Persistence of Memory? Waning Experiences, Anchored Expectations

Rui Yu

December 10, 2015

Adviser: Prof. Ricardo Reis [1]

Abstract

I reestimate a learning-from-experience model using survey reported inflation expectations and confirm that personal inflation experiences strongly predict inflation expectations. I use the model to study the effect of demographics and history on aggregate inflation expectations. Demographic shocks exert little effect on aggregate inflation expectations and population aging does not explain better anchoring of inflation expectations. On the other hand, high inflation in the 1970s raised and unanchored inflation expectations, an effect that surprisingly increased over time. The downweighting of past inflation experiences accounts for three-fifths of the better anchoring of inflation expectations during the Great Moderation.

[1] A sincere thank you to Prof. Ricardo Reis for advising this project. I thank Prof. Lena Edlund, my discussant, Evan Munro, and fellow writers in w4999 seminar. I also thank Keshav Dogra, Prof. Susan Elmes, Shaowen Luo, Prof. Serena Ng, and Prof. Daniel Rabinowitz for help in the process.

1 Introduction

Malthus (1794) argued that population is critical for economic outcomes. But little is still known about the role of demographics, or traits of populations, on macroeconomic phenomena. Inflation and inflation expectations are examples of such critical macroeconomic variables (Friedman 1968, Woodford 2003). Recent research links personal experiences to macroeconomic expectations, producing expectations that differ across demographic groups. A natural but unexplored question is whether demographics, defined hereafter as the population age distribution, affect inflation expectations.

Do demographic shocks affect aggregate inflation expectations? In a similar vein, does demographic aging in the U.S. explain better anchoring of inflation expectations during the Great Moderation? I replicate a learning-from-experience model (Malmendier and Nagel 2013) using survey reported inflation expectations from the Michigan Surveys of Consumers. Using the model, I answer these questions by simulating the effect of demographic shocks and aging on aggregate inflation expectations. In addition, I exploit the historical component of learning-from-experience. What is the legacy of high inflation in the 1970s on current expectations? Does downweighting past histories explain stable inflation expectations? I run counterfactual simulations to directly address these questions.

A growing literature has sought to explain the moderation of macroeconomic volatility in the U.S., with little convergence (Stock and Watson 2003, Campbell and Hercowitz 2005, Ramey and Vine 2006, Arias, Hansen, and Ohanian 2007). One feature of the so-called Great Moderation is the unprecedented stability, or well-anchoredness, of inflation expectations. Inflation expectations are well-anchored if long-run expectations are invariant to short-run fluctuations in inflation. The roots of this phenomenon are still unclear due to three obstacles to research. One, hypotheses need to match the specific timing of better anchored expectations - i.e., why now rather than at another time? Two, a model of expectations formation that is compatible with and can test the hypothesis might not exist. Three, data on inflation expectations are sparse, and even when available it is unclear how to evaluate the anchoredness of expectations.

< *Figure* 1 >

Jaimovich and Siu (2009) find, by exploiting variation across the G7 nations, that demographic makeup of the workforce strongly explains business cycle volatility. They extend their results and find that demographic aging in the U.S. explains a substantial portion of the Great Moderation. Could demographic change also affect inflation expectations? Figure 1 plots (1) perceived autocorrelation of inflation, a measure of anchoredness of expectations, and (2) the average age of surveyees in the Michigan Surveys of Consumers. Better anchoring of expectations seems to coincide with aging of the surveyees. Jaimovich and Siu's hypothesis is appealing, because the timing of aging matches up well with the start of the Great Moderation.

Malmendier and Nagel (2013) build a learning-from-experience model that predicts individual inflation expectations based on lifetime inflation history. The model proposes that generations with higher inflation experiences form higher inflation expectations and older populations tend to be less responsive to new experiences, relative to the young. The model directly relates micro-level demographic variation to aggregate inflation expectations. Alternatively, learning-from-experience also posits that history affects inflation expectations. Thus, learning-from-experience model can test two breeds of hypotheses, demographics versus history, to explain the trend of aggregate inflation expectations. Counterfactual simulations on the model can then isolate the incremental contribution of demographics or history on aggregate expectations

I estimate the model using micro-level survey responses about inflation expectations from the Michigan Surveys of Consumers. The survey data is a rotating panel covering 1966-2013, and it captures a sufficient amount of demographic information for analysis. Learning-from-experience exploits the variation in inflation expectations across cohorts and time in these surveys for estimation, setting it apart from models limited to aggregate time-series. In addition, the AR(1) structure of learning-from-experience produces an intuitive measurement of the anchoredness of both individual and aggregate inflation expectations.

Therefore my analysis overcomes the three obstacles to studying the roots of the moderation in inflation expectations. I re-estimate a learning-from-experience model using survey responses from the Michigan Surveys of Consumers and obtain almost identical estimates to Malmendier and Nagel's. I study the effect of demographic shocks on inflation expectations by running simulations, but find that demographic shocks have little effect on expectations. In counterfactual simulations with no population aging, I find that demographic change does not explain better anchoring of inflation expectations. On the other hand, high inflation in the 1970s had long-lasting effects on aggregate expectations. It is the downweighting of past inflation history, in turn, that explains a substantial amount of the better anchoring of inflation expectations.

Section 2 is a literature review, section 3 introduces the learning-from-experience model, section 4 describes the data, section 5 reports estimation results followed by a brief discussion of robustness checks, section 6 runs four counterfactual simulations, section 7 is a general discussion of results, and section 8 concludes.

2 Literature Review

Macroeconomic expectations, particularly inflation expectations, play a critical role in macroeconomics and policy. The perception of inflation by firms and agents affects prices, influencing actual inflation as embodied in the Phillips Curve (Romer 2012). Inflation expectations may affect the financial decisions of individuals, as expected inflation affects perceived interest rates and in turn influences financial decisions such as whether or not to buy a house. Inflation expectations thus affect individual and macroeconomic outcomes (Woodford 2003). Naturally, inflation expectations figure prominently to central banks seeking to maintain price and output stability (Bernanke 2007).

But inflation expectations stubbornly differ across demographic groups, generations, and time. Empirical evidence from surveys on expectations documents persistent heterogeneity, which does not sit well with conventional models with agents having the same expectations. Bryan and Venkatu (2001) document heterogeneity in inflation expectations along gender. De Bruine et al. (2010) expand the analysis and find persistently higher inflation expectations by individuals who are female, poorer, single, and less educated. Piazzesi, Salomo, and Schneider (2013) find significant divergence in inflation expectations between young and old households in the 1970s and show that the heterogeneity accounts for differences in financial decision making. Studies on heterogeneous expectations link differences to demographic and individual traits.

However, little agreement exists on how agents form inflation expectations and why expectations differ. Carroll (2003) suggests that macroeconomic news diffuse through the economy slowly through an epidemiological model. Thus, professional forecasters update their expectations rapidly while others do not, leading to disagreement about inflation expectations. Mankiw and Reis (2002) propose a "sticky information" model where agents rationally choose to update information intermittently. The model accounts for observed heterogeneity in inflation expectations (Mankiw, Reis, and Wolfers 2006).

Malmendier and Nagel (2013) propose a learning from experience model where a weighted average of lifetime inflation experiences predict inflation expectations. The weights have two features. First, individuals weigh recent and past experiences differently. Experiences during formative years or teachings by parents may shape adult decision, a process represented by higher weights on past experiences. On the other hand, it may be possible that memory fades over time, so recent experiences hold more weight. The model allows the data to speak; data shapes the behavior of memory. Second, older individuals weigh recent experiences lower than younger individuals. Using survey reported inflation expectations from 57 years of the Michigan Surveys of Consumers, they find that differences in lifetime inflation experiences strongly predict inflation expectations. However, their model does not disentangle cohort effects from cohort-specific inflation experiences. One advantage of learning from experience is that it bridges observed heterogeneity and demographic factors by hypothesizing a channel through which demographics affect expectations.

Learning from experience follows a small but growing literature incorporating personal experiences to economic decision making. The psychology literature separates information from experiences and information from summaries, arguing that personal experiences, particularly recent ones, have a greater effect than information extracted from summaries (Hertwig et al. 2004). Choi et. al (2009) find that individuals earning high returns on his/her 401(k) are associated with higher 401(k) saving rates. Greenwood and Nagel (2009) show that younger mutual fund managers tend to invest more in tech stocks relative to older managers, suggesting that experience shapes risk aversion. Studies find that corporate leaders who experienced the Great Depression tend to shy away from external financing or leverage (Malmendier and Tate 2005, Graham and Narasimhan 2004). Most closely related is Malmendier and Nagel (2011) who find that individuals surveyed by the Survey of Consumer Finances who experienced low stock market returns in their lifetimes report higher risk aversion, are less likely to own stocks, and are more pessimistic about future stock returns.

Learning from experience draws on adaptive learning models popular in macroeconomics. Adaptive learning posits that individuals' expectations can be approximated by a simple forecasting rule that uses available historical data (Bray 1982, Sargent 1993). Malmendier and Nagel posit that individuals recursively estimate inflation using an AR(1) model and data from lifetime inflation history. Each recursion includes the most recent inflation experience, but a gain parameter controls and weights the relative impact of new information on the AR(1) model. A simple adjustment to the gain parameter, detailed later in this paper, yields a recursive formulation of the ordinary least squares (Evans and Honkapohja 2001). Their specification is close to Honkapohja and Mitra (2003) where memory of past data gets downweighted over time. The advantage of learning from experience that depend on individual history is it yields heterogeneous inflation expectations that do not necessarily converge, matching observations about inflation expectations.

I focus on an unexplored demographic component of the learning-from-experience model. Since experiences are weighted according to age, their model suggests a link between demographics and inflation expectations. This perspective taps into a thin but timely vein of research on demographics and macroeconomic performance. Jaimovich and Siu (2009) use demographic trends and business cycle volatility from the G7 nations and find that the age composition of the workforce strongly accounts for business cycle volatility. They also find that one-fifth to one-third of recent moderation in the U.S. business cycle can be explained by demographic factors. Interestingly, they suggest that greater numbers of old people increase volatility, cutting slightly against Malmendier and Nagel's model where older individuals tend to stabilize macroeconomic expectations. Shimer (1998) attributes what were declining unemployment rates in the U.S. to demographic aging, based on the observation that unemployment is highest among the young. Given the aging population in the U.S., exploring the effect of demographic changes on inflation expectations opens a new path through which demographics may exert real macroeconomic effects.

In particular, I explore one facet of the widely reported macroeconomic stability before the financial crisis in the U.S., known as the Great Moderation (Stock and Watson 2003). One feature of the Great Moderation are well-anchored inflation expectations after decades of volatile expectations (Davis 2012). A common refrain among central banks is that sophisticated monetary policy has improved central bank credibility, leading to inflation expectations invariant to short-run inflation surprises. However, Learning from Experience suggests that older individuals have a lower gain parameter, or they react less to recent inflation surprises. Is it possible that well-anchored inflation expectations are due to demographic aging?

3 Learning from Experience

Learning-from-experience posits that individuals' lifetime inflation experiences predict inflation expectations. Formally, the model is a recursive formulation of weighted least squares that estimates an AR(1) model to predict inflation. A simple intuition for learning-from-experience is that when an individual experiences inflation, she adds that new experience to a personal dataset of past, $lifetime$, inflation rates. That personal dataset is weighted, so that past and recent inflation experiences differ in importance. The individual then fits a line to the personal dataset, and uses the line to forecast expected inflation in the next period. For estimation, time period is a quarter.

Formally, Malmendier and Nagel (2013) propose that an individual born at time s forecasts inflation for time $t+1$ as if with an AR(1) model:

$$\pi^e_{t+1,s} = \alpha_{t,s} + \phi_{t,s}\pi_t + \epsilon_{t+1} \tag{1}$$

Set $\mathbf{b}_{t,s}$=($\alpha_{t,s}$, $\phi_{t,s}$)' and $\mathbf{x}_t = (1, \pi_t)'$. The coefficients are estimated recursively:

$$\mathbf{b}_{t,s} = \mathbf{b}_{t-1,s} + \gamma_{t,s}\mathbf{R}^{-1}_{t,s}\mathbf{x}_{t-1}(\pi_t - \mathbf{b}'_{t-1,s}\mathbf{x}_{t-1}) \tag{2}$$

where t is time and s is birth year of individual and

$$\mathbf{R}_{t,s} = \mathbf{R}_{t-1,s} + \gamma_{t,s}(\mathbf{x}_{t-1}\mathbf{x}'_{t-1} - \mathbf{R}_{t-1,s}) \tag{3}$$

$(\pi_t - \mathbf{b}'_{t-1,s}\mathbf{x}_{t-1})$ measures the degree and direction of the inflation surprise at period t. $\gamma_{t,s}$ controls how much an individual updates the forecast given an inflation surprise at time t. A low $\gamma_{t,s}$ implies nearly unchanged expectations while a high $\gamma_{t,s}$ implies expectations sensitive to inflation surprises.

Key to this model is that individuals ignore experiences before birth in their forecast. Since each cohort has a unique set of lifetime inflation experiences, this allows $\mathbf{b}_{t,s}$ to vary across time and cohort, producing heterogeneous inflation expectations in one period. Malmendier and Nagel show that this heterogeneity in expectations is due to either (1) differences in individuals' perceived long-run, mean inflation, $\mu_{t,s} = \frac{\alpha_{t,s}}{1-\phi_{t,s}}$ or (2) from perceived persistence $\phi_{t,s}$, or autocorrelation, of deviation of inflation from this mean. Malmendier and Nagel specify:

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

Where θ is a parameter that determines the shape and weights of past experiences. $\gamma_{t,s}$ is a gain parameter that decreases with age, reflecting how older individuals already possess a large set of experiences and thus practices less updating to recent experiences. Crucially, setting $\gamma_{t,s} = 1$ whenever $t - s \geq \theta$ implies that inflation experiences before and around birth are ignored. Evans and Honkapohja (2001) show that setting $\gamma_{t,s} = 1/t$ results in a recursive formulation of ordinary least squares. In the same vein, Malmendier and Nagel's specifications for the gain parameter results in a recursive formulation of weighted least squares. The authors begin recursion "at some point in the distant past" and argue that initial conditions get downweighted fast enough so that they are irrelevant. I make a slight adjustment during estimation, due to the singularity of the $\mathbf{R}_{t,s}$ matrix, which is detailed in Appendix B.

< *Figure* 2 >

Rather than assuming whether recent or past experiences hold more weight, the model allows the data to shape the behavior of memory. Consider Figure 2, which plots the implied weights on inflation experiences for a 50-year old individual. For θ=1, individuals weigh all historical data since birth equally. $\theta < 1$ implies experiences in the distant past receive more weight while $\theta > 1$ means recent experiences hold more weight.

The psychology literature (Hertwig et al. 2004) posits that individuals use information from personal experiences and or from summary information, which includes news or historical information. The authors allow such summary information to influence expectations by adding a common factor f_t into the learning-from-experience model:

$$\pi^e_{t+1|t,s} = \beta \mathbf{b}'_{t,s}\mathbf{x}_t + (1-\beta) f_t \tag{5}$$

Expected inflation is a weighted average of learning from experience and a common factor to the time period. The β represents the contribution of lifetime experiences to inflation expectations. The common factor captures any effect that is common to a time period by varies across time. Included in the factors are, for example, the opinions of professional forecasters or available historical information. For estimation, time fixed effects replace the common factor:

$$\pi^e_{t+1|t,s} = \beta \mathbf{b}'_{t,s} x_t + \mathbf{d}'\mathbf{D}_t + \epsilon_{t,s} \tag{6}$$

The authors also add disturbances $\epsilon_{t,s}$ which are assumed to be uncorrelated with $\tau_{t+1|t,s}$. The disturbances are allowed to be correlated within a cohort across time or between cohorts in one time period.

The authors estimate β and θ.

4 Data

Following Malmendier and Nagel, I estimate the learning-from-experience model by using 46 years of cross sectional data from the Michigan Surveys of Consumers and the Survey of Consumer Attitudes and Behavior. The Michigan Surveys of Consumers is a monthly survey of consumer expectations in operation since 1978. Each month, the survey center telephones and interviews approximately 500 adults in the contiguous United States. Interviewees are selected by random telephone sampling aided by a proprietary sampling mechanism that approximates a representative sample of households. An independent subset of interviewees, approximately 40 percent, are reinterviewed six months later, resulting in a rotating panel design.

The survey primarily asks interviewees about expectations of their own and broader economic conditions. In addition, a few questions ask about current personal economic conditions and demographic information of interviewees including age, race, assets, income, and education. The sum of weighted responses to five core questions, normalized to a base, is the Index of Consumer Sentiment widely reported in the media. The survey responses from 1978 are available from the Thomson Reuters/University of Michigan Survey Research Center.

Of interest are responses to the question:

"By about what percent do you expect prices to go (up/down) on average during the next 12 months? Respondents give a percent estimate of future inflation.

While the official survey began in 1978, an earlier version of the survey that asks about inflation expectations began in 1954 as the Survey of Consumer Attitudes and Behavior (SCAB). These surveys are available from the Inter-University Consortium on Social and Political Research (ICSRP). However, the earliest versions of this survey from 1954-1965 only elicit categorical responses about whether prices were expected to rise or fall, but do not have continuous estimates. Malmendier and Nagel calculate continuous responses from categorical responses for these early surveys, an exercise omitted in my analysis. Therefore, I only use survey responses from 1966-2013 that give continuous responses.

I use historical data on the quarterly Consumer Price Index (CPI) from Shiller (2005). Data extend from 1872-2013.

Detailed procedure for data collection and preparation is in Appendix A. Following Malmendier and Nagel, I subset the data to those that who were 24 to 75 years old at the time of the survey. In addition, for each survey period and cohort, I calculate the mean weighted inflation expectation. If there are more than one survey in a period, I take the mean inflation expectations of all surveys in a quarter for each cohort.

$< Figure\, 3 >$

Using the Michigan Surveys, Figure 3 plots four quarter moving averages of inflation expectations of different age groups, shown as deviations from the mean. The figure depicts three hallmark observations of inflation expectations. One, inflation expectations are heterogeneous, with almost 0.3 deviations in expectations between individuals under 40 and over 60 in the late 1970s. Two, the inflation shock of the 1970s affected young people's expectations immediately, while old people's expectations reacted slowly. A similar observation occurs in the near-deflationary episode in the 2000s. This supports a feature of the weights in the learning from experience model, where older people discount information from recent experiences. Finally, young people began with above average expectations that trended down until they became below average in the 2010s, and vice-versa for old respondents. This supports learning from experience; young individuals in the 2010s lived through relatively stable inflation history while old people today experienced high and volatile inflation in the 1970s.

$< Figure\, 4 >$

Figure 4 plots the evolution of the age distribution of respondents to the Michigan Surveys. Note that the survey only interviews those over 18, skewing the distribution, particularly in the 1960s. In addition, while Malmendier and Nagel's estimation only includes respondents who were 25 to 74 years old, I include all in Figure 4 to capture the full age distribution over time. The trend is striking and clear. It charts the post-war demographic evolution in the U.S. The plot in the 1960s excludes the then recent post-war baby boom, since most of the baby boomers were likely under 18 and ineligibile for the survey. But in all later decades the baby boomers are represented by a visible bump, which ages over time. Later, in the 2010s the plot appears left-skewed. Simple cross-sections confirm well-documented facts about demographic aging in the U.S.

$< Figure\, 5 >$

Figure 5 plots log differenced annualized inflation rates by quarter, from 1873 to 2013. I took, as do Malmendier and Nagel, four quarter moving average of inflation to smooth fluctuations. Of note is modern inflation history after World War II. After decades of low inflation, the oil shocks in the 1970s began a period of persistent and unstable high inflation. Only starting in the mid 1980s does inflation subside, begining the so-called Great Moderation of low and stable inflation rates. During the Great Recession, there was a possibility of deflation and the figure captures both rapid disinflation as well as increasingly unstable inflation.

5 Estimation Results

I estimate the learning from experience model using non linear least squares. Note that the Michigan Surveys of Consumers solicit expectations about inflation one year later, while the inflation series and the learning-from-experience forecast uses quarterly data. Therefore the authors employ multi-period forecasts by iterating on the AR(1) model at time t and producing forecasts for $t+1$, $t+2$, $t+3$, and $t+4$. The final learning-from-experience forecast is the mean of the four forecasts.

$< Table\, 1 >$

Table 1 reports results of the estimation, withholding the time fixed effects. It is worth emphasizing that Malmendier and Nagel do not indicate the significance levels of their estimates, so I also omit them from my table. Column (1) reports Malmendier and Nagel's results. Their estimate for θ, 3.144 (cl. s.e. 0.257), suggests that recent experiences receive higher weight than past experiences, and the weights decrease slightly faster than linearly over time. Surprisingly, when Malmendier and Nagel (2011) fit a similar model to experienced stock returns and stock market participation in a separate dataset, they receive almost identical sequence of weights on past data. The sensitivity parameter β, 0.675 (cl. s.e. 0.079), suggests that for every 1% point difference in learning-from-experience forecasts between two individuals, their inflation expectations differ, on average, by 0.672% points.

There may be serial correlation in the error terms within a cohort across time or correlation in the errors within a cohort in a time period. Malmendier and Nagel use standard errors that are robust to two-way clustering, with clusters by cohort and quarter. But calculating two-way clustered standard errors for a non-linear model is computationally costly (Cameron et al. 2009) relative to my analysis, which does not ultimately depend on these standard errors. My errors are non-linear least square estimators, but are not clustered. This does not affect my actual parameter estimates for θ or β, but the standard errors will potentially be biased downward. Care needs to be taken when interpreting the significance of my estimates. Although this is one drawback to the replication, it is immaterial to the interpretation of simulated counterfactuals considered later.

Malmendier and Nagel do not report the significance of their estimates. Although they mention that the estimate of β is significant, they do not indicate at what level. Since our β is a non-linear least squares estimator, I assume that β is asymptotically normally distributed. The authors' estimate of β is significant at the 1% level under this assumption. The gain parameter θ is also a non-linear least squares estimator, but Malmendier and Nagel never assess its significance.

My replication results are nearly identical to those of Malmendier and Nagel. Column (2) of Table 1 reports my estimates. My θ, 3.070 (s.e. 0.263) and β, 0.630 (s.e. 0.057), are close to the authors' original estimates in column (1), with my estimate of β falling in a 95% confidence interval around Malmendier and Nagel's estimate. Appendix C delves into the inference procedure for these parameters.

My estimate for θ confirms that experiences decline in importance over time. The sensitivity parameter β, 0.632, suggests that for every 1% point difference in learning-from-experience forecasts between two individuals, their inflation expectations differ, on average, by 0.632% points. The replication results show that 52.9% of the variation in inflation expectations can be explained by learning-from-experience. While my replicated model has explanatory power, it does not fit the data as well as Malmendier and Nagel's estimates. I cannot compare the standard errors, as Malmendier and Nagel use two-way clustered errors while my standard errors are non-linear least square estimates. Given that the disturbances can be correlated within a cohort and across time, my standard errors may be underestimating the error. Appendix C analyzes the possible source of differences between our results.

While Malmendier and Nagel use inflation expectations data from 1965 to 2009, I re-estimate the model using the most recent Michigan Surveys data on inflation expectations, which range from 1965q2 to 2013q1. They are reported in column (3) of Table 1. The estimate of θ using recent data is slightly higher than in column (2), but the unrobust standard errors are identical. The sensitivity β falls slightly but is more precise relative to (2).

Malmendier and Nagel consider possible omitted variable bias in their model. They focus on the possibility that their results are driven by different inflation rates in age-specific consumption baskets. That is, lifetime inflation experiences may be correlated with recent age-specific inflation rates. The authors re-run estimation controlling for differences between the overall CPI inflation rate and a CPI constructed from a common consumption basket for the elderly obtained from the Bureau of Labor Statistics. They find that their new results are not statistically different from original estimates, suggesting that their results are robust to age-specific inflation rates.

The authors in addition explore the common factor captured by the time fixed effects. They first consider the possibility that the time fixed effects capture the opinions of professional forecasters by replacing the time fixed effects with professional forecasts from the Survey of Professional Forecasters (SPF). They obtain similar estimates for β and θ, suggesting that SPF forecasts align closely with the common factors captured by time fixed effects. Another possibility they consider is that the common factor is a result of socializing with others about inflation expectations. Malmendier and Nagel rerun the model but replacing the mean learning-from-experience forecast across all age groups instead of the time fixed effects. They obtain almost identical estimates as with learning-from-experience, suggesting that the time fixed effect is absorbing the effect of socializing on inflation expectations.

The analysis depends on estimating personal inflation experiences based on birth date and headline inflation rates. But a prominent literature has suggested that headline inflation masks considerable heterogeneity in actual, individually experienced inflation (Hobijn et al 2009). Therefore approximating individual experiences with macroeconomic history is imprecise. Maderia and Zafar (2012) exploit demographic information available in the Michigan Surveys, such as gender, ethnicity, and education levels, to construct group-specific inflation rates. They estimate a similar learning-from-experience model using these group-specific inflation rates to more precisely predict individual-level inflation experiences. They find that such an adjustment has little extra explanatory power, suggesting that headline inflation rates sufficiently capture inflation experiences. Their result follows other work that confirms different individual inflation experiences do not explain observed heterogeneity in inflation expectations (McGranahan and Paulson 2006).

However, this class of models does not have complete identification. While Malmendier and Nagel show that learning-from-experience strongly explains inflation heterogeneity, there is unconvincing analysis on whether the memory of high inflation experiences is *causing* individuals to give higher inflation expectations. They claim that time fixed effects in the model "rule out that omitted macroeconomic variables or any other unobserved effects common to all individuals bias the estimation results." They do not disentangle other effects that may bias their estimates. Formally, their assumption that the $\epsilon_{t,s}$ is uncorrelated with the learning-from-experience forecast is problematic.

For example, there may be cohort-level effects. Since lifetime inflation experiences are determined by birth, any cohort-level effect will be correlated with the learning-from-experience forecast. It is fairly possible that this cohort effect, which can capture generational attitudes toward economic decisions, can affect how individuals form expectations. A cohort effect, by definition, varies across individuals in one time period. Therefore the time fixed effects Malmendier and Nagel place do not cleanly identify the model

It is worth noting that, since I make predictions using the learning-from-experience model, I do not depend entirely on the robustness of Malmendier and Nagel's identification.

6 Simulation

In this section I run simulations on the learning-from-experience model. The following counterfactuals will assess the relative importance of demographics and history in aggregate inflation expectations. Particular attention is given to how well the model explains stable inflation expectations during the Great Moderation.

It is possible that individuals consult others when forming expectations. Since a demogaphic shock may remove people who could have influenced others' beliefs, simulating predictions based on the model described by equation (6) will be inaccurate. I estimate Malmendier and Nagel's learning-from-experience model with social learning. It is based on the assumption that as individuals share opinions, their opinion tends toward to the average opinion (DeGroot 1974). Denote $\tau_{t+1|t,s}=\mathbf{b}'_{t,s}\mathbf{x}_t$ the learning-from-experience forecast at time t for individual of cohort s. Let $\bar{\tau}_{t+1|t}$ therefore be the average learning-from-experience forecast across all individuals in a time t. Learning-from-experience with social learning is therefore:

$$\tilde{\pi}^e_{t+1|t,s} = \gamma\tau_{t+1|t,s} + (1-\gamma)\bar{\tau}_{t+1|t} + (1-\gamma)\epsilon_{t,s} \tag{7}$$

Column (4) of Table 1 reports estimates. The gain parameter θ is higher than the one reported in column (3), which is a pattern Malmendier and Nagel also find. They compare estimates, but do not perform inference to verify that the variation is insignificant.

I examine five aggregate properties of inflation expectations in simulation: (1) the inflation expectation, (2) inflation expectations with social learning (3) the perceived autocorrelation, or mean-reversion, (4) the perceived mean, (5) disagreement about inflation expectations, and (6) disagreement about inflation expectations with social learning. Variables (1) to (4) are simply the weighted mean across all surveyees at time t. Appendix C reviews the construction of these variables.

The perceived autocorrelation (3), ϕ, is the tendency for expectations to return to the perceived long-run mean. Thus, it is one measure of the anchoredness of inflation expectations. For positive values, a low ϕ denotes well-anchored expectations while a high ϕ represents unanchored expectations. It is difficult to interpret a negative ϕ. I treat a more negative phi to mean unanchored expectations, because in that case expectations can change dramatically from short-run fluctuations. The aggregate perceived autocorrelation is simply the mean perceived autocorrelation across individuals at time t.

A note on language. For clarity, I will refer to the baseline, unsimulated time series produced by learning-from-experience as the "actual series." The actual series are time series predictions by the learning-from-experience model using (1) actual U.S. inflation history and (2) actual distribution of the surveyees.

6.1 Demographic Shock

What is the effect of demographic shocks on aggregate inflation expectations? If there are two individuals, the first born in year s and the second in $s+j$, only the first individual experiences events between s and $s+j$. Learning-from-experience implies that they may form different inflation expectations since their experiences differ. It is possible, then, that a demographic shock that removes either individual could change aggregate inflation expectations.

I define five demographic groups: those younger than 16, 16 to 30, 31 to 45, 46 to 60, and 61 to 74 year olds. The demographic shock permanently removes one cohort. For example consider a shock that removes all 40 year olds at time t. After the shock, there are no longer 40 year olds at time t. At time $t+1$, there are no 41 year olds in the sample, but there are 40 year olds at $t+1$ since those aged 39 at t have aged one year. Since Malmendier and Nagel excluded respondents older than 74, my demographic set excludes them also. I do include groups of people younger than 25, even though they are not in the authors' analysis. They enter the population after some time. For example, although a 1 year old does not respond to the survey, the 1 year old does enter the surveyee pool decades later.

I set the shock at 1988 for the simulation considered here. This is because it is approximately the beginning of Alan Greenspan's tenure at the Federal Reserve, thus signaling a change in the monetary policy regime. The timing may be informative, because some have attributed low inflation and stable expectations in the past two-and-a-half decades to prudent monetary policy (Davis 2012). The goal of setting demographic shocks at 1988 is to explore an alternative explanation for the path of inflation expectations during the Great Moderation.

Figures 6 to 11 show the effect of removing each cohort on aggregate properties of inflation expectations. Each figure corresponds to the variables of interest described early in this section. Overall, these demographic shocks did not substantially affect the variables studied, or the effect dies out by 2013.

$$< Figure\, 6 >$$

$$< Figure\, 7 >$$

According to Figure 6 and 7, aggregate expectations stayed nearly unchanged. At most, removing the <16 years old cohort increased inflation expectations by 0.01% points by 2013. With social learning, the change in expectations is of a slightly greater magnitude than with learning-from-experience predictions. By 2013, removing ¡16 year olds increased inflation expectations with social learning by 0.32% points. In contrast, removing those between 30-45 reduced inflation expectations by 0.6% points by 2013. However, none of these changes is relatively substantial.

$$< Figure\, 8 >$$

Similarly, Figure 8 shows that removing the cohort younger than 16 increases the perceived autocorrelation of inflation by 0.04% points while removing those aged 31 to 45 decreases it by 0.05% points. This loosely suggests that the young comprise the component with the best anchored expectations and those aged 31-45 is the component with the worst anchored inflation expectations. Since better anchoring of inflation expectations is a welcome but puzzling recent phenomenon (Davis 2012), these demographic shocks suggest that better anchoring is associated with different demographic groups.

6.2 Demographic Aging

Learning-from-experience posits that older individuals, already with more experiences, are less impressionable to new inflation surprises. Formally, the gain parameter, $\gamma_{t,s}$, declines with age. Given the coincidence of demographic aging in the U.S. captured by Figure 1 and better anchored inflation expectations, a natural question is to ask whether there exists a link between the two. That is, does demographic aging explain (1) lower and stable inflation expectations and (2) better anchored inflation expectations according to learning-from-experience?

The relevant counterfactual for comparison is if the U.S. did not experience demographic aging. I take the population distribution at 1985 and fix it for every time period after 1985. Individuals still age in this counterfactual, but the population as a whole, or the average age, has stopped aging. 1985 is what some consider as the start of the Great Moderation (Stock and Watson 2003). Jaimovich and Siu (2009), in their counterfactual analysis of demographic change and the Great Moderation, set the age distribution of the population fixed after 1985 for this reason.

$$< Figure\ 12 >$$

Figure 12 illustrates the change in the six variables of inflation expectations with the demographic "freeze" at 1985. Aggregate inflation expectations predicted by learning from experience and with social learning barely changes; by 2013 aggregate expectations and expectations with social learning in the simulation are 0.06% points and 0.08% points lower, respectively. Against this counterfactual, aging population since 1985 is associated with higher inflation expectations, at most up to 0.15 percentage points by 2010. While the absolute change is modest, it is evidence against the role of demographic aging in lowering inflation expectations.

$$< Figure\ 13 >$$

Surprsingly, perceived autocorreation fell. For better illustration, Figure 13 compares the time series of perceived correlation under actual demographics and under the demographic "freeze." In 2013 for example, a one-percentage point increase in inflation in the past quarter is associated with a 0.13 percentage point increase in inflation expectations, compared to 0.05 percentage point increase in a world without population aging. In the counterfactual world where there is no demographic aging after 1985, inflation expectations became more anchored, being nearly perfectly anchored for a few years in the 2000s. According to learning-from-experience, demographic aging since 1985 does not explain better anchored inflation expectations; in fact, demographic aging seems to exert the opposite effect. The result is unusual, since it was hypothesized that an aging population is less impressionable to new inflation experiences.

6.3 Legacy of 1970s

Learning-from-experience accepts a demographic distribution and an inflation history to predict aggregate inflation expectations. The prior simulations investigated the role of demographic shocks and shifts on inflation expectations, but found they have a minor impact on inflation expectations. I investigate the effect of history on predictions by learning-from-experience.

What is the effect of experiencing above-average inflation in the 1970s on inflation expectations today? Since the learning-from-experience model places some weight on past events, events in the 1970s should have a measurable effect on expectations in the future.

The relevant counterfactual is the time series of inflation expectations and its properties if the 1970s actually had average inflation rates. For simulation, I propose a hypothetical situation where between 1965 and 1985 the average inflation rate in this country equals that of the U.S. from 1985-2013. Essentially, in the counterfactual, the U.S. did not have historically above-average inflation experiences. The experience of, for example, West Germany is roughly analogous to this counterfactual. I produce a counterfactual inflation history for this country, π_t^{new} where:

$$\pi_t^{new} = \begin{cases} \pi_t^{US} \bullet \frac{\overline{\pi}^{US}_{1985-2013}}{\overline{\pi}^{US}_{1965-1985}} & \text{if } t = 1965, ..., 1985 \\ \pi_t^{US} & \text{otherwise} \end{cases}$$

The effect of this rescaling is that the mean inflation in this new series from 1965 to 1985 is the same as that from 1985 to 2013. The constructed and actual inflation series are shown in Figure 14.

$$< Figure\ 14 >$$

I can now use the learning-from-experience model to generate aggregate inflation expectations for this migrant group. Comparing expectations for this migrant group to that of the general population (excluding the group) shows us the effect of above-average inflation in the 1970s on current inflation expectations.

$$< Figure\ 15 >$$

Figure 15 shows the difference between aggregate expectations in this new counterfactual inflation history and expectations with historical inflation. Unsurpringly, inflation expectations before 1985 are substantially lower, because actual inflation in those years is scaled lower. More interesting is the time series of expectations after 1985, when inflation returns to normal, so there is no diifference in inflation between the counterfactual and actual series. Aggregate inflation expectations with learning-from-experience and with social learning show that migrant inflation expectations remain lower than that of the general population, up to 1% points lower a decade later in 1995. But both expectations converge, and by 2013, counterfactual expectations are 0.07% points lower than the general population.

There is evidence the 1970s unanchored inflation expectations, even decades later. Under the counterfactual inflation series, perceived autocorrelation steadily falls through the 1970s. It is not, however, obvious that autocorrelation should fall by rescaling inflation lower. Even after 1985, when inflation returns to normal rates, autocorrelation continues to fall relative to actual perceived autocorrelation. In the mid 2000s perceived autocorrelation is lower by 0.3% points. These counterfactual inflation expectations are better anchored for every year except for 2006Q1 to 2008Q2, when the autocorrelation is negative. By 2013 in this counterfactual, a one percentage point increase in inflation decreases inflation expectations by -0.02 percentage points. This is compared to the perceived autocorrelation with actual inflation history, where a one percentage point increase in inflation in 2013 leads to a 0.13 percentage point increase in inflation expectations.

Disagreement in inflation expectations decreases for almost every year in the counterfactual. By 2013, the standard deviation in inflation expectations is -0.04% points lower.

6.4 Downweighting

My estimate of θ suggests that past inflation experiences decline in importance over time. The precise reason for this is undetermined. It could be that past experiences are forgotten and thus carry little weight in expectations. Alternatively, it can reflect the possibility of parametric drift, where past data no longer has predictive content today. One question is to ask what the consequence of downweighting is on current inflation expectations.

The relevant counterfactual are a time series of inflation expectations where past data does not get downweighted. I alter learning-from-experience by setting the gain parameter, $\gamma_{t,s}$=1/t, which is a recursive formulation fo ordinary least squares (Evans and Honkapohja 2005).

$$< Figure\ 16 >$$

Figure 16 plots the difference in inflation expectations between predictions where past data does not get downweighted the actual series. Both aggregate mean inflation expectations as well as expectations with social learning predict that expectations would be below actual expectations from 1966Q2 to 1997Q3. Expectations without downweighting past experiences are higher, albeit with fluctuations, than the actual series after 1997Q3. There is nearly no difference - about 0.04 percentage points - by 2013 between the two series.

But removing downweighting increases the perceived autocorrelation of inflation expectations. While without downweighting the autocorrelation of expectations fluctuate around actual autocorrelation before 1986Q2, after that the simulated autocorrelation trends higher. By 2013, a one percentage point increase in inflation raises inflation expectations by 0.44 percentage points, substantially higher than 0.13 percentage points for a one percent increase predicted with downweighting of past data. In the counterfactual world where individuals do not downweigh past data, inflation expectations are substantially less anchored.

The result suggests that downweighing past data has explanatory power for the recent anchoring of inflation expectations. I follow the method employed in Jaimovich and Siu (2009) for a simple accounting exercise to approximate the magnitude of this effect. I compare the counterfactual autocorrelation and actual correlation predicted by learning-from-experience. From 1985, had individuals not downweighted past experiences, perceived autocorrelation would have fallen 0.202 percentage points by 2013. This is in comparison to 0.529 percentage points observed with downweighting. Therefore the downweighting of inflation experiences explains (0.529-0.202)/0.529 or 61.8 percent of the improved anchoring in inflation expectations since 1985.

7 Discussion

I estimate a learning-from-experience model in which individuals use their lifetime inflation experiences to form inflation expectations. Using survey-reported inflation expectations from 1966 to 2013, I confirm key findings by Malmendier and Nagel (2013). First, the learning-from-experience model explains 52.3% of the heterogeneity in inflation expectations. Second, my estimate of the gain parameter implies that individuals downweigh past experiences when forming inflation expectations. In fact, when Malmendier and Nagel (2011) estimate a similar learning-from-experience model to explain stock market participation using experienced stock market returns, they receive almost identical sequence of weights on past experiences (Malmendier and Nagel 2013). Third, a 1% point difference in learning-from-experience forecast is associated with a 0.829% point difference in overall inflation expectations on average. In the model, this implies that differences in inflation experiences, even when holding summary information costant, is correlated with inflation expectations.

Using the replicated model, I study the response of aggregate variables to demographic shocks. Setting the shock at 1988, the beginning of a new monetary policy regime that also coincided with stabilizing inflation in the U.S., demographic shocks might shed life on the path of inflation expectations during the Great Moderation. Unexpectedly, aggregate expectations, with or without social learning, changed little from any of the demographic shocks. It is surprising, because the U.S. experienced high inflation in the 1970s that fell in the mid-1980s, which by the learning-from-experience model should show up as generational differences in inflation expectations. It is possible that after decades of low inflation post-1988, past data gets sufficiently downweighted. In this case, generational differences in experience narrows over time, and expectations converge so that demographic shocks exert little change on aggregate expectations.

The Great Moderation coincided with the aging of the baby-boomers in the U.S. Jaimovich and Siu (2009) show that this demographic shift strongly explains the reduction in business cycle volatility in the U.S. Since learning-from-experience posits that older individuals are less impressionable to new information, there is a possibility that demographic aging can explain better anchoring of inflation expectations. Surprisingly, a simulation that "freezes" the age distribution after 1985 produced lower inflation expectations and better anchoring. This implies that population aging since 1985 had actually raised and unanchored, though minimally, inflation expectations. According to the model, inflation expectations fell and became better anchored in the U.S. despite, not because of, demographic aging.

This surprising result is probably due to the specific history of inflation in the U.S. While older individuals in the model are less responsive to new information, people still use lifetime inflation history to forecast inflation. Stopping population aging is essentially reducing the proportion of old people. Therefore, reducing the ranks of older people at 1985 reduces the number of people who experienced high inflation in the 1970s. If it is true that the 1970s had raised and unanchored inflation expectations, then it is reasonable for expectations to stabilize when we remove the old cohort who experienced the 1970s.

I also create an alternative inflation history where the U.S. did not experience high inflation in the 1970s. Using the model, I predicted aggregate inflation expectations using this alternate series. As expected, reducing inflation produced lower inflation expectations between 1965-1985. After 1985, expectations remain lower than actual expectations, and even a decade later in 1995 counterfactual inflation expectation is 1% point lower than the baseline. There is a slow convergence back to actual inflation expectations, such that by 2013 inflation expectations using alternative inflation history is nearly identical to the actual expectation. This implies that high inflation in the 1970s raised inflation expectations decades afterward, with the effect slowly fading out by 2013.

Not only did the 1970s unanchor inflation expectations, but its effect surprisingly increased with time. Whereas high inflation seemed to have raised perceived autocorrelation by 0.1% points in the early 1980s, by 2000s inflation in the 1970s raised autocorrelation by 0.3% points. This result that the effect of history increases, rather than diminishes, over time is puzzling and difficult to reconcile with the model. However, it must be true that given enough time, actual and simulated perceived autocorrelations must converge. This is because at some point everyone who experienced the years before 1985 must die, and the 1970s will become irrelevant to aggregate inflation expectations.

Another result of this simulation is that disagreement in inflation expectations decreases by rescaling inflation in the 1970s. What this implies is that a modest portion of recent disagreement in inflation expectations can be explained by those who experienced high inflation in the 1970s. This is a sensible finding, since it attributes disagreement to specific episodes of inflation history. People who experienced the 1970s tend to have higher expectations than those that did not, leading to disagreement.

The final simulation examined the effect of downweighting past experiences on aggregate inflation expectations. Downweighting past data does not seem to exert a clear effect on aggregate expectations, perceived mean inflation, or disagreement. It does, however, have the effect of anchoring inflation expectations substantially. The results imply that beginning in the mid 1980s, the downweighting of past inflation experiences better anchored inflation expectations. A simple accounting exercise confirms that roughly three-fifths of the better anchoring of inflation expectations since 1985 can be explained by downweighting past experiences.

This result may be counterintuitive, because downweighting past experiences necessarily implies putting more importance to recent data. Why does downweighting improve anchoring if downweighting must make an individual more sensitive to recent events? One possible explanation is that recent inflation itself has become less persistent, a feature observed in Figure 17, so expectations have become more stable. Downweighting induces higher sensitivity to this less persistent recent inflation. In turn, this translates into lower perceived autocorrelation.

The inevitable challenge of inflation expectations is simultaneity bias. Learning-from-experience posits that actual inflation predicts inflation expectations. But inflation expectations clearly affects actual inflation, as expectations may drive price-setting decisions (Romer 2003). This requires a system of two equations that are determined simultaneously. Estimating one equation without the other, say regressing inflation expectations on inflation, potentially leaves errors that are correlated with a regressor, namely inflation. This biases estimates.

$< Figure\ 17 >$

Consider Figure 17. It compares the perceived autocorrelation from learning-from-experience to a 10-year rolling window estimate of the autocorrelation of actual, not expected, inflation rates. Just as perceived autocorrelation has fallen, so has actual autocorrelation of inflation. While actual autocorrelation began declining in the early 1990s while perceived autocorrelation only start declining in the 2000s, the two series roughly move together. Since my analysis requires regressing inflation expectations on inflation, simultaneity potentially poses a serious problem for interpreting coefficients on the AR(1) model.

But the setup of the learning-from-experience model may sidestep the simultaneity bias. The model posits that at time t, an individual forecasts inflation at $t + 1$ and her forecast uses available lifetime inflation rates until the end of period t-1. In other words at t the individual uses data up to and including π_t, which she just experienced in its entirety, to form π^e_{t+1}. The simultaneity problem arises if π^e_{t+1} affects π_t. But this is improbable, as beliefs cannot change events that have already happened. π^e_t may affect π_t, but π_t cannot affect π^e_t in learning-from-experience, because timing of the model dictates that π^e_t is formed before π_t even happens, thus averting simultaneity bias.

It remains, however, that inflation expectation affects inflation. My analyses only consider the effect of inflation on inflation expectations, and not the opposite effect described for example by the Phillips Curve (Romer 2002). In fact, changes in perceived autocorrelation of inflation are likely to influence autocorrelation of inflation (Milani 2007). This can affect the accuracy of simulations. For example, a demographic shock can change inflation expectations, but later expectations may affect actual future inflation rates. In turn, the change in actual future inflation causes future inflation expectations to change as well. Running counterfactual simulations on learning-from-experience alone does not account for the response of inflation to inflation expectations. Ultimately, the value of the simulations considered here is the implications of the learning-from-experience model on demographics and history.

8 Conclusion

My analysis shows that the learning-from-experience model does not predict substantial changes in aggregate inflation expectations due to demographic shocks. Demographic aging since 1985 also does not explain better anchoring of inflation expectations. Rather, history matters; high inflation in the 1970s raised and unanchored inflation expectations for decades afterward. I estimate that three-fifths of the better anchoring in inflation expectations can be accounted for by the downweighting of personal experiences.

A more realistic simulation than the ones considered here would account for the effect of inflation expectations on inflation. Such a modest, but more realistic, addition would help verify the conclusions found here. In addition, future interest in learning-from-experience ought to disentangle the effect of history and demographics. Changing demographics necessarily changes experienced inflation. It is unclear from the simulations which effect, history or demographics, is driving the often surprising results.

The central appeal of learning-from-experience models is that experience is ubiquitous. Everyone experiences. It is intuitively appealing that most people are swayed by personal experience in economic decisions, an effect learning-from-experience quantifies. Moreover, an advantage of learning-from-experience is it links micro-level behavior to aggregate outcomes. Models based on experience potentially can explain other macroeconomic phenomena besides inflation expectations.

But why does personal experience matter? Learning-from-experience documents how personal experience predict expectations, and that there is evidence experience explains the differences in subjective expectations. Their model does not, however, explain why individuals weigh personal experience more than information gleaned from summaries.

9 Appendix A - Data

This section describes the procedure to prepare data for analysis.

Before 1978, the Michigan Surveys of Consumers was the Survey of Consumer Attitudes and Behavior (SCAB), a quarterly survey from 1954 to 1977 with gaps in between. In addition, Malmendier and Nagel use inflation expectations included in the winter surveys of the Survey of Consumer Finances. Both are available from the ICSPR. We omit, for this analysis, SCAB surveys from 1954 to 1965, because those only solicit categorical responses of whether the price level will "go up", "stay the same", or "go down". There are, in total, 24 surveys containing inflation expectations from 1966 to 1977. The data format and coding for these surveys are vague and inconsistent. Non-trivial effort thus went into compiling and organizing the data.

A few SCAB surveys re-interviewed individuals several quarters after the initial interview. As a result, a survey may contain price expectations from many quarters. I omit data from any such re-interviews in this analysis, consistent with the frequency of birth cohorts and the cross section of inflation expectations given in Malmendier and Nagel (2013).

SCAB surveys before spring 1977 solicited percent estimates of inflation, but responses are coded as ranges. For ranges with an upper bound, I imputted a percent estimate by drawing from the uniform distribution of that range. For example, if respondents believed that the price level will rise by 1-2%, I imputted an estimate drawn from the uniform distribution [1%, 2%]. For ranges without an upper bound, I generated a percent estimate drawing on a $\chi^2(1)$ distribution scaled so that all draws are between the lower bound, $x_l\%$ and $x_l + 50\%$. The distribution approximates, albeit imperfectly, the distribution of data in surveys with actual percent responses. I reasonably assume every respondent who thought the price level will not change has a percent estimate of 0%.

Since learning-from-experience depends on cohort information, we extract birth year information from the surveys. For observation without birth years, I calculated the implied birth year given by age. Respondents with no birth year or age were omitted.

It is critical, then, to identify the correct age of the respondents to the survey. Before fall 1974, the survey does not contain the age of the respondent, but indicates the age of the head and wife, as well as whether the respondent is the head, wife, or other member of the household. I coded the age respective to the respondent, but omitted observations where the respondent was another member of the household. I drop all observations that were coded incorrectly.

There is a well-documented literature on calculating price expectations in the Michigan Surveys. We follow, as do Malmendier and Nagel, Curtin (1996) to correct several features of the data.

If respondents indicated that they expect prices to go up or down but do not give a percent estimate, I impute a percent response by drawing from the empirical distribution of percent estimates from that survey. For example, suppose half of the respondents thought prices will go up and gave a percent estimate estimated a 2% increase. Then of those who believed prices will rise but did not give an estimate, I would impute a 2% estimate. This approach includes respondents who believed prices will rise but will rise at the same rate as before.

Before February 1980, respondents who believed prices will fall did not give a percent estimate. I follow Curtin (1996) and code -3% for those respondents. Malmendier and Nagel note that these cases make up less than 2% of all responses.

In surveys before March 1982, some respondents misinterpreted ”same” response to mean that prices will rise at the same rate as before, not that the price level will stay the same. Curtin (1996) provides a model to estimate what fraction of “same” responses before March 1982 actually meant “up”:

$$SameError_t = 0.007 \bullet PctUp_t$$

So the new, adjusted fraction of “up” responses is:

$$PctUpAdj_t = SameError_t \bullet PctSame_t + PctUp_t$$

We correct these "same" responses to "up" and draw a percent estimate from the distribution of "up" responses in that period.

Curtin (1996) provides a truncation range to remove outliers in price expectations. Frequencies of price expectations provided by Malmendier and Nagel show little evidence the authors dropped outliers. I therefore do not use the truncation range provided by Curtin.

The Michigan Surveys of Consumers after 1978 provides a set of weights to approximate adults living in the U.S. However, data before 1978 have inconsistent weights. Some surveys have multiple weights and others have none. Since the model used a weighted mean price expectations by cohorts and quarter, weights complicated replication.

For surveys without weights, I weigh each respondent equally. For surveys with more than one weight, I prefer Head-Wife weights that approximate a representative sample of adults in the U.S. If Head-Wife weights are unavailable but other weights exist, I use other weights provided.

After correcting the data for missing information, I adjust the weights for the new frequency of observations. When I took mean expectations by cohort and time, I adjust the weights based on the frequencies in cohort and time.

We use log annualized inflation rates since 1877. Malmendier and Nagel's inflation series has removed substantial variation, suggesting that they are taking a moving average representation. Indeed, the series, after taking four-quarter moving averages, match closely to the authors' graphic distribution. The four quarter moving average representation effectively removes seasonal fluctuations in the data.

Less clear is if the authors use the raw or smoothed series for estimation. Estimation ran both inflation series. I only received estimates and time series close to the authors' results with the raw inflation. Therefore estimation and analysis used the raw, rather than the smoothed, inflation series.

10 Appendix B - Weighting Matrix Singularity

I highlight an estimation problem in Malmendier and Nagel's specifications. This is due to the singularity of the covariance matrix for degenerate values of the gain parameter. I begin recursion soon after birth with the initial condition that $\mathbf{R}_{t,s}$ is an identity matrix and $\mathbf{b}_{t,s} = (0,1)'$. Therefore I posit that in the initial condition, inflation expectations are simply the realized inflation rate of the previous period.

Consider a learning from experience forecast at any time $t^* < \theta + s$. Therefore $\gamma_{t^*,s} = 1$. But then

$$\mathbf{R}_{t^*,s} = \mathbf{x}_{t^*-1}\mathbf{x}'_{t^*-1}$$

Since $\mathbf{x}_{t^*-1} = (1, \pi_{t^*-1})'$

$$\mathbf{R}_{t^*,s} = \begin{pmatrix} 1 & \pi_{t^*-1} \\ \pi_{t^*-1} & \pi^2_{t^*-1} \end{pmatrix}$$

It is trivial to show that

$$\mathbf{R}^{-1}_{t^*,s} = \frac{1}{\pi^2_{t^*-1} - \pi^2_{t^*-1}} \begin{pmatrix} \pi^2_{t^*-1} & -\pi_{t^*-1} \\ -\pi_{t^*-1} & 1 \end{pmatrix}$$

is undefined. The model is undefined for the very first rounds of recursion considered here. For the purpose of estimation, I therefore start the recursion when $t - s = \theta$

11 Appendix C - Replication Robustness

I use a 95% confidence interval to evaluate my replicated results to those of the authors, an approach reviewed by Cumming and Maillardet (2006). Formally, my estimate of β falls within a 95% confidence interval of Malmendier and Nagel's β estimate.

I omit statistical inference for θ, since the authors also do not interpret on θ. To confirm my replication for θ, I explore the source of variation by considering replication in two parts, data and estimation.

Learning-from-experience is simply an adjusted formulation of ordinary least squares. Therefore, a simple way of veriying my replicated model is to adjust the model so that it produces ordinary least squares estimates and compare those estimates to that from a software regression output. The adjustment is trivial. Rather than specifying the gain parameter as is described, I set $\gamma_{t,s} = 1/t$. Evans and Honkapohja, as noted, show that this produces ordinary least squares estimates. I use this adjusted model to estimate an AR(1) using inflation from 1877q1-2013q1.

I confirm the results by separately obtaining AR(1) OLS estimates using a statistical software.

$$< Table\ 2 >$$

Table 2 displays the coefficients from this exercise. These coefficients are nearly identical, which confirm that my replicated model in fact is the adjusted learning-from-experience model.

Therefore I attribute the difference between my and Malmendier and Nagel's results to differences in data preparation. There are not enough details in Malmendier and Nagel's paper to perfectly replicate data preparation. Consider Table 1, the estimates of the model. Malmendier and Nagel have substantially fewer observations. In Figure 1, the cross-section of inflation expectations, expectations across age groups differ significantly in the 1970s, whereas a similar figure by the authors show that expectations started off fairly similar in the 1970s. Ambiguities as well as departures on the authors' part to conventional preparation of the Michigan Surveys abound. Appendix A describes them in detail.

12 Appendix D - Simulation Variables

This appendix details the variables examined for each simulation. In the calculation of means, I apply $\omega_{t,i}$, the implied weight on the respondent provided by the Michigan Surveys. The weight guarantees that the survey sample approximates the adult population in the U.S.

Define (1) aggregate inflation expectations $\bar{\pi}^e_t$ at time t as:

$$\bar{\pi}^e_{t+q|t,s} = \sum_{i=1}^{N_t} \omega_{t,i} \frac{\pi^e_{t+q|i,t}}{N_t}$$

where i is an individual at time period t, N_t is number of individuals at t, and $\pi^e_{i,t}$ is the learning-from-experience prediction for person i described by equation (6). Aggregate inflation expectations is the mean inflation expectation predicted by learning-from-experience.

(2) is aggregate inflation expectations with social learning, denoted $\bar{\tilde{\pi}}^e_t$ at time t:

$$\bar{\tilde{\pi}}^e_{t+q|i,t} = \sum_{i=1}^{N_t} \omega_{t,i} \frac{\tilde{\pi}^e_{t+q|i,t}}{N_t}$$

where i is an individual at time period t, N_t is number of individuals at t, and $\pi^e_{i,t}$ is the learning-from-experience prediction for person i described by equation (7). Aggregate inflation expectations with social learning is simply the mean inflation expectations predicted by learning-from-experience with social learning.

The perceived autocorrelation (3), ϕ, is the tendency for expectations to return to the perceived long-run mean. Thus, it is one measure of the anchoredness of inflation expectations. For positive values, a low ϕ denotes well-anchored expectations while a high ϕ represents unanchored expectations. It is difficult to interpret a negative ϕ. I treat a more negative phi to mean unanchored expectations, because in that case expectations can change dramatically from short-run fluctuations. Formally, aggregate perceived autocorrelation, $\bar{\phi}_t$ is:

$$\bar{\phi}_t = \sum_{i=1}^{N_t} \omega_{t,i} \frac{\phi_{i,t}}{N_t}$$

where i is an individual at time period t, N_t is number of individuals at t, and $\phi_{t,s}$ is the coefficient on the first lag inflation in the AR(1) for person i in equation (1). Aggregate perceived autocorrelation is the mean perceived autocorrelation estimate derived from the AR(1) component of learning-from-experience.

The perceived mean inflation (4), μ is the long-run inflation rate based on estimates of the AR(1). Taking equation (1) and setting inflation at long-run equilibrium, $\mu_{t,s}$, we have:

$$\mu_{t,i} = \alpha_{t,i} + \phi_{t,i}\mu_{t,i}$$

Rearranging, I have

$$\mu_{t,i} = \frac{\alpha_{t,i}}{1 - \phi_{t,i}}$$

Thus aggregate long-run mean inflation is:

$$\bar{\mu}_t = \sum_{i=1}^{N_t} \omega_{t,i} \frac{\mu_{i,t}}{N_t}$$

where i is an individual at time period t and N_t is the number of individuals at t.

The disagreement in inflation expectations (5) is the standard deviation of inflation expectations predicted by learning-from-experience. Similarly, the disagreement in inflation expectations with social learning (6) is the standard deviation of expectations predicted by learning-from-experience with social learning.

References

[1] Arias, A., Hansen, G. D., and Ohanian, L. E. (2007). Why have business cycle fluctuations become less volatile?. *Economic theory*, 32(1), 43-58.

[2] 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.

[3] Bray, M. (1982). Learning, Estimation and the Stability of Rational Expectations Equilibriua. *Journal of Economic Theory* 26, 318-339.

[4] Bryan, M. and Venkatu, G. (2001). The Curiously Different Inflation perspectives of Men and Women. *Federal Reserve Bank of Cleveland Economic Commentary,* 1-4.

[5] Carroll, C. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics*, 118(1) 269-298.

[6] Choi, J. et. al (2009). Reinforcement Learning and Savings Behavior. *Journal of Finance*, 64(6), 2515-2534.

[7] Curtin, R. T. (1996). Procedure to Estimate Price Expectations, available at http://www.sca.isr.umich.edu/fetchdoc.php?docid=24772.

[8] Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business and Economic Statistics*, 29(2).

[9] Campbell, J. R., and Hercowitz, Z. (2005). The role of collateralized household debt in macroeconomic stabilization (No. w11330). National Bureau of Economic Research.

[10] Cumming, G., and Maillardet, R. (2006). Confidence intervals and replication: Where will the next mean fall?. *Psychological methods*, 11(3), 217.

[11] De Bruin, W. et. al (2010). Expectations of Inflation: The Role of Demographic Variables, Expectation Formation, and Financial Literacy. *Journal of Consumer Affairs*, 44(2), 381-402.

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

[13] Graham, J., Hazarika, S., and Narasimhan, K. (2011). Financial Distress in the Great Depression. *Financial Management*, 40(4), 821-844.

[14] Greenwood, R., and Nagel, S. (2009). Inexperienced Investors and Bubbles. *Journal of Financial Economics*, 93(2), 239-258.

[15] Hertwig, R., Barron, G., Weber, E., and Erev, I. (2004). Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychological Science*, 15(8), 534-539.

[16] Hobijn, B., Mayer, K., Stennis, C., and Topa, G. (2009). Household inflation experiences in the US: a comprehensive approach. Unpublished paper, Federal Reserve Bank of New York.

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

[18] Jaimovich, N. and Siu, H. (2009). The Young, the Old, and the Restless: Demographics and Business Cycle Volatility. *American Economic Review*, 99(3), 804-826.

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

[20] Madeira, C., and Zafar, B. (2012). Heterogeneous inflation expectations, learning, and market outcomes (No. 536). Staff Report, Federal Reserve Bank of New York.

[21] Malmendier, U., and Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*. *Quarterly Journal of Economics*, 126(1), 373-416.

[22] Malmendier, U., and Nagel, S. (2013). Learning from Inflation Experiences. *Unpublished Manuscript, UC Berkley.*

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

[24] Malthus, T. R. (1798). An Essay on the Principle of Population.

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

[26] Mankiw, N. G., R. Reis, and J. Wolfers. (2004). Disagreement About Inflation Expectations. In *NBER Macroeconomics Annual 2003*, Volume 18 (pp. 209-270). The MIT Press.

[27] McGranahan, Leslie and Anna Paulson (2006), Constructing the Chicago Fed Income Based Economic Index - Consumer Price Index: Inflation Experiences by Demographic Group: 1983-2005, Chicago Fed Working Paper, WP2005-20.

[28] Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics*, 54(7), 2065-2082.

[29] Over, M. (1992). The macroeconomic impact of AIDS in Sub-Saharan Africa.

[30] Piazessi, M., J. Salomo, and M. Schneider. (2013) Trend and Cycle in Bond Premia. *Unpublished Manuscript, Stanford University.*

[31] Ramey, V. A., and Vine, D. J. (2006). Declining volatility in the US automobile industry. *American Economic Review*, 1876-1889.

[32] Romer, D. (2012). *Advanced Macroeconomics*. McGraw-Hill.

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

[34] Shimer, R. (1999). Why is the US unemployment rate so much lower?. *NBER Macroeconomics Annual*, Volume 13 (pp. 11-74). MIT Press.

[35] Stock, J. H., and Watson, M. W. (2003). Has the business cycle changed and why?. *NBER Macroeconomics Annual*, Volume 17 (pp. 159-230). MIT press.

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

13 Tables and Figures

	(1) M&N (2013)	(2) Replication	(3) Extended	(4) Social Learning
Gain Parameter θ	3.144	3.076	3.178	3.765
(two-way clustered s.e.)	(0.257)	-	-	-
(s.e.)	-	(0.263)	(0.263)	(0.102)
Sensitivity β	0.675	0.632	0.629	-
(two-way clustered s.e.)	(0.079)	-	-	-
(s.e.)	-	(0.057)	(0.055)	-
Sensitivity γ	-	-	-	0.647
(s.e.)	-	-	-	(0.066)
Time	1966q2-2009q4	1966q2-2009q4	1966q2-2013q1	1966q2-2013q1
Adj. R^2	0.637	0.529	0.523	0.336
#Obs.	7650	7921	8582	8582

Table 1: **Estimates from AR(1) Learning-from-Experience Model**: The estimates are from a non-linear regression of one-year survey-reported inflation expectations on learning-from-experience forecasts. Malmendier and Nagel (2013) withhold significance levels of estimates, an approach adopted in my analysis. Malmendier and Nagel calculate standard errors that are clustered two-way by time and cohort. Due to the computational cost of calculating two-way standard errors for non-linear parameters (Cameron et al. 2009), I only include unrobust standard errors.

	Replicated Recursive OLS	Software Outputted OLS
Autocorrelation ϕ	0.3852	0.3852
	(0.0394)	-
Intercept α	1.4154	1.4162
	(0.3459)	-

Table 2: **Replication Robustness**: Comparing estimates of AR(1) model using an OLS-adjusted learning-from-experience model and OLS estimates from Stata. The estimates are based on inflation data from 1871q1 to 2013q1.

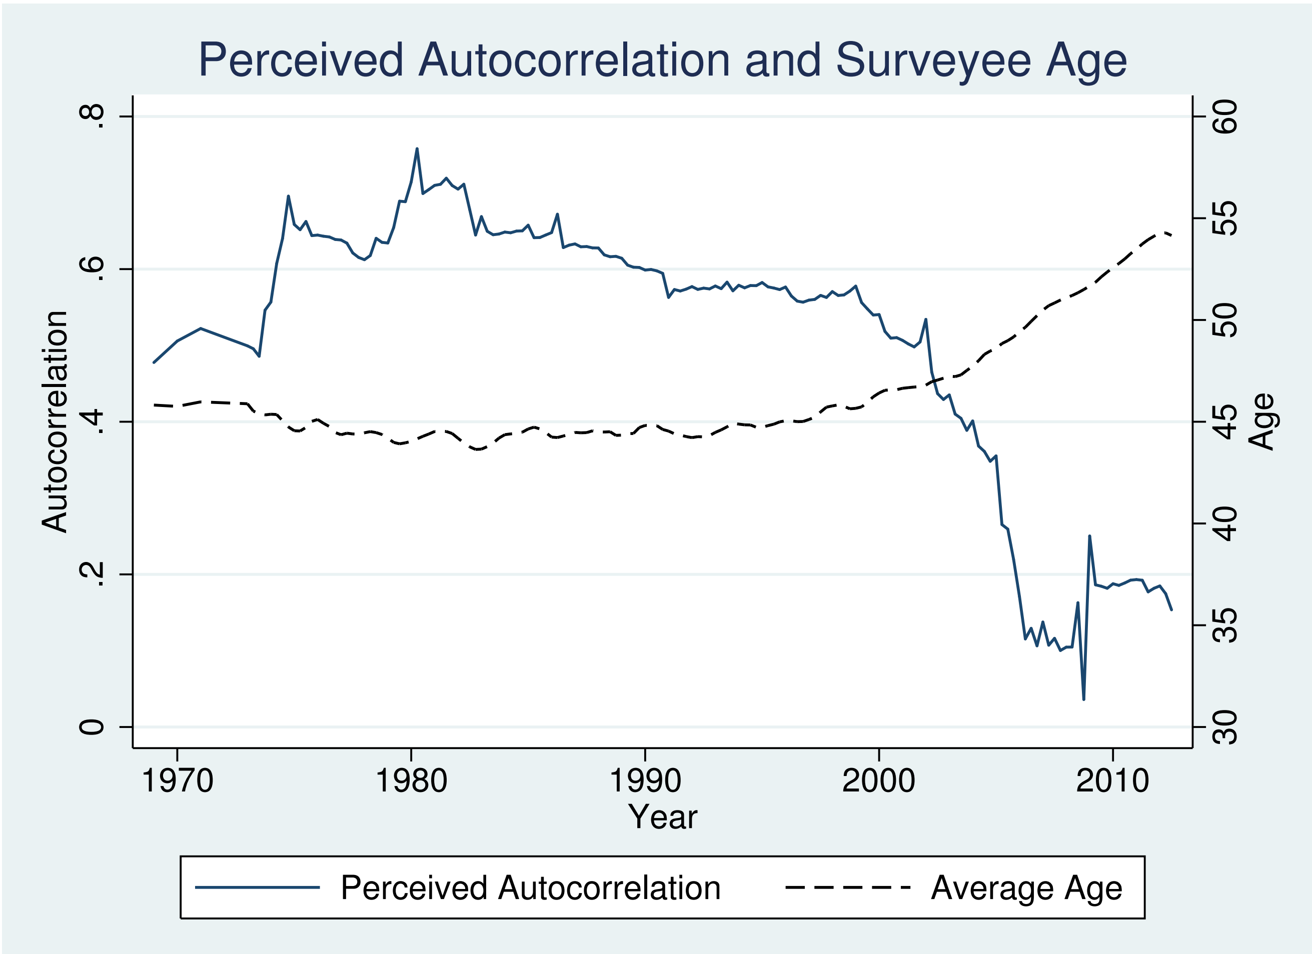

Figure 1: **Perceived Autocorrelation and Surveyee Age**: Solid line is perceived autocorrelation of inflation ($\bar{\phi}_t$) predicted by learning-from-experience AR(1) model. Dashed line is average age of respondents to the Michigan Surveys. Note that Malmendier and Nagel (2013) only use respondents who were aged 25 to 74 at the time of their survey interview. Both series above are constructed from respondents aged 25 to 74.

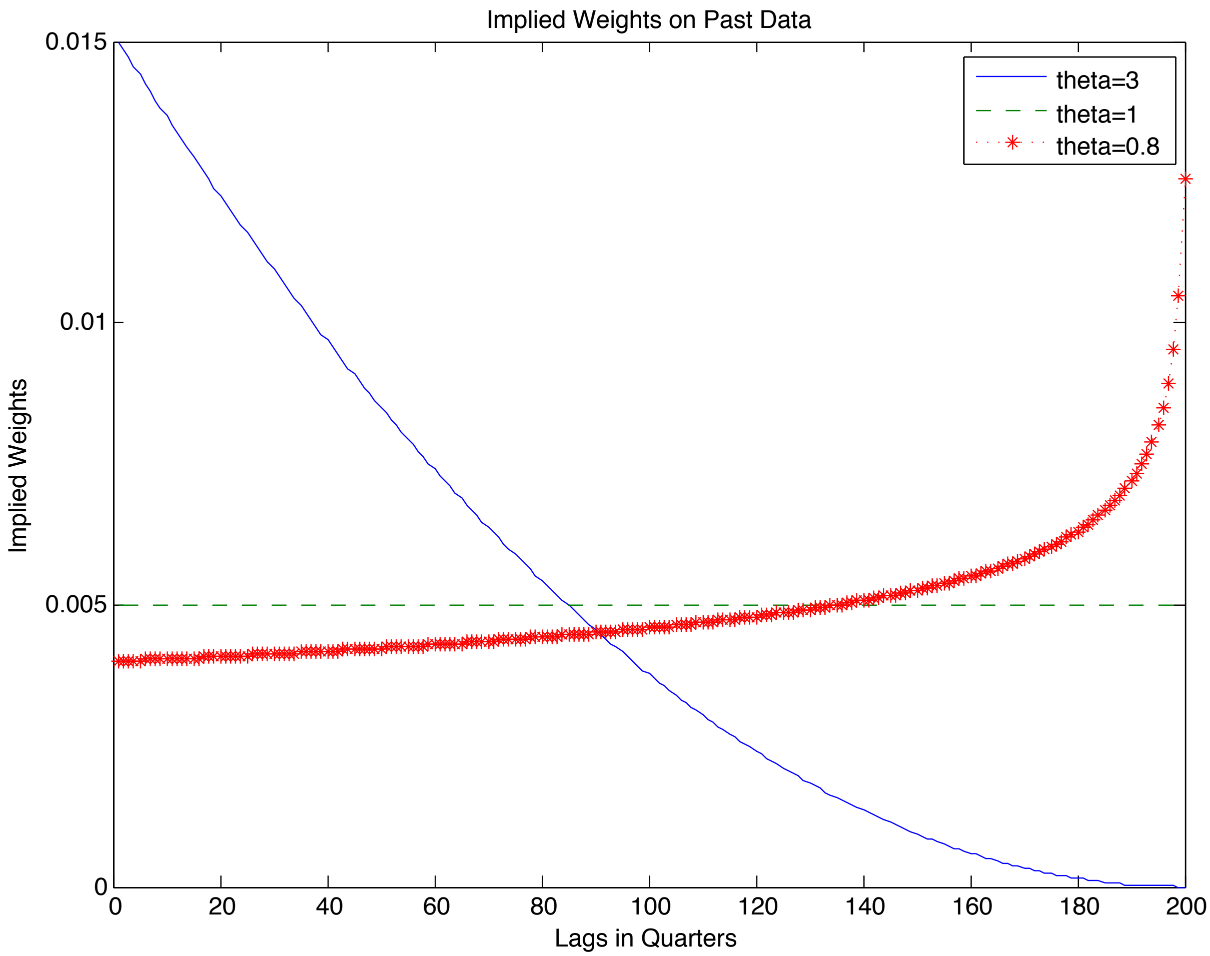

Figure 2: **Implied Weights on Past Data:** Implied weights on past experiences based on parameter values. This figure is a replication of Figure 2 in Malmendier and Nagel (2013)

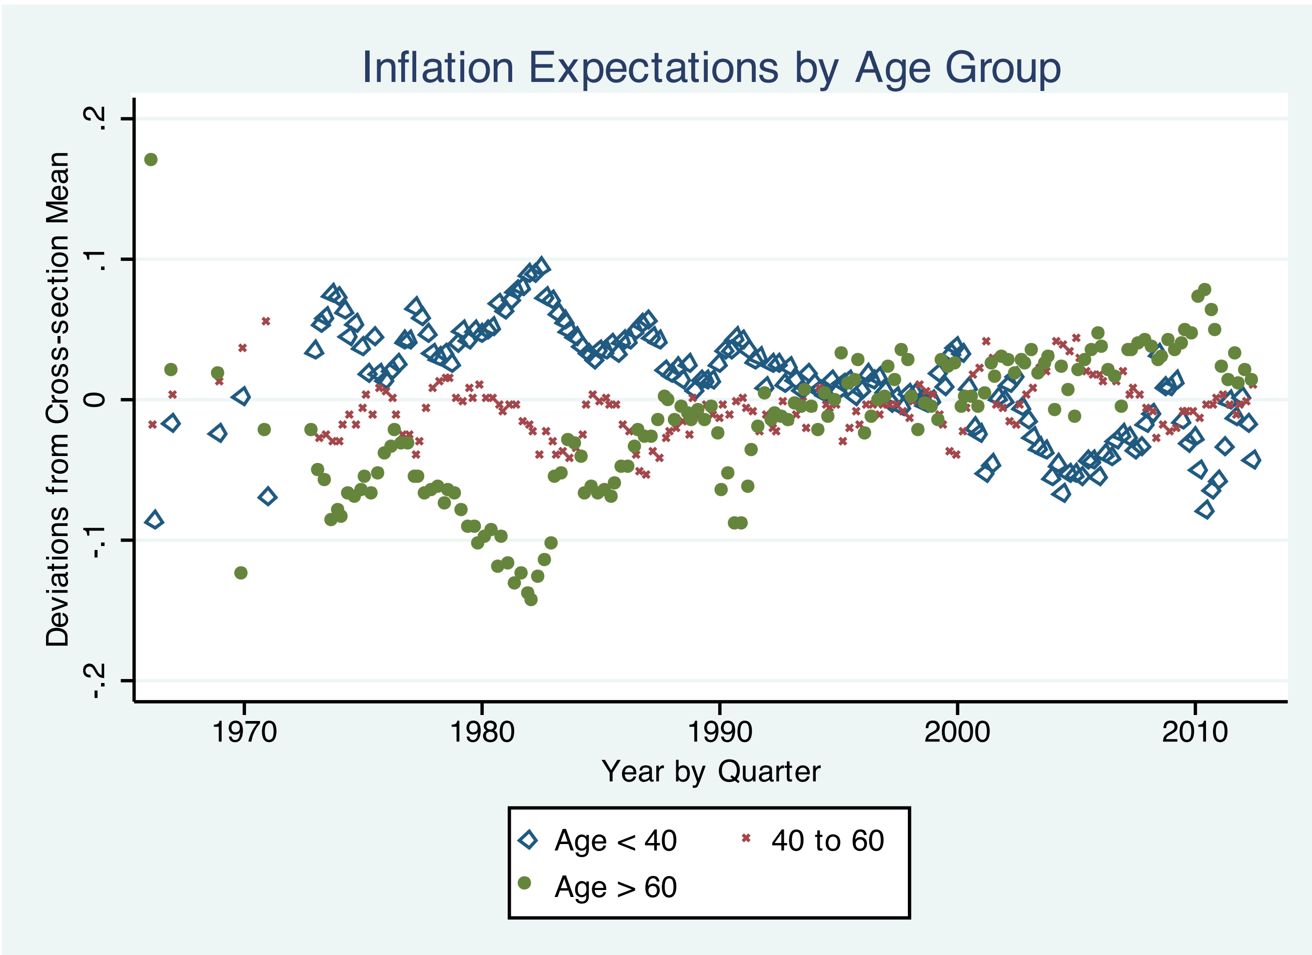

Figure 3: **Inflation Expectations by Age Group**: Four-quarter moving average representation of mean inflation expectations by age group, depicted as deviations from the cross-sectional mean expectation. Data from the Michigan Surveys of Consumers.

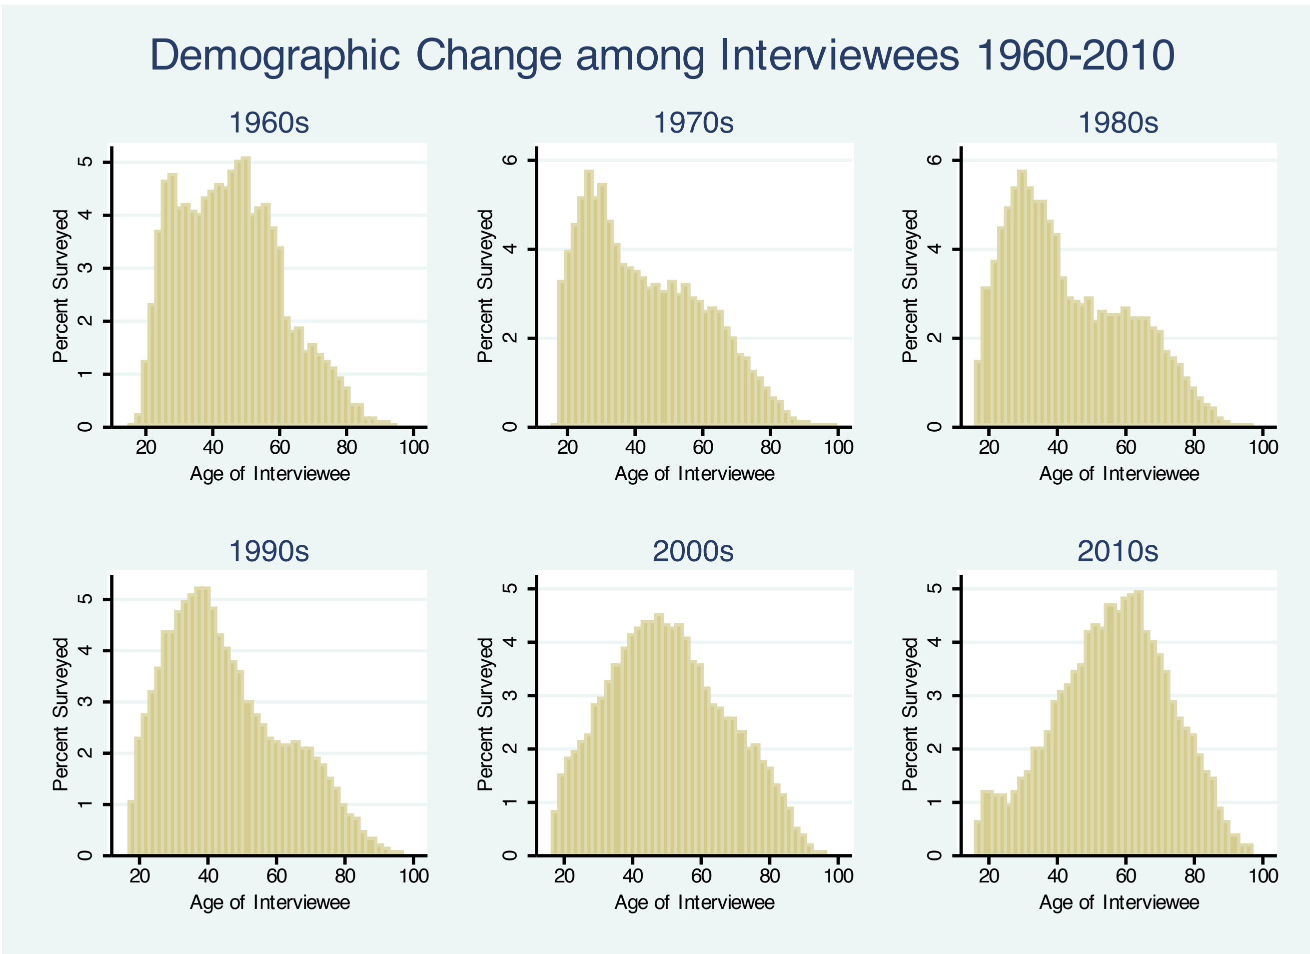

Figure 4: **Demographic Change among Interviewees 1960-2010**: Age distribution of respondents to the Michigan Surveys by decade. Note that Malmendier and Nagel (2013) only use respondents who were aged 25 to 74 at the time of their survey interview. However, above figure is constructed using all respondents to the survey, to capture the full dynamic of demographic aging. The survey does not interview individuals under 18.

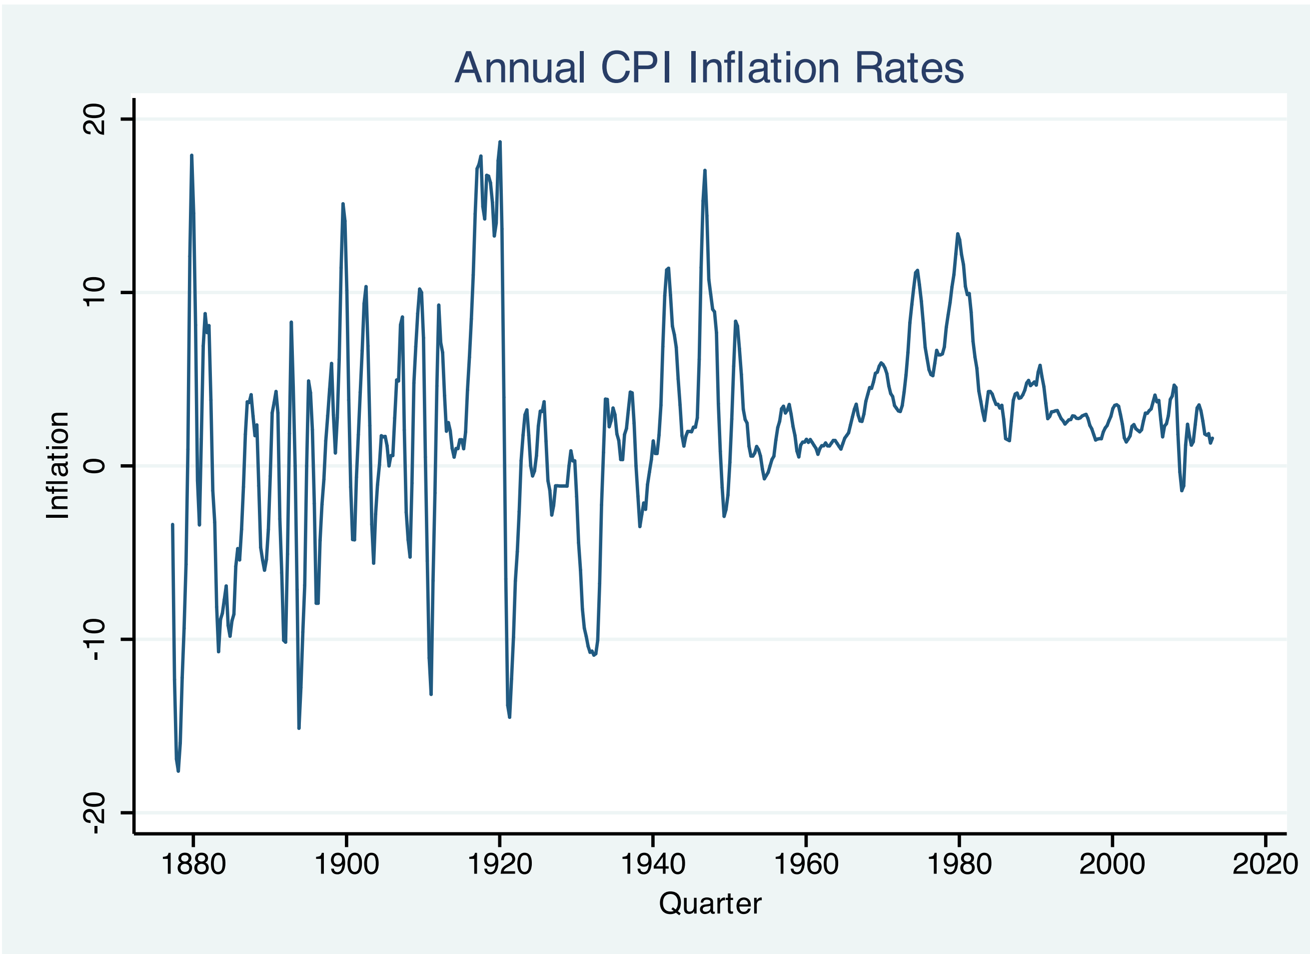

Figure 5: **Annual CPI Inflation Rates**: Four-quarter moving average of annual CPI inflation rates.

Figure 6: **Change in Aggregate Inflation Expectations Due to Demographic Shock at 1988**: Change is shown as the difference between aggregate inflation expectations with demographic shock and actual expectations, both predicted by learning-from-experience. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The aggregate inflation expectation is simply the average predicted expectation from learning-from-experience. Individual titles correspond to age group removed at shock.

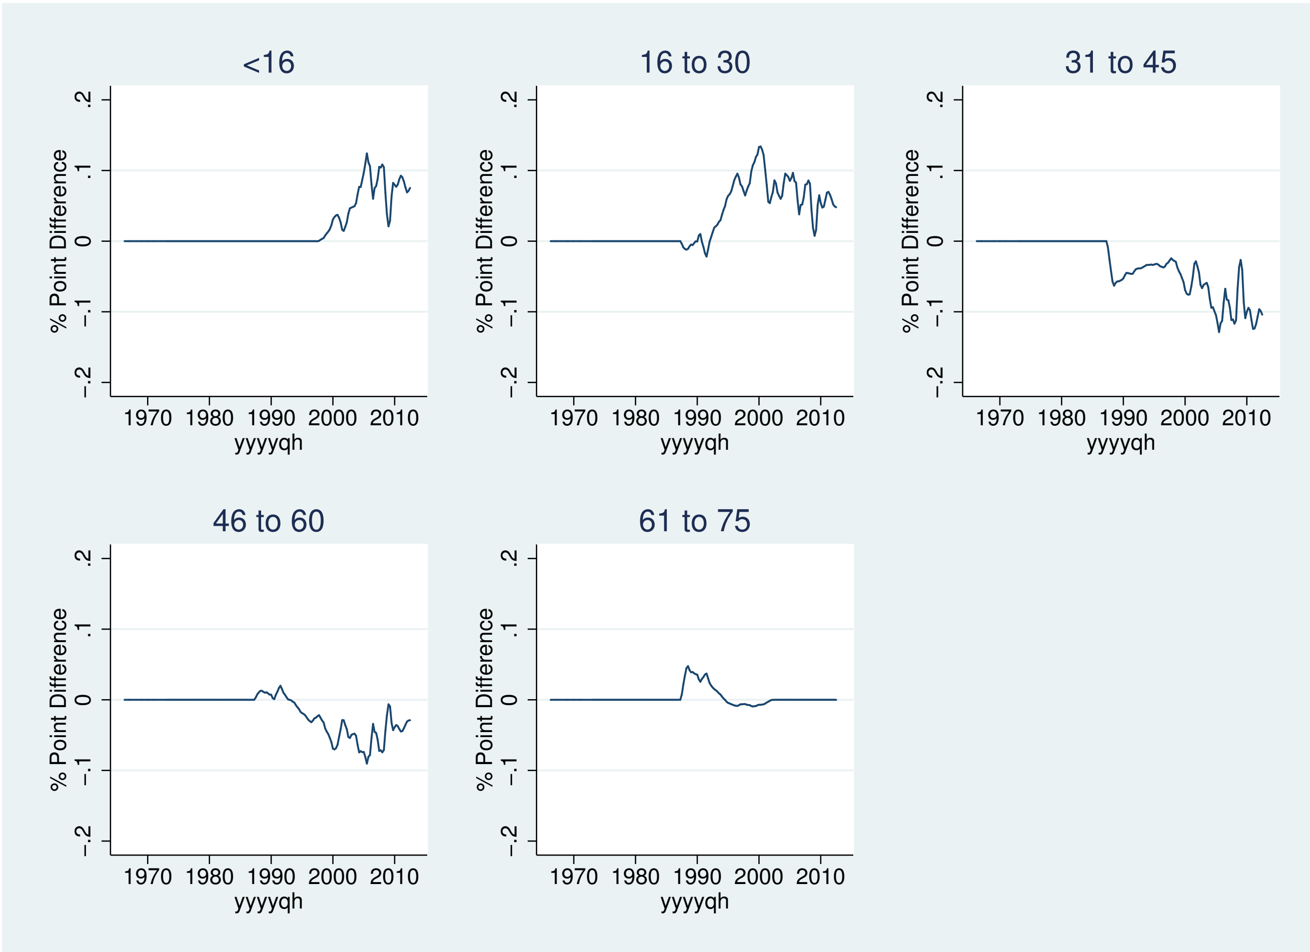

Figure 7: **Change in Aggregate Inflation Expectations with Social Learning Due to Demographic Shock at 1988**: Change is shown as difference between aggregate inflation expectations with demographic shock and actual expectations, both predicted by learning-from-experience with social learning. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The aggregate inflation expectation is simply the average predicted expectation from learning-from-experience with social learning. Individual titles correspond to age group removed at shock.

Figure 8: **Change in Aggregate Perceived Autocorrelation of Inflation Due to Demographic Shock at 1988**: Change is shown as difference between aggregate perceived autocorrelation of inflation due to demographic shock and actual perceived autocorrelation, both predicted by learning-from-experience. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The aggregate perceived autocorrelation of inflation is simply the average perceived autocorrelation at t from learning-from-experience. Titles correspond to age group removed at shock.

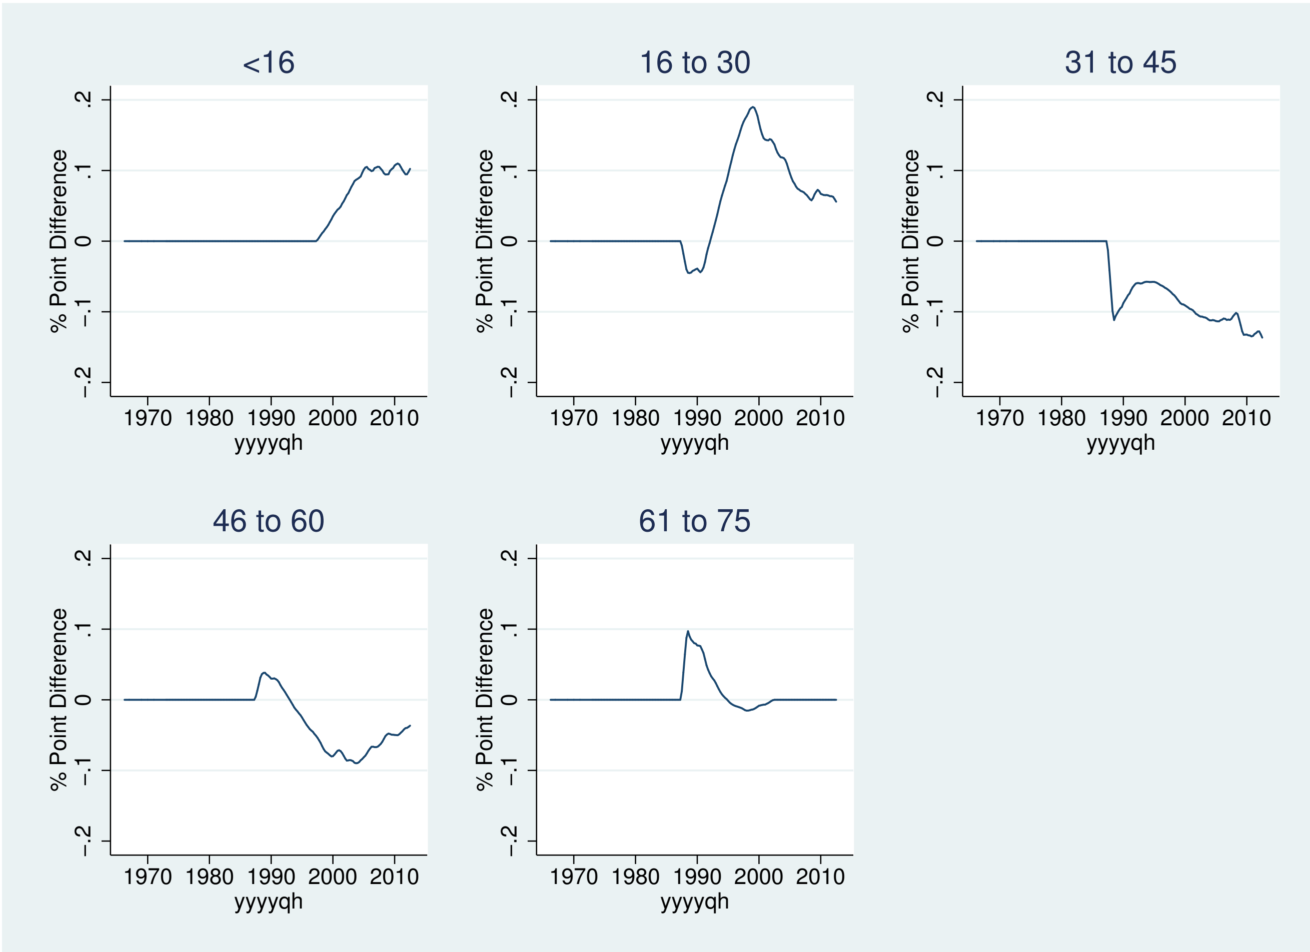

Figure 9: **Change in Perceived Mean Inflation Due to Demographic Shock at 1988**: Change is shown as difference between aggregate perceived mean of inflation from demographic shock and actual perceived mean, both predicted by learning-from-experience. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The perceived mean of inflation is simply the average perceived mean at t from learning-from-experience. Individual titles correspond to age group removed at shock.

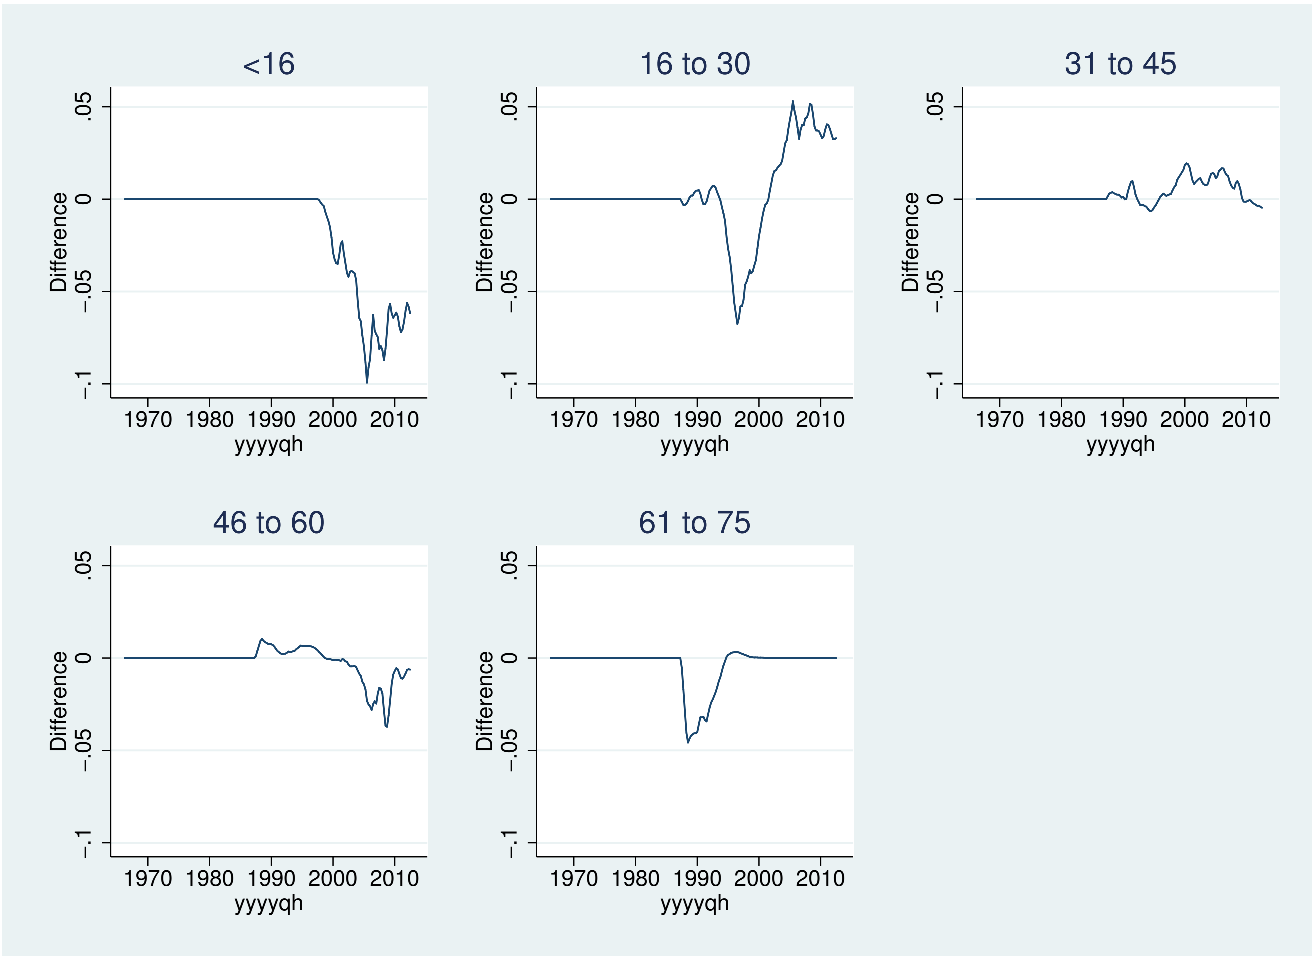

Figure 10: **Change in Disagreement about Inflation Expectations due to Demographic Shock**: Change is shown as difference between disagreement in inflation expectations due to demographic shock and actual disagreement, both predicted by learning-from-experience. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The perceived mean of inflation is simply the average perceived mean at t from learning-from-experience. Individual titles correspond to age group removed at shock.

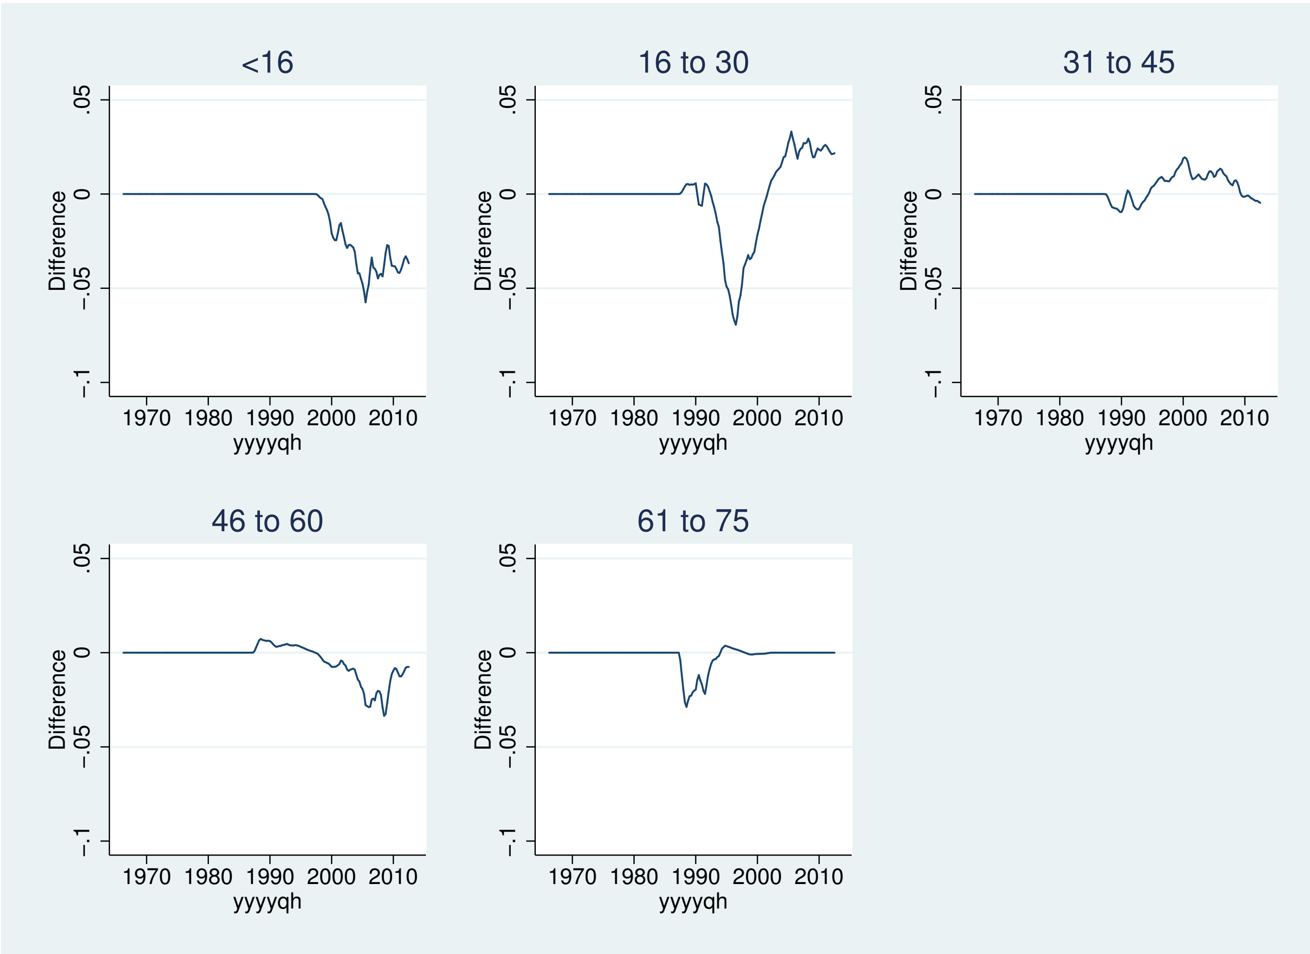

Figure 11: **Change in Disagreement About Inflation Expectations Due to Demographic Shock with Social Learning**: Change is shown as difference between disagreement in inflation expectations due to demographic shock and actual disagreement, both predicted by learning-from-experience with social learning. The shock occurs at 1988. To remove high-frequency variation, differences are shown as four-quarter moving averages. The perceived mean of inflation is simply the average perceived mean at t from learning-from-experience. Individual titles correspond to age group removed at shock.

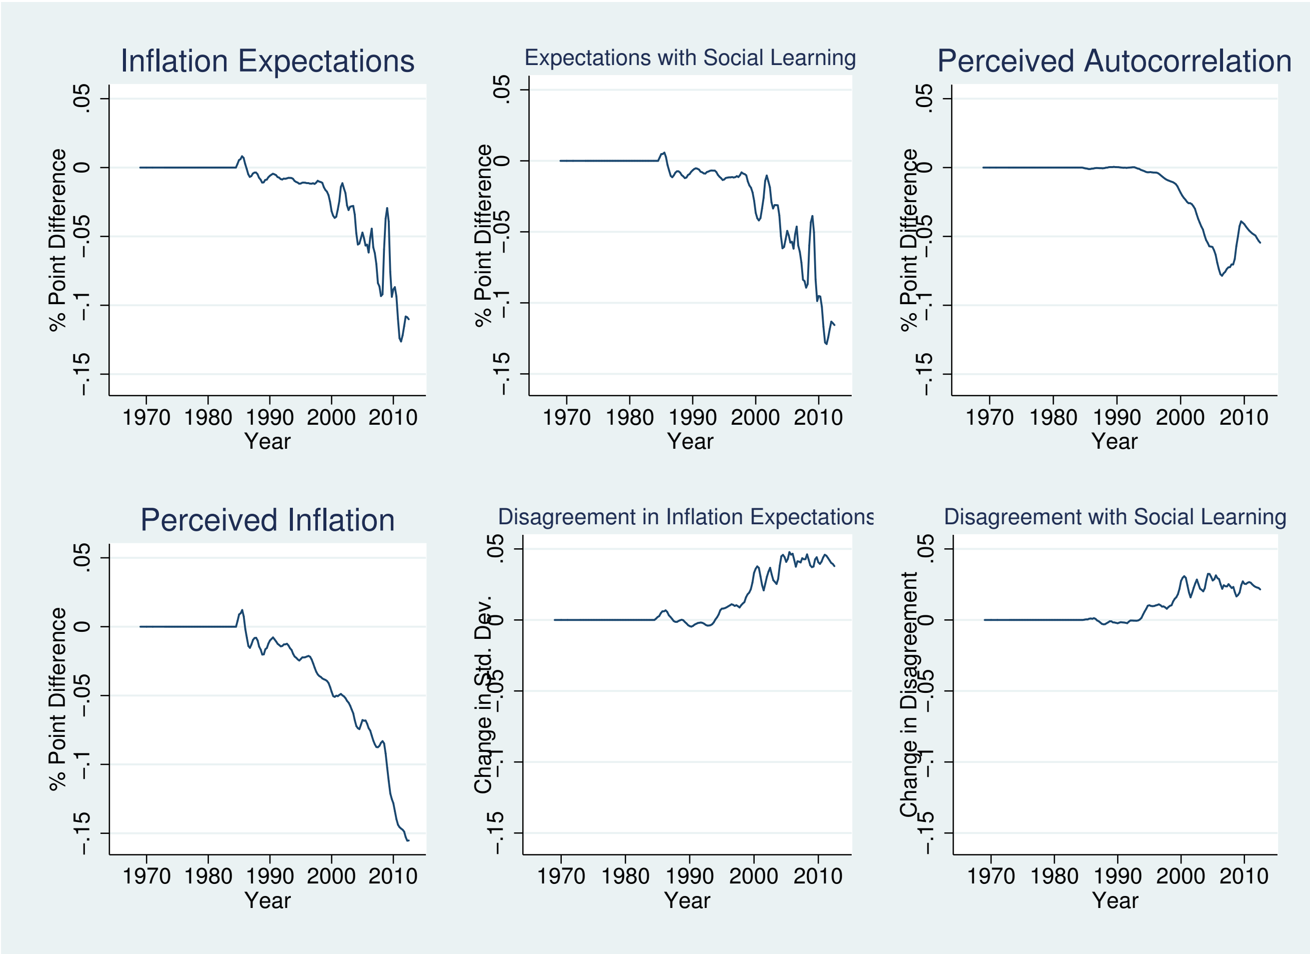

Figure 12: **Change in Aggregate Properties of Inflation Expectations with Non-Aging Demographic Distribution After 1985**: Change is shown as difference between simulated series and actual series. Inflation expectations are average inflation expectations at time t predicted by learning-from-experience. Expectation with social learning is average inflation expectations at t predicted by learning-from-experience with social learning. Perceived autocorrelation is mean perceived autocorrelation of inflation in AR(1). Disagreement in inflation expectations is the standard deviation of expectations predicted by learning-from-experience. Disagreement with social learning is the standard deviation of expectations predicted by learning-from-experience with social learning. To remove high-frequency variation, differences are shown as four-quarter moving averages.

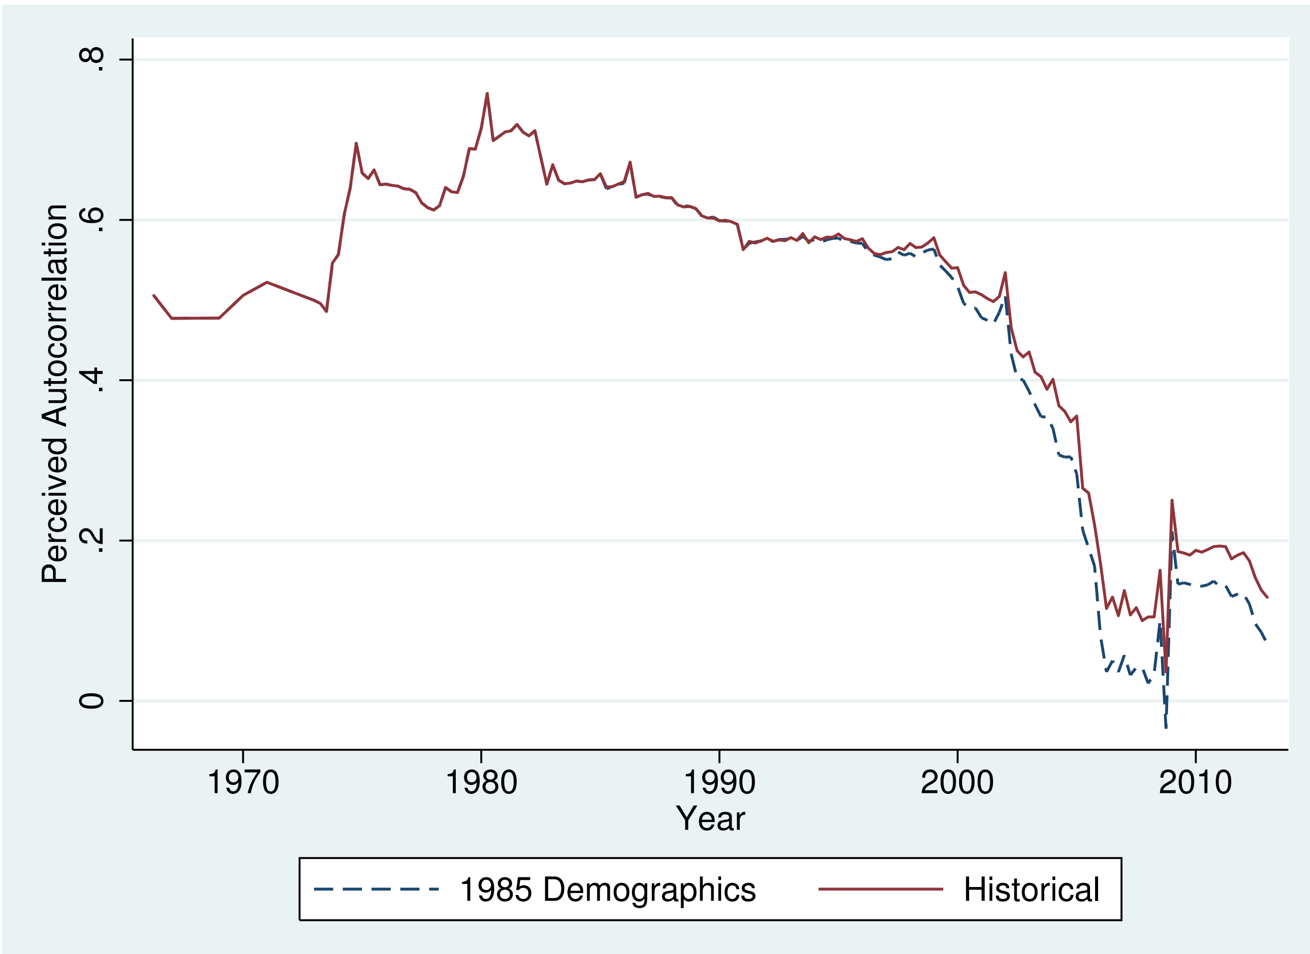

Figure 13: **Comparison of Perceived Autocorrelation**: Both series are the means at time t of perceived autocorrelation of inflation predicted by learning-from-experience model. The solid line is actual perceived autocorrelation while the dotted line is perceived autocorrelation with non-aging demographic distribution at 1985.

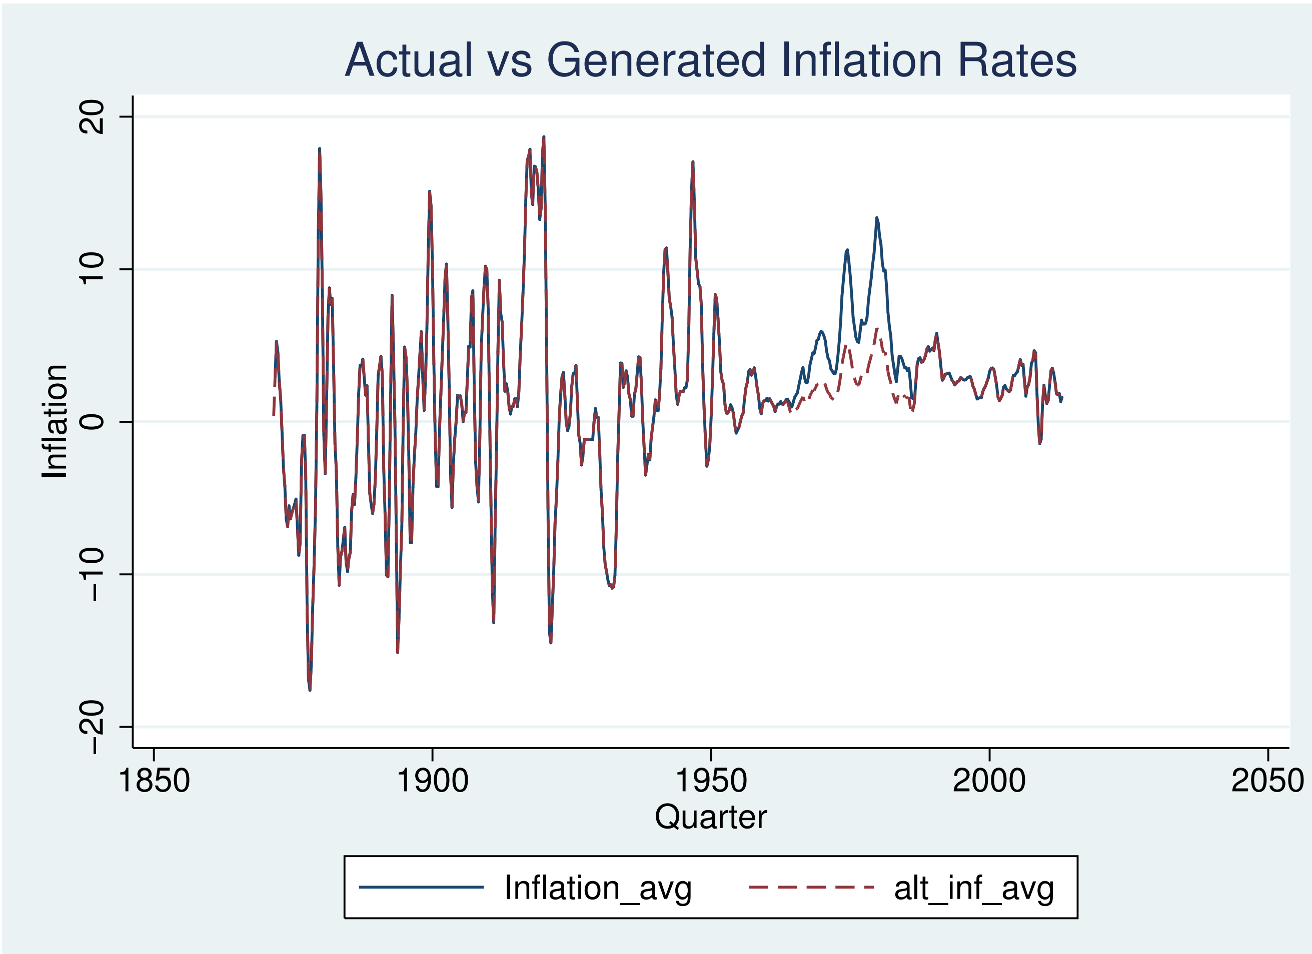

Figure 14: **Actual Versus Generated Inflation Rates**: Solid line is actual annualized CPI inflation rates. Dashed line is alternative, annualized CPI inflation rates rescaled from 1965 to 1985. The alternative inflation series was created by rescaling inflation in 1965-1985 so that average inflation in that period equals the average inflation from 1985-2013.

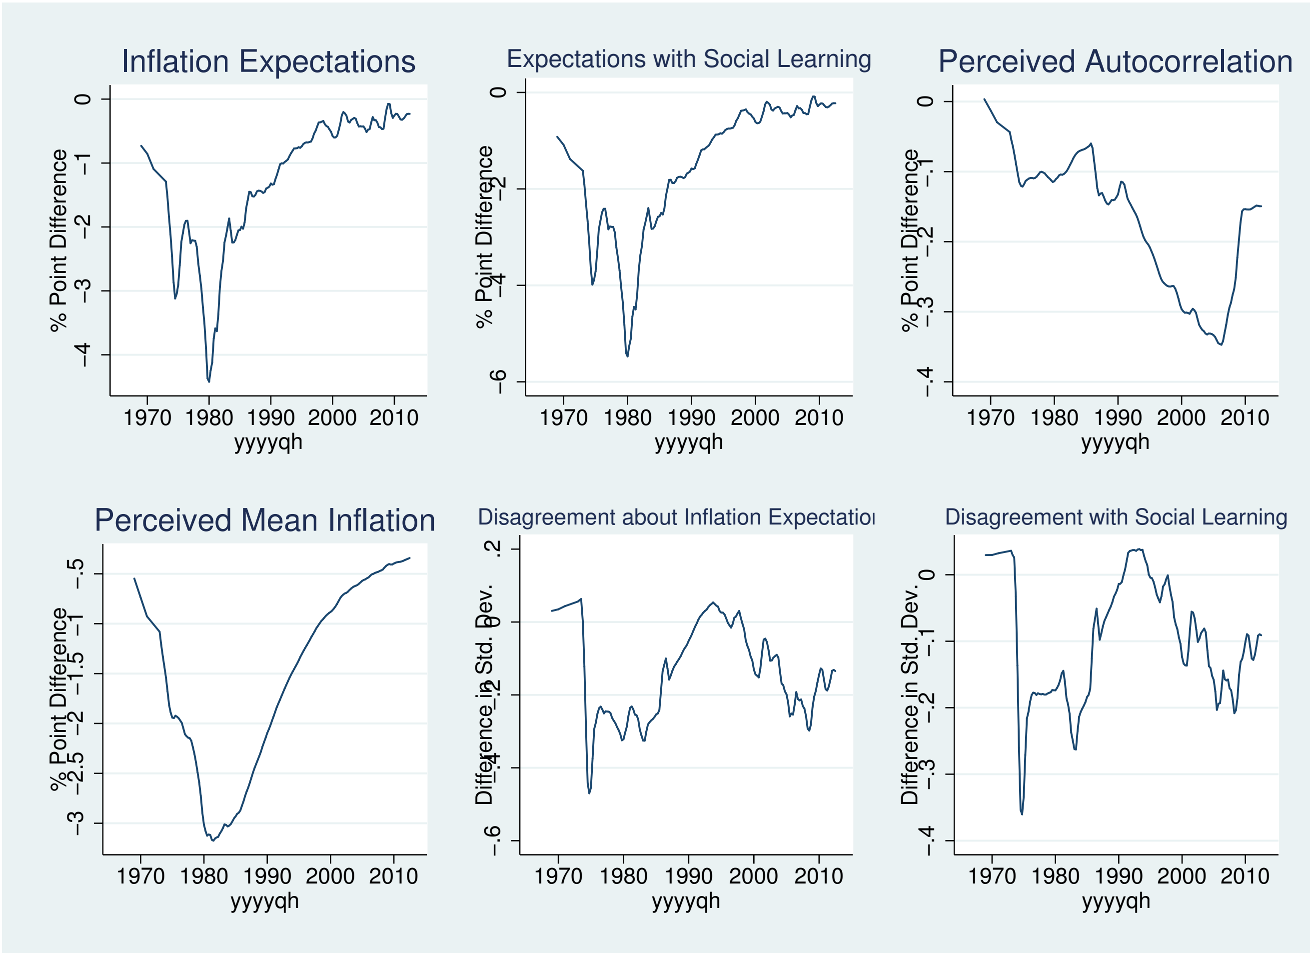

Figure 15: **Change in Aggregate Properties of Inflation Expectations with Alternate Inflation Series**: The alternative inflation series is produced by rescaling inflation from 1965-1985 so that the average in that period is the same as average inflation after 1985. Change is shown as difference between simulated series and actual series. Inflation expectations are average inflation expectations at time t predicted by learning-from-experience. Expectation with social learning is average inflation expectations at t predicted by learning-from-experience with social learning. Perceived autocorrelation is mean perceived autocorrelation of inflation in AR(1). Disagreement in inflation expectations is the standard deviation of expectations predicted by learning-from-experience. Disagreement with social learning is the standard deviation of expectations predicted by learning-from-experience with social learning. To remove high-frequency variation, differences are shown as four-quarter moving averages.

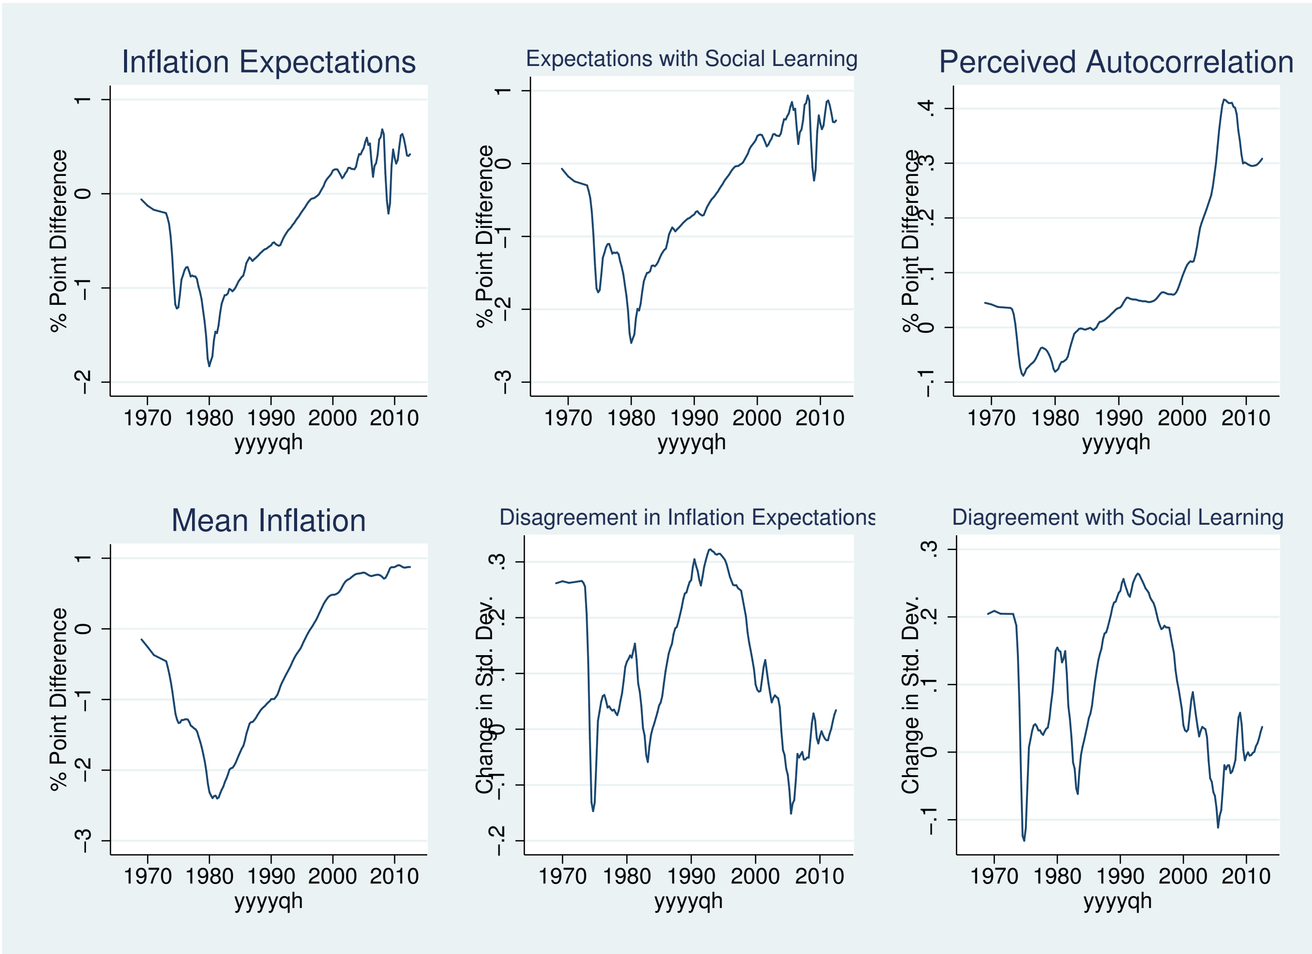

Figure 16: **Change in Aggregate Properties of Inflation Expectations without Downweighting Past Experiences**: Change is shown as difference between simulated series and actual series. Inflation expectations are average inflation expectations at time t predicted by learning-from-experience. Expectation with social learning is average inflation expectations at t predicted by learning-from-experience with social learning. Perceived autocorrelation is average at time t of perceived autocorrelation of inflation in AR(1). Disagreement in inflation expectations is the standard deviation of expectations predicted by learning-from-experience. Disagreement with social learning is the standard deviation of expectations predicted by learning-from-experience with social learning. To remove high-frequency variation, differences are shown as four-quarter moving averages.

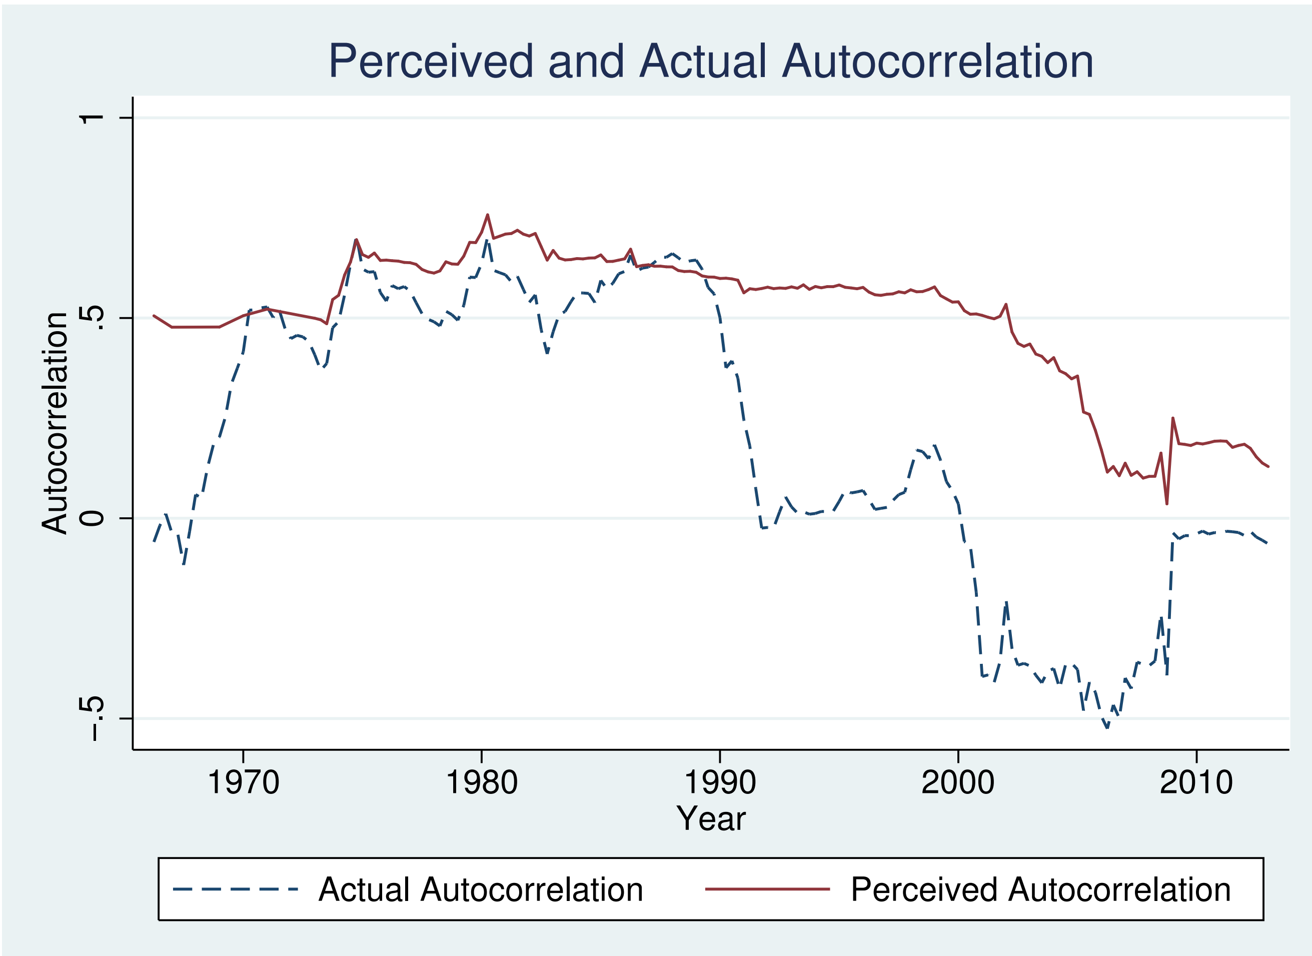

Figure 17: **Comparison of Perceived and Actual Autocorrelations**: The dotted line is actual autocorrelation of inflation, as estimated using an AR(1) model fit to inflation rates with a 10-year rolling window. The solid line is perceived autocorrelation as predicted by learning-from-experience.