

SPEECH

A Review of the Experience of Fielding the Survey of Consumer Expectations

June 24, 2014

James McAndrews, Executive Vice President and Director of Research

Remarks at the Barclays Global Inflation Conference, New York City

As prepared for delivery

Good afternoon. It is a pleasure to be here today to talk with you about some of the research being conducted by the New York Fed on the measurement of household expectations, including households' expectations of inflation. 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. To begin, I'd like to acknowledge some colleagues of mine who have managed our household survey project and conducted much of the research that I'll speak about today—namely Olivier Armantier, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar.

In 2012, my colleague Simon Potter spoke to this conference about the approach that researchers at the New York Fed adopted in rolling out a household survey of expectations. Today, I will report on the survey, which is the result of more than six years of research, including a pilot survey that was conducted for about five years, and the published survey that we ran for a year prior to its first publication in January 2014.

Starting in the fall of 2007, the New York Fed has conducted a long-term research project to improve our understanding of existing measures of consumer expectations, and to explore the feasibility of expanding and potentially improving these measures. We developed and then tested new questions on expectations, designed measures of individual forecast uncertainty, and collected information on expectations and decisions for a broader range of household behavior than is currently gathered by existing surveys. In general, we wanted to understand better the process of forming expectations and updating those expectations, as well as the links between expectations and consumer behavior.

In pursuing this project, we collaborated with the RAND Corporation, other Federal Reserve Banks, academic economists, and psychologists with expertise in survey design. We conducted in-depth cognitive interviews and fielded psychometric survey modules utilizing RAND's American Life Panel, and various convenience samples created by researchers at Carnegie Mellon. Finally, we conducted many experimental surveys as part of the American Life Panel internet survey. Incorporating the lessons we learned from these initiatives led to the final design and implementation of our Survey of Consumer Expectations.

Our goal throughout has been to collect timely and high-quality information on consumer expectations and decisions. We were particularly interested in filling in what we viewed as gaps in existing sources of information on household behavior, including data on household finance, labor, and housing market expectations. In addition, we decided that a rotating panel structure, one that tracks the same individuals over time, was necessary for many research and policy applications. Having a panel reduces the variability induced by changes in composition month to month and increases the precision of any measured changes in expectations.

Before we implemented the survey, our proposal was reviewed and approved by the Bank's senior management. Next, we issued a request for proposals and received three bids from independent survey organizations to field our survey design. In the end, we chose to work with the Demand Institute, a partnership of the Conference Board and Nielsen.

The resulting Survey of Consumer Expectations is a nationally representative monthly internet-based survey. Its rotating panel consists of about 1,200 household heads. We find that the internet approach allows us a great deal of flexibility to ask new questions, and it makes it considerably easier to pose probabilistic questions and to run experiments. It is also—and this is an important consideration—the most cost-effective mode of collection; at the same time, there is also some evidence of higher response accuracy to personally sensitive questions. Our sample is based on the Conference Board's Consumer Confidence Survey sample, which is drawn from U.S. postal addresses. Respondents are compensated with a payment of \$15 per survey, and surveys vary in length from fifteen to thirty minutes. Our first-time response rate is about 60 percent, and we find that respondents have somewhat greater-than-average education and income. The average response rate among repeat respondents is more than 80 percent. We use weights based on the American Community Survey to construct summary statistics for our public releases.

The survey consists of a core monthly module on expectations about macroeconomic and household-level variables and an additional quarterly module that allows us to address special topics. Our core module contains questions about survey respondents' expectations for inflation, wage growth, home prices, various commodities, household income and spending, taxes, government debt, credit access, and job search. So far, quarterly modules have included questions on housing and employment,

work history, job search, retirement, financial literacy, consumption and savings behavior, student loans, and various field experiments. These quarterly modules are developed by Federal Reserve staff in collaboration with academic consultants and the Nielsen team. First, however, new questions are evaluated using cognitive interviews, and every new module is pilot-tested before being put into the field.

The Bank began publishing the results of the Survey in January 2014, after having fielded the Survey since December 2012. We introduced it through a series of posts on our *Liberty Street Economics* blog; eventually, after a suitable embargo period, we plan to release to the public all the micro data from the Survey.

Approach to and Benefits of Eliciting Probabilistic Beliefs

As I mentioned earlier, an important innovation of our Survey is that, in addition to asking the respondent for a point forecast, it uses probabilistic questions to elicit a respondent's density forecast. By doing so, we aim to collect a more accurate and complete representation of individuals' subjective expectations, and the degree of uncertainty that they attach to their forecasts.

Our method for eliciting density forecasts builds on a large body of empirical work, led by economist Charles Manski, which involves the use of probabilistic survey question formats to measure a respondent's beliefs about an uncertain future event. Our specific question format was further tested during the experimental pilot phase of the project. Our density forecast question asks respondents to assign probabilities to pre-determined intervals or bins for future outcomes. For example, as shown on **Slide 2**, for year-ahead inflation, we ask for the "percent chance" that inflation will be between 0 percent and 2 percent, 2 percent and 4 percent, 4 percent and 8 percent, etc., with similar bins used for deflation. In a similar fashion, we solicit respondents' density forecasts for year-ahead earnings growth, house price change expectations, and three-year ahead inflation.

For each individual respondent, the resulting density forecast enables us to construct individual measures of central tendency (such as the density mean or median) and uncertainty (such as the density interquartile range). To compute these, we use each individual's responses to the probabilistic questions to parametrically estimate the underlying forecast density function as proposed by Engelberg, Manski, and Williams in a 2009 *Journal of Business and Economic Statistics* article. More specifically, when a respondent assigns a positive probability to three or more bins, we assume an underlying generalized beta distribution; and for fewer than three bins, we assume a triangular distribution. An example of such a fit is shown on **Slide 3**.

Our findings over the past five years, based on different samples and different survey platforms, indicate that individuals are as willing and able to respond to well-written probabilistic questions as they are to traditional attitudinal questions on the same subject. Moreover, they do so sensibly, with responses to probabilistic questions having internal consistency and measurement reliability. The uncertainty expressed in consumers' density forecasts is reliably related to other measures of uncertainty. For example, individuals who express higher levels of uncertainty in their subjective probability distribution are more likely to report a range when they are originally asked for their point forecast, and the width of this self-reported range is positively correlated with measured uncertainty. In the case of inflation, we find that uncertainty about future inflation is negatively related to self-assessed responsibility for investment decisions, planning horizons for financial decisions, and the respondent's performance on a financial literacy measure.

In addition to providing a measure of uncertainty, an important advantage of density forecasts over point forecasts is that they remove ambiguity about which (if any) measure of central tendency an individual's point forecast corresponds to. Measures of central tendency derived from density forecasts are strongly correlated with point forecasts, and they have very similar average values; however, for roughly half of the responses, the point forecast falls in either the top or bottom quartile of the individual's forecast density, away from the center. Thus, interpersonal comparisons of point forecasts reflect not just differences in beliefs but also differences in what distribution statistic is being reported as point forecast by different survey participants. Density forecasts therefore allow for a more accurate measurement of disagreement among forecasters by using a common measure of central tendency (for example, the mean or the median of individuals' subjective probability distribution).

Findings on Monthly Expectations

The Survey of Consumer Expectations provides a broad range of measures that are used routinely in our continuous assessment of the economy. We publish the monthly findings of our survey on the second Monday of each month on the New York Fed's website.

The Survey page is organized into three sections: inflation, labor market, and household finance. As shown in **Chart 1**, inflation expectations have been quite stable over the past twelve months, at both the one-year and the three-year ahead horizons. Our preferred measure is based on the median of the individual respondents' density means, but we also report the median point forecast as well as the interquartile range of the density means across respondents to highlight the dispersion of expectations across consumers. I will return to this aspect later.

We also report home price change expectations and price change expectations for various commodities. In our latest survey in May 2014, these expectations remained stable for the most part. Interestingly, medical care price change expectations have been

gradually declining since December 2013, dropping from around 11 percent to 9.5 percent. The expected cost of a college education has also been dropping since the start of 2014. These results are shown in **Charts 2 and 3**.

For the labor market, we report the median expected earnings growth over the next twelve months. This figure has been stable at around 2 percent over the sample period, with a small increase to 2.4 percent over the winter that has since been reversed, as seen in **Chart 4**. We also track workers' perceived chances of losing their current job, of leaving their current job voluntarily, and of finding another job in three months should they lose their current one. Quits, layoffs, and the job-finding rate are important flow measures, and workers' expectations about these events enable us to gauge households' perceptions of future labor market conditions. As I will mention later, these expectations seem to be meaningfully related to job search behavior and other economic decisions by households.

Finally, the household finance section contains time series of expected household income growth, expected household spending growth, changes in the amount of taxes that households expect to pay, as well as perceived credit availability—both relative to one year ago and looking ahead to one year from now. As **Chart 5** indicates, household income growth expectations declined slightly to 2.3 percent in May, but remained in the middle of the narrow band (2.0 – 2.6) observed over the last twelve months. On the other hand, perceptions of credit availability continued to improve slightly in May, as shown in **Chart 6**.

For most of our charts, the website gives additional demographic detail, breaking down each time series by income, education, age, numeracy, and region of the country. Further, for our main measures (inflation, home price changes, and earnings growth) we also report the time series of forecast uncertainty. One interesting finding—consistent with our prior research—is that forecast uncertainty about inflation is significantly higher for low education, low income, low numeracy respondents. These findings can be seen in **Chart 7**.

Summary of Past Research

Let me now briefly describe some of the basic research we have carried out through our consumer surveys. We have focused on two broad questions to date. First, we have investigated the connection between inflation expectations and behavior in a financially incentivized investment field experiment. In this study, survey respondents were asked for their inflation expectations. Subsequently, within the course of the same survey, they were asked to choose between two possible financial investments, one that gives a fixed nominal return after twelve months, and another that yields a return indexed by inflation, again after one year. Respondents were told that two survey participants would be drawn at random and be paid according to their choice in the investment after twelve months.

Theory predicts that respondents would be more likely to pick the inflation-protected investment if they expect higher inflation. Indeed, we find a strong association between expectations and actual choices in the experiment. Furthermore, those who express higher uncertainty about future inflation are also more likely to choose the inflation-protected investment. Both patterns are in accordance with the theory. Interestingly, the connection between expectations and experimental choices is weaker for respondents with low education, numeracy, and financial literacy. Our research thus shows that consumer expectations elicited through surveys are correlated in a meaningful way with actual behavior. We also find that expectations inform behavior in other contexts: for instance, workers who express a higher perceived chance of losing their current job over the next twelve months also search harder for a new job and exhibit a drop in spending plans relative to the present over the same time horizon.

The second area of research related to our consumer survey has focused on the formation and updating of inflation expectations. Through another field experiment embedded in our American Life Panel survey modules, we have examined how respondents update their expectations after receiving information about either past food price inflation or future inflation forecasts by professional forecasters. Before the information treatment, and again after the provision of information, respondents were asked for their inflation expectations. We also asked for respondents' priors about the information that was provided, to see whether those who have larger perception gaps are more likely to update their expectations, and if they would do so by a larger amount.

We found that respondents revise their expectations in the direction predicted by theory after receiving new information. Further, the extent of revisions is correlated with the size of their perception gaps. Finally, respondents are more receptive to the new information when the uncertainty expressed in their baseline expectations is greater. All three findings are consistent with rational, Bayesian updating of expectations.

Moving beyond average effects, we also found that our information treatment leads to a significant decline in the cross-sectional dispersion of inflation expectations, and causes the distribution of inflation expectations to converge toward its center—particularly for those who had high baseline expectations and were less informed ex ante. In fact, average revised expectations were nearer to actual realized CPI inflation as a result of our intervention. This is an encouraging result: it suggests that policymakers could partially influence the high-expectation right-tail of the inflation expectations distribution through public information campaigns in the spirit of our information treatments.

Why Heterogeneity for Consumers, and Why Are Expectations High?

An important and well-known distinctive feature of survey inflation expectations of individual consumers and firms—when compared with that of professional forecasters—is the much larger dispersion across individuals’ beliefs about future inflation and the greater right-skewedness, or long right-tail, of this distribution.

A number of different explanations have been proposed for the larger heterogeneity in beliefs, the larger proportion of extreme high responses, and the subsequent higher mean inflation forecast among consumers, summarized in **Slide 11**. First, in our research, we have found a tendency among a subset of respondents to think about specific prices, and especially those most salient to them, in coming up with a forecast. These often tend to be prices of goods that changed the most or are most volatile, such as food and gas prices. As a result, respondents who think about specific prices tend to report higher inflation expectations. Respondents who instead think more about a broader overall measure of inflation tend to report lower forecasts.

Several alternative explanations that have been proposed in the recent literature for the large heterogeneity in beliefs focus on differences in information sets or in expectation-formation processes. These include sticky information models, in which new information is slow to diffuse through the population (Mankiw and Reis, 2002), perhaps because agents only probabilistically pay attention to experts or to news (Carroll, 2003). They also include noisy information models, in which agents form expectations based on noisy private signals (Woodford, 2001). Other explanations are based on the use of some form of adaptive learning (Evans and Honkapohja, 2001), switching between different prediction rules (Branch, 2004); or learning from lifetime inflation experiences (Malmendier and Nagel, 2013; Madeira and Zafar, 2012).

In our information-based experiment that I discussed earlier, we found that cross-sectional disagreement (variance) in expectations falls after our information treatments, which is consistent with a sticky-information model and points to the importance of private information and information constraints. We also found evidence of gender differences in updating behavior, which in turn may reflect differences in gathering and evaluating new information. Numeracy, financial literacy, and education also appear to play a role in explaining dispersion across respondents, as seen in **Chart 8**. In our research, we have found extreme inflation expectations, especially high values, to be associated with lower financial literacy and numeracy skills, and with being relatively under-informed about past inflation or inflation-related facts.

Finally, there exists suggestive evidence of respondents forecasting under asymmetric loss, with respondents appearing to be averse to *under*-estimating inflation. Whereas it is common to assume the mean is being reported—implying that respondents generate forecasts under symmetric square loss—the evidence suggests that some individuals’ forecasts are biased away from the mean, suggesting forecasting using an asymmetric loss function. Heterogeneity in loss aversion would contribute to the cross-sectional heterogeneity of inflation expectations. Moreover, asymmetric loss would generate a mean bias that will vary in size with the uncertainty of individual expectations, and hence is likely nonconstant over time. This evidence points to another potential advantage of density forecasts: their reporting may be less affected by loss aversion than simple point forecasts.

While there are several possible explanations for the large dispersion of inflation forecasts among consumers, this heterogeneity is typically higher than that observed in various surveys of professional forecasters. In closing, I would like to speculate on a few potential reasons as to why the dispersion measured among professional forecasters may be too low relative to a “true” dispersion in expectations.

First, asymmetric incentives may induce a desire to conform. If professional survey respondents have beliefs that are far from the consensus, they might face an asymmetric payoff: a small positive payoff for being correct, but a decidedly negative payoff for being incorrect. This may cause a tendency to converge on a consensus belief. Similarly, in pricing TIPS, a trader with a nonconformist view faces costs in taking that view, so even if ultimately proven correct, the costs of establishing such a position may deter the nonconformist view’s from being reflected in market prices. Second, professional forecasters may be overly exposed to similar “world views,” sharing similar data, reports, and analyses. This may lead to an information aggregation bias of sorts, in which forecasters share the same information set, which is an aggregate of finer information, and which may result in the loss of some important information.

The next charts describe the median four-quarter-ahead CPI forecast and the dispersion around that median for the Blue Chip survey and the U.S. Survey of Professional Forecasters. While there is significant overlap in the two sets of respondents, there is an important difference between the two surveys: respondents are anonymous in the U.S. Survey of Professional Forecasters, whereas they are identified by name in the Blue Chip survey. The asymmetric payoffs described earlier may therefore have more of an impact in the Blue Chip survey than in the U.S. Survey of Professional Forecasters, leading to less dispersion in the former than in the latter.

Chart 9 shows that the median forecast is substantially the same across the two surveys. However, as seen in **Chart 10**, there is some evidence of lower heterogeneity among forecasters in the Blue Chip survey compared with the Survey of Professional Forecasters, especially in recent periods, which, in retrospect, were periods of high fundamental uncertainty in the economy. In both periods, during the run up to the financial crisis and its aftermath, most forecasters were mistaken about future growth rates and inflation rates by relatively large amounts. Both surveys of professional forecasters exhibit significantly less dispersion than

the Survey of Consumer Expectations—even when focusing only on “sophisticated” respondents with high education and numeracy.

This suggests that there may be some scope for the sort of mechanisms I described, which may induce the “fundamental” heterogeneity in beliefs to be suppressed in surveys of professional forecasters. Consequently, household surveys may add important information on “fundamental” heterogeneity of beliefs relative to professional surveys, as the latter may reflect, in part, relativistic payoffs or shared information sets which produce an excessive consensus of views.

Thank you for listening to me today; I would be happy to take some questions on our research on inflation expectations.

References

- Armantier, O., S. Nelson, G. Topa, W. van der Klaauw, and B. Zafar. 2012. “The Price is Right: Updating of Inflation Expectations in a Randomized Price Information Experiment.” Federal Reserve Bank of New York *Staff Reports*, no. 543, revised January 2013.
- Armantier, O., W. Bruine de Bruin, G. Topa, W. van der Klaauw, and B. Zafar. 2011. “Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?” Federal Reserve Bank of New York *Staff Reports*, no. 509.
- Branch, W. A. 2004. “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations.” *The Economic Journal* 114 (July): 592-621.
- Bruine de Bruin, W., W. van der Klaauw, J. S. Downs, B. Fischhoff, G. Topa, and O. Armantier. 2012. “The Effect of Question Wording on Consumers’ Reported Inflation Expectations.” *Journal of Economic Psychology* 33, no. 4 (August): 749-57.
- Bruine de Bruin, W., W. van der Klaauw, and G. Topa. 2011. “Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices.” *Journal of Economic Psychology* 32, no. 5 (October): 834-45.
- Carroll, C. D. 2003. “Macroeconomic Expectations of Households and Professional Forecasters.” *Quarterly Journal of Economics*, 118, no. 1 (February): 269-298.
- Engelberg, J., C. F. Manski, and J. Williams. 2009. “Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters”, *Journal of Business and Economic Statistics* 27, no. 1 (February): 30-41.
- Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton, New Jersey: Princeton University Press.
- Madeira, C., and B. Zafar. 2012. “Heterogeneous Inflation Expectations, Learning, and Market Outcomes.” Federal Reserve Bank of New York *Staff Reports*, no. 536, revised June 2014.
- Malmendier, U., and S. Nagel. 2013. “Learning from Inflation Experiences.” University of California, Berkeley, working paper.
- Mankiw, N. G., and R. Reis. 2002. “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *The Quarterly Journal of Economics*, 117, no. 4, (November): 1295-328.
- Woodford, M. 2001. “Imperfect Common Knowledge and the Effects of Monetary Policy”, National Bureau of Economic Research working paper, no. w8673.