

Simon M Potter: The advantages of probabilistic survey questions

Keynote address by Mr Simon M Potter, Executive Vice President of the Markets Group of the Federal Reserve Bank of New York, at the IT Forum and RCEA Bayesian Workshop, Rimini, Italy, 19 May 2016.

* * *

Accompanying slides can be found on the Bank of Finland's website: [Slides](#) (PDF).

I would like to thank Marco Del Negro, Giorgio Topa, and Wilbert van der Klaauw for their excellent assistance in the preparation of these remarks, Luis Armona, Daniele Caratelli and Joseph Fiorica for able research assistance, and colleagues in the Federal Reserve System for their insightful comments and suggestions.

Good afternoon. It's a pleasure to talk with you about a topic that intrigues me as a former academic researcher, and that's of significant importance to me in my current role as head of the New York Fed Markets Group and as manager of the System Open Market Account for the Federal Open Market Committee (FOMC): the measurement of policy-relevant expectations through surveys.

It is entirely fitting to talk about this topic here in Italy, given the Bank of Italy's long history of research and experimentation on expectations data. This history includes extensive work on data gathered from surveys using a probabilistic question format – the focus of my talk today. Since 1989, the Bank of Italy's Survey of Household Income and Wealth has used this question format to ask households about their expected labor or pension earnings in the next year. The findings from these surveys have provided new insights into consumer behavior, such as the impact of subjective earnings uncertainty on precautionary saving.¹ The Bank of Italy was also a pioneer in using probabilistic questioning to elicit expectations from businesses in its 1993 Italian Survey of Investment in Manufacturing. Data from this survey have been used to analyze the impact of uncertainty about future product demand on the investment decisions of Italian manufacturing firms.²

As always, the views expressed are my own and do not necessarily reflect those of the Federal Reserve Bank of New York or of the Federal Reserve System.

Why are expectations data useful?

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 are therefore determined not only by preferences but also by expectations for future outcomes. Typically, neither preferences nor expectations are directly observed – a condition that poses a significant challenge to understanding economic behavior. Although different combinations of preferences and expectations may be consistent with the same observed behavior, they could result in different responses to shocks or to alternative policies and could therefore have different implications for public policy.³ The standard way to address this problem within economic models has been to assume rational expectations – that is, expectations that are consistent with equilibrium behavior in the model. In nearly all cases considered to date, this leads to the conclusion that expectations of economic agents are identical or will converge over time.

While this approach has given us invaluable insights in a multitude of models and applications, it faces a major empirical problem: in practice, economic agents hold widely

¹ Guiso, Jappelli, and Terlizzese (1994).

² Guiso and Parigi (1999).

³ Manski (2002).

differing views about the future path of economic variables, as I will illustrate throughout these remarks. Thus, recent research has focused on alternative assumptions regarding the nature of expectations held by economic agents. Measuring subjective expectations directly can inform our modeling assumptions on the information sets available to agents, and on the nature of the expectations formation and updating process.⁴

There are many Bayesians in the audience today. However, for the benefit of those who are not as familiar with Bayesian terminology, let me quickly define what I mean by “subjective expectations” or beliefs. This term refers to the personal probability distributions that individuals hold over uncertain events. Apart from the laws of probability and rules about how beliefs are updated with the arrival of information, there are no restrictions on these probability distributions, hence the use of the word “subjective.”⁵ Further, individuals might act according to their subjective beliefs without being able to fully articulate in standard forms their underlying probability distribution. Thus, a need arises to elicit these views in order to understand behavior more fully.

In monetary policymaking, central bankers have long pointed out the importance of measuring the expectations of financial market participants, households, and firms – especially with regard to inflation and the central bank’s so-called “reaction function” to changes in the economic outlook. For example, measuring market participants’ expectations for the policy rate path sheds light on both the central bank’s effectiveness in communicating its policy intentions and the transmission of anticipated future monetary policy actions into financial conditions. Inflation expectations of households, firms, and other economic agents influence a variety of decisions, including those related to consumption, saving, investment, and the setting of prices and wages, as well as the determination of the real interest rates that influence these decisions. The aggregation of these choices in turn determines realized inflation in equilibrium. Thus, inflation expectations are a crucial determinant of actual inflation and, as such, need to be carefully monitored. Further, the effectiveness of monetary policy and central bank communication relies on longer-run inflation expectations being well-anchored, making their reliable measurement very important for policymakers. More generally, expectations by consumers and businesses about a number of household, firm-level, and economy-wide outcomes are increasingly useful inputs into a variety of forecasting and “now-casting” models.

In addition to tracking measures of expectations implied by asset prices, central banks today rely increasingly on survey-based measures of subjective expectations. At the New York Fed, we conduct a diverse set of surveys to measure policy-relevant expectations. Today I will talk about the surveys of households and market participants we conduct: the Survey of Consumer Expectations (SCE), the Survey of Primary Dealers (SPD), and the Survey of Market Participants (SMP). These surveys highlight the innovative and cutting-edge research conducted at the New York Fed and represent the culmination of the Bank’s decade-long investment in the collection of better policy-relevant expectations data. As I will now explain, much of this investment has been directed toward introducing probabilistic survey questions, but we have also learned other important lessons about question design and survey mode, particularly regarding the measurement of inflation expectations which I will also discuss. In some cases such as results of the SPD and SMP the underlying probabilistic nature of survey questions is immediately apparent, in the case of the SCE as I will describe the underlying probabilistic question is not so apparent for many of the results.

⁴ Data on subjective expectations can also be used to improve the efficiency of model parameter estimates, refine posterior estimates of agents’ unobserved types, and help researchers estimate latent variables in macroeconomic models.

⁵ Of course it is also possible that survey respondents have beliefs that do not satisfy these conditions. I will discuss some evidence on this.

Why use probabilistic survey questions?

A key feature of the New York Fed surveys is their focus on collecting rich quantitative data on subjective expectations. Extending a practice with a longer tradition in the field of psychology and in surveys of professional forecasters, economists, and other financial experts, our surveys rely heavily on a probabilistic question format to elicit the likelihood respondents assign to different future events. This approach builds on a large and growing body of research, led by econometrician Charles Manski, that has demonstrated survey respondents' willingness and ability to answer questions expressed in this way.⁶

Traditionally, surveys ask respondents to express their answers in one of a few ways: for example, using a Likert scale (specifying whether an event is “very likely,” “likely,” or “not too likely”), selecting the outcome they consider most likely (“Do you expect the unemployment rate to be higher or lower than today?”) or giving a point response (“By what percent do you think prices will increase over the next year?”).⁷ A problem with the Likert scale is that it does not allow comparisons across respondents, because different respondents may interpret the scales differently. A well-known illustration of the problem was provided by Daniel McFadden, who looked at a health status measure that used a Likert scale ranging from “excellent” to “poor.” While 62 percent of Danish men reported their health to be excellent, only 14 percent of French men did so, despite their enjoying two more years of life expectancy. A problem with questions asking for the outcome considered most likely – that is, questions asking for the modal outcome – is the ambiguous information content of the responses. For example, in the case of an event with a binary outcome (where the event either occurs or not), the modal outcome can be assigned a probability that ranges from 50 percent to 100 percent. Finally, a problem with questions asking for a point forecast – “what do you think will be” or “by how much do you expect” – is that it is not clear whether the estimate individuals provide represents a mean, mode, median, or something else. Research also suggests that some respondents provide a point estimate that reflects varying degrees of loss aversion rather than just a pure forecast.

Among the advantages of questions asking directly for the likelihood or “percent chance” of different outcomes are ease of interpretation, comparability across respondents, and the ability to measure subjective uncertainty. The consistency of a respondent's answers can be checked using the laws of probability, and importantly, probabilistic questions permit elicitation of the quantitative measures of belief and uncertainty typically required in estimating economic models.

Researchers can use a variety of methods to check the validity and information content of probabilistic subjective expectations data. First, one can look at response rates, reported ease and clarity of the questions used to elicit expectations, the use of focal points in responses, and the internal consistency of responses. Second, one can consider whether reported expectations vary with observable characteristics in a predictable way.⁸ Third, whenever possible and assuming a stationary environment, one can check for relevance of reported expectations, looking at whether they are meaningfully related to future or past realizations. Fourth, one can ask whether subjective expectations help predict actual behavior. Finally, one can look at whether agents update their reported expectations in sensible ways upon receipt of relevant information.

Based on these criteria, a literature that has grown rapidly since the mid-1990s has demonstrated the validity and information content of responses to probabilistic questions

⁶ See Manski (2004), Hurd (2009), Attanasio (2009), and Delavande, Giné, and McKenzie (2011a, 2011b) for recent overviews.

⁷ Likert (1932).

⁸ For instance, researchers can consider whether mortality expectations vary in expected ways with age, education, and risky behaviors.

about a broad spectrum of economic concepts. This evidence extends beyond the United States and Europe to developing countries.⁹ Thus, probabilistic questions, with appropriate instructions and visual aid tools, have been successfully used not only in the Survey of Professional Forecasters and the U.S. Health and Retirement Survey (and its sister surveys in Europe and Japan) but also in surveys of fisherman and farmers in India, aspiring migrants in Tonga, poor households in rural Colombia, and junior high school students in Mexico.¹⁰ Some of these questions ask respondents to assign probabilities to various ranges of outcome – so-called *density forecasts*.¹¹ Generally, people across a wide range of countries and walks of life appear to be able to answer probabilistic questions, and their expectations are found to be useful in forming predictions of future economic outcomes.

The survey of consumer expectations

Turning now to the development and use of expectations surveys at the New York Fed, let me first discuss our Survey of Consumer Expectations (SCE). Over the last ten years, we extensively tested and eventually implemented a new survey of households, with the goal of collecting timely and accurate information at high frequency on U.S. consumers' expectations and decisions on a broad variety of topics. The SCE, launched in June 2013, was designed to fill gaps in existing data collections about consumer expectations and outcomes, to provide a more integrated data approach, and to take advantage of state-of-the-art survey techniques. The SCE is implemented as a monthly, nationally representative survey of about 1,300 household heads.¹² Recently the Bank of Canada implemented its own version of the SCE, fielded at a quarterly frequency.

The SCE has various components, which are outlined on Slide 2. First, respondents are asked a core monthly module of questions concerning their expectations about various macroeconomic and household-level variables. These questions cover inflation expectations and expectations regarding changes in home prices and the prices of various specific goods and services, such as gasoline, food, rent, medical care, and a 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, including changes in earnings, the perceived probability of losing their current job or leaving it voluntarily, and the perceived probability of finding a new job. Moreover, respondents are asked about the expected change in their household's overall income and spending. As I will describe in more detail below, these expectations questions are fielded at various time horizons and with various formats, including both point forecasts as well as density forecasts based on probabilistic questions. The second component of the SCE contains a supplementary "ad hoc" module each month on special topics.¹³ Finally, SCE respondents also fill out longer surveys each

⁹ See Delavande, Giné, and McKenzie (2011a) for a comprehensive review.

¹⁰ See Mahajan et al. (2008), Giné and Klonner (2007), McKenzie, Gibson, and Stillman (2013), Attanasio, Meghir, and Vera-Hernández (2005), Attanasio and Kaufmann (2009), and Delavande, Lee, and Yoong (2012).

¹¹ More precisely, respondents are asked to assign probabilities to a set of mutually exclusive and exhaustive ranges of continuous outcomes.

¹² The SCE survey instrument is fielded by the Demand Institute, a nonprofit organization jointly operated by the Conference Board and Nielsen.

¹³ 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 cover credit access, job search and retirement, and spending. Topics addressed so far in the "floating" supplement include the Affordable Care Act, savings from lower gas prices, student loans, family leave, and use of insurance products. Together, the core monthly module and the monthly supplement take about fifteen minutes.

quarter on various topics. These are up to thirty minutes in length and are separate from the monthly survey.¹⁴

Before giving more details about the SCE questions and some background on the research underpinning them, I would like to briefly point out two other important design elements of the SCE. First, responses are collected using an online survey tool. While it has become harder to obtain nationally representative samples with telephone and mail surveys, for which response rates have been falling, the opposite has been true for online surveys. Moreover, some evidence suggests that respondents are more likely to answer financial questions posed in online surveys. In a recent randomized survey experiment conducted by colleagues at the New York Fed and the Dutch central bank,¹⁵ the share of respondents reporting their inflation expectations in online interviews slightly exceeded the share reporting their expectations in face-to-face interviews.

In addition to being more cost effective, online surveys, when fielded at high frequency, provide important flexibility by allowing questions to be added or changed at short notice. It is thus possible to collect information on new economic and financial developments, such as the impact of recent changes in the price of oil. Other attractive features of online surveys are the ease of incorporating graphics and other visual refinements and the ease of conducting randomized interactive experiments to analyze the updating of expectations and the links between expectations and behavior.

A second key design element of the SCE is that it uses a rotating-panel sample design, with approximately 1/12 of respondents rotating in each month, after which they stay in the panel for up to 12 months and then rotate out. In addition to keeping the sample representative over time, this rotational structure allows us to track the changes in individuals' responses over time. In other words, we can compute changes in expectations over time after differencing out individual idiosyncratic effects. This design contrasts with repeated cross-sectional surveys in which an entirely new sample is drawn each month. The panel structure of our survey, with a more stable sample composition, therefore reduces volatility in summary statistics due to a changing sample and increases the signal-to-noise ratio. Finally, it is possible to use the panel to link short-term expectations to actual realizations, and measure individual-level forecast accuracy and its implications in real time.

Designing survey questions for household inflation expectations

The launch of the SCE followed an extensive testing phase, dubbed the Household Inflation Expectations Project (HIEP), which was initiated in 2006 to explore the feasibility of implementing a new survey of consumer expectations with a focus on inflation expectations.¹⁶ The HIEP had several objectives: to improve existing measures of consumer expectations, to clarify the process of expectations formation and updating, and to study the links between reported expectations and actual behavior in a variety of realms. The HIEP set up a working group – composed of New York Fed and Federal Reserve System economists, academic economists, behavioral psychologists, and survey design experts – that devised, conducted, and analyzed a series of cognitive interviews and experimental surveys to explore various dimensions of the planned new survey.

In particular, the HIEP analyzed the information content of the inflation expectations questions in the University of Michigan Survey of Consumers, the primary source of

¹⁴ Most of these surveys are repeated at a yearly frequency. The SCE currently contains quarterly surveys on the housing market; the labor market; informal work participation; and consumption, saving, and assets. A subset of these surveys is wholly or partially designed by other Federal Reserve Banks.

¹⁵ Bruine de Bruin et al. (2016).

¹⁶ See van der Klaauw et al. (2008) and Armantier et al. (2013) for overviews.

information on household inflation expectations at the time; tested alternative wording of potential inflation expectations questions; studied the feasibility of eliciting individual uncertainty about future outcomes for inflation, house prices, and earnings growth; and introduced a panel dimension to the experimental data collection effort in order to study the persistence of inflation expectations and their responsiveness to inflation surprises.

The HIEP considered different potential time horizons for its inflation expectations questions. For the short-term expectations, we settled on the one-year-ahead horizon – similar to the Michigan survey. We also examined the feasibility of asking for inflation expectations over a longer time horizon. The Michigan survey asks respondents by what percent per year they “expect prices to go up or down on the average, during the next 5 to 10 years”.¹⁷ We replicated this question in our experimental surveys, and found that the Michigan question elicits a mixture of interpretations, with some respondents using a 10-year horizon and others thinking about a 5-year horizon.¹⁸

From the results of our cognitive interviews and experimental surveys, we decided in the SCE to elicit medium-term inflation expectations at the one-year two-year-forward horizon. For example, in this month’s edition of the SCE we ask respondents what they expect the rate of inflation to be “over the 12-month period between May 2018 and May 2019.” Our testing suggests that respondents understand this format better than that of the 5–10 year Michigan question; consumers are better able to provide their expectations over a specific time period in the future, and there is little value in asking them about a longer time horizon. In addition, we believe that the one-year two-year-forward horizon adopted in the SCE is better suited to measure inflation expectations at the medium-term horizon that matters most for central bankers, since monetary policy is expected to exert its full effect within that time frame.¹⁹

With regard to question wording, it is important to note that the Michigan Survey asks about changes in “prices in general.” The HIEP tested three alternative wordings of a potential inflation expectations question: we presented respondents with questions about the change in “prices in general” (the wording used in the Michigan survey), the change in “prices you pay,” and the “rate of inflation.” Compared with the “rate of inflation,” the phrases “prices in general” and especially “prices you pay” induced respondents to think more about their personal price experiences and to focus on specific price changes for individual items, such

¹⁷ The Michigan survey’s precise question format is as follows: First, respondents receive the question “What about the outlook for prices over the next 5 to 10 years? Do you think prices in general will be higher, about the same, or lower, 5 to 10 years from now?” Those who respond “stay the same” are then asked whether they mean that prices will go up at the same rate as now, or that prices in general will not go up during the next 5 to 10 years. Those who indicate that they mean prices will go up at the same rate are then given the same follow-up questions as those who answer that they believe prices will be higher 5 to 10 years from now. Respondents who answer that they expect prices to be higher [lower] 5 to 10 years from now receive the question “By about what percent per year do you expect prices to go up [down] on the average, during the next 5 to 10 years?” Only respondents who give a response over 5 percent are then asked the clarifying follow-up question “Would that be [x] percent per year, or is that the total for prices over the next 5 to 10 years?” Respondents who answer “total” are then asked for a “per year” amount.

¹⁸ Further, the clarifying question used in the Michigan survey that asks whether respondents meant their response to reflect price changes per year or over the entire time period induced significant revisions. The follow-up question is administered only to respondents who give expectations over 5 percent, thus failing to correct misinterpretations among those who gave lower responses. In our experimental surveys, we instead administered the follow-up to everyone. We found that if we administered the follow-up question only to those giving responses over 5 percent, then the median long-term expectation would be 4.3 percent instead of 3.7 percent in our test sample. Thus, we have reasons to believe that the 5 percent follow-up question in the Michigan Survey leads to a systematic overestimation of actual average responses.

¹⁹ In contrast, for the Survey of Primary Dealers discussed later, we added a question in 2007 that directly elicited uncertainty over CPI inflation from 5 to 10 years ahead. See Potter (2011, 2012).

as food or gasoline.²⁰ In related research, we have shown that respondents tend to express more extreme inflation expectations when they are prompted to think of specific prices in answering the question rather than overall inflation.²¹ In the SCE, we therefore ask directly for expectations about the rate of inflation or deflation.²²

Probabilistic questions on the SCE

As I noted earlier, a key feature of our survey is its use of probabilistic questions. Following the current practice in the literature for the elicitation of expectations about continuous variables, we tested questions that elicited density forecasts by asking respondents to assign the percent chance that the value of interest would fall within different pre-specified ranges or “bins.” For instance, for inflation, we ask for the subjective probability distribution over a range of possible future inflation outcomes. This is shown in Slide 3.²³

For each respondent, we fit a flexible parametric probability density function (PDF) to their bin probabilities. We can then compute several moments of interest from the fitted density forecast for each individual respondent, including measures of central tendency such as the density mean or median, measures of uncertainty, such as the interquartile range, or IQR, and measures of skewness.²⁴

As I already noted, there are several important advantages to eliciting a density rather than a simple point forecast. In a world characterized by pervasive uncertainty, density forecasts provide a comprehensive representation of respondents’ views about possible future outcomes for the variables of interest. Importantly, density forecasts enable us to compute the degree of uncertainty associated with each respondent’s estimate. The measurement of forecast uncertainty is important for policymakers in order to assess the extent to which inflation expectations remain well-anchored. In addition, by eliciting density forecasts, we can focus on a specific measure of central tendency and use that for cross-respondent comparison and aggregation purposes. Research indicates that by forcing respondents to consider the likelihood of different outcome ranges, density-based expectations measures,

²⁰ In analyzing individual changes in expectations (from one month to the next) expressed by our SCE respondents, we find that changes in inflation expectations are little correlated with changes in expectations about future gasoline price changes.

²¹ Bruine de Bruin et al. (2011a) carried out two studies to see whether individuals who think about specific price changes in forming their inflation expectations report more extreme and dispersed expectations because they focus on more extreme specific price changes. In the first study, the researchers show that those who are asked to report *any* or *the largest* individual price change tend to recall more extreme ones than those who are asked to report *the average* price change, and subsequently report more extreme inflation expectations. In the second study, the researchers show that among those who are asked about inflation expectations without first being asked about individual price changes, about half nonetheless think about individual items. Those who do so then report more extreme and dispersed inflation expectations.

²² In our open-ended cognitive interviews and experimental surveys, we found that respondents generally are familiar with the term “inflation” and have a good understanding of the concept of inflation (Bruine de Bruin et al. 2011a, 2012). Other studies have found that members of the general public are familiar with the term “inflation” and have a basic understanding of what it means (Leiser and Drori 2005; Svenson and Nilsson 1986; Williamson and Wearing 1996).

²³ An advantage of eliciting density forecasts in this way – as opposed, for example, to eliciting values of the cumulative distribution function – is that answers are less likely to violate the laws of probability, such as monotonicity of the distribution function. A simple instruction and visual tool showing a running total as probabilities are assigned to bins leads to a negligible number of cases in which probabilities do not add up to 100 percent.

²⁴ Following the approach described in Engelberg, Manski, and Williams (2009), we fit a generalized beta distribution to the bin responses of each individual respondent whenever the respondent assigns positive probability to three or more bins. When respondents only assign positive probability to one or two bins, we fit a uniform or a triangular distribution, respectively.

such as the mean, are less prone to outliers than are one-off answers to point expectations, resulting in smaller cross-sectional dispersion.²⁵

Eliciting probabilistic beliefs and density forecasts presents a number of challenges: much research is still needed on the best way to elicit such beliefs – for example, with regard to the format of the questions; the number, width, and location of bins; whether to tailor bins to individual respondents; and how to estimate the underlying continuous densities. The elicitation of such beliefs involves many survey design choices, and the extent to which the findings depend on the way in which the information is collected is not yet fully known.

In line with earlier findings in the literature, our research shows that survey respondents are able and willing to provide density forecasts.²⁶ In our experimental surveys, fielded as part of RAND's American Life Panel, we observe negligible non-response rates to the bin probability questions. Further, when respondents are asked point forecast questions, some elect to provide a range, in order to express a degree of uncertainty in their forecast. We find that the use of ranges in point forecast questions, as well as the widths of reported ranges, are positively correlated with the number of bins to which respondents assign positive probability in density forecast questions, and with measures of uncertainty based on the fitted density estimates. Respondents also rate the density forecast questions only slightly harder to answer than the corresponding point forecast ones. Finally, respondents' estimated forecast uncertainty is meaningfully associated with month-to-month revisions in inflation expectations: greater uncertainty is significantly correlated with larger revisions – qualitatively consistent with Bayesian updating.

What is the best way to aggregate the information we elicit regarding probabilistic beliefs and density forecasts? For binary outcomes, such as the perceived percent chance of losing one's job over the next three months, we report the mean probability across respondents. For density forecasts, we can compute several aggregate measures. Slide 4 shows a stylized depiction of how we approach this. As our aggregate measure of expected inflation, we report the median of individual density means. We pick the density mean, which represents an individual's expectation for the rate of inflation, as our measure of central tendency, and use the median to aggregate across respondents, since the median is less sensitive to outliers. We can also look at medians of various percentiles of the individual distributions.²⁷

We measure aggregate forecast uncertainty as the median across respondents of the individual IQRs. We also monitor the probability that respondents attach to extreme inflation outcomes by computing, for instance, the median probability of deflation: in other words, the median of the individual subjective probabilities assigned to outcomes with inflation less than zero. Finally, to characterize the variation in inflation expectations across respondents, we compute the dispersion of expected values: the IQR of individual density means. Again, we can look at other percentiles of the distribution across respondents to detect possible clustering and polarization.²⁸ We publish at a monthly frequency many of the aggregate measures of inflation expectations, as well as measures for a broad range of labor market

²⁵ Delavande et al. (2011b).

²⁶ See Bruine de Bruin et al. (2011b).

²⁷ By tracking the same individuals over time, we can in principle study the extent to which respondents update their expectations in response to new information. For instance, if respondents receive information that deviates from their priors, their density forecasts could initially exhibit greater uncertainty and even become bimodal until they converge to a new value for their central tendency.

²⁸ For example, in the immediate aftermath of the financial crisis, we detected a certain amount of polarization in the distribution of inflation expectations, with some consumers expecting relatively high inflation, and some expecting deflation.

and financial expectations, on our survey website, shown on Slide 5. In addition to providing national figures, we show trends for different demographic subgroups.

Current trends and additional research

So far, we have discussed the elicitation of density forecasts for a single variable of interest. In addressing various policy questions, however, we may be interested in the subjective joint probability distributions for two or more variables. For instance, if we are interested in households' expectations about real wage growth, we may want to elicit the joint distribution over future outcomes for both inflation and nominal wage growth. Measuring the joint distribution of expectations over income and spending growth may also be very useful.

This is an area at the frontier of current research, and more work is needed to test alternative ways to elicit joint subjective distributions. At the New York Fed, we have experimented with questions asking for density forecasts for one variable, conditional on ranges of outcomes for the other variable of interest. These conditional probability distributions, together with the outcome probabilities for the conditioning variable, can be used to recover the subjective joint distribution for each respondent. For instance, in the New York Fed's Survey of Primary Dealers (SPD), we have fielded questions to elicit the joint probability distribution for the size of the Federal Reserve's balance sheet and the state of the economy.

We might also be interested in exploring the cross-section of density forecasts across respondents, for example, demographically. As I noted earlier, the SCE is based on a representative sample of U.S. household heads, and the data highlight the rich heterogeneity of expectations held by consumers. For example, Slides 6 and 7 show the recent trends in median three-year-ahead inflation expectations overall and sorted by the education level of respondents. Slide 8 shows trends in inflation uncertainty, also by education. As is common in surveys of inflation expectations, more highly educated respondents typically report lower inflation expectations, and express less uncertainty in their density forecasts. We observe similar patterns when the responses are broken out by income, numeracy, or financial literacy. These differences could reflect heterogeneity in information sets, in individual experiences, or in the way people process information about inflation.²⁹

Next, I'll share two experiments we conducted during the development phase of the SCE in order to analyze the validity and information content of our inflation expectations questions. In the first experiment, which was financially incentivized,³⁰ survey respondents were first asked for their inflation expectations. They were then asked a series of ten questions offering a choice between two investments, one yielding a nominal return that varies with inflation, the other yielding an inflation-protected payoff with increasingly large amounts over the ten questions. When choosing between the two investments, the respondent seeking to maximize her payoff *should switch* investments at most once, and if so, should switch from the one yielding a nominal return to the inflation-protected investment sooner, the higher her inflation expectations.³¹

We found that the choices in the financially incentivized experiment were indeed meaningfully related to respondents' inflation expectations, on average. Differences in reported uncertainty were also associated with different choices in the experiment, in the direction predicted by theory. Deviations from expected utility maximization were more likely

²⁹ The observed heterogeneity of expectations is important for policy purposes, since we find that agents with different characteristics act on their inflation beliefs and update their expectations in distinct ways.

³⁰ Armantier et al. (2015).

³¹ This is an agent's optimal choice under expected payoff maximization. We also show that the respondent should switch sooner to the inflation-protected investment, the higher her risk aversion and the more uncertainty she expresses in her inflation expectations.

to occur for respondents with low education, as well as those with low numerical and financial literacy. We were able to repeat the experiment with the same respondents six months later. We found that individual changes in expectations were significantly associated with changes in how early respondents switched between investments, and that not only the direction but also the magnitude of these changes was consistent with a simple expected utility maximization framework. To our knowledge, these results provide the first direct evidence regarding a meaningful link between survey-based inflation expectations and actual behavior.³²

In the second experiment, we studied the formation process of inflation expectations by looking at whether and how survey respondents updated their inflation expectations when they received relevant new information.³³ Specifically, we first elicited inflation expectations. We then randomly provided information to a subset of respondents on either past food price inflation, or professional economists' forecasts of future inflation from the SPF. Finally, we re-elicited respondents' inflation expectations, and studied whether they were correlated with the information content of the signal they received.³⁴

We found that respondents, on average, updated their inflation expectations in response to the information provided, and they did so sensibly and in a direction consistent with Bayesian updating – with larger revisions for less informed respondents and for those with greater baseline uncertainty. Further, there was significant heterogeneity both in how fully informed respondents were about objective inflation measures and in their updating behavior. This finding points to the potential importance of allowing for heterogeneous information-processing rules in our economic models. These findings are also consistent with existing sticky-information models of expectations formation, since cross-sectional disagreement falls after the provision of information.³⁵ Results from the experiment also indicate that expectations about changes in the “prices you pay,” like the similar Michigan survey question about “prices in general,” are more responsive to information about food prices than expectations about the rate of inflation. This finding is consistent with our observation that wording based on “prices” causes respondents to focus more on price changes in their own consumption basket and to report expectations that are more correlated with gas and food price changes.³⁶

Finally, we have been exploiting the rotating panel nature of the SCE to better understand the changes in medium-term consumer inflation expectations that has occurred since July 2015. As I mentioned earlier, the panel structure of the SCE enables us to analyze the changes at the individual level, that is, for the same respondent – thus abstracting from

³² This financially incentivized experiment provides a joint test of our question's success in measuring a respondent's inflation expectations, and of whether consumers act on their expectations according to theory.

³³ Armantier et al. (forthcoming).

³⁴ Before providing the information, we also asked respondents for their expectations about the information provided in each treatment (e.g., what they thought past food price inflation had been), in order to control for respondents' priors about the information that they would receive.

³⁵ See, for instance, Mankiw and Reis (2002) or Carroll (2003).

³⁶ In our research, we have also been studying the extent to which subjective expectations in the SCE are associated with individual outcomes. For instance, for a given individual, the perceived probability of making large purchases over the next 4 months (home appliances, electronics, furniture, home repairs, improvements or renovations, autos or other vehicles) is highly correlated with actual purchases of those items 4 months later. Household spending growth expectations (over the next 12 months) are also significantly associated with self-reported actual spending growth 12 months later. In the labor market realm, the perceived probability of finding a job over the next 3 months accurately predicts actual transitions from unemployment to employment, in those 3 months. And the expected arrival rate of job offers is predictive of actual job offers over a 4-month horizon. So the evidence suggests that the subjective expectations we measure are indeed informative of future events, with individuals more likely to experience an event typically reporting a higher expectation of the event occurring.

possible changes in expectations coming from changes in sample composition. In particular, we looked at the group of respondents who completed the survey both in September 2015 and in January 2016, and compared their three-year-ahead inflation expectations as measured by the individual density means across the two surveys.³⁷ As shown in Slide 9, we found that the entire distribution of medium-term inflation expectations shifted to the left, indicating a widespread decline in expectations. Median expected inflation among these repeat respondents declined by 53 basis points, compared to 39 basis points without controlling for the changing sample composition.

The May results of the SCE showed an increase of 29 basis points from April in the median of the individual means (see Slide 6) and this can also be examined by holding the sample constant using the panel structure. In this case we find an increase of 27 basis points from April and the distribution shifting to the right, although the statistical significance of the change is more marginal than the decline over the longer period.³⁸ As can be seen in Slide 9 there is a wide dispersion in views about future inflation held by consumers.³⁹ This sort of analysis highlights the value of the panel nature of the SCE in yielding more robust measures of expectations.

Let me turn next to the surveys of market professionals conducted by the New York Fed.

Probabilistic questions on the surveys of primary dealers and market participants

Ahead of each meeting of the FOMC, the Trading Desk of the New York Fed asks a group of financial market participants, including both primary dealers and buy-side investors, to provide their views on a range of economic and financial indicators. The Desk has surveyed the primary dealers for over a decade, and since 2011 it has made the questions and results of the Survey of Primary Dealers (SPD) publicly available.⁴⁰ Beginning in 2014, the Desk also launched the Survey of Market Participants (SMP), which solicits views from active investment decision makers, such as mutual, pension, and hedge fund managers as well as corporate treasurers.

In both the SPD and the SMP, many of the questions are tailored toward matters that are of policy relevance at the time of the survey. Further, we use the fact that the respondents are market professionals with a good understanding of finance and economics to ask direct and at times technically complex questions. In this regard, the surveys contrast with the SCE, which uses language that is accessible to respondents with no economic or financial training.⁴¹

³⁷ See Armantier et al. (2016).

³⁸ Using a standard test of the difference between two distributions, the change from September to January is significant at the 1 percent level. The change from April to May has a “p-value” of 16 percent.

³⁹ This dispersion implies that random samples of small size (less than one thousand) will exhibit variability even if the underlying population distribution of inflation expectations remains constant.

⁴⁰ Correia-Golay, Friedman, and McMorrow (2013) provide an introduction to the SPD.

⁴¹ Many surveys of market participants do not use their technical sophistication to the extent necessary to produce robust measurement of expectations. For example, currently many surveys ask market participants to identify the FOMC meeting at which the next increase in the fed funds target range will most likely occur. This question implicitly assumes that the next rate change will be an increase; moreover, eliciting expectations about the most likely meeting when the number of future meetings is large is particularly uninformative. For example, there are five FOMC meetings remaining this year, so the “most likely” meeting could be one with 21 percent probability. In the SPD and SMP, we are careful to elicit the probabilities of the full range of possible outcomes. For example, the current mean probability that the next change in the fed funds target will be an increase is around two-thirds.

Probability versus point forecasts

For many policy relevant variables, both the SPD and the SMP elicit respondents' views on the entire distribution of future outcomes, and not only on their modal projections.⁴² Slide 10 provides a striking example of why this is important. The figure shows the projections for the federal funds rate at the end of 2017 (top panels) and 2018 (bottom panels) at three different times: December 2015, January 2016, and March 2016.

For each panel, we show the point forecast (red vertical line) averaged across respondents in both the SPD and SMP. This point forecast likely corresponds to the modal path for the federal funds rate at the end of each year.⁴³ The green lines depict the density forecasts, averaged across respondents, and the blue vertical lines show the mean projections implied by these density forecasts, which I will refer to as the “pdf-implied mean” in the remainder of this talk.⁴⁴ Note that the average modal path is usually close to, but does not match, the average mode of the marginal distributions for the policy rate. To the extent that participants are responding about the mode of the joint distribution, as opposed to the mode of each marginal distribution, this would not be surprising.⁴⁵

The striking feature of Slide 10 is the contrast between the steadiness of modal point forecasts and the dramatic changes in density forecasts. Modal forecasts for the policy rate in both year-end 2017 and 2018 barely moved from December to January, and edged slightly down in March. Had we just elicited point forecasts, we would have concluded that survey respondents' views had not changed very much. In contrast, the pdf-implied mean forecast decreased substantially from December to January, and then increased slightly from January to March. More importantly, density forecasts shifted from being slightly skewed toward higher rates in December to being skewed toward lower rates in January. In fact, density forecasts became bimodal in January, and in March, the “low rate” mode became the dominant one for 2018.

I should stress that the change in the density forecasts from December to January occurred following substantial changes in how the survey questions elicited respondents' probability-weighted expectations. In December, and in many preceding surveys, both the SPD and SMP asked respondents for the unconditional probabilities that they assigned to various rate outcomes. In January, however, survey respondents were asked to assign probabilities to rate outcomes at the end of 2017 and 2018 *conditional* on two different scenarios: “returning to the zero lower bound (ZLB) at some point in 2016–2018,” or “not returning to the ZLB at any point during 2016–2018.” Respondents were also asked for their subjective view of the probability of each scenario, a topic that I will discuss later. We changed the survey both

⁴² This is the case for some of the variables of interest, such as the level of the policy rate or, in the past, the projected level of the balance sheet.

⁴³ As noted earlier, questions that elicit a “most likely” value, or mode, can be difficult to interpret. This is particularly true of questions about the path of a variable, for example, the federal funds rate. If respondents are asked for their modal path, then technically, they should give the mode of the joint distribution. For example, the “mode” for policy rate projections could be the mode across all paths that the federal funds rate could follow from 2016 to 2018. It is an open question whether survey participants respond in this way, or respond with the mode of each marginal distribution.

⁴⁴ A number of explanations are in order. The survey asks for probabilities (marginal distributions) associated with bins for the federal funds rate. The green line is the density implied by these probabilities, assuming a uniform distribution within each bin, and assuming that the end bins are truncated (since this is the policy rate, truncation of the lower open bin is quite natural; less so for the higher open bin, but very little probability is placed there by the average distribution). For visual reasons, the density forecast is plotted by connecting the mid-points of each bin. The mean is that implied by this density forecast.

⁴⁵ In addition, even if participants were reporting the mode of the marginal distribution, the mode of the average distribution does not necessarily match the average of the modes, as is well understood. Of course, there is always the possibility that participants may not provide consistent answers when asked about the modal forecast and the entire distribution.

because we thought this was a more effective approach to elicit unconditional distributions, and because we were interested in the scenario probabilities. Note that from January to March the survey question did not change, but we added an explicit “negative rate” bin.

In addition to eliciting density forecasts conditional on various scenarios, the surveys have sometimes asked for joint probability distributions for different variables. This is helpful in assessing how market participants expect the FOMC to respond to various economic scenarios. In June 2013, for instance, the SPD asked for the probability distribution for the size of the Federal Reserve’s securities holdings – which we call the System Open Market Account, or SOMA, portfolio – at the end of 2014, conditional on three mutually exclusive outcomes for the unemployment rate at the end of 2013: less than 7.3 percent, between 7.3 and 7.5 percent, and greater than 7.5 percent.⁴⁶ For reference, the SOMA portfolio was about \$3.2 trillion in June of 2013. Clearly, the idea behind this question is to see how market participants expect the FOMC’s asset purchase programs to evolve in response to new data about the labor market.

Slide 11 shows the three conditional distributions for the projected change in the size of the portfolio relative to the June 2013 level, as well as the unconditional distribution, formed by using the probabilities of the unemployment scenarios. The results show that respondents thought that monetary policy decisions would be strongly dependent on labor market outcomes, a result that is in line with the September 2012 FOMC statement.⁴⁷

The probability of the SOMA portfolio increasing by more than \$800 billion to over \$4 trillion moved from about 30 percent conditional on unemployment being less than 7.3 percent to about 55 percent conditional on unemployment being greater than 7.5 percent. The probability of the SOMA increasing by more than \$1.3 trillion quadrupled between these two scenarios, from about 5 percent to 20 percent. In the end, unemployment was 6.7 percent by the end of 2013, and the year-end 2014 level of the SOMA portfolio reached \$4.25 trillion: a \$1.1 trillion change relative to June 2013. If we interpolate the bins using a uniform distribution, then this outcome is around the 81st percentile of the distribution conditional on the good labor market outcome and in the 75th percentile of the unconditional distribution.

If we were to judge on the basis of our survey of primary dealers, the outcome of the asset purchase program was on the high side but not surprisingly so. However, we don’t know how other market participants assessed the implications for the size of asset purchases from the September 2012 FOMC statement, particularly going into the June 2013 FOMC meeting. Many may recall that the June FOMC press conference continued something called the “taper tantrum.” As former Chairman Ben Bernanke writes in his memoir, “What explained markets’ strong reaction and why did it surprise us? In retrospect, I think our view of market expectations was too dependent on our survey of securities dealers.”⁴⁸ Importantly, it is possible that the size of future asset purchases was seen as smaller by primary dealer respondents than by the marginal investor.⁴⁹ Following this experience, we introduced the Survey of Market Participants to better capture the diversity of views in the market. In reaction to the taper tantrum, former Governor Jeremy Stein examined the large effects that can be produced if beliefs are dispersed and if ex-ante marginal investors receive information

⁴⁶ The specific question was: “Please indicate the percent chance you attach to the dollar level of the SOMA portfolio falling within the following ranges at year-end 2014 for each of three hypothetical unemployment rate scenarios.” See the June 2013 SPD survey.

⁴⁷ “If the outlook for the labor market does not improve substantially, the Committee will continue its purchases of agency mortgage-backed securities, undertake additional asset purchases, and employ its other policy tools as appropriate until such improvement is achieved in a context of price stability.”

⁴⁸ Bernanke (2015, 549).

⁴⁹ To quote Bernanke again, “In effect our PhD economists surveyed their PhD economists. It was a little like looking in a mirror. It didn’t tell us what rank-and-file traders were thinking”.

that leads them to strongly revise their views about how much risk to take. The outcome could be that the new marginal investor driving market prices has a very different set of beliefs and risk appetite.

Market rates and probability forecasts

It is well understood in theory why the market-implied path of the policy rate usually differs from survey measures of policy expectations.⁵⁰ Even if the marginal investor's probability distribution for future outcomes coincided with that of the average survey respondent, she would want additional compensation for bearing risk in states of the world that are unfavorable to her. It is conceivable in the current environment that outcomes in which the federal funds rate would return to the zero lower bound, or more generally remain low, are not good news for the overall portfolio of the marginal investor. Thus, in assessing the price at which she would be willing to trade interest rate derivatives like federal funds futures, she would weight those "negative" states of the world, such as states involving a return to the zero lower bound, more than her subjective probability distribution would imply – and the opposite would hold true for "positive" states of the world. In other words, in this environment the risk premium compensation she requires would pull the market-implied path downward, away from survey expectations.

While risk premia, especially at this juncture, are likely very important, looking only at pointforecasts may overemphasize their role when these point forecasts coincide with modal forecasts. In particular, when the implied policy rate from federal funds or Eurodollar futures changes, while the point forecasts from surveys do not, it is customary to attribute all the change in the market rates to risk premia.

The top two panels of Slide 12 compare the market-implied federal funds path for year-end 2016, 2017, and 2018 with the surveyed modal path projections, averaged across respondents.⁵¹ It is clear that the modal forecasts are much higher than market-implied rates, especially for 2017 and 2018, when the gap is larger than 100 basis points. Moreover, market-implied rates fell from December to January, while modal forecasts moved little between December and January, and fell slightly in March.

The bottom left panel of Slide 12 shows the risk premium derived from a standard econometric model.⁵² These are negative, which is consistent with the gap between market implied rates and survey expectations. However, the premium at short maturities is about as large, if not larger, than that at long maturities – a finding that is inconsistent with the fact that the gap at short maturities is a lot smaller. Moreover, the estimated premium becomes less negative from December to January, implying that market-implied rates should, *ceteris paribus*, have increased, while in fact they fell.

Changes in the probability distribution are, however, consistent with the changes in market-implied rates. The bottom right hand panel shows the pdf-implied means for the three different horizons, two of which were shown before in Slide 10. These are also higher than market-implied rates, but are closer than the point forecasts, especially in January. Most importantly, they fall significantly from December to January, in line with the change in market rates. Pdf-implied means changed little between January and March. However, as we know from Slide 10, this masks substantial changes in the underlying density forecasts, with the lower rate mode becoming more prominent.

⁵⁰ See e.g. Crump et al. (2014), and more recently Brodsky et al. (2016a).

⁵¹ Market-implied rates were derived from futures prices at the same time as responses were received from the surveys.

⁵² Adrian, Crump, and Moench (2013).

Finance models make a variety of assumptions that allows them to extract risk premia from asset prices. Once that is done, they can obtain the underlying so-called “physical” probability distribution of the marginal investor – that is, the actual probability she assigns to various outcomes. Brodsky et al. propose a complementary approach, called “tilting,” which is, in a way, the reverse of the finance approach.⁵³ It starts from the survey probability distribution and asks: How much do I have to “tilt” the original distribution – that is, how much probability mass do I have to shift around – to obtain the market-implied path as its mean?

There are many ways to tilt or alter the survey distribution such that its mean is the market rate. Our approach is to require that the tilted distribution is as close as possible to the original survey distribution, where the notion of distance is one that is commonly used in statistics, and is based on a measure called relative entropy or Kullback-Leibler information criterion (KLIC). This relative entropy approach is nothing new. It was introduced to the economic forecasting literature by Robertson, Tallman, and Whiteman (2005), and used more recently by Altavilla, Giacomini, and Costantini (2014), among others.

So, what do the tilted survey distributions look like, and how large are the discrepancies with the original probability distributions? Slide 13 shows the results for year-end 2018 using two survey dates: December, shown in the left panel, and January, in the right. For each panel, the red line shows the probability distribution averaged across respondents. The vertical gold and gray lines show the survey-implied mean and the market-implied value for the policy rate, respectively. The blue line shows the tilted distribution.⁵⁴ We see that the amount of tilting required for the January survey is very small in comparison to that needed for the December survey. For December, the tilted distributions are very different from the original distributions, with the probabilities in the lowest bin roughly tripling.

How do we interpret the difference between the original survey distribution and the tilted distribution? If we assume that the marginal investor and the average survey respondent share the same probability distribution, this difference loosely captures the additional premium or discount associated with future interest rate outcomes. Under this interpretation, the cost of insuring against low policy rate outcomes fell from December to January, because less tilting is required in January than in December, a conclusion consistent with the change in estimated model-based risk premia.

Alternatively, the difference might reflect the disparity between the subjective probability distributions of the marginal investor and the average survey respondent – and therefore measure the dispersion in beliefs between the two. Under this interpretation, the marginal investor in the futures market was much more pessimistic – in other words, placed more probability on low rate outcomes – than the average survey respondent in December, while by January the average survey respondent revised her beliefs downward. Note that the marginal investor in the futures market likely changes over time, and this change can be an additional source of fluctuations in the gap between market-implied rates and mean projections.

Unfortunately, we have no direct measure of the beliefs of the marginal investor. But we can obtain evidence on whether the dispersion in beliefs fell from December to March by looking at the dispersion of beliefs among survey respondents. One measure of such dispersion in beliefs is the extent to which the KLIC varies across respondents. A high (low) KLIC means that the respondent’s forecast distribution requires substantial (little) tilting to deliver the market-implied rate. Slide 14 shows, in the bars outlined in light blue, the KLIC distribution

⁵³ Brodsky et al. (2016b).

⁵⁴ Note that this distribution is constructed by tilting the average of respondents’ probability distributions. Alternatively, we could have tilted the probability distribution for each respondent and then averaged across these tilted distributions. While the two approaches do not, in general, deliver the same answer, in this case they almost do.

across respondents for the 2018 projections in December, and in the dark blue bars, the distribution for the 2018 projections in January. It is clear that dispersion in beliefs fell dramatically between the two surveys. The right panel shows that dispersion increased slightly in March but remains below the December levels.

Belief heterogeneity

Slide 14 already illustrates that forecasts are very different across survey respondents. I'll now provide further evidence of how pervasive and important heterogeneity in beliefs is.

Recall that in both the recent March and January surveys, respondents were asked about their subjective probability of "returning to the ZLB at some point in 2016–2018." Slide 15 plots the distribution across respondents of this probability. It shows that respondents belonged to three groups, roughly speaking: a few who believed this probability to be low (less than 15 percent of respondents), a majority who thought the probability was about one in four (between 20 and 35 percent of respondents), and a sizable minority who assessed this probability as fairly large (more than 40 percent of respondents).

This heterogeneity in the probability assigned to a return to the ZLB translates of course into heterogeneity in the probability distribution for the federal funds rate. The left panel of Slide 16 shows the smoothed density forecasts for the year-end 2017 federal funds rate for all March 2016 survey respondents. The density shown in the bottom left panel of Slide 10 is simply the average of the densities shown in Slide 16. Two features stand out. First, respondents display very different views on future policy rates, consistent with the evidence in Slide 15. Second, the bimodality of the distribution shown earlier in Slide 10 is not simply the result of aggregation, but is a feature of many individual density forecasts. In fact, the majority of respondents appear to have a "low rate" model, which mostly lies in the 0 to $\frac{1}{4}$ percent range, and a "higher rate" mode, which varies substantially across participants and ranges from about 1 to 3 percent.

The right panel of Slide 16 shows the cross-sectional distribution of the modal path projections. The cross-sectional distribution of point forecasts gives the misleading impression that there is broad agreement among respondents on the policy rate at year-end 2017, when in fact this is not the case, as we can see from the left panel. This misleading impression arises from the fact that when interpreting the modal projections one often assumes that the density is unimodal – an assumption that is clearly violated here.

Conclusion

In conclusion, I would like to re-emphasize three points. Survey respondents can answer probabilistic questions. Probabilistic questions produce measures of subjective expectations that are superior to point forecasts or most likely outcomes. Heterogeneity of beliefs is pervasive and can have important implications for policymakers. Thank you for your attention.

References

- Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110, no. 1: 110–38.
- Altavilla, Carlo, Raffaella Giacomini, and Riccardo Costantini. 2014. "Bond Returns and Market Expectations." *Journal of Financial Econometrics* 12, no. 4: 708–29.
- 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. Forthcoming. "The Price Is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics*.
- Armantier, Olivier, Wilbert van der Klaauw, Giorgio Topa, and Basit Zafar. 2016. "[Who Is Driving the Recent Decline in Consumer Inflation Expectations?](#)" Federal Reserve Bank of New York *Liberty Street Economics (blog)*, January 25.
- Attanasio, Orazio. 2009. "Expectations and Perceptions in Developing Countries: Their Measurement and Their Use." *American Economic Review Papers and Proceedings* 99, no. 2: 87–92.
- Attanasio, Orazio, and Katja Kaufmann. 2009. "Educational Choices, Subjective Expectations, and Credit Constraints." NBER Working Paper no. 15087, July.
- Attanasio, Orazio, Costas Meghir, and Marcos Vera-Hernández. 2005. "Elicitation, Validation, and Use of Probability Distributions of Future Income in Developing Countries." Paper prepared for the 2005 Econometric Society Meeting.
- Bernanke, Ben S. 2015. *The Courage to Act: A Memoir of a Crisis and Its Aftermath*. New York: W. W. Norton and Company.
- Brodsky, Bonni, Marco Del Negro, Joseph Fiorica, Eric LeSueur, Ari Morse, and Anthony Rodrigues. 2016a. "[How Do Survey- and Market-Based Expectations of the Policy Rate Differ?](#)" Federal Reserve Bank of New York *Liberty Street Economics (blog)*, April 7.
- Brodsky, Bonni, Marco Del Negro, Joseph Fiorica, Eric LeSueur, Ari Morse, and Anthony Rodrigues. 2016b. "[Reconciling Survey- and Market-Based Expectations for the Policy Rate.](#)" Federal Reserve Bank of New York *Liberty Street Economics (blog)*, April 8.
- 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, Maarten van Rooij, Federica Teppa, Klaas de Vos. 2016. "[Measuring Expectations of Inflation: Effects of Survey Mode, Wording, and Opportunities to Revise.](#)" De Nederlandsche Bank working paper no. 506.
- 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.
- Carroll, Christopher. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118, no. 1: 269–98.

- Correia-Golay, Ellen, Steven Friedman, and Michael McMorrow. 2013. "[Understanding the New York Fed's Survey of Primary Dealers](#)." Federal Reserve Bank of New York *Current Issues in Economics and Finance* 19, no. 6: 1–8.
- Crump, Richard, Emanuel Moench, William O'Boyle, Matthew Raskin, Carlo Rosa, and Lisa Stowe. 2014. "[Survey Measures of Expectations for the Policy Rate](#)." Federal Reserve Bank of New York *Liberty Street Economics* (blog), December 5.
- Delavande, Adeline, Jinkook Lee, and Joanne K. Yoong. 2012. *Harmonization of Cross-National Studies of Aging to the Health and Retirement Study: Expectations*. Santa Monica, Calif.: RAND Corporation.
- Delavande, Adeline, Xavier Giné, and David McKenzie. 2011a. "Measuring Subjective Expectations in Developing Countries: A Critical Review and New Evidence." *Journal of Development Economics* 94, no. 2 (March): 151–63.
- Delavande, Adeline, Xavier Giné, and David McKenzie. 2011b. "Eliciting Probabilistic Expectations with Visual Aids in Developing Countries: How Sensitive Are Answers to Variations in Elicitation Design?" *Journal of Applied Econometrics* 26, no. 3 (April–May): 479–97.
- 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.
- Giné, Xavier, and Stefan Klöner. 2007. "Technology Adoption with Uncertain Profits: The Case of Fibre Boats in South India." Mimeo, World Bank.
- Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1994. "Earnings Uncertainty and Precautionary Saving." In Albert Ando, Luigi Guiso, and Ignazio Visco, eds, *Saving and the Accumulation of Wealth: Essays on Italian Household and Government Saving Behavior*. Cambridge: Cambridge University Press.
- Guiso, Luigi, and Giuseppe Parigi. 1999. "Investment and Demand Uncertainty." *Quarterly Journal of Economics* 114, no. 1 (February): 185–227.
- Hurd, Michael. 2009. "Subjective Probabilities in Household Surveys." *Annual Review of Economics* 1 (September): 543–64.
- Leiser, David, and Shelly Drori. 2005. "Naive Understanding of Inflation." *Journal of Socio-Economics* 34, no. 2: 179–98.
- Likert, Rensis. 1932. "A Technique for the Measurement of Attitudes." *Archives of Psychology* 22, no. 140: 1–55.
- Mahajan, Aprajit, Alessandro Tarozzi, Joanne Yoong, and Brian Blackburn. 2008. "Bednets, Information, and Malaria in Orissa." Mimeo, Stanford University.
- Mankiw, Gregory, and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117, no. 4: 1295–328.
- Manski, Charles. 2002. "Identification of Decision Rules in Experiments on Simple Games of Proposal and Response." *European Economic Review* 46, no. 4–5: 880–91.
- Manski, Charles. 2004. "Measuring Expectations." *Econometrica* 72, no. 5 (September): 1329–76.
- McKenzie, David, John Gibson, and Steven Stillman. 2013. "A Land of Milk and Honey with Streets Paved with Gold: Do Emigrants Have Over-Optimistic Expectations about Incomes Abroad?" *Journal of Development Economics* 102: 116–27.
- Potter, Simon. 2011. "[Improving Survey Measures of Inflation Expectations](#)." March 30, 2011, speech at the Forecasters Club of New York.

Potter, Simon. 2012. "[Improving the Measurement of Inflation Expectations](#)." June 7, 2012, speech at the Barclays 16th Global Inflation-Linked Conference, New York.

Robertson, John C., Ellis W. Tallman, and Charles H. Whiteman. 2005. "Forecasting Using Relative Entropy." *Journal of Money, Credit, and Banking* 37, no. 3: 383–401.

Stein, Jeremy C. 2013. "Yield-Oriented Investors and the Monetary Transmission Mechanism." Proceedings of the symposium *Banking, Liquidity, and Monetary Policy*.

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.

Williamson, Maureen R., and Alexander J. Wearing. 1996. "Lay People's Cognitive Models of the Economy." *Journal of Economic Psychology* 17, no. 1 (February):3–38.