Research Data Reports

FEDERAL RESERVE BANK OF BOSTON™

No. 16-4

The 2014 Survey of Consumer Payment Choice: Technical Appendix

Marco Angrisani, Kevin Foster, and Marcin Hitczenko

Abstract:

This document serves as the technical appendix to the 2014 Survey of Consumer Payment Choice administered by the RAND Corporation. 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.

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/wp/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 and the RAND Corporation. 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 and the RAND Corporation: 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 2016

CONSUMER PAYMENTS Research Center

Contents

1 Introduction

This document serves as the technical appendix for the 2014 Survey of Consumer Payment Choice (SCPC) administered by the RAND Corporation (RAND). 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 RAND. In 2014, in addition to RAND, the CPRC contracted with the Dornsife Center for Social and Economic Research at the University of Southern California (USC) for additional observations. As the 2014 SCPC paper (Greene, Schuh, and Stavins 2016) describes results for data collected by RAND only, so too does this document. A separate document discussing the USC dataset and how it compares to the RAND dataset will be released in the future, although many of the descriptions provided below naturally apply to the USC version of the SCPC as well.

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 2014 SCPC data collected by RAND. 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 from the 2013 version to the 2014 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. In Section 7, we present the statistical methodology used for estimating and comparing population estimates. Finally, Section 8 builds on these results by conducting a variety of hypothesis tests, the results of which are given in Section 9.

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.

2.2 Unit of Observation

The SCPC uses the individual consumer as both the sampling unit and the observation unit. This choice stands in contrast to those of 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 that 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 non-bill payments made in a month. It may be difficult for one household member to accurately report the behavior of other household members, and, even if asked, household members may not feel comfortable sharing such information with one another at such a level of detail. Therefore, it is most appropriate to ask the individidual 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 more closely associated with the household than with an individual. Many such payments are scheduled to be automatic and often come out of joint accounts or pooled resources. 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 2015b). Treating such a sample as representative of all consumers may lead to overestimation of the number of bills paid. To accurately measure bills, it might be better 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, or in person.

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 members of RAND's American Life Panel (ALP) are given internet access upon recruitment into the panel. The survey instrument is the MMIC survey system, developed by the RAND Corporation.[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 2014 SCPC median time to completion was 29.5 minutes and the middle 50 percent of respondents completed the survey in 20.9 to 56.3 minutes. 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

[1] MMIC stands for Multimode Interviewing Capability. More information on MMIC is available at http://www.rand.org/labor/mmic.html.

2005).

2.4 Public Use Datasets

Users who are interested in downloading the original, unprocessed datasets can obtain these from the RAND Corporation's website. The Boston Fed SCPC website contains a link to the RAND data download site. Interested users must create a username and password to download data from the RAND 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, and randomized variables have not been unrandomized. 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.[2]

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 of all new variables not found in the original dataset are described in the companion document, "SCPC Data User's Guide: 2014" (Foster 2015), 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 the RAND and the Boston Fed datasets, and it can be used to merge the raw, uncleaned data from RAND with the Boston Fed's processed dataset. In addition, `prim_key` can be used to merge the SCPC dataset with any other RAND American Life Panel survey dataset.

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 2014 SCPC are the same as or similar to those in the 2013 version, although changes are introduced every year

[2] http://www.bostonfed.org/economic/cprc/SCPC

either to collect new information or to improve the data collection process for the same information. This section describes the changes to the questionnaire from 2013 to 2014.

Uniquely in 2014, data were collected through two separate survey instruments: the 2013 SCPC and an additional survey dubbed "Module B." A description of the questions in Module B can be found in the 2013 Technical Appendix (Angrisani, Foster, and Hitczenko 2015), and the questionnaire is available on the SCPC website. In 2015, Module B was not administered. However, some questions from Module B were moved to the main survey for the 2014 SCPC.

As in most years, interest in new topics led to the addition of new questions. In 2014, these questions related primarily to the awareness and use of virtual currencies as well as to the identification of underbanked consumers. Because there was a concerted effort to reduce the average time it takes to complete the SCPC, more questions were removed from the survey than were added. The removed questions were chosen based on a combination of the measured time it took respondents to answer them as well as on how often the questions were used in research. Finally, certain questions were edited in the hope of improving the quality of the reported data. These edits may involve a simple change of the question text, a change in the survey logic associated with the question, or a change in the location of the question within the questionnaire.

Tables 1 − 4 display changes to the questionnaire according to the following types of changes:

1. Table 1: Questions moved from Module B to the 2014 SCPC.
2. Table 2: Questions that appeared in the 2013 SCPC or Module B, but are not in the 2014 SCPC.
3. Table 3: Questions that are new to the SCPC.
4. Table 4: Questions that were edited from 2013 to 2014. If a question was changed in the 2013 SCPC from a previous version, then that change remained in effect in the 2014 SCPC, unless stated otherwise.

Sections 3.1 − 3.7 describe changes to the questionnaire, by the six sections designated in the written form of the questionnaire.

Table 1: Questions moved from Module B to the 2014 SCPC.

Variable ID	Question description
fr001_a	In your household, how much responsibility do you have for these tasks? Paying monthly bills.
fr001_b	In your household, how much responsibility do you have for these tasks? Regular shopping for the household.
fr001_d	In your household, how much responsibility do you have for these tasks? Making decisions about saving and investments.
fr001_e	In your household, how much responsibility do you have for these tasks? Making other financial decisions.
as005_a−as005_d	How would you rate the security of each type of debit card transaction?
ph004	Have you, or anyone you know well (family, friends, neighbors, coworkers, etc.), ever been a victim of what you consider to be identity theft?

Table 2: Questions deleted for the 2014 SCPC.

Variable ID	Question description
as004_a–as004_f	How do you rate the overall SECURITY of the following methods of making a payment?
as012	Please tell us which payment characteristic is most important when you decide which payment method to use.
as005_e	How would you rate the security of each type of debit card transaction?... Using a debit card on a voice phone
pa007	At what type of financial institution is your primary savings account?
pa026_a	Have you set up any of the following methods of accessing your current bank accounts?... Mobile banking
pa026_b–pa026_g	Using your mobile phone, have you done any of the following in the past 12 months?
pa126_b–pa126_g	Using your mobile phone, have you ever done any of the following?
pa033_a, _b, _c, _e, pa032, pa049	In the past 12 months, have you used the following methods to access your account?
pa016_a, _b	When you get cash [from an ATM/back at a retail store], what [kind of plastic card/method] do you use most often?
pa027_a–_g	Do you have any of the following payment methods with contactless payment technology?
pa051_a–_h	In the past 12 months, have you made any of the following types of mobile payments?
pa048	Please tell us how your non-bank online payment service is funded.
pa044	In the past 12 months, have you used a non-bank online payment service to make a purchase or pay another person?
pu100	Were any of the payments you reported in the previous questions made for both your household and some other organization?
pu012	What interest rate do you pay on [your credit card that has the largest revolving balance]?
ph005_a, _c–_e, _g	Have you ever entered any of the following information on an Internet web site or sent the information in an e-mail message?
ph012	During the past 12 months, did you pay for anything in cash to receive a discount?

Table 3: New questions in the 2014 SCPC.

Variable ID	Question description
pa055_a, _b	In the past 12 months, did you use any of the following financial services? [List of underbanked services]
pa120_a, _b	Have you heard of [Bitcoin/any other virtual currency]?
pa121_a, _b	Do you have or own any [Bitcoin/other virtual currencies]?
pa122_a, _b	Have you ever had or owned any [Bitcoin/other virtual currencies]?
pa123_a, _b	How much [Bitcoin/other virtual currencies] do you have or own?
pa124_a, _b	In the past 12 months, have you used [Bitcoin/other virtual currencies] to make a payment or transaction?

Table 4: Questions that were edited from 2013 to 2014.

Variable ID	Question description	Description of change
pa034	If you are given a choice while completing a debit card purchase, do you prefer to enter your PIN or give your signature?	Added new response option "Neither one/I don't like PIN or signature."
pa040_d	In the past 12 months, have you used any of the following payment methods, even once? Certified check	Added skip logic so that only bank account adopters see this question.

3.1 SCPC Section I: Preliminaries

There were no changes to any of the questions in Section I: Prelimaries.

3.2 SCPC Section II: Financial responsibility

In the 2013 survey, the financial responsibility questions appeared in the Module B survey. These questions were moved into the main SCPC survey for 2014. Table 1 describes these four questions (fr001_a, fr001_b, fr001_d, fr001_e).

3.3 SCPC Section III: Assessment of Characteristics

This section notes any changes to the questionnaire in Section III: Assessment of Characteristics. The changes were as follows:

- Dropped the table of questions asking respondents to rate the security of various means of making a payment (`as004_a`–`as004_f`).
- Moved the table of questions asking respondents to rate the security of making various types of debit card transactions (`as005_a`–`as005_d`) from the Module B survey into the main survey in 2014.
- Dropped the question asking respondents to rate the security of making a debit card transaction on a voice phone call (`as005_e`).
- Dropped the question asking what payment characteristic is most important when deciding which payment method to use (`as012`).

3.4 SCPC Section IV: Payment Adoption

This section notes any changes to the questionnaire in Section IV: Payment Adoption.

In the module on bank account adoption and use, the changes were as follows:

- Dropped the question asking the financial institution of the primary savings account (`pa007`).
- Added a new response option to the question asking about debit card authorization preferences (`pa034`).
- Dropped the question asking if the respondent has set up mobile banking (`pa026_a`).
- Dropped the questions asking if the respondent has done any mobile banking in the past 12 months (`pa026_b`–`pa026_g`).
- Dropped the questions asking if the respondent has ever done any mobile banking (`pa126_b`–`pa126_g`).
- Dropped the questions asking if the respondent has accessed their bank account in various ways (`pa033_a`, `_b`, `_c`, `_e`, `pa032`, `pa049`).
- Added two questions to mimic the FDIC measurement of the "underbanked" population (`pa055_a`, `_b`).

In the module on cash, the changes were as follows:

- Dropped the question asking what plastic card is used when getting cash from an ATM (`pa016_a`).
- Dropped the question asking what method is used when getting cash back at a retail store (`pa016_b`).

In 2014, there is a new Bitcoin/virtual currency module. All the questions in this section are new.

- Have you heard of [Bitcoin/any other virtual currency]? (`pa120_a`, `_b`)
- Do you have or own any [Bitcoin/other virtual currencies]? (`pa121_a`, `_b`)
- Have you ever had or owned any [Bitcoin/other virtual currencies]? (`pa122_a`, `_b`)
- How much [Bitcoin/other virtual currencies] do you have or own? (`pa123_a`, `_b`)
- In the past 12 months, have you used [Bitcoin/other virtual currencies] to make a payment or transaction? (`pa124_a`, `_b`)

In the module on adoption of all other methods or technologies adopted, the changes were as follows:

- Dropped the table of questions about contactless payment technologies (`pa027_a`–`pa027_g`).
- Dropped the table of questions asking whether the respondent made any mobile payments in the past 12 months (`pa051_a`–`pa051_h`).
- Dropped the question asking how the non-bank online payment service is funded (`pa048`).
- Dropped the question asking whether the respondent made a purchase using a non-bank online payment service in the past 12 months (`pa044`).
- Added skip logic to the question asking whether the respondent used a certified check in the past 12 months (`pa040_d`). We will only ask this question to bank account adopters.

3.5 SCPC V: Payment Use

This section notes any changes to the questionnaire in Section V: Payment Use. The changes were as follows:

- Dropped the question asking whether any reported payments were made for both the household and for some other organization (`pu100`).
- Dropped the question about the interest rate on the credit card with the largest revolving balance (`pu012`).

3.6 SCPC VI: Payment History

This section notes any changes to the questionnaire in Section VI: Payment History. The changes were as follows:

- Included a Module B question asking whether the respondent or anyone he or she knows well has ever been a victim of what they consider to be identity theft (`ph004`).
- Dropped the table of questions asking whether the respondent had ever entered any personal information on a website or in an email (`ph005_a`, `_c`, `_d`, `_e`, `_g`).
- Dropped the question asking whether the respondent had paid for anything in cash in the past 12 months in order to receive a discount (`ph012`).

3.7 SCPC VII: Demographics

There were no changes to any of the questions in Section VII: Demographics.

4 Data Collection

This section describes various aspects of the data collection for the SCPC, with a primary focus on the 2014 version. Once the survey instrument has been finalized, the collection of data involves two general steps: sample selection and administration of the survey. The strategies and philosophies adopted by the CPRC in each step are outlined below. In addition, summary statistics related to survey completion 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).

Unlike in previous years, the respondents to the 2014 SCPC originate from two different panels: RAND's American Life Panel (ALP) and the Center for Economic and Social Research's Understanding America Study (UAS). The ALP, which had been the sole source of respondents for all previous CPRC surveys dating back to 2008, provided around 60 percent of respondents. The UAS, by contrast, was established in 2013, with the 2014 SCPC being the first CPRC survey administered to its panelists. The use of two survey panels allows the CPRC to compare population estimates based on the ALP to those based on a panel that was not only built through a different methodology, but also has considerably less experience with surveys in general and the SCPC in particular. A report comparing the results of the two respondent sources and drawing the implications for accurate population estimates will be released in the future. However, for the purpose of studying trends in payment behavior, it is easier to do so by restricting analysis to the ALP respondents, thereby eliminating from the analysis the effect of any possible differences between the two samples. Below, we describe data collection strategies and statistics for the ALP sample only.

4.1 American Life Panel

The ALP commenced in 2003 as a panel of approximately 500 members, with the original intent to study the methodological issues of internet-based surveys among the older population. As a result, until 2006 all recruits into the ALP were over the age of 40. Since then, the ALP has expanded to include individuals between the ages of 18 and 39 and has grown considerably in size. At the time of the 2014 SCPC sample selection (end of September 2014), there were 8,230 panelists.

There are several pathways that lead individuals into the ALP, but from a survey methodological point of view these condense into two general recruiting strategies. The first strategy involves recruiting volunteers from households that are not yet represented in the ALP. Traditionally, RAND has done this by gathering volunteers from other, already-established panels, such as the University of Michigan Internet Panel Cohort (http://www.sca.isr.umich.edu/) and the National Survey Project Cohort (terminated in 2009). Potential subjects have also been recruited via address-based sampling directly by RAND. Most notably, in 2011, around 2,000 panel members from ZIP code areas with high percentages of Hispanics and low-income households (referred to as the "Vulnerable Population Cohort") were added to the ALP. The second strategy involves asking individuals already in the ALP to recommend acquaintances or fellow household members to participate in ALP-distributed surveys. As of 2014, members who were lone representatives of their households represented 77 percent of the ALP cohort.

ALP members remain in the panel unless they formally ask to be removed or stop participating in surveys over a prolonged period of time. At the beginning of each year, RAND 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. Between September 2013 and September 2014, 68 individuals left the ALP, and 1,238 new respondents were added.

In its early stages, the ALP was, understandably, not demographically representative of the U.S. population of adults. First, due to the early intentions of researchers developing the panel, the panel prior to 2006 was composed exclusively of individuals above the age of 40. In addition, as the panel was expanded, members recruited directly from already-existing panels were recruited on a voluntary basis, with recruitment rates ranging from around 30 percent to approximately 50 percent. Even if the source panels had been representative, nonuniform eagerness to join the ALP across demographic strata could have easily produced a biased cohort. Finally, expanding the panel by inviting household members likely skewed the demographic composition further. Nevertheless, as the ALP has been growing in size, its overall representativeness relative to the Current Population Study (CPS) with respect to a variety of demographic variables has been improving. More information about the American Life Panel can be found at the website http://mmic.rand.org/alp.

4.2 SCPC Sample Selection

The SCPC was originally conceived as a longitudinal panel. The benefits of a longitudinal panel, namely, the added power associated with tracking trends at the individual level, have been well discussed (Baltagi 2008; Duncan and Kalton 1987; Frees 2004; Lynn 2009). Thus, for many research agendas, it is advantageous to base results on a longitudinal panel, rather than on a sequence of cross-sectional studies. As a result, one of the primary goals of SCPC sample selection in each year of its existence has been the preservation of the longitudinal structure.

The planned sample size of the 2008 SCPC was 1,000 respondents. The limitations of the ALP size at the time of sample selection in 2008 (1,113 individuals) forced a virtual census of the ALP. In each year from 2009 to 2012, everyone who had completed the SCPC in the previous year was invited to participate again, in order to maximize the size of the longitudinal panel. High retention rates led to a four-year panel from 2009 to 2012 of 1,515 individuals. Respondents with no prior experience with the SCPC were also added each year, and the details of this process can be found in the Technical Appendices of the corresponding years (Angrisani, Foster, and Hitczenko 2013; 2014). However, it is important to note that in 2012 an effort to pair the SCPC with the first full version of the Diary of Consumer Payment Choice (DCPC) led to the addition of 1,111 new respondents to the 2,065 respondents with previous experience. By design, many of the new respondents represented demographic strata that were poorly, if at all, represented in the pool of respondents with previous SCPC experience.

Analysis of the 2012 SCPC data made it clear that the added demographic coverage of the new respondents had a beneficial effect on estimating population parameters. Specifically, population estimates of payment instruments whose adoption and use were expected to relate to those demographic strata largely represented by the new respondents (young, low-income, Hispanic) showed statistically significant changes (Hitczenko 2015a). The presumed benefit of improved coverage led the CPRC to put a greater emphasis on preserving coverage in the 2013 sample by focusing on recruiting many of the respondents newly added in 2012 who were responsible for the increased coverage. Because a lower budget allowed for an expected sample size of around 2,000 individuals in 2013, this meant that fewer longtime panelists, who tended to skew toward older individuals with higher incomes, were recruited. The size of the longitudinal panel dropped.

The 2014 SCPC was administered to individuals from two survey vendors, a strategy that was partly motivated by a desire to see how robust population estimates are to the choice of panel. In order to make comparisons across the two panels efficiently, it was necessary to reduce the number of respondents from the ALP so as to better balance the number of respondents from each vendor. Therefore, in 2014, the target number of respondents from the ALP was 1,750, more than 10 percent lower than in years other than 2012 and 2008. As in 2013, the respondent selection strategy involved maximizing the longitudinal panel subject to expected sample representativeness with respect to 15 strata, based on race, age, and income:

Table 5: Strata used in 2014 SCPC sample selection.

Stratum	Race	Age	Income	Stratum	Race	Age	Income
1	White	18−39	<$30K	10	Non-white	18−39	<$30K
2	White	18−39	$30K-$60K	11	Non-white	18−39	≥$30K
3	White	18−39	≥$60K	12	Non-white	40−55	<$30K
4	White	40−55	<$30K	13	Non-white	40−55	≥$60K
5	White	40−55	$30K-$60K	14	Non-white	56+	<$30K
6	White	40−55	≥$60K	15	Non-white	56+	≥$30K
7	White	56+	<$30K				
8	White	56+	$30K-$60K				
9	White	56+	≥$60K				

The first step was to determine the necessary number of respondents to invite from each stratum so that, after factoring in nonresponse, the expected sample of 1,750 has strata composition as close as possible to the U.S. population, as determined by the 2014 Current Population Survey Annual Social and Economic Supplement (CPS), administered in March. Individuals are selected for each stratum according to experience with the SCPC, with preference given to people who have more experience. Likely due to the emphasis on preserving population frequencies in 2013, all 1,132 individuals who participated from 2009 to 2013 were invited to participate in 2014. Additionally, as Table 6 indicates, most of those individuals who were invited to participate in the 2014 SCPC had some prior SCPC experience. Only 11 individuals who had never taken the SCPC before were invited in 2014, resulting in 10 respondents with no prior experience.

ALP members who are selected for a survey receive an email message with a request to visit the ALP webpage and fill out the survey's online questionnaire. Anyone who logs on to the survey is considered a participant in the survey, no matter how much of the survey he or she completes. Naturally, not everyone will participate. Table 6 provides the participation rates for individuals as new and existing SCPC panelists from 2012 to 2014.

Table 6 indicates that retention rates among individuals who had taken the SCPC at some prior point were quite high. Around 87 percent of those who had participated before agreed to participate in 2014, a rate that is generally consistent with those of previous years (it is little lower than in 2013, but a little higher than in 2012). As a result, there remains a sizeable contingent of respondents with several consecutive years' worth of experience. Indeed, as Figure 1 shows, there are 1,036 individuals in the six-year panel from 2009 to 2014 and an additional 315 individuals who participated in all three years from 2012 to 2014. The

Table 6: The sources of the 2012, 2013, and 2014 SCPC respondents. "Repeat" refers to those who had also participated in previous editions of the SCPC, while "New" refers to those who had not.

2012 SCPC Recruitment			
Respondent Type	# Eligible	# Participated	Participation Rate
Repeat	2,473	2,065	83.5
New	1,197	1,111	92.8

2013 SCPC Recruitment			
Respondent Type	# Eligible	# Participated	Participation Rate
Repeat	1,989	1,812	91.1
New	395	277	70.1

2014 SCPC Recruitment			
Respondent Type	# Eligible	# Participated	Participation Rate
Repeat	2,059	1,799	87.3
New	11	10	90.9

participation rate among new respondents in 2014 was 90.9 percent, falling slightly below the corresponding rate from 2012 but above the corresponding rate from 2013. The final 2014 sample of ALP respondents consisted of 1,809 individuals.

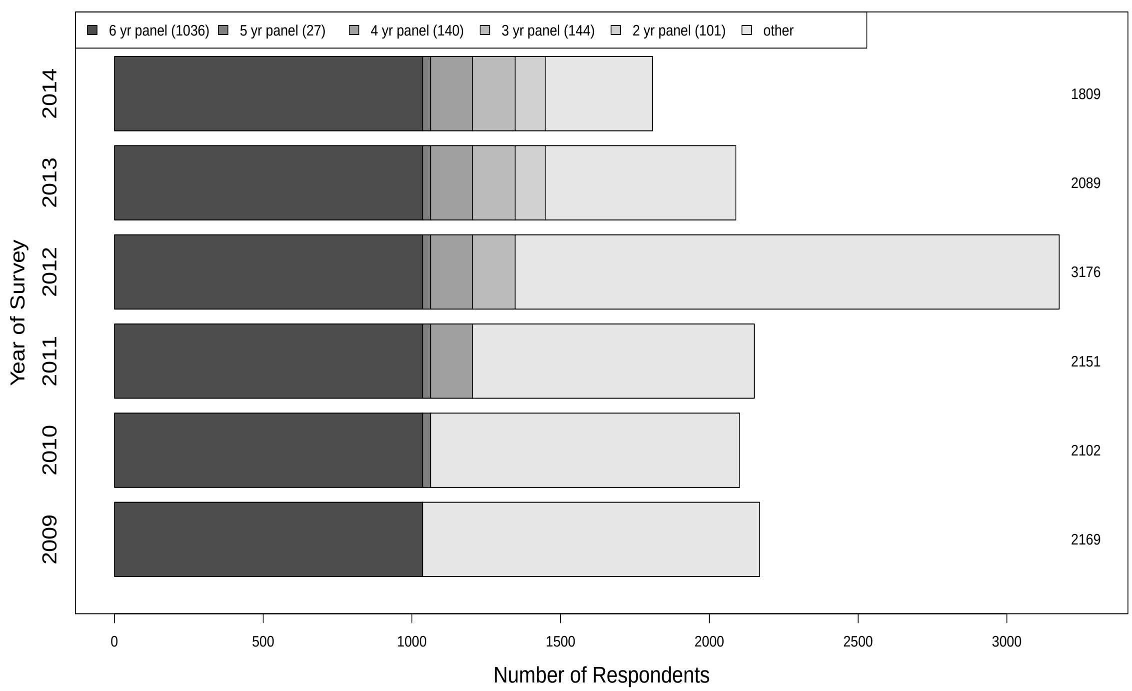

Figure 1: The annual composition of the SCPC respondents.

4.3 Survey Completion

Each year, the SCPC is fielded in the fall with the goal of having most of the surveys completed in the month of October. The desire to standardize this response period is three-fold. First, 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. Second, from an economic point of view, the month of October is a reasonably representative month with respect to certain economic variables such as employment or sale volumes; it includes no major holidays and falls between summer and winter. Although we ask respondents for responses in a "typical" month, it is possible that recent behavior may influence responses. Finally, the Diary of Consumer Payment Choice (DCPC) was also administered in October (a pilot version in 2010 and 2011 and the full version in 2012 and 2015), and responses from both surveys can be linked more easily if they correspond to the same period of economic activity.

As mentioned previously, selected individuals receive an invitation to take the SCPC survey via email. The email is sent to everyone simultaneously, and the day on which this occurs is the "release date" of the survey. The respondent is offered a $20 financial incentive to complete the survey. Each respondent can begin the survey at any point after receiving the invitation. The time of starting is defined as the time when the individual first logs on to the survey, and the time of completion is defined to be the day when the respondent logs off for the final time. 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. Because our analysis utilizes data from everyone who ever participated (logged on), these distinctions are not vital to further analysis or results. Individuals who have not logged on after a few weeks are given reminders to do so with follow-up emails.

Figure 2 shows the proportion of surveys completed by each calendar day within each of the years from 2009 to 2014. This plot shows that, while in 2009 the survey was not released until the second week of November, the release date in the past four years has generally been within a few days of the beginning of October. While most previous surveys were released in the last few days of September, the 2014 version was released around a week later, on October 6, 2014. To the extent that taking the survey near the beginning of a month influences responses, one should be aware of this when comparing the 2014 results to those of previous years. However, despite this delay, as in the past four years, close to 90 percent of surveys were completed by the end of October. In every year, only about 2 percent of individuals never log off.

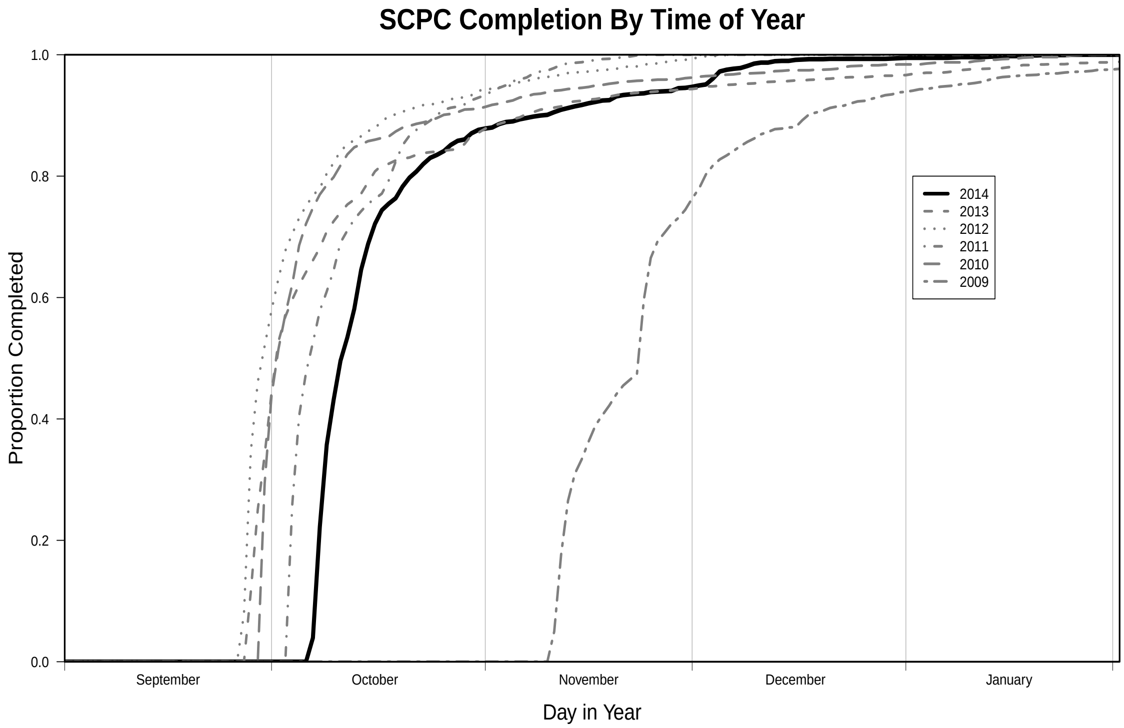

Figure 2: The proportion of respondents who completed the survey as a function of the date within the year.

Figure 3, which shows the proportion of surveys completed as a function of the number of days since the survey was distributed for the 2009−2014 versions, gives a better sense of the distribution of days until completion. Except for 2009, the distribution of completion rates from the time of release is very similar across years. From 2010 to 2014, all of the surveys were completed by over 50 percent of the respondents within two days of its being made available and by 91 percent within a month. In 2009, while 90 percent of the respondents had completed the survey after a month, only about 18 percent had done so after two days.[3]

An important aspect of the SCPC time-series data made evident by the completion data relates to the relatively wide range of dates within a year during which surveys are taken.

[3]The 2009 SCPC went into the field on Tuesday, November 10, 2009. The fact that the following day was a public holiday (Veterans Day on November 11, 2009) might explain why few respondents answered the survey after a day.

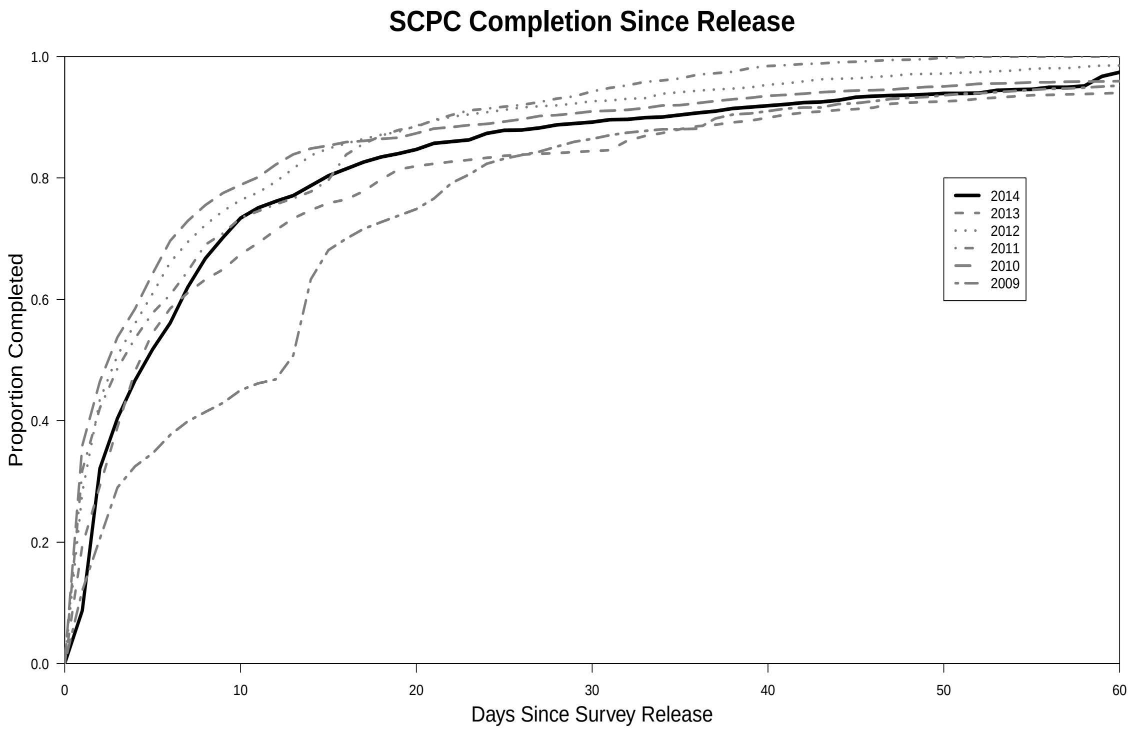

Figure 3: The proportion of respondents who completed the survey as a function of the number of days since the survey was received. The spike at 14 days for 2009 is likely the result of an email reminder sent out two weeks after the survey was distributed. This spike can be seen in Figure 2 as well.

Although approximately 80 percent of surveys are completed within two or three weeks of the release date, as Figure 3 makes clear, the range of completion dates for the remaining surveys spans a period of months. What is more, the later release of the 2009 survey ensures that there is little overlap in the completion periods for the SCPC in this and the following years. As a result, comparisons across years may be influenced by differences due to seasonal behavior as well as by general trends across years. For example, if typical behavior changes in November due to the ensuing holiday season, payment use responses in the 2009 SCPC may reflect this, while those in the other years will not. This type of temporal gap is 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 issues of comparability. An effort to minimize this seasonal effect has been a motivating factor in the the consistent timing of the release in the past three surveys near the end of September.[4]

[4]The 2012 Diary of Consumer Payment Choice was administered over a strict calendar time period (September 29−November 2, 2012) and is linked to the SCPC, so the SCPC was consistently launched at the end of September or beginning of October from 2011 to 2013.

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[5]. Figure 4 indicates that from 2009 to 2012 the survey was getting longer, with the median completion time going from 30 minutes in 2009 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 considerably shorter than all but the 2009 version. Interestingly, despite having the shortest median time, the 2014 survey also had the highest percentage of individuals who took over two hours to complete the survey. As the completion time is calculated by the amount of time that has passed between the initial log on and the final log off, this could be a result of respondents effectively taking the SCPC not all at once, but in smaller increments over time.

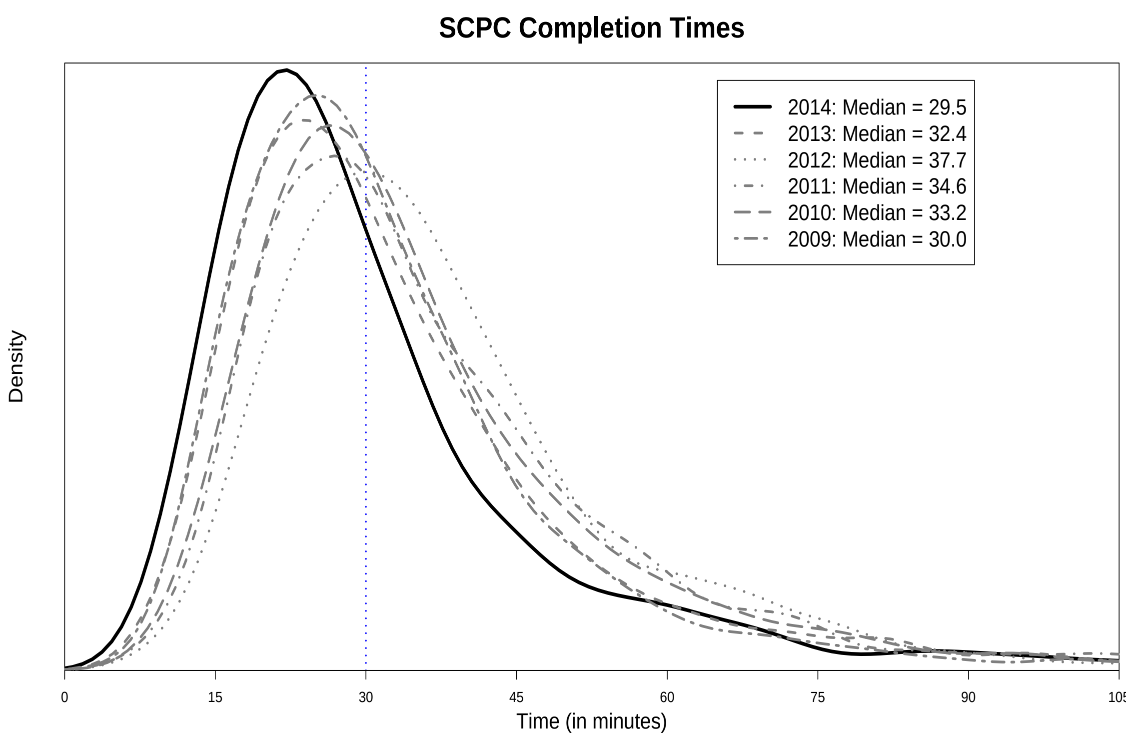

Figure 4: The proportion of respondents who completed the survey as a function of time. The vertical line at 30 minutes represents the intended average length of completion.

[5]The distribution is highly skewed to the right, since completion time is defined as the difference in minutes between the time of first log-in to the survey and the last log-out. A log-out requires responding to the very last question in the survey. Individuals who take breaks while taking the survey will thus have long completion times. In addition, as noted above, around 2 percent of individuals never log out of the survey.

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 2014 SCPC has around 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 0.4 to 2.5 percent, as shown in Table 7. Because of those who fail to complete the survey, there is evidence that the response rates generally drop the farther along one goes in the survey. However, 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.

Table 7: Non-response rates for eight questions in the 2014 SCPC.

Question	fr001_a	as003a4	pa001_a	pa050	pa053	pa024	ph006	de011
Section in SCPC	II	III	IV	IV	IV	IV	VI	VII
Non-response rate (%)	0.44	1.00	1.88	1.33	1.49	2.10	2.10	2.49

5 Sampling Weights

5.1 Post-Stratification

An important goal of the SCPC is to provide estimates of payment statistics for the entire population of U.S. consumers over the age of 18. As mentioned in Section 4, the ALP is a collection of volunteers from a variety of existing databases. A direct implication of this fact is that any SCPC sample will not be a probability sample, making probability-based weighting to generate population-wide inferences impossible. Nevertheless, recent work by Wang et al. (2009) suggests that nonrepresentative polling can provide relatively accurate estimates with appropriate statistical adjustments.

The aforementioned evolution of the ALP, as well as the CPRC's focus on preserving the longitudinal aspect of the sample, suggests that the SCPC sample itself is not necessarily representative of the U.S. population of consumers. Table 8 shows the unweighted sample proportions for a set of chosen demographic categories for various renditions of the SCPC along with the weighted ones for the 2014 SCPC sample. As discussed in Section 4, a concerted effort was made in 2013 to improve representativeness by effectively exchanging respondents with a long history of participation for new respondents from under-represented strata. This sampling strategy manifests itself in a significant improvement in the unweighted distributions in 2013 as compared with the U.S. population. Males, the young, non-whites and those with lower household incomes are considerably better represented in the 2013 SCPC than in previous years. The 2014 unweighted sample also seems to better match the U.S. population than the 2012 unweighted sample does, although perhaps not to the extent of the 2013 sample. This may be due to the drop in sample size, which makes random fluctuations in participation more influential in terms of overall composition. As in previous years, education level, which is not one of the demographic variables used in sample selection, has the worst unweighted representation, with the ALP constituting a better-educated group than the U.S. population as a whole. This discrepancy results partly from the fact that education is not a variable used to match the sample to the U.S. population in the sample selection process.

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). The improved coverage and better unweighted matching of the sample to the CPS results in sampling weights with lower variance, as the standard deviation of the weights changes from 41.9 and 44.8, respectively, in 2011 and 2012, to 34.2 in 2013 and 39.4 in 2014. The increase in variability of the weights from 2013 to 2014 is largely due to a decrease in sample size of around 15 percent.

Table 8: Unweighted percentages for various marginal demographics in the 2012, 2013, and 2014 SCPC sample, as well as weighted percentages for the 2014 SCPC. The weighted values are based on CPS data.

Demographics		Unweighted 2012 SCPC	Unweighted 2013 SCPC	Unweighted 2014 SCPC	Weighted 2014 SCPC
Gender	Male	43.6	46.3	47.0	48.1
	Female	56.4	53.7	53.0	51.9
Age	18−24	3.0	4.0	2.8	4.8
	25−34	15.7	20.6	20.9	26.1
	35−44	13.1	17.3	16.7	15.1
	45−54	22.3	21.1	19.7	18.9
	55−64	26.0	20.7	20.9	16.7
	65 and older	20.0	16.4	19.0	18.5
Race	White	85.5	77.4	79.9	75.0
	Black	8.2	11.3	10.3	11.7
	Asian	1.8	2.5	2.4	2.5
	Other	4.4	8.7	7.5	10.8
Ethnicity	Hispanic	7.3	17.2	12.5	17.4
Education	No HS diploma	2.7	3.6	2.7	6.9
	High School	15.9	16.8	14.5	35.0
	Some College	36.8	37.9	39.0	28.7
	College	25.2	24.6	25.2	16.9
	Post-graduate	19.4	17.1	18.7	12.5
Income	$< \$25K$	17.0	23.3	19.4	24.1
	$\$25K - \$49K$	24.7	27.7	26.0	23.8
	$\$50K - \$74K$	21.6	19.5	21.9	20.2
	$\$75K - \$99K$	14.5	10.8	11.3	10.5
	$\$100K - \$124K$	9.7	7.7	9.2	9.6
	$\$125K - \$199K$	9.0	8.1	8.8	9.0
	$\geq \$200K$	3.5	2.9	3.4	2.8

5.2 Raking Algorithm

Sampling weights are generated by RAND, using a raking algorithm (Deming and Stephan 1940; Gelman and Lu 2003). 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. Continuous variables such as age and income are recoded as categorical variables by assigning each to one of several disjoint intervals. For example, Table 8 shows six classifications for age and seven classifications for income. 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. Table 9 shows the variables used in weighting as well as the levels within each variable. 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. More precisely, the weighting algorithm is performed using the 31 pairs of demographic variables shown in Table 9.

Table 9: The set of weighting variables. "M" stands for male, and "F" stands for female. The highest income brackets for single households were combined to avoid small cell sizes.

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

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

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

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

The socio-economic variables chosen for the raking procedure result from recent research conducted by RAND regarding the sampling properties of weights based on different demographic factors. First, a new imputation algorithm for all possible socio-demographic variables was developed to allow for weights based on a wider range of consumer information. 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. 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 9 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 better correct 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 9 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.

Because the ALP 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 largely attributed 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 ALP 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 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 carefully examine the data 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 data cleaning.

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 2014. Nevertheless, just as in Angrisani, Foster, and Hitczenko (2015), the methodologies are described in detail for all variables except those relating to dollar values reloaded and stored on prepaid cards, which were removed from the 2013 SCPC. It should be noted that all procedures are applied retroactively to the data of previous years, so data variables from the 2008 – 2013 surveys may have different values from those in previous data releases. The edited variables are used for analysis by the CPRC, most notably to generate population estimates provided in the SCPC tables. However, both edited and unedited data are released to the public. A guide to 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, RAND 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 ALP. For household income and household size, attributes that could easily change within a year, values are imputed by RAND through logistic regression models for the purpose of creating post-stratification weights. The imputations are used only 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 (Hitczenko (2013); see Daamen and de Bie (1992) and Friedman, Herskovitz, and Pollack

(1994) for general discussion on anchoring effects), but not of economic significance. 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 the development of 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 for 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 takes the most liberal approach in that all responses are kept as given for as much of the sequence as possible. At all subsequent levels, inconsistent responses 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.

Identification of data that are possible, but very unlikely, is much more difficult, especially for economic variables such as cash holdings and value of assets, as it involves comparing a data entry within the context of heterogeneity of behavior within the population. 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.

Below we outline the considerations and economic motivations in cleaning several different variables and provide adopted algorithms for each. We focus on identifying potentially invalid data entries in the right tails, as these are most likely be influential data points, or those whose inclusion or exclusion in any inferential analysis causes a significant difference in estimates (Bollen and Jackman 1990; Cook and Weisberg 1982). To the degree possible, the procedures adopted by the CPRC rely on economic intuition to identify potentially invalid data entries in the right tails. 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. Section 6.3 provides the details of which survey variables are edited and how new and old versions of such variables are named and identified. 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 certain cases, new data patterns have made previous editing strategies ineffective. In such cases, we update the algorithm or fall back on simpler strategies. As noted above, the raw (uncleaned) data are available, so researchers are free to preprocess the data as they see fit.

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 natural frequencies of use. 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 are 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 determined to be 300, representing 10 payments per day for 30 days. 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 each other. 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.

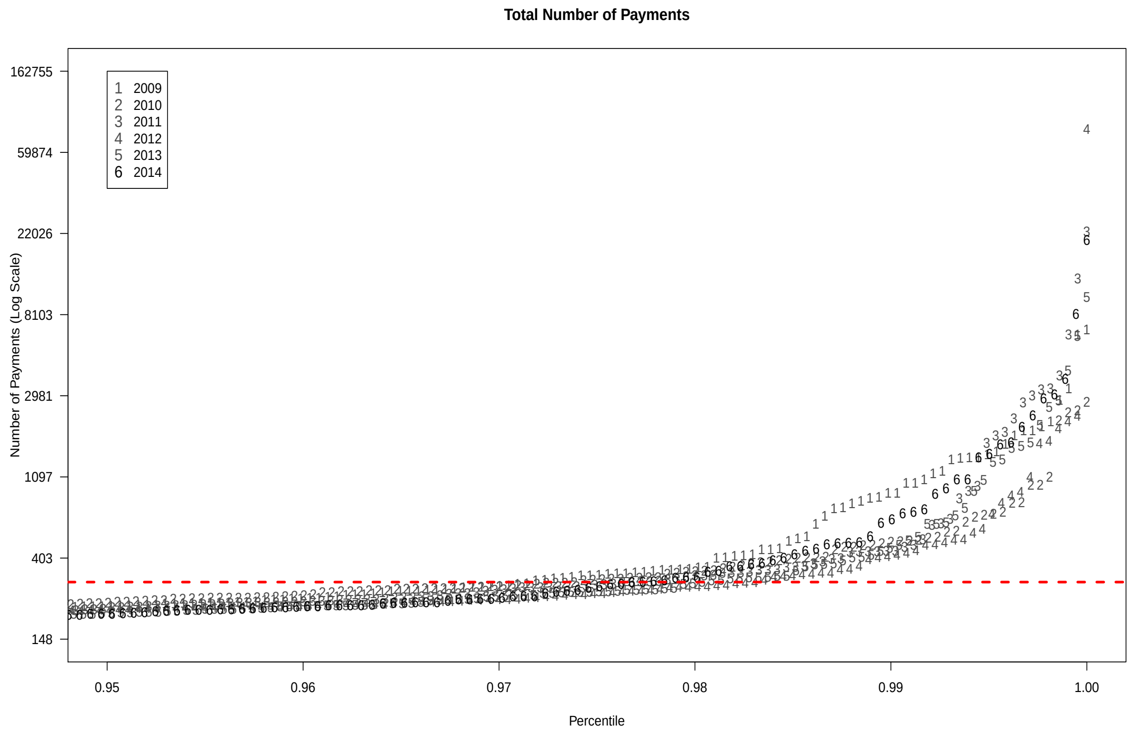

Figure 5: The log-values of the largest 5 percent of the total monthly payments data plotted against the percentiles for past six years of data. The dashed line represents 300 monthly payments.

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 $\mathcal{P}_{it}$ be the bundle adopted by individual i. For example, $\mathcal{P}_{it} = \{1, 2\}$ 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

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

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

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

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 $j \in \mathcal{P}_{it}$. 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 is 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). It is evident that despite the use of 300 as the cleaning parameter, the algorithm allows individuals to have more payments. 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 50, although the changes to the number of payments made are relatively small.

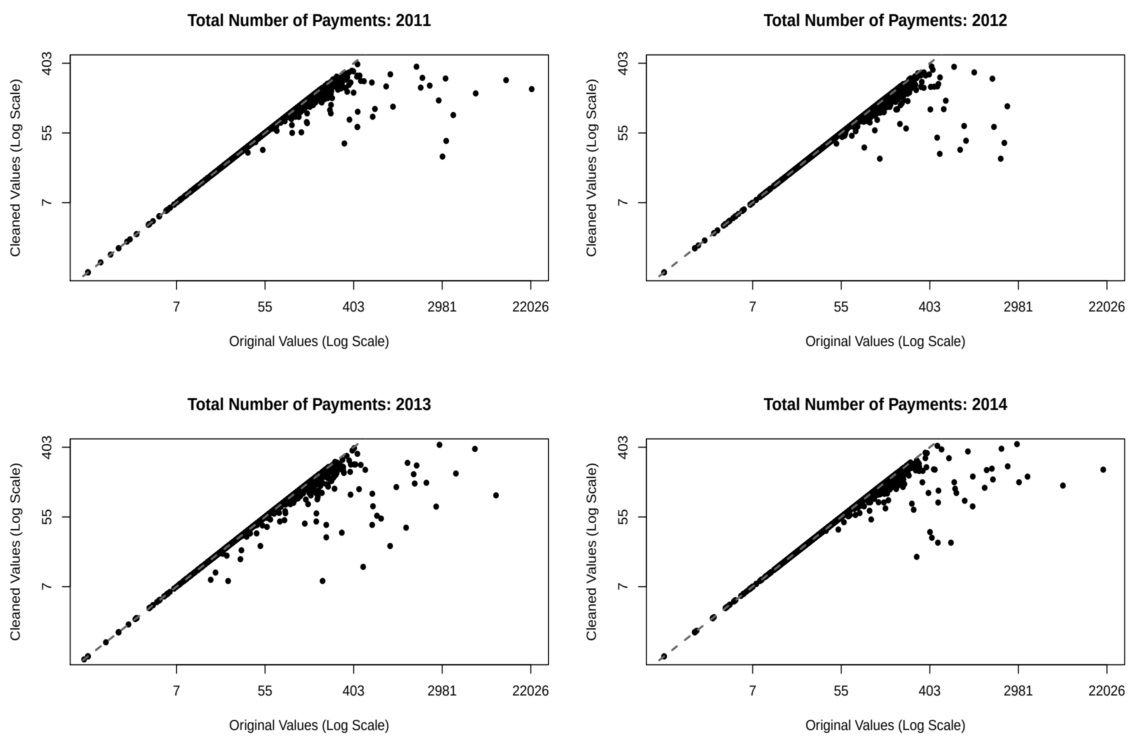

Figure 6: The log-values of the cleaned total monthly payments data plotted against the log-values of the original values.

Algorithm 1 Preprocessing: Number of Monthly Payments

```
for i = 1 : N do
  Determine P_it
  for j ∈ P_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
```

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. Cash withdrawal since the 2009 SCPC is reported as a combination of four separate variables: frequency of withdrawal at primary and all other locations and typical dollar amount per withdrawal at primary and all other locations. Because reported dollar amounts correspond to typical values, which could represent the mean, the median, or the mode, the value determined by multiplying the reported frequency and the dollar amount does not necessarily correspond to the average total cash withdrawal either for primary or for all other locations. In preprocessing the cash withdrawal values, data for primary and all other locations are treated separately. The editing process, revised for the 2011 and 2012 data, is described below.

Assuming that N independent individuals report positive cash withdrawal in a typical month, let $C_{it} = A_{it}F_{it}$, 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 than those who take out smaller values suggests a negative correlation between the two variables. Thus, the joint distribution will take the form

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

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

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

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

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

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

Algorithm 2 Preprocessing: Monthly Cash Withdrawal

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

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

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

In practice, μ and Σ are not known and must be estimated from the data. We use lower-case notation, such as $w_i = (\log(a_i), \log(f_i))$, to represent the actual values observed in any given survey year, and estimate the bivariate mean and covariance with $\hat{\mu}$, the sample mean, and $\hat{\Sigma}$, the sample covariance. The entire procedure is outlined in Algorithm 2. Figure 7 shows the result of the heretofore outlined cleaning algorithm applied to the 2012 cash withdrawal data from the primary source. The plot shows an ellipse corresponding to the 98 percent confidence interval for any observation from the Log-Normal distribution defined by the parameters estimated from the sample. Via the preprocessing, all points outside this region are moved to the nearest point on the ellipse.

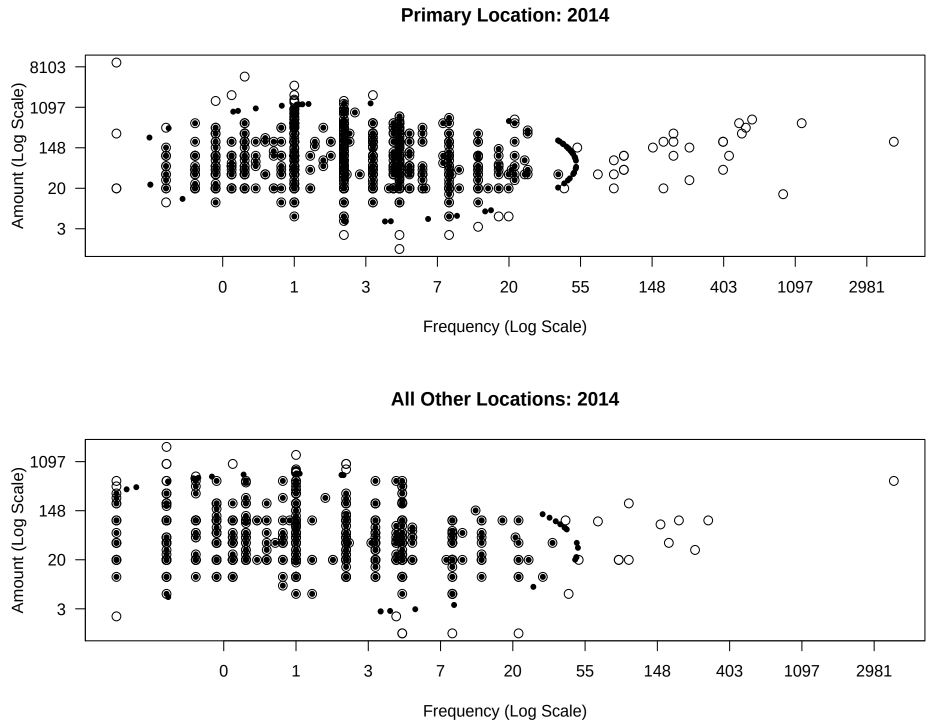

Figure 7: A diagram of the cleaning algorithm for cash withdrawal data in 2014. Circles represent original data and filled-in points represent the cleaned data (both plotted on the log-scale).

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 $\log(f_i) < \hat{\mu}_F$.

6.2.3 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 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 for the procedure or delve into any details.

Figure 8 shows the distribution of the right tails of cash holdings for each of the two variables. As indicated, this cleaning procedure results in no edits to the cash holdings on person. The maximum reported values for the 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.

With respect to cash holdings at home, a datapoint corresponding to \$600,000 in 2012 was winsorized to \$100,000, which was the next highest value and the highest reported value in the other years. No changes to datapoints were made for 2014.

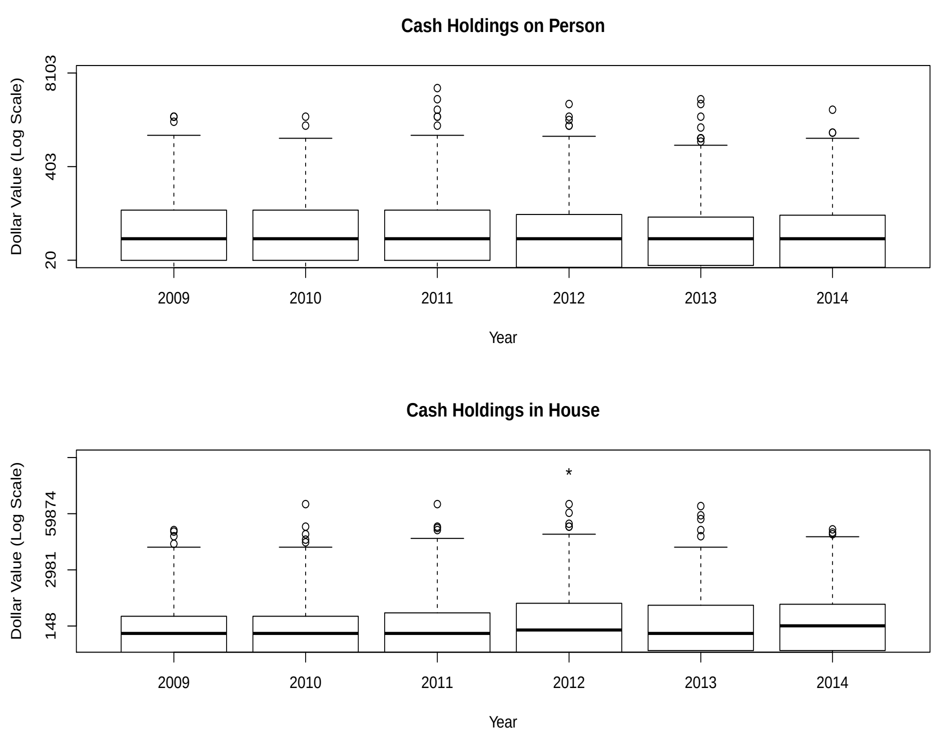

Figure 8: Boxplots of right tails of cash holdings. The asterisk represents the only edited value.

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 variables that are defined as functions of edited variables are created using edited data. Perhaps most importantly, all variables relating to payment use from "csh_typ," which defines the number of cash payments, to "paper_typ," which defines the number of payments made with cash, check, or money order, to "tot_pay_typ," which defines the total number of monthly payments, are aggregates of the lowest-level entries for payment use. All statistics for such 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 classified by an "_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_amnt_1st" holds all edited entries for the dollar value of cash withdrawn from the primary location, while "csh_amnt_1st_unedited" defines the unedited version of the data. Table 10 lists all variables that are edited by the CPRC.

Table 10: Summary of edited variables. "Underlying variables" are any survey or created variables that 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.
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 details the model that provides a framework for achieving both of these goals. 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 category $j = 1, \ldots, J$ 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 the entire population, the identifier i will range from 1 to the total number of consumers over the years in question. However, within the sample, the respondent identifier i ranges from 1 to N, where N represents the total number of unique respondents in all six years. Let w_{it} designate the survey weight of person i in year t. J will naturally vary with the area of application and, for the 2014 SCPC, $T = 7$, with the years counted starting from 2008. Taking the function $1_{[t=x]}$ to be 1 when $t = x$ and 0 otherwise, a natural model for the population means is

$$Y_{ijt} = \mu_{j1} 1_{[t=1]} + \mu_{j2} 1_{[t=2]} + \ldots + \mu_{jT} 1_{[t=T]} + \epsilon_{ijt}, \tag{1}$$

where ϵ_{ijt} are mean 0 random variables with $\text{Var}(\epsilon_{ijt}) = \sigma^2_{jt}$ and $\text{Cov}(\epsilon_{ijt}, \epsilon_{i'j't'}) = \rho_{jtt'}$ for $i = i'$ and $j = j'$. This model is focused on estimating the population means, $\mu_j = [\mu_{j1}\ \mu_{j2}\ \ldots \mu_{jT}]^T$, and it can correspond to a variety of underlying processes on the micro-economic scale. For example, in the context of typical monthly payments, such a model could correspond to a process in which each person conducts a random number of total transactions, where the totals are statistically dependent for each consumer across years. Then, the payment option used for each transaction is chosen independently according to some set of probabilities that are also allowed to vary from year to year.

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 category j in year t, and let $N_{jtt'}$ represent the number of respondents who gave responses for category 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 category 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

$$\hat{\mu}_j = \left(\mathbf{X}_j^T \mathbf{W}_j \mathbf{X}_j\right)^{-1} \mathbf{X}_j^T \mathbf{W}_j \mathbf{Y}_j. \tag{2}$$

Before we proceed, note that the population estimates calculated from the model, given in (2), 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 category j in year t and t', then

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

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

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

and

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

7.1 Standard Errors and Covariances

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 2010 − 2014 SCPC are available at http://www.bostonfed.org/economic/cprc/SCPC.

The standard errors themselves give a sense of how much faith we have that the estimates are accurate given the stratum weights. Larger standard errors will denote more uncertainty in the true population values. As the standard error tables show, it is generally true that the standard errors in the year 2008 are considerably higher than those of the later years. This is so primarily because the sample size grew considerably from 2008 to 2009, giving a more accurate picture of the average behaviors.

The off-diagonal elements of $\hat{\mathbf{\Lambda}}_j$ correspond to the $\text{Cov}(\hat{\mu}_{jt}, \hat{\mu}_{jt'})$, which, when divided by $\sqrt{\text{Var}(\hat{\mu}_{jt})\text{Var}(\hat{\mu}_{jt'})}$, yield a correlation. This correlation reflects the extent to which estimates based on the samples within the assumed sampling scheme relate to one another. If the samples for two years did not include any of the same individuals, independence across individuals would imply that the correlations would be zero. However, as there is overlap, one expects positive correlations between estimates for two different years.

As an example, consider the results for the population average number of typical weekly debit card uses conditional on credit card adoption (*ccu*) and the proportion of the population that adopts credit cards (*cca*). For the data from the past three years, the correlation matrices for the two statistics are given by

$$\text{Corr}(ccu_{12,13,14}) = \begin{bmatrix} \mathbf{0.8} & 0.36 & 0.34 \\ 0.36 & \mathbf{0.8} & 0.46 \\ 0.34 & 0.46 & \mathbf{0.9} \end{bmatrix} \quad \text{and} \quad \text{Corr}(cca_{12,13,14}) = \begin{bmatrix} \mathbf{0.013} & 0.37 & 0.44 \\ 0.37 & \mathbf{0.012} & 0.46 \\ 0.44 & 0.46 & \mathbf{0.013} \end{bmatrix},$$

where the diagonal values in bold represent standard errors.

7.2 Functions of Population Means

While the most interesting population parameters are the μ_{jt} in (1) themselves, we are also interested in some variables that are functions of these population parameters. Perhaps the two most interesting 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 the averages rather than the average of the individual growth rates. We thus let

$$g_{jt} = \frac{\mu_{j,t+1} - \mu_{jt}}{\mu_{jt}} \tag{3}$$

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

$$s_{jt} = \frac{\mu_{jt}}{\sum_{k=1}^{J} \mu_{kt}} \tag{4}$$

be the share of category j in year t.

The macroeconomic definitions used in (3) and (4) should be contrasted with their microeconomic alternatives. The former involve defining individual shares for each category, $s_{ijt} = \frac{y_{ijt}}{\sum_{k=1}^{J} y_{ikt}}$ and estimating s_{jt} by applying (1) and (2) 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 category 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 (1) represents a population mean 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 used 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 taken 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 2014 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 thus 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.

8 Hypothesis Tests for Temporal Changes in Consumer Payments

Knowledge of $\hat{\mu}_j$ and $\hat{\mathbf{\Lambda}}_j$ for all $j = 1, \ldots, J$ also allows one to make inferences and test hypotheses about the population across the different years. In the following subsections, we delineate and conduct a variety of hypothesis tests with the general goal of assessing changes in population estimates across years.

Sections 8.1 – 8.3 provide the methodology for three different types of hypothesis tests. The applications for the SCPC data are found below in Section 9. The hypothesis tests generally compare 2013 estimates to 2014 estimates, although a few compare the growth rate from

2012 to 2013 to the growth rate from 2013 to 2014. Test results themselves are organized by subject matter in Section 9.

8.1 Hypothesis Tests for Means

Perhaps the most basic assertion one would like to make is the degree to which the population means change over time. Therefore, in the context of the model outlined above, we consider the following hypotheses:

$$H_o : \mu_{jt} = \mu_{jt'} \qquad H_a : \mu_{jt} \neq \mu_{jt'}.$$

In order to do so, we need to estimate $\mathrm{Var}(\mu_{jt'} - \mu_{jt})$, which we do by estimating the identity:

$$\mathrm{Var}(\mu_{jt'} - \mu_{jt}) = \mathrm{Var}(\hat{\mu}_{jt'}) + \mathrm{Var}(\hat{\mu}_{jt}) - 2\mathrm{Cov}(\hat{\mu}_{jt'}, \hat{\mu}_{jt})$$

with

$$\hat{\mathrm{Var}}(\mu_{jt'} - \mu_{jt}) = \hat{\mathbf{\Lambda}}_j[t', t'] + \hat{\mathbf{\Lambda}}_j[t, t] - 2\hat{\mathbf{\Lambda}}_j[t, t'].$$

Now, under the null hypothesis, the test statistic

$$Z = \frac{\hat{\mu}_{jt'} - \hat{\mu}_{jt}}{\sqrt{\hat{\mathrm{Var}}(\mu_{jt'} - \mu_{jt})}}$$

is approximately distributed as a standard normal distribution. This fact allows us to calculate p-values and accordingly accept or reject the null hypotheses.

8.2 Hypothesis Tests for Growth Rates

In addition to changes in population means, tests for the significance of the change in the growth rates of the means from one year to the next are developed. With the growth rate in a given year t defined as in (3), $\Delta_{jt} = g_{j,t+1} - g_{jt}$ is the change in growth rates over two consecutive years, which, written in terms of the means, takes the form

$$\Delta_{jt} = \frac{\mu_{j,t+1}}{\mu_{jt}} - \frac{\mu_{jt}}{\mu_{j,t-1}}.$$

Of course, Δ_{jt} is a nonlinear function of the means, which means that conducting a hypothesis test is no longer as simple. However, the delta method (Casella and Berger 2002) allows one to approximate the distribution of $\hat{\Delta}_{jt}$ by approximating the relationship between Δ_{jt} and the μ_{jt} through linearization. Since $\hat{\mu}_{jt}$ are close to normally distributed, a linear function of these variables will also be normally distributed. Let $f_{jt}(\cdot)$ be the function that maps the vector μ_j to $\mathbf{\Delta}_{jt}$ and let $[\partial \mathbf{f}_{jt}]$ be the 1×3 vector such that the i^{th} element is $\frac{\partial f(\mu_j)}{\partial \mu_{ji}}$. Then, if the $\hat{\mu}_j$ are asymptotically normally distributed, the delta method tells us that

$$\hat{\Delta}_{jt} \rightarrow_D N(f(\hat{\mu}_{jt}), [\partial \mathbf{f}_{jt}]\hat{\mathbf{\Lambda}}_j[\partial \mathbf{f}_{jt}]^T),$$

where $\rightarrow_D$ indicates a convergence in distribution as the sample size gets larger.

With this result, the test for the null hypothesis

$$H_o : \Delta_{jt} = 0 \qquad H_a : \Delta_{jt} \neq 0,$$

relies on calculating the statistic

$$z = \frac{\hat{\Delta}_{jt}}{\sqrt{[\partial \mathbf{f}_{jt}]\hat{\mathbf{\Lambda}}_j[\partial \mathbf{f}_{jt}^T]}}$$

and using the normal distribution to calculate a p-value. While the assumption of normality of the resulting Δ_{jt} is only an approximation, it is likely to be a poor one if μ_{jt} or $\mu_{j,t-1}$ is small (near 0). In this case, the approximation of local linearity used in the delta method is not a good one, and the assumed distribution of Δ_{jt} does not match the real one, which will be more skewed than a normal density curve. This means that the p-value calculated from the above process might be a poor approximation of reality.

8.3 Hypothesis Tests for Shares

From an economic standpoint, it is not just the level of use of each payment method but also the relative prevalence of payments made by a particular payment method that matters. The relative prevalence, in many ways, most directly gets at the heart of a consumer's choice of payment method. One can view each individual as needing to make some (random) number of payments over the course of a period of time, including for bills, groceries, and other fairly regular payments, along with other, less predictable payments. Given these necessary payments, it is up to the consumer to decide how to execute each transaction. The decision reflects a variety of factors such as convenience, cost, and acceptance of the payment method, which is why the prevalence of payments is important to economists. The level of use or growth rate will not reflect these aspects of the decision, since a decrease in use in terms of frequency per month could actually correspond to an increase in prevalence if the total number of payments decreased.

There are two statistics that can be used to measure prevalence. The first statistic is the relative growth differential (RGD), which measures the difference between the growth rate in the use of a particular payment option and the overall growth rate in the total number of payments. After some simple algebra, the RGD for payment option j from year t to t' is

$$G_{jtt'} = \frac{\mu_{jt'}}{\mu_{jt}} - \frac{\sum_{k=1}^{J} \mu_{kt'}}{\sum_{k=1}^{J} \mu_{kt}}. \tag{5}$$

The second commonly used statistic is the share differential (SD), defined to be the difference in the percentage of all payments made by payment option j in two years. The mathematical form is

$$S_{jtt'} = \frac{\mu_{jt'}}{\sum_{k=1}^{J} \mu_{kt'}} - \frac{\mu_{jt}}{\sum_{k=1}^{J} \mu_{kt}}. \tag{6}$$

In each case, the statistics of interest are nonlinear functions of the μ_{jt} and are evidently dependent, making hypothesis testing more complicated. Again, the delta method is used, although now it involves a joint, multi-variable hypothesis test. As mentioned above, normal approximations to growth rates can be poor when the means are close to 0. The share differential will not have this problem in this scenario, because the denominator, as the mean number of monthly payments, will be large, making the linear approximation inherent in the delta method a good one. For this reason, share differential is adopted as a preferred measure of relative prevalence.

Below, the methodology for the multivariate delta method hypothesis test (Casella and Berger 2002), as applied to the share differentials, is explained. For simplicity of notation, let S_j stand for $S_{jtt'}$ in the following paragraphs. The necessity of a multivariate test is due to the clear dependence between S_j and $S_{j'}$. In fact, $S_J = -\sum_{j=1}^{J-1} S_j$. This issue of dependence means that the joint hypothesis test takes the form

$$H_o : S_1 = S_2 = \ldots = S_{J-1} = 0 \qquad H_a : S_j \neq 0 \text{ for at least one } j.$$

Now, let $\hat{\mathbf{S}} = \left[\hat{S}_1 \ \hat{S}_2 \ \ldots \ \hat{S}_J\right]^T$, and let $\mathbf{h}(\mu_t, \mu_{t'})$ be the function that maps the population means to the share differential statistics with $[\partial\mathbf{h}(\mu_t, \mu_{t'})]$, the matrix of partial derivatives $\frac{\partial h(\mu_t,\mu_{t'})}{\partial \mu_{j,k}}$ for $k = t, t'$ and $j = 1, \ldots, J$. Now, letting $\hat{\mathbf{\Lambda}}_{tt'}$ be the data estimate of the covariance of $[\mu_{1t} \ldots \mu_{Jt} \ \mu_{1t'} \ldots \mu_{Jt'}]^T$, the multivariate version of the delta method tells us that

$$\hat{\mathbf{S}} \quad \to_D \quad N\left(\mathbf{h}(\hat{\mu}_t, \hat{\mu}_{t'}), [\partial\mathbf{h}(\mu_t, \mu_{t'})]\hat{\mathbf{\Lambda}}_{tt'}[\partial\mathbf{h}(\mu_t, \mu_{t'})]^{-1}\right).$$

For simplicity of notation, let

$$\mathbf{C}_{tt'} \quad = \quad [\partial\mathbf{h}(\mu_t, \mu_{t'})]\hat{\mathbf{\Lambda}}_{tt'}[\partial\mathbf{h}(\mu_t, \mu_{t'})]^T.$$

The matrix $\mathbf{C}_{tt'}$ estimates the variances and covariances of the sample statistics $S_{jtt'}$ for $j = 1, \ldots, J$. Given this approximate multivariate normal distribution of dimension J, it is known that under the null hypothesis, the statistic

$$Z \quad = \quad \hat{\mathbf{S}}_{tt'}^T \mathbf{C}_{tt'}^{-1} \hat{\mathbf{S}}_{tt'}$$

will be approximately Chi-square distributed with J-1 degrees of freedom. Therefore, $Z \sim \chi^2_{J-1}$, a fact that can be used to calculate a p-value corresponding to the hypothesis.

Of course, such a test provides insight only into whether the collection of share differentials is significantly different from the vector $\mathbf{0}$, but it is impossible to attribute the cause of the rejection to any particular payment method. However, one can consider whether the exclusion of any choice would make the relative share differentials of the remaining $J-1$ choices consistent with the null hypothesis. Determining the joint 95 percent confidence intervals under the null hypothesis and studying the range of values observed within this interval for each payment choice provides some insight into this. In the case of a normal distribution and a null hypothesis that $S_j = 0$, this turns out to correspond to the one-dimensional 95 percent confidence interval for each option.

In addition to the one-dimensional 95 percent confidence intervals, it is useful to calculate the one-dimensional p-value for each observed share differential under the hypothesis that $S_j = 0$. While there is no straightforward way to determine which choice will result in the most similar set of all possible $J-1$ share differentials based on the calculated p-values and confidence intervals, choices corresponding to lower p-values and larger distances from the center of the confidence intervals, especially as they correspond to higher shares in the two years, are good candidates.

9 Hypothesis Test Results

In this section, we provide in tabular form the results of hypothesis tests relating to several key economic variables. The statistical foundation is detailed in Section 8. The tests are organized according to concept, namely, adoption of instruments, use of payment instruments, and miscellaneous tests. As discussed previously, the SCPC considers payments in terms of payment instruments and type of transaction. Because certain instruments are naturally grouped together due to similarity, as is the case for transaction types, some hypothesis tests are related to broader groups of each. Specifically, we consider instruments as paper (cash, check, and money order), plastic (credit, debit, and prepaid cards), or online (online banking bill payment and bank account number payments). Similarly, we consider transactions as bills (automatic bill payments, online bill payments, in-person bill payments), online payments, or in-person nonbill payments (retail payments, payments for services, and person-to-person payments). Finally, we consider payment instruments by types of assets and liabilities.

9.1 Adoption of Payment Instruments

Table 11: Adoption rates of payment instruments.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Cash	1.00	1.00	0.00	0.40	0.69
Check	0.83	0.83	-0.00	-0.15	0.88
MO	0.21	0.18	-0.03	-2.24	0.03
Debit	0.78	0.79	0.01	0.57	0.57
Credit	0.70	0.72	0.02	1.23	0.22
Prepaid	0.50	0.52	0.02	0.89	0.37
OBBP	0.54	0.57	0.03	1.88	0.06
BANP	0.60	0.62	0.02	1.30	0.19
Income	0.18	0.17	-0.01	-0.55	0.58

Table 12: Note: "MO" represents money orders, "OBBP" represents online banking bill payments, and "BANP" represents bank account number payment.

Table 13: Adoption rates of payment instrument groups.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Paper	1.00	1.00	-0.00	-0.25	0.80
Card	0.95	0.95	0.01	0.63	0.53
Electronic	0.77	0.79	0.02	1.36	0.17

9.2 Use of Payment Instruments

9.2.1 Changes in Mean Number of Uses

Table 14: Mean number of payments per month by instrument.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Cash	17.89	16.91	-0.98	-1.20	0.23
Check	5.71	5.02	-0.69	-2.58	0.01
MO	0.34	0.30	-0.04	-0.64	0.52
Debit	21.15	20.37	-0.79	-0.89	0.37
Credit	15.31	15.37	0.07	0.10	0.92
Prepaid	0.68	0.70	0.01	0.07	0.95
OBBP	2.98	3.40	0.41	2.01	0.04
BANP	3.26	3.49	0.22	1.32	0.19
Income	0.57	0.54	-0.03	-0.34	0.74
Total	67.90	66.09	-1.81	-1.00	0.32

Table 15: Mean number of payments per month by instrument group.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Auto. Bill	6.75	7.21	0.46	0.98	0.33
Online Bill	6.93	6.74	-0.19	-0.54	0.59
Other Bill	8.56	8.10	-0.46	-1.08	0.28
Online	3.88	3.60	-0.28	-1.14	0.25
Retail	23.86	22.88	-0.98	-1.16	0.25
Service	14.91	14.32	-0.60	-1.02	0.31
P2P	3.01	3.24	0.23	0.84	0.40
Total	67.90	66.09	-1.81	-1.00	0.32

Table 16: Mean number of payments per month by transaction type.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Paper	23.96	22.30	-1.66	-1.82	0.07
Card	37.14	36.43	-0.71	-0.59	0.55
Electronic	6.25	6.88	0.63	2.06	0.04
Total	67.90	66.09	-1.81	-1.00	0.32

Table 17: Mean number of payments per month by groups of transaction types.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Bill	22.24	22.06	-0.19	-0.22	0.83
Online	3.88	3.60	-0.28	-1.14	0.25
In Person	41.78	40.44	-1.34	-1.07	0.29
Total	67.90	66.09	-1.81	-1.00	0.32

9.2.2 Changes in Growth Rates

Table 18: Growth rates of monthly use by instrument.

	Growth Rate 2012 – 2013	Growth Rate 2013 – 2014	Difference	z-stat	p-value
Cash	-2.94	-5.48	-2.54	-0.33	0.74
Check	-12.84	-12.09	0.75	0.10	0.92
MO	-34.33	-11.21	23.13	0.88	0.38
Debit	2.56	-3.72	-6.28	-0.85	0.39
Credit	2.71	0.44	-2.28	-0.28	0.78
Prepaid	-19.27	1.48	20.75	0.65	0.52
OBBP	-7.12	13.78	20.90	1.94	0.05
BANP	0.11	6.83	6.73	0.78	0.44
Income	4.77	-4.90	-9.68	-0.42	0.67
Total	-1.44	-2.67	-1.23	-0.27	0.79

Table 19: Growth rates of monthly use by transaction type.

	Growth Rate 2012 – 2013	Growth Rate 2013 – 2014	Difference	z-stat	p-value
Auto. Bill	10.38	6.83	-3.55	-0.29	0.77
Online Bill	3.75	-2.72	-6.47	-0.72	0.47
Other Bill	-6.61	-5.37	1.24	0.16	0.87
Online	-5.41	-7.29	-1.87	-0.16	0.88
Retail	-1.09	-4.09	-3.00	-0.49	0.63
Service	-2.97	-4.00	-1.03	-0.16	0.87
P2P	-9.89	7.65	17.54	1.39	0.17
Total	-1.44	-2.67	-1.23	-0.27	0.79

Table 20: Growth rates of monthly use by instrument groups.

	Growth Rate 2012 – 2013	Growth Rate 2013 – 2014	Difference	z-stat	p-value
Paper	-6.19	-6.92	-0.73	-0.12	0.91
Card	2.12	-1.91	-4.03	-0.71	0.48
Electronic	-3.48	10.15	13.63	1.77	0.08
Total	-1.44	-2.67	-1.23	-0.27	0.79

Table 21: Growth rates of monthly use by groups of transaction types.

	Growth Rate 2012 – 2013	Growth Rate 2013 – 2014	Difference	z-stat	p-value
Bill	1.27	-0.84	-2.11	-0.33	0.74
Online	-5.41	-7.29	-1.87	-0.16	0.88
In Person	-2.45	-3.21	-0.76	-0.15	0.88
Total	-1.44	-2.67	-1.23	-0.27	0.79

9.2.3 Changes in Share

Table 22: Share of monthly payments by instrument. See Figure 9 for marginal distributions.

	Shares in 2013	Shares in 2014	Difference
Cash	26.35	25.59	-0.76
Check	8.41	7.59	-0.81
MO	0.50	0.46	-0.04
Debit	31.15	30.81	-0.34
Credit	22.54	23.26	0.72
Prepaid	1.01	1.05	0.04
OBBP	4.40	5.14	0.74
BANP	4.81	5.28	0.47
Income	0.84	0.82	-0.02
Chi-stat			14.48
p-value			0.07

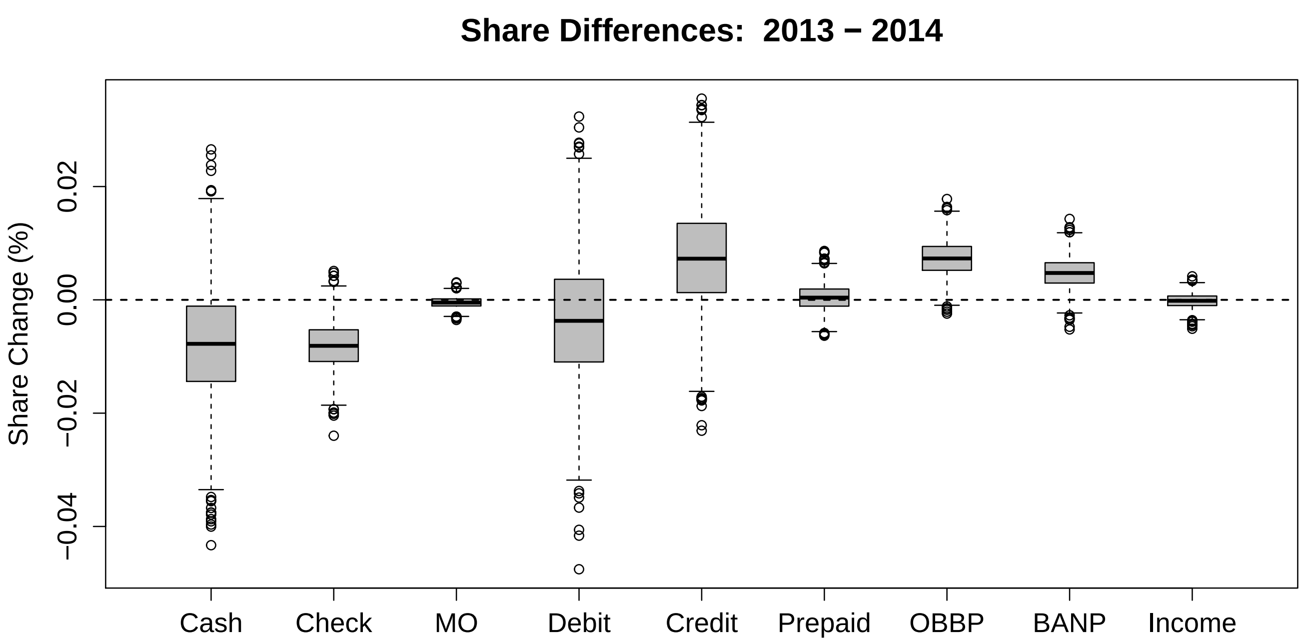

Figure 9: Share of monthly payments by instrument.

Table 23: Share of monthly payments by transaction type. See Figure 10 for marginal distributions.

	Shares in 2013	Shares in 2014	Difference
Auto. Bill	9.94	10.91	0.97
Online Bill	10.21	10.20	-0.01
Other Bill	12.61	12.26	-0.35
Online	5.71	5.44	-0.27
Retail	35.13	34.62	-0.51
Service	21.96	21.66	-0.30
P2P	4.43	4.90	0.47
Chi-stat			4.32
p-value			0.63

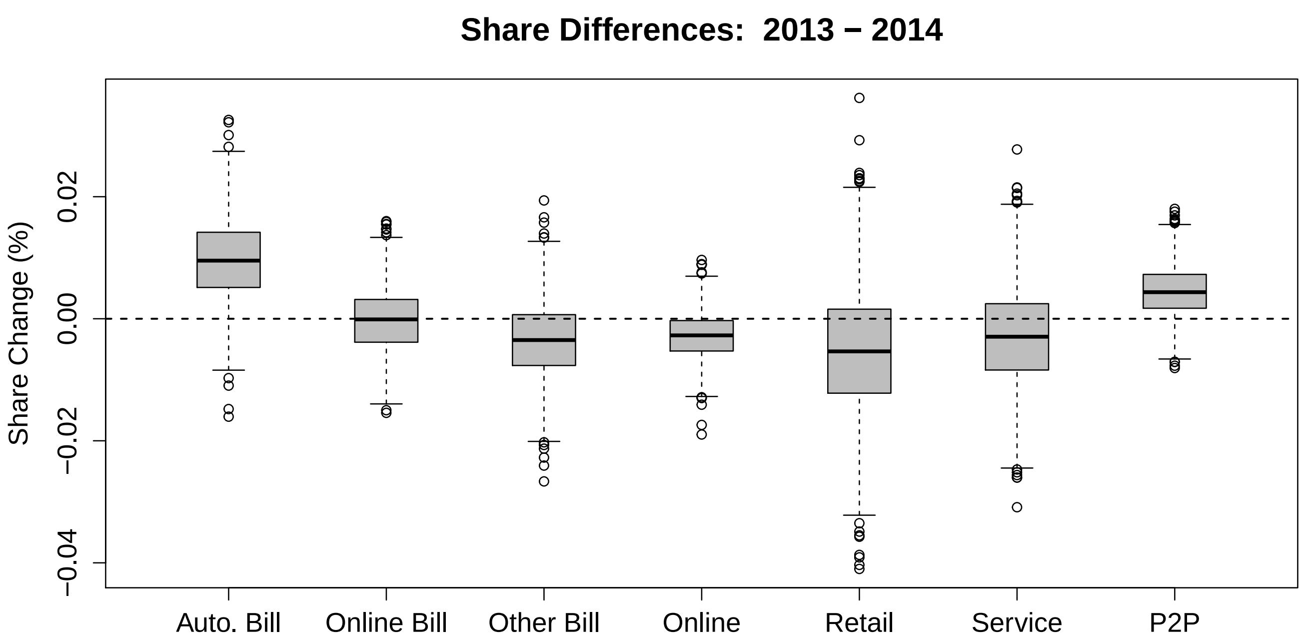

Figure 10: Share of monthly payments by transaction type.

Table 24: Share of monthly payments by instrument groups. See Figure 11 for marginal distributions.

	Shares in 2013	Shares in 2014	Difference
Paper	35.58	33.99	-1.59
Card	55.15	55.52	0.38
Electronic	9.28	10.49	1.21
Chi-stat			7.66
p-value			0.02

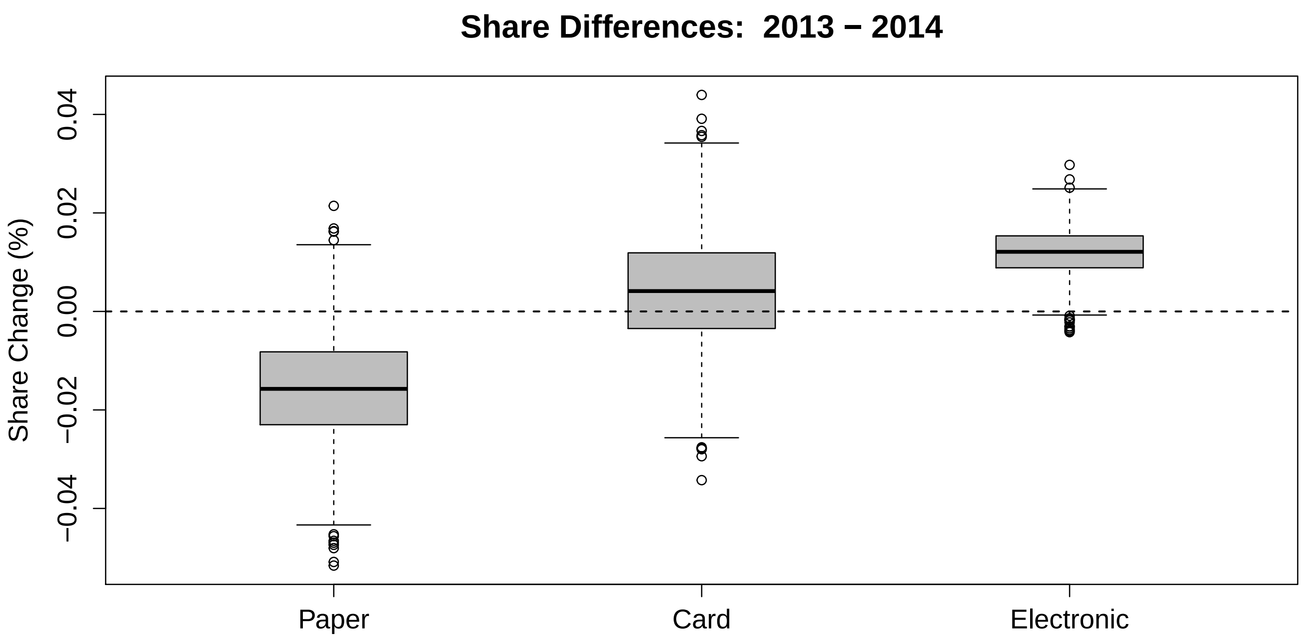

Figure 11: Share of monthly payments by instrument groups.

Table 25: Share of monthly payments by groups of transaction types. See Figure 12 for marginal distributions.

	Shares in 2013	Shares in 2014	Difference
Bill	32.76	33.38	0.62
Online	5.71	5.44	-0.27
In Person	61.53	61.18	-0.34
Chi-stat			0.79
p-value			0.67

Table 26: Share of monthly payments by asset/liability type. See Figure 13 for marginal distributions.

	Shares in 2013	Shares in 2014	Difference
Currency (cash)	26.57	25.80	-0.77
Demand deposits	49.17	49.22	0.05
Other Assets	1.53	1.52	-0.00
Liabilities	22.73	23.45	0.72
Chi-stat			0.88
p-value			0.83

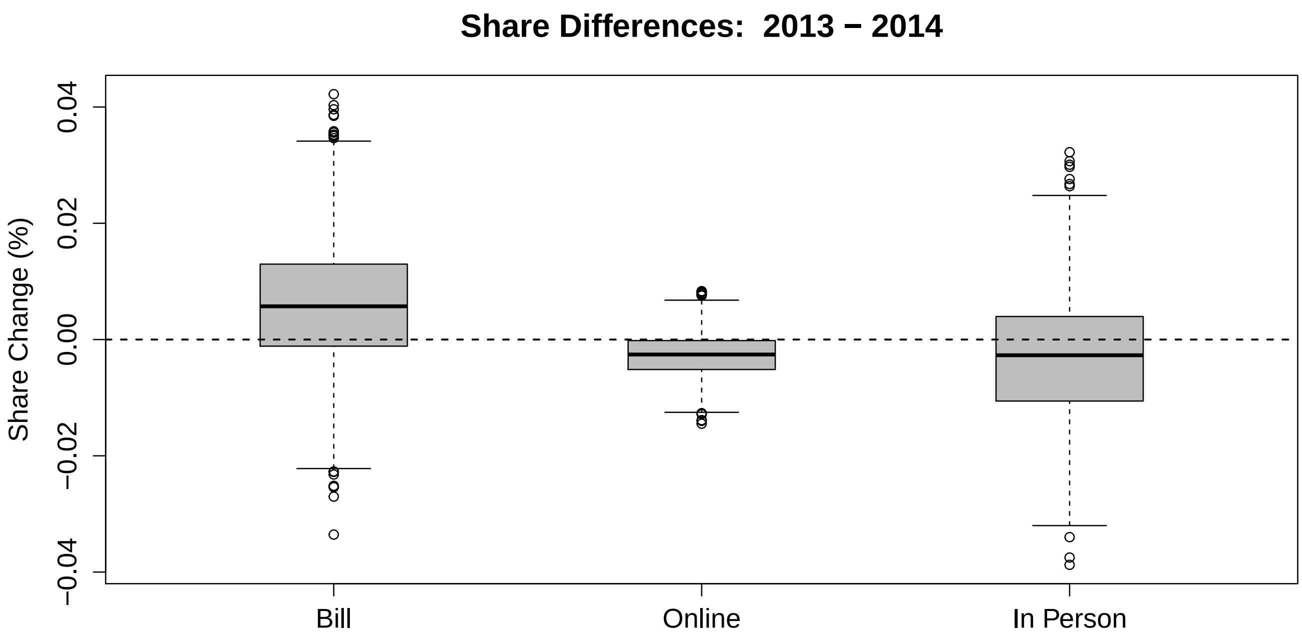

Figure 12: Share of monthly payments by groups of transaction types.

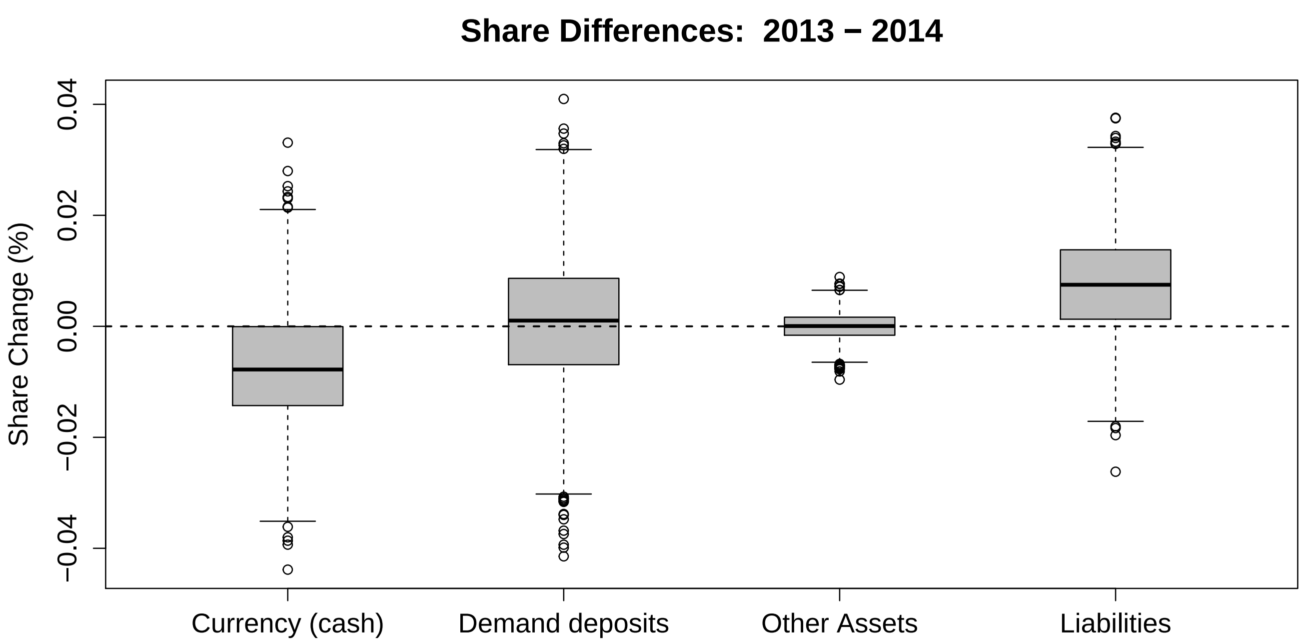

Figure 13: Share of monthly payments by groups of asset types.

9.3 Miscellaneous Variables

Table 27: Preferred method of authorization of debit cards.

	Shares in 2013	Shares in 2014	Difference
Prefer Pin	51.81	56.56	4.75
Prefer Signature	23.96	20.89	-3.07
Indifferent	24.22	22.55	-1.67
Chi-stat			9.91
p-value			0.01

Table 28: Use of cash. "Value" refers to the total dollar value of withdrawals per month, "Amount" refers to the amount withdrawn per withdrawal, and "Frequency" refers to number of monthly withdrawals. Cash holdings are excluding large value holdings (top 2 percent).

	Level in 2013	Level in 2014	Difference	z-stat	p-value
All Sources					
Value	684.63	547.99	-136.64	-2.40	0.02
Amount	124.29	119.18	-5.11	-0.92	0.36
Frequency	6.49	5.62	-0.87	-1.90	0.06
Value: Primary	558.14	450.55	-107.59	-2.25	0.02
Amount: Primary	130.33	124.67	-5.66	-0.94	0.35
Frequency: Primary	4.81	4.01	-0.80	-2.26	0.02
Value: Secondary	131.49	100.27	-31.22	-1.88	0.06
Amount: Secondary	41.48	42.45	0.97	0.25	0.80
Frequency: Secondary	1.73	1.65	-0.09	-0.52	0.60
Cash in Wallet	64.45	63.19	-1.26	-0.28	0.78
Cash in Wallet (w/out Large Values)	61.67	57.90	-3.77	-1.05	0.29
Cash in House	460.68	285.57	-175.11	-1.65	0.10
Cash in House (w/out Large Values)	61.67	57.90	-3.77	-1.05	0.29
Cash Holdings	507.62	339.92	-167.70	-1.64	0.10
Cash Holdings (w/out Large Values)	229.14	206.91	-22.23	-1.63	0.10

Table 29: Ownership rates of payment accounts.

	Level in 2013	Level in 2014	Difference	z-stat	p-value
Bank Account	0.91	0.92	0.00	0.37	0.71
Checking Account	0.90	0.91	0.00	0.26	0.80
Savings Account	0.74	0.75	0.00	0.21	0.83
Nonbank Payment Account	0.55	0.56	0.01	0.68	0.50

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.

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.

Casella, George, and Roger L. Berger. 2002. *Statistical Inference*. California: Thomson Learning.

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

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

CPS. 2010. "Current Population Survey." http://www.census.gov/cps/.

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. 2015. "SCPC User's Guide: 2014." 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.

Greene, Claire, Scott Schuh, and Joanna Stavins. 2016. "The 2014 Survey of Consumer Payment Choice: Summary Results." Federal Reserve Bank of Boston Research Data Report 16-xx.

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

Hitczenko, Marcin. 2015a. "Estimating Population Means in the 2012 Survey of Consumer Payment Choice." Federal Reserve Bank of Boston Research Data Report 15-2.

Hitczenko, Marcin. 2015b. "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.

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.

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.