Learning from Inflation Experiences*

Ulrike Malmendier†
UC Berkeley, NBER and CEPR

Stefan Nagel‡
Stanford University, NBER and CEPR

February 10, 2012

Abstract

How do individuals form expectations about future inflation? We propose that past inflation experiences are an important determinant absent from existing models. Individuals overweigh inflation rates experienced during their life-times so far, relative to other historical data on inflation. Differently from adaptive-learning models, experience-based learning implies that young individuals place more weight on recently experienced inflation than older individuals since recent experiences make up a larger part of their life-times so far. Averaged across cohorts, expectations resemble those obtained from constant-gain learning algorithms common in macroeconomics, but the speed of learning differs between cohorts.

Using 54 years of microdata on inflation expectations from the Reuters/Michigan Survey of Consumers, we show that differences in life-time experiences strongly predict differences in subjective inflation expectations. As implied by the model, young individuals place more weight on recently experienced inflation than older individuals. We find substantial disagreement between young and old individuals about future inflation rates in periods of high surprise inflation, such as the 1970s. The experience effect also helps to predict the time-series of forecast errors in the Reuters/Michigan survey and the Survey of Professional Forecasters, as well as the excess returns on nominal long-term bonds.

*We thank Ralph Koijen, Monika Piazzesi, Christina Romer, and workshop participants at Berkeley, the European Central Bank, the Federal Reserve Board, Harvard, the University of Mannheim, Ohio State, Pompeu Fabra, Stanford, the NBER EFG, NBER Behavioral Finance, and American Economic Association meetings for helpful comments and Sith Chaisurote for excellent research assistance.

†Department of Economics and Haas School of Business, University of California, 501 Evans Hall, Berkeley, CA 94720-3880, ulrike@econ.berkeley.edu

‡Stanford University, Graduate School of Business, 518 Memorial Way, Stanford, CA 94305, Nagel_Stefan@gsb.stanford.edu

1 Introduction

How do individuals form expectations about future inflation? The answer to this question is of central importance to policy-makers in the arena of monetary economics and to individual households making financial and consumption decisions alike. Despite a large volume of research on the determinants of expectation formation, there is still little convergence on the best model to predict inflation expectations (see Mankiw, Reis, and Wolfers (2003); Blanchflower and Kelly (2008)). Both the "stickiness" of inflation rate changes (Sims (1998), Mankiw and Reis (2006)) and the heterogeneity in the formation of expectations remain hard to reconcile with existing models.

In this paper, we propose that a key ingredient missing from existing theories is individuals' personal experiences of past inflation. We argue that, when forming inflation expectations, individuals put a higher weight on realizations experienced over their life-times than on other available historical data. Such experience-based learning is related to the adaptive-learning approach in macroeconomics, but it differs in one key respect: data realized during individuals' life-times carries higher weight than other historical data. Averaged across cohorts, the resulting expectations resemble those obtained from constant-gain learning algorithms commonly used in macroeconomics; but between cohorts, learning speed and beliefs are heterogeneous. We use the heterogeneity in subjective expectations between individuals to estimate learning rules without having to rely on aggregate time-series information about average expectations to fit the learning parameters.

The experience hypothesis carries a rich set of implications for the formation of inflation expectations. First, beliefs are heterogeneous. Individuals who have lived through high-inflation periods forecast higher future inflation than individuals who experienced low inflation during their life-times so far. Second, young individuals place more weight on recent inflation rates than older individuals since recent experiences make up a larger part of their life-times so far. Third, learning dynamics are perpetual. Beliefs keep fluctuating and do not converge in the long-run, as weights on historical data diminish when old generations disappear and new generations emerge.

We test the experience-based model using 54 years of microdata on inflation expectations from the Reuters/Michigan Survey of Consumers (MSC). Our empirical framework employs linear regression-based forecasting rules similar to those used in the adaptive learning literature, in particular Marcet and Sargent (1989), but with the twist that we allow individuals to overweigh data realized during their life-times so far. Specifically, individuals use inflation rates experienced in the past to recursively estimate an AR(1) model of future inflation. The learning-from-experience mechanism is implemented by allowing the gain, i.e., the strength of updating in response to surprise inflation, to depend on age. For example, young individuals react more strongly to an inflation surprise than older individuals who have more data accumulated in their life-time histories. A gain parameter determines how fast these gains decrease with age as more data accumulates. We estimate this gain parameter empirically by fitting the learning rule to individuals' inflation expectations as reported in the MSC. The empirical estimate reveals how people weight their inflation experiences when forming their beliefs about future inflation.

The availability of microdata is crucial for our purpose as it allows us to identify the experience effect from cross-sectional heterogeneity. Our identification strategy relies on time-variation in cross-sectional differences of inflation experiences and relates it to time-variation of cross-sectional differences in inflation expectations. Moreover, the time-variation in cross-sectional differences allows us to employ time dummies in our estimations and thus to separate the experience effect from time trends or other time-specific determinants of inflation expectations that affect all individuals. For example, our analysis does not assume that past inflation experiences are the only influence on people's subjective beliefs about future inflation. Rather, with the inclusion of time dummies, we account for the possibility that individuals draw on the full inflation history that is available at that point in time. Our estimation isolates the incremental explanatory power of life-time experiences over and above the explanatory power of the full time-series of historical inflation data. More generally, the time dummies absorb any variation in inflation expectations that is common to all individuals. For example, individuals might rely, to some extent, on the published forecasts of professional forecasters, which could contain additional information over and above the univariate history of inflation rates. The inclusion of time dummies rules out that any such omitted macroeconomic variables bias the estimation results. This is an important difference to other models of belief formation, such as adaptive learning models, where parameters are fit to aggregate time-series of expectations (e.g., median expectations), making it difficult to rule out that such unobserved effects that are common to all individuals bias the estimation results.

Our estimation results show that learning from experience has an economically important effect on inflation expectations. Individuals of different ages differ in their inflation expectations, and these differences are well explained by differences in their inflation experiences. The heterogeneity in expectations is particularly pronounced following periods of high surprise inflation. For example, in the late 1970s and early 1980s, the average inflation expectations of individuals under the age of 40 exceeded those of older individuals above age 60 by several percentage points, consistent with the fact that the experience of younger individuals was dominated by the high-inflation years of the 1970s, while the experience of older individuals also included the low-inflation years in the 1950s and 1960s. This discrepancy faded away only slowly by the 1990s after a many years of moderate inflation. Our model explains this difference as the result of younger individuals perceiving inflation to be higher, on average, and, to be more persistent when inflation rates were high until the early 1980s, but to be less persistent when inflation rates had dropped subsequently.

Our estimates of the gain parameter imply that recent inflation experiences receive relatively higher weight than experiences earlier in life, though experiences from 20 to 30 years ago still have some long-run effects for older individuals.

We also explore the aggregate implications of learning from experience for the time series of average inflation expectations. We show that if one averages experience-based expectations across cohorts at each point in time, the average learning-from-experience forecast matches the average survey expectations closely. The similarity is remarkable because our estimation did not utilize any information about the level of the average survey expectations, only information about cross-sectional differences between cohorts, Hence, learning from experience helps to simultaneously predict both the cross-section and time-series of inflation expectations. We also show that the average learning-from-experience forecast can be approximated very well with constant-gain learning algorithms that are popular in macroeconomics. The constant-gain parameter that best matches our estimated learning-from-experience weights, $\gamma = 0.0175$ turns out to be quantitatively very similar to the gain that Orphanides and Williams (2005) and Milani (2007) have estimated by fitting the parameter to macroeconomic data and aggregate survey expectations (0.0183 and 0.02 respectively). As with the survey data, this similarity is remarkable because we did not calibrate learning-from-experience rule to match the average level of inflation expectations or any macroeconomic data.

Learning, and learning in boundedly rational fashion in particular, implies that forecast errors should be predictable, at least in sample, but possibly also out of sample. Consistent with this implication, we find that the learning-from-experience forecasts contain information that can be used to predict forecast errors in the level of average MSC inflation expectations in sample as well as out of sample. Furthermore, the forecast error predictability arising from our model is not limited to the non-professional forecasters in the MSC. We also show that the same predictor variable helps predict forecast errors in the Survey of Professional Forecasters and the excess returns on nominal long-term bonds (which could reflect the inflation forecast errors of bond market investors).

Our paper connects to several strands of literature. There is a large literature in macroeconomics analyzes the formation of expectations. While it is well understood at least since Keynes (1936) that macroeconomic outcomes and asset prices depend in crucial ways on the expectations of economic actors, we know less about how economic agents form their subjective beliefs about the future. The literature on adaptive learning (see Bray (1982);

Sargent (1993); Evans and Honkapohja (2001)) views individual agents as econometricians who make forecasts based on simple forecasting rules estimated on historical data. Yet, there is little direct empirical evidence on the actual forecasting rules employed by individuals, even though understanding the formation of inflation expectations, and macroeconomic expectations more generally, is likely to be of first-order importance for macroeconomic policy (Bernanke (2007)).

Conceptually, our approach is related to bounded-memory learning in Honkapohja and Mitra (2003) in that memory of past data is lost. However, while bounded-memory learning agents are homogeneous, the memories of agents in the experience-based model differs depending on their age.

There is a small, but growing literature that looks at heterogeneity in expectations formation with microdata. Building on early work by Cukierman and Wachtel (1979), Mankiw, Reis, and Wolfers (2003) examine the time-variation in dispersion in inflation expectations, and they relate it to models of "sticky" information. Carroll (2003) further investigates the sticky information model, but with aggregate data on inflation expectations. Branch (2004), Branch (2007), and Pfajfar and Santoro (2010) estimate from survey data how individuals choose among competing forecasting models. Piazzesi and Schneider (2010) incorporate data survey data on heterogeneous subjective inflation expectation in asset pricing, while Piazzesi and Schneider (2011) use data on subjective interest rate expectations and a model with adaptive learning. Shiller (1997) and Ehrmann and Tzamourani (2009) examine the relation between cross-country variation in inflation histories and the public's attitudes towards inflation-fighting policies. Our paper contributes to this literature by demonstrating that learning from experience plays a significant role in expectations formation and produces both heterogeneity in expectations and gradually fading memory over time.

Our analysis is related to earlier empirical findings of Malmendier and Nagel (2011), who show that past stock-market and bond-market experiences predict future risk taking of individual investors. Their data, the Survey of Consumer Finances, however, did not allow them to determine whether these effects are driven by beliefs (e.g., experiences of high stock returns make individuals more optimistic) or by endogenous preference formation (e.g., experiences of high stock returns make individuals less risk averse or lead to other changes in "tastes" for certain asset classes). The data used in this paper measures directly individual expectations and thus allows to focus specifically on the beliefs channel. Interestingly, the weighting of past experiences implied by the learning-from-experience rules estimated in this paper matches very closely the weighting scheme estimated from a completely different data source in Malmendier and Nagel (2011). Evidence consistent with learning-from-experience effects is also presented in Greenwood and Nagel (2009) and Vissing-Jorgensen (2003), who show that young mutual fund managers and young individual investors in the late 90s were more optimistic about stocks, and in particular technology stocks, than older investors, consistent with young investors being more strongly influenced by their recent good experience with technology stocks. Vissing-Jorgensen also points out that there is age-heterogeneity of inflation expectations in the late 1970s and early 1980s. Kaustia and Knüpfer (2008) and Chiang, Hirshleifer, Qian, and Sherman (2011) find that investors' participation decision and bidding strategies in initial public offerings is influenced by extrapolation from previously experienced IPO returns.

The rest of the paper is organized as follows. Section 2 introduces our experienced-based learning framework and estimation approach. Section 3 discusses the data set on inflation expectations. Section 4 presents our core set of results on learning-from-experience effects in inflation expectations. In Section 5, we look at the implications of our results at the aggregate level. Section 6 concludes with some final thoughts.

2 Learning from experience

Consider two individuals, one is member of the cohort born at time s, and the other belongs to the cohort born at time $s+j$. At time $t > s+j$, how do they form expectations of next period's inflation, π_{t+1}? The essence of the learning-from-experience hypothesis is that when these two individuals forecast π_{t+1}, they place different weights on recent and distant historical data, reflecting the different lengths of the inflation histories they have experienced in their own lives so far. The younger individual, born at $s+j$, has experienced a shorter data set, and is therefore more strongly influenced by recent data. As a result, the two individuals may produce different forecasts at the same point in time.

Our analytical framework builds on the forecasting rules proposed in the adaptive learning literature, in particular Marcet and Sargent (1989). (See also Sargent (1993) and Evans and Honkapohja (2001).). The key departure from the standard adaptive-learning models is that we allow individuals to put more weight on data experienced during their lifetimes than on other historical data. Thus, the adaptive component of forecasting gives rise to cross-sectional differences in expectations between different cohorts, depending on their life-time inflation experiences.

We model the perceived law of motion that individuals are trying to estimate as an AR(1) process, as, for example, in Orphanides and Williams (2004):

$$\pi_{t+1} = \alpha + \phi \pi_t + \eta_{t+1}. \tag{1}$$

Individuals estimate $b \equiv (\alpha, \phi)'$ recursively from past data following

$$b_t = b_{t-1} + \gamma_t R_t^{-1} x_{t-1} (\pi_t - b'_{t-1} x_{t-1}) \tag{2}$$

$$R_t = R_{t-1} + \gamma_t (x_{t-1} x'_{t-1} - R_{t-1}), \tag{3}$$

where the recursion is started at some point in the distant past. (We will see below that, in our specific setting, past data gets downweighted sufficiently fast that initial conditions do not exert any relevant influence.)

The sequence of gains γ_t in the recursive algorithm determines the degree of updating when faced with an inflation surprise. For example, with $\gamma_t = 1/t$, the algorithm represents a recursive formulation of an ordinary least squares estimation that uses all data available until time t with equal weights (see Evans and Honkapohja (2001)). With γ_t set to a constant, it represents a constant-gain learning algorithm, which weights past data with exponentially decaying weights. Our key modification of the standard learning framework is that we let the gain parameter depend on the age $t-s$ of the members of the cohort s. As a result, individuals of different age can be heterogeneous in their forecasts and they adjust their forecasts to different degrees in response to surprise inflation. Given the perceived law of motion in equation (1), these cross-sectional differences can arise from two sources: first, from differences in individuals' perception of the mean, $\mu = \alpha(1-\phi)^{-1}$, and, second, from differences in the perception of persistence, ϕ, of deviations of recent inflation from this perceived mean.

Specifically, we consider the following decreasing-gain specification,

$$\gamma_{t,s} = \begin{cases} \frac{\theta}{t-s} & \text{if } t-s \geq \theta \\ 1 & \text{if } t-s < \theta, \end{cases} \tag{4}$$

where $\theta > 0$ is a constant parameter that determines the shape of the implied function of weights on past experienced inflation observations. We let the recursion start with $\gamma_{t,s} = 1$ for $t-s < \theta$, which implies that data before birth is ignored. (As explained below, our econometric specification does allow for all available historical data to affect the forecast, but isolates the effect of data realized during individuals' life-times on expectation formation.) This specification is the same as in Marcet and Sargent (1989) with one modification: the gain here is decreasing in age, not in time, and individuals use only data realized during their life-times.

Figure 1 illustrates the role of the parameter θ, Conceptually, we want to allow for the possibility that experiences in the distant past have a different influence than more recent experiences. For example, the memory of past episodes of high inflation might fade away over time, as also implied by standard models such as constant-gain learning. Alternatively, high-inflation experiences at young age, perhaps conveyed through the worries of parents, might leave a particularly strong impression and have a lasting impact on the formation of beliefs about future inflation. The top graph of Figure 1 presents the sequences of gains γ as a function of the age of the individual for different values of θ. Regardless of the value of θ, gains decrease with age. This is a sensible assumption in the context of the learning-from-experience hypothesis. Young individuals, who have experienced only a small set of historical data, presumably have a higher gain than older individuals, who have experienced a longer data history, and for whom a single inflation surprise observation should have a weaker marginal impact on their estimates of the the inflation process parameters. The down-weighting of past data is also consistent with the corresponding assumption in constant-gain learning models. There, the motivations for assuming such down-weighting of past data are, for example, that indviduals believe that a structural break may have occurred or that they perceive the parameters of the inflation process to be time-varying. Our model adds additional micro-foundation to that assumed pattern.

The top graph of Figure 1 also illustrates that the higher θ is, the slower is the rate at which the gains decrease with age and, hence, the less weight is given to observations that are more distant in the past. The latter implication is further illustrated in the bottom graph of Figure 1, which shows the implied weights on past inflation observations as a function of the time lag relative to current time t for the example of 50-year (200 quarters) old individual. For $\theta = 1$, all historical observations since birth are weighted equally. For $\theta > 1$, instead, weights on earlier observations are lower than those on more recent observations. With $\theta = 3$ very little weight is put on observations in the first 50 quarters since birth towards the right of the bottom graph.

In other words, our gain parametrization is quite flexible in accommodating different weighing schemes. The weights can be monotonically increasing, decreasing, or flat. An additional advantage of the decreasing-gain specification in equation (4) is that, for appropriate choices of the weighting parameters. it produces weight sequences that are virtually identical to those in Malmendier and Nagel (2011). (See Appendix D.) This allows us to compare the

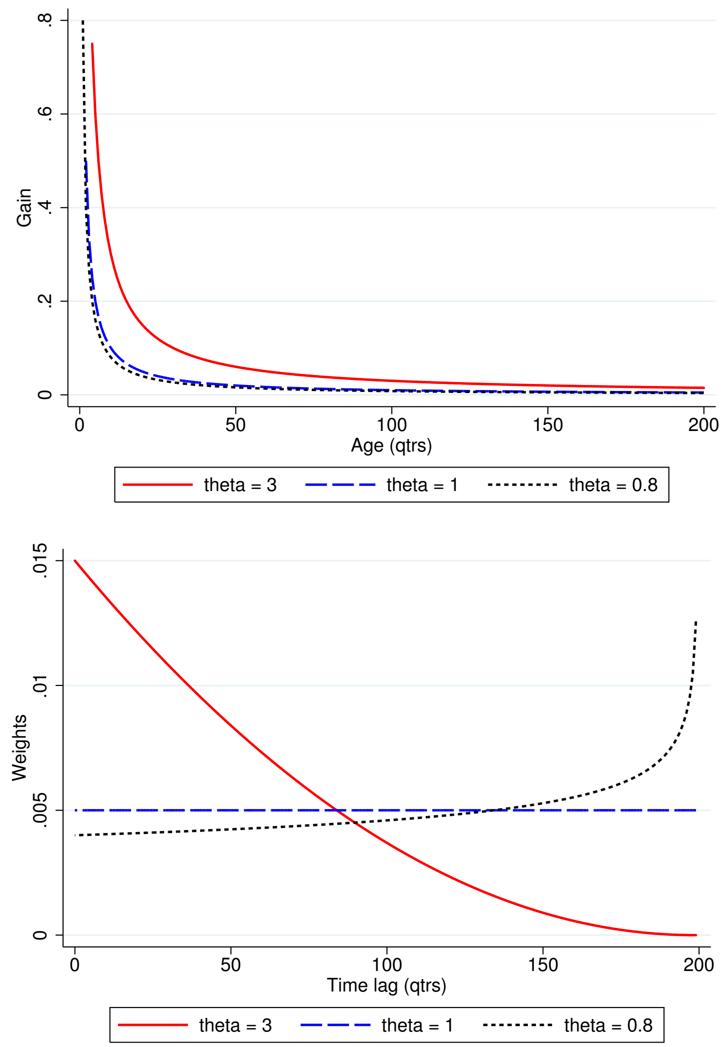

Figure 1: Examples of gain sequences (top) and associated implied weighting of past data (bottom)

experience-based weights implied by our estimates of θ from inflation expectations data, with the earlier evidence in Malmendier and Nagel (2011) where the weighting scheme is estimated from data on portfolio allocations.

In addition to the influence of past inflation experiences, we allow other information sources to affect the formation of inflation expectations. Let $\pi_t^h = h^{-1} \sum_{h=0}^{h-1} \pi_{t-i}$ denote the h-period average inflation rate (with both π_t and π_t^h measured at annual rates). Let $\pi_{t+h|t,s}^h$ denote the forecast of the average (annualized) inflation rate over the next h periods made by cohort s at time t, where subscript $|t,s$ denotes that a forecast was made using information available to agents of cohort s at time t and where the superscript h denotes the forecast horizon. Individuals' one-step ahead adaptive learning forecast of the experience-based component of inflation is obtained as $\tau_{t+1|t}^1 = b_t' x_t$, and multi-period forecasts of the epxerience based-component $\tau_{t+h|t}^h$ are obtained by iterating on the forecasting model at the time-t estimates of the model parameters. We capture the influence of information sources other than experienced inflation by assuming

$$\pi_{t+h|t,s}^h = \beta \tau_{t+h|t,s}^h + (1-\beta) f_t^h. \tag{5}$$

That is, the subjective expectation is a weighted average of the learning-from-experience component $\tau_{t+h|t,s}^h$ and an unobserved common component f_t^h of individuals' h-period forecasts. This unobserved component f_t^h could represent any kind of forecast based on common information available to all individuals at time t, such as the opinion of professional forecasters or the representation of their opinions in the news media (e.g., as in Carroll (2003)). Alternatively, f_t^h could capture a common component of individual forecasts that is driven by all available historical data. In either case, the coefficient β captures the incremental contribution of life-time experiences $\tau_{t+h|t,s}^h$ to $\pi_{t+h|t,s}^h$ over and above thes common components. Thus, we do not assume that individuals only use data realized during their life-times, but isolate empirically the incremental effect of life-time experiences on expectations formation.

Empirically, we estimate the following modification of equation (5):

$$\tilde{\pi}^h_{t+h|t,s} = \beta \tau^h_{t+h|t,s} + \delta^{h\prime} D_t + \varepsilon^h_{t,s}, \tag{6}$$

where $\tilde{\pi}^h_{t+h|t,s}$ denotes measured inflation expectations from survey data. In this estimating equation, we absorb the unobserved f^h_t with a vector of time dummies D_t. We also add the disturbance $\varepsilon^h_{t,s}$, which we assume to be uncorrelated with $\tau^h_{t+h|t,s}$, but which is allowed to be correlated over time within cohorts and between cohorts within the same time period. It captures, for example, measurement error in the survey data and idiosyncratic factors influencing expectations beyond those explicitly considered here. We use this specification to jointly estimate θ and β with non-linear least squares. (Recall that $\tau^h_{t+h|t,s}$ is a non-linear function of θ.)

The presence of time dummies in Eq. (6) implies that we identify β and θ, and hence the learning-from-experience effect on expectations, from cross-sectional differences between the subjective inflation expectations of individuals of different ages, and from the evolution of those cross-sectional differences over time. The cross-sectional identification allows to rule out confounds affecting prior work, which has estimated adaptive learning rules from aggregate data, e.g., time-series of mean or median inflation expectations. Under the prior approach, it is hard to establish whether the time-series relationship between inflation expectations and lagged inflation rates truly reflects adaptive learning rules, or whether the expectations implied by adaptive learning just happen to be highly correlated with the expectations implied by some other formation mechanism (e.g., rational expectations). In contrast, the model of experience-based learning makes a clear prediction about the cross-section: Expectations should be heterogeneous by age, and for young people they should be more closely related to recent data than for older people. Moreover, we can estimate the gain parameter θ that determines the learning speed from this cross-sectional heterogeneity. This provides a new source of identification for the learning speed in adaptive learning algorithms.

Finally, it is worth emphasizing that, despite the formal similarities in learning algorithms between adapative and experience-based models (other than the dependence on age), the underlying interpretation is different. In the adaptive learning literature, the use of relatively simple learning algorithms is motivated by the fact that economic agents face cognitive and computational constraints which limit their ability to use optimal forecasts. The algorithms are viewed as an approximation of the "rules of thumb" that practitioners and individuals might employ to form their expectations. The focus of much of the adaptive learning literature is on the conditions under which such simple learning rules can lead to convergence to rational expectations. Our objective, instead, is to use the simple recursive least-squares learning framework as a starting point for an empirical investigation of individuals' actual forecasting rules. Correspondingly, we depart from the standard adaptive learning algorithms and introduce age-dependece in order to allow for learning-from-experience effects.

3 Data

To estimate the learning-from-experience model, we use long-term historical data on the consumer price index (CPI). Our survey data starts in 1953. In order to fully capture experienced inflation, even for the oldest individuals in the survey sample, we need inflation data stretching back 75 years before that date. We use CPI data from Shiller (2005), available on Robert Shiller's website from 1871 until the end of 2009. to calculate annualized quarterly log inflation rates. To illustrate the long-run variation in inflation rates, Figure 2 shows annual inflation rates from this series.

The inflation expectations microdata is from the Reuters/Michigan Survey of Consumers (MSC), conducted by the Survey Research Center at the University of Michigan. These surveys were administered since 1953, initially three times per year, then quarterly from 1960 through 1977, and monthly since 1978 (see Curtin (1982)). We obtain data for surveys conducted from 1953 to 1977 from the Inter-university Consortium for Political and Social Research (*ICPSR*) at the University of Michigan. From 1959 to 1971, the questions of the winter-quarter Survey of Consumer Attitudes were administered as part of the Survey of

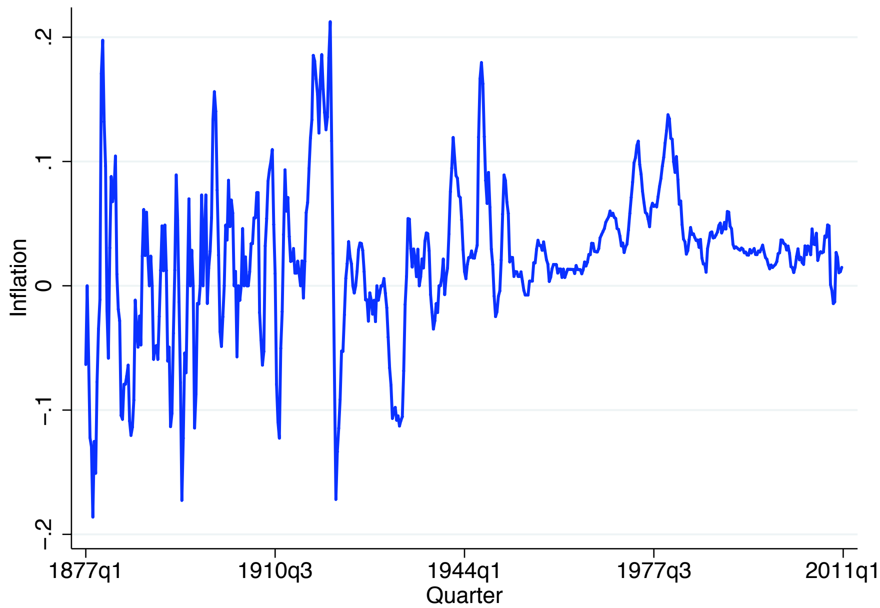

Figure 2: Annual CPI inflation rates

Consumer Finances (*SCF*), and so we obtain those data from the *SCF* files at *ICPSR*. The data from 1978 to 2007 is available in from the University of Michigan Survey Research Center.

In most periods, survey respondents are asked two questions about expected inflation, one about the expected direction of future price changes ("up", "same", or "down") and one about the expected percentage of price changes. In many periods, consumers are asked these two questions both for a one-year horizon and for a five-to-ten year horizon. Our analysis aims to make quantitative predictions and thus focuses on percentage expectations about future inflation, typically for the one-year horizon. Figure 3 highlights the periods in which we have percentage expectations data for the one-year horizon. Those quarters are shaded in light grey. Quarters in which the survey asked only the categorical questions about the expected direction are shaded in dark grey. In those quarters we impute percentage responses from the categorical responses. The imputation procedure is described in detail in Appendix B.

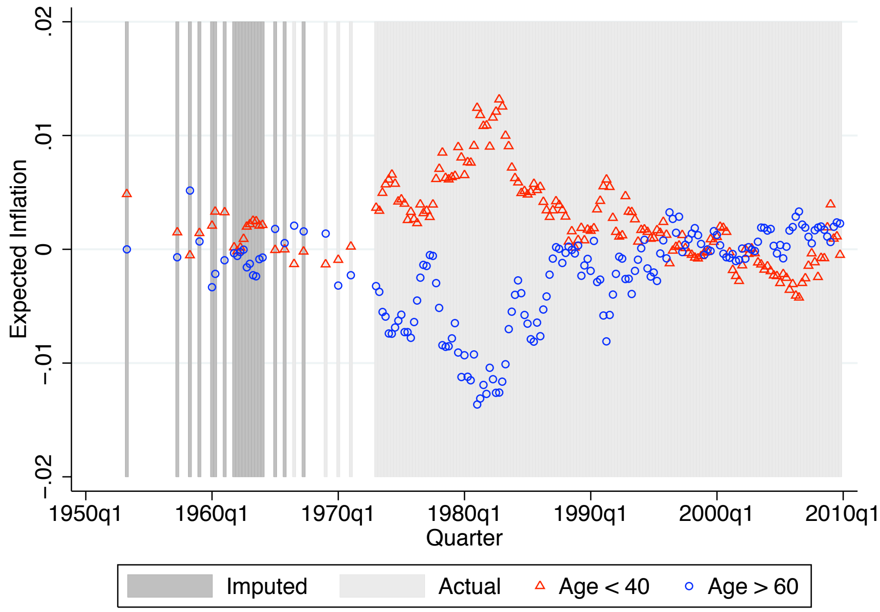

Figure 3: Four-quarter moving averages of mean one-year inflation expectations of young (age < 40) and old (age > 60) in excess of the full-sample cross-sectional mean expectation. Percentage forecasts are available in light shaded periods, they are imputed from categorical responses in dark shaded periods, and unavailable in unshaded periods.

Since our learning-from-experience hypothesis predicts that inflation expectations should be heterogeneous across different age groups, we aggregate the data at the cohort level, i.e., by birth year. For each survey month and each cohort, we compute the mean inflation expectations of all members of the cohort. In the computation of this mean, we apply the sample weights provided by the MSC. If multiple monthly surveys are administered within the same quarter, we average the monthly means within each quarter to make the survey data compatible with our quarterly inflation rate series.

We restrict our sample to respondents whose age ranges from 25 to 74. This means that for each cohort we obtain a quarterly series of inflation expectations that covers the time during which members of this cohort are from 25 to 74 years old.

To provide some sense of the variation in the data, Figure 3 plots the average inflation expectations of young individuals (averaging across all cohorts that are in the age range from

25 to 39) and old individuals (averaging across cohorts that are in the age range 61 to 75), relative to the full-sample cross-sectional mean expectation at each point in time. To better illustrate lower frequency variation, we plot the data as four-quarter moving averages. The dispersion across age groups widens to almost 3 percentage points (pp) during the high inflation years of the 1970s and early 1980s. The fact that young individuals at the time expected higher inflation is consistent with the learning-from-experience story: The experience of young individuals around 1980 was dominated by the recent high-inflation years, while older individuals' experience also included the modest inflation rates of earlier decades. For younger individuals, with a smaller set of experienced inflation data points, these recent observations exert a stronger influence on their expectations. As we show below, differences in the perception of inflation persistence between young and old matter as well, not just differences in the level of inflation rates they experienced in the past.

4 Estimation of learning-from-experience effects from expectations heterogeneity

We now estimate the learning-from-experience effects by fitting the estimating equation (6) and the underlying AR(1) model to the MSC inflation expectations data, using nonlinear least squares on the data aggregated at the cohort level. We relate survey expectations measured in quarter t to learning-from-experience forecasts $\tau^h_{t+h|t,s}$, where we assume that the data available to individuals in constructing $\tau^h_{t+h|t,s}$ are quarterly inflation rates until the end of quarter $t-1$. To account for possible serial correlation of residuals within cohorts and correlation between cohorts within the same time period, we report standard errors that are robust to two-way clustering by cohort and calendar quarter.

Table 1 the estimation results. Using the full sample, our estimate of the gain parameter in column (1) is $\theta = 3.006$ (s.e. 0.249). Comparing this estimate of θ with the illustration in Figure 1 one can see that the estimate implies weights that are declining a bit faster than

Table 1: Explaining heterogeneity inflation expectations with learning from experience

Each cohort is assumed to recursively estimate an AR(1) model of inflation, with gain decreasing with age, using quarterly annualized inflation rate data up to the end of quarter $t-1$. The table reports the results of non-linear least-squares regressions of one-year inflation expectations in quarter t on these learning-from-experience forecasts. Standard errors reported in parentheses are two-way clustered by time and cohort. The sample period runs from 1953 to 2009 (with gaps).

	Baseline		Restricted		
	(1)	(2)	(3)	(4)	(5)
Gain parameter θ	3.006	3.097	3.991	4.192	3.006
	(0.249)	(0.272)	(0.609)	(0.445)	-
Sensitivity β	0.647	0.650	0.670	1.000	1.000
	(0.074)	(0.077)	(0.083)	-	-
Time dummies	Yes	Yes	No	No	No
Imputed data included	Yes	No	No	No	No
Restrictions			$f_t = SPF_{t-1}$	$\beta = 1$	$\beta = 1$ $\theta = 3.006$
Adj. R^2	0.636	0.634	-	-	-
RMSE	0.0148	0.0152	0.0189	0.0192	0.0195
#Obs.	8165	7600	7400	7600	7600

linearly. The estimation results in column (1) also show that there is a strong relationship between the learning-from-experience forecast $\tau^h_{t+h|t,s}$ and measured inflation expectations, captured by the sensitivity parameter β, which we estimate to be 0.647 (s.e. 0.074). This magnitude of the β parameter implies that when two individuals differ in the weighted-average inflation experienced during their life time by 1 pp, their one-year inflation expectations differ by 0.647 pp on average.

The presence of the time dummies in these regressions is important to rule out that the estimates pick up effects unrelated to learning from experience. If individual expectations were unaffected by inflation experiences – for example, because all individuals learned from the same historical data set in the same way, applying the same forecasting rules — then all the effect of historical inflation rates, including "experienced" inflation rates, on current forecasts would be picked up by the time dummies and β would be zero. The fact that β is not equal to zero is direct evidence that differences in experienced inflation histories are correlated with differences in expectations. The positive β-estimate also implies that recent observations exert a stronger influence on expectations of the young, because their set of experienced historical inflation rates comprises only relatively few observations.

To check whether the imputation of percentage responses from categorical responses has any influence on the results, we re-run the estimation without the imputed data, using only those time periods in which percentage responses are directly available. The results are presented in column (2). As can be seen, whether or not imputed data is used has little effect on the results. We estimate a similar gain parameter, $\theta = 3.097$ (s.e. 0.272), and a similar sensitivity paramenter $\beta = 0.647$ (s.e. 0.074).

Interestingly, the weighting of past inflation experiences implied by the point estimates of θ is similar to the weighting implied by the estimates obtained in Malmendier and Nagel (2011) by relating data on household asset allocation to experienced risky asset returns.[1] This is quite remarkable since the data on inflation expectations is drawn from a completely different data set, and since we look at beliefs about inflation rather than asset allocation choices. Despite those differences, the dependence on life-time macroeconomic history in both cases seems to imply a similar weighting of experienced data, suggesting that a common expectations-formation mechanism may be driving all of these results.

One possible alternative theory for these (time-varying) age-related differences in inflation expectations is that different age groups consume different consumption baskets, and that individuals form inflation expectations based on the (recent) inflation rates they observe on their age-specific consumption baskets. The concern would be that these inflation differentials between age-specific consumption baskets could be correlated with differences in age-specific learning-from-experience forecasts that we construct. In other words, inflation differentials

[1] The weighting function in Malmendier and Nagel (2011) is controlled by a parameter λ which relates to θ as $\theta \approx \lambda + 1$ (see Appendix D), and which is estimated to be in the range from 1.1 to 1.9 depending on the specification.

between age-specific consumption baskets could be a correlated omitted variable. To address this issue, we re-run the regressions in Table 1 controlling for differences between inflation rates on consumption baskets of the elderly and overall CPI inflation rates. We measure the inflation rates of the elderly with the experimental CPI for the elderly series (CPI-E) provided by the Bureau of Labor Statistics. The results reported in Appendix C show that this does not affect our results. The cross-sectional differences that we attribute to learning-from-experience effects are not explained by differences in age-specific inflation rates.

Columns (3) to (5) in Table 1 consider restricted versions of the estimating equation (5). The specification in (5) with time dummies is useful to cleanly demonstrate the existence of the learning-from-experience effect with a test of the null hypothesis $\beta = 0$. It also allows estimation of θ purely from cross-sectional differences, removing potentially confounding unobserved factors that also affect expectations. On the other hand, it would be useful to know to what extent variation in the levels of expectations rather than just cross-sectional differences can be explained with the learning-from-experience forecasts.

The specification in column (3) explores which factors may be captured by the time dummies in (5). A likely possibility is that individuals put some weight on the opinions of professional forecasters when these forecasts get disseminated in the media. To check this, we remove the time dummies and intercept and use the Survey of Professional Forecasters (SPF) forecast of quarter $t-1$ as the common factor f_t in (6). We further restrict the coefficient on the SPF to be $1-\beta$ so that individuals' expectation is a weighted average of the learning-from-experience forecast and the SPF forecast. Without the time dummies, the estimation now uses information about levels in inflation expectations, not just cross-sectional differences, and so we remove the imputed data, because our imputation is only designed to impute cross-sectional differences, but not levels. The number of observations in column (3) is further slightly lower than in column (2) because SPF forecasts are not available in a few quarters early in the sample. As column (3) shows, replacing the time dummies with the SPF has little effect on the estimate of β compared with columns (1) and (2). With 3.991

(s.e. 0.272), the estimate of θ is higher, though. The shape of the weights on past inflation data implied by this point estimate is still quite similar to the shape implied by the estimates in columns (1) and (2). Since this regression is run without intercept, the adj. R^2 is not a useful measure of fit, but the RMSE shows that replacing time dummies with the SPF leads to a rather moderate decrease in explanatory power. Thus, the SPF seems to capture much of the factors absorbed by the time dummies in columns (1) and (2).

The specification in column (4) completely removes the unobserved factor f_t in (6) by restricting $\beta = 1$. Thus, it checks to what extent cross-sectional differences as well as average levels of inflation expectations can be explained with the learning-from-experience forecast alone. This is the most parsimonious specification, as it leaves only the parameter θ to be estimated. Remarkably, with 4.192 (s.e. 0.445) the estimate of θ is very close to the estimate in column (3). Judging by the RMSE, the fit is almost as good, too.

Column (5) explores the explanatory power of the learning-from-experience forecasts when θ is set to the point estimate from column (1), which one could regard as the cleanest estimate, as the time dummies in column (1) removed potentially confounding unobserved factors. We also restrict $\beta = 1$, so this column simply reports the fit at these parameters, without any further estimation. The RMSE is only slightly higher than in column (4). This underscores that the higher θ in column (4) does not lead to big differences in the resulting learning-from-experience forecasts.

To get a better sense of the extent to which learning-from-experience effects explain cross-sectional differences in inflation expectations, Figure 4 presents some plots of fitted values for different age groups. Panel (a) is based on the baseline estimates from column (1) in Table 1, Panel (b) reports the fitted values from the restricted model in column (4).

For the purpose of these plots, we average inflation expectations and the fitted values within the same young (age < 40) and old (age > 60) categories that we used earlier in Figure 3. Since our baseline estimation with time dummies focuses on cross-sectional differences, we plot the inflation expectations and fitted values of these subgroups after subtracting the full-

(a) Baseline

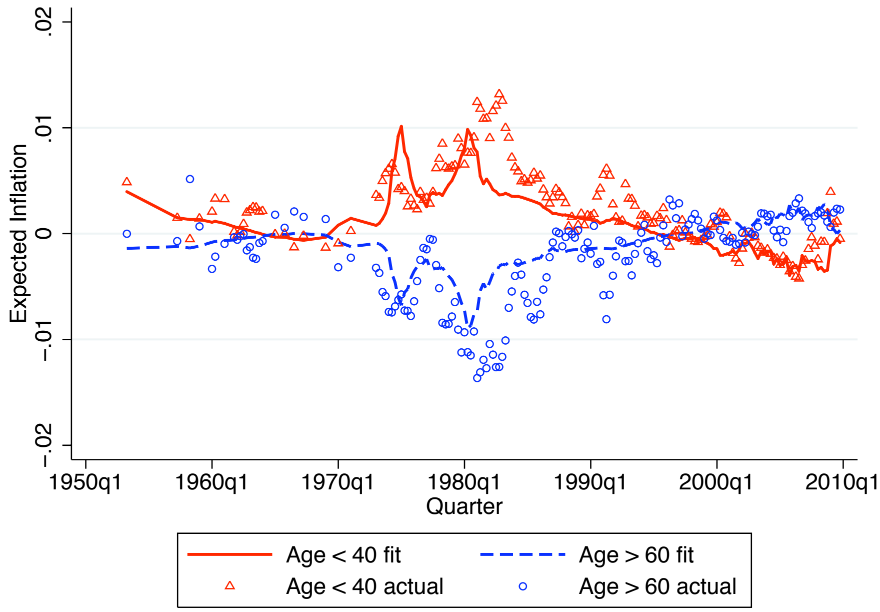

(b) Restricted: $\beta = 1$

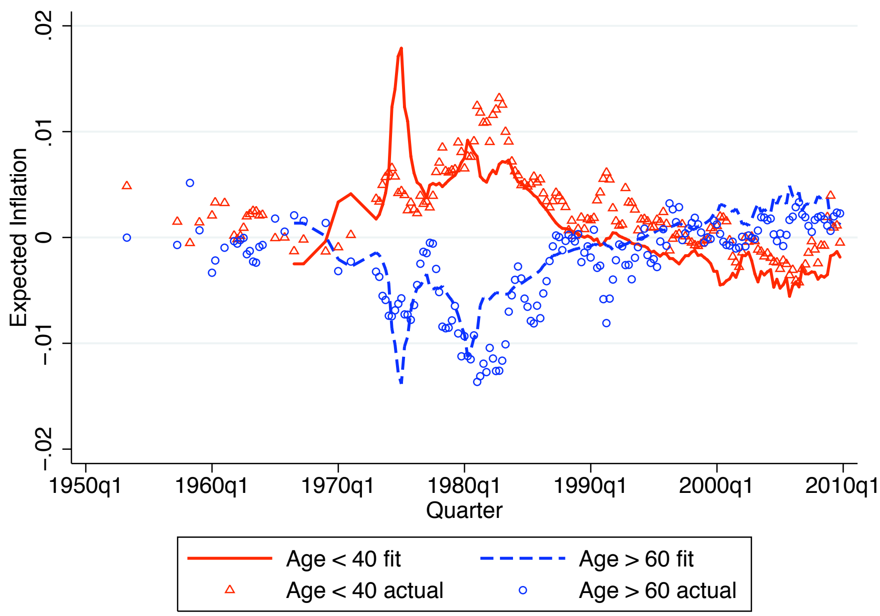

Figure 4: AR(1) model: Comparison of 4-quarter moving averages of actual and fitted 1-year inflation expectations for young and old in excess of the full-sample cross-sectional mean expectation. Panel (a) corresponds to column (1) and Panel (b) corresponds to column (4) in Table 1.

sample mean each period. Thus, the plots focus on cross-sectional differences. To eliminate high-frequency variation, we show 4-quarter moving averages for both actual and fitted values. Fitted values are drawn as lines, raw inflation expectations are shown as triangles (young) or circles (old).

The plot shows that the experience-based model does a good job of explaining the differences in inflation expectations between young and old. In particular, it accounts, to a large extent, for the large difference in expectations between young and old in the late 1970s and early 1980s. These plots also highlight that the unrestricted baseline model in Panel (a) and the restricted model in Panel (b) do roughly equally well overall in explaining age-heterogeneity in inflation expectations, but some features of the data are fit better by the baseline model, while some are better fit by the restricted model. For example, the mean reversion in the old-young gap in the early 1980s is too fast in the baseline model. The restricted model does better on this dimension. The restricted model, however, produces a spike in the old-young gap after the first inflation shock of the 1970s that is much bigger than the gap found in the data. The baseline model does better on this dimension.

Figure 5 reports the persistence and conditional mean inflation perceived by young and old over the course of the sample, as implied by the estimate of θ from Table 1, column (1). The figure shows that there has been first in increase and then a dramatic decline in the perceived persistence and the perceived mean inflation rates. Young individuals' views about mean and persistence are much more volatile than older individuals' views, as they are more strongly influenced by recent data. For example, our estimates imply that at the end of the sample period, young individuals' inflation expectations are well anchored at low expected inflation rates, as the perceived persistence is close to zero. Older individuals perceived inflation persistence, however, is still substantially above zero.

(a) Autocorrelation ϕ

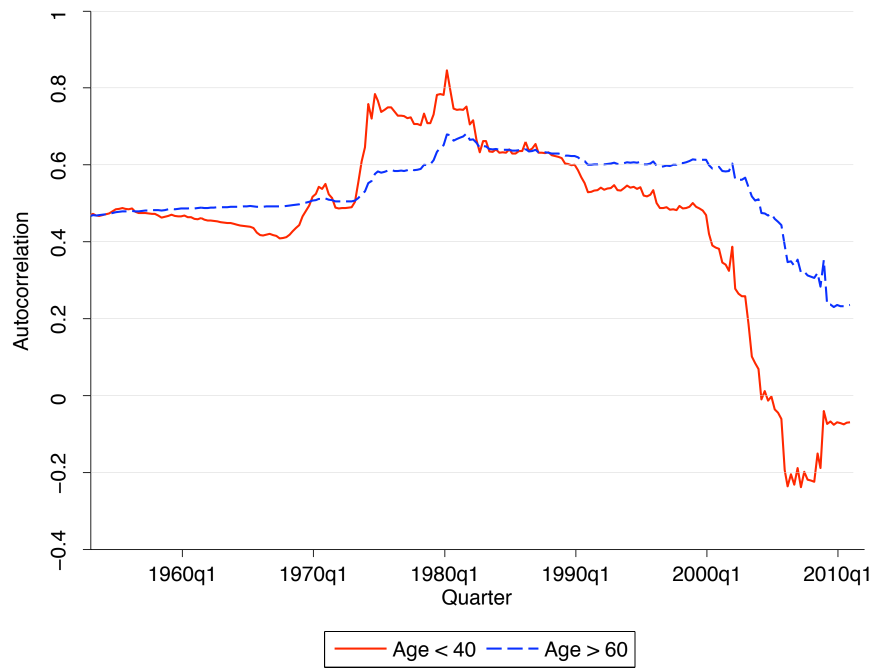

(b) Mean μ

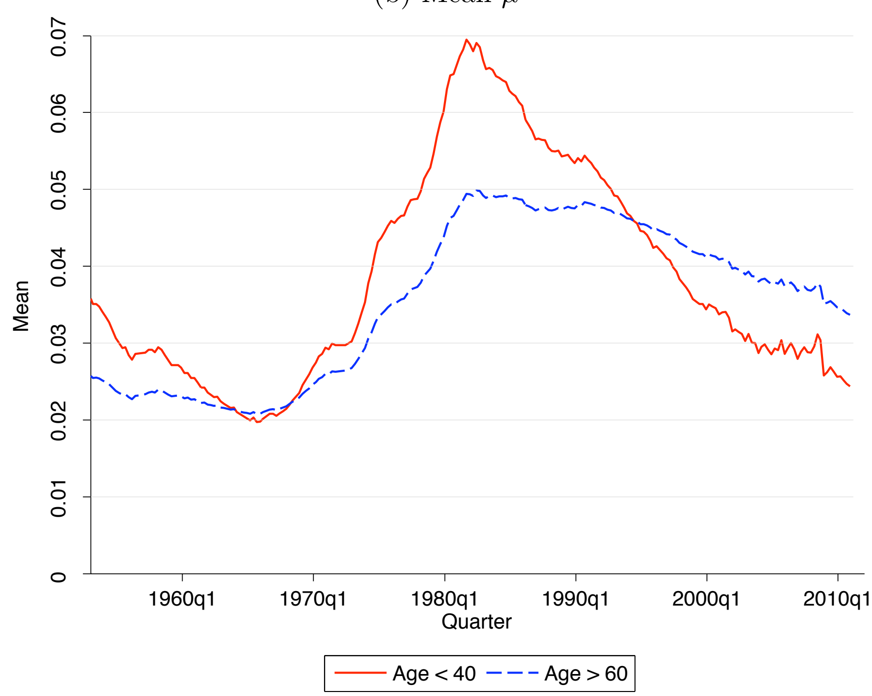

Figure 5: Learning-from-experience AR(1) model estimates (with $\theta = 3.006$) of autocorrelation (top) and mean inflation (bottom) for young and old.

5 Implications for inflation expectations in the aggregate

So far we have focused on understanding to what extent experienced inflation can help understand the formation of inflation expectations at the cohort level. From a macroeconomic perspective it would also be interesting to see to what extent the learning-from-experience mechanism, based on the estimates of θ from cross-sectional heterogeneity, helps explain inflation expectations in the aggregate. In this section we show that the learning-from-experience forecasts at the cohort level aggregate to average forecasts that closely resemble those from constant-gain algorithms that are popular in macroeconomics. We also show that one can extract components from the experience-based forecasts that are useful in predicting forecast errors in the Michigan survey (MSC) and the Survey of Professional Forecasters (SPF), as well as the returns on long-term bonds.

5.1 Approximating learning-from-experience with constant-gain learning

In our learning-from-experience framework, individuals update their expectations with *decreasing* gain: as individuals age, their experienced set of data expands and their expectations react less to a given inflation surprise than those of younger individuals. However, older individuals leave the population at some point and are replaced by younger ones. Hence, at any given point in time, there is a time-specific distribution of gains in the population, but to the extent that the age distribution is relatively stable, the average gain should be approximately constant. Therefore, the average forecast across all age groups can be approximated by a *constant*-gain learning algorithm where updating takes place in the same way as laid out in equations (1) to (3), but with the decreasing gain in (4) replaced by a constant gain, and with a single "representative" agent.

How well this works can be seen by comparing the average weights on past inflation data implied by the cohort-level learning-from-experience rules with the weights implied by constant-gain learning. The solid line in Figure 6 plots the average of implied weights on past inflation with learning from experience, where the average is taken (equal-weighted) across all cohorts alive in the population at a point in time. The implied weights are based on our point estimate of $\theta = 3.006$ from Table 1, column (1). We then look for a constant gain so that the weights on past data implied by this constant-gain algorithm minimize the squared deviations from the average learning-from-experience weights. The result is a constant gain of $\gamma = 0.0175$, with implied weights as shown by the dashed line. The figure shows that the weighting of past data is very similar. Thus, the implications of learning from experience for expectations formation in aggregate are likely to be very similar to those of the constant-gain learning algorithms that are common in macroeconomics (see, e.g., Orphanides and Williams (2005), Milani (2007), Evans and Honkapohja (2001)).[2]

There are two important differences, though. First, the motivation for the loss of memory of past data is different. In constant-gain learning, the gradual loss of influence of past data is typically motivated as a concern on part of agents that past data is not relevant anymore due to structural changes and time-variation in the parameters of the perceived law of motion. While these concerns may also be relevant in the learning-from-experience framework and lead to $\theta > 1$ so that recent data receives a higher weight than data realized earlier in life, learning from experience comes with the additional feature that memory of past data is lost as old generations die and new ones are born. In aggregate, data in the distant past would be downweighted even if each individual weighted all life-time experiences equally.

Second, as we demonstrated in the previous section, the gain parameter of the learning-from-experience rule can be estimated from cross-sectional data. Our estimate of θ is not fitted to aggregate expectations. The time dummies in our estimation absorb all variation in the cross-sectional average expectation, and so θ is identified from cross-sectional information only. In light of the fact that we did not employ aggregate expectations in estimation of θ and we did not calibrate θ to achieve the best fit to realized future inflation, it is remarkable that the constant gain $\gamma = 0.0175$ in Figure 6 that best matches the weights implied by our estimate of θ is virtually the same as the gains that seem to be required to match aggregate

[2]Cross-sectional heterogeneity in expectations between different cohorts could matter for other macroeconomic implications, though; see, e.g., Piazzesi and Schneider (2010).

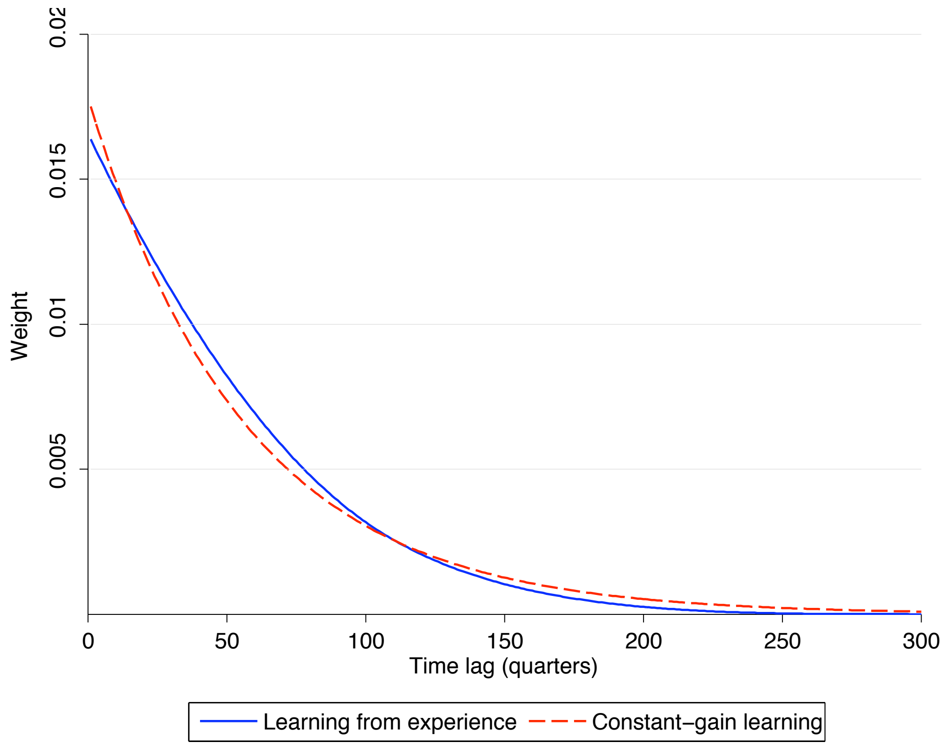

Figure 6: Implied aggregate weights for past inflation observations under learning from experience (equal-weighted average of weights across age groups at point estimate of $\theta = 3.006$ from Table 1, Panel A, column (1)) compared with implied weights under constant-gain learning by a single agent (with gain $\gamma = 0.0175$ that minimizes squared deviations from the aggregated learning-from-experience weights).

expectations and macro time-series data. For example, Milani (2007) estimates a DSGE model with constant-gain learning and obtains an estimate of 0.0183 that results in the best fit of the model to the macro time series employed in estimation. Orphanides and Williams (2005) choose a gain of 0.02 to match the time series of inflation forecasts from the Survey of Professional Forecasters (SPF). Thus, our estimation from cross-sectional heterogeneity between different cohorts brings in new additional data that provides "out-of-sample" support for values of the gain parameter in this range. This is particularly important because the identification of the learning speed in macro models from macro data is econometrically difficult (Chevillon, Massmann, and Mavroeidis (2010)).

5.2 Explaining the level of average inflation expectations

Figure 7 explores how well the average learning-from-experience forecast tracks the average 1-year survey expectations (i.e., the data we used in the estimation in Table 1 is now averaged across all cohorts each quarter). Since our imputation of percentage responses only targeted cross-sectional differences, but not the average level of percentage expectations, we omit all periods from these regressions in which we only have categorical inflation expectations data.

It is apparent that the average learning-from-experience forecast (calculated with $\theta = 3.006$ from Table 1, Panel A, column (1)), shown as the solid line, tracks the average survey expectations closely. It is important to keep in mind that this is by no means a mechanical result. Our estimation of θ used only cross-sectional differences in survey expectations between cohorts. It did not utilize any information about the level of the average survey expectation. Therefore, it could have been possible, in principle, that the θ that fits cross-sectional differences produces average forecasts that fail to match the level of average expectations. As the figure shows, though, we find that the two match well.

We also compare the average learning-from-experience forecast to a constant-gain-learning forecast (with $\gamma = 0.0175$ as in Figure 6), shown as the dashed line. Not surprisingly, given how similar the weights on past inflation data are for the two expectations-formation mechanisms (see Figure 6), the forecasts are almost indistinguishable. This provides further support for the idea that at the aggregate level, the learning-from-experience expectations formation mechanism can be approximated well with constant-gain learning.

Next, we compare the average learning-from-experience forecast to a sticky-information forecast. Sticky information, as in Mankiw and Reis (2002) and Carroll (2003) induces stickiness in expectations, and it is possible that our estimation of the learning-from-experience rule might be picking up some of this stickiness in expectations. We calculate sticky-information inflation expectations as in Carroll's model as a geometric distributed lag of current and past quarterly SPF forecasts of one-year inflation rates.[3] We set the weight parameter $\lambda = 0.25$

[3]We use the 1-year inflation forecasts that the SPF constructs from median CPI inflation forecasts for each

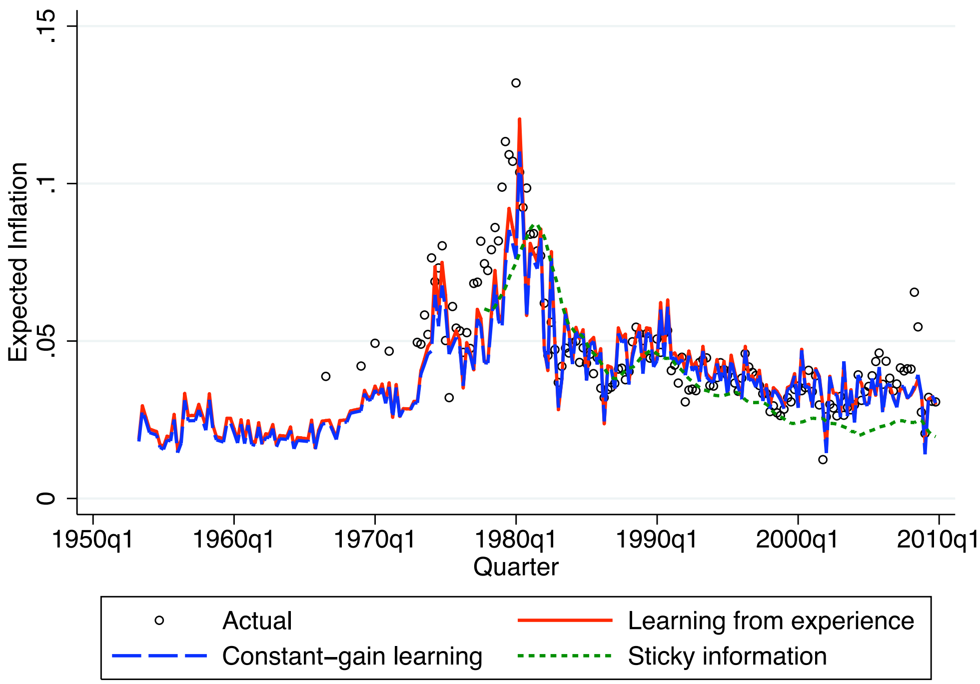

Figure 7: Average 1-year survey expectations (actual) compared with average learning-from-experience forecasts and with constant-gain and sticky-information forecasts.

as in Mankiw and Reis (2002) (Carroll (2003) estimates $\lambda = 0.27$). The resulting sticky-information forecast is shown as the short-dashed line in Figure 7.

In addition to the informal graphical comparison in Figure 7, Table 2 reports the results from a regression of quarter t average survey expectations on the learning-from-experience forecast in quarter t. Column (1) shows that with 0.893 the coefficient on the learning-from-experience forecast is very close to one, and less than one standard error away from it. With 57.0% the adj. R^2 is high. This is another confirmation of the fact that the learning-from-experience forecast tracks the actual average survey expectations very closely. Not surprisingly, given the similarity of average learning-from-experience forecasts and constant-gain learning forecasts, using the constant-gain learning forecast as explanatory variable in column (2) produces almost identical results. The explanatory power of the sticky-information

of the four quarters ahead. Before 1981Q3, when the CPI inflation forecast series is not available, we use the GDP deflator inflation forecast series.

Table 2: Explaining mean inflation expectations

OLS regressions with quarterly data from 1973Q1 to 2009Q4 (with gaps). The dependent variable is the forecast of one-year inflation made during quarter t, averaged across all cohorts. Newey-West standard errors (with 5 lags) are shown in parentheses.

	(1)	(2)	(3)	(4)
Learning-from-experience forecast	0.893			0.707
	(0.124)			(0.130)
Constant-gain-learning forecast		0.943		
		(0.144)		
Sticky-information forecast			0.877	0.385
			(0.202)	(0.145)
Intercept	0.009	0.008	0.011	-0.001
	(0.005)	(0.005)	(0.006)	(0.005)
Adj. R^2	0.570	0.557	0.597	0.713
#Obs.	172	172	129	129

forecast in column (3) is lower. Adding the sticky-information forecast as an explanatory variable along with the learning-from-experience forecast in column (4) lowers the coefficient on the learning-from-experience forecast a little, but the effect is small. This shows that the learning-from-experience forecast does not just pick up the sticky-information effect of Mankiw and Reis (2002) and Carroll (2003).

5.3 Predictability of forecast errors

Adaptive learning may lead to predictable and persistent forecast errors (from the econometrician's perspective). If such forecast errors do not cancel out in the aggregate, they can influence macroeconomic outcomes. We therefore now turn our attention to the question whether we can link the learning-from-experience behavior to predictability of level of average forecast errors.

That learning-from-experience can lead to predictable and persistent forecast errors can be seen in the following simplified example. Consider first the simple mean model with a time-varying mean, $\pi_{t+1} = \mu_t + \eta_{t+1}$, as the true as well as the perceived model of inflation. The average one-step ahead learning-from-experience forecast results in the forecast error

$$\pi^1_{t+1|t} - \pi_{t+1} = \mu_{|t} - \mu_t - \eta_{t+1}. \tag{7}$$

Now consider an econometrician who analyzes subjective expectations data ex-post with data available. If μ_t is equal to a constant μ, the econometrician can, with a sufficiently large sample (which is not restricted to the $[s,t]$ interval that learning-from-experience agents in cohort s are learning from), approximately observe the true mean. Any fluctuations of $\mu_{|t}$ around μ translate, predictably and one-to-one, into forecast errors. Regressing $\pi^1_{t+1|t} - \pi_{t+1}$ on $\mu_{|t}$ would yield a coefficient of one, with the second term in (7) absorbed by the intercept. If μ_t is time-varying, $\mu_{|t}$ is likely to have positive correlation with μ_t which lowers the regression coefficient. Of course, it is also possible that agents have some biases in their forecasts that influence the coefficient upwards.

In the case of a true and perceived AR(1) model for inflation with time-varying parameters, $\pi_{t+1} = \mu_t + \phi_t(\pi_t - \mu_t) + \eta_{t+1}$, the situation is more complicated. The one-step ahead forecast error in the average learning-from-experience forecast is given by

$$\pi^1_{t+1|t} - \pi_{t+1} = \mu_{|t}(1 - \phi_{|t}) - \mu_t(1 - \phi_t) + \phi_{|t}\pi_t - \phi_t\pi_t. \tag{8}$$

If μ_t and ϕ_t are constant, regression of $\pi^1_{t+1|t} - \pi_{t+1}$ on $\mu_{|t}(1 - \phi_{|t})$, $\phi_{|t}\pi_t$, and π_t produces a coefficient of one on the first two variables, and a coefficient of ϕ on the third. The second term in (8) is absorbed by the intercept. If μ_t and ϕ_t are time-varying, this can result in lower coefficients on the first two variables, just like in the simple mean model above, but, in addition, regression coefficients here can also be impacted by correlation between the various terms in (8).

We now run these regressions in our data. We work with 1-year ahead forecasts ($h = 4$ quarters). The multi-period expression corresponding to the right-hand side of equation (8)

Table 3: Predictability of average forecast errors

OLS regressions with quarterly data from 1973Q1 to 2009Q4 (with gaps). The dependent variable is the forecast of 1-year inflation made during quarter t, averaged across all cohorts, minus the inflation rate realized over the 12 months following the interview month. Newey-West standard errors (with 5 lags) are shown in parentheses. Out-of-sample (OOS) forecasts for the OOS tests at the bottom of each panel are constructed recursively, with an initial minimum window size until 1976Q3 (20 observations), except for column (3), where the initial window extends until 1989Q4.

	Full sample		Post-1989	SPF
	(1)	(2)	(3)	(4)
Mean component	2.553	2.432	1.526	3.203
	(0.750)	(0.735)	(0.864)	(1.004)
AR component		-0.380	-0.839	-0.049
		(0.318)	(0.648)	(0.337)
Lagged inflation		0.090	0.199	-0.062
		(0.112)	(0.098)	(0.112)
Intercept	-0.073	-0.068	-0.036	-0.096
	(0.023)	(0.023)	(0.024)	(0.031)
Adj. R^2	0.222	0.243	0.158	0.288
#Obs.	152	152	80	148
OOS RMSE with constant only	0.019	0.019	0.018	0.020
OOS RMSE with constant and predictor(s)	0.017	0.017	0.016	0.017
Diebold-Mariano one-tailed p-value	0.019	0.011	0.044	0.052

can be obtained by iterating on the AR(1) model. The three predictors in this multi-period case are $\mu_{|t}\left(1-\sum_{i=1}^{4} i^{-1}\phi_{|t}^{i}\right)$, which we label as the mean component, $\left(\sum_{i=1}^{4} i^{-1}\phi_{|t}^{i}\right)\pi_t$, which we label as the AR component, and π_t. The $\mu_{|t}$ and $\phi_{|t}$ parameter estimates are averages of the parameter estimates across all cohorts at time t, where we computed the cohort-level estimates from the learning-from-experience rule with $\theta = 3.006$ as in Table 1. In the computation of the average survey expectation on the left-hand side (from which we subtract the realized four-quarter inflation rate π_{t+4}^{4}), we take care to first align individuals' reported expectations with realized inflation rates by interview month, i.e., we align it with the inflation rates realized over the 12 months following the interview month.

Table 3 presents the results. As column (1) shows, there is a strong positive relationship between the mean component of the learning-from-experience forecast at time t and average inflation forecast errors of the participants in the Michigan survey during the forecast period t to $t+4$. The coefficient estimate of 2.553 (s.e. 0.750) is greater than one, which suggests that the regression picks up not only forecast errors induced by learning, but also other errors over and above the error induced by learning. The point estimate is just about two standard errors above one, though, so a coefficient of one is still within the likely range of possible values that one might find if one had a larger sample. The adj. R^2 of 22.2% indicates that predictability of forecast errors is substantial. Column (2) adds the AR component and the lagged inflation rate π_{t-1} as predictors,[4] but both of these are not significant, neither statistically nor in terms of their incremental explanatory power. To check whether all the predictability is driven by the high-inflation periods around 1980, the regression reported in column (3) is run with the sample restricted to the post-1989 period. The coefficient on the mean component is lower, but the adj. R^2 of 15.8% still indicates substantial predictability. Evidently, the forecast error predictability is not just limited to the high-inflation periods.

Another interesting issue is to what extent inflation expectations of $professional$ forecasters mirror the predictability that we find in individuals' forecast errors in the MSC. Individuals' forecast errors may be more significant for individuals' decisions (e.g., household investment decisions, labor market choices), while professional forecasts may be more relevant for asset pricing in financial markets. For this reason, column (4) reports results from a regression where we use the forecast error from the SPF as dependent variable. The results are similar to those with the MSC data in column (2): A large coefficient on the mean component, close-to-zero coefficients on the AR component and lagged inflation, and an adj. R^2 greater than 20%. Thus, the forecasts of professionals exhibit similar forecast error predictability.

Our focus so far has been on tests of in-sample predictability. To check for predictability induced by learning along the lines discussed above, this is the appropriate perspective.

[4]As before, we assume here that forecasts in quarter t are made with information up to end of quarter $t-1$, and so $\mu_{|t}$, $\phi_{|t}$, and the lagged inflation rate are also calculated from inflation rates up to quarter $t-1$.

Learning does not necessarily induce predictability of forecast errors out-of-sample (OOS), though (although it might, to the extent that rationality is bounded and individuals discard information, as in learning from experience, or use information in suboptimal ways, or work with misspecified models). In addition to shedding light on individuals' expectations formation mechanism, exploring OOS predictability would also have the potential practical implication that it could help to extract better inflation forecasts from the Michigan survey data by removing some predictable errors in real time.

To provide some perspective on the OOS predictability of forecast errors in the Michigan survey, the bottom rows of Table 3 report (pseudo) OOS test results. The prediction for the forecast error in period t to $t+4$ is constructed from estimates of a regression using data from the start of the sample up to quarter t. We use an initial window until 1976Q3 (20 observations) for the first prediction, with the exception of column (3), where the initial window extends until 1989Q4. We report the root mean squared error (RMSE) from this OOS prediction exercise for two specifications: one regression with only a constant, and one with the predictors included. In column (1), including the mean component of the learning-from-experience forecast in addition to the constant lowers the OOS RMSE to 0.017 from 0.019. To check the significance of this difference, we calculate the Diebold and Mariano (1995) statistic (with Newey-West adjustment). We obtain a p-value of 0.019, indicating evidence for OOS predictability. Adding additional predictors in column (2) has little effect. Out-of-sample predictability is also evident in the late sample in column (3) and the SPF in column (4).

5.4 Predictability of bond excess returns

As an alternative way of assessing whether the predictability of forecast errors is pervasive among macroeconomic forecasters and financial market participants and not just confined to the individuals in the MSC sample, we now examine excess returns on nominal long-term bonds. The tests with bond market returns have the additional benefit that we can use data that extends further back in time, because we only need inflation and return data, but not survey data for these tests.

For default-free bonds, an identity connects realized (log) returns in excess of the one-period risk-free rate from holding an n period bond from t to $t+1$ as follows (see, e.g., Piazzesi and Schneider (2011)):

$$rx_{t+1}^{(n)} = (n-1)(f_t^{(n-1,n)} - i_{t+1}^{(n-1)}), \tag{9}$$

where $rx_{t+1}^{(n)}$ denotes the excess return, $f_t^{(n-1,1)}$ is the time-t forward interest rate rate for the period starting at $t+1$ to $t+n$ and $i_{t+1}^{(n-1)}$ is the time $t+1$ yield yield of an $n-1$ period bond. Taking subjective expectations, $\hat{E}_t[.]$, of (9),

$$\hat{E}_t[rx_{t+1}^{(n)}] = (n-1)(f_t^{(n-1,n)} - \hat{E}_t[i_{t+1}^{(n-1)}]). \tag{10}$$

Taking objective expectations of (9),

$$E_t[rx_{t+1}^{(n)}] = (n-1)(f_t^{(n-1,n)} - E_t[i_{t+1}^{(n-1)}]). \tag{11}$$

If we assume, for simplicity, that investors price bonds with zero risk premia so that the expectations hypothesis holds under investors' subjective beliefs and hence $\hat{E}_t[rx_{t+1}^{(n)}] = 0$, then, substituting this into (10) and then into (11) yields objectively expected excess returns

$$E_t[rx_{t+1}^{(n)}] = (n-1)(\hat{E}_t[i_{t+1}^{(n-1)}] - E_t[i_{t+1}^{(n-1)}]), \tag{12}$$

i.e., objective expected excess returns are driven by deviations of investors' subjective expectations of future n period yields from objective expectations. These subjective expectations of future yields are in turn likely to be driven by subjective expectations of future inflation.[5]

[5]One way of making the link between yield expectations and inflation expectations explicit would be to combine a factor model of bond yields, most simply a single-factor model in which all bond yields are linear in the short-term interest rate, with an interest-rate policy rule under which the short-term interest rate is a

Suppose $\hat{E}_t[i_{t+1}^{(n-1)}] = \psi \hat{E}_t[\pi_{t+1}]$ and $E_t[i_{t+1}^{(n-1)}] = \psi E_t[\pi_{t+1}]$ for some constant ψ. Then,

$$E_t[rx_{t+1}^{(n)}] = \psi(n-1)(\hat{E}_t[\pi_{t+1}] - E_t[\pi_{t+1}]), \tag{13}$$

i.e., the predictability of bond excess returns is linked to the predictable component of inflation forecast errors $\hat{E}_t[\pi_{t+1}] - E_t[\pi_{t+1}]$. The more investors' subjective expectations of higher inflation (and hence higher future bond yields) exceed those under objective expectations, the higher the objectively expected excess returns.

For this reason, we now investigate whether we find predictability patterns in bond returns that are similar to those in survey forecast errors. Since only the mean component of the learning-from-experience forecast emerged as an economically and statistically significant predictor of survey forecast errors in Table 3, we focus on this single predictor here.

To measure long-term bond returns we use a return series of U.S. Treasury Bonds with maturities between 61 and 120 months from the Fama Bond database at the Center for Research in Security Prices ($CRSP$) and we construct excess returns by subtracting the 1-month T-Bill return (from Ibbotson Associates). We use quarterly returns as well as returns compounded to annual returns in our return-prediction regressions.

The results are presented in Table 4. The OLS coefficient estimate with quarterly returns in column (1) is 0.891 (s.e. 0.343), which yields an adj. R^2 of 2.1%. For return-prediction regressions this is an economically significant and plausible R^2. With annual returns, the magnitudes of coefficient and standard errors roughly quadruple and the adj. R^2 rises to 7.7%.

The predictor variable in these regressions is highly persistent, and its innovations (which are closely related to innovations in inflation) are contemporaneously correlated with long-term bond returns. Under these circumstances, it is well known that inference based on conventional OLS t-statistics leads to hypothesis tests that reject the null of no predictability too frequently in finite samples (Stambaugh (1999)). For this reason, we construct confidence

function of current inflation

Table 4: Predictability of bond excess returns

Quarterly and annual regressions of long-term U.S. Treasury bond returns in excess of 1-month Treasury Bill returns on the mean component of the learning-from-experience forecast (calculated with $\theta = 3.006$). Quarterly and annual bond returns are calculated by compounding monthly returns. The regression with annual returns uses non-overlapping windows. The sample period runs from 1952Q1 to 2010Q4. The table shows OLS estimates along with a 90% Bonferroni confidence interval following Campbell and Yogo (2006) for the coefficient on aggregate experienced inflation.

	Quarterly (1)	Annual (2)
Intercept	-0.019	-0.075
OLS s.e.	(0.008)	(0.033)
Mean component of learning-from-experience forecast	0.891	3.614
OLS s.e.	(0.343)	(1.401)
Campbell-Yogo 90% Bonferroni CI	[0.154, 1.384]	[0.830, 6.624]
Adj. R^2	0.021	0.077
AR order of predictor by BIC	2	1
90% CI for largest AR root	[1.001, 1.008],	[0.935, 1.002]
#Obs.	236	59

intervals using the methods of Campbell and Yogo (2006). Campbell and Yogo use local-to-unity asymptotics to achieve a better approximation of the finite-sample distribution in cases when the predictor variable is persistent. Their construction of the confidence interval uses the Bonferroni method to combine a confidence interval for the largest autoregressive root of the predictor variable with confidence intervals for the predictive coeffficient conditional on the largest autoregressive root.

As Table 4 shows, the Campbell-Yogo confidence interval for the regression coefficient of the predictor variable do not include zero, and they are approximately centered around the OLS point estimate.[6] This indicates that there is statistically reliable evidence in favor

[6] At the bottom of the table, we also report the estimated autoregressive lag length for the predictor variable, as determined by the Bayesian Information Criterion (BIC), as well as a confidence interval for its largest autoregressive root. These are among the inputs to Campbell and Yogo's construction of confidence intervals. The confidence intervals for the largest autoregressive root contain an explosive root. This is similar to the dividend-price ratio regressions in Campbell and Yogo (2006), and it underscores the potential importance of

of predictability. The mean component of the learning-from-experience forecast thus not only predicts the forecast errors in survey expectations from the MSC, but it also helps predict bond excess returns, which indicates that the learning-from-experience expectations-formation mechanism may be relevant for understanding expectations formation of bond market investors, too.

Unlike for the survey expectation forecast errors, there is, however, no evidence of out-of-sample predictability. OOS regressions with a constant (i.e., predicting simply with the past average return) yield a slightly lower OOS RMSE than regressions that include the mean component of learning-from-experience forecasts as a predictor. This suggests that bond market investors might be better than the respondents in the MSC and SPF in avoiding out-of-sample predictable forecast errors. As a caveat, though, it is difficult to interpret the out-of-sample results in return prediction regressions. Lack of OOS predictability is a common feature of return prediction regressions, and, as discussed in Campbell and Thompson (2008), OOS tests have low power to detect predictability.

6 Discussion and conclusion

Our empirical analysis shows that individuals' inflation expectations differ depending on the characteristics of the inflation process experienced during their life times. Differences in the experienced mean inflation rate and the persistence of inflation shocks generate (time-varying) differences in inflation expectations between cohorts. Younger individuals' set of experienced data is dominated by recent observations, while older individuals draw on a more extended historical data set in forming their expectations.

This learning-from-experience expectations-formation mechanism can explain, for example, why young individuals forecasted much higher inflation than older individuals following the high inflation years of the late 1970s and early 1980s. This is due to a combination of a high mean rate of inflation and high persistence in the short data set experienced by young

accounting for the persistence of the predictor variable in testing for predictability.

individuals at the time. Learning-from-experience also provides an alternative and complementary mechanism to the sticky information hypothesis in Mankiw and Reis (2002) and Carroll (2003) that contributes to the high level of disagreement about inflation expectations around that time noted in Mankiw, Reis, and Wolfers (2003).

For the most recent periods towards the end of our sample in 2010, our results suggest that individuals perception of the persistence of inflation shocks is close to zero, particularly for young individuals. This suggests that unexpected movements in the inflation rate are currently unlikely to move inflation expectations much. As argued in Roberts (1997), Orphanides and Williams (2005), and Milani (2007), these changes in individuals' perceptions of persistence are also likely to influence the persistence of inflation rates.

Even though the learning-from-experience framework is substantially different from more conventional representative-agent applications of learning in that it generates heterogeneity in inflation expectations, its implications for the average level of inflation expectations are similar to those resulting from representative-agent constant-gain learning algorithms that are popular in macroeconomics (see, e.g., Orphanides and Williams (2005); Milani (2007)). There are, however, two important differences.

First, the learning-from-experience theory provides an alternative motivation for a constant-gain learning at the aggregate level. With learning-from-experience, information in the distant past is discarded not only because individuals believe that structural shifts and parameter drift could occur, but also because individuals' memory is bounded: Memory of macroeconomic history is lost as new generations emerge whose subjective beliefs are shaped by relatively recent experience. This is an additional reason why learning dynamics may be perpetual, without convergence in the long-run.

Second, in the learning-from-experience framework, the heterogeneity between cohorts can be exploited to estimate the parameter controlling the gain in individuals' learning rule, and hence the speed of updating in response to inflation surprises, from cross-sectional differences alone, without using information about the level of average inflation expectations. This is useful, because identifying the gain from macro data seems to be difficult. In light of this, it is remarkable that our estimate of the speed of updating, averaged across cohorts, are quantitatively similar to those obtained in earlier work in macroeconomics that estimated the speed of updating to fit macroeconomic time-series or aggregate survey expectations.

Appendix

A Michigan Survey data

The inflation expectations data is derived from the responses to two questions, the first is categorical, while the second one elicits a percentage response. For example, for 1-year expectations the two questions are:

1. "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are right now?"
2. "By about what percent do you expect prices to go (up/down) on average during the next 12 months?"

As outlined in Curtin (1996), some adjustments to the raw data are necessary to address some known deficiencies. We follow Curtin's approach, which is also the approach used by the Michigan Survey in constructing its indices from the survey data:

For respondents who provided a categorical response of "up" ("down"), but not a percentage response, we drew a percentage response from the empirical distribution of percentage responses of those who gave the same categorical response of "up" ("down") in the same survey period. Prior to the February 1980 survey, respondents were not asked about percentage expectations if they responded (in the categorical first part of the question) that they expected prices to decline. We assign a value of -3% to these cases before February 1980. In most survey periods, they account for less than 2% of observations.

Starting in March 1982 the administrators of the Michigan survey implemented additional probing, which revealed that the categorical response that prices will remain the "same" was often misunderstood as meaning that the inflation rate stays the same. We use the adjustment factors developed in Curtin (1996) to adjust a portion of "same" responses prior to March 1982 to "up", and we assign a percentage response by drawing from the empirical distribution of those observations in the same survey period with a categorical response of "up".

B Imputation of percentage expectations from categorical responses

In the early years of the Michigan survey, only categorical responses about prices going "up", "down", or stay the "same" were elicited, but no percentage responses. We nevertheless attempt to use the information in those surveys in our analysis of percentage expectations by imputing percentage responses from the categorical information. We do so by estimating the relationship between categorical responses, the dispersion of categorical responses, and percentage responses in those periods in which we have both categorical and percentage response data. We conjecture that the average percentage response of individuals in an age group should be positively related to the proportion of "up" responses and negatively to the proportion of "down" responses.

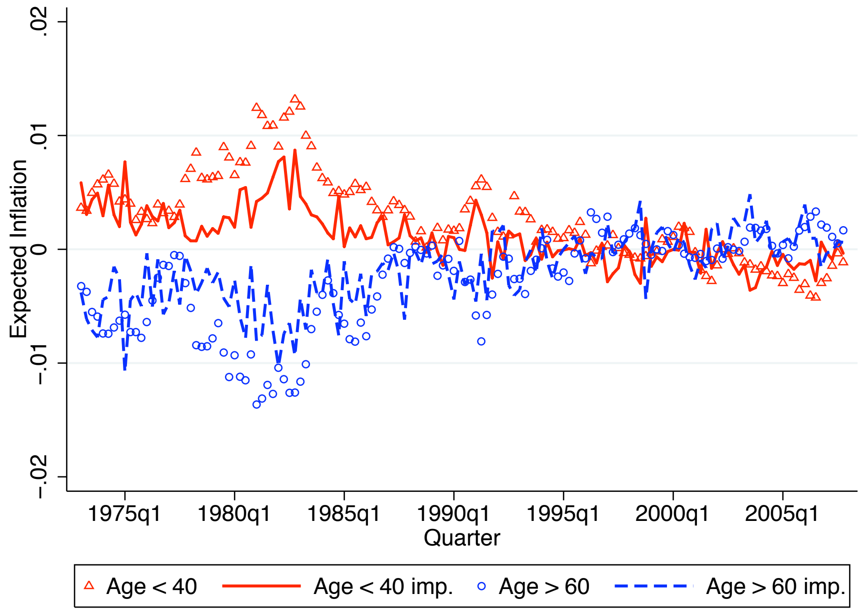

Figure A.1: Actual and imputed 1-year inflation expectations in excess of the full-sample cross-sectional mean.

We first calculate the proportion of "up" and "down" responses, $p_{t,s}^{up}$ and $p_{t,s}^{down}$, within each cohort s at time t (in this case t denotes a calendar month). We then run a pooled regression of measured percentage inflation expectations, $\hat{\pi}_{t+1|t,s}^{e}$, on $p_{t,s}^{up}$ and $p_{t,s}^{down}$, including a full set of time dummies, and obtain, for one-year expectations, the fitted values

$$\tilde{\pi}_{t+h|t,s}^{h} = \text{...time dummies...} + \underset{(0.001)}{0.052} p_{t,s}^{up} - \underset{(0.004)}{0.069} p_{t,s}^{down} \quad (R^2 = 35.3\%)$$

with standard errors in parentheses that are two-way clustered by quarter and cohort.

Because we employ time dummies in our main analysis, our main concern here is whether the imputed expectations track well cross-sectional differences of expectations across age groups, rather than the overall mean over time, and so we also estimate the relationship between percentage expectations and categorical responses with time dummies included in the regression.

Figure A.1 illustrates how the imputed percentage expectations compare with the actual expectations in the time periods in which we have both categorical and percentage expectations data. To focus on cross-sectional differences between age groups, the figure shows the average fitted and actual values (in terms of four-quarter moving averages) for individuals below 40 and above 60 years of age after subtracting the overall cross-sectional mean expectation in each time period.

Table A.1: Controlling for age-specific inflation rates

The estimation is similar as in Table 1, but with the experimental CPI for the elderly interacted with age included as control variable. The sample runs from 1984Q1 to 2009Q4, the period for which lagged 12-month inflation rates from the experimental CPI for the elderly is available. Standard errors in parentheses are two-way clustered by time and cohort.

	(1)	(2)
Gain parameter θ	2.702	3.901
	(0.286)	(0.633)
Sensitivity β	0.455	0.511
	(0.092)	(0.105)
Age		0.000
		(0.000)
Age$\times(\pi_{t-1}^{Elderly} - \pi_{t-1})$		-0.001
		(0.002)
Time dummies	Yes	Yes
Adj. R^2	0.244	0.245
#Obs.	5350	5350

C Controlling for age-specific inflation rates

We re-run the regressions from Table 1 with controls for age-specific inflation-rates. We measure the inflation rates of the elderly from the experimental CPI for the elderly series (CPI-E) provided by the Bureau of Labor Statistics. We calculate annualized quarterly log inflation rates from the CPI-E, similar to our calculation of overall CPI inflation rates. We then include in our regressions the differential between the CPI-E and CPI inflation rates, $\pi_{t-1}^{Elderly} - \pi_{t-1}$, interacted with age.

Table A.1 presents the results. The inflation series based on the CPI-E is only available from the end of 1983 onwards, and so the sample in this table is restricted to 1984Q1 to 2009Q4. As a basis for comparison, we therefore first re-run the regression without the additional age-dependent inflation control on this shorter sample. The results in column (1) show that the estimate of the gain parameter is similar to the earlier estimate in Table 1, but the sensitivity parameter β is estimated to be lower than before. Its magnitude is still statistically, as well as economically significant, though. In column (2) we add the interaction term between age-related inflation differentials and age, as well as age itself (the $\pi_{t-1}^{Elderly} - \pi_{t-1}$ variable itself without the interaction is absorbed by the time dummies). We obtain a small and insignificantly negative coefficient on the interaction term, which is not consistent with the idea that inflation expectations of the elderly are positively related to the inflation rates on the consumption basket of the elderly. Including age and the interaction term does, however, have some effect on the estimates for θ.

D Implied weighting of past data with learning from experience

The learning-from-experience algorithm in our analysis implicitly weights past observations in almost exactly similar fashion as the (ad-hoc) weighting function in Malmendier and Nagel (2011). Moreover, the parameter θ that controls the strength of updating in the framework here maps into the parameter that controls the weighting function in Malmendier and Nagel (2011). This makes the results easily comparable. For simplicity, we illustrate the connection between the two weighting schemes in the case of the simple mean model, where an agent tries to estimate the mean. But an analogous result applies in the AR(1) case or other regression-based forecasts.

Consider an individual of age $t-s$ making an inflation forecast at time t. The weighting function in Malmendier and Nagel (2011) implies that this individual forms a weighted average of past inflation, where the inflation rate observed at time $t-k$ (with $k \leq t-s$) gets the following weight:

$$\omega_{t,s}(k) = \frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.1}$$

This implies that the most recent observation, i.e. time-t inflation, π_t, receives the weight

$$\omega_{t,s}(0) = \frac{1}{\sum_{j=0}^{s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.2}$$

For comparison, in the learning-from-experience algorithm, the forecast $\tau^1_{t+1|t,s}$ is a weighted average of the prior-period forecast and π_t,

$$\tau^1_{t+1|t,s} = (1-\gamma_{t-s})\,\tau^1_{t|t-1,s} + \gamma_{t,s}\pi_t. \tag{A.3}$$

which implies that the most recent observation carries the weight $\tilde{\omega}_{t,s}(0) = \gamma_{t,s}$. Iterating, one finds that earlier observations receive the weight

$$\tilde{\omega}_{t,s}(k) = \begin{cases} \gamma_{t,s} & \text{for } k=0 \\ \gamma_{t-k,s}\prod_{j=0}^{k-1}(1-\gamma_{t-j,s}) & \text{for } k>0 \end{cases}. \tag{A.4}$$

We now show that both weighting schemes are equivalent if the gain sequence is chosen to be age-dependent in the following way:

$$\gamma_{t,s} = \frac{1}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}} \tag{A.5}$$

We present a proof by induction. First, the choice of $\gamma_{t,s}$ in (A.5) implies that $\tilde{\omega}_{t,s}(0) =$

$\omega_{t,s}(0)$. It remains to be shown that if $\tilde{\omega}_{t,s}(k) = \omega_{t,s}(k)$, then $\tilde{\omega}_{t,s}(k+1) = \omega_{t,s}(k+1)$ (with $k+1 \leq t-s$). Thus, assume that

$$\tilde{\omega}_{t,s}(k) = \frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.6}$$

Then, from Eq. (A.4),

$$\begin{aligned}
\tilde{\omega}_{t,s}(k+1) &= \gamma_{t-s-k-1}\frac{(1-\gamma_{t-s-k})}{\gamma_{t-s-k}}\tilde{\omega}_{t,s}(k) \\
&= \frac{\left[\sum_{j=0}^{t-s-k}\left(\frac{t-s-j}{t-s-k}\right)^{\lambda}\right]-1}{\sum_{j=0}^{t-s-k-1}\left(\frac{t-s-j}{t-s-k-1}\right)^{\lambda}}\frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\left[\sum_{j=0}^{t-s-k}\left(\frac{t-s-j}{t-s-k}\right)^{\lambda}\right]-1}{\sum_{j=0}^{t-s-k-1}\left(\frac{t-s-j}{t-s-k}\right)^{\lambda}}\frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\sum_{j=0}^{t-s-k-1}\left(\frac{t-s-j}{t-s-k}\right)^{\lambda}}{\sum_{j=0}^{t-s-k-1}\left(\frac{t-s-j}{t-s-k}\right)^{\lambda}}\frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \omega_{t,s}(k+1),
\end{aligned}$$

where for the third-to-last equality we multiplied numerator and denominator by $\left(\frac{t-s-k-1}{t-s-k}\right)^{\lambda}$. This concludes the proof.

Finally, we show that the gain sequence (A.5) can be approximated by

$$\gamma_{t,s} \approx \frac{\lambda+1}{t-s},$$

i.e., by the gain specification in (4) with $\theta = \lambda + 1$. To see this write the gain in (A.5) as

$$\gamma_{t,s} = \frac{(t-s)^{\lambda}}{\sum_{j=0}^{t-s}(t-s-j)^{\lambda}}.$$

Focusing on the denominator of this expression, note that if one were to make the increments j infinitesimally small (instead of being discrete steps of 1), the denominator would become $\int_0^{t-s} x^{\lambda}dx = \frac{1}{\lambda+1}(t-s)^{\lambda+1}$. Therefore, in this limiting case of infinitesimal increments, we get

$$\gamma_{t,s} = \frac{\lambda + 1}{t - s}.$$

In our case with quarterly increments, this approximation is, for all practical purposes, virtually identical with the true gain sequence in (A.5).

References

Bernanke, B. S. (2007). Inflation Expectations and Inflation Forecasting. Speech at NBER Monetary Economics Workshop, available at http://www.federalreserve.gov/newsevents/speech/Bernanke20070710a.htm.

Blanchflower, D. G. and R. Kelly (2008). Macroeconomic Literacy, Numeracy and the Implications for Monetary Policy. working paper, Bank of England.

Branch, W. A. (2004). The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations. *Economic Journal 114*, 592–621.

Branch, W. A. (2007). Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations. *Journal of Economic Dynamics and Control 31*, 245–276.

Bray, M. (1982). Learning, Estimation, and the Stability of Rational Expectations Equilibria. *Journal of Economic Theory 26*, 318–339.

Campbell, J. Y. and S. B. Thompson (2008). Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies 21*, 1509–1531.

Campbell, J. Y. and M. Yogo (2006). Efficient Tests of Stock Return Predictability. *Journal of Financial Economics 81*, 27–60.

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics 118*(1), 269–298.

Chevillon, G., M. Massmann, and S. Mavroeidis (2010). Inference in Models with Adaptive Learning. *Journal of Monetary Economics 57*, 341–351.

Chiang, Y.-M., D. Hirshleifer, Y. Qian, and A. E. Sherman (2011). Do Investors Learn from Experience? Evidence from Frequent IPO Investors. *Review of Financial Studies 24*, 1560–1589.

Cukierman, A. and P. Wachtel (1979). Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence. *American Economic Review 69*(4), 595–609.

Curtin, R. T. (1982). Indicators of Consumer Behavior: The University of Michigan Surveys of Consumers. *Public Opinion Quarterly 46*(3), 340–352.

Diebold, F. X. and R. S. Mariano (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statistics 13*, 253–263.

Ehrmann, M. and P. Tzamourani (2009). Memories of High Inflation. working paper, European Central Bank.

Evans, G. W. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics.* Princeton, NJ: Princeton University Press.

Greenwood, R. and S. Nagel (2009). Inexperienced Investors and Bubbles. *Journal of Financial Economics 93*, 239–258.

Honkapohja, S. and K. Mitra (2003). Learning with Bounded Memory in Stochastic Models. *Journal of Economic Dynamics and Control 27*, 1437–1457.

Kaustia, M. and S. Knüpfer (2008). Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions. *Journal of Finance 63*, 2679–2702.

Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money.* London: Macmillan.

Malmendier, U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking? *Quarterly Journal of Economics 126*, 373–416.

Mankiw, G. N. and R. Reis (2002). Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics 117*(4), 1295–1328.

Mankiw, N. G. and R. Reis (2006). Pervasive Stickiness. *American Economic Review Papers and Proceedings 96*, 164–169.

Mankiw, N. G., R. Reis, and J. Wolfers (2003). Disagreement About Inflation Expectations. In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2003.*

Marcet, A. and J. Sargent, Thomas (1989). Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory 48*, 337–368.

Milani, F. (2007). Expectations, Learning, and Macroeconomic Persistence. *Journal of Monetary Economics 54*, 2065–2082.

Orphanides, A. and J. C. Williams (2004). Imperfect Knowledge, Inflation Expectations and Monetary Policy. In B. S. Bernanke and M. Woodford (Eds.), *The Inflation Targeting Debate*. Chicago, IL: University of Chicago Press.

Orphanides, A. and J. C. Williams (2005). The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations. *Journal of Economic Dynamics and Control 29*, 1927–1950.

Pfajfar, D. and E. Santoro (2010). Heterogeneity, Learning and Information Stickiness in Inflation Expectations. *Journal of Economic Behavior and Organization 75*, 426–444.

Piazzesi, M. and M. Schneider (2010). Inflation and the Price of Real Assets. working paper, Stanford University.

Piazzesi, M. and M. Schneider (2011). Trend and Cycle in Bond Premia. working paper, Stanford University.

Roberts, J. M. (1997). Is Inflation Sticky? *Journal of Monetary Economics 39*, 173–196.

Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics.* Oxford, UK: Clarendon Press.

Shiller, R. J. (1997). Why Do People Dislike Inflation? In C. D. Romer and D. H. Romer (Eds.), *Reducing Inflation: Motivation and Strategy*, pp. 13–65. Chicago: University of Chicago Press.

Shiller, R. J. (2005). *Irrational Exuberance* (second edition). Princeton, NJ: Princeton University Press.

Sims, C. A. (1998). Stickiness. *Carnegie-Rochester Conference Series on Public Policy 49*, 317–56.

Stambaugh, R. F. (1999). Predictive Regressions. *Journal of Financial Economics 54*, 375–421.

Vissing-Jorgensen, A. (2003). Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions. In *NBER Macroeconomics Annual*.