

Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?

Olivier Armantier Wändi Bruine de Bruin Giorgio Topa
Wilbert van der Klaauw Basit Zafar ¹

September 2014

ABSTRACT

We compare the inflation expectations reported by consumers in a survey with their behavior in a financially incentivized investment experiment. The survey is found to be informative in the sense that the beliefs reported by the respondents are correlated with their choices in the experiment. More importantly, we find evidence that most respondents act on their inflation expectations showing patterns consistent with economic theory. Respondents whose behavior cannot be rationalized tend to have lower education, numeracy and financial literacy. These findings help confirm the relevance of inflation expectations surveys, and provide support to the micro-foundations of modern macroeconomic models.

Running Head: Inflation Expectations and Behavior

JEL Codes: C83, C93, D12, D84, E31

Keywords: Inflation Expectations, Surveys, Experimental Economics

¹ We thank Charles Bellemare, Richard Blundell, Michael Bryan, Jeff Dominitz, Julie Downs, Baruch Fischhoff, Jordi Gali, Charlie Holt, Eric Johnson, Arthur Kennickell, Chuck Manski, Rosemarie Nagel, Athanasios Orphanides, Onur Ozgur, Simon Potter, Robert Rich, Justin Wolfers and Ken Wolpin for their advice on this project, as well as Sandy Chien, Tim Colvin, Daniel Forman, Peter Fielding, Daniel Greenwald, Tanya Gutsche, Mandy Holbrook, Scott Nelson and Bas Weerman for their help with conducting the research. Thanks to Nicolas Treich for his help on the theory. We also would like to thank seminar participants at the Federal Reserve Bank of New York, the Federal Reserve Bank of Boston, the University of Salento, the University of Virginia, the University of Maastricht, Science Po. Paris, the Banque de France, the Bank of England, the University College London, the 2010 ESA conference, the 2011 IMEBE conference, the 2011 Florence Workshop on Behavioral and Experimental Economics, the 2011 LeeX International Conference on Theoretical and Experimental Macroeconomics, the Third French Econometric Conference, the 2011 workshop on Recent Advances on the Role of Beliefs on Decision Theory, the 2012 FUR XV International Conference, and the 2012 SED conference. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System

1. Introduction

Expectations, and inflation expectations in particular, are at the center of monetary policy and much of modern macroeconomic theory (Woodford 2005, Galí 2008, Sims 2009). Although the academic debate about expectations formation is still open,² macroeconomic models are generally built on the assumption that agents maximize expected utility under a well defined distribution representing their beliefs. In a nominal world (i.e. with prices denominated in money), any intertemporal decision requires inflation expectations to convert nominal into real returns. As a result, theory predicts that inflation expectations are a key determinant of behavior. In particular, households should take expected inflation into consideration when deciding on large durable purchases, saving, managing debt, mortgage (re)financing, or wage negotiations.

In general equilibrium, these intertemporal decisions in turn affect real economic activity and therefore realized inflation. Because households in aggregate are an important driver of economic activity,³ they play a major role in this transmission mechanism. The importance of this mechanism and the role played by households are now well recognized both in academic and in central banking circles (Bernanke 2004, 2007). It is therefore generally agreed that one of the first steps to controlling inflation consists in actively managing the public's expectations.⁴

Because of the role played by inflation expectations, accurate measurements of the public's beliefs are important to scholars and policy makers. In particular, macroeconomists increasingly

² Since Muth (1961) and Lucas (1972) most macroeconomic models assume full information rational expectations. Over the past 20 years, with mounting empirical evidence rejecting the full information rational expectations hypothesis, several alternatives have been introduced including adaptive learning (Sargent 1999), rational inattention (Sims 2006), sticky information (Mankiw and Reis 2002), noisy information (Woodford 2003), or more generally imperfect information (Coibion and Gorodnichenko 2013). Regardless of how expectations formation is formalized, these models all assume that agents act on their inflation beliefs.

³ For instance, consumer expenditure represented 71% of U.S. GDP in 2011.

⁴ Bernanke (2004) argues that “an essential prerequisite to controlling inflation is controlling inflation expectations.”

use outside estimates of inflation expectations as an input to their models.⁵ Central banks also need accurate measures of inflation expectations to calibrate monetary policy. In addition, to monitor the effectiveness of its communication, a central bank needs to regularly assess the consistency of the public's beliefs with policy objectives.⁶

Existing measures of inflation expectations may be partitioned into two broad categories depending on whether they are direct or indirect. Indirect measures are inferred from either financial instruments (such as TIPS, the Treasury Inflation-Protected Security), the term structure of interest rates, or past realizations of inflation rates. Direct measures are obtained from surveys in which consumers, businesses or professional forecasters are asked to self-report their subjective beliefs about future inflation. In the U.S., such surveys include the monthly Reuters/University of Michigan Survey of Consumers, the Livingston Survey, the Conference Board's Consumer Confidence Survey and the Survey of Professional Forecasters. In addition, because of the role played by households in determining aggregate demand, several central banks around the world are now conducting inflation expectations surveys of consumers.⁷

Each of these direct and indirect measures has potential weaknesses. While market based estimates rely on strong modeling assumptions (e.g. about risk and liquidity premia), surveys of professionals are often based on small samples and may be biased by strategic misreporting.⁸

⁵ Examples include Roberts (1995), Carroll (2003), Mankiw, Reis and Wolfers (2003), Nunes (2010).

⁶ Observe that what academic economists and policy makers need are not necessarily accurate predictors of future inflation, but accurate measures of the public's true beliefs. Of course, the public's true beliefs may be unbiased predictors. However, unbiasedness is not a requirement for inflation expectations to be informative about the economic decisions the public makes.

⁷ Central banks that survey consumers about their inflation expectations include the Bank of England, the European Central Bank, the Bank of Australia, the Bank of Japan, the Sveriges Riksbank, and the Reserve Bank of India.

⁸ For instance, the Survey of Professional Forecasters conducted by the Federal Reserve Bank of Philadelphia currently consists of 45 respondents on average. Moreover, it has been argued that, because of strategic and reputational considerations, professional forecasters may have incentives to misreport their beliefs (Ehrbeck and Waldmann 1996, Ottaviani and Sørensen 2006).

Because of the absence of direct financial incentives, inflation surveys of households may suffer from a “cheap talk” problem leading to noisy and possibly uninformative responses.⁹ Moreover, although inflation is consistently ranked among the top economic concerns in public-opinion polls (Shiller 1997), the impact of inflation expectations on behavior is still not well-understood. In particular, because households may face more significant risks (related to e.g. employment or health), the extent to which they account for future inflation when making decisions is unclear.

In this paper, we examine whether consumers act on the inflation expectations they report in a survey. Because of possible confounds, it is difficult to address this issue with field (i.e. non-experimental) data. For instance, the timing of a durable good purchase could be influenced by the respondent’s inflation expectations, but also by time discounting or liquidity constraints (which may be difficult to measure). Instead, we propose to use controlled experimental methods to isolate the impact of inflation expectations on behavior. More specifically, we compare the behavior of consumers in a financially incentivized investment experiment with the beliefs they self report in an inflation expectations survey. As further explained below, the survey and experiment were fielded twice with the same respondents at a five-month interval. Our analysis may be decomposed in two parts.

First, we evaluate the extent to which inflation expectations surveys of consumers are informative by looking at the correlation between reported expectations and experimental choices. In essence, we are conducting what the survey literature refers to as a “construct validity” exercise, which is a key requirement to validate a survey question (Carmines and Zeller 1991).¹⁰

⁹ For a discussion of this issue see e.g. Keane and Runkle (1990), Manski (2004), or Pesaran and Weale (2006).

¹⁰ For instance, a survey question aimed at eliciting understanding of HIV risks is validated by examining how responses are correlated with risky sexual activities (Bruine de Bruin et al. 2007). An alternative validation approach would be to compare survey responses with inflation expectations elicited with a financially incentivized mechanism

More generally, our validation exercise is relevant to the literature on subjective expectations in which expectations surveys are often used to complement choice data in order to better explain education, labor, retirement or health decisions (see e.g. Wolpin and Gonul 1985, van der Klaauw and Wolpin 2008, or Pesaran and Weale 2006 and van der Klaauw 2012 for reviews). For this approach to be effective, however, a key requirement is that the survey responses truly capture the agent's unobserved beliefs. We propose to test this hypothesis and evaluate the information content of an expectations survey by using controlled experimental methods.

Second, we subject the data to a stricter analysis by examining the extent to which inflation expectations and experimental decisions comply with economic theory. More precisely, we exploit the panel structure of the data to investigate whether or not the survey respondents who report different inflation beliefs in each survey also modify their experimental choices in a way that can be rationalized. In other words, we conduct a simple yet formal test of one of the basic assumptions underlying most modern macroeconomic models. Note that we are not trying to assess how consumers generally take into account expected inflation when making decisions in their everyday life. Instead, our objective is to evaluate whether, when given an explicit opportunity to do so, consumers act on their inflation expectations. Therefore, this study should only be considered a first step in establishing empirically the role played by inflation expectations in shaping economic behavior.

The results indicate that the inflation expectations survey is informative. Indeed, stated beliefs and experimental decisions are found to be highly correlated and consistent, on average, with

such as a “proper scoring rule” (see e.g. Savage 1971). There is no guarantee, however, that such an approach would be valid. Indeed, the respondents' wealth is likely to depend on future inflation, which creates a stake in the event predicted. As shown by (e.g.) Karni and Safra (1995), incentivized beliefs elicitation techniques are only incentive compatible when the respondent has no stake in the event predicted (the so called “no stake” condition).

payoff maximization. Furthermore, we find that respondents who change their inflation expectations from one survey to the next, also tend to adjust their decisions in the experiment in a way consistent (both in direction and magnitude) with economic theory. Finally, our results suggest that respondents whose behavior is difficult to rationalize are less educated and score lower on a numeracy and financial literacy scale. We consider these findings to lend support to the use of survey measures of inflation expectations for monetary policy, and to provide direct empirical evidence on one of the basic assumptions underlying most macroeconomic models.

The remainder of the paper is organized as follows. The survey and the respondents are described in Section 2. The design of the experiment is presented in Section 3. The responses to the inflation expectations questions and the choices made in the experiment are analyzed separately in Section 4. In Section 5, we test whether stated beliefs about future inflation are informative about experimental choices. In Section 6, we exploit the panel structure of the data to study how respondents who change their predictions from one survey to the next adjust their experimental choices. Our final comments are provided in Section 7.

2. The Survey and the Respondents

The survey is part of an ongoing effort by the Federal Reserve Bank of New York, with support from academic economists and psychologists at Carnegie Mellon University. The general goal of this initiative is to better understand how the public forms and updates beliefs about future inflation, and to develop better tools to measure consumers' inflation expectations (Bruine de Bruin et al. 2010a). The survey consists of two sets of questions. The first set, which is analyzed in this paper, examines the link between self-reported beliefs and economic behavior. The

second set, which is analyzed separately in Armantier et al. (2014), investigates how individual consumers revise their inflation expectations after being exposed to new information.

2.1 The Respondents

The survey, which includes the experiment, was conducted over the internet with RAND's American Life Panel (ALP). Our target population consists of individuals 18 or older who participated in the Michigan Survey between November 2006 and July 2010 and subsequently agreed to participate in the ALP.¹¹ Out of a total sample of 972 individuals invited to participate in the survey, 771 did so, for a response rate of 79.3%. Those who completed the first survey were invited to participate in the second survey, of which 734 did so, implying a response rate of 95.2%. The first survey was fielded between July 20, 2010 and August 17, 2010. The second survey was fielded roughly five months later, between January 3rd, 2011 and February 9, 2011. Respondents received \$20 for each completed survey. As explained in the next section, respondents were also eligible to earn extra money if they completed the experiment. Although respondents were allowed to skip questions, those who tried to do so received a prompt encouraging them to provide an answer.

As indicated in Table 1 (column "All Data"), respondents reported an average age of 52.1, with a median of 53. In total, 57% of the respondents were female, 15% had no more than a high school diploma, while 21% possessed a post graduate degree (i.e. beyond a Bachelor degree). The median reported income range was \$60-\$75k, with 42% of the respondents reporting incomes over \$75k. Thus, our respondents are not representative of the U.S. population (e.g. they are

¹¹ The Michigan survey is a monthly telephone survey with 500 respondents, consisting of a representative list assisted random-digit-dial sample of 300, and 200 respondents who were re-interviewed from the random-digit-dial sample surveyed six months earlier. Our target population is further restricted to ALP members who participated in at least one ALP survey within the preceding year, or were recruited into the ALP within the past year.

older, wealthier and more educated). In our analysis linking reported beliefs and experimental choices we therefore control for demographic characteristics. The average and median time taken to complete the survey and experiment were respectively 42 and 26 minutes with no notable differences between survey 1 and 2. There is, however, substantial heterogeneity across respondents. While fewer than 1% of the respondents completed the survey in less than 9 minutes, other respondents took a considerable amount of time between the moments they opened and finished the survey (more than a week for 5% of the respondents).

2.2 Procedure

Both surveys had a similar structure.¹² As explained in more detail below, respondents first reported their expectations for future inflation. Then, they were asked to explain what information they used to form their reported inflation expectations (not analyzed here). The experiment was presented next. After answering questions about how they update their beliefs about future inflation (not analyzed here), respondents completed measures of numeracy, financial literacy, and willingness to take risk.

Three features of the design are worth noting. First, up to the experiment, the respondents were asked the same questions in the same order in both surveys. Second, the inflation expectations questions were asked before the incentivized experiment and the respondents did not know about the experiment before they reach it. Thus, the presence of the experiment should not bias the inflation beliefs reported by respondents. Third, there is a risk that some respondents may see the experiment not as a financially incentivized investment decision, but as a way to double check

¹² The complete list of questions asked in the first survey may be found in the supplemental material available online at <https://sites.google.com/site/olivierarmantier/>.

their response to the point prediction question. If so, their experimental choices would not reflect a utility maximization process, but simply a rationalization ex-post of the inflation prediction they previously reported. To lower this risk, more than 30 questions were asked between the inflation expectations questions and the experiment. Although we cannot rule out the possibility that this form of ex-post rationalization played a role, we find evidence that most respondents exhibited behavior inconsistent with ex-post rationalization in the sense that their experimental choices did not map directly to their inflations expectations (see Section 5.1 for more details).

2.3 Reported Beliefs

Each respondent was randomly assigned to one of two “expectation” treatments. In the “Inflation” treatment, respondents were asked directly about their expectations for the “rate of inflation.” In the “Price” treatment, respondents were asked about their expectations for the “prices of the things I usually spend money on.” In both treatments and both surveys point predictions were elicited for the same time horizon: between now and 12 months from now. In addition to point estimates, the respondents in both expectation treatments were asked to report probabilistic beliefs for a range of inflation outcomes. More specifically, respondents were asked to state the percent chance that, over the next 12 months, the “rate of inflation” or “changes in prices” would be within the following intervals:]-12% or less, [-12%,-8%], [-8%,-4%],[-4%,-2%], [-2%,0%], [0%,2%], [2%,4%], [4%,8%], [8%,12%], [12% or more. Respondents could press a button to see the sum of the probabilities entered so far in order to verify that their answers added to 100%. If it was not the case, they were prompted to go back and make the appropriate changes.

Following Engelberg, Manski and Williams (2009), a generalized beta distribution is fitted to each respondent’s stated probabilistic beliefs. We then generate two variables that will be used in

the econometric analysis. The first, the “Estimated Expected Prediction,” is the mean of the respondent’s beta distribution. The second, equal to the variance of the respondent’s fitted distribution, is assumed to capture the respondent’s “inflation uncertainty.”

2.4 Numeracy, Financial Literacy and Self-Reported Risk Tolerance

Six questions were asked in the first survey to measure the respondent’s numeracy and financial literacy. The numeracy questions were drawn from Lipkus, Samsa, and Rimer (2001), while the questions about financial literacy were slightly adapted from Lusardi (2007).¹³ We created a variable taking integer values between zero and six depending on the number of correct answers the respondent gave to these questions. As indicated in Table 1 (column “All Data”) respondents answered 4.5 questions correctly on average with a median of 5. There is, however, some heterogeneity across respondents: While 29.8% answered every question correctly, 11.0% got less than half of the answers right.

In each survey, respondents were also asked to assess their willingness to take risk regarding financial matters using a qualitative scale ranging from 1 (Not willing at all) to 7 (very willing). This instrument has been shown to produce meaningful measures of risk preferences. In particular, Dohmen et al. (2011) find that the risk tolerance self-reported on this qualitative scale is consistent with the risk preference elicited with a financially incentivized lottery-type experiment developed by Holt and Laury (2002). Other studies using this measure of risk attitude include Bonin et al. (2007), and Caliendo, Fossen and Kritikos (2010). As indicated in Table 1 (column “All Data”) the average reported risk tolerance across the two surveys is 3.3 with a

¹³ Here is an illustration of the type of questions we asked: “If you have \$100 in a savings account, the interest rate is 10% per year and you never withdraw money or interest payments, how much will you have in the account after two years?”. The other numeracy and financial literacy questions may be found in the supplemental material.

median of 3. One third of the respondents selected a rating of 1 or 2, thereby reflecting substantial aversion to risk. In contrast, one respondent out of four indicated a high tolerance toward risk by selecting a rating of 5, 6 or 7. This distribution is generally consistent with those obtained in previous work using the same measure. Furthermore, our risk attitude measure appears to be generally stable over time. Indeed, we find a correlation of 0.822 between the risk tolerances reported by the same respondents across the two surveys.¹⁴

3. The Financially Incentivized Experiment

3.1 Experimental Design

As shown in Appendix A where the experimental instructions are reported, the experiment consists of 10 questions with real monetary consequences.¹⁵ For each question the respondent is asked to choose between two investments. Each investment produces a specific revenue payable 12 months later.¹⁶ Investment B produces a fixed dollar amount while investment A is indexed on future inflation. More specifically, the respondent's earnings under investment A depend on what realized inflation will be over the next 12 months. The possible earnings under investment A as a function of realized inflation were presented to the respondents as in Table 2, where the "rate of inflation" was defined as the annual U.S. CPI rounded to the nearest percentage point.¹⁷

¹⁴ Reporting different risk tolerance in each survey is not necessarily a violation of theory. Indeed, the qualitative measure reflects both the respondents risk preference and the nature of the risks they face. While standard theory assumes that the former is stable, the latter may have evolved in the five months separating the two surveys.

¹⁵ The Appendices are available online at <https://sites.google.com/site/olivierarmantier/> .

¹⁶ Observe that, regardless of the investment selected, a respondent can only receive money 12 months after the experiment. As a result, time preference cannot influence the choice between the two investments.

¹⁷ Although related to studies on portfolio choices (e.g. Dominitz and Manski 2007), our experiment has a distinctive feature as it focuses on uncertain returns linked to future inflation prospects.

As indicated in Appendix A, investment A remains the same in each of the 10 questions. In contrast, the revenue produced by investment B varies across questions. We conducted two treatments by changing the order in which investment B was presented to respondents. In the “Ascending scale” treatment, the earnings of investment B increase in increments of \$50 from \$100 in question 1 to \$550 in question 10. In the “Descending scale” treatment the earnings of investment B decrease in increments of \$50 from \$550 in question 1 to \$100 in question 10. To simplify, we only refer to the “Ascending” treatment for the remainder of Section 3.

With respect to predicted behavior, observe in Appendix A that an expected payoff maximizer with an inflation expectation within $[0\%,9\%]$, say 5%, should first select investment A for the first 4 or 5 questions (the respondent is indifferent between the two investments in question 5 as they both produce \$300 in expectation), and then switch to investment B for the remaining 5 questions. Likewise, it is easy to see that, regardless of risk attitude, an expected utility maximizer should switch investments at most once and in a specific direction (i.e. from investment A to B). The analysis conducted in the next section therefore focuses on a respondent’s “switching point.” We only define this switching point for respondents whose behavior may be rationalized, that is for respondents who switch at most once from investment A to investment B. For these respondents, the switching point is set equal to the number of questions for which the respondent selected investment A. So, for both experimental treatments, the switching point can take integer values between 0 (the respondent always selects investment B) and 10 (the respondent always selects investment A).¹⁸

¹⁸ Observe that, although the question addressed is different, the structure of our experiment is akin to the experiment of Holt and Laury (2002) in which risk attitude is measured by the number of questions after which a respondent switches from a safer to a riskier lottery. Unlike Holt and Laury (2002), however, the switching point does not have a direct interpretation here.

The participants were informed that two respondents would be paid in each of the two surveys according to their choices in the experiment. Once a survey was completed, we randomly picked one of the ten questions, and two survey participants who completed the experiment. Twelve months later, these two participants were paid according to the investment choice they made for the selected question. Although the amounts a respondent could earn were substantial compared to traditional lab experiments (i.e. up to \$600), the respondents were not able to calculate exactly their odds of being selected for payment since the exact number of participants was unknown at the time the experiment was conducted. Note also that motivating subjects by paying a few of them large amounts with small probabilities is a method used in several lab and field experiments (e.g. Harrison, Lau and Williams 2002, Dohmen et al. 2011).

Finally, each respondent was randomly assigned to one of the four possible treatment combinations (i.e. either the “Price” or the “Inflation” treatment, and either the “Ascending” or the “Descending” treatment). Once assigned to a treatment, a respondent remains in the same treatment in the two surveys.

3.2 Economic Considerations

Although presented in terms of terminal payoffs to facilitate the respondents’ comprehension, investments A and B both have an economic interpretation. Investment A corresponds to the following scenario: “an agent borrows \$5,000 for 12 months at a rate equal to the inflation rate, and invests the \$5,000 for 12 months in an account that earns a fixed annual rate of 11%.” Investment B corresponds to the following scenario: “an agent borrows \$5,000 for 12 months at a rate equal to the inflation rate, and invests the \$5,000 for 12 months in an inflation protected

account that earns an annual rate equal to the inflation rate plus k %, where k varies in increments of 1% from 2% in question 1 to 11% in question 10. ”

In nominal terms, investment B earns $\$5,000 * k$ while investment A earns $\$5,000 * (0.11 - i)$, where i denotes the inflation rate over the next 12 months. If expressed in real terms, investment B earns $X = \$5,000 * k / (1 + i)$, while investment A earns $\$5,000 * (0.11 - i) / (1 + i) = \alpha X - \beta$, where $\alpha = 1.11/k > 1$ and $\beta = \$5,000$. It is then easy to see that the variance of the earnings with respect to inflation is always lower with investment B whether one expresses earnings in nominal or in real terms. We can then derive three propositions that will help us assess whether the behavior observed in the experiment is consistent with expected utility. To do so we consider an expected utility framework and we assume throughout the paper that the agent has a utility function defined over experimental income, $U(\cdot)$, which is invariant over time, thrice differentiable, strictly increasing, and satisfies the von Neumann Morgenstern axioms.¹⁹

Proposition 1: *If investment A and investment B have the same expected return then a risk-averse agent prefers investment B to investment A.*²⁰

The proposition therefore shows that, all else equal, and in particular when the distributions of beliefs are identical, a risk-averse (respectively risk-loving) agent has a lower (respectively, higher) switching point than a risk neutral agent. For instance, consider a respondent who expects the inflation rate to be 4% over the next 12 months. In question 6 the two investments produce the same expected return of \$350. If this respondent is risk averse (respectively risk

¹⁹ It is unclear whether the propositions remain valid when expected utility is relaxed. As a result, we will only be able to test whether a respondent violates expected utility, not whether the behavior of violators could be rationalized under an alternative non-expected utility model.

²⁰ The proofs of all the propositions are provided in Appendix B.

loving) then he should select the safer (respectively riskier) option in question 6, that is, he should select investment B (respectively investment A).

We now generalize Proposition 1 by showing that, all else equal, and in particular when the belief distributions are identical, a more risk-averse agent has a lower switching point.

Proposition 2: *If a risk averse agent is indifferent between investment A and investment B then, all else equal, a more risk averse agent (in the classical sense of Pratt 1964) prefers investment B to investment A.*

Proposition 2 therefore allows us to rationalize differences in behavior observed in the experiment. Consider for instance two agents who share the same beliefs about future inflation and report an expectation of 4%. Furthermore, assume the first agent selects a switching point of 5 while the second agent selects a switching point of 3. Under expected utility, we can rationalize this difference in behavior by a difference in risk aversion, whereby the second agent is more risk averse than the first.

For Proposition 3, we restrict our attention to HARA (hyperbolic absolute risk aversion) utility functions. This class of utility functions is quite general as it encompasses CRRA (constant relative risk aversion) and CARA (constant absolute risk aversion) utility functions. In addition, virtually all the utility functions used in practice (e.g. exponential, logarithmic, power) belong to the HARA family.

Proposition 3: *If a risk-averse agent with a HARA utility function is indifferent between investment A and investment B, then the agent prefers investment B to investment A for any increase in risk (in the classical sense of Rothschild and Stiglitz 1970).*

Proposition 3 therefore shows that if a risk-averse agent is indifferent between the two investments for a given belief distribution, then the agent should strictly prefer the safer investment (i.e. investment B) for any variance increasing (but mean preserving) change in his belief distribution. In other words, all else equal, a risk averse agent should switch from investment A to investment B earlier when the risk associated with investment A increases.

4. Responses to the Survey

Out of the 771 respondents who answered at least one of the two surveys, a total of 81 respondents (23 in survey 1 only, 24 in survey 2 only, and 34 in both surveys) failed to report a point prediction and/or to provide an answer to all the 10 questions in the experiment. Out of the 688 (respectively 676) remaining respondents in survey 1 (respectively survey 2), 598 or 86.9% (respectively 615 or 91.0%) had rationalizable answers as they switched at most once from investment A to investment B during the course of the 10 questions. Furthermore, 82.8% (1,004 out of 1,213) of the total number of rationalizable answers are due to the same 502 repeat respondents who provided rationalizable answers to both surveys. Note that the ratio of non-rationalizable responses (11.1% across the two surveys) is lower than those typically obtained in the literature on measuring risk attitude using the Holt and Laury's instrument. In particular, Holt and Laury (2002) and Eckel and Wilson (2004) report that respectively 25% and 15% of their respondents made non-rational choices.

Perhaps not surprisingly, the respondents with missing data and multiple switching points have specific characteristics. Indeed, a comparison of the columns "Group 1" and "Group 2" with the columns "Group 3" and "Group 4" in Table 1 indicates that the first two groups score significantly lower on our scale of numeracy and financial literacy, are more likely to be a

female, they have lower income and lower education, and they provide higher and more volatile point predictions. As indicated in the last column of Table 1, however, a probit regression in which the dependent variable is equal to 1 when a respondent provides rationalizable answers in both surveys reveals that only the measure of numeracy and financial literacy remains significant after controlling for other respondent's characteristics. More specifically, we find that a respondent is 6% more likely to make rationalizable choices in both surveys for each additional numeracy and financial literacy question he answers correctly. As we shall see later on, low numeracy and low financial literacy are characteristics shared not only by respondents with non-rationalizable choices, but also by respondents whose choices, although rationalizable (in the sense that they switch at most once from investment A to B), are not consistent with theory.

Before examining the possible link between the respondents' inflation expectations and their behavior in the experiment, we look separately at the responses to the inflation expectations questions and the choices made in the experiment.

4.1 Responses to the Point Predictions Questions

In Figure 1, we plot for each survey the distribution of the inflation point predictions over the next 12 months combined across the "Inflation" and the "Price" treatments. As we can see, both distributions have similar shapes with the same mode (2% to 4%), the same median (3%) and the same interquartile range (3%). Observe, however, that the distribution of point predictions for both treatments shifts to the right in survey 2, thereby indicating an increase in inflation expectations between the five months that separate the two surveys. A Wilcoxon signed-rank test confirms that the average point prediction in survey 1 (4.1%) is significantly lower (P-value=0.034) than the average point prediction in survey 2 (4.8%). Finally, note that the

distributions of point predictions are similar to those obtained in other inflation expectation surveys we conducted during the same period of time with different but comparable respondents (Bruine de Bruin et al. 2010b). Thus, we find no evidence that the inflation expectations reported in each survey were affected by the presence of the subsequent experiment. This result is not surprising since respondents only learned of the experiment after they reported their beliefs.

We plot in Figure 2 the distribution of the individual differences in point predictions across the two surveys. To do so we calculate for each of the 502 repeat respondents with rationalizable answers in both surveys the difference between the point prediction she/he made in survey 2 and the point prediction she/he made in survey 1. Although the mode of the distribution is centered on zero, the majority of respondents reported different predictions in the two surveys. More precisely, between the five months that separate the two surveys, 74% of the respondents revised their point prediction by at least $\pm 0.5\%$, and 27% by more than $\pm 4.0\%$.²¹ Finally, note that, consistent with the increase in average point prediction observed in Figure 1, the distribution reported in Figure 2 has more weight in the positive domain. In fact, only 28% of the respondents revised their point prediction downward in survey 2.

A year after we conducted each survey, i.e. in August 2011 and February 2012, the *Bureau of Labor Statistics* reported that the CPI over the past 12 months was respectively 3.8% and 2.9%. Our respondents therefore made relatively accurate forecasts in the first survey as their average point prediction (4.1%) is not significantly different from realized inflation. In contrast, with an average point prediction of 4.8%, respondents significantly overestimated inflation in the second

²¹ In Armantier et al. (2014) we find that respondents update their inflation expectations in a way consistent in direction (but not necessarily in magnitude) with Bayesian learning. Further, we find systematic differences across demographic groups. In particular, women, low income, less educated, lower financial literacy, and older respondents appear to be more responsive to new information.

survey and incorrectly predicted an increase in inflation. Similarly, we find that the respondents were relatively well (respectively poorly) calibrated in the first (respectively second) survey as the realization of the CPI fell within the inter-quartile range (calculated from the distribution fitted to each respondent's probabilistic beliefs) of 78.3% (respectively 47.5%) of our respondents. The two surveys therefore provide mixed evidence about the ability of consumers to forecast accurately future inflation. Without a longer panel, however, our results should be considered anecdotal evidence. Furthermore, recall that the focus of this paper is on the link between reported beliefs and economic behavior, not on the forecasting ability of consumers.

We now explore whether the responses to the point prediction question are affected by the “price” versus “inflation” treatment to which a respondent is assigned. Recall that roughly half of the respondents are asked about their expectations for the “prices of things I usually spend money on”, while the other half is asked about the “rate of inflation”. We plot in Figure 3 the distribution of responses for the two expectation treatments. In both surveys, the different distributions exhibit a similar pattern. Note, however, that consistent with previous studies we conducted (e.g. Bruine de Bruin et al. 2010b), the question about the “prices of things I usually spend money on” yields higher average predictions than the question about the “rate of inflation” question. More specifically, the average point prediction for the “prices of things I usually spend money on” is 4.29% in survey 1 and 4.91% in survey 2, while the average point prediction for the “rate of inflation” question is 3.80% in survey 1 and 4.58% in survey 2. These differences, however, are well within one standard deviation (between 5% and 6% across expectation treatments) and a Mann-Whitney test fails to identify a significant difference between expectation treatments (P-value equals 0.18 in survey 1 and 0.25 in survey 2).

Having looked at point predictions, we now conclude this section with a brief analysis of probabilistic beliefs. A respondent's point prediction is generally consistent with his probabilistic beliefs. In particular, most point predictions (80.6% in survey 1 and 84.4% in survey 2) fall within the non-parametric bounds proposed by Engelberg et al. (2009) for the mean of the probabilistic beliefs. Further, there is a high correlation (0.759 in survey 1 and 0.780 in survey 2) between a respondent's point prediction and his "estimated expected prediction" (the mean of the respondent's fitted distribution). The tight connection between the two variables may also be appreciated in Appendix C (Figures 1.1 to 1.3) as we find little differences when we replicate Figures 1 to 3 using the estimated expected predictions instead of the point predictions. As to the variance of a respondent's fitted distribution, the "inflation uncertainty," observe in Figures 1.2 and 2.2 in Appendix C that its distribution shifted slightly to the right between the two surveys. In other words, it seems that there was both an increase in inflation expectations and an increase in inflation uncertainty between the five months that separate the two surveys.

4.2 Choices Made in the Experiment

The distribution of switching points in each survey is plotted in Figure 4. We can see that although respondents make use of all possible switching points, most choices (61% in survey 1 and 55% in survey 2) are concentrated between 4 and 7. Note also that the distribution of switching points shifts to the left in survey 2. A Wilcoxon signed-rank test confirms that this difference across surveys is in fact highly significant (P-value = 6.5E-4). As we shall see next, this shift toward lower switching points is consistent with the fact that respondents reported higher point predictions in the second survey.

We plot in Figure 5 the distribution of individual differences in switching points for the 502 respondents who made rationalizable experimental choices in both surveys. As with the point predictions, we can see that 78% of the respondents chose a different switching point in each survey. Furthermore, note that 48% of the 502 respondents who made rationalizable choices in both surveys selected a strictly lower switching point in survey 2. In the next section, we will therefore be able to exploit the fact that most respondents change their predictions and switching points across the two surveys to test whether the direction and the magnitude of those changes are consistent with expected utility theory.

Finally, we explore whether the choice of switching point is influenced by the treatment combination to which a respondent is assigned. Recall that our sample is segmented in four groups depending on which expectation treatment (i.e. “Price” or “Inflation”) and which experimental treatment (i.e. “Ascending” or “Descending”) a respondent is assigned to. We plot in Figure 6 the distribution of switching points for each treatment combination. None of these distributions seems to exhibit a distinguishable pattern. This absence of treatment effect is confirmed by a series of Mann-Whitney tests (the P-values range from 0.21 to 0.77).

5. The Link between Beliefs and Behavior

5.1 Are Point Predictions Informative about Experimental Choices?

We now turn our attention to the correlation between the respondents’ point predictions and their switching points. In Figure 7, we plot for each switching point between 0 and 10 the average point prediction across the respondents who selected that switching point. For instance, we can see that the respondents who always selected investment B, and who therefore have a switching

point equal to 0, reported an average point prediction of 9.3% in survey 1 and 10.2% in survey 2. Observe first in Figure 7 that there is a generally monotonic decreasing relationship between the reported beliefs and the switching points. Furthermore, note that this relationship is very similar in both treatments. This result therefore supports the hypothesis that our inflation expectations survey is informative, in the sense that the beliefs the respondents reported correlate well, on average, with their choices in the financially incentivized experiment.

We also plot in Figure 7 a risk-neutral band indicating the range of beliefs that would rationalize each switching point under risk neutrality. For instance, if a risk-neutral agent selects a switching point equal to 5, then his point prediction should belong to the interval [3.5%, 5.5%]. As shown in Proposition 1, switching points below (respectively, above) the risk-neutral band may be rationalized under risk aversion (respectively, risk loving). We can see in Figure 7 that, on average, respondents exhibited behavior consistent with risk neutrality, although a number of switching point and average prediction combinations are close to the risk averse frontier.

This does not imply, however, that respondents systematically behaved as if risk-neutral. In fact, the box plot in Figure 8 reveals that most respondents are outside the risk-neutral band. More precisely, we find that in survey 1 (respectively in survey 2) 41%, 32% and 27% (respectively 37%, 41% and 22%) of the respondents behaved as if risk averse, risk neutral, and risk loving. This finding provides some evidence against the hypothesis that most respondents simply chose a switching point as a way to rationalize ex-post the inflation prediction they previously reported. Under this hypothesis, a respondent uses Table 2 to identify the switching point that corresponds to his inflation forecast (that is, the respondent finds his point prediction in Table 2 and switches investment when reaching the question with the same dollar amount). Ex-post rationalization

would therefore produce switching points exclusively within the risk-neutral band which is not what we see in the data. Furthermore, as shown in the next section, we find evidence of expected utility maximization for most respondents within the risk neutral band as their experimental choices are consistent with both their point predictions and their reported risk attitudes.²²

The general trends observed in Figure 7 seem to be robust. In particular, instead of the average point prediction, we plot in Figures 7.1 and 7.2 (reported in Appendix C) the median point prediction in one case, and the “estimated expected prediction” (calculated with the respondents’ reported probabilistic beliefs) in the other case. Although slightly flatter, these two additional figures display a similar relationship between point predictions and switching points. In Figures 7.3 to 7.6 (reported in Appendix C) we reproduce Figure 7 with the data collected in each of the four treatment combinations. Although the average point predictions are somewhat more volatile across switching points than in Figure 7 (as can be expected given the reduction in sample sizes) the general trend does not vary substantially across treatment combinations.

To confirm these observations statistically, we estimate a series of ordered probit models in which the dependent variable is the respondents’ switching points. Table 3 shows the results of these estimations for each survey. In Model 1, the parameter associated with the variable “Point prediction” is highly significant and negative.²³ This result therefore confirms that the respondents’ reported beliefs are informative about their decisions in the incentivized experiment.

We also find that the parameter associated with the self-reported measure of risk attitude is

²² Of course, a more sophisticated form of ex-post rationalization in which a respondent misreports his risk attitude at the end of the survey simply to reconcile his stated inflation expectation with his experiment choice is observationally indistinguishable from expected utility behavior and thus cannot be ruled out.

²³ We do not report the marginal effects from the ordered probit regressions because, unlike binary probit models, they are not directly interpretable. Instead, we report in Appendix D the outcome of simple linear regressions. These additional regressions not only confirm the robustness of the results presented in this section, but they also provide a sense of the relative effect of each explanatory variable.

positive and significant. In other words, consistent with Proposition 2, respondents who report being more risk-averse tend to select lower switching points, while respondents who report being more risk-loving have higher switching points. Furthermore, the parameter associated with inflation uncertainty is significant and negative. Respondents with more diffuse beliefs, therefore, tend to switch investment earlier. According to Proposition 3, this result may be rationalized under expected utility if respondents exhibit risk aversion (which is the case for many respondents). Finally, note that none of the treatment dummies is significantly different from zero in Model 1, thereby supporting the absence of treatment effects.

To confirm the robustness of the results, we estimated several additional specifications. In Model 2 of Table 3, we augment the specification by including demographic variables. Observe that the parameters estimated in Model 1 remain essentially unchanged. Once we control for the variables in Model 1, we find that none of the demographic variables plays a role in explaining when a respondent switches from investment A to investment B. In Model 3 of Table 3, we replace the point prediction by the “estimated expected prediction” (i.e. the mean of the Beta distribution fitted to the respondent reported probabilistic beliefs). Once again the parameters previously estimated remain essentially unchanged. The parameter associated with the “estimated expected prediction” in Model 3 is similar, both in sign and magnitude, to the parameter associated with the point prediction variable in Model 1. In fact, a log-likelihood ratio test reveals that the two parameters are statistically indistinguishable at the usual significance levels (the P-value is 0.182 in survey 1 and 0.354 in survey 2).

Finally, three of the four treatment combination dummies were interacted with the point prediction in Model 4. As indicated in Table 3, none of the corresponding parameters is found to

be significantly different from zero at the usual significance levels. In other words, we do not find statistical evidence that the slope of the relationship between the point predictions and the switching points varies significantly across treatments. This absence of treatment effect may seem somewhat surprising. Indeed, the payments in the experiment depend on a measure of inflation (the CPI), while respondents in the “price” treatment are asked to make a prediction about the “prices of things I usually spend money on.” As a result, one could have expected a weaker relationship between point predictions and experimental choices in the “price” treatment than in the “inflation” treatment (where respondents are asked to state their beliefs about inflation). Note, however, that this lack of treatment effect is consistent with the fact that we failed to identify significant differences between the point predictions (respectively the switching points) reported in the “inflation” and “prices” treatments.²⁴

5.2 Deviations from Risk Neutrality and Reported Risk Tolerance

Although the estimates in Table 3 confirm that reported beliefs and experimental decisions are correlated, they fail to assess the extent to which choices in each experiment comply with standard theory. Under expected utility, the switching point of a respondent can deviate from risk neutrality only because of his risk attitude and subjective beliefs about the investment’s riskiness, as determined by the shape of his inflation distribution.

²⁴ We also estimated several alternative models not reported here. For instance, we estimated a model in which i) a “price” and “ascending” dummy variable were introduced instead of interactions between the “expectation” and “experimental” treatment variables, ii) the reported risk aversion is defined as a set of seven dummy variables, iii) respondents with extreme switching points (i.e. 0 or 10) are excluded, iv) the reported risk attitude was interacted with other explanatory variable (e.g. the estimated variance), v) the time taken to complete the survey enters non-linearly in order to test the hypothesis that the respondents with the shortest completion time did not answer the survey with as much attention as other respondents, vi) the sample is limited to the 502 respondents that made rationalizable choices in both surveys. These models produced similar conclusions to the ones we present here.

To test this hypothesis, we calculate each respondent's pair of "risk neutral switching points," that is, the pair of switching points the respondent should have selected if risk neutral given his stated point prediction. For instance, consider a respondent in the ascending scale treatment who made a point prediction of 5%. If risk neutral, this respondent should switch to investment B after selecting investment A in the first 5 or 6 questions (the respondent is indifferent between the two investments in question 5 as they both earn \$400). In that case, the respondent's pair of "risk neutral switching points" is $\{5,6\}$. We then define the respondent's "deviation from risk neutrality" as the difference between the respondent's actual switching point and his pair of risk neutral switching points. For instance, if in the example above the respondent selected a switching point of 4 (respectively 8), then her deviation from risk neutrality is $-1 = \text{Max}\{4-5, 4-6\}$ (respectively $2 = \text{Min}\{8-5, 8-6\}$). The deviation from risk neutrality is therefore an integer between -10 and 10.

Using Proposition 2, we can then classify respondents depending on their "revealed risk attitude." In what follows, we will refer to a respondent as exhibiting risk aversion, risk loving, or risk neutrality when his deviation from risk neutrality is respectively strictly negative, strictly positive or equal to zero. Said differently, a respondent exhibits risk aversion, risk loving, or risk neutrality when his pair of point prediction and experimental choice is located below, above or within the risk neutral band in Figure 7. Note that a respondent's revealed risk attitude should not be confused with his reported risk tolerance. The former is fully characterized by the respondent's point prediction and experimental choice, while the latter is the respondent's self assessment in the survey about his general willingness to take risk in financial matters.

To get a complete picture of the factors that influence the deviations from the risk neutral choices, we report in Table 4 the estimates produced by four ordered probit models. The dependent variable is a respondent's "deviation from risk neutrality" in Model 1, and the absolute value of this deviation in Model 2. In the last two models, the dependent variable is a respondent's deviation from risk neutrality, but we restrict the sample to respondents with revealed risk aversion in Model 3, and to respondents with revealed risk loving in Model 4. The last two models will help us gauge whether the same factors influence respondents who exhibit risk aversion and risk loving. Based on the theory, the hypothesis to be tested with these models is that the reported risk tolerance and the inflation uncertainty are the only two explanatory variables with statistical power to explain the deviation from risk neutrality.

The results reported in Table 4 (Model 1) reveal a strong positive relationship between a respondent's reported risk tolerance and his/her deviation from risk neutrality. As indicated in Model 3 and 4, this effect applies both to respondents that exhibit risk aversion and risk loving. In other words, we find that, all else equal, respondents who report being more risk averse (respectively, risk loving) select lower (respectively, higher) switching points in the experiment thereby revealing higher levels of risk aversion (respectively, risk loving). This result is confirmed in Model 2, as the effect of reported risk tolerance on the absolute deviation from risk neutrality is now identified as being U shaped. Consistent with our hypothesis, we therefore find that reported risk tolerance plays a role in explaining why a respondent did not behave as if risk neutral in the experiment.

Inflation uncertainty does not appear to affect choices in Model 1. This result, however, is misleading. Indeed, observe in Model 2 that inflation uncertainty has a significant positive impact on the absolute deviation from risk neutrality. The difference between the results obtained

in Model 1 and Model 2 may be explained by the fact that inflation uncertainty has an opposite effect on respondents depending on whether their experimental choices reveal risk aversion or risk loving (see Model 3 and 4). Indeed, in line with Proposition 3, we find that, all else equal, respondents who exhibit risk aversion (respectively, risk loving) have a lower (respectively, higher) switching point when their inflation uncertainty is higher. Consistent with our hypothesis, we therefore find that deviations from risk neutrality may be explained by the respondents' reported risk tolerance and their subjective uncertainty about future inflation.

In contrast with our hypothesis, however, other variables also have explanatory power. In particular, observe in the last three models in Table 4 that the parameters associated with the respondent's education level, as well as numeracy and financial literacy are highly significant. Furthermore, the results from Model 3 and 4 indicate that these variables have an opposite effect on respondents who behave as if risk averse and those who behave as if risk loving. In other words, it seems that all else equal, and in particular after controlling for reported risk attitudes, the experimental choices of respondents with lower numeracy, financial literacy, and education level tend to deviate further away from risk neutral behavior. This finding has several possible interpretations. Respondents with lower numeracy, financial literacy or education i) may be more likely to make optimization errors when they try to select an investment consistent with their beliefs, ii) they may not fully understand the experiment (although their experimental choices are rationalizable), iii) they may report point predictions (respectively risk tolerance) in the survey that are not a true reflection of their inflations beliefs (respectively risk attitude), or iv) they are not expected utility maximizers.

Finally, we find that all else equal, and in particular after controlling for reported risk attitude and inflation uncertainty, respondents who took longer to complete the survey have a behavior in the experiment that is more likely to deviate from that of a risk neutral optimizer. This result also has several possible interpretations. It could reflect the fact that, when these respondents finally reach the experiment, the predictions they made at the beginning of the survey do not accurately represent their beliefs anymore (some respondents took more than 10 days between the time they started and completed the survey and experiment). Alternatively, it is possible that respondents who took a longer to complete the survey did not make predictions and choices in the experiment with the same attention as the other respondents. Finally, it could reflect a lack of comprehension, whereby respondents who took longer to complete the survey may have needed the extra time because it was more difficult for them to understand and answer the questions.

6. Changes in Predictions and Changes in Experimental Choices across Surveys

We now exploit the panel structure of the data to look at the predictions and experimental choices of a given respondent across the two surveys. To do so, we concentrate on the 502 repeat respondents who made rationalizable choices in both surveys. Our general objective in this section is to address the following question: When a respondent reports different inflation expectations in survey 2 than in survey 1, does this respondent also adjust his experimental choices in a way consistent with theory? We decompose the analysis in two parts: First we look at whether the *direction* of the adjustments across surveys may be rationalized; then, we study whether the *magnitude* of the adjustments is consistent with a simple expected utility model.

6.1 Is the Direction of Adjustments Consistent with Theory?

6.1.1 Changes in Revealed Risk Attitudes

The first test we conduct is based on the assumptions made in Section 3.2 that an agent's utility is defined over income and invariant over time. In that case, an expected utility maximizer should not act as if risk averse in one survey and as if risk loving in the other survey.²⁵ If a respondent exhibits such behavior, then we say that the respondent's revealed risk attitudes are inconsistent across the two surveys. In contrast, switches from revealed risk neutrality to revealed risk aversion or risk loving will not be considered inconsistent with theory. Indeed, because experimental choices are discrete, we cannot exclude the possibility that respondents who make decisions consistent with risk neutrality are in fact slightly risk averse or risk loving.

We report in Figure 9 the distributions of revealed risk attitudes in both surveys. In more than 50% of the cases a respondent with a specific revealed risk attitude in survey 1 (i.e. risk averse, risk loving or risk neutral) exhibits the same revealed risk attitude in survey 2. Furthermore, note that only 6% of the repeat respondents have inconsistent revealed risk attitudes. More precisely, 16 out of the 214 (respectively 14 out of the 123) respondents who made experimental choices consistent with risk aversion (respectively risk loving) in survey 1, behaved as if risk loving (respectively risk averse) in survey 2.

To identify whether these respondents have specific characteristics, we estimate a probit model in which the dependent variable is equal to 1 when a respondent's revealed risk attitudes are

²⁵ The claim does not necessarily hold if the respondent's i) utility is defined over final wealth (instead of income), ii) wealth changes between the two surveys, and iii) utility function is convex in some regions and concave in others. We note, however, that concave-convex utility functions are non-standard. Instead, the utility functions commonly used in practice (such as those in the general HARA family discussed in Section 3.2) are concave (or convex) everywhere, in which case the claim holds even under i) and ii).

inconsistent across the two surveys. In addition to demographic variables we control for the absolute difference between the risk tolerance reported in surveys 1 and 2. Thus, we can test whether inconsistent revealed risk attitudes are accompanied by large changes in reported risk tolerance. Similarly, we control for the difference in inflation uncertainty across surveys and the total time taken to complete both surveys.

The results reported in Table 5 (Model 1) indicate that inconsistent revealed risk attitudes are driven by three factors: the respondent's i) financial literacy and numeracy, ii) education level and iii) total completion time. In particular, we find that respondents with at most a high school diploma are 15.3% more likely to exhibit inconsistent revealed risk attitudes than respondents with some college but no more than a Bachelor degree. Likewise, respondents who answer correctly half of the numeracy and financial literacy questions are 14.4% more likely to exhibit inconsistent revealed risk attitudes than respondents who answer all six questions correctly. Finally, observe that changes in reported risk tolerance or in inflation uncertainty across surveys do not explain inconsistent switches in revealed risk attitudes.

6.1.2 Correlation between Changes in Point Predictions and Changes in Switching Points

As illustrated in Figure 7, we should expect a negative correlation between changes in inflation expectations and changes in switching points: All else equal, a respondent whose inflation expectations increase (decrease) from one survey to the next should select a lower (higher) switching point in survey 2 than in survey 1. We call this first order effect, the “expectation effect.” There is, however, a possible second order effect, a “risk effect,” for respondents who are not risk neutral. Indeed, all else equal, and in particular if expected inflation remains the

same, an agent may be led to select a different switching point if the shape of his beliefs distribution changes from one survey to the next. For instance, if the inflation uncertainty of a risk-averse agent increases (decreases), then, as shown in Proposition 3, the agent will be led to select a smaller (higher) switching point. The first order “expectation effect” and the second order “risk effect” can go in the same or in opposite directions. In the latter case, the exact correlation between changes in predictions and changes in switching points depends on the relative strength of the two effects. Simulations, however, indicate that the first order “expectation effect” almost always dominates.²⁶

Taking these two effects into consideration, we now define a new variable to study how changes in switching points across surveys respond to changes in point predictions. Let $\{S_1; S_2\}$ and $\{P_1; P_2\}$ denote a respondent’s pairs of switching points and point predictions in surveys 1 and 2. We say that this respondent has an “inconsistent correlation” when the following three conditions are all satisfied:

- 1) $(S_2 - S_1) * (P_2 - P_1) \geq 0$
- 2) $\{S_1; S_2\} \neq \{1; 1\}$ and $\{S_1; S_2\} \neq \{11; 11\}$
- 3) $|S_2 - S_1| > 1$ or $|P_2 - P_1| > 1\%$

The first condition reflects a violation of the first order “expectation effect” just described (e.g. an increase in point prediction should result in a decrease in switching points). Condition 2 excludes respondents who pick the same switching point either equal to 1 or equal to 11 in both surveys. Indeed, regardless of the point predictions made in survey 1 and 2, it is always possible

²⁶ The simulations are based on the expected utility model presented in Section 6.2. For almost all levels of risk aversion, these simulations suggest that the “risk effect” only dominates when expected inflation remains virtually identical across the two surveys, while the distribution of beliefs change from extremely concentrated in one survey to nearly uniform in the other survey.

to find a specific risk attitude to rationalize the behavior of respondents who consistently pick the same extreme choice in both surveys. Note, however, that only 3.2% of the 502 repeat respondents exhibit such behavior.

The last condition characterizes an “indifference zone” in which changes in point predictions and changes in switching points do not exceed respectively $\pm 1\%$ and ± 1 . Respondents with inconsistent correlations should fall outside this indifference zone. In contrast, three types of respondents fall within the indifference zone. First, the 8.4% (42 out of 502) respondents who provide exactly the same prediction and the same switching point in both surveys. Second, the respondents with about the same prediction in both surveys who make slight changes in switching points because of the second order “risk effect.” Finally, the indifference zone reflects the discreteness of the experiment. Indeed, recall that the choice of a switching point may be rationalized by a range of predictions. For instance, a risk neutral agent could have the same inflation expectation of 4% in both surveys, but selects a switching point of 6 in one survey and a switching point of 7 in the other survey as he is indifferent between the two investments in question 6. Likewise, a risk neutral agent could select the same switching point of 6 in both surveys, and change his prediction from 4% to 5%. Note, however, that the way in which the indifference zone is defined may be considered somewhat conservative. Indeed, the behavior of a respondent with $S_2=S_1$ and $|P_2-P_1|=2\%$, or $|S_2-S_1|=2$ and $P_2=P_1$ is rational under risk neutrality. However, such a respondent would fall outside our indifference zone.

To illustrate the adjustments across surveys, we plot in Figure 10 the difference between a respondent’s point predictions across the two surveys (X-axis) and the difference between the respondent’s switching points across surveys (Y-axis). The area of each bubble in Figure 10

reflects the number of respondents at that point. In addition, the indifference zone defined in the previous paragraph is indicated in Figure 10.

As mentioned earlier, we generally expect a negative correlation between the adjustment in predictions and the adjustment in switching points. If this is the case, then respondents have consistent correlations and they should be located either in the upper left or bottom right quadrants of Figure 10. As we can see, this is overwhelmingly the case. More precisely, a total of 102 out of 502 repeat respondents (20.3%) exhibit inconsistent correlations. Furthermore, we find that the correlation coefficient between the adjustment in predictions and the adjustment in switching points is -0.49 .²⁷

To identify whether respondents with inconsistent switching points have specific characteristics, we again estimate a probit model in which the dependent variable is equal to 1 when the respondent exhibits an inconsistent correlation. The results reported in Table 5 (Model 2) reveal that respondents with lower education, respondents who scored low on our numeracy and financial literacy scale, and respondents who took more time to complete the survey are more likely to exhibit inconsistent correlations. More specifically, we find that, compared to a respondent with the highest possible numeracy and financial literacy score, a respondent who answered half the questions correctly is 16% more likely to exhibit an inconsistent correlation. Furthermore, compared to a respondent who went to college up to a Bachelor degree, a respondent without a college education (respectively with a post graduate education) is 15% more likely (respectively 10% less likely) to exhibit an inconsistent correlation.

²⁷ A perfectly linear correlation should not be expected because the difference in switching points across surveys is bounded between -10 and 10.

6.2 *Is the Magnitude of Adjustments Consistent with Theory?*

We have just established that when respondents change their predictions from one survey to the next, most of them also adjust their experimental choices in a direction consistent with the theory. In this section, we explore whether the magnitude of the adjustments in behavior may be rationalized under a simple expected utility model.

6.2.1 Methodology

Consider respondent i who selects a switching point $S_{i,1}$ in survey 1 and $S_{i,2}$ in survey 2. We assume that this respondent possesses a specific power utility function over income of the form $U(x, \theta_i) = x^{\theta_i}$, with $\theta_i > 0$. The parameter θ_i , which is assumed to be time invariant, therefore characterizes the agent's risk attitude. When $0 < \theta_i < 1$ the agent is risk averse with a constant relative risk aversion parameter equal to $(1 - \theta_i)$. The agent is risk neutral when $\theta_i = 1$ and risk loving when $\theta_i > 1$. Finally, we assume that the generalized beta distribution fitted to the respondent's reported probabilistic beliefs in a given survey accurately represents the respondent's true inflation beliefs. The expectation with respect to the distribution elicited in survey 1 (respectively survey 2) is then denoted $E_{i,1}[\cdot]$ (respectively $E_{i,2}[\cdot]$). Our approach to evaluate the magnitude of respondent i 's behavior adjustments proceeds in three steps.

Step 1: The first step consists of inferring a range of risk attitude parameters $[\bar{\theta}_i, \underline{\theta}_i]$ that rationalizes the switching point $S_{i,1}$. More specifically, denote $E_{i,1}[U(A, \theta_i)]$ the respondent's expected utility under investment A. Recall that the payments under investment B increase in increment of \$50 from \$100 in question 1 to \$550 in question 10. We therefore denote $R(S) =$

$\$50 * (1 + S)$ the revenue (in nominal terms) generated by investment B in question $S \in \{1, \dots, 10\}$ of the experiment. Furthermore, we assume that $R(0) = \$50$ and $R(11) = \$600$.²⁸

In this framework, the switching point $S_{i,1}$ actually selected by the agent should satisfy:

$$E_{i,1}[U(R(S_{i,1}), \theta_i)] \leq E_{i,1}[U(A, \theta_i)] \leq E_{i,1}[U(R(S_{i,1} + 1), \theta_i)].$$

In other words, given his probabilistic beliefs and his risk attitude parameter, the respondents should have no strict incentive to switch earlier or later than he actually did in the experiment.²⁹

We can then define the pair $\{\underline{\theta}_i, \bar{\theta}_i\}$ that satisfies:

$$E_{i,1}[U(R(S_{i,1}), \underline{\theta}_i)] = E_{i,1}[U(A, \underline{\theta}_i)] \quad \text{and} \quad E_{i,1}[U(R(S_{i,1} + 1), \bar{\theta}_i)] = E_{i,1}[U(A, \bar{\theta}_i)]$$

It must then be the case that, given the probabilistic beliefs expressed by the respondent, any risk attitude parameter $\theta \in [\underline{\theta}_i, \bar{\theta}_i]$ can rationalize the choice of switching point $S_{i,1}$. Finally, we denote $\hat{\theta}_i = (\underline{\theta}_i + \bar{\theta}_i) / 2$ the midpoint of the risk attitude parameters.

Step 2: In step 2, we use the pair $\{\underline{\theta}_i, \bar{\theta}_i\}$ inferred in step 1 and the probabilistic beliefs expressed by the respondent in survey 2 to predict a range of possible switching points $[\underline{S}_{i,2}, \bar{S}_{i,2}]$. These predicted switching points are defined by the pair of inequalities:

$$E_{i,2}[U(R(\underline{S}_{i,2}), \underline{\theta}_i)] \leq E_{i,2}[U(A, \underline{\theta}_i)] < E_{i,2}[U(R(\underline{S}_{i,2} + 1), \underline{\theta}_i)]$$

$$E_{i,2}[U(R(\bar{S}_{i,2}), \bar{\theta}_i)] \leq E_{i,2}[U(A, \bar{\theta}_i)] < E_{i,2}[U(R(\bar{S}_{i,2} + 1), \bar{\theta}_i)]$$

²⁸ Assuming that $\$0 \leq R(0) < \50 and $R(11) > \$600$ does not affect the results meaningfully.

²⁹ To illustrate, consider a respondent who opts for investment A in the first question and then switches to investment B for the remaining nine questions. In that case, we have $S_{i,1} = 1$, and this decision is rational in our framework if θ_i is such that $E_{i,1}[U(100, \theta_i)] \leq E_{i,1}[U(A, \theta_i)] \leq E_{i,1}[U(150, \theta_i)]$.

As we shall see, $(\underline{S}_{i,2} - \bar{S}_{i,2})$ may be equal to 0, 1 or 2 depending on the values taken by the pair $\{\underline{\theta}_i, \bar{\theta}_i\}$ and by the respondent's reported belief distribution in survey 2.

Step 3: The last step consists of comparing the predicted range of switching points $[\underline{S}_{i,2}, \bar{S}_{i,2}]$ to $S_{i,2}$, the respondent's actual choice of switching point in survey 2. More precisely, we calculate a new variable, the "adjustment precision," as the minimum distance between $S_{i,2}$ and $[\underline{S}_{i,2}, \bar{S}_{i,2}]$. The adjustment precision is therefore an integer between -10 and 10 and equal to 0 when $S_{i,2} \in [\underline{S}_{i,2}, \bar{S}_{i,2}]$. Observe that an adjustment precision equal to 0 does not simply imply that the respondent choice of switching point in survey 2 is consistent with our model. Instead, it implies that the magnitudes of all adjustments the respondent made are consistent with our model. In other words, an adjustment precision equal to 0 requires the distribution of beliefs expressed in survey 1 and 2, as well the experimental choices made in survey 1 and 2, to all be consistent with our simple expected utility model.

6.2.2 Results

Step 1 of our methodology produces a specific pair of risk attitude parameters $\{\underline{\theta}_i, \bar{\theta}_i\}$ for each of the 502 repeat respondents. We find that the distribution of the mid-points, $\hat{\theta}_i$, of these risk attitude parameters has a mean of 0.975 and a standard deviation of 0.342. This distribution is consistent with the results presented in Section 5. In particular, the average $\hat{\theta}_i$ is close to 1, thereby reflecting the fact that, on average, respondents in survey 1 were found to behave as if risk neutral (see Figure 7). Furthermore, the relatively large standard deviation of $\hat{\theta}_i$ is consistent with the substantial heterogeneity in revealed risk attitudes illustrated in Figure 8.

Our methodology also produces a range of predicted switching points $[\underline{S}_{i,2}, \bar{S}_{i,2}]$ for each respondent. The width of this range is equal to either 0, 1 or 2 for respectively 30%, 64% and 6% of the respondents. Now imagine the respondents in survey 2 made their choices in the experiment randomly by assigning the same probability to each possible switching point. In that case, we would find that on average 16% of the switching points chosen randomly in step 2 fall within the range of predicted switching points $[\underline{S}_{i,2}, \bar{S}_{i,2}]$ obtained in step 1. This statistic will therefore serve as a benchmark to evaluate how well our simple model explains the respondents' experimental choices in survey 2.

We plot in Figure 11 the distribution of adjustment precisions in survey 2. This distribution is highly concentrated around 0. In fact, 41.2% of our respondents have an adjustment precision of 0, therefore clearly exceeding the 16% benchmark. Accounting for small errors respondents may have made when reporting inflation expectations or when making experimental choices, we find that 63.7% (respectively 78.8%) of the respondents have an adjustment precision in the interval $[-1,1]$ (respectively $[-2,2]$). In other words, for most of the respondents, the magnitude of behavior adjustment across surveys is consistent, or nearly consistent, with our simple expected utility model. Several factors may explain why the remaining respondents have large differences between predicted and actual switching points: i) their behavior in the experiment is inconsistent with expected utility, ii) they are expected utility maximizers, but their utility function is different than the one assumed in our simple model, iii) they incorrectly report their inflation expectations, or iv) their attitude toward risk changed between surveys.

To identify the characteristics of respondents whose prediction and switching point adjustments cannot be explained by our model, we estimate an ordered probit model in which the dependent variable is the absolute value of the adjustment precision. As we can see in Table 5 (Model 3), the usual three variables are statistically significant: low numeracy and financial literacy, lower education, and high completion time. We also find that respondents with an income greater than \$75k are better able to adjust their behavior across surveys. This result may be explained by the fact that these respondents are also more likely to report the same prediction and the same switching point across surveys.

7. Discussion

In this paper, we compare the inflation expectations individual consumers reported in a survey with their choices in a financially incentivized investment experiment. Our results show that, on average, there is a tight correspondence between stated beliefs and behavior in the experiment. Across respondents, we find a substantial amount of heterogeneity in behavior that can be explained to a large extent by the respondents' self reported risk tolerances. Furthermore, when considering changes in beliefs for the same individual over time, we find the adjustments in experimental behavior to be mostly consistent (both in direction and magnitude) with standard economic theory. Finally, the respondents whose behavior is difficult to rationalize (either because they switch multiple times in the experiment, or because their behavior is inconsistent with theory) tend to exhibit specific characteristics: they score lower on a numeracy and financial literacy scale, they are less educated and they take longer to complete the survey. We now conclude with a brief discussion of the implications of our results.

Historically, economists have been skeptical about expectations surveys (see Manski 2004 for a discussion). One of the main concerns is that, unlike choices, subjective beliefs are unobservable. As a result, nothing guarantees that surveys can capture the true beliefs of each respondent. By showing that the beliefs reported in an inflation survey are informative, in the sense that they correlated with a decision with financial consequences, our analysis provides evidence to support the empirical validity of inflation expectations surveys. More generally, our study illustrates how controlled experiments can be used to establish the information content of an expectations survey.

More importantly, by showing that respondents generally act on their stated inflation beliefs in a way consistent with standard economic theory, we provide some evidence to support one of the key assumptions underlying macroeconomic models whereby forward looking agents make economic decisions that are influenced by their beliefs about future inflation. The generality of this finding, however, should not be overstated for at least two reasons. First, our surveys concentrate exclusively on households and therefore ignore other important economic actors (e.g. firms, professional investors). Nevertheless, we note that, on aggregate, households are a key component of economic activity with considerable influence on aggregate demand and therefore equilibrium prices. Second, our experimental results do not necessarily imply that the decisions made by consumers in their daily life (e.g., about investments, savings, wages, or large durable purchases) are systematically influenced by the beliefs they hold about future inflation. As argued in the introduction, it is very difficult to test this hypothesis directly with real life observations without other confounding factors (e.g. time discounting, liquidity constraints) playing a substantial role. What we have shown is that, when presented with an opportunity to act on their inflation beliefs, most consumers responding to a survey make decisions consistent

with the expectations they report. Our results are therefore a first step in understanding how inflation prospects affect actual economic behavior.

More generally, this paper fits in a broader research agenda aimed at collecting new survey data to better measure and understand how agents form and update their subjective expectations about macroeconomic variables of interest (see e.g. Branch 2004, Eusepi and Preston 2011). Furthermore, although economists implicitly assume that expectations influence behavior, it is important to confirm empirically the extent to which this is the case in a variety of contexts. For example, policy makers need to understand how the decision of households to default on their home mortgages may depend on their expectations of future house price appreciation or declines. Likewise, it is important to know the extent to which job search efforts in the labor market are affected by workers' expectations about their ability to find jobs in various occupations or locations, as well as their expectations for wage growth on a current job. As discussed earlier, this paper provides direct evidence at the individual level that inflation expectations can affect economic decisions, which is part of the micro-foundations of modern macroeconomic models.

Beyond testing the assumption that agents act on their beliefs, it is also important to identify the characteristics of those who fail to do so. Indeed, this may be the only way for policy makers to address the possible negative externalities generated by agents who are ill-equipped to make sound economic choices. In that respect, our finding that consumers with low numeracy and financial literacy are less likely to act on their reported beliefs in accordance with expected utility theory is particularly relevant in the context of the rapidly growing literature on the role of numeracy and financial literacy in economic decisions. For instance, financial illiteracy has now clearly been linked with lower wealth accumulation (Behrman et al. 2012), insufficient

retirement planning (Lusardi and Mitchell 2007), and inadequate portfolio diversification (Goetzmann and Kumar 2008). In its 2009 report, the “President’s Advisory Council on Financial Literacy” identified numeracy and financial literacy as prime determinants of the 2008 subprime mortgage crisis. These views are supported by the results of Gerardi, Goette and Meier (2011) who find a large negative correlation between numeracy and mortgage delinquency in the U.S. Our study underlies the importance of financial education by providing further evidence about the role played by financial literacy on sound economic choices. Furthermore, by identifying a breakdown of the connection between beliefs and actions, our results suggest a specific channel through which financial literacy affects economic behavior.

Olivier Armantier (corresponding Author. olivier.armantier@ny.frb.org), Giorgio Topa, Wilbert van der Klaauw and Basit Zafar: Federal Reserve Bank of New York, 33 Liberty Street, New York, NY, 10013, USA.

Wändi Bruine de Bruin, Leeds University Business School, Leeds, S2 9JT, UK.

REFERENCES

- Armantier O., Bruine de Bruin W., van der Klaauw W., Nelson S., Topa G. and B. Zafar (2014): “The Price is Right: Belief Updating in a Price Information Experiment,” *Review of Economics and Statistics*, forthcoming.
- Behrman J., Mitchell O., Soo C. and D. Bravo (2012): “How Financial Literacy Affects Household Wealth Accumulation”, *AER Papers and Proceedings*, 102(3), 300-304.
- Bernanke B. (2004): “The Economic Outlook and Monetary Policy” speech at the Bond Market Association Annual Meeting, New York, New York.

- Bernanke B. (2007): "Inflation Expectations and Inflation Forecasting," speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts.
- Bonin H., Dohmen T., Falk A., Huffman D. and U. Sunde (2007): "Cross-Sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes," *Labour Economics*, 14, 926–937.
- Branch W. (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *Economic Journal*, Vol. 114, 592-621.
- Bruine de Bruin W., Downs J.S., Fischhoff B. and C. Palmgren (2007): "Development and Evaluation of an HIV/AIDS Knowledge Measure for Adolescents Focusing on Misunderstood Concepts," *HIV/AIDS Prevention in Children and Youth*, 8, 35-57.
- Bruine de Bruin W., Potter S., Rich R., Topa G. and W. van der Klaauw (2010a): "Improving Survey Measures of Household Inflation Expectations," *Current Issues in Economics and Finance*, Vol. 16(7), Aug/Sept.
- Bruine de Bruin W., van der Klaauw W., Downs J.S., Fischhoff B., Topa G. and O. Armantier (2010b): "Expectations of Inflation: The Role of Financial Literacy and Demographic Variables," *Journal of Consumer Affairs*, 44, 381-402.
- Caliendo M., Fossen M. and A. Kritikos (2010): "The Impact of Risk Attitudes on Entrepreneurial Survival", *Journal of Economic Behavior & Organization*, 76-1, pp 45-63.
- Carmines E. and R. Zeller (1991): *Reliability and Validity Assessment*, Thousand Oaks, CA: Sage Publications.
- Carroll C. (2003): "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, CXVIII, 269–298.
- Coibion O. and Y. Gorodnichenko (2013): "What Can Survey Forecasts Tell us about Information Rigidities?" *Journal of Political Economy*, 120-1, 116–159
- Dohmen T., Falk A., Huffman D., Sunde U., Schupp J. and G. Wagner (2011): "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences," *Journal of the European Economic Association*, Vol. 9, 3, 522–550.
- Dominitz J. and C. Manski (2007): "Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study", *Journal of the European Economic Association*, Vol. 5 no2-3, 369–379.
- Eckel E. and R. Wilson (2004): "Is Trust a Risky Decision?" *Journal of Economic Behavior & Organization*, 55, pp 447-465.
- Ehrbeck T. and R. Waldmann (1996): "Why are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics*, 111.
- Engelberg J., Manski C. and J. Williams (2009): "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Business and Economic Statistics*, 27, pp. 30-41.
- Eusepi S. and B. Preston (2011): "Expectations, Learning and Business Cycle Fluctuations," *American Economic Review*, Vol. 101, No. 6.

- Gali, J. (2008): *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.
- Gerardi K., Goette, L. and S. Meier (2011): “Financial Literacy and Subprime Mortgage Delinquency: Evidence from a Survey Matched to Administrative Data,” *Federal Reserve Bank of Atlanta*, working paper.
- Goetzmann W. and A. Kumar (2008): “Equity Portfolio Diversification,” *Review of Finance*, 12, 433-463.
- Gollier C. and J. Pratt (1996): “Weak Proper Risk Aversion and the Tempering Effect of Background Risk,” *Econometrica*, 64, 1109-1123.
- Harrison G., Lau M. and M. Williams (2002): “Estimating Individual Discount Rates in Denmark: A Field Experiment,” *American Economic Review*, 92, 1606–1617.
- Holt C. and S. Laury (2002): “Risk Aversion and Incentive Effects,” *American Economic Review*, 2002, 92(5), pp. 1644–55.
- Jewitt I. (1989): “Choosing Between Risky Prospects: The Characterization of Comparative Statics Results, and Location Independent Risk,” *Management Science*, 35: 60-70.
- Karni E. and Z. Safra, (1995): “The Impossibility of experimental Elicitation of Subjective Probabilities,” *Theory and Decision*, 38, 313-320.
- Keane M. and D. Runkle (1990): “Testing the Rationality of Price Forecasts: New Evidence from Panel Data,” *American Economic Review*, Vol. 80, No. 4, pp. 714-735.
- Lipkus I., Samsa G. and B. Rimer (2001): “General Performance on a Numeracy Scale Among Highly Educated Samples,” *Medical Decision Making*, 21, 37–4.
- Lucas R. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4, 103-124.
- Lusardi A. (2007): “Financial Literacy: An Essential Tool for Informed Consumer Choice?” presented at the conference “Understanding Consumer Credit: A National Symposium on Expanding Access, Informing Choices, and Protecting Consumers,” organized by the Harvard Joint Center for Housing Studies, Harvard University.
- Lusardi A. and O. Mitchell (2007): “Financial Literacy and Retirement Preparedness. Evidence and Implications for Financial Education,” *Business Economics*, Jan, 35-44.
- Mankiw G. and R. Reis (2002): “Sticky Information versus Sticky Prices: a Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117, 1295-1328.
- Mankiw G., Reis R. and J. Wolfers (2003): “Disagreement about Inflation Expectations,” NBER Macroeconomics Annual.
- Manski C. (2004): “Measuring Expectations,” *Econometrica*, Vol. 72, No. 5, 1329–1376.
- Muth J. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315-335.
- Nunes R. (2010): “Inflation Dynamics: The Role of Expectations,” *Journal of Money, Credit and Banking*, Volume 42, Issue 6, pages 1161–1172.
- Ottaviani M. and P. Sørensen (2006): “The Strategy of Professional Forecasting,” *Journal of Financial Economics*, 81, 441-466.

- Pesaran M.H. and M. Weale (2006): "Survey Expectations," in *Handbook of Economic Forecasting*. Graham Elliott, Clive W.J. Granger, Allan Timmermann.
- Pratt J. (1964): "Risk Aversion in the Small and in the Large," *Econometrica*, 32-1/2, p 122-136.
- Roberts J. (1995): "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, 27, 975–984.
- Rothschild M. and J. Stiglitz (1970): "Increasing Risk: I. A Definition", *Journal of Economic Theory*, 2, 225-243.
- Savage L.J. (1971): "Elicitation of Personal Probabilities and Expectation," *Journal of the American Statistical Association*, 66, 783-801.
- Sargent T. (1999): *The Conquest of American Inflation*. Princeton University Press.
- Shiller R. (1997): "Why Do People Dislike Inflation?" in *Reducing Inflation: Motivation and Strategy*, Chicago: University of Chicago Press, pp. 13-65.
- Sims C. (2006): "Rational Inattention: Beyond the Linear-Quadratic Case," *American Economic Review*, 96 (2), 158–163.
- Sims C. (2009): "Inflation Expectations, Uncertainty and Monetary Policy," BIS Working Paper No. 275.
- van der Klaauw W. (2012): "On the Use of Expectations Data in Estimating Structural Dynamic Choice Models," *Journal of Labor Economics*, Vol. 30(3), 521-554.
- van der Klaauw W. and K. Wolpin. (2008): "Social Security and the Savings and Retirement Behavior of Low Income Households," *Journal of Econometrics*, Vol. 145(1-2), 21-42.
- Wolpin K. and F. Gonul (1985): "On the Use of Expectations Data in Micro Surveys: The case of Retirement." Report on the Employment and Training Administration, U.S. Department of Labor, Washington, DC, March.
- Woodford M. (2003): "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps*, edited by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford. Princeton, NJ: Princeton Univ. Press.
- Woodford M. (2005): "Central Bank Communication and Policy Effectiveness," Proceedings, Federal Reserve Bank of Kansas City, August issue, 399-474.

Table1: Descriptive Statistics						
	All Data	Group 1 Missing Data for Point Prediction or Experiment	Group 2 Non-Rationalizable b Experimental Choices	Group 3 Rationalizable Experimental Choices	Group 4 Rationalizable Choices in Both Surveys	Probit ^a Probability to be in Group 4
Age	52.126 (14.049)	51.765 (16.794)	51.291 (13.657)	52.519 (13.736)	52.592 (13.422)	-0.001 (0.001)
Gender (female)	57.3%	73.0%	68.2%	54.4%	52.6%	-0.043 (0.032)
Education: No more than High School	15.1%	17.4%	18.6%	14.0%	12.8%	0.051 (0.048)
Education: More than Bachelor	20.9%	23.5%	13.25%	21.6%	22.5%	0.017 (0.038)
Income greater than \$75k	41.8%	23.5%	29.1%	44.9%	47.4%	0.032 (0.031)
Numeracy and Financial Literacy ^c	4.502 (1.481)	3.3487 (1.692)	3.629 (1.553)	4.746 (1.335)	4.986 (1.234)	0.058*** (0.010)
Reported Risk Tolerance ^d	3.335 (1.586)	3.601 (1.823)	3.308 (1.745)	3.368 (1.567)	3.403 (1.548)	0.011 (0.010)
Point prediction	5.382 (6.903)	9.394 ^e (12.842)	8.705 (6.387)	4.777 (3.190)	4.563 (3.167)	—
Number of Observations	1,479 (100%)	115 (7.8%)	151 (10.2%)	1,213 (82.0%)	1,004 (67.9%)	1,479

^a Marginal effects are reported. The endogenous variable equals one when a respondent is in group 4. The log Likelihood is -269.79. The pseudo R² is 0.080.

^b A respondent is said to make non rationalizable choices when he switches more than once from investment A to investment B in the ascending scale treatment.

^c The variable takes integer values between 0 and 6 depending on the number of correct answers the respondent gave to the six questions asked to measure numeracy and financial literacy.

^d Self reported willingness to take risk regarding financial matters on a scale from 1 (Not willing at all) to 7 (very willing).

^e This statistic is based on 85 observations.

Table 2: Earnings under investment A												
Rate of inflation	-1% or less (deflation)	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10% or more
Earnings	\$600	\$550	\$500	\$450	\$400	\$350	\$300	\$250	\$200	\$150	\$100	\$50

Table 3: Factors Influencing the Choice of the Switching Point
 Ordered probit estimations where the dependent variable is the switching point
 (an integer between 0 and 10) for each respondent

	Model 1		Model 2		Model 3		Model 4	
	Survey 1	Survey 2	Survey 1	Survey 2	Survey 1	Survey 2	Survey 1	Survey 2
Point prediction	-0.108*** (0.009)	-0.104*** (0.008)	-0.108*** (0.009)	-0.104*** (0.009)	—	—	-0.065*** (0.017)	-0.091*** (0.016)
Estimated Expected Prediction ^d	—	—	—	—	-0.090*** (0.012)	-0.096*** (0.010)	—	—
Reported Risk Tolerance ^a	0.188*** (0.029)	0.183*** (0.028)	0.191*** (0.030)	0.184*** (0.029)	0.163*** (0.030)	0.165*** (0.028)	0.193*** (0.030)	0.183*** (0.029)
Numeracy and Financial Literacy Score ^b	0.003 (0.035)	0.005 (0.032)	-0.002 (0.035)	0.001 (0.032)	-0.001 (0.035)	0.011 (0.032)	-0.002 (0.035)	-0.001 (0.031)
Inflation Uncertainty ^c	-0.007*** (0.002)	-0.004*** (0.001)	-0.007*** (0.002)	-0.004*** (0.001)	-0.005*** (0.002)	-0.007*** (0.001)	-0.007*** (0.002)	-0.004*** (0.001)
Log of Time Taken to Complete the Survey	0.007 (0.024)	-0.012 (0.019)	0.008 (0.024)	-0.008 (0.019)	0.009 (0.024)	-0.014 (0.019)	0.008 (0.024)	-0.007 (0.019)
Age	—	—	0.003 (0.003)	0.001 (0.003)	0.004 (0.003)	0.001 (0.003)	0.003 (0.003)	0.001 (0.003)
Gender (female)	—	—	-0.060 (0.089)	-0.094 (0.088)	-0.101 (0.088)	-0.126 (0.088)	-0.063 (0.089)	-0.094 (0.088)
Income greater than \$75k	—	—	-0.027 (0.106)	-0.095 (0.087)	0.046 (0.089)	-0.059 (0.087)	-0.035 (0.089)	-0.099 (0.088)
Education: No more than High School	—	—	0.124 (0.125)	0.052 (0.127)	-0.001 (0.124)	0.036 (0.128)	0.128 (0.125)	0.070 (0.127)
Education: More than Bachelor			0.078 (0.106)	0.147 (0.105)	0.043 (0.106)	0.081 (0.105)	0.077 (0.107)	0.177 (0.105)
“Prices” * “Ascending”	-0.130 (0.119)	-0.108 (0.120)	-0.129 (0.119)	-0.108 (0.120)	-0.146 (0.121)	-0.069 (0.121)	-0.022 (0.152)	0.045 (0.170)
“Prices” * “Descending”	0.050 (0.119)	0.026 (0.118)	0.053 (0.116)	0.033 (0.118)	-0.013 (0.116)	0.034 (0.118)	0.150 (0.157)	0.031 (0.152)
“Inflation” * “Descending”	-0.007 (0.118)	0.104 (0.116)	-0.011 (0.118)	0.091 (0.119)	0.039 (0.118)	0.096 (0.119)	0.056 (0.148)	0.199 (0.165)
Point prediction * “Prices” * “Ascending”	—	—	—	—	—	—	-0.028 (0.025)	-0.033 (0.026)
Point prediction * “Prices” * “Descending”	—	—	—	—	—	—	0.025 (0.024)	-0.002 (0.022)
Point prediction * “Inflation” * “Descending”	—	—	—	—	—	—	-0.018 (0.027)	-0.026 (0.023)
N	589	607	589	607	589	607	589	607
Log Likelihood	-1,227.0	-1,266.9	-1,225.5	-1,265.3	-1,232.9	-1,265.2	-1,224.4	-1,277.7

The standard deviations are robust and clustered at the treatment combination level. Significance: * = 10%, ** = 5%, *** = 1%.

^a Self reported willingness to take risk regarding financial matters on a scale from 1 (Not willing at all) to 7 (very willing).

^b The variable takes integer values between 0 and 6 depending on the number of correct answers the respondent gave to the six questions asked to measure numeracy and financial literacy.

^c As explained in Section 2.3, this variable is equal to the standard deviation of the Beta distribution fitted to the respondent’s reported probabilistic beliefs.

^d As explained in Section 2.3, this variable is equal to the mean of the Beta distribution fitted to the respondent’s reported probabilistic beliefs.

Table 4: Explaining the Deviations from Risk Neutrality
 Ordered Probit estimations based on the difference between a respondent
 switching point and her/his pair of “risk neutral switching points.”

	Model 1 ^a		Model 2 ^b		Model 3 ^c		Model 4 ^d	
	Dependent = Deviation		Dependent = Absolute Deviation		Dependent = Deviation when Deviation < 0		Dependent = Deviation when Deviation > 0	
	Survey 1	Survey 2	Survey 1	Survey 2	Survey 1	Survey 2	Survey 1	Survey 2
Reported Risk Tolerance ^e	0.247*** (0.030)	0.196*** (0.028)	-0.550*** (0.128)	-0.543*** (0.124)	0.258*** (0.056)	0.220*** (0.053)	0.356*** (0.059)	0.217*** (0.059)
Square of Reported Risk Tolerance	—	—	0.077*** (0.017)	0.071*** (0.017)	—	—	—	—
Inflation Uncertainty ^g	0.002 (0.002)	0.001 (0.011)	0.004** (0.002)	0.004*** (0.001)	-0.011*** (0.004)	-0.003** (0.002)	0.010*** (0.003)	0.008*** (0.003)
Numeracy and Financial Literacy Score ^f	-0.034 (0.035)	-0.057* (0.031)	-0.146*** (0.036)	-0.118*** (0.033)	0.234*** (0.059)	0.191*** (0.054)	-0.175*** (0.068)	-0.239*** (0.067)
Log of Time Taken to Complete the Survey	-0.036 (0.025)	-0.029 (0.019)	0.075*** (0.024)	0.070*** (0.020)	-0.097** (0.036)	-0.064** (0.031)	0.087* (0.048)	0.091** (0.041)
Gender (female)	0.101 (0.087)	0.011 (0.088)	0.043 (0.091)	0.007 (0.093)	0.206 (0.148)	-0.071 (0.152)	-0.014 (0.184)	0.236 (0.206)
Age	0.004 (0.003)	0.001 (0.003)	-0.001 (0.003)	0.004 (0.003)	0.001 (0.006)	0.006 (0.006)	0.005 (0.007)	0.011 (0.007)
Income greater than \$75k	-0.092 (0.088)	-0.143 (0.087)	-0.125 (0.092)	-0.048 (0.092)	0.185 (0.149)	0.306** (0.152)	-0.131 (0.184)	-0.106 (0.205)
Education: No more than High School	0.298** (0.125)	0.085 (0.126)	0.535*** (0.128)	0.591*** (0.129)	-0.467** (0.218)	-0.666*** (0.203)	0.820*** (0.217)	0.950*** (0.244)
Education: More than Bachelor	0.019 (0.105)	0.102 (0.105)	-0.273** (0.112)	-0.350*** (0.116)	0.463** (0.187)	0.487** (0.198)	-0.575** (0.257)	-0.831*** (0.287)
“Prices” * “Ascending”	0.003 (0.121)	-0.039 (0.122)	0.106 (0.127)	0.086 (0.128)	-0.110 (0.206)	-0.115 (0.211)	0.088 (0.255)	0.330 (0.284)
“Prices” * “Descending”	0.180 (0.120)	0.164 (0.119)	0.022 (0.127)	-0.001 (0.126)	-0.086 (0.216)	-0.066 (0.214)	0.021 (0.251)	-0.286 (0.270)
“Inflation” * “Descending”	-0.138 (0.118)	-0.173 (0.118)	0.002 (0.123)	-0.084 (0.125)	-0.141 (0.194)	-0.078 (0.197)	-0.230 (0.268)	-0.422 (0.310)
N	589	607	589	607	242	134	161	134
Log Likelihood	-1,295.6	-1,259.7	-1,030.1	-1,008.2	-369.8	-204.1	-255.4	-213.9

The standard deviations are robust and clustered at the treatment combination level. Significance: * = 10%, ** =5%, *** = 1%.

^a In Model 1 the dependent variable is the difference between a respondent actual switching point and his/her pair of “risk neutral switching points.”

^b In Model 2 the dependent variable is the absolute value of the difference between a respondent actual switching point and his/her pair of “risk neutral switching points.”

^c In Model 3 the dependent variable is the difference between a respondent actual switching point and his/her pair of “risk neutral switching points,” but the sample is restricted to respondents who behaved as if risk averse (i.e. the deviations from risk neutrality is strictly negative).

^d In Model 4 the dependent variable is the difference between a respondent actual switching point and his/her pair of “risk neutral switching points,” but the sample is restricted to respondents who behaved as if risk loving (i.e. the deviations from risk neutrality is strictly positive).

^e Self reported willingness to take risk regarding financial matters on a scale from 1 (Not willing at all) to 7 (very willing).

^f The variable takes integer values between 0 and 6 depending on the number of correct answers the respondent gave to the six questions asked to measure numeracy and financial literacy.

^g As explained in Section 2.3, this variable is equal to the standard deviation of the Beta distribution fitted to the respondent’s reported probabilistic beliefs.

Table 5: Are Changes in Predictions and Changes in Switching Points across Surveys Consistent with Theory?

	Model 1 ^a	Model 2 ^b	Model 3 ^c
	Probit (Marginal Effects) Dependent = Inconsistent Revealed Risk Attitudes	Probit (Marginal Effects) Dependent = Inconsistent Correlation	Ordered Probit Dependent = Absolute value of Adjustment Precision
Absolute Difference in Reported Risk Tolerance across Surveys ^d	-0.009 (0.014)	0.032 (0.022)	-0.027 (0.060)
Absolute Difference in Inflation Uncertainty across Surveys ^e	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Numeracy and Financial Literacy Score ^f	-0.048*** (0.014)	-0.054*** (0.014)	-0.153*** (0.042)
Log of Total Time Taken to Complete the Two Surveys	0.013** (0.005)	0.027*** (0.007)	0.054** (0.021)
Gender (female)	-0.026 (0.024)	-0.062 (0.073)	0.089 (0.100)
Age	-0.001 (0.001)	-0.001 (0.001)	0.001 (0.004)
Income greater than \$75k	-0.009 (0.025)	-0.054 (0.036)	-0.203** (0.101)
Education: No more than High School	0.153** (0.060)	0.150** (0.064)	0.722*** (0.148)
Education: More than Bachelor	-0.052 (0.027)	-0.096** (0.038)	-0.353*** (0.125)
“Prices” * “Ascending”	-0.026 (0.021)	-0.054 (0.047)	-0.146 (0.144)
“Prices” * “Descending”	0.025 (0.032)	-0.035 (0.049)	0.098 (0.141)
“Inflation” * “Descending”	-0.028 (0.035)	-0.020 (0.048)	-0.105 (0.137)
N	495	495	495
Log Likelihood	-88.5	-222.1	-791.5

The standard deviations are robust and clustered at the treatment combination level. Significance: * = 10%, ** =5%, *** = 1%.

^a A respondent is said to have “inconsistent revealed risk attitudes” when he behaves as if risk averse in one survey and as if risk loving in the other survey.

^b A respondent is said to have an “inconsistent correlation” when his pair of switching points and his pair of point predictions satisfy all the conditions described in Section 6.1.2.

^c The “adjustment precision” is the difference between a respondent’s switching point in survey 2 and his pair of predicted switching points (see Section 6.2).

^d Self reported willingness to take risk regarding financial matters on a scale from 1 (Not willing at all) to 7 (very willing).

^e As explained in Section 2.3, this variable is equal to the standard deviation of the Beta distribution fitted to the respondent’s reported probabilistic beliefs.

^f The variable takes integer values between 0 and 6 depending on the number of correct answers the respondent gave to the six questions asked to measure numeracy and financial literacy.

Figure 1: Distribution of Point Predictions




Figure 2: Distribution of Individual Differences in Point Predictions (N=502)




Figure 3: Histogram of Point Predictions for each Survey Treatment




Figure 4: Distribution of Switching Points




Figure 5: Distribution of Individual Differences in Switching Points (N=502)




Figure 6: Histogram of Switching Points for each Treatment Combination




Figure 7: Choices and Predictions




Figure 8: Choices and Dispersion of Predictions (N=502)




Figure 9: Distribution of Revealed Risk Attitudes in each Survey (N=502)




Figure 10: Prediction and Behavior Adjustment (N=502)




Figure 11: Adjustment Precision in Survey 2

