Federal Reserve Bank of New York
Staff Reports

An Overview of the Survey of Consumer Expectations

Olivier Armantier
Giorgio Topa
Wilbert van der Klaauw
Basit Zafar

Staff Report No. 800
November 2016

FRBNY Staff REPORTS

An Overview of the Survey of Consumer Expectations
Olivier Armantier, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar
Federal Reserve Bank of New York Staff Reports, no. 800
November 2016
JEL classification: C81, D80, E31

Abstract

This report presents an overview of the Survey of Consumer Expectations, a new monthly online survey of a rotating panel of household heads. The survey collects timely information on consumers' expectations and decisions on a broad variety of topics, including but not limited to inflation, household finance, the labor market, and the housing market. There are three main goals of the survey: (1) measuring consumer expectations at a high frequency, (2) understanding how these expectations are formed, and (3) investigating the link between expectations and behavior. This report discusses the origins of the survey, the questionnaire design, the implementation of the survey and the sample, and computation of various statistics that are released every month. We conclude with a discussion of how the results are disseminated, and how the (micro) data may be accessed.

Key words: survey, expectations, inflation, measurement

Armantier, Topa, van der Klaauw, Zafar: Federal Reserve Bank of New York (e-mail: olivier.armantier@ny.frb.org, giorigo.topa@ny.frb.org, wilbert.vanderklaauw@ny.frb.org, basit.zafar@ny.frb.org). The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the Deutsche Bundesbank. All errors are the authors'.

1. Introduction

The importance of compiling high-quality data on the expectations held by economic agents has been increasingly recognized in both academic research and policymaking. Most economic decisions involve uncertainty, and should therefore be determined not only by preferences but also by expectations about the future. Expectations should drive a variety of households' economic choices, including those related to saving, investment, purchases of durable goods, wage negotiations, etc. The aggregation of these choices in turn determines macroeconomic outcomes, including realized inflation, in equilibrium. Given the role of households in aggregate as an important driver of economic activity, monitoring and managing consumers' expectations have become primary goals of policy makers, and are central components of modern monetary policy (Woodford, 2004; Bernanke, 2007; Gali, 2008; Sims, 2009).[1] The effectiveness of monetary policy and central bank communication relies on longer-run inflation expectations being well-anchored, making their measurement important for policymakers. More generally, expectations by consumers about a number of personal and macroeconomic outcomes are increasingly useful inputs into a variety of forecasting models.

Effective monitoring and managing of expectations requires measuring consumer expectations and understanding of how these expectations are formed. The New York Fed's Survey of Consumer Expectations (SCE) was developed with precisely these goals in mind. It collects timely, high frequency information on consumer expectations and decisions on a broad variety of topics. Its overall goal is to fill gaps in existing data sources (such as the Reuters/University of Michigan Survey of Consumers, the Federal Reserve Board's Survey of Consumer Finances, and the Bureau of Labor Statistics' Consumer Expenditure Survey) on household expectations and behavior, providing a more integrated data approach. The SCE aims to cover a broad range of economic outcomes, including inflation, household finance, the labor market, the housing market, as well as special topics as the need arises for policy or research analysis.[2]

The SCE is designed as a rotating panel, which enables researchers and policy-makers to follow the same individuals over time, reducing sample composition changes and thus sampling volatility in survey responses from month to month. The panel structure of the SCE is intended to provide both valuable input into the evaluation of the economic outlook and policy formulation, and a valuable resource for the research community to understand better how

[1] In particular, Bernanke (2004) argued that "an essential prerequisite to controlling inflation is to control inflation expectations."

[2] National surveys of public (inflation) expectations are now conducted in multiple countries. These include the Reuters/University of Michigan Survey of Consumers, the Livingston Survey, the Conference Board's Consumer Confidence Survey and the Survey of Professional Forecasters in the US. Other central banks that survey consumers about their inflation expectations include the Bank of England, the European Central Bank, the Bank of Japan, the Reserve Bank of India, and the Sveriges Riksbank. Since 2015, the Bank of Canada has implemented a largely comparable version of the SCE, fielded at a quarterly frequency. Since 2013, the Federal Reserve Board has conducted the Survey of Household Economics and Decision making (SHED), which elicit some expectations about the economic well-being of U.S. households.

consumers form and update their expectations, as well as the links between expectations and behavior. For instance, the data allow one to study how expectations about house prices and interest rates affect consumers' choices regarding buying or renting a home, or regarding the type of mortgage used to purchase a home. Data on expectations about the likelihood of finding a job and about future wage earnings may be used to analyze workers' job search behavior or retirement decisions. Researchers may also study how inflation expectations shape consumers' spending and saving behavior. Collecting such data for the same households over time enables researchers to study potential interactions between decisions and expectations across many different domains of consumer behavior. Finally, the data can be used to study the evolution of a diverse, but related set of expectations, the way they co-vary over time at the individual level, and identify any (structural) breaks in this relationship.

A key feature of the SCE is its reliance on a probabilistic question format to elicit the likelihood respondents assign to different future events. In addition to questions asking respondents for *point forecasts* – for example, in the case of year-ahead inflation, "*What do you expect the rate of [inflation/deflation] to be over the next 12 months?*" – for several continuous outcomes, we ask for *density forecasts*. That is, the percent chance the respondent assigns to different future possible values of that variable. In the case of future inflation, for example, respondents are asked for the likelihood that future inflation will fall within different pre-specified intervals. These density forecasts allow respondents to express uncertainty regarding their expectations. For binary outcomes (where the event either occurs or not), eliciting the percent chance associated with the event fully identifies the underlying subjective distribution. By doing so, the SCE extends a practice with a longer tradition in the field of psychology and in surveys of professional forecasters, economists, and other financial experts. Our approach also builds on a large and growing body of economic research, led by Charles Manski, that has demonstrated survey respondents' willingness and ability to answer questions expressed in this way.

Finally, the SCE is implemented as a monthly internet survey in order to provide more flexibility for question design and more real-time capabilities for data collection. An internet platform enables the researcher to design more user-friendly questions, with the help of visual aids and other tools that make it easier for respondents to understand and answer a specific question. An internet platform also makes it easier to develop and field new questions on special topics at a short notice. For example, we designed and fielded special surveys to help assess consumer responses in spending and inflation expectations to sudden large gas price declines, and to elicit beliefs regarding the early impact of the Affordable Care Act on future health care spending, prices and coverage.

1.1. Survey Overview

The SCE started in June 2013, after a six-month initial testing phase.[3] It is a nationally representative, internet-based survey of a rotating panel of about 1,300 household heads, where a household head is defined as the person in the household who owns, is buying or rents the home. The survey is conducted at a monthly frequency. New respondents are drawn each month to match various demographic targets from the American Community Survey (ACS), and stay on the panel for up to twelve months before rotating out of the panel. The survey instrument is fielded on an internet platform designed by *The Demand Institute*, a non-profit organization jointly operated by *The Conference Board* and *Nielsen*. The respondents for the SCE come from the sample of respondents to the Consumer Confidence Survey (CCS), a mail survey conducted by The Conference Board. In turn, the respondents for the CCS are selected using from the universe of U.S. Postal Service addresses. From that universe, a new random sample is drawn each month, stratified only by Census division.

The SCE has various components. First, it includes a core monthly module on expectations about various macroeconomic and household level variables.[4] Respondents are asked about their inflation expectations, as well as expectations regarding changes in home prices and in prices of various specific spending items, such as gasoline, food, rent, medical care and college education. The core survey also asks for expectations about unemployment, interest rates, the stock market, credit availability, taxes and government debt. In addition, respondents are asked to report their expectations about several labor market outcomes that pertain to them, including changes in their earnings, the perceived probability of losing their current job (or leaving their job voluntarily), and the perceived probability of finding a job. Finally, respondents are asked about the expected change in their household's overall income and spending. As described in more detail below, these expectations questions are fielded at various time horizons and with various formats, including both point and density forecasts.

Second, the SCE contains a supplementary "ad-hoc" module each month on special topics. Three such modules are repeated every four months, leaving three "floating" supplements per year on topics that are determined as the need arises. The three repeating supplements are on credit access, labor market, and spending. Topics covered so far in the "floating" supplement include (but are not limited to) the Affordable Care Act, student loans, workplace benefits such as childcare and family leave, and use of insurance products. Together, the core monthly module and the monthly supplement take about 15 minutes to complete.

Finally, SCE respondents also fill out longer surveys (up to 30 minutes in length, and separately from the monthly survey) each quarter on various topics. Most of these surveys are repeated at a yearly frequency. Since each SCE panelist stays in the panel for up to 12 months, these annual

---

[3] As discussed below, the SCE was preceded by an extensive feasibility study over the period 2006-2013, using experimental surveys on the RAND Corporation's American Life Panel. Early findings from that study are described in van der Klaauw et al. (2008).

[4] Each entering cohort is also administered a module with demographic questions about the respondent and their household.

surveys can be used as independent repeated cross-sections, although they obviously can be linked to the monthly core survey panel responses. The SCE currently contains quarterly surveys on the housing market; the labor market; informal work participation; consumption, saving, and assets. A subset of these surveys is designed in part or wholly by other Federal Reserve Banks.

2. Questionnaire Design

The 1990s represented a period of significant change in the way economists elicit expectations through surveys. Traditionally researchers measured expectations through verbal and qualitative questions, asking respondents whether they expect that an event will occur or not, or asking whether they think it is 'very likely,' 'fairly likely,' 'not too likely,' or 'not at all likely' that a specific event will occur. In addition to the limited information captured due to the coarseness of choice options or due to an inability to express uncertainty altogether, a major drawback of this traditional approach concerns the lack of inter- and intra-personal comparability of responses.

Led and inspired by the work of Charles Manski, himself building on the early work of Juster (1966) and a longer tradition of collecting such data in cognitive psychology, economists began to elicit probabilistic expectations during the 1990s. It quickly became clear that, with some guidance, survey respondents were able and willing to answer probabilistic expectations questions.[5] Moreover, with a fixed numerical scale, responses are interpersonally comparable, and have been found to be better predictors of outcomes. A growing number of large scale surveys in the U.S. and abroad now use probabilistic formats to elicit expectations for a wide range of events, and these formats have also been successfully implemented in surveys associated with laboratory and field experiments, including several conducted in less developed countries. Manski (2004) reviews research eliciting probabilistic expectations in surveys and assesses the state of the art at that time. Delavande (2011, 2014) provides a more recent review.

The questionnaire design of the SCE builds on this previous research and was informed, in large part, by our experiences with the Household Inflation Expectations Project (HIEP), wherein we fielded surveys every six weeks on RAND's American Life Panel (ALP) going back to 2007 (see Bruine de Bruin et al., 2010a). The HIEP questionnaires were developed in collaboration with RAND, other Federal Reserve Banks, academic economists, psychologists and survey design experts. During the 2006-2012 period, the HIEP conducted in-depth cognitive interviews, fielded psychometric surveys on the ALP as well as on various "convenience" (i.e. non-random) samples, and administered a number of experimental consumer surveys in the ALP. In addition to testing probabilistic question formats, we also experimented with alternative question wordings, especially for questions related to inflation (discussed in more detail below). The

[5] Usually before asking such questions respondents are provided a brief introduction or explanation of basic probabilistic ideas through examples.

findings from this project formed the basis for the creation of the SCE (see van der Klaauw et al., 2008; Armantier et al. 2013).

During the SCE's experimental or development phase, between December 2012 and June 2013, we further sharpened and tested the questionnaire. The process involved conducting additional cognitive interviews with a small auxiliary sample to identify potential interpretations of the questions. After reading the questions out loud, interviewees were instructed to think out loud while generating their answers. This allowed us to gauge if in fact interviewees interpreted the questions the way we had intended them to. When necessary, we modified the wording of questions accordingly. Pilot surveys were also conducted and the data were analyzed to make sure the questions were eliciting meaningful variation.

Of particular importance were our findings in the SCE questionnaire development phase regarding the elicitation of expectations. In both the HIEP and SCE experimental phase we tested a large set of probabilistic questions. Respondents in our surveys showed a consistent ability and willingness to assign a probability (or "percent chance") to future events. Unlike simple point forecasts, probabilistic expectations of binary outcomes as well as density forecasts for continuous outcomes provide a valuable measure of individual uncertainty. Moreover, we found their responses to be largely internally consistent in terms of simple laws of probabilities. Finally, we conducted an experiment which confirmed that the densities elicited are informative (Armantier et al. 2013). For a review of this work see van der Klaauw et al. (2008), and Bruine de Bruin et al. (2010a, 2011b).

Having discussed the SCE questionnaire development process, we now turn to the actual questions, focusing on the SCE core survey. The quantitative questions can be broadly divided into three categories: (1) questions that elicit expectations of binary outcomes (such as the likelihood of the US stock market being higher in 12 months); (2) questions that elicit pointwise expectations for continuous outcomes (such as the rate of inflation over the next 12 months); and (3) questions that elicit respondents' probability densities for forecasts of continuous outcomes. Besides these questions, the survey also includes some qualitative questions, including questions in which respondents are asked to answer using a (e.g.) 7 points rating scale. The full questionnaire is available here.

The bulk of the survey elicits near-term expectations, that is, expectations regarding outcomes over either the next 12 months and/or over the next 3 months. The ad-hoc supplements also contain questions at the four month horizon. However, in certain instances, such as inflation or home prices, expectations are also elicited for the medium-term, that is, three years out.

We next describe the rationale for how we elicit some of these expectations.

2.1. Expectation Elicitation of Continuous Outcomes

For some continuous outcomes, we elicit both point and density forecasts. We begin by illustrating the format of our point forecast questions, and describe the density forecast questions in section 2.3 below.

Inflation Expectations Inflation expectations are elicited using the following two-stage format. Respondents are first asked: "*Over the next 12 months, do you think that there will be inflation or deflation? (Note: deflation is the opposite of inflation)*". Depending on their response to this, they are next asked for a point estimate: "*What do you expect the rate of [inflation/deflation] to be over the next 12 months? Please give your best guess*".

Note that we directly ask respondents for the rate of "inflation". This differs from the widely used approach of avoiding the term "inflation" in consumer surveys. Most existing inflation expectations questions, such as those posed on the Michigan Survey (Curtin, 2006), ask for expected changes in "prices" as follows: "*During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*" followed by the response options "Go up," "Stay the same," and "Go down". Those who respond "go up" or "go down" are then asked to give a specific point estimate.

We prefer asking respondents for "inflation" because our prior research in the HIEP suggests that the way the Michigan Survey question is phrased induces mixed interpretations, with some respondents thinking about specific prices they pay and others thinking about the overall rate of inflation. As the former tend to think more of salient price changes, they are more likely to provide extreme responses (Bruine de Bruin et al., 2011a). We found that asking about the rate of inflation directly reduces ambiguity in question interpretation and yields, we believe, more reliable, more inter-personally comparable responses. Moreover, it is consistent with the concept of forward inflation expectations of interest to central banks.

One source of hesitation in asking consumers about inflation may be the concern that this is a relatively complex concept. However, evidence suggests that consumers tend to have a basic understanding of what it means (Leiser and Drori, 2005; Svenson and Nilsson, 1986). Cognitive interviews conducted during the HIEP similarly indicate that the vast majority of consumers have a good understanding of the concept of inflation (van der Klaauw et al. 2008; Bruine de Bruin et al. 2012). Furthermore, our research indicates that consumers act on their reported inflation expectations in sensible ways (Armantier et al., 2015) and that they update their inflation expectations meaningfully when provided with arguably inflation-relevant information (Armantier et al. 2016).

In addition, during the period May 2013- September 2015, we asked SCE survey respondents the following question: "*On a scale of 1 to 7, how well would you say you understand what "inflation" means? (where 1 means "I don't know what "inflation" means", and 7 is labelled as "I know exactly what "inflation" means")*". Of the 5,182 first-time respondents to whom this question was fielded, only 44 (less than 0.9%) choose 1 (no understanding on the scale). In fact, 82% of respondents choose 5 or higher, suggestive of inflation being a fairly well-understood concept. This of course does not mean that households find it easy to express inflation in quantitative terms. To investigate this, we asked our survey respondents the following: "*On a scale of 1 to 7, how easy is it for you to express the rate of inflation as a number? (where 1 is "Very easy" and 7 is "Very difficult")*". Of the 5,179 respondents, 82% reported 5 or lower, with just 5.6% choosing 7 on the scale. Overall, this is pretty convincing evidence that the vast majority of consumers understands the concept of inflation and is able to express it numerically.

Expectations for Other Continuous Outcomes We next present the wording for one other representative question. Earnings expectations, for those employed, are asked as follows:

Please think ahead to 12 months from now. Suppose that you are working in the exact same job at the same place you currently work, and working the exact same number of hours. What do you expect to have happened to your earnings on this job, before taxes and deductions?

- *increase by 0% or more*
- *decrease by 0% or more*

This is then followed by: "*By about what percent do you expect your earnings to have [increased/ decreased]?*"

Note that respondents are not presented with a "stay the same" option when they are asked for a directional change. Instead the instructions specify that they can enter a zero change by either picking the increase or decrease option. When experimenting with question wording, we found that a substantial proportion of respondents would choose "stay the same" when that was presented as one of the options. Upon probing them further, we found that a substantial proportion of them had changes in mind that were bigger than 0.5% in magnitude, suggesting that the mere availability of choosing the "stay the same" option was leading several respondents to pick it. In our analysis, we found that the distribution of responses to a given question with and without the "stay the same" option to be noticeably different. More research is clearly needed to understand which elicitation method is best at recovering the true underlying subjective belief. However, given that respondents can always enter zero and hence we arguably elicit more information without the "stay the same" option, we chose to ask for changes without the "stay the same" option. In addition to the questions asking for point forecasts for inflation and earnings which we just discussed, we use this same question format for a wide range of other outcomes.

2.2. Elicitation of Expectations for Binary Outcomes

The survey includes several probabilistic questions that elicit the likelihood (or percent chance) of a certain event. These questions were preceded with an introduction on the use of percentages: "*In some of the following questions, we will ask you to think about the percent chance of something happening in the future. Your answers can range from 0 to 100, where 0 means there* *is absolutely no chance, and 100 means that it is absolutely certain. For example, numbers like: 2 and 5 percent may indicate "almost no chance"; 18 percent or so may mean "not much chance"; 47 or 52 percent chance may be a "pretty even chance"; 83 percent or so may mean a "very good chance"; 95 or 98 percent chance may be "almost certain".*"

For example, those unemployed and actively looking for work are asked: "*What do you think is the percent chance that within the coming 12 months, you will find a job that you will accept, considering the pay and type of work?*"

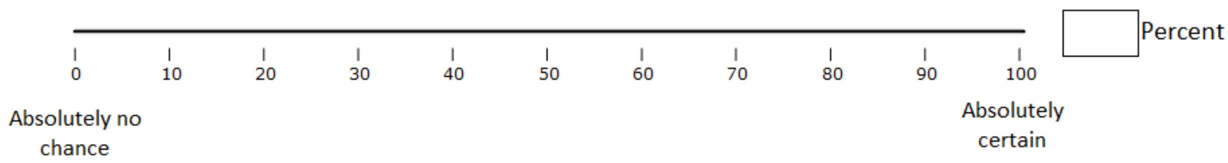

Respondents could either enter a number (on 0-100 scale) directly into the box, or click anywhere on the sliding scale. To prevent respondents from anchoring their response, no marker appears on the scale until the respondent clicks somewhere on it.

2.3. Elicitation of Forecast Densities

Relative to existing consumer surveys of expectations, one of our innovations is that we also elicit consumers' subjective probability distribution for certain continuous outcomes, such as future inflation, earnings, and home prices. These density data allow us to construct individual measures of central tendency (e.g., the density mean or median), uncertainty, and perceived tail risks (such as the probability of extreme positive or negative outcomes).

Our density questions follow a format similar to that of the Survey of Professional Forecasters and the Bank of Italy's Survey of Household Income and Wealth. Respondents are presented with various pre-defined non-overlapping bins that exhaust the whole range of values that the random variable may take, and are then asked for the percent chance that the variable would take values in each of those intervals, with the reminder that numbers need to add up to 100%. The density forecast for year-ahead national home price changes, for example, is elicited as follows:

And in your view, what would you say is the percent chance that, ***over the next 12 months,*** *the average home price nationwide will...*

increase by 12% or more	________________	*percent chance*
increase by 8% to 12%	________________	*percent chance*
increase by 4% to 8%	________________	*percent chance*
increase by 2% to 4%	________________	*percent chance*
increase by 0% to 2%	________________	*percent chance*
decrease by 0% to 2%	________________	*percent chance*
decrease by 2% to 4%	________________	*percent chance*
decrease by 4% to 8%	________________	*percent chance*
decrease by 8% to 12%	________________	*percent chance*
decrease by 12% or more	________________	*percent chance*
Total	***XXX***	

As respondents enter their answers, they can see what the total sums to. Respondents who nevertheless give answers that do not add up to 100% receive the notice "*Please change the numbers in the table so they add up to 100.*"

We use each individual's responses to the probabilistic questions to parametrically estimate the underlying forecast density function, following Engelberg, Manski and Williams (2009). We describe this estimation in more detail in section 5.2 below. Based on the probability density function for each respondent, we compute corresponding density means and medians. Further, we use the density Inter-Quartile Range (IQR), the difference between the third and the first quartiles, as a measure of individual forecast uncertainty. We choose this measure because the IQR is less sensitive than, say, the standard deviation to small variations in the tails of the estimated density.

3. Implementation

3.1. Sample Design

The SCE sample design is based on that of The Conference Board's monthly *Consumer Confidence Survey (CCS)*. The CCS is a mail survey that uses an address-based probability sample design to select a new random sample each month based on the universe of U.S. Postal Service addresses.[6] The latter is derived from the files created by the U.S. Postal Service and represents near-universal coverage of all residential households in the United States. It is updated monthly to ensure up-to-date coverage of U.S. households. The targeted responding CCS sample size is approximately 3,000 completed questionnaires each month from household heads. Questionnaire instructions define the latter as "*the person in your household who owns, is buying or rents this home*".[7]

The SCE sampling frame (or sampling population) consists of CCS respondents who expressed ability and a willingness to participate in the SCE based on their answers to two questions included at the end of the CCS questionnaire. The first asks "*Do you have access to the Internet and an email address?*" Those who answer "YES" are then asked the question:

---

[6] The CCS random sample of household addresses is drawn after first stratifying geographically by 9 census division to provide a proportionate geographic distribution. To ensure proportional representation in the sample of respondents, the CCS uses weights based on gender, income, geography, and age.

[7] This definition is similar to that used in the Current Population Survey (CPS), in the ACS, and by Census more generally: there, the "reference person" in the household (or the "householder") is the person who owns or rents the unit of residence. Note that the instructions state that "*if that person is not available or unable, please have an adult aged 18 or older who lives in this household complete the survey*". In a representative month (March 2013), 96% (95% weighted) of CCS respondents are household heads, with very little variation across gender, age, income and race/ethnicity. Note also that this definition does not exclude the possibility that a household may have multiple "co-household heads."

"You may be eligible to participate in a survey about your perceptions of the economy, employment, finances, and related topics. This is a paid survey that would be conducted monthly for up to 12 months. You would receive $15 for each completed survey.[8] *If selected, we would email you a web address where you could respond to the online questionnaire.*

> *Would you be interested in participating in this monthly, paid survey?*
>
> *Yes, my email address is:* ____________________________________
> *No, I am not interested in participating*

As discussed in more detail below, on average 53% of CCS respondents in a given month express willingness to participate in a new online survey.[9] Of those interested in participating, approximately 300 to 320 are invited within the following two months to join the SCE internet panel, of which about 150 to 180 actually end up joining. A stratified random sampling approach is used to draw new CCS respondents into the SCE, with strata based on income, gender, age, race/ethnicity, and census division,[10] and weights chosen to maximize representativeness of the SCE panel.

3.2. Data Collection

The goal of the survey is to capture consumers' expectations over a given month. To do so, the survey is sent to respondents in three batches throughout the month. Specifically, each month, the pool of respondents is partitioned into three batches of roughly equal sizes. In general, the first, second and third batch receives an email invitation to fill out the survey on the 2nd, 11th and 20th of the month, respectively. On occasion, this schedule is amended by a day or two to reflect holidays or shorter months (i.e., February). If they have not yet completed the survey, respondents in each batch receive two reminders by email, 3 and 7 days after their initial invitation. On rare occasions, a third reminder is sent to the first and second batch on an ad hoc basis (e.g., if the response rate is perceived to be lower than usual). Survey responses for all three batches are collected until the last day of the month.

In 2014, the median respondent in each batch completed the survey 3 days after receiving her invitation. In Figure 1, we plot the number of surveys completed each day of a typical month (April 2016).[11] Although not uniformly distributed, the completion of surveys is spread out throughout the month, with three major peaks on the days after each batch receives its invitation to fill out the survey, and smaller peaks on the days on or after which each batch receives a reminder to fill out the survey.

---

[8] As explained below, before July 2013, some respondents received a letter with a different amount.
[9] The monthly average of 53% was based on CCS responses during the December 2012-September 2015 period.
[10] We distinguish between 8 household income groups, 5 age groups, 5 race and ethnicity groups, and 9 census divisions.
[11] Figure A1 in Appendix A shows the median and 25th and 75th percentiles of the daily frequency of responses over all months from December 2012 to September 2015, thus combining months with different dates for the invitations and reminders.

Each month, the panel of household heads invited to answer the survey consists of roughly 300 new respondents and 1,100 "repeat" respondents (i.e., respondents who have completed at least one survey within the past eleven months). The new respondents invited to answer the survey for the first time are randomly allocated to one of the three batches. A few days before they are to receive the invitation to fill out the survey, the new respondents are first contacted by mail and by email to welcome them into the panel. These letters inform the respondent about the nature, the number, the duration and the timing of the surveys they will be asked to complete over the next 12 months.[12] The new respondents are also told about the payment they will receive for each survey completed, and they are given access to a website where they can find additional information and ask questions to the help desk.

Figure 1: Number of Surveys Completed over April 2016

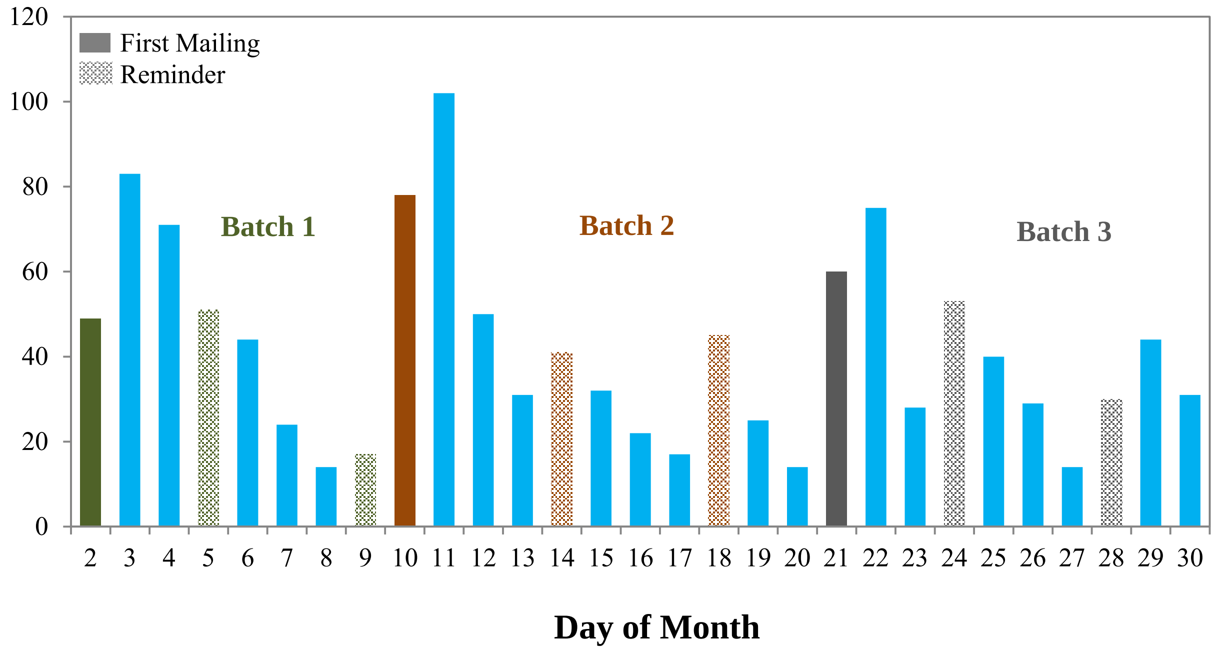

Note: The full bars in green, brown and grey represent the day respondents from (respectively) batch1, 2 and 3 are invited to fill out the survey. The shaded bars in green, brown and grey represent the day respondents from (respectively) batch1, 2 and 3 receive a reminder to complete the survey.

At the beginning of each month, repeat respondents (i.e., respondents who have already completed at least one survey in the past) are partitioned in two groups: the "skippers" (i.e., those who failed to complete the survey in the previous month) and the "non-skippers." The wide majority of repeat respondents are non-skippers (93% in 2014). Skippers are assigned randomly

[12] We experimented with sending a welcome email only (and no mail) the new respondents. However, that led to a noticeable decline in the response rate, suggesting that the welcome mail lent greater credibility to the survey. Thus, we reverted back to new respondents receiving both a welcome mail and email.

to one of the three batches. The assignment procedure for non-skippers is designed so that i) there is an equal number of non-skippers in each batch, and ii) non-skippers in each batch have (roughly) the same average number of days between the completion of two consecutive surveys. On the first of each month, non-skippers are ranked according to the number of days since they completed the survey in the previous month and partitioned into terciles. The first tercile (i.e., the respondents who completed the survey most recently) is assigned to batch 3, the second to batch 2, and the third to batch 1.[13]

Any respondent invited after July 2013 has been paid $15 for each monthly survey completed. We settled on this amount after testing whether the amount paid for each completed survey impacts the response rate. Specifically, we had three groups of respondents between December 2012 and July 2013. During their 12-month tenure, each group was randomly assigned to be paid $10, $15 or $20 for each survey completed. The response rate in the first month was 61%, 66% and 56% in the $10, $15 and $20 group, respectively. Further, 28%, 37% and 32% of the respondents in the $10, $15 and $20 group (respectively) completed all 12 surveys.[14] Thus, we concluded that a payment of $15 per survey was the most cost-effective.

Respondents can be removed from the panel if they fail to respond to the monthly surveys they are invited to fill out. This is the case in particular for respondents who do not complete the first survey they are invited to fill out. Otherwise, if a respondent does not complete the monthly survey in three consecutive months, then the respondent is dropped from the panel and no longer invited to fill out any additional surveys. Twelve months after their first survey completion, every respondent is rotated out of the panel.

We now turn to the issue of survey participation. Most of the non-response occurs in the first month. Out of the 3,582 household heads we invited to participate in the survey in 2014, 1,647 (or 46%) failed to complete the first survey and were therefore not re-invited. Once in the panel, however, attrition drops rapidly. Indeed, we can see in Figure 2, that while 26% of first time respondents fail to complete a second survey, the response rate after the second month is essentially flat. In particular, observe in Figure 2 that 58% of the respondents who entered the panel in 2014 completed all 12 surveys

[13] Prior to February 2016, the allocation procedure for non-skippers was also applied to skippers. As a result, skippers were found predominantly in batch 1 (because skippers completed their last survey more than 30 days ago). Because skippers may have specific unobserved characteristics, we were concerned that the response rate and the survey responses from batch 1 would be different from those of batch 2 and 3. Thus, we decided to allocate skippers randomly across the three batches.

[14] The lower response rate for the $20 group may be due to the fact that these respondents were pulled from an older CCS sample.

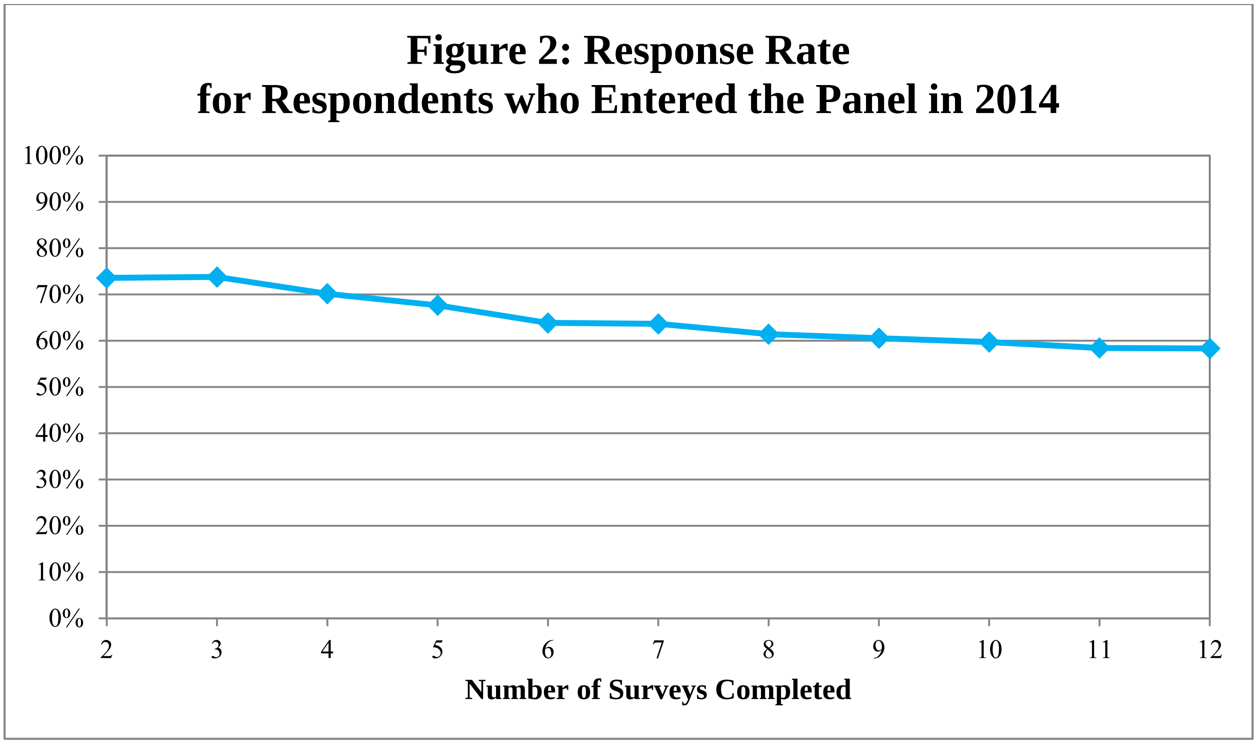

4. Panel Representativeness

The representativeness of our panel of respondents depends on a number of factors including the composition of: (i) the sample of CCS respondents who reported having access to the internet and email and who are willing to participate in our survey; (ii) the sample of invited and interested CCS respondents who actually chose to enter our panel by completing their first SCE survey; and (iii) the sample of SCE participants who continue to participate in our panel after entry.

As discussed earlier, the CCS target population is the U.S. population of household heads, defined as the person who owns, is buying or rents the home. As shown in Table A1 of Appendix A, average characteristics of household heads who participated in the CCS during October 2013-September 2015 are largely comparable to those in the 2013 and 2014 American Community Surveys (ACS). The main difference in sample composition between the CCS and ACS concerns the age distribution, with younger (older) household heads being somewhat underrepresented (overrepresented) in the CCS -- a common feature of mail surveys.

The SCE sampling frame consists of CCS respondents who reported having access to the internet and email, and expressed a willingness to join a new online survey. Columns 1, 2 and 3 of Table 1 report characteristics of respectively, all CCS respondents, CCS respondents who reported to have access to the internet and an email address, and the subset of those who indicated an interest in joining an online panel survey, during the October 2013-September 2015 period. As shown in the table, relative to CCS respondents overall, those with internet access were somewhat younger, with especially those over age 60 underrepresented, and marginally more likely to be male and white or Asian. Those with internet access also were more likely to have family incomes exceeding $50,000, more likely to have young children, to have slightly higher household incomes, and to reside in the West. Those who expressed an interest in joining an online survey had very similar average characteristics as the CCS respondents with internet access, but compared to CCS respondents overall were even more likely to be younger and to have a child under age 12 in the household.

Instead of demographic characteristics, Table 2 shows average responses to the standard set of CCS consumer sentiment questions. While differences are generally remarkably small, compared to CCS respondents overall, those with internet access and email on average are slightly more positive and optimistic about current and future business conditions, job availability and income, and expect slightly lower inflation. We find the same pattern for those interested in joining an online survey, except that the differences are slightly larger in magnitude. Overall, in terms of consumer sentiment, we find the pool of interested CCS-respondents to be quite similar to CCS respondents overall.

Table 1: Sample Comparisons – CCS and SCE Survey Respondents

	Full CCS sample (N=64133)	CCS respondents with internet and email (N=50089)	CCS respondents who consented (N=26439)	SCE respondents (N=3853)
	(1) (in %)	(2) (in %)	(3) (in %)	(4) (in %)
Age <30	3.7	4.2	5.9	11.7
Age 30-39	11.3	13.2	17.0	19.0
Age 40-49	15.7	17.7	20.1	18.8
Age 50-59	23.9	25.3	25.0	20.6
Age 60+	45.5	39.5	31.9	29.9
Female	47.7	47.0	47.9	48.1
Income <$15,000	8.3	4.7	5.6	8.5
$15,000-$24,999	9.9	7.0	7.3	11.3
$25,000-$34,999	10.3	8.7	8.6	9.9
$35,000-$49,999	15.2	14.7	13.8	13.1
$50,000-$74,999	19.7	21.3	20.6	21.0
$75,000-$99,999	13.1	15.3	15.2	13.5
$100,000-$124,999	9.4	11.1	11.6	7.3
$125,000 or more	14.0	17.2	17.2	15.4
New England	5.0	5.2	4.7	4.3
Middle Atlantic	13.6	13.6	13.4	12.9
EN Central	17.6	17.1	17.3	14.4
WN Central	7.7	7.5	7.1	7.6
South Atlantic	19.9	20.1	20.8	20.4
ES Central	5.8	5.2	4.9	5.1
WS Central	9.2	9.3	9.2	11.4
Mountain	6.9	7.3	7.1	8.8
Pacific	14.5	15.2	15.4	15.1
Northeast	18.5	18.5	17.9	17.2
Midwest	25.2	24.6	24.6	22.0
South	34.9	34.5	35.2	36.9
West	21.4	22.3	22.5	23.9
Mean Household Size	2.4	2.5	2.6	2.5
Any Child Under 12	16.4	18.6	23.0	23.0
Asian	3.3	3.6	3.5	3.5
Black	9.1	8.2	9.7	10.4
White	82.1	83.5	82.1	81.8
Other	4.1	4.1	4.9	4.4
Has Internet	78.1	100.0	100.0	100.0
Is Interested	41.2	52.7	100.0	100.0

Each number in the table is the percentage of the sample that falls in that category.

Table 2: Sample Comparisons – CCS and SCE Survey Respondents

	Full CCS sample (*N*=64133)	CCS respondents with internet (*N*=50089)	CCS respondents who consented (*N*=26439)	SCE respondents (*N*=3853)
	(1)	(2)	(3)	(4)
	(in %)	(in %)	(in %)	(in %)
General business conditions in Area				
Good	23.2	24.8	25.1	24.9
Normal	54.5	54.2	53.7	53.7
Bad	21.9	20.5	20.7	21.0
General business conditions in area in 6 months*				
Better	17.2	18.3	19.8	19.4
Same	71.1	70.6	69.1	70.0
Worse	11.2	10.7	11.0	10.4
Job availability in area*				
Plenty	16.3	18.0	18.7	19.1
Not so many	54.8	55.4	53.9	53.6
Hard to get	27.9	25.6	26.6	26.6
Job availability in area in 6 months*				
More	15.4	16.1	17.1	17.2
Same	67.2	67.8	66.1	67.1
Fewer	16.5	15.6	16.4	15.4
Family income in 6 months*				
Higher	14.3	16.6	19.9	22.3
Same	73.5	72.2	68.6	67.1
Lower	11.8	10.8	11.2	10.5
Increase in prices over next 12 months*				
2% or lower	21.1	22.21	22.7	23.8
3-4%	31.0	32.4	31.7	31.7
5-6%	21.5	21.5	21.8	21.8
7% or more	25.6	23.4	23.2	22.4
Expected change in interest rates**				
Mean (in %)	3.7	3.8	3.8	3.8
Expected change in stock prices**				
Mean (in %)	3.1	3.1	3.1	3.1

Each number in the table is the percentage of the sample that falls in that category.
*: remainder category is the small proportion of missing/invalid responses.
**: Averages for responses based on Likert scale, ranging from 1 (increase) to 5 (decrease). All statistics based on CCS surveys from October 2013 to September 2015.

Turning now to the SCE sample, as discussed earlier in drawing a sample of new panel members each month from among those who expressed an interest in joining an online panel, we use a stratified sampling procedure that attempts to account for differential SCE survey participation and attrition rates across different demographic groups, in terms of income, gender, age, race/ethnicity, and census division.[15]

Of those SCE volunteers newly invited to the SCE, on average 53% actually participate, with this proportion ranging between 48% and 60% during the October 2013-September 2015 period.[16] As shown in the fourth column of Tables 1 and 2, CCS respondents that end up participating in the SCE have (unweighted) demographic characteristics and consumer sentiment that are very similar to those of CCS volunteers (those who consent to being contacted for online surveys) and CCS respondents overall. Given that the pool of CCS respondents already was highly representative of the US population of household heads (as shown in Table A1), the similarity between SCE and CCS respondents indicates that our stratified sampling procedure in inviting CCS respondents was largely effective. This is further exemplified by the notable difference between the CCS and SCE samples in the age distribution of respondents. Reflecting the efficacy of our pre-stratification approach to inviting CCS consenting respondents, SCE participants are somewhat younger than CCS respondents, and in fact have an age distribution of household heads that is very comparable to that in the ACS.

While the previous comparison was concerned with how SCE *entrants* compared to CCS participants, we finally assess the representativeness of SCE respondents overall. That is, how representative are SCE respondents in a typical cross-section? The sample of SCE respondents each month of course reflects not only their initial recruitment into the panel, but also their continued participation over time. Table 3 reports the means and standard deviations of the monthly average sample characteristics of SCE respondents during the period October 2014 to September 2015 period. The first column of the table shows that the average (unweighted) characteristics of respondents in the SCE are very similar to those of SCE entrants (shown earlier in column 4 of Table 1), but they are slightly older and have slightly higher incomes on average, reflecting differences in survey participation rates after entering the SCE panel. The relatively small standard deviations reported in the first column further indicate that the sample composition of SCE participants each month is highly stable over time. This of course is not surprising given that SCE respondents constitute a panel, with approximately 90% of respondents in a given month participating again in the following month.

[15] That is, in inviting SCE volunteers we not only oversample those less likely to consent, but also those less likely to accept our invitation and those more likely to attrite from our panel once entered or to occasionally skip surveys, before completing the 12 month survey period.

[16] Newly invited SCE volunteers are only provided a one-time opportunity to join the SCE panel in the month for which they are first invited. Those who do not participate in the first month are no longer considered for future participation in the SCE. Note that the first-time participation rates listed include respondents with invalid or inactive email addresses.

As mentioned earlier, to account for any remaining differences between the SCE and ACS (for example due to differential sample attrition or skipping behavior), we apply weights to make our sample representative of the population of U.S. household heads. The weights are based on four individual characteristics (income, education, region and age) with targets based on the Census population estimates derived from the American Community Survey for that calendar year.[17]

Column 2 of Table 3 shows the means of the monthly weighted average demographic characteristics of SCE respondents, while column 3 shows the distribution of these characteristics in the 2013 ACS. A comparison indicates that weighting is highly successful in making the SCE samples comparable to the population of household heads in the U.S. overall.

[17] The weights applied to the survey responses are obtained using the "RIM" (Random Iterative Method) weighting method Sharot (1986). This method essentially uses minimum least squares to find the set of weights that minimize the distance between the marginal distribution in the sample and that in the population – given by the demographic targets. The weights are constructed through an iterative procedure that minimizes the distance between sample frequencies and population proportions sequentially along each dimension (demographic characteristic) separately, then iterates till the weights converge. Target statistics from the American Community Survey (ACS) are updated each year based on the most recent ACS release. We distinguish between 4 income groups (up to $30k, between $30k and $50k, between $50k and $100k, and above $100k), 3 education groups (up to high school, some college, and college graduate and above), 4 census regions (North East, Midwest, South and West), and 5 age groups (less than 30, 30 to 39, 40 to 49, 50 to 59 and 60 and above).

Table 3: SCE Sample Composition

	SCE Average (STDEV) monthly unweighted sample proportions (in %)	SCE Average monthly weighted sample proportions (in %)	2013 ACS proportions (in %)
	(1)	(2)	(3)
Sample Size	24	24	
Age <30	10.3 (1.2)	10.9	10.8
Age 30-39	17.5 (0.5)	16.9	16.9
Age 40-49	18.2 (0.6)	19.3	19.2
Age 50-59	21.4 (1.3)	20.6	20.8
Age 60+	32.6 (0.8)	32.3	32.4
Female	47.4 (1.0)	50.0	49.9
Up to HS	12.3 (1.0)	37.2	36.7
Some college	33.9 (1.8)	31.2	31.3
College Grad	53.8 (1.9)	31.5	32.0
Income under $50,000	37.7 (1.2)	48.3	47.9
$50,000-$99,999	36.4 (1.3)	30.6	29.7
$100,000 or more	26.0 (1.4)	21.1	22.5
New England	4.4 (0.7)	4.4	4.8
Middle Atlantic	13.6 (1.5)	13.6	13.2
EN Central	15.7 (2.1)	16.1	15.5
WN Central	6.4 (1.3)	6.1	7.0
South Atlantic	20.0 (1.5)	20.6	19.7
ES Central	5.1 (0.4)	6.2	6.1
WS Central	10.2 (1.6)	10.7	11.5
Mountain	8.9 (0.6)	7.9	7.1
Pacific	15.7 (0.9)	14.5	15.1
Northeast	18.0 (2.1)	18.0	18.0
Midwest	22.1 (1.4)	22.2	22.5
South	35.3 (2.7)	37.4	37.3
West	24.6 (1.2)	22.4	22.2

All SCE statistics based on surveys from October 2013 to September 2015.
Mean Proportion reported in the cells.
Standard deviations of sample proportions across months reported in brackets in first column.

4.1. Learning and Experience

One common feature of survey panels that deserves some discussion is learning. As respondents continue to participate in the survey and answer the same questions over time, participation in the survey may potentially affect responses through learning. For example, after seeing a question covering a certain topic for the first time, respondents may pay more attention to the media, or may simply think more about the topic, perhaps in anticipation of receiving the question again in a future survey. Alternatively, the respondent may become more familiar and comfortable with the question formats. If such learning effects exist and affect responses in a systematic way, then changes over time in a respondent's responses may not capture true changes in beliefs.

In analyses we conducted we find at best modest evidence of such effects in our panel. For instance, Figure 3 focusses on the density mean of short-term inflation expectations elicited for the respondents who entered the panel in 20014.[18] Specifically, we first take the absolute difference from one month to the next in each respondent density mean. Then, we plot the median of these absolute differences across respondents for each month of participation the panel. For instance, we can see that the short term inflation expectations density mean elicited for the median respondent changed by nearly 1.4 percentage points between the first and the second survey she completes. After that, the median respondent reports different beliefs from one month to the next (as should be expected), but the magnitude of the month-to-month changes remain relatively stable. Thus, most of the learning occurs within the first few months of the respondent's participation to the panel.

Figure 3: Median Month to Month Absolute Deviation in Density Mean for Short Term Inflation Expectations

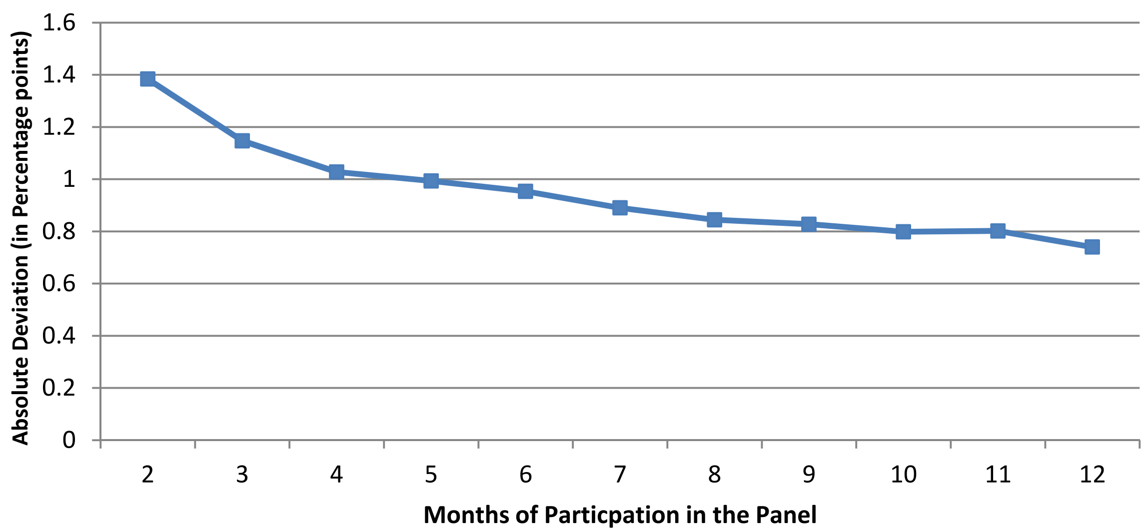

Perhaps most importantly, the design of the panel, with a constant in- and out- flow of respondents each month, assures a stable survey tenure distribution, so the extent of learning and

[18] To avoid selection effects due to respondents who repeatedly fail to complete the survey and rotate out of the panel quickly, we focus here exclusively on respondents who stay in the panel for at least 6 months.

experience (and any associated impact on responses), is constant over time. As a result, month-to-month changes in median responses should capture real changes in population beliefs.

5. Computation and Reporting of SCE statistics

To summarize and present our survey findings we report median responses overall and by demographic characteristics. The median is a robust measure of central tendency that is less sensitive to the presence of outliers than the mean.[19] Using a robust summary measure is important as we do not delete or recode outliers in the SCE. In addition to the median, for some survey questions we also report the 25th and 75th percentiles of the distribution of responses with the difference between the two quartiles, the IQR, representing a measure of dispersion or disagreement among respondents.

5.1. Quantile Interpolation

A common feature of response behavior when a survey question asks for a numerical response is the use of rounding. When asked about past or expected future changes in percentage terms, almost all respondents appear to round to the nearest integer value. Accordingly, when tracking changes in survey responses over time, it would be common to see either no change in the computed raw median or a sudden abrupt change of one or more percentage points. In case of grouped or rounded responses, it is therefore more informative to compute instead the median (and other quantiles of the distribution of responses) using an interpolation method. Interpolated medians will better capture shifts in the frequencies of responses around the median.[20] The same issue applies to other quantiles of the underlying distribution including the first and third quartiles.

To compute interpolated quantiles we use the symmetric linear interpolation approach proposed by Cox (2009).[21] We provide in Appendix B details about the procedure. We have compared Cox's procedure to other interpolation methods, including simple linear interpolation of the CDF

[19] See Huber (1981) on robust statistics and estimation. Robust methods provide automatic ways of detecting, down-weighting (or removing), and flagging outliers, largely removing the need for manual screening and deletion of outliers.

[20] For example, consider two points x and y (with $x<y$) and two different *empirical* cumulative frequency distributions. The first empirical distribution attains the values 0.4 and 0.51 in x and y, while the second empirical distribution attains the values 0.49 and 0.6 in x and y. When the raw median is defined as the first value at which the cumulative distribution reaches or exceeds 0.5, the two empirical distributions both have the median of y. However, one may expect the median of the *underlying continuous* distribution to be closer to y for the first distribution and closer to x for the second distribution.

[21] In Stata the procedure is implemented using the *iquantile* module. See Cox (2009).

(asymmetric) and the Harrell-David procedure.[22] Computed quantiles, and month-to-month changes in quantiles, are generally very similar.

5.2. Density Estimation

In addition to point forecasts and probabilities of binary events, we ask respondents in the SCE for their density forecasts of various continuous variables. As discussed in section 2.3, we elicit these by asking individuals to assign probabilities to ranges or intervals of possible future realizations. In addition to future inflation (at the one- and three-year horizons) we elicit density forecasts for year-ahead national home price growth and, for those employed, year-ahead earnings growth (holding the job and the number of hours fixed).

In reporting and analyzing such density forecasts we focus on two summary measures: the density mean and density IQR, defined as the difference between the third and first quartile. To compute the density mean and density quartiles of each individual's reported density, we use the reported bin probabilities to fit an underlying parametric density following the approach adopted by Engelberg, Manski and Williams (2009). This approach is explained in detail in Appendix C.

Once fitted, the estimated density parameters are used to compute each individual respondent's "density mean" and "density quartiles". The mean represents the expected value, so in case of the inflation density forecast we refer to the computed density mean as the respondent's "expected inflation rate". Similarly, we use the estimated parameters to compute density quartiles with the difference between a respondent's 75th and 25th percentiles (the IQR) measuring her "uncertainty". When we aggregate across respondents, we obtain the *median density mean* (and the *median density quartiles*) which we use predominantly in our reports (as discussed in the next section).

An important and unique strength of the SCE is its ability to provide quantitative measures of overall uncertainty among respondents and changes therein over time. In our SCE releases, we report the (non-interpolated) median of the respondents' IQRs as a summary measure of overall uncertainty in expectations. This measure should not be confused with our measures of disagreement of expectations among respondents. The latter is measured by the IQR of respondents' point forecasts, or the IQR of respondents' density means, with both measuring dispersion in beliefs across respondents, while our uncertainty measure captures average forecast uncertainty among respondents.

5.3. Reporting of multiple medians

For several expectation questions, we solicit both point forecasts and density forecasts. For example, respondents are asked how much they expect the average home price to change nationwide over the next 12 months. They are also asked for the percent chance that, over the

[22] In Stata the procedure is implemented using the *hdquantile* module. See http://ocair.org/files/presentations/Paper2005_06/UseIM.pdf.

next 12 months, the average home price nationwide will increase (decrease) by 12% or higher; between 8% and 12%; between 4% and 8%; between 2% and 4%; between 0 and 2%. As explained earlier, the latter bin probabilities are then used to fit the respondent's underlying density of beliefs about year-ahead home price changes.

One would expect the respondent's point forecast to represent some summary statistic of the central tendency of her density, such as the density mean or median. While this often appears to be the case, with point forecasts largely tracking density means (as well as density medians), for a nontrivial subset of respondents, reported point forecasts correspond to values in the tails of the respondent's density forecast. Similar findings were reported by Engelberg et al. (2009) for professional forecasters. An important advantage of using the density mean is that it captures the same measure across respondents, while this might not the case for point forecasts, which for some respondents may represent the density mean, while for others the density median or mode, or some other moment of the respondent's forecast distribution. For this reason in our monthly reporting of SCE findings, we place more emphasis on the *median density mean*, although we include both medians (of point forecasts and density means) in our interactive charts.

6. Dissemination of the Data

The monthly SCE findings are released on the second Monday of every month. The release takes the form of a press release (see here) as well as a set of interactive charts that show the trends in the different variables, for the overall sample as well as various subgroups (such as by age or census regions). The underlying chart data are made available at the same time.

To facilitate the use of these data by researchers and policy-makers, the micro data for the monthly survey are also released on the SCE webpage with a 9-month lag. Open-ended responses or sensitive information (such as the respondent's zip code) are not released.

The SCE is still a project in its infancy and the process of setting up webpages for the other data that are collected under the SCE umbrella (either as part of the ad-hoc modules or the quarterly surveys) is ongoing. The SCE Credit Access Survey, which provides information on consumers' experiences and expectations regarding credit demand and credit access every four months, is available here. As in the case of the monthly survey, the micro data are made public with a 9-month lag. The annual SCE Housing survey, which provides rich and high-quality information on consumers' experiences, behaviors, and expectations related to housing, can be accessed here; the corresponding micro data are released with an 18-month lag. Interested readers should check the data page of the Center of Microeconomic Studies for latest products related to the SCE.

REFERENCES

Armantier, Olivier, Wändi Bruine de Bruin, Simon Potter, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2013. "Measuring Inflation Expectations." *Annual Review of Economics* 5: 273-301.

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2015. "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?" *International Economic Review* 56, no. 2 (May): 505-36.

Armantier, Olivier, Wilbert van der Klaauw, Scott Nelson, Giorgio Topa, and Basit Zafar. 2016. "The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics*, Vol. 98, No. 3, pages 503-523

Bernanke B. (2004): "The Economic Outlook and Monetary Policy," speech at the Bon Market association Annual Meeting, New York, NY.

Bernanke B. (2007): "Inflation Expectations and Inflation Forecasting," speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts.

Bruine de Bruin, Wändi, Simon Potter, Robert Rich, Giorgio Topa, and Wilbert van der Klaauw.2010a., "Improving Survey Measures of Household Inflation Expectations." *Current Issues in Economics and Finance*, Vol. 16. No. 7.

Bruine de Bruin, Wändi, Wilbert van der Klaauw, and Giorgio Topa. 2011a. "Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices." *Journal of Economic Psychology* 32, no. 5 (October): 834-45.

Bruine de Bruin, Wändi, Charles F. Manski, Giorgio Topa, and Wilbert van der Klaauw. 2011b. "Measuring Consumer Uncertainty about Future Inflation."*Journal of Applied Econometrics* 26, no. 3 (April-May): 454-78.

Bruine de Bruin, Wändi, Wilbert van der Klaauw, Giorgio Topa, Julie S. Downs, Baruch Fischhoff, and Olivier Armantier. 2012. "The Effect of Question Wording on Consumers' Reported Inflation Expectations. " *Journal of Economic Psychology* 33, no. 4: 749-57.

Bruine de Bruin, Wändi, Wilbert van der Klaauw, Julie S. Downs, Baruch Fischhoff, Giorgio Topa, and Olivier Armantier. 2010b. Expectations of inflation: The role of financial literacy and demographic variables. Journal of Consumer Affairs, 44, 381-402.

Bryan M, Venkatu G. 2001. "The Demographics of Inflation Opinion Surveys." Federal Reserve Bank of Cleveland *Economic Commentary*, 2001 (October): 1–4.

Cox, Nicholas, (2009), IQUANTILE: Stata module to calculate interpolated quantiles, http://EconPapers.repec.org/RePEc:boc:bocode:s456992.

Curtin, R. 2006. "Inflation Expectations: Theoretical Models and Empirical Tests." Paper presented at the National Bank of Poland Workshop on The Role of Inflation Expectations in Modelling and Monetary Policy Making, Warsaw, February 9–10.

Delavande, Adeline. 2014. "Probabilistic Expectations in Developing Countries." *Annual Review of Economics* 6, 1-20.

Delavande, Adeline, Xavier Giné, and David McKenzie. 2011. "Measuring Subjective Expectations in Developing Countries: A Critical Review and New Evidence." *Journal of Development Economics* 94, no. 2 (March): 151-63.

Engelberg Joseph, Charles Manski, and Jared Williams. 2009. "Comparing the point predictions and subjective probability distributions of professional forecasters," *Journal of Business and Economic Statistics* 27, no.1: 30–41.

Gali, J. (2008): *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

Huber, Peter J. 1981. *Robust Statistics*. New York: Wiley.

Juster, Thomas. 1966. "Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design." *Journal of the American Statistical Association*, 61, 658–696.

Leiser, David and Shelly Driori. 2005. "Naïve understanding of inflation." *The Journal of Socio-Economics* 34, 179-198.

Manski, Charles. 2004. "Measuring Expectations." *Econometrica* 72, No. 5 (September), 1329–1376.

Sharot, T. (1986). Weighting survey results. In *Journal of the Market Research Society*, 28: 269-284.Sims C. (2009): "Inflation Expectations, Uncertainty and Monetary Policy," BIS Working Paper No. 275.

Svenson, Ola, and Goran Nilsson. 1986. "Mental Economics: Subjective Representations of Factors Related to Expected Inflation." *Journal of Economic Psychology* 7*:* 327-49.

Van der Klaauw, Wilbert, Wändi Bruine de Bruin, Giorgio Topa, Simon Potter, and Michael Bryan. 2008. "Rethinking the Measurement of Household Inflation Expectations: Preliminary Findings." Federal Reserve Bank of New York *Staff Reports*, no. 359, December.

Woodford M. (2004): "Inflation Targeting and Optimal Monetary Policy," *Federal Reserve Bank of St. Louis Economic Review*, 86(4), 15-41.

Appendix A: Additional Tables and Figures

Table A1: Comparison of Consumer Confidence Survey (CCS) and American Community Survey (ACS) Samples

	Full CCS sample*	2013 ACS	2014 ACS
	(in%)	(in%)	(in%)
Age <30	3.7	10.8	10.6
Age 30-39	11.3	16.9	16.9
Age 40-49	15.7	19.2	18.7
Age 50-59	23.9	20.8	20.6
Age 60+	45.5	32.4	33.2
Female	47.7	49.9	49.9
Up to HS	NA	36.7	36.3
Some college	NA	31.3	31.2
College Grad	NA	32.0	32.5
Income under $15,000	8.3	13.0	12.6
$15,000-$24,999	9.9	10.9	10.6
$25,000-$34,999	10.3	10.3	10.1
$35,000-$49,999	15.2	13.7	13.4
$50,000-$74,999	19.7	17.9	17.8
$75,000-$99,999	13.1	11.8	12.0
$100,000-$124,999	9.4	7.9	8.1
$125,000 or more	14.0	14.6	15.5
New England	5.0	4.8	4.8
Middle Atlantic	13.6	13.2	13.2
EN Central	17.6	15.5	15.4
WN Central	7.7	7.0	7.0
South Atlantic	19.9	19.7	19.7
ES Central	5.8	6.1	6.1
WS Central	9.2	11.5	11.6
Mountain	6.9	7.1	7.2
Pacific	14.5	15.1	15.1
Northeast	18.5	18.0	18.0
Midwest	25.2	22.5	22.4
South	34.9	37.3	37.4
West	21.4	22.2	22.2

Each number in the table is the proportion of the sample that falls in that category.
* CCS averages are unweighted averages, based on 64,133 CCS respondents during the October 2013-September 2015 period.

Figure A1

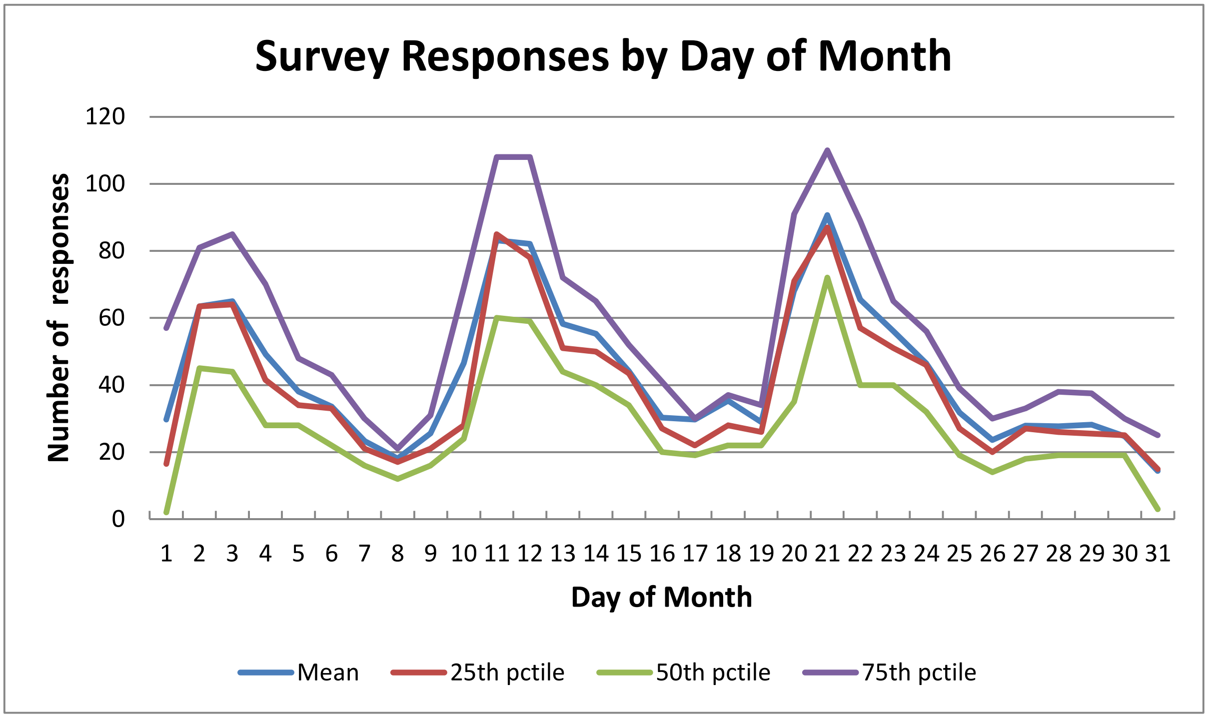

Survey Responses by Day of Month
120
100
80
60
40
20
0
Number of responses
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
Day of Month
Mean
25th pctile
50th pctile
75th pctile

Appendix B: Quartile Interpolation

The main idea behind the approach proposed by Cox (2009) to interpolate the cumulative distribution (or quantile) function is the following: rather than linearly interpolating $\Pr(X<x)$ or $\Pr(X\leq x)$, the average of the two, the *mid-distribution function* $\Pr(X<x)+0.5\Pr(X=x)$, is interpolated. More specifically, here is a brief description of the approach. First, for all observed values of x compute the cumulative proportions, symmetrically considered, as $\mathrm{CDFS}(x) = \Pr(X\leq x) - 0.5\Pr(X=x)$. Then to compute the median, determine the values of x observed with positive frequency with cumulative frequency CDFS that surround 0.5, defined as L (the smaller of the two) and H, and compute CDFS(L) and CDFS(H). Then the linearly interpolated median m equals

$$m = L + (H-L).[0.5 - \mathrm{CDFS}(L)]/[\mathrm{CDFS}(H)-\mathrm{CDFS}(L)].$$

Similarly for other quantiles, for example the third quartile, we identify the values of x observed with positive frequency with mid-distribution function values closest around 0.75, and in the equation above replace 0.5 by 0.75. When applying sample weights, the CDFS values are computed by calculating frequencies $\Pr(X\leq x)$ as sums of the relative weights (normalized to have mean 1) corresponding to all observations below or at x.

Appendix C: Density Estimation

We follow the approach proposed by Engelberg, Manski and Williams (2009) to fit a parametric distribution for each respondent based on the probabilities he reported for each possible density interval. We assume the underlying distribution to have a generalized Beta distribution when the respondent assigns positive probability to three or more outcome intervals We assume an isosceles triangular distribution when the respondent puts all probability mass in two intervals and a uniform distribution when the respondent puts all probability mass in one interval.

The generalized Beta distribution is a flexible four-parameter unimodal distribution that allows different values for its mean, median and mode and has the following functional form:

$$f(x)= \begin{cases} 0 & \text{if } x<l \\ (x-l)^{\alpha-1} (r-x)^{\beta-1} / B(\alpha , \beta)(r-l)^{\alpha + \beta -1} & \text{if } l \leq < x \leq r \\ 0 & \text{if } x>r \end{cases}$$

where $B(\alpha , \beta)=\Gamma(\alpha)\Gamma(\beta)/ \Gamma(\alpha ,\beta)$.

It uses two parameters (α and β) to describe the shape of the distribution and two more (l and r) to fix the support of the distribution. To fit a unique Beta distribution requires a respondent to have assigned positive probability mass to at least three (not necessarily adjacent) intervals.[23]

The triangular distribution, for cases where a respondent assigns positive probability to exactly two adjacent bins, has the shape of an isosceles triangle whose base includes the interval with the highest probability mass and part of the adjacent interval. Thus the triangle is anchored at the outer bound of the interval with probability mass above 50%.[24] Its density has the functional form:

$$f(x) = \begin{cases} \frac{4}{(r-l)^2}(x-l), & l \leq x \leq \frac{l+r}{2} \\ \frac{4}{(r-l)^2}(r-x), & \frac{l+r}{2} \leq x \leq r \\ 0 & elsewhere. \end{cases}$$

[23] In fitting a generalized beta distribution to a respondent's bin probabilities we use a minimum distance procedure that minimizes the distance between the empirical and estimated parametric distribution. We fix l and r to be the minimum and maximum bound of the positive-probability intervals, unless the corresponding bin is open-ended, in which case l and/or r are estimated together with α and β. In the latter case we restrict l to be greater or equal to -38 and restrict r to be at most 38. The sample statistics that we report are generally not sensitive to the choice of the imposed lower and upper bound.

[24] This rule only applies to the case of two adjacent intervals of equal width where neither interval is open-ended. In case of two adjacent intervals with unequal width, the support of the triangle is assumed to include the smaller-width bin in its entirely if its probability exceeds 40%, and includes the larger-width bin entirely otherwise, with the triangle covering only part of the adjacent bin. In the former case the triangle would be anchored at the outer bound of the narrower bin, and in the latter at the outer bound of the wider bin. In all cases where one of the two bins represents an open-ended interval (the left or right- tail of the distribution), the base always includes the inner closed-end bin, with the triangle anchored by the inner-most bound of the two intervals.

With the triangle being anchored at one of the outer bounds (l or r), there is only one parameter (either l or r) to fit, which fixes the center and height of the triangle.[25] Note that an isosceles triangle is symmetric, so the mean, median, and mode are identical to each other.

Densities are not fitted for respondents who put positive probability in only two bins that are non-adjacent, or for whom the probabilities do not sum to 100. Such respondents comprise less than 2% of our sample.

[25] In case of two adjacent bins with equal width, no estimation is required as the support of the triangle now fully includes both intervals, with the triangle anchored at the left most and right most interval bounds.