Median Time (minutes)
Oct 6 - Oct 13	772	53.7
Oct 14 - Oct 20	101	49.9
Oct 21 - Nov 03	135	39.9
Nov 04 - Nov 17	133	40.7
Nov 18 - Dec 01	183	37.9
Dec 02 - Jan 10	105	39.1

Source: Authors' calculations.

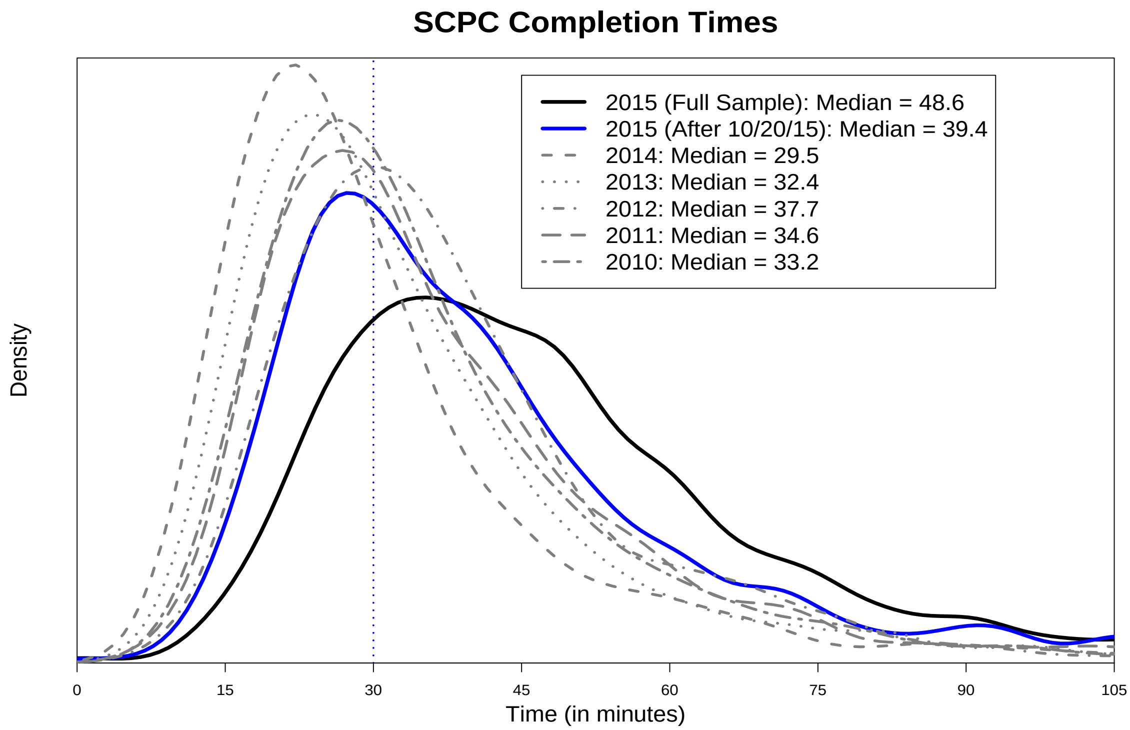

Figure 4: The proportion of respondents who completed the survey as a function of time. The vertical line at 30 minutes represents the intended average length of completion.

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 2015 SCPC has over 200 survey variables, although the survey itself is administered with a 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 ranging from 1.32 to 2.03 percent, as shown in Table 5. Interestingly, the range of nonresponse rates is not as variable as in 2014 and does not follow as clear a pattern of increasing as one moves further along in the survey. Nevertheless, even for the later questions, 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. Overall, 96 percent of respondents answered all eight of the selected questions. Item response rates are very similar across the 2014 and 2015 SCPCs.

Table 5: Nonresponse rates (%) for eight questions in the 2014 and 2015 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
2014 SCPC	0.44	1.00	1.88	1.33	1.49	2.10	2.10	2.49
2015 SCPC	1.61	1.40	1.40	1.32	1.54	1.75	2.03	1.89

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 is not 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 composition for variables with a lot of heterogeneity across demographic groups, as is the case with many payments variables (Stavins 2016), can lead to bias if not accounted for.

Table 6 shows the unweighted and weighted marginal proportions of various demographic groups in the 2015 SCPC along with unweighted ones for the 2014 RAND sample. Overall, the 2015 unweighted sample seems to be slightly farther from target levels than the 2014 unweighted sample. In particular, the 2015 sample has greater under-representation of males, those over the age of 65, and Hispanics. Another large discrepancy is found among the race category "Other," where the 2015 has a much higher percentage of the sample in this group. Of course, this is explained by the inclusion of the Native American subset in 2015, which makes up 5.6 percent of the 2015 sample. It is not surprising that the 2014 sample matches the general U.S. population better, as sample selection relied on stratified quota sampling, targeting a proportional number of responses within each strata. Nevertheless, the 2015 UAS sample does an adequate job of matching the U.S. population of adults, due largely to the address-based recruitment adopted by the CESR.

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.

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 7 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

Table 6: 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 2014 SCPC**	**Unweighted 2015 SCPC**	**Weighted 2015 SCPC**
Gender	Male	47.0	45.0	48.2
	Female	53.0	55.0	51.8
Age	18−24	2.8	4.0	6.7
	25−34	20.9	17.4	23.3
	35−44	16.7	21.3	16.4
	45−54	19.7	20.6	17.7
	55−64	20.9	22.2	16.7
	65 and older	19.0	14.5	19.2
Race	White	79.9	75.8	76.3
	Black	10.3	9.8	13.2
	Asian	2.4	2.8	4.1
	Other	7.5	11.6	6.4
Ethnicity	Hispanic	12.5	8.4	13.2
Education	No HS diploma	2.7	4.6	9.0
	High School	14.5	15.5	32.8
	Some College	39.0	38.8	28.3
	College	25.2	24.6	17.0
	Post-graduate	18.7	16.5	13.0
Income	$< \$25K$	19.4	22.1	21.8
	$\$25K - \$49K$	26.0	21.7	24.1
	$\$50K - \$74K$	21.9	19.7	19.2
	$\$75K - \$99K$	11.3	13.3	11.7
	$\$100K - \$124K$	9.2	9.4	8.6
	$\$125K - \$199K$	8.8	10.0	11.0
	$\geq \$200K$	3.4	3.8	3.7

Source: Authors' calculations.

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 were estimated by averaging over the samples. The combination of variables in Table 7 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 7 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 7: 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 $\times$ 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 $\times$ Ethnicity	
M, White	M, Other
F, White	F, Other

Gender $\times$ 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 $\times$ 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, 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 2015 SCPC, income is the most commonly missing variable. Nevertheless, only 22 individuals had some demographic variable imputed for weight calculation.

We describe the raking algorithm in greater detail below. Let $d = 1, \ldots, 4$ represent the four bivariate demographic categories presented in Table 7. 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.

The weighting algorithm begins by assigning base weights that are intended to account for the presence of the Native American subset, effectively downweighting the contribution of those individuals. The initial weights are such that the relative weight of the Native American subset matches the 1.7 percent found in the general population, which, in the case of the 2015 SCPC, sample works out to:

$$w_i^{(0,4)} = \begin{cases} 0.26, & \text{if Native American} \\ 1.04, & \text{if General Population} \end{cases}$$

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 2015 SCPC data converges. 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 487 unique weights, meaning 487 out of a possible 660 strata implied by the demographics in Table 7 are represented in the 2015 sample. It is worth noting that due to the trimming of weights, the Native American subsampample carries more weight than in the regular population: the total weight given to Native Americans in the sample is 3.1 percent versus 1.7 percent in the general population. The increase in variability of the weights from 2014 to 2015 is smaller than would be expected from the decrease in sample size. A reduction of sample size from around 1,800 to around 1,250 would generally lead to an increase in standard errors of around 1.2. Instead, the variability increases by a factor of 1.07 (from 0.83 to 0.89).

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 over-represented and others under-represented, 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"). There were no changes in the statistical methodologies used to edit the data prior to analysis in 2015. Nevertheless, the methodologies are described in detail for all edited variables featured in the 2015 SCPC.

The CPRC uses the edited variables in its analysis, most notably to generate population estimates provided in the SCPC tables. However, in 2015, even after editing, the data corresponding to two variables measuring the frequency of cash withdrawals was so puzzling that the CPRC decided not to publish any estimates of 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: $2011 - 2012$ (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 $t = 2008, \ldots, 2014$. 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 top 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. 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.

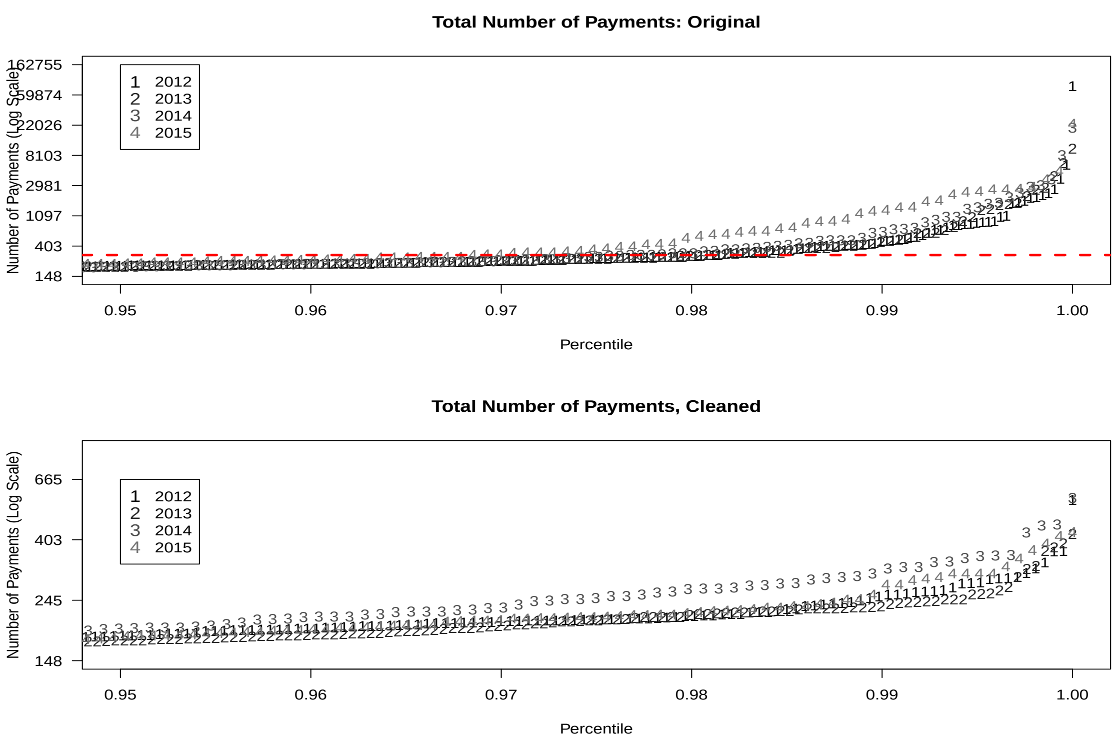

Figure 5: The log values of the largest 5 percent of the total monthly payments data before and after processing for the past four years.

Source: Authors' calculations.

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 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 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. A feature of this algorithm is that, although it uses 300 as 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 20, although the changes to the number of payments made are relatively small.

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 estimates for all economic concepts based on cash withdrawal frequency in 2015.

6.2.2.1 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

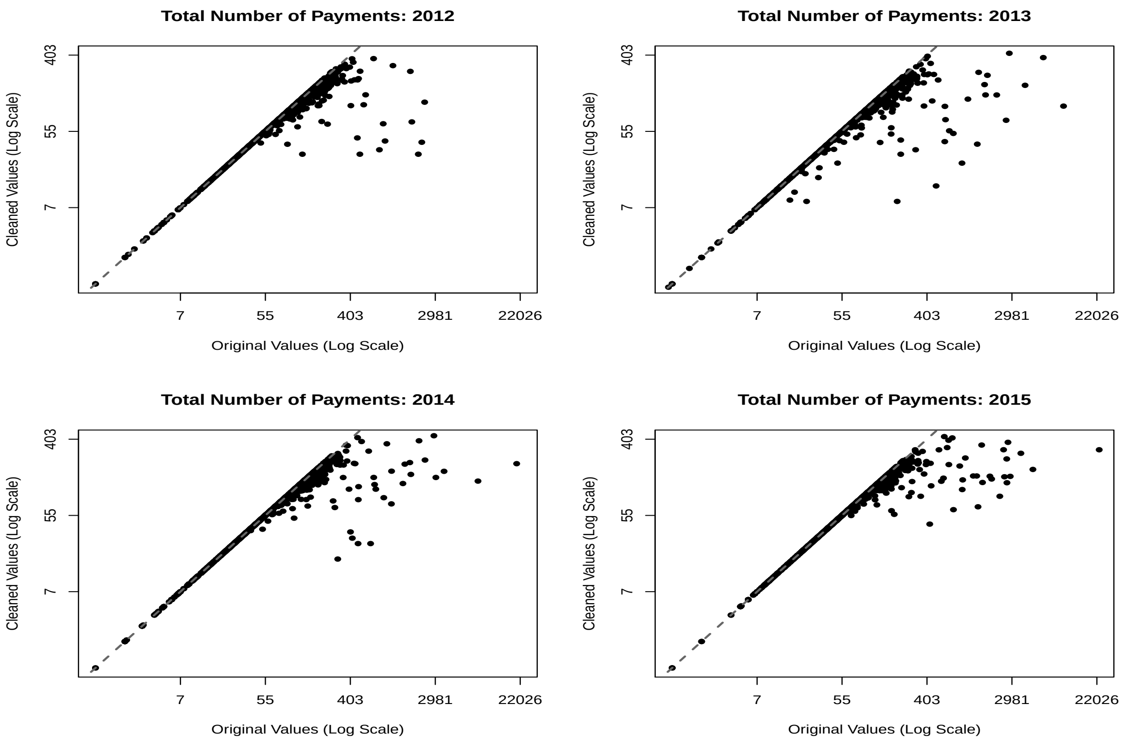

Figure 6: The log values of the cleaned total monthly payments data plotted against the log values of the original values.

Source: Authors' calculations.

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 $C_i \sim$ Log-Normal(μ_W, σ_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

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
```

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 $\mathrm{Exp}(0.5)$ 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, ..., N
Estimate μ̂ = mean(w_i) and Σ̂ = var(w_i) from sample statistics of the w_i
Calculate Λ̂ such that Λ̂^T Λ̂ = Σ̂^{-1}
Calculate q_.98 based on μ̂ and Σ̂
for i = 1, ..., N do
    Calculate z_i = Λ̂^T (w_i − μ̂)
    Calculate d_i^2 = ||z_i||^2
    if d_i^2 ≤ q_.98 then
        Calculate z_k^new
        Calculate w_k^new = μ̂ + Λ̂^{-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) < μ̂_A and log(f_i) < μ̂_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 dollar 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 $\log(f_i) < \hat{\mu}_F$.

6.2.2.2 Discussion of Cash Withdrawal Frequencies

Even after the editing has been done, the cleaned variable for cash withdrawal frequency has a vastly different distribution than in previous years. In particular, it is characterized by a much fatter right tail: 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 six 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. While there were other observed increases in the right tail, the jump from 2014 to 2015 is of a much greater magnitude.

Figure 8 shows the reported frequency values for the 80th through the 95th percentiles for data collected from 2013 to 2015, including data from both panels in 2014. There is a notable discrepancy between data coming from the two panels, especially for the primary location. Percentiles are nearly identical for the 2013 and 2014 ALP data, while the 2014 and 2015 UAS data differ not only from the ALP data, but also from each other. In the case of the primary location data, the upper two deciles of the distributions are fundamentally different between the two panels. While the divergence among secondary locations only occurs at the 94th percentile, this too represents a significant portion of the population. The fact that there is a significant difference within the UAS panel data from 2014 to 2015 for the primary location withdrawal frequencies is also unexpected.

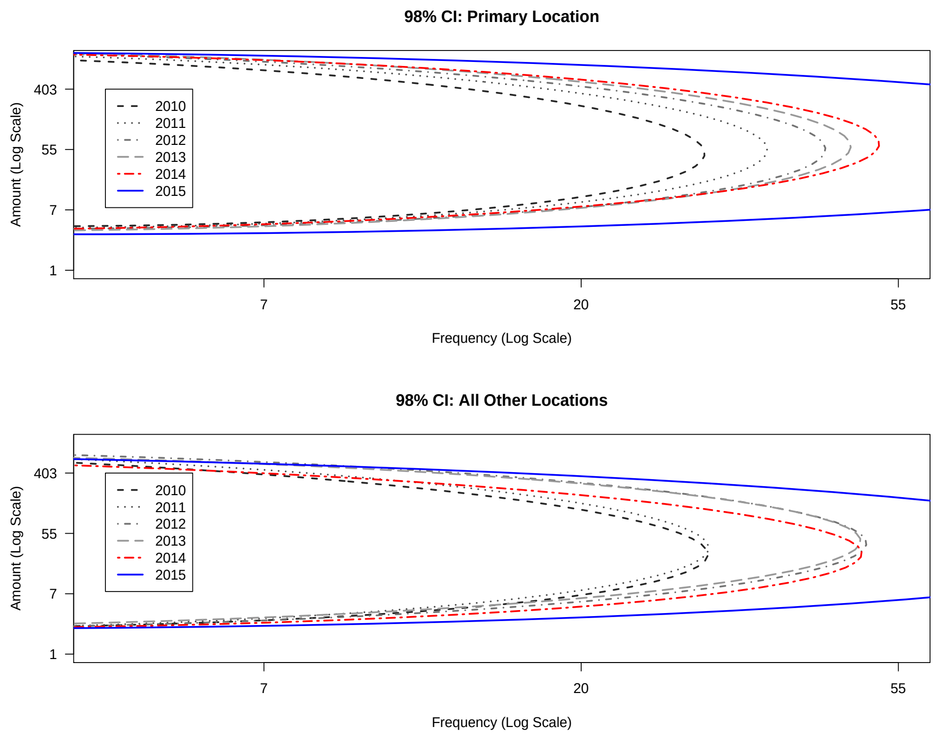

Figure 7: 98th percentiles of estimated bivariate Normal distribution for cash withdrawal data (on the log scale) from 2010 to 2015.

Source: Authors' calculations.

Our analysis focused first on identifying whether certain respondent characteristics could explain the change across panels and years in observed distributions. Demographics alone do not explain the differences. Further analysis was done to see whether the extremely high values in 2015 can be attributed to a subset of respondents who were less diligent in taking the survey for some reason. Dichotomies were developed based on the amount of time it took to complete the survey, the number of contradictory statements made in other parts of the survey, and the frequency on which respondents chose to report cash withdrawals. All models failed to account for the observed differences. The distributions of reporting frequency were very similar across the two years: in 2014, 19, 54, and 27 percent chose week, month, and year, respectively, while comparable numbers for 2015 were 19, 50, and 31 percent.

The only editing methodology that would have produced estimates of population means with

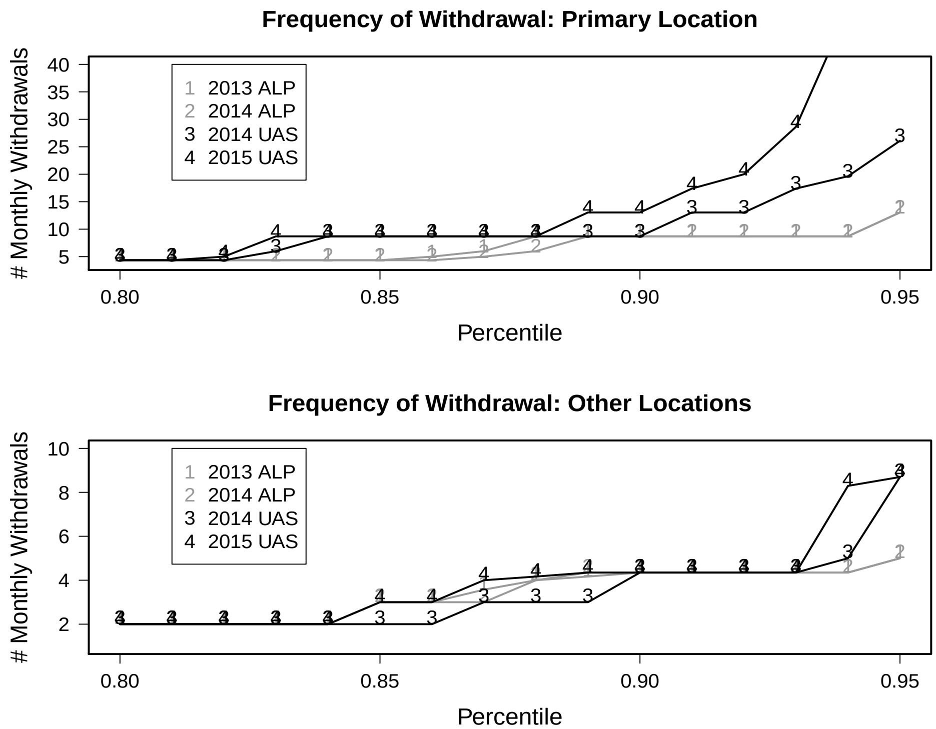

Figure 8: Percentiles of reported cash withdrawal frequencies from 2013 to 2015 in the ALP and UAS panels.

reasonable confidence would require aggressive cleaning to remove essentially the entire right tails of the data. For example, restricting mean estimates to those individuals who had a total cash withdrawal total below $5,000 (a reasonable amount, for which little economic basis exists to set as an upper limit) still had the estimate for the number of monthly cash withdrawals from a primary source in 2015 as over 50 percent higher than in 2014 (7.0 vs. 4.0).

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 did not collect reliable data in 2015. 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 the total dollar value withdrawn per month. Nevertheless, the raw and cleaned data are released in the official dataset.

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 9 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 five years range from $2,000 to $5,000. These values are large, and it is certainly plausible that an input error caused $20.00 to be coded as $2,000. At the same time, the reported values are plausible and the presence of other observations of this magnitude suggests that there is not enough evidence to edit these values.

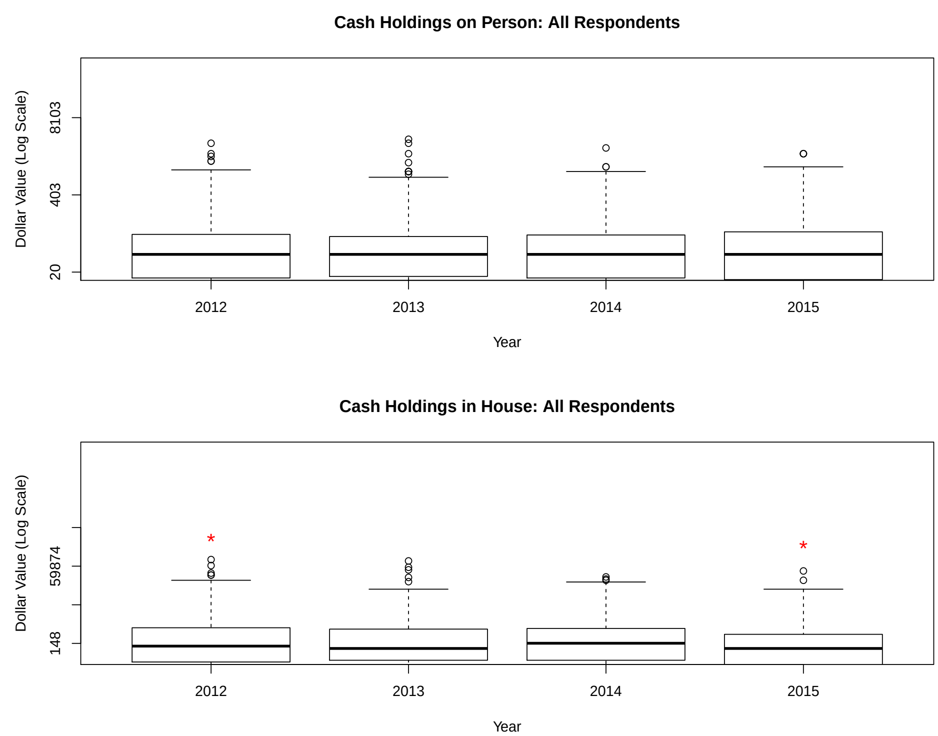

Figure 9: Boxplots of right tails of cash holdings. The asterisk represents the only edited value. Source: Authors' calculations.

As in the previous three years, there are no changes to the reported cash holdings on person, with the maximum reported value being $2,000. For cash holdings at home, one item was identified as improbable: a reported value of $350,000 was adjusted to $41,311. This adjusted value is still higher than any cash holding reported in the previous two years. The second highest reported in 2015 value was $20,000.

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 8 lists all variables that are edited by the CPRC.

Table 8: 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

An important goal of the data collection in the SCPC is to produce estimates of consumer payment behavior for the entire population of U.S. consumers, especially 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. For the 2015 UAS alone, $T = 1$ and $N = 1{,}429$. 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, 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 – 2015 SCPC are available at http://www.bostonfed.org/economic/cprc/SCPC.

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 10, 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), show 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.

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

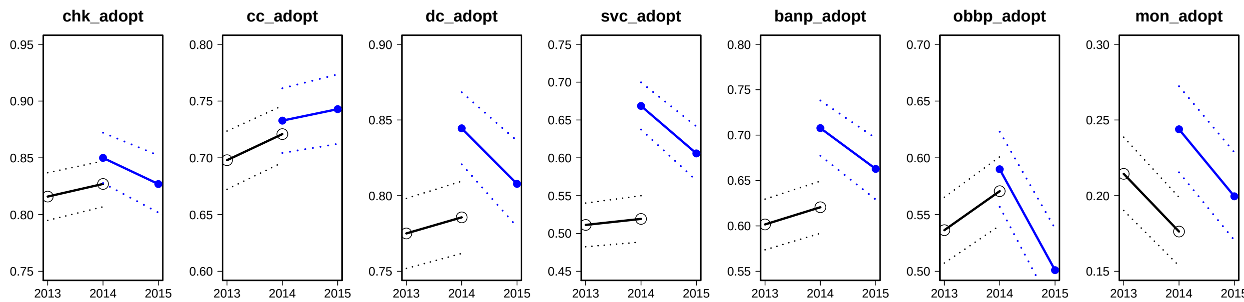

Figure 10: Estimates of payment instrument adoption rates and 95 percent confidence intervals based on ALP data in 2013 – 2014 and UAS data in 2014 – 2015.

Source: Authors' calculations.

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 micro-economic 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 2015 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.

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. 2017. "SCPC User's Guide: 2015." 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.