Online Appendix
for
"Learning from Inflation Experiences"

Ulrike Malmendier*
UC Berkeley, NBER and CEPR

Stefan Nagel†
University of Michigan, NBER and CEPR

September 8, 2014

A Implied weights on 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

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

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

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

Thus, iterating further on (OA.1) and (OA.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 $k \in \{0, 1, ..., t-s\}$,

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

†Ross School of Business and Department of Economics, University of Michigan, 701 Tappan St., Ann Arbor, MI 48019, e-mail: stenagel@umich.edu

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

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

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

For comparison, the weighting function in Malmendier and Nagel (2011) assigns observations at time $t-k$ (with $0 \leq k \leq t-s$) the weight[2]

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

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:

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

We present a proof by induction. First, the choice of $\gamma_{t,s}$ in (OA.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

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

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

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

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

[2] 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.

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

Finally, as last step, we show that the gain specification in (OA.5) is approximately equivalent to the gain specification (4) that we use in the empirical analysis. Focusing first on the denominator of (OA.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 of infinitesimal increments, we get

$$\gamma_{t,s} = \frac{\lambda+1}{t-s}. \qquad \text{(OA.8)}$$

In our case with quarterly increments, this approximation is virtually identical to the true gain sequence in (OA.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 (OA.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

We adjust the raw Reuters/Michigan Survey of Consumers (MSC) data for several known deficiencies following Curtin (1996), 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).

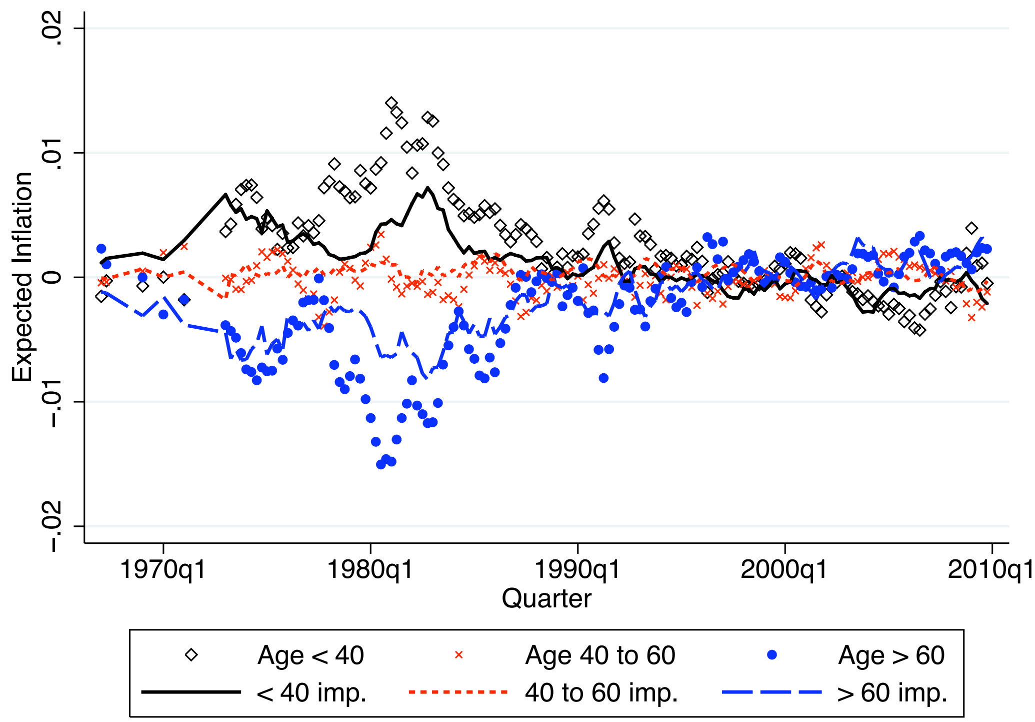

Figure OA.1: Four-quarter moving averages of actual and imputed mean one-year inflation expectations of young individuals (below 40), mid-aged individuals (between 40 and 60), and old individuals (above 60), shown as deviations from the cross-sectional mean expectation.

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

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

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 OA.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 Different starting points for experience accumulation

In the analysis in the main paper, we allow all realizations since birth to affect an individual's experience. One may wonder whether individuals truly use the whole time series since their birth, or rather starting from a later age when they begin to pay attention. Vice versa, one may wonder whether experience effects start from a point in time even before birth, reflecting the experience of the parents transmitted to the children.

In this Section of the Appendix, we illustrate that our results are not sensitive to the exact choice of starting data and discuss why this is the case.

Table OA.1 repeats the baseline estimation from Table 1, column (i), of the main paper, but with different starting points for the set of experienced inflation rates. In column (i), experiences are assumed to accumulate from 10 years before birth, in column (ii) from 10 years after birth.

This change has little effect in terms of explanatory power. As in our main results in Table 1 of the main paper, the resulting learning-from-experience forecast is strongly positively related to inflation experiences. The fit (adj. R^2) is slightly better than in the baseline regression in Table 1, but not much.

When starting pre-birth (column (i)), the gain parameter θ is higher, which implies a greater degree of downweighting of early experiences. Evidently, the estimation adapts to the pre-birth starting point by a greater downweighting of data in the early part of the experienced data set, which effectively gets back to implied weights that look quite similar to those in the baseline estimation, with relatively little weight on early life experiences and very small weights on pre-birth data.

When starting post-birth (column (ii)), the estimate of θ is lower, implying less downweighting of experiences in the early part of the experience set. Since experiences accumulate only from age 10 onwards, there is automatically a higher emphasis on recent data compared with the baseline estimation in Table 1 of the main paper, which results in less need of downweighting through a high θ. As a result the estimation generates implied weights that look relatively similar to the baseline estimation, with low weights on early life experiences and higher weights for recent experiences.

Table OA.1: Learning-from-experience model: Varying the starting point for experienced inflation

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 from 10 years before (column (i)) or 10 years after birth (column (ii)) 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).

	Start -10yrs (i)	Start +10yrs (ii)
Gain parameter θ	3.363	2.137
	(0.293)	(0.174)
Sensitivity β	0.892	0.500
	(0.073)	(0.060)
Time dummies	Yes	Yes
Imputed data included	Yes	Yes
Adj. R^2	0.638	0.637
RMSE	0.0148	0.0148
#Obs.	8215	8215

Why are the results rather insensitive to the choice of starting point for experience accumulation? The intuition is that, by treating θ as a parameter to be estimated, our specification has sufficient flexibility to adapt to different starting points without much difference in the fit of the model.

E Controlling for age-specific inflation rates

We re-run the regressions from Table 1 in the main paper 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 OA.2 presents the results. The inflation series based on the CPI-E is only available

Table OA.2: 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

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

F Reduced-form evidence on age heterogeneity in the responses to inflation surprises

One possible implication of learning from experience is that younger people revise their expectations more strongly in response to inflation surprises, which can be tested by regressing each cohort's expectations revisions on the inflation surprise they experienced and its interaction with age:

$$\tilde{\pi}_{t+1|t,s} - \tilde{\pi}_{t|t-1,s} = (\alpha_0 + \alpha_1(\tfrac{Age}{100}))(\pi_t - \tilde{\pi}_{t|t-1,s}) + \delta' D_t + \epsilon_{t,s}. \tag{OA.9}$$

The dependent variable is the change in mean inflation expectations of cohort s from period $t-1$ to t. Inflation surprise is calculated as the difference between the realized inflation rate from the end of period $t-1$ to the end of period t, denoted as π_t, minus the expectation of cohort s in period $t-1$, denoted as $\tilde{\pi}_{t|t-1,s}$. D_t includes age and time dummies.[3] Interacting the surprise term with age allows to capture age heterogeneity in the reaction to the inflation surprise. If younger individuals react more strongly to inflation surprises, consistent with the learning-from-experience hypothesis, the interaction coefficient should be negative.

There are two caveats with such an analysis. First, in the case of AR(1) processes, a stronger response of younger cohorts to inflation surprises is not necessarily implied by experience-based learning because of the effect of experience on the autocorrelation coefficient.

[3] As in our main analysis, we include time dummies to control for factors other than past inflation experience, including any omitted macroeconomic variables or other unobserved effects common to all individuals.

If the perceived law of motion were an IID process, and individuals used simply experienced inflation rates to estimate the mean, the above intuition would hold: A positive inflation surprise would always lead to an upward revision of expectations; a negative inflation surprise would always lead to a downward revision; and in both cases the effect would weaken with age. In case of an AR(1) perceived law of motion, however, an inflation surprise also leads to a revision in the estimate of the autocorrelation coefficient, and the resulting revision of expectations need not be in the same direction as the inflation surprise. The only way to properly account for both channels is to estimate the learning rule under an assumed perceived law of motion, as we do in the main analysis of the paper.

A second caveat with the reduced-form approach in equation (OA.9) is that the estimates may be biased due to measurement error. As with any survey data, inflation expectation $\tilde{\pi}$ may be mismeasured. In equation (OA.9), inflation expectations are included in the construction of both the explanatory and the dependent variables, which could induce spurious correlation between the variables on the right-hand and the left-hand side. A significant coefficient estimate may thus be (partly) due to spurious correlation.[4] The interaction term between inflation surprise and age adds an additional layer of complications because age could be correlated with the magnitude of the measurement error, e.g., if younger people tend to give more noisy responses to questions of macroeconomics outlook.

Taking into account these two caveats, it is nevertheless useful to check to what extent we can describe the data with the simple updating model in equation (OA.9).

We apply equation (OA.9) to quarterly inflation data and survey data that provides individuals' one-year-ahead (i.e., four-quarter) inflation forecasts. As in the main analysis in the paper, we assume that individuals surveyed at various points during quarter t have access to quarterly inflation rates until the end of quarter $t-1$.[5] 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 OA.3 reports the results of the estimation. We find that the coefficient on the interaction of the surprise with age has the hypothesized negative sign and is highly statistically significant. In terms of magnitude, an age difference of 50 years implies a forecast revision that is about 0.13. For comparison, with our estimates from the main analysis of $\theta = 3$, we get a difference in gains between a 25- and a 75-year old of 0.08.

As discussed above, these estimates are not directly comparable, though, because our main analysis takes into account that individuals also perceive some persistence in inflation rates. In other words, the magnitude of the coefficient on the surprise variable alone picks up some of the effects induced by autocorrelation of inflation and individuals' perception of it. To illustrate this confound, consider the extreme case where individuals are convinced that $\phi \approx 1$, and their forecast during quarter t is the inflation rate observed from $t-5$ to $t-1$.

[4] This is not an issue in the main analysis of the paper. In the specification that we use there, measurement error affects only the dependent variable of the estimating equation. The explanatory variable depends on past realized inflation, not on survey data. Hence, any measurement error would end up in the residual.

[5] To economize on notation, we do not explicitly highlight the quarterly equivalent of the inflation surprise variable $\pi_t - \tilde{\pi}_{t|t-1,s}$, which is $\pi_t - \tilde{\pi}_{t|t-4,s}$, i.e., the realized inflation rate from $t-4$ to t minus the expectation of cohort s during quarter $t-4$ about the inflation rate from $t-4$ to t.

Table OA.3: Age-Heterogeneity in reaction to inflation surprise

Dependent variable is the change in the one-year mean inflation expectation of cohort s from quarter $t-4$ to t. Explanatory variables are the inflation surprise, measured as the realized inflation rate from $t-5$ to $t-1$ minus the expectation of cohort s during quarter $t-4$, and its interaction with age (divided by 100 for sake of readability). Standard errors in parentheses are two-way clustered by time and cohort. The number of observations is smaller than in Table 1 in the main analysis. The reason is that we can only use observations for which a four-quarter lagged observation is available to calculate the forecast revisions and inflation surprise terms.

	(i)	(ii)
$\frac{\text{Age}}{100} \times$ Surprise	-0.270	-0.264
	(0.083)	(0.088)
Surprise	1.108	1.106
	(0.048)	(0.051)
Age dummies	Yes	Yes
Time dummies	Yes	Yes
Imputed data included	Yes	No
Adj. R^2	0.623	0.623
#Obs.	7475	7154

Now add a highly persistent but possibly tiny between-cohort heterogeneity in forecasts, so that each cohort's forecast is the observed inflation rate plus a small $a_{t,s}$. In this case, the regression in Table OA.3 would yield a coefficient of 1 on the surprise term and zero on the interaction term.

Thus, while the interpretation of the coefficient estimates in Table OA.3 is not as clear as the interpretation of the fully estimated model in the main analysis, the empirical results confirm the basic intuition that learning from experience implies that the strength of forecast revisions in response to inflation surprises declines with age.

G Survey of Consumer Finances

We use the data from 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 at least 25 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.[6] 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 1977, 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. We impute the outstanding mortgage balances as follows: 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.

For our main analysis, we average assets and liabilities within cohort (and typically use the log value). Since 1983, the sample of the *SCF* is designed to oversample high-income households. Thus, 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).

H Inflation experiences and financial decisions: Household-level regressions

The analysis of households' financial decisions in the main paper (Table 3) focuses on cohort-level regressions, similar to our analysis of experience-based inflation expectations (Table 1). In this section, we re-do the analysis of households' financial decisions using household- rather than cohort-level data. The household-level analysis allows us to include additional household-level controls that, in prior literature, have been found to affect borrowing and mortgage choices, such as education, ethnicity, gender, family status, and financial variables.[7]

[6] As explained in the main paper, we include only primary residences to ensure comparability over time.

[7] For example, Cox and Jappelli (1993) find, in the 1983 SCF wave, that income, net worth, age, and urban status help predict the level of household debt (including its largest component, fixed-rate mortgage debt). In the household-level analysis above, we include all significant variables from Cox and Jappelli (1993) other than urban status, which is not available in the public-use version of the SCF. Relatedly, Coulibaly and Li (2009) analyze the choice between fixed-rate and variable-rate mortgages in the 1995-2004 SCF waves. They find that, in addition to the variables mentioned above, self-assessed risk aversion, income volatility, and mobility predict mortgage choice, consistent with Campbell and Cocco (2003). Risk aversion is available only in 1983 and 1989-2007, and we re-analyze this subsample with the additional risk aversion controls. Income volatility variable is calculated by occupation category, and we account for this with occupation dummies. (To merge the 1983 with the 1989 to 2007 occupation codes, we use the following mapping from 1983 codes (variable B3114) to 1989 to 2007 code (variable X7401): 2 and 3 (1983) to 1 (1989 to 2007); 4 and 1 to 2; 8 and 5 to 3;

Note that these additional controls are unlikely to affect our results as they would have to be correlated not only with age-related differences in inflation experiences but also with the evolution of these differences over time. Nevertheless, we replicate Table 3 with household-level regressions and household-level controls.

In the baseline specification in Table OA.4, we include time dummies, age group dummies (below 41, 41 to 50, 51 to 60, above 60), dummies for completed high school education, completed college education, race, gender, status as retired, being married, furthermore the number of children in the household, log income, log net worth, and a dummy for stock market participation (including in defined-contribution retirement accounts). Table OA.5 reports the regression for a shorter sample with additional risk-related controls (see footnote 7). Standard errors in these regressions are adjusted for multiple imputation of missing values as described in Malmendier and Nagel (2011). Further details on the data are also available in Malmendier and Nagel (2011).

Focusing first on Table OA.4, columns (i) to (iv) show Tobit regressions of the log of one plus the loan or investment amount on the learning-from-experience forecast and the controls. The results are exactly in line with the estimates we obtained from cohort-level data in Table 3 of the main paper. As in Table 3, we obtain a positive and statistically significant effect on fixed-rate mortgage amounts (column (i)) and a negative, but statistically insignificant coefficient for investment in long-term bonds (column (ii)). In the subsample of new loans, we also estimate a significantly positive effect of learning-from-experience inflation forecast on fixed-rate amounts (column (iii)), and a negative coefficient for variable-rate loans (column (iv)). Differently from the cohort-level regressions, the negative effect on the amount of new variable-rate loans is highly significant.

The household-level analysis also allows us to relate the learning-from-experience inflation forecast directly to the binary choice between variable- and fixed-rate loan. Column (v) reports the results from a Probit regression, using the same subsample of new loans. Consistent with our hypothesis, the learning-from-experience forecast is negatively related to the probability of getting a variable-rate (instead of a fixed-rate) loan.

The effects implied by the estimates are of the same order of magnitude as those in Table 3 in the same paper. For example, in column (i), a change from the 10th to 90th percentile in the learning-from-experience forecast—a difference of 3.2 percentage points—raises the log dollar amount by 2.25. To compare this to the coefficient estimate of 35.27 in column (i) of Table 3, which represents the effect of a unit change in the learning-from-experience forecast, we need to multiply by 31, resulting in an effect size of 69.8.

Using the shorter sample with the additional risk-related controls in Table OA.5 does not yield a substantial change compared with Table OA.4. Overall, the results that we obtained in the main paper from cohort-level data are robust to controlling for various demographic characteristics of households that are known to affect mortgage and investment choices.

---

6 to 4 and 5; 7 to 6.) The mobility question was asked only starting in 1995, and hence we cannot use it as a control variable.

Table OA.4: Inflation experiences and household nominal positions: Household-level regressions

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 sample is described in Malmendier and Nagel (2011). Each cohort is assumed to recursively estimate an AR(1) model of inflation, with $\theta = 3.044$, as in Table 1, column (i), of the main paper. In columns (i) and (ii), we use the resulting learning-from-experience forecast of inflation to explain log (1 + fixed-rate mortgage borrowing) and log(1+ long-term bond holdings) in Tobit regressions. Log mortgage borrowing in columns (iii) and (iv) comprises only new loans taken out or refinanced in the year in which the survey was carried out, and the sample is restricted to households with new loans. Column (v) reports a probit regression for the type of new loan taken out. Standard errors reported in parentheses are clustered by time period and adjusted for multiple imputation as described in Malmendier and Nagel (2011).

	Tobit (Log of 1+ Dollar Amount)				Probit (Indicator)
	Fixed-rate mortgages (i)	Long-term bonds (ii)	Fixed-rate new (iii)	Variable-rate new (iv)	New loan variable-rate? (v)
Coefficient estimates					
Learning-from-experience forecast	148.89	-3.61	116.15	-574.55	-34.54
	(33.70)	(22.35)	(53.77)	(191.22)	(11.82)
Log net worth	1.71	1.01	-0.17	1.04	0.06
	(0.08)	(0.10)	(0.12)	0.60	0.04
Log income	1.30	1.51	-0.08	2.68	0.15
	(0.18)	(0.14)	(0.29)	(1.67)	0.10
Magnitude of effects					
Average of fitted value at 90th pctile. minus fitted value at 10th pctile. of learning-from-exp forecast	$\Delta E(y\|x)$				Δ Prob.
	2.25	-0.05	2.52	-2.59	-0.19
	(0.70)	(0.30)	(1.25)	(0.48)	(0.02)
Demographics controls	Yes	Yes	Yes	Yes	Yes
Time dummies	Yes	Yes	Yes	Yes	Yes
Age group dummies	Yes	Yes	Yes	Yes	Yes
Sample	Full	Full	$\geq$ 1983 New loans only	$\geq$ 1983 New loans only	$\geq$ 1983 New loans only
Pseudo R^2	0.09	0.11	0.01	0.03	0.06
#Obs.	50,908	49,676	1,669	1,669	1,669

Table OA.5: Inflation experiences and household nominal positions: Household-level regressions with additional risk-related controls

The regressions in this table are the same as in Table OA.4, but with an additional set of risk-related controls: A self-assessed risk aversion measure and occupation dummies. The sample includes fewer survey waves (1983; 1989-2007), due to limited availability of the additional risk-related controls. Standard errors reported in parentheses are clustered by time period and adjusted for multiple imputation as described in Malmendier and Nagel (2011).

	Tobit (Log of 1+ Dollar Amount)				Probit (Indicator)
	Fixed-rate mortgages (i)	Long-term bonds (ii)	Fixed-rate new (iii)	Variable-rate new (iv)	New loan variable-rate? (v)
Coefficient estimates					
Learning-from-experience forecast	131.86	24.98	151.71	-743.65	-45.28
	(33.76)	(23.04)	(50.33)	(170.31)	(10.84)
Log net worth	1.76	0.98	-0.23	1.34	0.08
	(0.12)	(0.04)	(0.12)	0.64	0.04
Log income	0.89	1.21	-0.17	3.20	0.18
	(0.24)	(0.18)	(0.26)	(1.61)	0.09
Magnitude of effects					
Average of fitted value at 90th pctile. minus fitted value at 10th pctile. of learning-from-exp forecast	$\Delta E(y\|x)$				Δ Prob.
	1.24	0.27	2.49	-2.34	-0.17
	(0.44)	(0.26)	(0.89)	(0.56)	(0.02)
Demographics controls	Yes	Yes	Yes	Yes	Yes
Time dummies	Yes	Yes	Yes	Yes	Yes
Age group dummies	Yes	Yes	Yes	Yes	Yes
Additional risk-related controls	Yes	Yes	Yes	Yes	Yes
Sample	1983; $\geq$ 1989	1983; $\geq$ 1989	1983; $\geq$ 1989 New loans only	1983; $\geq$ 1989 New loans only	1983; $\geq$ 1989 New loans only
Pseudo R^2	0.08	0.14	0.01	0.03	0.07
#Obs.	28,571	28,571	1,526	1,526	1,526

I Out-of-sample predictions

Another advantage of 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.

In representative-agent macro models with adaptive learning, the identification of structural parameters is difficult, and the problems are magnified if the parameters of the learning rule are unknown and need to be estimated (Chevillon, Massmann, and Mavroeidis (2010)). 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 in the main paper, we have shown that the estimates of the gain parameter θ are most precise when we focus purely on cross-sectional variation by employing time dummies (Table 1 in the main paper). 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 OA.2 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 OA.2 shows, the learning-from-experience rule performs slightly better than the constant-gain rule.[8] 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.

[8] Note that the gap between the learning-from-experience model and the constant-gain model varies over time. For example, it shrinks between 2000 and 2005, indicating better performance of the constant-gain model during that period.

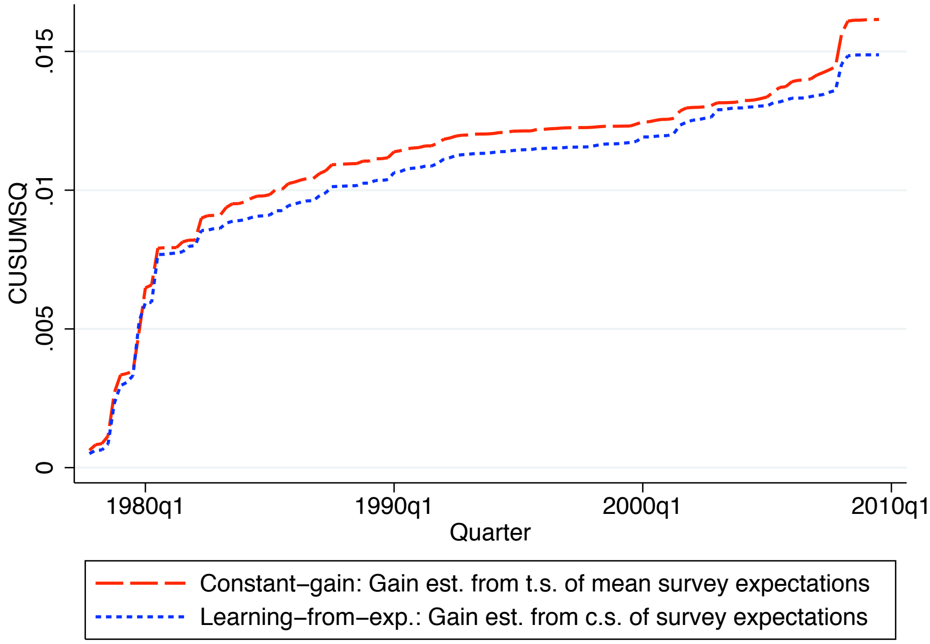

Figure OA.2: Pseudo-out-of-sample cumulative sum of squared errors in predicting mean inflation expectation. For the constant-gain model, the gain parameter γ is estimated over expanding windows, where the first extends from 1953Q4 to 1977Q4. For each window, mean expectations data until quarter $t-1$ 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 with expanding estimation windows.

References

Campbell, J. Y. and J. F. Cocco (2003). Household Risk Management and Optimal Mortgage Choice. *Quarterly Journal of Economics 118*(4), 1449–1494.

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

Coulibaly, B. and G. Li (2009). Choice of Mortgage Contracts: Evidence from the Survey of Consumer Finances. *Real Estate Economics 37*(4), 659–673.

Cox, D. and T. Jappelli (1993). The Effect of Borrowing Constraints on Consumer Liabilities. *Journal of Money, Credit and Banking 25*(2), 197–213.

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.

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