Learning from Inflation Experiences[*]

Ulrike Malmendier[†]
UC Berkeley, NBER and CEPR

Stefan Nagel[‡]
Stanford University, NBER and CEPR

April 4, 2013

Abstract

How do individuals form expectations about future inflation? We propose that personal experiences play an important role. Individuals adapt their forecasts to new data but overweight inflation realized during their life-times. Young individuals update their expectations more strongly in the direction of recent surprises than older individuals since recent experiences make up a larger part of their lives so far. We find support for these predictions using 57 years of microdata on inflation expectations from the Reuters/Michigan Survey of Consumers. Differences in life-time experiences strongly predict differences in subjective inflation expectations. Learning from experience explains the substantial disagreement between young and old individuals in periods of high surprise inflation, such as the 1970s. It also explains household borrowing and lending behavior, including the choice of fixed versus variable-rate investments and mortgages. The loss of distant memory implied by learning from experience provides a natural microfoundation for models of perpetual learning, such as constant-gain learning models.

1 Introduction

How do individuals form expectations about future inflation? The answer to this question is of central importance both for monetary policy and for the financial decision-making of individuals. Policy makers would like to better understand the formation of inflation expectations in order to improve their inflation forecasts and the resulting policy choices (Bernanke (2007)). The financial decisions of individuals, e.g., in the housing market, are sensitive to their perception of real interest rates, which in turn are influenced by their perception of inflation. Hence, inflation expectations influence current real expenditure and macroeconomic outcomes (e.g., Woodford (2003)).

Despite a large volume of research, there is still little convergence on the best model to predict inflation expectations (see Mankiw, Reis, and Wolfers (2003); Blanchflower and Kelly (2008)). The "stickiness" of inflation rates (Sims (1998), Mankiw and Reis (2006)) and the empirical heterogeneity in the formation of expectations remain hard to reconcile with existing models. Consider the simple time-series plot of inflation expectations from the Reuters/Michigan Survey of Consumers (MSC) in Figure 1.[1] The figure plots the expectations of young individuals (age below 40), middle-aged individuals (age 40 to 60), and older individuals (age above 60), expressed as deviations from the cross-sectional mean expectation in each period. The figure shows that the dispersion in beliefs between young and old individuals can be large, reaching almost 3 percentage points during the high-inflation years of the 1970s and early 1980s. The graph also shows repeated reversals in relative beliefs, with the old having lower inflation expectations than the young in the 1970s and 1980s, but higher inflation expectations in the late 1960s, mid-1990s, and late 2000s. These patterns are unexplained in existing models. Similar concerns apply to the formation of beliefs about other macro-economic variables and their influence on aggregate dynamics, as discussed, for example, in Fuster, Laibson, and Mendel (2010) and Fuster, Hebert, and Laibson (2011).

In this paper, we argue that individuals' personal experiences play an important role in

---

[1]We will discuss the figure and underlying data in more detail in section 3.1.
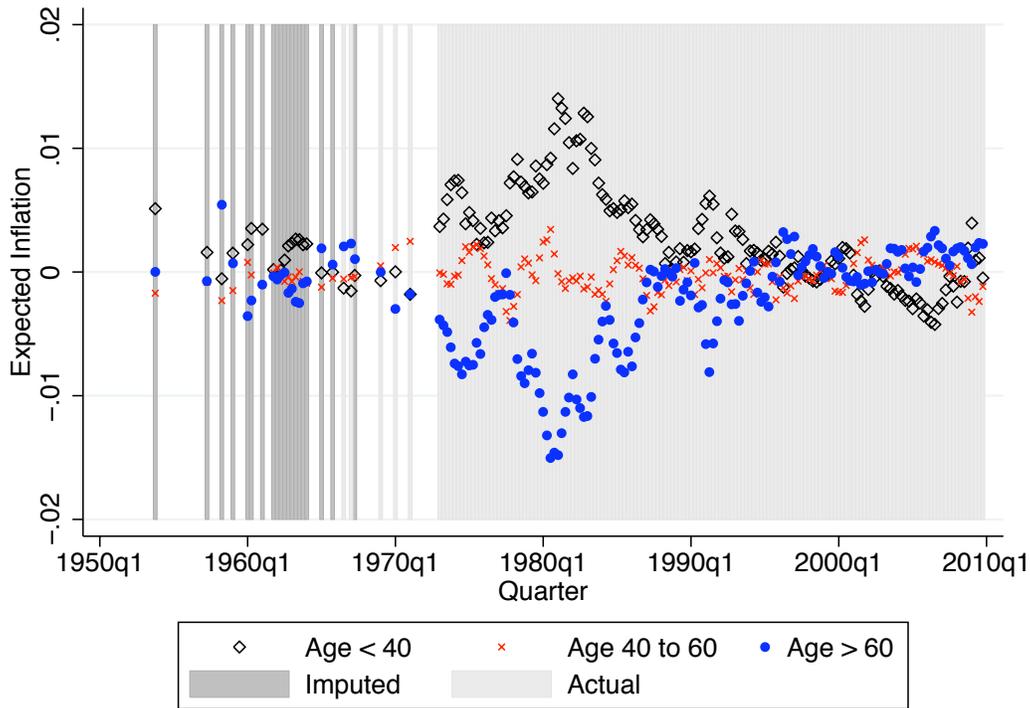
Figure 1: Four-quarter moving averages of mean one-year inflation expectations of young individuals (below 40), mid-aged individuals (between 40 and 60), and old individuals (above 60), shown as deviations from the cross-sectional mean expectation. Percentage forecasts are available from the survey in periods shaded in light grey; they are imputed from categorical responses in the periods shaded in dark grey; and they are unavailable in unshaded periods.

shaping expectations that is absent from existing models. When forming macroeconomic expectations, individuals put a higher weight on realizations of macroeconomic data experienced during their life-times compared with other available historical data. As a result, 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.

Such *learning from experience* carries two central implications. First, expectations are history-dependent. Cohorts that have lived through periods of high inflation for a substantial amount of time have higher inflation expectations than individuals who have mostly experienced low inflation. Second, beliefs are heterogeneous. Young individuals place more weight

2

on recent data than older individuals since recent experiences make up a larger part of their life-times so far. As a result, different generations tend to disagree about the future.

Both effects have been noticed in practice. During the high-inflation period of the late 1970s, the Chairman of the Federal Reserve Paul Volcker remarked: "An entire generation of young adults has grown up since the mid-1960's knowing only inflation, indeed an inflation that has seemed to accelerate inexorably. In the circumstances, it is hardly surprising that many citizens have begun to wonder whether it is realistic to anticipate a return to general price stability, and have begun to change their behavior accordingly."[2] Both effects are also visible in Figure 1: Following years of high inflation, young people expect much higher inflation going forward than older people.

We estimate a learning-from-experience model using 57 years of microdata on inflation expectations from the MSC. In our empirical framework, we assume that individuals employ regression-based forecasting rules as 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 inflation. Learning from experience is implemented by allowing the gain, i.e., the strength of updating in response to surprise inflation, to depend on age. 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 the gain parameter empirically by fitting the learning rule to the cohort-level inflation expectations from the MSC.

The availability of microdata is crucial for our estimation. Our identification strategy exploits time-variation in the cross-sectional differences of inflation experiences between cohorts and relates it to time-variation in the cross-sectional differences of inflation expectations. This identification from cross-sectional heterogeneity allows us to include time dummies in

---

[2]See Volcker 1979, p. 888; quoted in Orphanides and Williams (2005b).

our specifications to separate the experience effect from time trends or any other time-specific determinants. For example, individuals might also draw on the full inflation history available at a given time, or on published forecasts of professional forecasters. With the inclusion of time dummies, the experience coefficient isolates the *incremental* explanatory power of individual life-time experiences over and above such common time-specific factors. In other words, our empirical approach allows us to rule out that omitted macroeconomic variables or any other unobserved effects common to all individuals bias the estimation results. This is a key distinction from other models of belief formation, such as adaptive learning models, where parameters are fit to aggregate time-series of (mean or median) expectations.

Our estimation results show that past experiences have an economically important effect on inflation expectations. Individuals of different ages differ significantly in their inflation expectations, and these differences are well explained by differences in their inflation experiences. The heterogeneity is particularly pronounced following periods of high surprise inflation. Consider again the strong divergence in expectations between younger and older cohorts during the late 1970s and early 1980s displayed in Figure 1. The higher expectations of younger individuals are consistent with their experience being dominated by the high-inflation years of the 1970s, while older individuals also experienced the low-inflation years of the 1950s and 1960s. The discrepancy faded away only slowly by the 1990s, after many years of moderate inflation. Our model explains this difference as the result of younger individuals perceiving inflation to be (i) higher on average and (ii) more persistent when inflation rates were high until the early 1980s, but less persistent when inflation rates dropped subsequently. Our estimates of the gain parameter further imply that when individuals weight their accumulated life-time experiences, recent data receives higher weight than experiences earlier in life, though experiences from 20 to 30 years ago still have some measurable long-run effects.

Going one step further, we also link the effect of experiences on beliefs to actual household decision-making. The experience-induced disagreement about future inflation leads to disagreement about real interest rates on assets and liabilities with nominally fixed long-term

4

interest rates. Consistent with this idea, we show that the learning-from-experience model predicts the borrowing and investment decisions of households in the Survey of Consumer Finances. Households that forecast higher inflation according to the learning-from-experience model are more inclined than other individuals to borrow using fixed-rate mortgages, but not using variable-rate mortgages. They are also less likely (but with only marginal statistical significance) to invest in long-term bonds. Taken together, our results show that households with higher experience-induced inflation expectations tilt their exposure towards liabilities with nominally fixed rates rather than assets with nominally fixed rates. The effects are economically large. For instance, a one percentage point difference in the learning-from-experience forecast of one-year inflation affects the mortgage balance by as much as a one-standard-deviation change in log income.

Finally, we link learning from experience to aggregate expectations, as existing macro models with adaptive learning are commonly fit to the time series of aggregate (mean) inflation expectations. We show that the cross-sectional average of the fitted learning-from-experience forecasts at each point in time 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.

Learning from experience thus provides a natural micro-foundation for constant-gain learning algorithms that are popular in macroeconomics. In fact, the *average* learning-from-experience forecast can be approximated quite closely with constant-gain learning: The constant-gain parameter that best matches the weighting of past data implied by our estimated learning-from-experience model, $\gamma = 0.0180$, is quantitatively very similar to those that Orphanides and Williams (2005a) and Milani (2007) have estimated by fitting the parameters to macroeconomic data and aggregate survey expectations, 0.02 and 0.0183, respectively. This similarity is remarkable because we did not calibrate the learning-from-experience rule to match the average level of inflation expectations or any macroeconomic data.

At the same time, learning from experience and constant-gain learning differ fundamentally in their motivation. The down-weighting of past data in constant-gain learning models is typically motivated as the response to structural changes in macroeconomic time series. Learning from experience, instead, is based on the notion that memory of past data is lost as older generations die and new ones are born. We show that the structural-change based explanation is hard to reconcile with the empirical evidence: households discount past inflation data much less than the degree of structural change in the time-series of inflation would call for. Moreover, macroeconomic time series, such as GDP and inflation, differ in the degree of structural change, but these differences does not seem to be reflected in the gains implied by survey expectations.

Learning from experience also differs from existing models of adaptive learning in terms of identification. The gain parameter in our estimations is identified purely from cross-sectional variation. The econometric advantage, compared to using aggregate time-series data to identify the gain parameter, is twofold. First, using a new dimension of data to pin down the parameters of individuals' learning rule helps alleviate the concern about non-standard (or, boundedly rational) learning models involving too many degrees of freedom and, hence, not being falsifiable, as expressed in Sargent (1993) and Marcet and Nicolini (2003). Second, empirically, the identification from cross-sectional data offers some advantage in capturing the dynamics of survey expectations. We show that estimating the gain parameter purely from cross-sectional variation provides a better out-of-sample fit to mean inflation expectations than estimating a constant-gain rule from time-series data, even though the estimation within the learning-from-experience framework does not use any data on the level of mean expectations.

Our paper connects to several strands of literature. A large macroeconomic literature analyzes the formation of expectations. While the crucial role of expectations for macroeconomic outcomes and asset prices is well understood at least since Keynes (1936), the empirical knowledge how economic agents form their beliefs about the future is more limited. The liter-

ature 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. Fuster, Laibson, and Mendel (2010) and Fuster, Hebert, and Laibson (2011) propose a model of "natural expectations" and demonstrate its ability to match hump-shaped dynamics in different economic time series.

A small, but growing literature looks at heterogeneity in expectations with microdata. Building on early work by Cukierman and Wachtel (1979), Mankiw, Reis, and Wolfers (2003) examine the time-variation in the dispersion of inflation expectations and relate it to models of "sticky information." Carroll (2003) further investigates the sticky-information model, but with aggregate data on inflation expectations. Bryan and Venkatu (2001) provide evidence of gender-specific heterogeneity in inflation expections as well as other demographics-based differences. Branch (2004), Branch (2007), and Pfajfar and Santoro (2010) estimate from survey data how individuals choose among competing forecasting models. Relatedly, Piazzesi and Schneider (2012) show that the disagreement about future inflation between younger and older households in the late 1970s, and the implied heterogeneity in perceived real interest rates, helps understand quantities of household borrowing and lending, shifts in household portfolios, and the prices of real assets at the time. Piazzesi and Schneider (2011) use data on subjective interest-rate expectations in 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 the formation of expectations and helps explain the observed heterogeneity in expectations as well as the notion of gradually fading memory over time.

Conceptually, our approach is related to bounded-memory learning in Honkapohja and Mitra (2003) in that memory of past data is lost or, in our case, underweighted. However, while bounded-memory learning agents are homogeneous, the memories of agents in the learning-from-experience model differ depending on their age. Our approach also relates to

the boundedly-rational learning model in Marcet and Nicolini (2003), where agents learn differently during stable and during instable periods. Our approach pins the gain in learning to agents' age rather than a threshold rule that is chosen to fit the data.

Our paper is also related to the empirical findings in Malmendier and Nagel (2011), who show that past stock-market and bond-market experiences predict future risk taking of investors. Interestingly, the weighting of past inflation experiences estimated in this paper matches very closely the weighting scheme estimated from a completely different data source, the Survey of Consumer Finances, in Malmendier and Nagel (2011). However, the data used in Malmendier and Nagel (2011) did not allow 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 identifies the beliefs channel.

Evidence consistent with learning-from-experience effects is also presented in Vissing-Jorgensen (2003), who shows that young retail investors with little investment experience had the highest stock return expectations during the stock-market boom in the late 1990s. Similar evidence on experience effects exist for mutual fund managers who experienced the stock market boom of the 1990s (Greenwood and Nagel (2009)) and for CEOs who grew up in the Great Depression (Malmendier and Tate (2005) and Malmendier, Tate, and Yan (2011)). In addition, 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 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 learning-from-experience framework and the estimation approach. Section 3 discusses the data set on inflation expectations and presents the core results on learning-from-experience effects in the

8

formation of beliefs about future inflation. Section 4 shows that learning-from-experience effects also help understand actual household decisions about bond investing and mortgage borrowing. In Section 5, we look at the implications of our results at the aggregate level. Section 6 concludes.

2   Learning from experience

Consider two individuals, one born at time s, and the other 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 they place different weights on recent and distant historical data, reflecting the different inflation histories they have experienced in their 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 algorithms proposed in the adaptive learning literature (see Marcet and Sargent (1989); also Sargent (1993) and Evans and Honkapohja (2001)). In this literature, these algorithms are meant to represent the "rules of thumb" that agents use in practice because they face cognitive and computational constraints. Our model departs from standard adaptive learning in that we allow individuals to put more weight on data experienced during their lifetimes than on other historical data. This gives rise to between-cohort differences in expectations. While much of the existing literature focuses on identifying the conditions under which these simple learning rules lead to convergence to rational-expectations equilibria, we focus on empirically estimating individuals' actual forecasting rules.

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 (2005b):

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

9

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

$$b_{t,s} = b_{t-1,s} + \gamma_{t,s} R_{t,s}^{-1} x_{t-1}(\pi_t - b_{t-1,s}' x_{t-1}), \tag{2}$$

$$R_{t,s} = R_{t-1,s} + \gamma_{t,s}(x_{t-1} x_{t-1}' - R_{t-1,s}), \tag{3}$$

where $x_t \equiv (1, \pi_t)'$ and the recursion is started at some point in the distant past.[3] The sequence of gains $\gamma_{t,s}$ determines the degree of updating cohort s applies when faced with an inflation surprise at time t. For example, with $\gamma_{t,s} = 1/t$, the algorithm represents a recursive formulation of an ordinary least squares estimation that uses all available data until time t with equal weights (see Evans and Honkapohja (2001)). With $\gamma_{t,s}$ set to a constant, it represents a constant-gain learning algorithm, which weights past data with exponentially decaying weights. The key modification of the standard adaptive learning framework is that we let the gain γ depend on the age $t - s$ of the members of cohort s. As a result, individuals of different age can be heterogeneous in their forecasts and adjust their forecasts differently in response to surprise inflation. Given the perceived law of motion in equation (1), these cross-sectional differences can arise from two sources: from differences in individuals' perception of mean inflation $\mu = \alpha(1 - \phi)^{-1}$ and from differences in the perception of the 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 inflation experiences. We let the recursion start with $\gamma_{t,s} = 1$ for $t - s < \theta$, which implies that data before birth is ignored.[4] This specification is the same as in Marcet

---

[3]We will see below (see Figure 2 and the empirical estimates determining $\gamma_{t,s}$) that past data gets downweighted sufficiently fast so that initial conditions do not exert any relevant influence.

[4]As explained below, our econometric specification does allow for individuals to use all available historical

and Sargent (1989) with one modification: the gain here is decreasing in age, not in time. The parameterization $\frac{\theta}{t-s}$ allows for experiences early in an individual's life-time history to have a different influence compared with more recent experiences. For example, the memory of past episodes of high inflation might fade away over time as an individual ages. Alternatively, high inflation experienced at young age, and perhaps conveyed through the worries of parents, might leave a particularly strong and lasting impression.

Figure 2 illustrates the role of θ. The top graph of Figure 2 presents the sequences of gains γ as a function of age (in quarters) for different values of θ. Regardless of the value of θ, gains decrease with age. This is sensible in the context of learning from experience: Young individuals, who have experienced only a small set of historical data, have a higher gain than older individuals, who have experienced a longer history and for whom a single inflation surprise should have a weaker marginal influence on expectations. The top graph also illustrates that the higher θ is, the higher is the gain 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 2, which shows the implied weights on past inflation observations as a function of the time lag relative to current time t in the case of a 50-year old individual. For $\theta = 1$, all historical observations since birth are weighted equally. For $\theta < 1$, observations early in life receive more weight, and for $\theta > 1$ observations early in life receive less weight than more recent observations. With $\theta = 3$, for example, very little weight is put on observations in the first 50 quarters since birth towards the right of the bottom graph. Hence, our gain parametrization is flexible in accommodating different weighing schemes. The weights can be monotonically increasing, decreasing, or flat. An additional advantage of the specification 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 details in Appendix A.) This allows us to compare the weights estimated from inflation expectations with those estimated from portfolio allocations in Malmendier and Nagel (2011).

---

data when forming expectations and isolates the incremental effect of data realized during individuals' life-times.
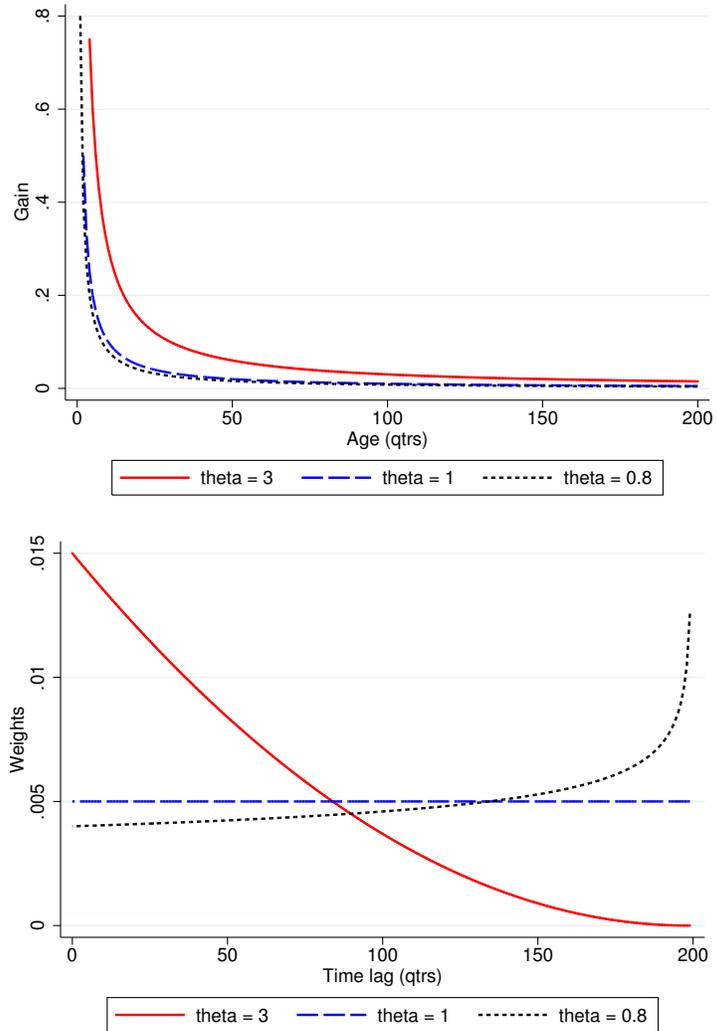
Figure 2: Examples of gain sequences (top) and associated implied weighting of past data (bottom) for an individual who is 50 years (200 quarters) old. The top panel shows the sequence of gains as a function of age. The bottom panel shows the weighting of past data implied by the gain sequence in the top panel, with the weights for most recent data to the left and weights for early-life experiences to the right.

In addition to past inflation experiences, we allow other information sources to affect inflation expectations. Let $\pi_{t+1|t,s}$ be the forecast of period $t+1$ inflation made by cohort s at time t. The learning-from-experience component of individuals' one-step-ahead forecast of inflation is obtained from (2) as $\tau_{t+1|t,s} = b'_{t,s} x_t$.[5] We capture the influence of information sources other than experienced inflation by assuming

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

That is, the subjective expectation is a weighted average of the learning-from-experience component $\tau_{t+1|t,s}$ and an unobserved common component f_t. The unobserved component f_t could be based on any kind of 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 could capture a component that is based on all available historical data. In either case, the coefficient β captures the incremental contribution of life-time experiences $\tau_{t+1|t,s}$ to $\pi_{t+1|t,s}$ over and above these 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):

$$\tilde{\pi}_{t+1|t,s} = \beta \tau_{t+1|t,s} + \delta' D_t + \varepsilon_{t,s}, \qquad (6)$$

where $\tilde{\pi}_{t+1|t,s}$ denotes measured inflation expectations from survey data. In this estimating equation, we absorb the unobserved f_t with a vector of time dummies D_t. We also add the disturbance $\varepsilon_{t,s}$, which we assume to be uncorrelated with $\tau_{t+1|t,s}$, but which is allowed to be correlated over time within cohorts and between cohorts within the same time period.

---

[5]In our estimation, we apply the AR(1) process in equation (1) to the quarterly inflation data, while the survey data provides individuals' forecasts of inflation over a one-year (i.e., four-quarter) horizon. Correspondingly, we employ multi-period learning-from-experience forecasts that we obtain by iterating, at time t, on the perceived law of motion in (1) using the time-t estimates $b_{t,s}$. To economize on notation, we do not explicitly highlight the multi-period nature of the forecasts here.

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 $\tau_{t+1|t,s}$ is a non-linear function of θ.)

The presence of time dummies in equation (6) implies that we identify β and θ, and hence the learning-from-experience effect, 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 rules out confounds that could have affected prior work that estimated adaptive learning rules from aggregate data, e.g., from the time-series of mean or median inflation expectations. With aggregate data, it is hard to establish whether the time-series relationship between inflation expectations and lagged inflation truly reflects adaptive learning or some other formation mechanism (e.g., rational expectations) that happens to generate expectations that are highly correlated. In contrast, the learning-from-experience model 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 strength of updating in response to new data from this cross-sectional heterogeneity. This provides a new source of identification for the gain parameter in adaptive learning schemes.

3   Inflation experiences and inflation expectations

We estimate the learning-from-experience effects by fitting the estimating equation (6) and the underlying AR(1) model to data on inflation expectations from the Reuters/Michigan Survey of Consumers.
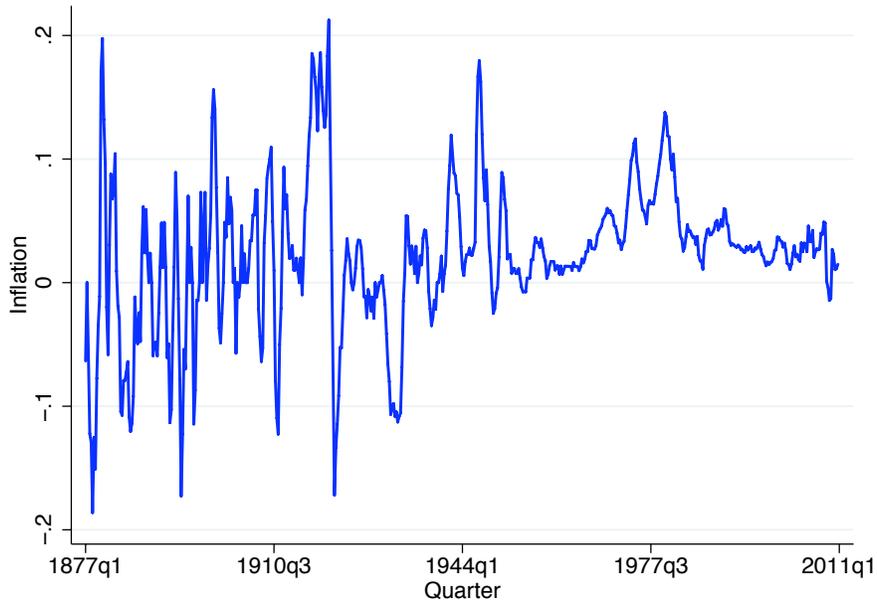
Figure 3: Annual CPI inflation rates

3.1 Data

We measure inflation experiences using long-term historical data on the consumer price index (CPI) from Shiller (2005). In order to fully capture inflation experiences during the lifetimes of all individuals in the survey data, including the oldest respondents in the earliest survey wave, we need inflation data stretching back 74 years before the start of the survey data in 1953. We obtain the time series of inflation data since 1872 (until the end of 2009) from Robert Shiller's website and calculate annualized quarterly log inflation rates. For illustration, Figure 3 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. This survey has been administered since 1953, initially three times per year, then quarterly from 1960 through 1977, and monthly since 1978 (see Curtin (1982)). We obtain the 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 Consumer Finances (SCF), also available at the $ICPSR$. The data from 1978 to 2009 is available from the University of Michigan Survey Research Center. Appendix B provides more detail on the data.

In most periods, the survey asks 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. Since our analysis aims to make quantitative predictions, we focus on the percentage expectations. However, for quarters in which the survey asks only the categorical questions about the expected direction, we are able to impute percentage responses from the distribution of the categorical responses. The imputation procedure is described in Appendix C. Figure 1 in the introduction highlights the periods in which we have percentage expectations data in light grey, and the quarters in which the survey asks only the categorical questions in dark grey.

Since the learning-from-experience hypothesis predicts that inflation expectations are 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 the sample to respondents aged 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.

Figure 1 provides some sense of the variation in the data. As mentioned in the introduction, the figure plots the average inflation expectations of young individuals (averaging, for the figure only, across all cohorts below 40 years of age), middle-aged individuals (ages 40 to 60), and older individuals (ages above 60), expressed as deviations from the cross-sectional mean expectation each period. To better illustrate lower-frequency variation, we plot the

16

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 hypothesis: The experience of young individuals around 1980 was dominated by high and persistent inflation in recent years, while the experience of older individuals also included the modest and less persistent inflation rates of earlier decades. For younger individuals, these recent observations exert a stronger influence on their expectations since their experience set contains a smaller number of data points.

3.2 Baseline results

We now fit the estimating equation (6) and the underlying AR(1) model using nonlinear least squares on the cohort-level aggregate data. We relate the inflation forecasts in the MSC to learning-from-experience forecasts. We assume that the data available to individuals who are surveyed (at various points) during quarter t are quarterly inflation rates until the end of quarter $t-1$. Since the survey elicits expectations about the inflation rate over the course of the next year, but the (annualized) inflation rates that serve as input to the learning-from-experience algorithm are measured at quarterly frequency, we require multi-period forecasts from the learning-from-experience model. We obtain these multi-period forecasts by iterating on the perceived AR(1) law of motion (1) at each cohort's quarter-t estimates of the AR(1) parameters α and ϕ (which are based on inflation data up to the end of quarter $t-1$). Hence, the one-year forecast that we relate to survey expectations is the average of the AR(1) forecasts of quarter $t+1$ to quarter $t+4$ annualized inflation rates. 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 presents the estimation results. In the full sample (column (i)), we estimate the gain parameter θ to equal 3.044 (s.e. 0.233). Comparing this estimate of θ with the

Table 1: Learning-from-experience model: Estimates from cohort data

Each cohort born at time s is assumed to recursively estimate an AR(1) model of inflation, with the decreasing gain $\gamma_{t,s} = \theta/(t-s)$ and 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 survey inflation expectations in quarter t (cohort means) on these learning-from-experience forecasts. Standard errors reported in parentheses are two-way clustered by time (quarter) and cohort. The sample period runs from 1953 to 2009 (with gaps).

	Baseline		Restricted		
	(i)	(ii)	(iii)	(iv)	(v)
Gain parameter θ	3.044	3.144	3.976	4.144	3.044
	(0.233)	(0.257)	(0.494)	(0.446)	-
Sensitivity β	0.672	0.675	0.664	0.693	0.672
	(0.076)	(0.079)	(0.084)	(0.111)	(0.093)
Time dummies	Yes	Yes	No	No	No
Imputed data included	Yes	No	No	No	No
Restrictions			$f_t = SPF_{t-1}$	$f_t = \bar{\tau}_{t+1|t}$	$f_t = \bar{\tau}_{t+1|t}$
					$\theta = 3.044$
Adj. R^2	0.637	0.635	-	-	-
RMSE	0.0148	0.0152	0.0189	0.0191	0.0194
#Obs.	8215	7650	7400	7650	7650

illustration in Figure 2 one can see that the estimate implies weights that are declining a bit faster than linearly. The estimation results also show that there is a strong relationship between the learning-from-experience forecast and measured inflation expectations, captured by the sensitivity parameter β, which we estimate to be 0.672 (s.e. 0.076). The magnitude of β implies that, when two individuals differ in their learning-from-experience forecast by 1 pp, their one-year inflation expectations differ by 0.672 pp on average.

To check whether the imputation of percentage responses from categorical responses affects our results, we re-run the estimation using only those time periods in which percentage responses are directly available. As can be seen in column (ii), not using the imputed data has little effect on the results. We estimate a similar gain parameter, $\theta = 3.144$ (s.e. 0.257),

and a similar sensitivity paramenter $\beta = 0.675$ (s.e. 0.079).

Figure 4 illustrates the extent to which learning-from-experience effects explain cross-sectional differences in inflation expectations. The figure shows both the raw survey data and fitted values based on the estimates in column (i) of Table 1. For the purpose of these plots, we average inflation expectations and the fitted values within the same categories of the young (age < 40), middle-aged (age between 40 and 60) and old (age > 60) that we used earlier in Figure 1. Since our baseline estimation with time dummies focuses on cross-sectional differences, we plot all time series as deviations from the respective population means, i.e., after subtracting their cross-sectional mean each period. To eliminate high-frequency variation, we show four-quarter moving averages for both actual and fitted values. Fitted values are drawn as lines, raw inflation expectations are shown as diamonds (young), triangles (middle-aged) or filled circles (old). The plot shows that the learning-from-experience model does a good job of explaining the age-related heterogeneity in inflation expectations. 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, including the double-spike. It also captures all of the low-frequency reversals in the expectations gap between older and younger individuals.

The presence of the time dummies in these regressions is important to rule out that the estimates pick up time-specific effects unrelated to learning from experience. If individual expectations were unaffected by heterogeneity in inflation experiences—for example, if all individuals learned from the same historical data applying the same forecasting rules—then β would be zero. The effect of historical inflation rates, including "experienced" inflation rates, on current forecasts would be picked up by the time dummies. The fact that β is significantly different zero is direct evidence that differences in experienced-inflation histories are correlated with differences in expectations. The significant β-estimate also implies that recent observations exert a stronger influence on expectations of the young since the set of historical inflation rates experienced by the young that enters into the construction of the learning-from-experience forecast comprises only relatively few observations.
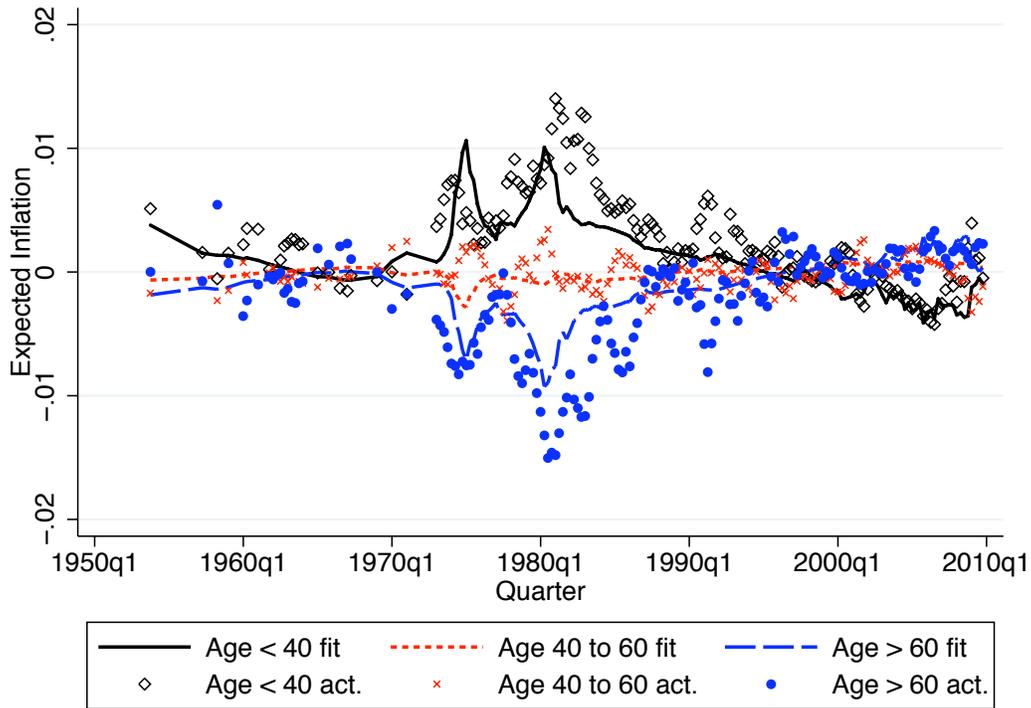
Figure 4: Comparison of four-quarter moving averages of actual and fitted one-year inflation expectations for individuals below age of 40, between 40 and 60, and above 60, shown as deviations from the cross-sectional mean expectation. The fitted values corresponds to column (i) in Table 1.

Figure 5 illustrates the stronger response of younger individuals to recent data. We plot the time series of the persistence and mean parameters for each age group over the course of our sample period. We calculate both parameters based on the θ estimate from Table 1, column (i). For the purpose of this plot, we average these perceived parameters within three age groups (below age of 40, between 40 and 60, and above 60.) The figure reveals that the perceived mean increased up to 1980 and then declined. The path of perceived persistence is flatter but also increases initially and then declines dramatically after 2000. Both graphs illustrate that the assessment of the mean and the persistence among young individuals are much more volatile than among older individuals, with middle-aged individuals lying in between. Our estimates also imply that at the end of the sample period, young individuals'
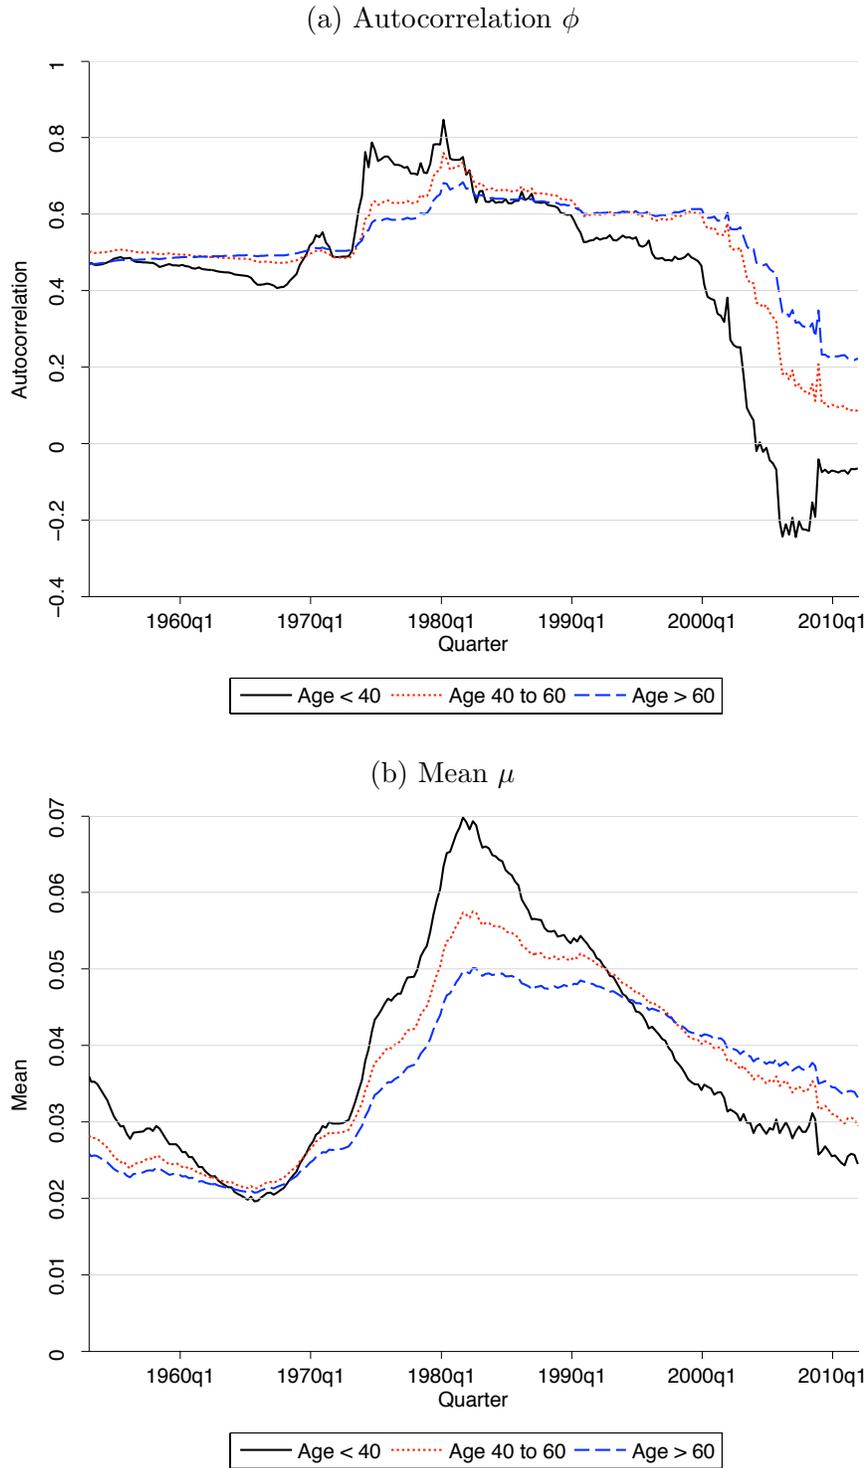
(a) Autocorrelation ϕ

(b) Mean μ

Figure 5: Learning-from-experience AR(1) model estimates (with $\theta = 3.044$) of autocorrelation (top) and mean inflation (bottom) for individuals below age of 40, between 40 and 60, and above 60.

perceived inflation persistence is close to zero. This implies that young individuals inflation expectations at the end of the sample period are well anchored in the sense that they would be relatively insensitive to a short period of higher-than-expected inflation. As Mishkin (2007) and Bernanke (2007) argue, better anchoring of inflation expectations plays an important role in explaining why the dynamics of inflation have changed in recent decades, and Figure 5 illustrates that learning-from-experience effects help understand the source of this improved anchoring. According to our estimates, older individuals' perceived inflation persistence, however, is still substantially above zero.

Finally, it is interesting to note that the weighting of past inflation experiences implied by the point estimates of θ is very similar to the weights estimated in Malmendier and Nagel (2011), who relate data on household asset allocation to experienced risky asset returns.[6] This is remarkable since the data on inflation expectations is drawn from a 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 histories seems to imply a similar weighting of past experiences. Taken together, our findings imply that a common expectations-formation mechanism may be driving both sets of results.

One possible alternative explanation for these findings is that individuals form inflation expectations based on (recent) inflation rates they observe in their age-specific consumption baskets. The concern would be that inflation differentials between age-specific consumption baskets could be a correlated omitted variable, i.e., could be correlated with the differences in age-specific learning-from-experience forecasts that we construct. 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. As reported in Appendix D, our results are unaffected. Appendix

---

[6]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 A), and which is estimated to be in the range from 1.1 to 1.9, depending on the specification.

D also reports a similar analysis with a gasoline price series to check whether age-specific sensitivity to gasoline price inflation drives the results. We find that this extension does not add explanatory power either, nor does it significantly affect our learning-rule parameter estimates. Hence, the cross-sectional differences that we attribute to learning-from-experience effects are not explained by differences in age-specific inflation rates.

3.3 Exploring the common factor

Our main estimating equation in (6) includes time dummies in order to cleanly identify the learning-from-experience effect with a test of the null hypothesis $\beta = 0$. The specification also allows us to estimate θ purely from cross-sectional differences, removing potentially confounding unobserved determinants of expectations. We now ask whether learning-from-experience forecasts can explain not only cross-sectional differences but also the level of expectations. We also explore the nature of the common factor f_t in the underlying structural model (5), which is absorbed by the time dummies in our estimating equation (6). For both steps, we re-estimate equation (6) without the time dummies and intercept, but instead with observable proxies for f_t.

First, we consider the possibility that f_t captures individuals' tendency to rely, to some extent, on the opinions of professional forecasters that get disseminated in the media. To check this, we specify f_t as the sum of the Survey of Professional Forecasters (SPF) forecast in quarter $t-1$ and a noise term η_t that is uncorrelated with the SPF forecast and the learning-from-experience forecast. Equation (5) becomes

$$\pi_{t+1|t,s} = \beta \tau_{t+1|t,s} + (1 - \beta) \, SPF_{t-1} + (1 - \beta) \, \eta_t. \tag{7}$$

Note that this estimating equation does not include time dummies and, hence, utilizes information about the levels of inflation expectations, not only cross-sectional differences. The estimation results are shown in column (iii) of Table 1. In the estimation, we remove the imputed data from the sample, as the imputation was only designed to impute cross-sectional

differences, not levels. The number of observations is further slightly lower than in column (ii) because SPF forecasts are not available in a few quarters early in the sample. As before, we work with one-year forecasts in the survey data, so we use the corresponding four-quarter, iterated version of the learning-from-experience forecast.

As column (iii) shows, replacing the time dummies with the SPF has little effect on the estimate of β compared with column (ii). With 3.976 (s.e. 0.612), the estimate of θ is higher, but the implied weighting of past inflation experiences remains quite similar to the weighting implied by the estimates in columns (i) and (ii). The standard error of β also remains similar while the standard error of θ doubles in column (iii), relative to columns (i) and (ii). The noisier θ estimate reflects that the removal of the time dummies leaves the noise term η_t in (7) in the regression residual. This effect of the noise term can also be seen in the increase in RMSE compared with columns (i) and (ii).[7] Nevertheless, the fact that the β-estimate in columns (iii) is virtually identical to those in columns (i) and (ii) indicates that SPF forecast captures much of the common component of f_t that could be correlated with the learning-from-experience forecast.

Another possibility is that the common component f_t in individuals' beliefs is the result of a social learning process in which individuals with different experienced inflation histories share their opinions, and, as a result, their beliefs have a tendency to converge to the average belief, as in DeGroot (1974). To explore this possibility, we represent f_t as the mean learning-from-experience forecast across all age groups, which we denote as $\bar{\tau}_{t+1|t}$, and a noise term. Now, equation (5) becomes

$$\pi_{t+1|t,s} = \beta \tau_{t+1|t,s} + (1 - \beta) \bar{\tau}_{t+1|t} + (1 - \beta) \eta_t. \tag{8}$$

Column (iv) reports the results for this model. The estimates are almost identical to those in column (iii). Evidently, the average learning-from-experience forecast is very close to the SPF

---

[7]Since this regression is run without intercept, the adjusted R^2 is not a useful measure of fit and we focus on the RMSE.

forecast and it, too, does a good job in absorbing the component of f_t that could be correlated with the learning-from-experience forecast. As in column (iii), though, θ is estimated with substantially higher standard errors than in the specification with time dummies in columns (i) and (ii).

In a last estimation, shown in column (v), we explore the fit after fixing θ at its more precise estimate from column (i), $\theta = 3.044$, which was more cleanly identified (due to the inclusion of time dummies). This specification allows us, on the one hand, to track both the cross-sectional and the time-series variation in inflation expectations induced by learning from experience, and, on the other hand, eliminates the potential confounds affecting the (noisier) estimates of columns (iii) and (iv), which are potentially subject to unobserved omitted correlated common factors. We find that the estimate of β is almost unchanged, and there is little deterioration in fit. The RMSE is only slightly higher than in column (iv). We will use this model specification further below in Section 5 when we explore time-variation in inflation expectations and the aggregate implications of learning from experience. We will also explore whether, relying on clean and precise cross-sectional estimates of θ might result in a better out-of-sample fit for specification (v) than the specifications in columns (iii) and (iv), despite the slightly worse in-sample fit. We return to this point further below in Section 5.4.

4 Inflation experiences and financial decisions

So far our results show that differences in inflation experiences generate differences in beliefs about future inflation. To what extent do these differences in beliefs affect the economic decisions of households? Since differences in inflation expectations generate disagreement about real rates of return on assets and liabilities with nominally fixed rates, households with higher experience-based inflation expectations should be more inclined to borrow and less inclined to invest at nominally fixed rates than households with lower experience-based

inflation expectations. We test this prediction by estimating the regression equation

$$y_{t,s} = \beta_1 \tau_{t+1|t,s} + \beta_2' X_t + \beta_3' A_t + \beta_4' D_t + \xi_{t,s}, \tag{9}$$

where $y_{t,s}$ is a measure of either fixed-rate liabilities or fixed-rate assets held at time t by people born in year s, and $\tau_{t+1|t,s}$ denotes the learning-from-experience forecast of inflation (constructed with the estimate of θ from Table 1, column (i)). X_t is a vector of individual characteristics, A_t a vector of age dummies, and D_t a vector of time dummies. We also add the disturbance $\xi_{t,s}$, which we allow to be correlated within cohorts over time and between cohorts within the same time period.

To estimate the effect of learning-from-experience forecasts on financial decisions, we turn to a different data source, the Survey of Consumer Finances (SCF), which provides detailed information on households' financial situation. We rely on the data set constructed by Malmendier and Nagel (2011), which comprises both the modern triennial SCF from 1983-2007 and older versions of the SCF from 1960-1977. For comparability with our baseline estimation in Table 1, we aggregate the microdata again at the cohort-level.[8] In each survey wave, we construct per-capita numbers of debt and bond holdings (in September 2007 dollars), as well as income and net worth, for all birth-year cohorts. We then run the estimation on the resulting cohort panel. Appendix E provides further detail about the data set and the construction of our variables.

Table 2 provides summary statistics for the key variables in our analysis. Households' primary fixed-rate liability is mortgage debt, shown in column (i). Prior to 1983, the SCF often provides mortgage information only for households' primary residence, not for other real estate owned by the household. To construct a measure that is consistent over time, we define fixed-rate liabilities as the sum of fixed-rate mortgage balances secured by the primary residence. On the asset side, we measure households' holdings of long-term bonds, shown in

---

[8]The cohort-level aggregation also serves to minimize the influence of outliers and erroneous zeros when analyzing ratios such as log(debt/income), or regressing log(debt) on log(income), and is common in the household finance literature that studies borrowing and savings decisions.

Table 2: Survey of Consumer Finances: Summary statistics of cohort aggregates

The *SCF* sample includes 19 surveys during the period from 1960 to 2007, and 18 of those have information on holdings of long-term bonds. The data on borrowing and bond holdings is aggregated to per-capita numbers at the cohort level.

	Log fixed-rate mortgages	Log long-term bonds	Log refi fixed-rate mortgages	Log refi variable-rate mortgages	Log income	Log net worth
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Mean	9.76	8.61	6.71	3.39	10.91	11.63
Std.dev.	1.29	1.85	3.45	3.72	0.41	0.89
p10	8.43	6.26	0.00	0.00	10.42	10.44
Median	10.05	8.73	8.21	0.00	10.93	11.62
p90	10.92	10.87	9.50	8.30	11.42	12.79
Sample	Full	Full	≥ 1983	≥ 1983	Full	Full

column (ii). This variable includes holdings through mutual funds and defined contribution (DC) accounts.

We also tabulate separately the summary statistics for mortgages that were newly taken out or refinanced in the same year during which the survey took place. (The survey is carried out from June to September.) We use these alternative outcome variables when focusing on the flow rather than level of liabilities. We split these (re-)financing volumes into fixed-rate and variable-rate (re-)financings, as shown in columns (iii) and (iv) of Table 2. The information whether mortgages have variable rates or fixed rates is only available starting in 1983, but variable-rate mortgages were largely non-existent in the U.S. prior to the 1980s (see Green and Wachter (2005)).

Finally, columns (v) and (vii) show family income and net worth, which we use as control variables. In the years before 1983, the coverage of household assets in the *SCF* is not as comprehensive as from 1983 onwards. For the sake of comparability over time, our measure of net worth uses only categories of assets and liabilities that are available in all survey waves: financial assets, defined as stocks, bonds, and cash, including mutual funds and DC accounts,

plus equity in the households' primary residence.

Table 3 presents the results of the estimations. In each column we regress the log of the respective cohort-level per-capita nominal position on the learning-from-experience inflation forecast, constructed using our point estimate of $\theta = 3.044$ from Table 1. We control for the logs of income and net worth. All regressions include dummies for the survey year and for age.

Column (i) shows the estimation results for the regression using the total size of households' fixed-rate mortgage positions as the outcome variable. We find that it is positively related to the learning-from-experience inflation forecast, and the point estimate of the coefficient is more than four standard errors above zero. As predicted, households whose experiences lead them to expect higher inflation and, hence, lower real interest rates take on more fixed-rate liabilities. The magnitude of the effect is large: a one percentage point difference in the learning-from-experience forecast corresponds to a 0.35 change in the log of the fixed-rate mortgage balance, which is between a third and a quarter of a standard deviation of the dependent variable (see Table 2). This magnitude is comparable to the variation associated with a one-standard-deviation change in log income.

In column (ii), we estimate the effect of inflation experiences on fixed-rate investment, i.e., the size of households' nominal bond positions. Here, the sign of the coefficient is negative, indicating that households with higher learning-from-experience forecasts of inflation take smaller positions in long-term bonds. The size of the coefficient is on the same order of magnitude as the coefficient in column (i), but the estimate is not statistically different from zero. Taken together, the results in column (i) and (ii) show that households with higher learning-from-experience inflation forecasts tilt their exposure to liabilities rather than assets with nominally fixed rates. As shown in columns (iii) and (iv), we also obtain similar results when restricting the sample period to 1983-2007, when the SCF data is of higher quality.

In columns (v) and (vi), we refine the analysis in two ways. First, we focus on the extent to which households have recently taken out new mortgages or refinanced them, rather than

28

Table 3: Inflation experiences and household nominal positions

The *SCF* sample includes 19 surveys during the period from 1960 to 2007, and 18 of those have information on holdings of long-term bonds. The data on borrowing and bond holdings is aggregated to per-capita numbers at the cohort level. Each cohort is assumed to recursively estimate an AR(1) model of inflation, with $\theta = 3.044$, as in Table 1, column (i). We use the resulting learning-from-experience forecast of inflation to explain log fixed-rate mortgage borrowing and log long-term bond holdings in OLS regressions. Log mortgage borrowing in columns (v) and (vi) comprises only loans taken out or refinanced in the year in which the survey was carried out. Standard errors reported in parentheses are clustered by time period.

	Fixed-rate mortgages (i)	Long-term bonds (ii)	Fixed-rate mortgages (iii)	Long-term bonds (iv)	Fixed-rate refi (v)	Variable-rate refi (vi)
Learn.-from-exp. forecast	35.27	-20.56	26.77	-9.07	132.71	-42.82
	(8.39)	(13.74)	(4.47)	(6.92)	(25.08)	(55.57)
Log income	0.92	0.45	0.60	0.02	1.23	2.60
	(0.16)	(0.25)	(0.13)	(0.13)	(1.19)	(1.29)
Log net worth	-0.10	1.09	0.18	1.18	-0.56	-1.79
	(0.15)	(0.13)	(0.06)	(0.10)	(0.69)	(0.94)
Time dummies	Yes	Yes	Yes	Yes	Yes	Yes
Age dummies	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Full	Full	≥ 1983	≥ 1983	≥ 1983	≥ 1983
Adj. R^2	0.617	0.852	0.856	0.915	0.485	0.243
#Obs.	950	900	450	450	450	450

the total mortgage position of the household. The flow variable addresses concerns about the "stickiness" of mortgage positions: Total fixed-rate mortgage balances include loans taken out many years ago and may change only slowly over time. Once a household has taken out a mortgage and bought a house, the mortgage balance cannot easily be adjusted. While a household can take out a second mortgage, buy a bigger house, or (since the 1980s) refinance with a variable-rate mortgage, there are frictions and indivisibilities that are likely to generate substantial stickiness. In contrast, for the volume of newly taken out or refinanced mortgages, this stickiness plays less of a role. Using the more recent data on re-financing, we are also able

to distinguish between fixed-rate and variable-rate mortgages. According to our hypothesis, households with higher learning-from-experience forecasts of inflation should be more likely to take out a fixed-rate mortgage in a new loan or refinancing, and less likely to take out a variable-rate mortgage.

The results in column (v) and (vi) are consistent with this prediction. We find that households with high learning-from-experience forecasts of inflation are significantly more likely to take out new fixed-rate mortgages and to re-finance at fixed rates. A one percentage point difference in the learning-from-experience forecast corresponds to a roughly 1.33 change in the log of the fixed-rate mortgage balance, which is more than a third of a standard deviation of the dependent variable according to Table 2. We also find that experience-based inflation expectations are negatively related to the volume of new variable-rate mortgages, but the estimated coefficient in column (vi) is not statistically significant, and the point estimate is small—a one percentage point difference in the learning-from-experience forecast corresponds only to about a ninth of the standard deviation of the dependent variable.

Overall, the findings in this section confirm that learning from inflation experiences not only affects the expectations of individuals, but also helps understand patterns in the financial decision-making of households. These results also speak to the relevance of expectations data from household surveys (as opposed to surveys of professional forecasters) for real economic variables. For the decisions that we study here—asset allocation to long-term bonds and mortgage financing—we estimate a significant effect of individuals' perceptions on aggregate outcomes. On the asset side, this might seem surprising because of the widespread delegation of portfolio management to professional fund managers. However, households control much of the asset allocation decisions in the economy when they allocate their wealth to mutual funds that are restricted to invest in certain asset classes, including the choice stocks versus bonds. On the liabilities side, the results might be less surprising as households are the main players as buyers and sellers in the residential real estate market and as borrowers in residential mortgage financing. In summary, the findings in Table 3 show that the learning-

from-experience model helps understand household behavior in these important markets.

5 Aggregate implications

Our analysis so far has focused on explaining heterogeneity in expectations and financial decisions based on heterogeneity in inflation experiences. We now test whether learning from experience also helps explain aggregate dynamics. We show that experience-based forecasts aggregate to average forecasts that closely resemble those from constant-gain algorithms in the existing literature, which have been shown to explain macroeconomic time series data. We argue that learning from experience provides a micro-underpinning for adaptive-learning models, but offers conceptual and econometric advantages in the identification of the structural parameters that pin down the learning rule.

5.1 Approximating constant-gain learning

We start from the model in (8), which allows for social learning. Averaging across all cohorts s in each period t, and denoting cross-sectional averages with an upper bar, we get

$$\bar{\pi}_{t+1|t} = \bar{\tau}_{t+1|t} + (1 - \beta) \, \eta_t. \tag{10}$$

Thus, apart from the noise term η_t, the mean expectation is pinned down by the mean learning-from-experience forecast across all age groups, $\bar{\tau}_{t+1|t}$. This mean learning-from-experience forecast behaves approximately as if it was generated from a *constant*-gain learning algorithm: While individuals update their expectations with *decreasing* gain, i.e., older individuals react less to a given inflation surprise than younger individuals, the gain $\gamma_{t,s}$ varies only by age $t - s$ and hence its average each period is constant (as long as the weight on each age group is constant over time). The average learning-from-experience forecast is an approximation (rather than an exact match) of a constant-gain learning forecast because the means of the surprise terms in (2) and (3) are not exactly identical to the surprises arising
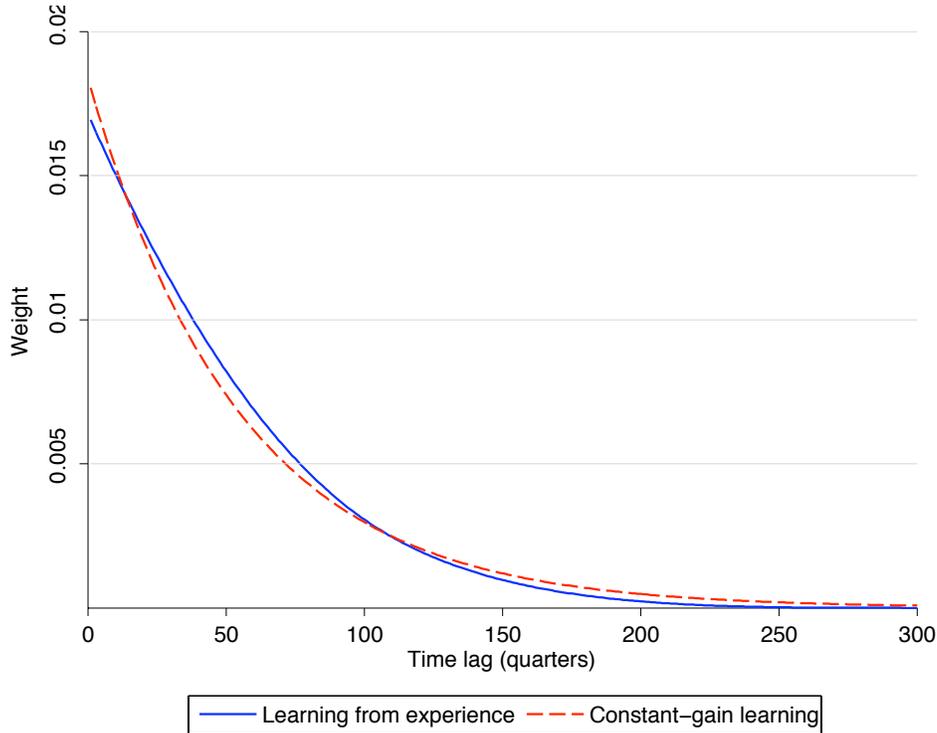
31

Figure 6: Comparison of implied mean weights on past inflation observations under learning-from-experience and constant-gain learning. The learning-from-experience weights for each lag are calculated for each age at the point estimate $\theta = 3.044$ from Table 1, column (i), and then averaged across all ages. The weights implied by constant-gain learning are calculated with gain $\gamma = 0.0180$, which minimizes squared deviations from the learning-from-experience weights shown in the figure.

in a constant-gain learning algorithm.

Figure 6 illustrates how well the approximation with a constant gain works. The figure compares the implied learning-from-experience weights on past inflation averaged across all age groups (solid line) with the weights implied by constant-gain learning (dashed line). We calculate the learning-from-experience weights based on our point estimate of $\theta = 3.044$ from Table 1, column (i), averaged across all ages from 25 to 74 (equally weighted). We then calculate the constant gain for which the constant-gain algorithm minimizes the squared deviations from the average learning-from-experience weights. The result is a constant gain of $\gamma = 0.0180$.

The figure shows that the weighting of past data is very similar. It is noteworthy that the deviation-minimizing constant gain $\gamma = 0.0180$ is virtually the same as the gain required to match aggregate expectations and macro time-series data. For example, Milani (2007) reports that an estimate of $\gamma = 0.0183$ provides the best fit of a DSGE model with constant-gain learning to macroeconomic variables. Orphanides and Williams (2005a) choose a gain of 0.02 to match the time series of inflation forecasts from the SPF. Our estimate of γ is, instead, chosen to match the weights implied by our estimate of θ from cross-sectional heterogeneity. We did not employ aggregate expectations data and we did not try to fit future realized inflation rates. Hence, our estimates from between-cohort heterogeneity provide "out-of-sample" support for the values of the gain parameter chosen to fit time-series data in previous literature. We conclude that the aggregate implications of learning from experience for the formation of expectations are very similar to those of the constant-gain learning algorithms, which have been shown to be a good description of macroeconomic dynamics (e.g., Sargent (1999), Orphanides and Williams (2005a), Milani (2007)).

5.2 Explaining aggregate expectations

We now test directly how well the learning-from-experience model matches aggregate survey expectations. Figure 7 shows both the time path of averages from the raw survey data (circles) and average experience-based forecasts (solid line), as before based on $\theta = 3.044$. Since our imputation of percentage responses only targeted cross-sectional differences, but not the average level of percentage expectations, we omit all periods in which we only have categorical inflation expectations data.

It is apparent from the figure that the average learning-from-experience forecasts track the average survey expectations closely. The good match is by no means mechanical: Our estimation of θ used only cross-sectional differences in survey expectations, but no information about the level of the average survey expectation. It could have been possible that the θ that fits cross-sectional differences produces a time path for average expectations that fails to
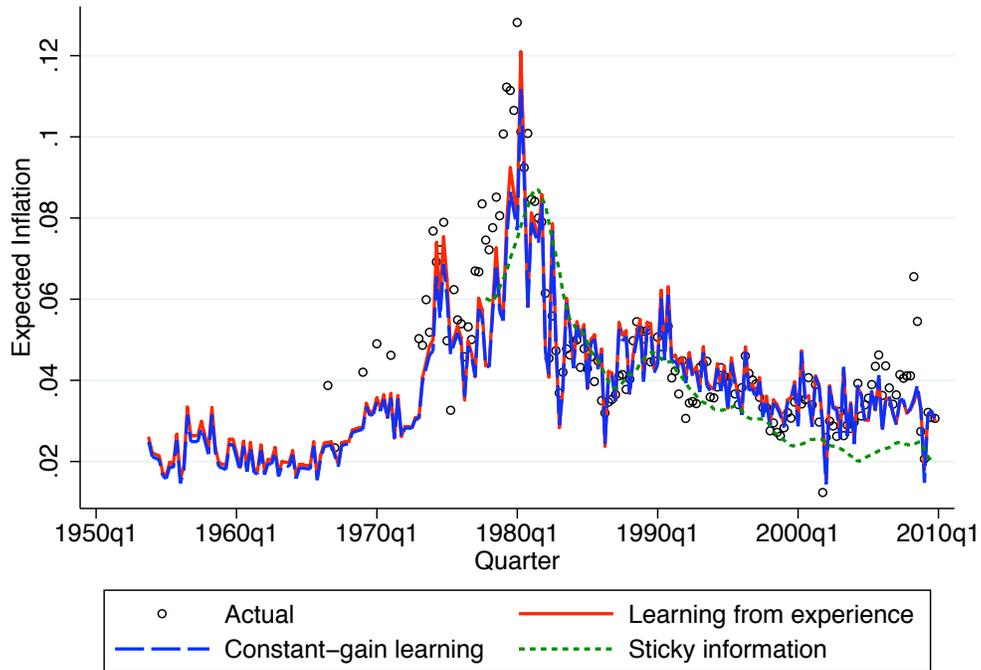
Figure 7: Average survey expectations (actual) compared with average learning-from-experience forecasts, constant-gain forecasts, and sticky-information forecasts.

match average survey expectations. As the figure shows, though, the two time paths match well.

Figure 7 also shows the time path of constant-gain-learning forecasts, using $\gamma = 0.0180$ from Figure 6. Not surprisingly, given that γ was chosen to minimize the distance in the implied weights, the forecasts are almost indistinguishable. This illustrates further that, at the aggregate level, the learning-from-experience expectations formation mechanism can be approximated well with constant-gain learning.

Finally, 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

Table 4: 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 five lags) are shown in parentheses.

	(i)	(ii)	(iii)	(iv)
Learning-from-experience forecast	0.887			0.695
	(0.120)			(0.132)
Constant-gain-learning forecast		0.931		
		(0.137)		
Sticky-information forecast			0.878	0.390
			(0.199)	(0.150)
Intercept	0.009	0.008	0.011	-0.000
	(0.005)	(0.005)	(0.006)	(0.005)
Adj. R^2	0.564	0.555	0.602	0.715
#Obs.	173	173	129	129

past quarterly SPF forecasts of one-year inflation rates.[9] We set the weight parameter $\lambda = 0.25$ as in Mankiw and Reis (2002) (and similar to $\lambda = 0.27$ estimated in Carroll (2003)). The resulting sticky-information forecast is shown as the short-dashed line in Figure 7. The graph illustrates that the learning-from-experience model helps to predict actual forecast data beyond the predictive power of the sticky-information model. For example, learning-from-experience forecasts track actual forecasts more closely during the peak around 1980 and also during the 2000s, though both models fail to match a few highly positive expectations towards the end of the the decade.

We evaluate the economic and statistical significance of this graphical impression in Table 4. We regress the average survey expectations in quarter t on the average forecast predicted by learning-from experience (column (i)), by constant-gain learning (column (ii)), and by the sticky-information model (column (iii)). The learning-from-experience model in (10) predicts a coefficient on the experience-based forecast of one, and column (i) shows that the

---

[9]We use the one-year inflation forecasts that the SPF constructs from median CPI inflation forecasts for each of the four quarters ahead. Before 1981Q3, when the CPI inflation forecast series is not available, we use the GDP deflator inflation forecast series.

estimated coefficient of 0.887 is close to one, and less than one standard error away from it. With 56.4% the adjusted R^2 is high. This confirms the informal graphical impression in Figure 7 that the learning-from-experience forecast closely tracks the actual average survey expectations. Not surprisingly, given the similarity of $\bar{\tau}_{t+1|t}$ and constant-gain learning forecasts using $\gamma = 0.0180$, the constant-gain learning forecast produces almost identical results. The explanatory power of the sticky-information forecast in column (iii) is also similar, only a bit lower and a bit noisier, and the adjusted R^2 is slightly higher. Most importantly, if we add the sticky-information forecast as an explanatory variable along with the learning-from-experience forecast (column (iv)), the coefficient on the learning-from-experience forecast becomes slightly smaller, but remains large (also relative to the sticky-information coefficient) and significant. Hence, we can conclude 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 A foundation of perpetual learning

Perpetual learning plays a centrol role in explaining macroeconomic dynamics, as emphasized, for example, by Sargent (1999), Orphanides and Williams (2005a), and Milani (2007). It is therefore important to identify the underlying reasons for perpetual learning; only then is it possible to predict the circumstances under which economic agents update with a high or low gain.

Despite their similarity at the aggregate level, models of experience-based learning and constant-gain learning differ fundamentally in their motivation for the down-weighting of past data, and resulting perpetual learning. The standard motivation in constant-gain models for the discounting of past data, and resulting perpetual learning, is the concern that structural changes or drifting parameters have rendered historical data from the distant past irrelevant for the estimation of current parameters of the perceived law of motion. Learning from experience attributes the down-weighting to memory of past data being lost as older generations

die and new ones are born.

The standard rationalization has been difficult to reconcile with the data. If individuals' concerns about structural change and parameter drift motivate the down-weighting of data in the distant past, the degree of down-weighting, and hence the gain, should depend on the nature of the stochastic process that individuals are trying to estimate. For example, individuals should use high gains for processes that are subject to frequent structural changes and substantial parameter drift. Empirically, however, there seems to be a disconnect between statistically optimal gains and the gains implied by survey data. Using a quadratic loss function, Branch and Evans (2006) find that the optimal gain that minimizes the loss from inflation forecast errors is substantially higher than the gain that delivers the best fit to the Survey of Professional Forecasters (SPF) inflation forecasts.

We show that a similar discrepancy exists in the Michigan Survey data. Figure 8 presents, on top, recursive estimates of the optimal gain that minimizes the MSE in constant-gain learning forecasts of one-year inflation rates. (The first estimation window extends from 1953Q4 to 1977Q4; subsequent ones expand quarter by quarter while keeping the starting point fixed.) For the most part, the optimal gain is relatively high, around 0.08, similar to the optimal gain calculated by Branch and Evans (2006). The gains implied by our learning-from-experience model fitted to MSC expectations data, instead, are much lower, similar to the SPF-based gain estimates in Branch and Evans (2006). We calculate these implied gains by estimating θ with expanding estimation windows, and then, as in Figure 6, picking the gain in a constant-gain learning rule that most closely matches the learning-from-experience model. Figure 8 also plots the gain estimates from fitting a constant-gain learning rule to mean inflation expectations from the MSC.[10] The estimated gains are slightly higher than the implied gains from the learning-from-experience rule, but the difference is small compared with the big gap to the statistically optimal gains. Apparently, both professionals (in Branch

---

[10]Since our imputation procedure is not suitable for imputation of mean expectations, only for cross-sectional differences, we use only non-imputed data in the estimation of the constant-gain learning rule. For comparability, we therefore also use non-imputed data in the estimation of the learning-from-experience model, i.e., as in column (ii) of Table 1 rather than column (i).
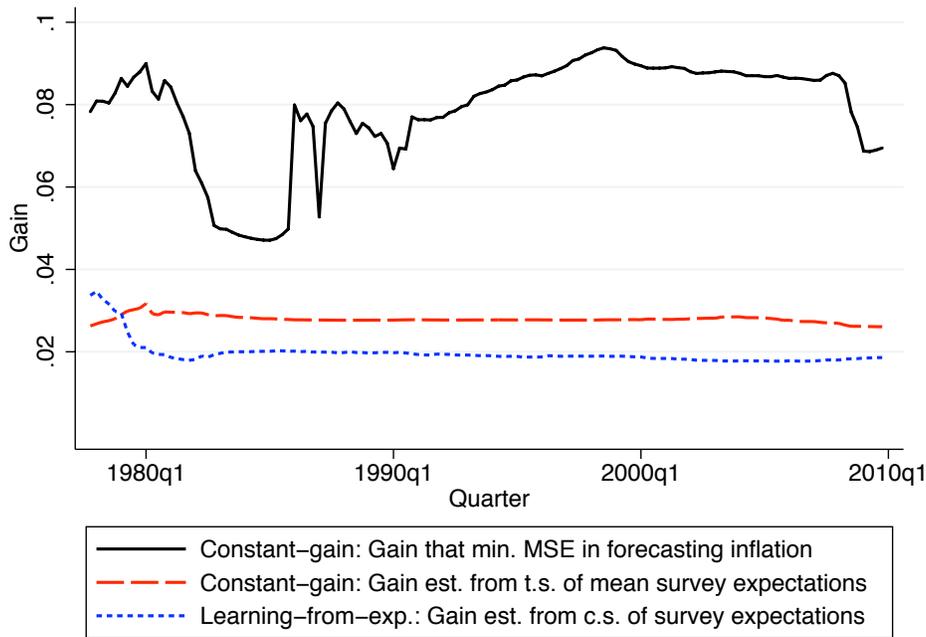
Figure 8: Recursive comparison of optimal gain with gains implied by constant-gain and learning-from-experience model estimates. The optimal gain (solid line) minimizes the MSE in one-year inflation forecasts, and it is estimated over expanding windows, where the first extends from 1953Q4 to 1977Q4. Gains implied by constant-gain (long-dashed line) and learning-from-experience (short-dashed line) models are estimated over the same windows, employing only non-imputed data. The constant-gain model is fitted to mean expectations, and the learning-from-experience model is fitted to cohort data as in column (ii) in Table 1, but with expanding estimation windows, and the estimate of θ is converted to an implied gain as in Figure 6.

and Evans (2006)) and households (in our analysis) discount past inflation data much less than the degree of structural change in the time-series of inflation would call for.

Moreover, the differences in the degree of structural change observed in different time series do not seem to be systematically related to the gains implied by survey expectations— at least in the relatively stable environment of post-war US data.[11] For example, Branch and Evans (2006) find that the statistically optimal gain for GDP forecasting is much lower than the optimal gain for inflation forecasting—indicating a lower degree of structural change in

---

[11]Marcet and Nicolini (2003) propose a model where individuals update their expectations with decreasing gain in stable periods, but they switch to a constant gain when a large prediction error is detected.

the GDP series—but the gains that each fit best the survey expectations on inflation and GDP are almost identical.

The lack of connection between the degree of structural change in macroeconomic time series and the gains in survey-data estimates suggests that alternative reasons besides structural change may account for the discounting of past observations that is evident in survey expectations data. Experience-based learning provides such an alternative explanation: memory of past data is lost as older generations die and new ones are born. Thus, even if individuals weighted all data realized during their lifetimes equally, the mean expectation in the population would resemble the forecast from a constant-gain learning model. However, differently from the traditional rationalization in constant-gain models, the experience-based explanation is consistent not only with the time-series of aggregate inflation expectations, but also with cross-sectional heterogeneity in inflation expectations and in borrowing/investing decisions, as we have shown above.

Note, however, that the memory loss implied by the finite life span of generations does not fully explain the level of the implied gain we estimate. Equally weighting each life-time observation of inflation would imply $\theta = 1$. Hence our estimate of $\theta = 3.044$ (in the full sample) implies additional down-weighting within a lifetime, for example a slow fading of experiences as time progresses (and limited absorption of experiences very early in life). The latter explanation appears to be more natural than individuals' concerns about structural change given the inconsistencies between the structural-change based rationalization and the data, as well as our evidence from cross-sectional heterogeneity that individuals' forecasts are shaped by their life-time experiences.

5.4 Out-of-sample predictions

Another advantage of the learning-from-experience model is that it makes predictions about cross-sectional heterogeneity in expectations. This in turn implies that empirically observed cross-sectional heterogeneity in expectations provides useful information that can

help estimate the parameters of individuals' learning rules. This is a key difference from representative-agent applications of adaptive learning models.

As Chevillon, Massmann, and Mavroeidis (2010) show, the identification of structural parameters in representative-agent macro models with adaptive learning is difficult, and the problems are magnified if the parameters of the learning rule are unknown and need to be estimated. Fitting the learning rule to the time path of mean or median survey expectations can help pin down the learning-rule parameters, but the estimates may be imprecise. Within the learning-from-experience model, the gain parameter θ can be identified from cross-sectional data. This brings in a new dimension of data that can help pin down the learning dynamics. We illustrate this point by comparing the out-of-sample fit of the different models.

In our estimations so far, we have shown that we obtain the most precise estimates of the gain parameter θ when we focus purely on cross-sectional variation by employing time dummies (Table 1). Here we show that this way of estimating the gain also yields the best pseudo-out-of-sample fit to the time-series of mean survey expectations. Figure 9 compares the pseudo-out-of-sample fit of the learning-from-experience model with the constant-gain learning model. We estimate the gain parameters in both models recursively, with expanding windows, where the first window extends from 1953Q4 to 1977Q4. For each window, expectations data until quarter $t-1$ is used to estimate the gain parameter (mean expectations in case of the constant-gain model, and cohort data as in column (ii) of Table 1 in case of the learning-from-experience model), and we then predict, based on this gain estimate and historical inflation data until $t-1$, the mean inflation survey expectation in quarter t. In case of the learning-from-experience model this prediction is given by $\bar{\tau}_{t+1|t}$ as in (10), but with θ estimated only from expectations data up to quarter $t-1$; in the case of constant-gain learning it is simply the fitted value of the constant-gain learning rule. The figure plots the cumulative sum of squared errors from these predictions from 1978Q1 to the end of the sample.

As Figure 9 shows, the learning-from-experience rule performs slightly better than the
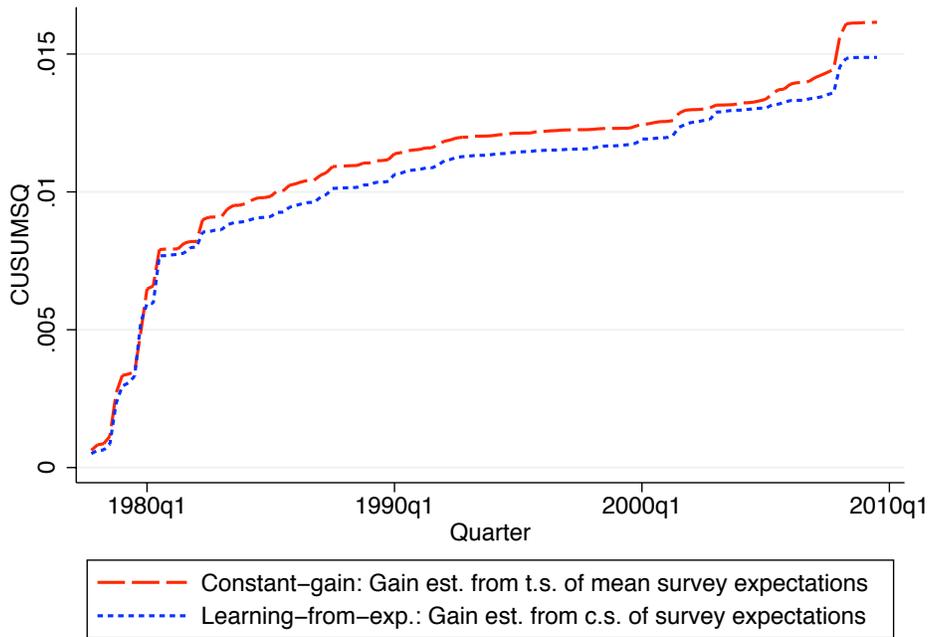
Figure 9: Pseudo-out-of-sample cumulative sum of squared errors in predicting mean inflation expectation. For the constant-gain model, the gain parameter γ is estimated over expanding windows, where the first extends from 1953Q4 to 1977Q4. For each window, mean expectations data until quarter $t - 1$ is used to estimate γ, and this estimate of γ is then used to predict, based on inflation data until $t - 1$, the mean inflation survey expectation in quarter t. The plot cumulates the sum of squared errors from this prediction. For learning from experience, the plots are constructed in similar fashion, but based on estimates of θ from cohort-panel expectations data as in column (ii) in Table 1, but with expanding estimation windows.

constant-gain rule. It achieves this advantage even though no information on the level of mean expectations is used in the estimation of θ, as all of it is absorbed by the time dummies. Evidently, the time dummies are helpful in absorbing common noise factors in the mean survey expectations that are otherwise obscuring the relationship between survey respondents' inflation expectations and historical inflation data.

6 Conclusion

Our empirical analysis shows that individuals' inflation expectations differ depending on 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. The experience of younger individuals is dominated by recent observations, while older individuals draw on a more extended historical data set in forming their expectations.

Such learning from experience 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. Both the mean rate of inflation and inflation persistence were particularly high in the short data set experienced by young individuals at the time. Learning-from-experience complements the sticky information explanation of expectations heterogeneity put forward in Mankiw and Reis (2002) and Carroll (2003) for the same time period.

For the more recent time period towards the end of our sample in 2010, our estimates imply that the 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 (2005a), and Milani (2007), these changes in individuals' perceptions of persistence are likely to influence, in turn, the persistence of inflation rates.

The learning-from-experience effects show up not only in data on inflation expectations, but also in financial decisions of households. The experience-induced disagreement between cohorts about future inflation implies disagreement about real interest rates. We find that households with higher experience-based forecasts of inflation are more inclined to borrow rather than invest at nominally fixed long-term interest rates.

The learning-from-experience framework differs from more conventional representative-agent applications of learning in that it generates heterogeneity in inflation expectations. Nevertheless, 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 (2005a); Milani (2007)). There are, however, two important differences.

First, the learning-from-experience theory points to a different and complementary reason why data in the distant past are down-weighted and learning dynamics are perpetual. While standard implementations of constant-gain learning motivate a gradual loss of memory with structural shifts and parameter drift, learning from experience implies that memory of macroeconomic history is lost as new generations emerge whose subjective beliefs are shaped by relatively recent experience.

Second, the learning-from-experience framework allows us to estimate individuals' reaction to inflation surprises (their gain parameter) from heterogeneity between cohorts. This opens up a new dimension of data as a source of identification. Remarkably, even though we rely purely on cross-sectional variation in survey expectations in estimating the gain, the implications of our estimates for time-variation in the beliefs of the average person are quantitatively similar to those obtained in earlier work in macroeconomics that estimated the gain to fit macroeconomic time-series and aggregate survey expectations.

The expectations heterogeneity generated by learning from experience has the potential to produce interesting macroeconomic effects. For example, Piazzesi and Schneider (2012) show that the disagreement about future inflation and real interest rates between younger and older households in the late 1970s helps understand *aggregate* quantities of household borrowing and lending and the prices of real assets at the time. Our findings help understand the reasons for this dispersion in inflation expectations.

Appendix

A Implied weighting of past data with learning from experience

We derive the weighting of past data implied by the learning-from-experience algorithm proposed in equations (1) to (4), and show that the implicit weights on past observations correspond almost exactly to 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 λ that controls the weighting function in Malmendier and Nagel (2011). This makes the results easily comparable.

Consider an individual born in year s who makes an inflation forecast at time t. We can rewrite (2) as

$$
\begin{aligned}
R_{t,s}b_{t,s} &= R_{t,s}b_{t-1,s} + \gamma_{t,s}x_{t-1}(\pi_t - b'_{t-1,s}x_{t-1}) \\
&= (R_{t,s} - \gamma_{t,s}x_{t-1}x'_{t-1})b_{t-1,s} + \gamma_{t,s}x_{t-1}\pi_t \\
&= (1 - \gamma_{t,s})R_{t-1,s}b_{t-1,s} + \gamma_{t,s}x_{t-1}\pi_t \\
&= (1 - \gamma_{t-1,s})(1 - \gamma_{t,s})R_{t-2,s}b_{t-2,s} + \gamma_{t-1,s}(1 - \gamma_{t,s})x_{t-2}\pi_{t-1} + \gamma_{t,s}x_{t-1}\pi_t, \quad (\text{A.1})
\end{aligned}
$$

where the second-to-last equality follows from (3). Similarly, we can rewrite (3) as

$$
\begin{aligned}
R_{t,s} &= R_{t-1,s} + \gamma_{t,s}(x_{t-1}x'_{t-1} - R_{t-1,s}) \\
&= (1 - \gamma_{t,s})R_{t-1,s} + \gamma_{t,s}x_{t-1}x'_{t-1} \\
&= (1 - \gamma_{t-1,s})(1 - \gamma_{t,s})R_{t-2,s} + \gamma_{t-1,s}(1 - \gamma_{t,s})x_{t-2}x'_{t-2} + \gamma_{t,s}x_{t-1}x'_{t-1}. \quad (\text{A.2})
\end{aligned}
$$

Thus, iterating further on (A.1) and (A.2), we can write $R_{t,s}b_{t,s} = X'\Omega Y$ and $R_{t,s} = X'\Omega X$, where X collects stacked x_{t-1-k} and Y collects stacked π_{t-k} for $k \in \{0, 1, ..., t-s\}$, and Ω is a diagonal matrix. Thus, $b_{t,s} = R_{t,s}^{-1}R_{t,s}b_{t,s} = (X'\Omega X)^{-1}X'\Omega Y$, i.e., the learning-from-experience scheme is a recursive version of a weighted-least squares estimator with weighting matrix Ω, and the diagonal elements of Ω can be expressed recursively as

$$
\widetilde{\omega}_{t,s}(k) = \widetilde{\omega}_{t,s}(k-1)\frac{1 - \gamma_{t-k+1,s}}{\gamma_{t-k+1,s}}\gamma_{t-k,s} \quad (\text{A.3})
$$

for $0 \leq k \leq t-s$ with initial condition $\widetilde{\omega}_{t,s}(-1) = \frac{\gamma_{t+1,s}}{1-\gamma_{t+1,s}}$.[12]

For comparison, the weighting function in Malmendier and Nagel (2011) assigns observa-

---

[12]The purpose of the initial value is solely to initialize the recursion; it plays no role in the weighting matrix Ω.

tions at time $t - k$ (with $0 \leq k \leq t - s$) the weight[13]

$$\omega_{t,s}(k) = \frac{(t - s - k)^{\lambda}}{\sum_{j=0}^{t-s}(t - s - j)^{\lambda}}. \tag{A.4}$$

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

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

We present a proof by induction. First, the choice of $\gamma_{t,s}$ in (A.5) implies that $\widetilde{\omega}_{t,s}(0) = \omega_{t,s}(0)$. It remains to be shown that if $\widetilde{\omega}_{t,s}(k) = \omega_{t,s}(k)$, then $\widetilde{\omega}_{t,s}(k+1) = \omega_{t,s}(k+1)$ (with $k + 1 \leq t - s$). Thus, assume that

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

Substituting (A.5) into (A.3) we obtain

$$
\begin{aligned}
\widetilde{\omega}_{t,s}(k+1) &= \frac{(t - k - 1 - s)^{\lambda}}{\sum_{j=0}^{t-k-1-s}(t - k - 1 - s - j)^{\lambda}} \left(\frac{\sum_{j=0}^{t-k-s}(t - k - s - j)^{\lambda}}{(t - k - s)^{\lambda}} - 1 \right) \widetilde{\omega}_{t,s}(k) \\
&= \frac{(t - k - 1 - s)^{\lambda}}{\sum_{j=0}^{t-k-1-s}(t - k - 1 - s - j)^{\lambda}} \frac{\sum_{j=0}^{t-k-1-s}(t - k - 1 - s - j)^{\lambda}}{(t - k - s)^{\lambda}} \widetilde{\omega}_{t,s}(k) \\
&= \frac{(t - k - 1 - s)^{\lambda}}{\sum_{j=0}^{t-s}(t - s - j)^{\lambda}} \\
&= \omega_{t,s}(k+1),
\end{aligned} \tag{A.7}
$$

where for the second-to-last equality we have used (A.6). This concludes the induction proof.

Finally, as last step, we show that the gain specification in (A.5) is approximately equivalent to the gain specification (4) that we use in the empirical analysis. Focusing first on the denominator of (A.5), note that if one makes the increments of the summation infinitesimally small, the denominator becomes $\int_0^{t-s} x^{\lambda} dx = \frac{1}{\lambda+1}(t - s)^{\lambda+1}$. Therefore, in the limiting case

---

[13]Note that the weighting scheme is presented slightly differently in Malmendier and Nagel (2011), with weights assigned to past data starting at $k = 1$, rather than at $k = 0$. This hardwires into the weighting a one-period lag that reflects the fact that investment choices measured during period t cannot fully reflect asset return experiences until the end of period t. The weighted average experienced return A_{it} in their notation is therefore equivalent to $b_{t-1,s}$ in the notation here. Our weighting here does not hardwire the time lag into the weighting scheme (which time lag is appropriate depends on the specifics of the empirical application). Instead, we simply relate, in our estimation, survey inflation expectations measured during quarter t to $b_{t-1,s}$ and inflation rates leading up to end of quarter $t - 1$. Also, in Malmendier and Nagel (2011) the summation term runs to $t - s - 1$, not $t - s$, but this makes no difference as $0^{\lambda} = 0$, and letting it run to $t - s$ is helpful when we take the limit to infinitesimal time increments below.

of infinitesimal increments, we get

$$\gamma_{t,s} = \frac{\lambda + 1}{t - s}. \tag{A.8}$$

In our case with quarterly increments, this approximation is virtually identical to the true gain sequence in (A.5). Hence, the (implicit) weights put on past observations in the learning-from-inflation experiences model here and in the experience-based model of stock-market investment in Malmendier and Nagel (2011) are approximately equivalent if the gain sequence is chosen to be age-dependent as defined in (4). Equation (A.8) also illustrates how the parameter θ that controls the strength of updating in the gain sequence maps into the parameter λ that controls the weighting function in Malmendier and Nagel (2011).

B Michigan Survey data

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

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?"

We follow Curtin (1996) to adjust the raw data for several known deficiencies, which is also the approach used by the Michigan Survey in constructing its indices from the survey data: (1) For respondents who provided a categorical response of "up" ("down"), but not a percentage response, we draw 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. (2) 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. (3) 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."

Also, in surveys before 1960, the age of the respondent was collected as a categorical variable. In those years we only have five or nine age groups. From 1960 onwards, the exact birth year was collected as age variable and we have 50 age groups (age 25 to 74).
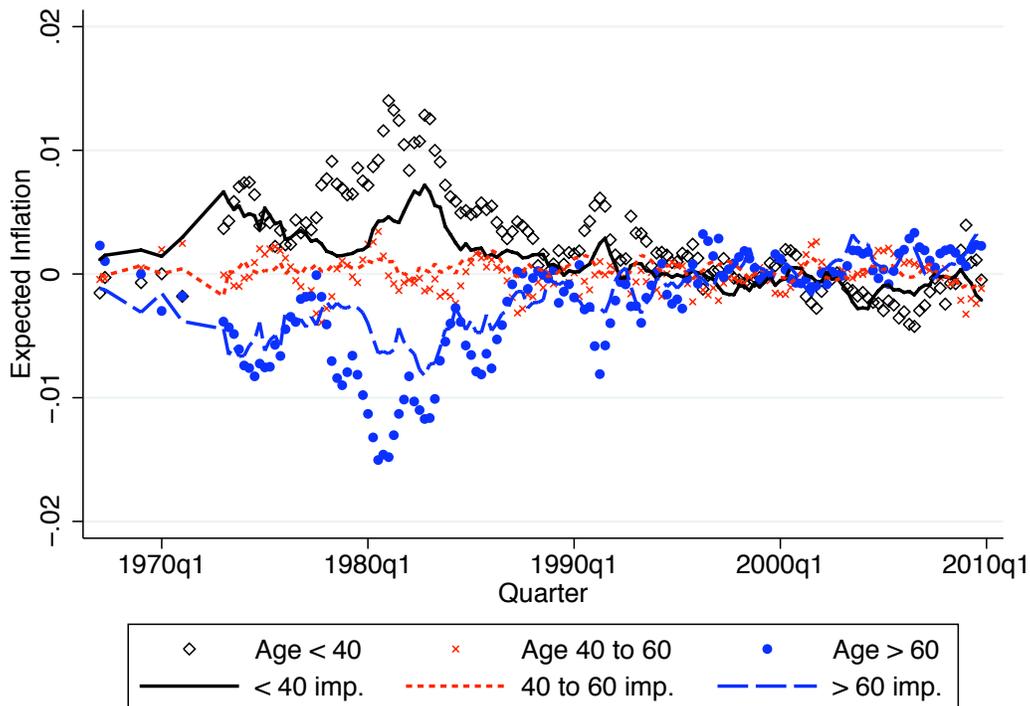
Figure A.1: Four-quarter moving averages of actual and imputed mean one-year inflation expectations of young individuals (below 40), mid-aged individuals (between 40 and 60), and old individuals (above 60), shown as deviations from the cross-sectional mean expectation.

C Imputation of percentage expectations from categorical responses

In the early years of the Michigan survey, only categorical responses about prices going "up," "down," or staying the "same" were elicited. We attempt to use the information in those surveys 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.

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, $\tilde{\pi}_{|t,s}$, on $p_{t,s}^{up}$ and $p_{t,s}^{down}$, including a

full set of time dummies, and obtain the fitted values

$$\tilde{\pi}_{|t,s} \quad = \quad ...\text{time dummies}... + \underset{(0.001)}{0.056}p_{t,s}^{up} - \underset{(0.004)}{0.072}p_{t,s}^{down} \quad (\text{adj. } R^2 = 52.9\%)$$

with standard errors in parentheses (two-way clustered by quarter and cohort). We employ time dummies, mirroring our main analysis: our concern here is whether the imputed expectations track well cross-sectional differences of expectations across age groups, rather than the overall mean over time.

Figure A.1 illustrates how the imputed expectations compare with the actual expectations in time periods in which we have both categorical and percentage expectations. As in Figure 1, we show four-quarter moving averages for individuals below 40, between 40 and 60, and above 60 years of age, after subtracting the period-specific cross-sectional mean.

D   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 using 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 difference between the CPI-E and CPI inflation rates, $\pi_{t-1}^{Elderly} - \pi_{t-1}$, interacted with age. These inflation rates are measured over four quarters leading up to quarter $t-1$ (calculating this difference term with quarterly inflation rates produces similar results).

Table A.1 presents the results. The inflation series based on the CPI-E is only available from the end of 1983 onwards. As a basis for comparison, we therefore first re-run the baseline regression without the additional age-dependent inflation control on the shorter sample from 1984Q1 to 2009Q4. The results in column (i) show that the estimate of the gain parameter is quite similar to the earlier estimate in Table 1. The sensitivity parameter β is estimated to be lower than before but it remains statistically as well as economically significant. In column (ii) we add the interaction term between age-related inflation differentials and age, as well as age itself. (The difference term $\pi_{t-1}^{Elderly} - \pi_{t-1}$ 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 have some effect on the estimates for θ. With 3.475, the point estimate is higher than in column (i), though the difference is not significant.

Column (iii) reports a similar test using the gasoline component of the CPI instead of the CPI-E. This series is available from the Bureau of Labor Statistics (with gaps) since 1935. Hence, we can use our full sample of survey inflation expectations from 1953 to 2009 (and there is thus no need to rerun the baseline estimation.) The estimates are directly comparable to those in column (i) of Table 1. As column (iii) shows, adding the interaction of $\pi_{t-1}^{Gas} - \pi_{t-1}$ with age does not have much effect. The interaction term is insignificantly

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

The estimation is similar to Table 1, but with interactions of age with the experimental CPI for the elderly in column (ii) and the CPI for gasoline in column (iii) included as control variable. In columns (i) and (ii), the sample runs from 1984 to 2009, the period for which lagged four-quarter inflation rates from the experimental CPI for the elderly are available. In column (iii) the sample extends from 1953 to 2009. Standard errors in parentheses are two-way clustered by time and cohort.

	(i)	(ii)	(iii)
Gain parameter θ	2.561	3.475	3.720
	(0.275)	(0.591)	(0.414)
Sensitivity β	0.408	0.432	0.599
	(0.085)	(0.092)	(0.089)
$\frac{\text{Age}}{100}$		-0.004	-0.009
		(0.003)	(0.002)
$\frac{\text{Age}}{100} \times (\pi_{t-1}^{Elderly} - \pi_{t-1})$		-0.491	
		(0.938)	
$\frac{\text{Age}}{100} \times (\pi_{t-1}^{Gas} - \pi_{t-1})$			-0.011
			(0.010)
Time dummies	Yes	Yes	Yes
Adj. R^2	0.245	0.246	0.639
#Obs.	5200	5200	8215

different from zero, β is largely unaffected, and the estimate of θ is only slightly higher than in Table 1.

E   Survey of Consumer Finances

We use the data set of Malmendier and Nagel (2011), which comprises both the modern triennial SCF starting in 1983 and earlier versions of the SCF extending back to 1960 if they contain data on asset holdings. The data includes the years 1960, 1962, 1963, 1964, 1967, 1968, 1969, 1970, 1971, 1977, 1983, 1986, 1989, 1992, 1995, 1998, 2001, 2004, and 2007. We restrict attention to households whose head is more than 24 years and less than 75 years old. For information about this data set in general, see Malmendier and Nagel (2011). Here we focus on the definitions of the variables that are specific to the analysis in this paper.

To measure households' fixed-rate liabilities, we consider all mortgages secured by a household's primary residence.[14] Outstanding mortgage balances on first and second mortgages (and, from 1983 onwards, additional loans secured by the primary residence) are available in all survey waves except 1977. In that year, we observe, separately for first and second mortgages, only the original loan amount, the annual percentage rate, the number of years left on the mortgage, and, only for the second mortgage, the year the loan was obtained. For first mortgages, we impute the year it was obtained by assuming a 30-year maturity, which was standard at the time (Green and Wachter (2005)). Then, for first and second mortgages, we calculate the remaining mortgage balance outstanding in 1977 by assuming an amortization schedule with fixed payments over the life of the mortgage.

On the asset side, holdings of long-term bonds include those through mutual funds and defined contribution retirement accounts. Holdings of financial assets include stocks, bonds, and cash-like investments such as deposit accounts. In 1964, holdings of long-term bonds are not reported separately, and hence our tests with long-term bond holdings data discard this survey wave.

Since 1983, the sample of the SCF is designed to oversample high-income households. The oversampling is helpful for asset allocation analyses since it provides a substantial number of observations on households with significant stock holdings, but it could also induce selection bias. When we aggregate asset holdings and borrowing data within cohorts, we apply the weights provided by the SCF that undo the overweighting of high-income households and that also adjust for non-response bias. The weighted estimates are representative of the U.S. population.

Finally, all income, wealth, and asset holdings variables are deflated into September 2007 dollars using the consumer price index (CPI-U until 1997 and CPI-U-RS thereafter).

---

[14]As explained in Section 4, we include only primary residences to ensure comparability over time.

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.

Branch, W. A. and G. W. Evans (2006). A Simple Recursive Forecasting Model. *Economics Letters 91*, 158–166.

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

Bryan, M. and G. Venkatu (2001, November). The Curiously Different Inflation Perspectives of Men and Women. *Federal Reserve Bank of Cleveland Economic Commentary*, 1–4.

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.

DeGroot, M. H. (1974). Reaching a Consensus. *Journal of the American Statistical Association 69*, 118–121.

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.

Fuster, A., B. Hebert, and D. Laibson (2011). Natural Expectations, Macroeconomic Dynamics and Asset Pricing. In D. Acemoglu and M. Woodford (Eds.), *NBER Macroeconomics Annual 2011*. Chicago, IL: University of Chicago Press.

Fuster, A., D. Laibson, and B. Mendel (2010). Natural Expectations and Macroeconomic Fluctuations. *Journal of Economic Perspectives 24*, 67–84.

Green, R. K. and S. M. Wachter (2005). The American Mortgage in Historical and International Context. *The Journal of Economic Perspectives 19*(4), 93–114.

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.

Malmendier, U. and G. Tate (2005). CEO Overconfidence and Corporate Investment. *Journal of Finance 60*(6), 2661–2700.

Malmendier, U., G. Tate, and J. Yan (2011). Overconfidence and Early-life Experiences: The Effect of Managerial Traits on Corporate Financial Policies. *Journal of Finance 66*(5), 1687–1733.

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. P. Nicolini (2003). Recurrent Hyperinflations and Learning. *American Economic Review 93*, 1476–1498.

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.

Mishkin, F. S. (2007). Inflation Dynamics. *International Finance 10*(3), 317–334.

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

Orphanides, A. and J. C. Williams (2005b). 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.

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

Piazzesi, M. and M. Schneider (2012). Inflation and the Price of Real Assets. 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.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton, NJ: Princeton University 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*. Princeton, NJ: Princeton University Press.

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

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

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.