

Madeira, Carlos; Zafar, Basit

Working Paper

Heterogeneous inflation expectations, learning, and market outcomes

Staff Report, No. 536

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Madeira, Carlos; Zafar, Basit (2012) : Heterogeneous inflation expectations, learning, and market outcomes, Staff Report, No. 536, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60826>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Heterogeneous Inflation Expectations, Learning,
and Market Outcomes

Carlos Madeira
Basit Zafar

Staff Report no. 536
January 2012

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Heterogeneous Inflation Expectations, Learning, and Market Outcomes

Carlos Madeira and Basit Zafar

Federal Reserve Bank of New York Staff Reports, no. 536

January 2012

JEL classification: C53, C81, E31, E37, D83, D84

Abstract

Using the panel component of the Michigan Survey of Consumers, we show that individuals, in particular women and ethnic minorities, are highly heterogeneous in their expectations of inflation. We estimate a model of inflation expectations based on learning from experience that also allows for heterogeneity in both private information and updating. Our model vastly outperforms existing models of inflation expectations in explaining the heterogeneity in the data. We find that women, ethnic minorities, and less educated agents have a higher degree of heterogeneity in their private information, and are also slower to update their expectations. In addition, we show that personal income forecasts are positively related to subjective inflation expectations. During the 2000s, consumers believe inflation to be more persistent in the short term, but temporary fluctuations in inflation have less effect on income and long-term inflation expectations. Finally, we find evidence that sticky expectations and the heterogeneity of new information received by consumers generate higher mark-ups and inflation.

Key words: inflation expectations, imperfect information, heterogeneous expectations, learning, sticky information

Madeira: Central Bank of Chile (e-mail: carlosmadeira2009@u.northwestern.edu). Zafar: Federal Reserve Bank of New York (e-mail: basit.zafar@ny.frb.org). The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System..

1 Introduction

Inflation expectations of agents are the main focus of modern macroeconomic models as well as monetary policy (Sims, 2009). Today, several surveys collect direct measures of inflation expectations of consumers, professional forecasters and experts.¹ However, controlling inflation expectations requires one to understand how they are formed. Interpretation of the data and policy outcomes is greatly affected by whether models assume rational expectations (Lucas, 1972) or some sort of bounded rationality (Mankiw and Reis, 2002; Carroll, 2003; Branch, 2004, 2007). For instance, the effects of monetary policy can be less powerful when agents doubt the commitment of central banks to fight future inflation (Orphanides and Williams, 2003). More recently, macroeconomic studies that use time series of median expectations such as Roberts (1997), Adam and Padula (2011) and Eusepi and Del Negro (2011) find that sticky prices are explained by non-rationalities in price expectations. Survey work has shown that individuals are not fully informed about future outcomes, and that there is substantial divergence among individuals' beliefs (Mankiw, Reis, and Wolfers, 2003). There is little work, however, trying to explain the heterogeneity of individuals' expectations, how they learn from new information, and the implications of this heterogeneity for market outcomes. This paper fills some of that gap.

We propose a model where agents provide inflation forecasts based on observable information – such as the previous inflation rates – and unobservable information, and study how they update their beliefs. Our model improves upon previous work by including idiosyncratic heterogeneity and dynamic updating of each agent's inflation expectations. For this purpose, we use the panel component of the Reuters/University of Michigan Survey of Consumers (1978-2009). Previous studies have mostly forgotten about the panel dimension of survey expectations (Keane and Runkle, 1990; Souleles, 2004; and Anderson, 2008, are exceptions). This complicates the interpretation of previous work in terms of learning, since only the aggregate evolution of beliefs is analyzed, while the actual updating of individuals is not studied. Instead of assuming a model of rational expectations or adaptive learning, our updating model is based on the assertion that individuals learn from their experiences. This is similar to the *learning-from-experience* model of Malmendier and Nagel (2011a,b), where agents are influenced by the macroeconomic data they experience during

¹In the U.S., such surveys include the monthly Reuters/University of Michigan Survey of Consumers, the Livingston Survey, the Conference Board's Consumer Confidence Survey and the Survey of Professional Forecasters. In addition, several countries, including the Euro area, now collect such survey data.

their lifetime. This learning model is motivated by empirical work that shows that individuals' risk-taking choices, portfolio holdings and expectations are correlated with their macroeconomic experiences (Vissing-Jorgensen, 2003; Greenwood and Nagel, 2009; Malmendier and Nagel, 2011b). In reduced-form regressions, we do find empirical support for this kind of learning. We find, for example, that during the 1980s – a period of high inflation – young respondents provide one-year ahead inflation expectations that are about 1.7 percentage points higher than those of older respondents; this is what one would expect to see if in fact expectations were based on experiences.

Inclusion of heterogeneous life experiences of inflation allows us to measure both heterogeneity of expectations and the effect of new information on agents' predictions. It is important to note that estimates from aggregate time series are biased, when individuals have different information sets (Keane and Runkle, 1990). Our model estimates show that individuals are highly heterogeneous in their expectations of inflation and differ in how quickly they update their information. In particular, women, ethnic minorities, and less educated agents are slower to update their expectations, giving a larger focus to previous life experience rather than to most recent events. These groups are also less likely to change their idiosyncratic private beliefs in the following semesters, and they also have a higher degree of heterogeneity in their private beliefs. The same demographic groups – women and less educated agents – have been found in the literature to report higher inflation expectations and to be less informed about objective measures of inflation.² A possible explanation for our finding could be that these demographic groups have different inflation experiences. However, we get qualitatively similar results if we use group-specific inflation series that take into account demographic differences in expenditures and prices, suggesting that the differences in updating and learning that we find are not primarily a result of different inflation experiences, but instead indicative of different information processing rules.³ Furthermore, we show that these features of heterogeneity across demographic groups and their dynamic updating explain a much larger component of agents' expectations than previous alternative explanations, such as rational or adaptive expectations. We also re-estimate our updating model, allowing for the coefficients to vary over time in order to control for changes in the macro-environment. We find that, over the

²See Bryan and Venkatu (2001), Bruine de Bruin et al. (2010), and Armantier et al. (2011).

³This is also in line with studies which argue that different experiences are unlikely to explain observed differences in inflation expectations (McGranahan and Paulson, 2006; Hobijn et al., 2009).

years, the heterogeneity of expectations for both short-term and long-term inflation has decreased substantially, which is consistent with studies that find inflation has become easier to predict in recent times (Stock and Watson, 2007).

The Michigan survey also collects data on subjective income growth rates of respondents. We use an updating model similar to that for updating of inflation expectations to model the revision of income growth expectations. We find that young and more educated households have greater heterogeneity of personal income growth forecasts, consistent with the larger dispersion in their observable earnings (Attanasio et al., 1999; Katz and Autor, 1999). As with inflation expectations, we find a decrease in the heterogeneity of personal income growth forecasts over time, consistent with the evidence that the change in the earnings' structure was largely predictable to individual agents (Primiceri and van Rens, 2009). Policy-makers are always concerned about the vicious cycle of inflation expectations feeding into wage demands. We do find that, over the period 1978-1985, households incorporated a great degree of their inflation forecasts in their income growth expectations. However, this tendency decreased after 1985 and in the 2000's there is no longer an effect of inflation expectations on income growth forecasts.

We also relate heterogeneity in inflation expectations to realization of future inflation. Models where consumers search for "best-bargains" show that heterogeneity of price expectations leads firms to charge higher prices (Benabou and Gertner, 1993). Heterogeneous information observed by consumers and slow updating of new information increases the monopoly power of firms, since there will be a higher mass of consumers accepting to buy at high prices. We test this hypothesis by analyzing the statistical relation between inflation and the inter-decile range of inflation expectations in the Michigan survey data. We also test whether the estimated measure of sticky expectations obtained from our model of price updating could cause inflation, since firms may be slower to change prices when consumers are slow to update their information set. We show that the heterogeneity of new idiosyncratic information and "sticky expectations" by consumers are statistically significant in explaining the next quarter's inflation, which lends empirical support to the negative effects of uncertainty and heterogeneous information on firm competition and price mark-ups.

Previous literature on inflation expectations has studied possible explanations for the heterogeneity of agents' beliefs. Souleles (2004) and Anderson (2008) find that females, racial minorities, and

lower income persons have larger forecast errors than average. Souleles (2004) exploits the panel dimension of the Michigan Survey of Consumers, and shows that expectations of consumer sentiment are biased and that forecast errors tend to be systematically heterogeneous and correlated with demographic characteristics. This systematic heterogeneity in forecast errors turns out to be important in explaining the excess sensitivity of household expenditure to sentiment data. Anderson (2008), as in the current paper, focuses on inflation expectations of households and exploits the panel dimension of the Michigan Survey of Consumers to study learning and accuracy of households' inflation expectations. While the goal of this paper is somewhat similar to ours, the two use different updating rules. Anderson (2008) uses an adaptive learning model and focuses on forecast accuracy. We use heterogeneous lifetime experiences (of inflation) in individuals' updating process and estimate a structural model of belief-updating, and link the heterogeneity in updating to market outcomes. Other studies look at the (cross-sectional) heterogeneity of inflation forecasts and explain it as a result of different lifetime inflation experiences (Malmendier and Nagel, 2011a), switching between different prediction rules (Branch, 2004, 2007), or rational inattention (Mankiw, Reis, and Wolfers, 2003). We show that while the fit of our model is similar to adaptive and rational expectation models when explaining the conditional mean of inflation expectations, our model vastly outperforms these models when trying to explain the other conditional moments of the expectations distribution, i.e., the heterogeneity in expectations and the individual updating between different time periods.

This paper is divided as follows. Section 2 discusses our model of expectations formation and outlines how we deal with both observable information and unobservable idiosyncratic beliefs. Section 3 summarizes the Michigan survey data and the historical inflation series we use in our work, and shows descriptive demographic patterns in the data. Section 4 discusses the results of the updating model described in Section 2 and analyzes differences in learning rules across demographic groups. In Section 5 we relate the effects of heterogeneous expectations and slow learning about new information to the pricing evolution of consumer goods, using a search model framework. Finally, Section 6 concludes the paper with a summary of the results and proposals for future research.

2 The model of expectation updating

2.1 Basic model

We denote $\pi_{t',i}^p(t)$ as the prediction for the annualized inflation observed at quarter $t' \geq t$ that agent i makes in quarter t . We assume agent i of cohort s learns about future inflation by using previous inflation experiences lived in his lifetime, $\pi_{t,s}^{life}(\theta)$ (as in Malmendier and Nagel, 2011a), and other public information available to everyone, $z_{t'}(t)$. Lifetime inflation experience can be thought as a weighted average of observed inflation rates in the previous life of the agent, with more recent experiences slowly adding to older ones. In this case public information can include all contemporary information generally known to the public, such as the last reported inflation rate. To simplify, let us assume a linear updating model for future inflation expectations based on $\pi_{t,s}^{life}(\theta)$ and $z_{t'}(t)$:

$$2.1) \pi_{t',i}^p(t) = \beta \pi_{t,s}^{life}(\theta) + z_{t'}(t) + \eta_{t',i}^p(t), \text{ with } t' \geq t,$$

where β denotes the importance attached to lifetime inflation experiences, and $\eta_{t',i}^p(t)$ is idiosyncratic private information. That is, agents' expectations are assumed to depend on lifetime inflation experience and private information. Note that while inflation is an aggregate event, there could be several sources of private information affecting individual agents' predictions. One, could be, for instance, how frequently agents read economic and financial news, if at all. Other examples could be how often the household head goes to the supermarket and notices price changes. Other factors can also be thought as privately observed. Poorer households are more likely to be aware of rent and food price inflation, while richer households should arguably be more aware of prices of durable and luxury goods. Older households could be more aware of health costs. Since the sources of private information differ markedly across households of different background, it is reasonable to assume that $\eta_{t',i}^p(t)$ is heteroskedastic. Our model assumes that the heteroskedasticity of private information changes both across demographic groups and time, since it is possible that households change their information sources over time.

We assume the idiosyncratic private information term, $\eta_{t',i}^p(t)$, follows an AR(1) process:

$$2.2) \eta_{t',i}^p(t) = \lambda \eta_{t',i}^p(t-1) + u_{t',i}(t), \text{ with } \lambda \leq 1.$$

The term λ informs us how slow individuals are to update their idiosyncratic opinions on inflation rates. One can interpret this term as a combination of both the innovation process in the information sources

of the agent and of the behavioral speed with which the agent actually updates his predictions.

It is assumed $u_{t',i}(t)$ is normally distributed $(u_{t',i}(t)) \sim N(0; \sigma_{u_i}^2(t' - t))$. $\sigma_{u_i}^2(t' - t)$ can be interpreted as a measure of the unexplained heterogeneity or dispersion in agents' beliefs about future inflation. It can also be denoted as "disagreement" in opinions, as in Mankiw, Reis, and Wolfers (2003) and Rich and Tracy (2006).

The indicator $(t' - t)$ in the variance term $\sigma_{u_i}^2(t' - t)$ is written to imply that the variance in the dispersion of opinions depends on the duration between the time at which the forecast is made and the event. The Michigan survey data collects information for two time horizons, 1 year and 5-10 years after the forecast. Therefore our model allows us to learn about the dispersion in expectations at different time horizons.

Lifetime inflation is updated by the agent as a weighted-average of observed inflation in previous quarters:

$$2.3.1) \pi_{t=s,s}^{life}(\theta) = \pi_{t-1}$$

$$2.3.2) \pi_{t,s}^{life}(\theta) = \pi_{t-1,s}^{life} + \gamma_{t-s}(\pi_t - \pi_{t-1,s}^{life}), t > s,$$

$$\text{with } \gamma_{t-s} = \frac{\theta}{t-s}.$$

In this model, θ denotes how rapidly agents include new information in their estimates of lifetime inflation, while β denotes how important lifetime inflation is in agents' expectations. A positive fixed θ would indicate that the gain sequence of learning is decreasing in age. This would imply that the learning from recently observed data is quicker and more volatile for young consumers. This formulation is consistent with behavioral results that imply young agents are more overconfident in the reliability of recent information relative to older agents (Barber and Odean, 2001; Vissing-Jorgensen, 2003; Greenwood and Nagel, 2009; Malmendier and Nagel, 2011a).

The model of expectation updating is therefore quite parsimonious and depends only on a vector of four parameters: $\varpi \equiv \{\beta, \theta, \lambda, \sigma_{u_i}^2(t' - t)\}$.

2.2 Estimation

To study the heterogeneity and updating of inflation expectations we use the panel component of the Michigan Survey of Consumer Expectations. In this survey respondents give, for two consecutive semesters, their subjective expectations of inflation in the next 12 months and inflation for the next 5-10 years.

The Michigan data has several demographic characteristics which allow us to measure observable heterogeneity in expectations updating across different groups. The empirical model therefore includes a vector of individual demographic characteristics, x_i , which will include income, education, race, and gender. Therefore we consider heterogeneity in average expectations by allowing the model to differ across x_i , i.e.: $\varpi = \varpi(x_i)$.

To estimate this model we assume a parametric form for the public information observed by rational agents, $z_{t'}(t)$:

$$3.1) \ z_{t'}(t) = \gamma[\pi_{t'}, \pi_{t-1}, d_t],$$

where π_{t-1} is the inflation observed in the previous quarter and $\pi_{t'}$ is the actual inflation rate observed in the future (which could be partially known by a signal observed by forward-looking rational agents). d_t is a dummy variable for the half-decade period in which the inflation forecast is being made. Therefore, $z_{t'}(t)$ allows us to measure how much agents approach the ideal rational agent. If agents are perfectly rational then the coefficient of $\pi_{t'}$ should be close to one, while the others should be close to zero.

A problem in the empirical model is that our observed panel only includes individuals for two periods distanced by six months. Therefore one cannot deal with the first period error term as an AR(1). We solve this by specifying the first period error term to be a purely idiosyncratic term $\eta_{t',i}^p(t) = u_{t',i}^1(t)$ and using the AR(1) process, $\eta_{t',i}^p(t) = \lambda\eta_{t',i}^p(t-1) + u_{t',i}^2(t)$, in the second period. This gives us two variance terms to estimate, $\sigma_{u_i^1}^2(t' - t)$ and $\sigma_{u_i^2}^2(t' - t)^4$. We also consider parametric forms for $\varpi = \varpi(x_{i,t})$:

$$3.2) \ \theta \text{ is a constant scalar,}$$

⁴ Another option is to impose $\sigma_{u_i^1}^2(t' - t) = \frac{\sigma_{u_i^2}^2(t' - t) \times (1 - \lambda^{2 \times a(i)})}{1 - \lambda^2}$, which is the steady-state variance for someone with $a(i)$ quarters of age. Assuming this form has no qualitative differences in our results. We prefer to report the two variance terms for each panel period due to its simplicity.

$$3.3) \beta = \exp(\alpha_\beta x_{i,t}),$$

$$3.4) \lambda = \frac{\exp(\alpha_\lambda x_{i,t})}{1 + \exp(\alpha_\lambda x_{i,t})},$$

$$3.5) \sigma_{u_i}^2(t' - t) = \exp(\alpha_{u_i,t} x_{i,t}) \text{ for } a = 1, 2,$$

where $x_{i,t} \equiv \{Female, Asian, Black, Hispanic, Young, Middle-aged, low-income, middle-income, years of education, half-decade, \pi_{t-1}, |\pi_{t-1} - \pi_{t-2}|\}$. For the case of the heterogeneity in inflation expectations, we also allow $\sigma_{u_i}^2(t' - t)$ to depend on the inflation change observed in the previous period, since previous studies find that individuals are more uncertain in periods of high and volatile inflation rates (Rich and Tracy, 2006).

3 Data

The Michigan Survey of Consumer Expectations has been conducted monthly by the University of Michigan between 1978 to the present day, based on telephone interviews of a sample of approximately 500 respondents, which is representative of the US demographic population. The survey incorporates a rotating sample design, where 40% of the monthly sample are re-contacts from six-months ago, and the remaining 60% are new respondents. Although this survey has been implemented since 1953, the panel data are only available after 1978. Rather surprisingly, few studies have exploited this feature of the Survey of Consumer Expectations; exceptions include Souleles (2004) and Anderson (2008).

In this survey respondents provide their subjective expectations of inflation in the next 12 months and inflation for the next 5-10 years. 12 month-ahead inflation expectations are elicited as follows:

During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?

followed by:

By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?

In addition, households are asked their forecast for their personal income growth over the next year:

During the next 12 months, do you expect your income to be higher or lower than during the past year?

followed by:

By about what percent do you expect your income to (increase/decrease) during the next 12 months?

Therefore, between the period 1978 to 2009, the survey allows us to measure the expectations of more than 85,000 individuals at two different points in time, with all households being re-interviewed after 6 months.

We also use long-term historical data on the Consumer Price Index (CPI) collected by Robert Shiller on his website. Based on this long time series we calculate the quarterly inflation rates observed in the US. Then we estimate the average lifetime inflation experienced by each birth cohort using a quarterly update interval for several values of θ . The updating rule considered in 2.3.2) is highly non-linear in the inflation rates of previous periods and the age of the respondents, requiring the algorithm to go over all the life inflation rates of each agent and compute a different weight for each period. To reduce the computation burden of this exercise we computed the life inflation series at 40 different values of θ and then used a linear interpolation rule to compute the life inflation at intermediate values.⁵

According to Judd (1998), approximating a function through linear interpolation between points gives consistent and shape-preserving estimates of the true function as the number of evaluation points increases to infinity. It is fundamentally required that the number of evaluation points exceeds the number of unknown parameters of the objective function. Since we use 40 points to approximate a function of one unknown parameter, it is reasonable to expect that the approximation error is small. There is a correlation of 99.99% between adjacent series of θ around 2.2 and 2.3, which represent the most likely values for this parameter. The correlation between life inflation at adjacent points is high, therefore there is little measurement error involved in this approximation.

⁵We chose $\{0, 0.5, 1, 1.25, 1.5, 1.65, 1.8, 1.9, 2, 2.05, (2.1: 0.025: 2.2), (2.2: 0.01: 2.3), (2.3: 0.025: 2.4), 2.45, 2.5, 2.6, 2.7, 2.85, 3, 3.25, 3.5, 4, 4.5, 5\}$, as the exact values of θ ., where $a :$ denotes an arithmetic progression in steps of a .

3.1 Descriptive analysis

Before estimation of the updating model described in Section 2, we show some descriptive patterns in the data. We retain the full sample for this purpose and do not restrict to the respondents who are re-surveyed. Figure 1 shows the median one-year ahead inflation expectations in the Michigan survey. Compared to realized one-year ahead inflation the median underestimates the realized inflation up to the early 1990s. After that the median expectation slightly overestimates the realized inflation. The visual depiction of the two series suggests that inflation expectations lag behind realized inflation, i.e., they seem to be anchored to realized inflation in the survey year. The figure also reports the 25th and 75th percentiles of the expectations distributions. The inter-quantile range – a measure of respondents’ disagreement – is quite large. Though the range is larger in periods of high inflation, the inter-quantile range is about 5% even in periods of low inflation. This indicates substantial heterogeneity in point forecasts of survey respondents.

To shed light on the correlates of inflation expectations, we regress the respondents’ point forecast of one-year ahead inflation onto a set of demographic variables in Table 1. The first and sixth columns show that female, Black, Hispanic, young, the less wealthy, and less educated respondents report higher expectations. These demographic differences in inflation perceptions are mostly similar to those documented by other studies (Bryan and Venkatu, 2001; Bruine de Bruin et al., 2010).

In columns (2) and (7) of Table 1, we include the annual rate of inflation prevalent at the time of the survey as well as the realized one-year ahead inflation (the quantity for which the respondent provides a forecast) as regressors along with the demographic variables. The coefficient on current inflation is about 0.4, while the magnitude of the coefficient on one-year ahead inflation is less than 0.1. This suggests that respondents are closer to adaptive expectations than to rational expectations.

Since both the inflation target as well as the US economy has changed since the early 80s, in columns (3)-(5) of the table, we restrict the sample to each of the decades. What stands out is that the disparity in forecasts by demographic variables is not constant over time. For example, the coefficient on females in 2000s is nearly half that in the earlier decades. Similarly, the disparity by race and ethnicity is smaller in the 2000s. The most notable estimate is that of the younger respondents. In the 1980s the mean estimate of the younger agents was 1.69 percentage points

higher than that of the older respondents. This disparity had decreased to 0.768 in the 1990s, and by the 2000s the young were making predictions statistically similar to those of the old. This pattern is consistent with our learning-from-experience updating model outlined in Section 2, which posits that individuals are more strongly influenced by data realized during their lifetimes than by other historical data. Because the 80s was a period of high inflation (Figure 1), it makes sense that the young report higher point forecasts in that period, since they should rely to a greater extent on extrapolation of recent inflation data if they form expectations based on experience.

Table 2 shows the heterogeneity by observables in updating of one-year ahead inflation expectations. The first six columns regress the absolute change in point forecasts between the two surveys onto a set of demographic and other control variables. Column (1) shows that females, minorities, young, lower-income, and less-educated agents make larger absolute revisions. These are the same demographic groups that report larger inflation forecasts and therefore have more to learn in order to approach less biased expectations. Column (2) of the table includes the absolute error in the respondent’s forecast (defined as the absolute gap between the respondent’s point forecast of one-year ahead inflation and actual realized one-year ahead inflation) and the realized change in inflation between the two surveys as additional regressors. Parameter estimates on both these regressors are positive and statistically significant, i.e., respondents who have worse forecasts in the first survey tend to make larger revisions, and respondents revise their beliefs more during periods of more variable inflation. Columns (3)-(5) of the table report the OLS estimates for each of the decades separately. Over time the heterogeneity in absolute revisions by demographics has decreased substantially, but continues to be statistically different from zero.

The last two columns report the OLS estimates of a regression of the absolute error in the respondent’s point forecast for one-year ahead inflation in each of the two surveys. Column (6) shows that the demographic groups who report larger point forecasts and revise more between the two surveys - females, minorities, young, the less wealthy and the less educated – also make larger forecast errors. The last column shows that, even when interviewed the second time, error patterns by demographics look similar. We include the respondent’s absolute error in the first survey as an explanatory variable in this specification. The estimates show a positive relationship between the absolute error in the first survey and the error in the second survey, i.e., there is persistence in forecast errors of respondents. These results are consistent with Souleles (2004) and

Anderson (2008) who also find that females, racial minorities, and low income respondents make larger forecast errors than average.

Why is it that certain demographic groups report larger point forecasts, make larger forecast errors, and revise more? It could be that females, lower income individuals, less educated, young, and minorities have different actual inflation experiences and hence report larger point forecasts. Also, groups facing more volatile inflation rates could show less persistence in their inflation expectations. However, we find that this explanation is unlikely and should play a minor role. The Chicago Fed IBEX 12 month inflation series (1983-2005)⁶ takes into account the different inflation experiences of various socioeconomic and demographic groups. The series uses Consumer Expenditure Survey (CEX) data and price data produced by the Bureau of Labor Statistics to construct a group-specific inflation rate, and includes the inflation rates for 42 distinct demographic groups with a monthly frequency between January of 1983 and December of 2005, which corresponds to a time series of 324 observations for each group. McGranahan and Paulson (2006) find that lower income and lower education groups have somewhat more variable inflation than higher income and higher education groups. We estimated the correlation of the inflation rates of each one of these demographic groups with the aggregate monthly inflation series of the Bureau of Labor Statistics (BLS). We find that the correlation of each demographic group's inflation rate with the BLS inflation is above 90% for all groups during the period 1983 to 2005.⁷ Therefore it is highly unlikely that these small differences in the inflation rate experienced by each group can explain the large heterogeneity of inflation expectations observed in the data. This evidence is consistent with other studies that show different inflation experiences are unlikely to explain the large differences in agents' inflation expectations (McGranahan and Paulson, 2006; Hobijn et al., 2009). Malmendier and Nagel (2011a) also show that group specific inflation rates have little significance in explaining cohort inflation expectations, once the lifetime weighted average inflation experience of the cohorts is accounted for.

Other possible explanations for demographic differences in inflation expectations include different expectations formation and information-processing rules. More specifically, there could either be

⁶Description of the series is available at http://www.chicagofed.org/webpages/research/data/ibex/ibex_inflation.cfm.

⁷In fact the correlation of each demographic group's specific inflation rate is above 95% with the BLS inflation and with the inflation specific rates of all the other 41 demographic groups, with the single exception of the group defined as "Food-stamp recipients". "Food-stamp recipients" are the single group exhibiting a specific inflation rate with a correlation of only 92% in relation to the BLS aggregate inflation and the inflation of the other groups.

demographic differences in heterogeneity of private information, or the speed at which different groups update their inflation expectations. Heterogenous updating of expectations has implications for steady-state inflation, fiscal deficits, and asset savings, and understanding the underlying channels is important for effective monetary policy. We next explore sources of demographic differences in updating explicitly in our model.

4 Interpreting the heterogeneity of prediction rules

4.1 Short-term inflation forecasts

We estimate the model of equations 2.1) and 2.2) by Maximum Likelihood⁸, using agents' forecasts at 1 year and 5-10 year horizons. Tables 3 and 4 summarize the empirical results for inflation expectations at the 1 year horizon.

The first two columns of Table 3 show the estimated coefficients and standard errors for the mean expectations process, $z_{it}(t)$, for β , and θ for the 1-year horizon. The model clearly rejects the hypothesis of rational expectations. This can be concluded by the estimated coefficient for future one-year ahead inflation being close to zero, while the coefficient for the inflation lag being close to 0.5. Therefore households are closer to adaptive expectations than to rational expectations. From estimates of β , it is also clear that women, ethnic minorities, lower income, and less educated agents are slower to update their expectations, since they give more importance to lifetime inflation. Our estimate of the life inflation update velocity θ is 2.41, while the estimate of the importance given to lifetime inflation, β , averages 0.24 for all individuals. Malmendier and Nagel (2011a) estimate a similar model, using cohort-average inflation expectations instead of the individual forecasts, finding an estimate of 2.485 for θ and 0.479 for β . However, their estimate of β is 0.332 for the period 1984-2009. Our estimate of β shows that lifetime inflation plays a smaller role in explaining expectations than the one described by Malmendier and Nagel (2011a) for the period 1953-2009.

The estimated coefficients in the top panel of Table 4 show that women, ethnic minorities, lower income, and less educated agents have a higher degree of heterogeneity in their expectations (i.e., larger estimates of $\log(\sigma_{u_i}^2)$). The data also shows that there is a higher dispersion (or disagreement) in the inflation predictions of households in periods of higher inflation and more volatile inflation

⁸It is also possible to estimate our model by using just the conditional moments of the mean, variance, and auto-correlation of the expectations. These GMM estimates do not require the normality assumption.

(as measured by the absolute change of inflation in the previous two quarters). We also find that Asians, Hispanics, and higher educated persons are slower to update their idiosyncratic opinions on inflation.

In Panel A of Table 5 we show how the heterogeneity has evolved over the years through the dummies for each half-decade. It is evident that the heterogeneity of inflation forecasts has decreased significantly over the years, particularly since the early 80s. This result is consistent with the evidence shown by Stock and Watson (2007), who find that inflation has been easier to forecast in the last two decades. However, the heterogeneity of inflation expectations has increased since 2005, perhaps as a consequence of the greater uncertainty due to the economic crisis.

Economists often worry that the degree of persistence of inflation shocks depends on how much expectations incorporate changes in the previous inflation rates (Orphanides and Williams, 2003). Our model estimates suggest that the lagged inflation term, π_{t-1} , is the most important determinant of inflation expectations. Therefore it is interesting to see how much the effect of past inflation has changed future inflation expectations over different periods and policy regimes. Next, we re-estimate our model but with the coefficient for the lagged inflation term changing at each half-decade. The evidence in Panel B of Table 5 suggests that households, in the 2000's, condition more strongly their expectations on the previously observed inflation. Again, this result is consistent with the findings of Stock and Watson (2007), who show that the use of the last observed inflation in the economic models' forecasts has increased substantially in the last decade.

4.2 Long-term inflation forecasts

Now we comment on the results for the expectations at a 5 year-horizon. In general the results are qualitatively similar to those at the 1 year horizon. The last quarterly inflation is a positive determinant of the 5-10 year inflation expectations (last two columns of Table 3). We also find that women, ethnic minorities, and less educated agents are slower to update their expectations and attach more importance to lifetime inflation. Again, our estimate of the life inflation update velocity θ is 2.68, which differs from the value of 1.889 estimated by Malmendier and Nagel (2011a). Our estimates implies that individuals update their lifetime inflation experience more quickly in order to forecast future long-term inflation. Our value of β averages 0.29 for all individuals, which is lower than their estimate of 0.482 for the period 1953-2009. However, our estimate of β is very

similar to the value of 0.269 estimated by Malmendier and Nagel for the period 1984-2009.

The fact that agents attach a positive value to the last quarter's inflation in their long-term expectations is an important result. The coefficient of 0.14 for previous inflation in the long-term expectations is smaller than the 0.48 coefficient for the one year horizon, indicating that agents do expect recent inflation shocks to die out as time passes. However, agents expect that part of the recent shock will be persistent even after 5 years. In general, due to the high persistence in price formation, one should expect that shocks to inflation affect expectations of short-term expectations. However, if agents believe that central banks have a credible target rate for inflation, last quarter's inflation should not explain long-term inflation expectations.

We also used the mean inflation of 5-10 years in the future ($\bar{\pi}_{t+4 \times 10, t+4 \times 5}$) and 5-10 years in the past ($\bar{\pi}_{t-4 \times 10, t-4 \times 5}$) as regressors. However, neither previous long-term inflation or future long-term inflation is significant. The fact that future long term inflation rates do not affect current long-term inflation expectations implies that agents are again rejected to be rational.

Also, heterogeneity within each demographic group is qualitatively similar to the expectations of inflation at 1 year horizon (lower panel of Table 4). Again we find that women, ethnic minorities, younger persons, lower income, and less educated agents have a higher degree of heterogeneity in their expectations. The data also show that there is a higher dispersion (or disagreement) in the inflation predictions of households in periods of higher inflation. However, the persistence of long-term inflation expectations behaves differently than in the short-term. Here, we find that females, Asians and Blacks are quicker to update their idiosyncratic opinions on inflation. We also show how the heterogeneity has evolved over the years through the dummies for each half-decade (Table 5). Again, it is clear that the heterogeneity of inflation forecasts decreased significantly over the years, and that the heterogeneity of inflation expectations has increased since 2005.

It is also noticeable that lower income agents have a higher heterogeneity of expectations both at short-term (top panel of Table 4) and long-term horizons (bottom panel of Table 4). This result fits well with previous research that finds lower income agents have poorer forecasts (Vissing-Jorgensen, 2003). However, another potential explanation for this is that lower income households consume different consumption baskets and may have a higher consumption share in items, such as food, that have more volatile prices at both the local level and at different time periods (Pollak, 1980; Mankiw, Reis, and Wolfers, 2003). We test for this second possibility by using the Chicago Fed IBEX 12

month inflation rate instead of the CPI. As described earlier, the IBEX is a monthly inflation measure designed to capture the inflation experiences of specific socioeconomic and demographic groups. Using this alternate inflation series does not change our results qualitatively, indicating that different inflation experiences cannot explain our results (results not reported here; available from the authors upon request).

Again, it is also interesting to see how the effect of past inflation on long-term inflation expectations has changed over different periods and policy regimes. Next, we re-estimate the updating model allowing for the coefficient for the lagged inflation term to change at each decade and half-decade. We show the coefficients for each half-decade in Panel B of Table 5. The evidence seems to point out that American households, during the late 1980s and the 2000s, did not incorporate short-term fluctuations in the previous inflation rate in their long-term inflation expectations. This is interesting, because it suggests that people were slow to react to the credibility of the new regime imposed by Volcker in the early 1980s. It is also a sign that consumers during the 2000's trust the ability of the Federal Reserve to revert short-term inflation fluctuations over the long term.

4.3 Personal income growth forecasts

Economists are often worried that inflation expectations could affect wage demands. We explore this issue by studying how the personal income growth point forecasts of the households in the next year relate to their inflation expectations. The first two columns of Table 6 regress the subjective income growth expectation reported in the first and second surveys on various demographic variables and controls, respectively. Few notable patterns stand out. Mean income growth expectations (as indicated by the constant terms) are negative. The young and middle-age respondents report economically and statistically significant large positive (or smaller negative) income growth expectations.

The elasticity of income growth expectations with respect to inflation expectations is 0.030. While it is fairly inelastic, the coefficient is very precise, suggesting that respondents do perceive a positive link between wage fluctuations and inflation. The third column of the table reports the OLS estimates of the regression of the *change* in income growth expectations on demographics; we find that, except for the Young dummy, no demographic variable is statistically significant. Since we are interested in the magnitude of revisions, the last two columns report the estimates of the absolute change in income expectations on the various covariates. Low income and young respondents revise

their income expectations more- this could be the result of their labor market experiences being more volatile. The last column shows that absolute revisions in income expectations, while positively correlated with absolute revisions of inflation expectations, are not correlated with realized changes in inflation between the surveys. That is, respondents rely more on changes in their subjective inflation forecasts than in actual changes in inflation between surveys.

We next estimate the same model as described in section 2 for the forecasts of personal income growth (i.e., now using income growth forecasts as the dependent variable instead of inflation expectations). However, we now use the subjective inflation expectation of the household for the next year, $\pi_{t+4,i}^p(t)$, as a regressor instead of the actual inflation observed in the following period. The results are shown in Tables 7 and 8. It is noticeable that, in both periods, agents' personal income growth forecasts rely more on their subjective inflation expectations than on the past inflation (Table 7), especially in the first period forecast. This is similar to our finding in the reduced-form regressions reported in columns (1) and (2) of Table 6. This makes a strong case for central banks to contain inflation expectations, since the estimates imply that rises in inflation expectations are tied to an expected increase in wages. Our model estimates also show that males, low and middle income, and more educated households rely more on their lifetime inflation experience than others (i.e., higher estimates of β). The data also shows that there is a higher dispersion in the personal income growth forecasts during periods of higher inflation volatility, although there is lower dispersion in forecasts during periods of higher inflation (Table 8).

It is also noticeable that young, middle-aged, low and middle income, and highly educated households have a higher heterogeneity of personal income growth forecasts than the average (Table 8). This is consistent with the life-cycle evidence. Most of the heterogeneity in the trend growth of personal income happens early in life (Attanasio et al., 1999). Moreover, evidence on the inequality of work earnings shows that the highly educated workers are the ones with more intra-group heterogeneity (Katz and Autor, 1999). Also, females, young and middle-aged, middle income, and highly educated households update their expectations more quickly (Table 8). This result makes sense since young and more educated households are the ones that have more to learn about the prospects of their future jobs. In Panel A of Table 9, we show how the heterogeneity has evolved over the years through dummies for each half-decade. It is clear that the heterogeneity of

personal income growth forecasts decreased significantly over the years. This result is consistent with the evidence found by Primiceri and van Rens (2009), who find that the increases in earnings' inequality over the last 25 years was the result of predictable changes in earnings.

We also experimented with using the actual realized inflation in the next year as a regressor instead of using the household's inflation expectation, but the results were qualitatively similar. The model with the actual future inflation as a regressor, however, had a significantly lower log-likelihood and therefore provided a worse fit for the data. This result makes sense, since it is reasonable to expect households to rely more on their subjective inflation forecast when forming their expectations for their own personal income growth. We also saw evidence consistent with this in the last column of Table 6.

Again, it is interesting to see how the relationship between inflation expectations and income growth expectations has changed over different periods and policy regimes. Next, we re-estimate the model, allowing for the coefficient for the inflation expectation term to change at each decade and half-decade. We show the coefficients for each half-decade in Panel B of Table 9. This time-varying coefficients' approach shows that, during the period 1978-1985, households incorporated a great degree of their inflation forecasts to determine their income growth expectations. However, this tendency decreased after 1985, and in the 2000's there is no longer an effect of inflation expectations on income growth forecasts. These are good news for policy-makers and central banks because it means that, in recent years, households no longer feed the vicious cycle of inflation expectations and wage demands.

4.4 Model Fit

As we discussed before, several explanations and models have been offered to explain the evolution of inflation expectations and its heterogeneity, including: a) rational expectations, b) adaptive expectations, and c) different life experiences. In relation to the previous alternatives in the literature, our model includes heterogeneity in the use of information both in terms of observable information (demographic groups attach different importance to their lifetime experiences and have different bias for their expectations) and unobservable information (people update information differently with some groups being "faster learners" than others). Therefore we can think of our model as essentially a model with heterogeneous information and dynamic updating features.

The question is how relevant are these additional features and how much do they add to our understanding relative to previous explanations?

For this purpose, we compare the Efron’s R-square and McFadden Pseudo R-square of our model with those of the alternative models. The Efron’s R-square is a measure of how much of the variability in the data is explained by the model and is built in relation to square deviations from the mean. The McFadden Pseudo R-square compares the likelihood of the model to that of the null model, which is a bivariate normal model of expectations that just considers a constant. This model is the equivalent of assuming people never update their expectations with time-varying information and that heterogeneity of expectations does not differ across demographic groups or over different years. Another way to think about these two metrics is that the Efron’s R-square considers only deviations from the mean, while the Pseudo R-square cares about how the model fits the whole probability distribution. In the case of the bivariate normal in our model, that translates into how well the model’s three parameters – the mean, variance, and the correlation between the private information in the two periods - fit the entire joint distribution of the data.

We compute both the (Efron’s) R-square and the Pseudo R-square for a model that also includes the information of future inflation in the next period, i.e., a purely forward-looking model. This model differs from rational expectations, because it considers that agents can be systematically biased and therefore there could be a significant constant differentiating expectations from future inflation. Therefore we can think of this alternative as a biased rational expectations model. We also report these statistics for a model that includes both a constant and the inflation rate observed in the last quarter, which can be thought of as a model of biased adaptive expectations, since it includes the possibility that a persistent bias could differentiate mean expectations from past inflation. Finally, we consider a model that allows for a constant and the individual lifetime inflation rate, as suggested by Malmendier and Nagel (2011a).

The R-squares and Pseudo R-squares of our model and its three alternatives are shown in Table 10. We focus on the R-square and Pseudo R-square of the models computed for the agents within the 25th and 75th percentiles of inflation expectations for both the 1 year and 5-10 years horizons and the personal income expectations for the next year. We choose to do this because parametric likelihood models are better at explaining the center of the distribution than their tails. In the case of inflation expectations, there are a considerable number of individuals reporting inflation rates

far from the historical experience of the last 35 years and those predictions are hard to explain by any economic model. Therefore, focusing on the population closer to the center of the distribution insures our analysis is not plagued by outliers.

Based on the Efron's R-square, shown in the first three columns of Table 10, we see that our model does slightly better than the alternatives for each of the three expectations, in particular for 1-year ahead inflation expectations, in explaining deviations from the mean. Amongst the alternatives, all three models perform equally well. However, our model has a much higher Pseudo R-square value than any of the alternatives. It is quite striking that these other models explain only a very small part of the individual-level heterogeneity of inflation predictions and the updating process between the first and second periods.

So, overall, all these models do a similar job in explaining the mean value of inflation expectations. However, our model differs a lot in how much of the heterogeneity it explains. This result confirms that demographic heterogeneity and differences in dynamic updating of information are an essential characteristic of inflation expectations.

5 Implications for market outcomes

5.1 Effects of heterogeneity and sticky expectations on inflation

The heterogeneity and persistence (or "stickiness") of inflation expectations should have strong effects on macroeconomic equilibrium. There are several search models where consumers look around for a good deal and decide whether to trade or not (Diamond, 1987; Benabou and Gertner, 1993). Benabou and Gertner (1993) show that cost uncertainty reduces the informativeness of prices by scrambling relative and aggregate variations. But if agents can become better informed, price competition will intensify. This is a similar intuition to that in the literature on learning and asset prices, where several researchers have found that the higher heterogeneity of opinions among agents creates an option price feature to the asset which translates into a price-drift (Harris and Raviv, 1993, Sims, 2009).

Here we stick to a framework similar to Benabou and Gertner (1993), but with major simplifications in order to illustrate the simple point of how consumers' information heterogeneity on prices affects the price level set by firms. We posit a very simple search model, where buyers accept to buy a

product from a store depending on how they think the overall price distribution is. Assume that a store with an homogeneous and indivisible product is trying to set its optimal price, p . Customers will buy its product if p is below their subjective reserve price, p_i , which is heterogeneously distributed, with cdf $F(p_i)$, depending on the private information about what the overall price for this product in the economy is, and whether p is a good-bargain or not. Therefore, the store maximizes its profit function given by:

$$4.1) \max_p p \int_p^\infty 1 \partial F(p_i) = p(1 - F(p))$$

In this model, when there is heterogeneous information observed by consumers on the upper side of the distribution, this increases the monopoly power of firms, since there will be a higher mass of consumers willing to buy at high prices. This should imply that heterogeneity of consumer expectations should have a positive effect on the price level and inflation rates. To observe this more clearly, assume that $F(p_i)$ is an uniform distribution between $[0.5 - h, 1.5 + h]$, where h represents a higher degree of heterogeneity regarding the overall price level in the economy. Now the profit function simplifies to $p(1 - \frac{p-0.5+h}{1+2h})$ and the optimal price simplifies to $p = \frac{1.5 + h}{2}$. Therefore the price level set by firms is clearly increasing in the heterogeneity of consumer profiles in the economy. While this model is simple, it is easy to see that the same conclusions would persist under a much more general framework. For instance, a similar proof could be easily worked out if we used another symmetric distribution of consumer reserve prices besides the uniform, such as a normal distribution.

Also, if agents take a long time to revise their expectations, this should also increase the period of time in which firms will be able to charge high prices without losing their consumers. Therefore higher degrees of "price stickiness" should also be associated with higher markups and more inflation. This is easy to see by making the following change to the simple model summarized in expression 4.1). Suppose that now there is a proportion g of consumers with old information and with distribution of reserve prices given by $[0.5 - h, 1.5 + h]$. g can be interpreted as the proportion of consumer who do not update their information and therefore keep the same reserve prices. There is also a proportion $1 - g$ of consumers with new and more precise information on the overall price level and therefore have lower heterogeneity $[0.5 - w, 1.5 + w]$, where $w < h$. This would change the profit function of the firm to:

$$4.2) \max_p p[(1 - g)(1 - \frac{p-0.5+w}{1+2w}) + g(1 - \frac{p-0.5+h}{1+2h})].$$

Now the new optimal price of a firm that faces consumers with new information is $p = g(\frac{1.5 + h}{2}) + (1 - g)(\frac{1.5 + w}{2})$. Since the heterogeneity of consumers who do not update their information is larger ($w < h$), then it is trivial to check that the price level of this model is increasing in the proportion of consumers who have "sticky expectations" with regards to the correct price level.

We now test the predictions of this simple framework, using our learning model estimated from the Michigan survey data. To measure heterogeneity of consumer expectations, we use the panel data nature of our model to obtain two measures of heterogeneity. The first measure of heterogeneity measures the variance of current inflation forecasts. We use the inter-decile range of the inflation forecasts of the 2nd round interviews, $IDR(\pi_{t+4,i}^p(t))$, as a measure of current heterogeneity in inflation expectations, which is denoted as *heterogeneity_t*. We also include one measure of sticky expectations, the median λ_i , which we denote as *sticky_{i,t}*. The variable λ_i shows how slow agents are to update their idiosyncratic expectations term and documents changes in these expectations over time. Our results are qualitatively similar if we use the the mean value of λ_i instead of the median.

We run a regression of next quarter inflation on $median[\pi_{t+4,i}^p(t)]$, *heterogeneity_t*, and *sticky_{i,t}*. We also condition our regression on other variables affecting inflation, such as inflation of the previous quarter, π_{t-1} , and the inflation volatility in the last semester, $|\pi_{t-1} - \pi_{t-2}|$. The results for the effects of the expectations at one year horizon and five-to-ten years on the next quarter's inflation are shown in the two panels of Table 11. It is easy to verify that both *heterogeneity_t* and *sticky_{i,t}* are highly significant in explaining the inflation in the next quarter, for both horizons. It is also noticeable that median inflation expectation, $median[\pi_{t+4,i}^p(t)]$, has no discernible effect on future inflation. Therefore, using only median/mean expectations and ignoring the heterogeneity and persistence of idiosyncratic beliefs neglects important information.

6 Conclusions

We propose a model where agents provide inflation forecasts based on observable information - such as the previous inflation rates - and unobservable information. In our model, upon receipt of new information, agents may update both the public information as well as their private information. Using the panel data of the Michigan Survey of Consumers, we show that individuals are highly heterogeneous in their updating of inflation expectations. However, over the years, the heterogeneity of expectations for both short-term and long-term inflation has decreased substantially. Also, in the recent decades, agents rely more on previous observed inflation to forecast future inflation rates. This result is consistent with studies that find inflation and earnings' structure became easier to predict more recently (Stock and Watson, 2007).

During the 2000's, the previous period inflation rate matters more for the one-year horizon inflation forecast than for long-term inflation expectations. This shows that although contemporary consumers expect inflation to be more persistent in the short-term, there is a greater confidence in the ability of the Federal Reserve to revert those shocks over the long term. In a similar way, personal income forecasts during the 2000's are less sensitive to subjective inflation expectations, i.e., the interplay between wages and inflation – a common feature of the wage-inflation spiral of the 1970s – has diminished in the last decade.

One notable finding is that individuals differ in how quickly they update their expectations of inflation. In particular, women, ethnic minorities, and less educated agents are slower to update their expectations, giving a larger focus to previous life experience rather than to recent events. These groups are also less prone to change their idiosyncratic private beliefs in the following semesters. This slowness in the updating of new information could explain why these groups systematically report inaccurate expectations. Finally, we relate how learning about inflation and belief heterogeneity is related to market equilibrium of consumer goods and financial markets. We show that the heterogeneity of new idiosyncratic information and "sticky expectations" by consumers can increase the next quarter's inflation. This evidence is consistent with models where consumers search for "best-bargains" (Benabou and Gertner, 1993).

However, our most important result is that our model vastly outperforms other models in explaining the heteroskedasticity and updating of agents' expectations. Expectation differences

across agents are large and persistent over time. Demographic heterogeneity and differences in dynamic updating of information are therefore an essential characteristic of inflation expectations and the most salient feature observed over the last three decades of expectations data.

This conclusion is relevant for improvements in future macro modelling of agents' reactions, since it shows heterogeneity is a much more essential feature of the data than the dichotomy between rational expectations versus backward looking expectations or adaptive updating. Recent research has shown that several structural macro models may not have a stable equilibrium when there is heterogeneity of inflation expectations and updating (Giannitsarou, 2003; Honkapohja and Mitra, 2006). This implies that standard monetary policy will not be able to make inflation converge to the best possible outcome (Orphanides and Williams, 2003). Also, heterogeneous learning dynamics imply that monetary and fiscal policy has different effects on agents' savings (agents that believe in higher future inflation will save and invest less), as well as on the steady-state rate of government deficits (Evans, Honkapohja and Marimon, 2001). Therefore, our finding that agents' learning about inflation is highly heterogeneous should have important implications for the simulation of realistic macro models and policy-making.

References

- [1] Adam, Klaus and Mario Padula, (2011), "Inflation dynamics and subjective expectations in the United States," *Economic Inquiry*, 49 (1), 13-25.
- [2] Anderson, Robert (2008), "US Consumer Inflation Expectations: Evidence Regarding Learning, Accuracy and Demographics," Centre for Growth and Business Cycle Research Discussion Paper 099.
- [3] Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber (1999), "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 17 (1), 22-35.
- [4] Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar (2011), "Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?" mimeo, Federal Reserve Bank of New York.
- [5] Barber, Brad and Terrance Odean (2001), "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116 (1), 261-292.
- [6] Benabou, Roland and Robert Gertner (1993), "Search with learning from prices: Does increased inflationary uncertainty lead to higher mark-ups?," *The Review of Economics Studies*, 60 (1), 69-93.
- [7] Branch, William (2004), "The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations," *Economic Journal*, 114 (497, 7), 592-621.

- [8] Branch, William (2007), "Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations," *Journal of Economic Dynamics and Control*, 31(1), 245-276.
- [9] Bruine de Bruin, Wandi, Wilbert van der Klaauw, Julie Downs, Baruch Fischhoff, Giorgio Topa, and Olivier Armantier (2010), "Expectations of Inflation: The Role of Demographic Variables, Expectation Formation, and Financial Literacy", *The Journal of Consumer Affairs*, 44(2): 381-402.
- [10] Bryan, Michael, and Guhan Venkatu (2001), "The Demographics of Inflation Opinion Surveys," *Federal Reserve Bank of Cleveland Economic Commentary*, October: 1-4.
- [11] Carroll, Christopher (2003), "Macroeconomic Expectations of Households and Professional Forecasters", *Quarterly Journal of Economics*, 118 (1), 269-298.
- [12] Diamond, Peter (1987), "Consumer Differences and Prices in a Search Model," *The Quarterly Journal of Economics*, 102 (2), 429-436.
- [13] Evans, George, Seppo Honkapohja, and Ramon Marimon (2001), "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules," *Macroeconomic Dynamics*, 5 (1), 1-31.
- [14] Eusepi, Stefano and Marco Del Negro (2011), "Fitting Observed Inflation Expectations," forthcoming in the *Journal of Economic Dynamics and Control*.
- [15] Giannitsarou, Chryssi (2003), "Heterogeneous Learning," *The Review of Economic Dynamics*, 6 (4), 885-906.
- [16] Greenwood, Robin and Stefan Nagel (2009), "Inexperienced Investors and Bubbles", *Journal of Financial Economics*, 93(2): 239-258.
- [17] Harris, Milton, and Artur Raviv (1993), "Differences of opinion make a horse race," *Review of Financial Studies*, 6, 473-506.
- [18] Hobijn, Bart, Giorgio Topa, Kristy Mayer, and Cartier Stennis (2009), "Whose Inflation Is It? Household Level vs. Aggregate Measures of Inflation", Manuscript, Federal Reserve Bank of New York.
- [19] Honkapohja, Seppo, and Kaushik Mitra (2006), "Learning Stability in Economies with Heterogeneous Agents," *Review of Economic Dynamics*, 9 (2), 284-309.
- [20] Judd, Kenneth (1998), "Numerical methods in Economics", MIT Press, Cambridge, Massachusetts.
- [21] Katz, Lawrence F., and David H. Autor (1999), "Changes in the Wage Structure and Earnings Inequality," in *Handbook of Labor Economics*, vol. 3A, edited by Orley Ashenfelter and David Card, Amsterdam: North-Holland.
- [22] Keane, M. and D. Runkle (1990), "Testing the rationality of price forecasts: new evidence from panel data", *The American Economic Review*, 80, 714-735.
- [23] Lucas, Robert Jr. (1972), "Expectations and the neutrality of money", *Journal of Economic Theory*, 4 (2), 103-124.
- [24] Malmendier, Ulrike, and Stefan Nagel (2011a), "Learning from inflation experiences," working paper, Stanford University and UC Berkeley.
- [25] Malmendier, Ulrike, and Stefan Nagel (2011b), "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *Quarterly Journal of Economics*, 126(1): 373-416.

- [26] Mankiw, N.Gregory and Ricardo Reis (2002), "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *Quarterly Journal of Economics*, 117 (4), 1295–1328.
- [27] Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2003), "Disagreement About Inflation Expectations," in *NBER Macroeconomics Annual 2003*, ed. by M. Gertler, and K. Rogoff.
- [28] 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.
- [29] Orphanides, Athanasios and John C. Williams (2003), "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," *Inflation Targeting*, Ben Bernanke and Michael Woodford, eds., Chicago: University of Chicago Press for NBER.
- [30] Primiceri, Giorgio, and Thijs van Rens (2009), "Heterogeneous life-cycle profiles, income risk and consumption inequality," *Journal of Monetary Economics*, 56 (1), 20-39.
- [31] Rich, Robert and Joseph Tracy (2006), "The Relationship between Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts," FRBNY Working Paper.
- [32] Roberts, John (1997), "Is Inflation Sticky?" *Journal of Monetary Economics*, 39, 173-196.
- [33] Sims, Christopher (2009), "Inflation Expectations, Uncertainty and Monetary Policy," BIS Working Paper No. 275.
- [34] Souleles, Nicholas (2004), "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 36 (1), 39-72.
- [35] Stock, James, and Mark Watson (2007), "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39 (1), 3–33.
- [36] Vissing-Jorgensen, Annette. (2003), "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions," in *NBER Macroeconomics Annual*.




Figure 1: The figure shows the median inflation expectations as well as the 25th and 75th percentiles of the corss-sectional data. Realized one-year ahead inflation also reported.

Table 1: Heterogeneity in 1-year Inflation Expectations by Various Demographics

	First Survey					Second Survey	
	All Years	All Years	1980	1990	2000s	All Years	All Years
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	0.950*** (0.0454)	0.963*** (0.0456)	1.240*** (0.0878)	0.968*** (0.0745)	0.584*** (0.0592)	1.054*** (0.0458)	1.074*** (0.0462)
Asian	0.0956 (0.186)	0.105 (0.188)	1.307*** (0.419)	-0.343 (0.287)	-0.245 (0.209)	0.731*** (0.190)	0.744*** (0.191)
Black	1.173*** (0.0850)	1.184*** (0.0852)	1.965*** (0.162)	1.012*** (0.137)	0.397*** (0.115)	0.918*** (0.0860)	0.944*** (0.0868)
Hispanic	1.062*** (0.121)	1.117*** (0.121)	1.804*** (0.270)	1.346*** (0.179)	0.0326 (0.140)	1.015*** (0.122)	1.057*** (0.123)
Young ^a	1.058*** (0.0713)	1.067*** (0.0716)	1.691*** (0.134)	0.768*** (0.119)	-0.0199 (0.104)	0.432*** (0.0725)	0.453*** (0.0732)
Mid-age	0.772*** (0.0601)	0.790*** (0.0606)	1.423*** (0.122)	0.483*** (0.101)	0.159** (0.0702)	0.408*** (0.0615)	0.424*** (0.0624)
Lowest Income tercile	0.886*** (0.0756)	0.922*** (0.0761)	0.779*** (0.191)	0.939*** (0.107)	0.752*** (0.0839)	0.724*** (0.0762)	0.724*** (0.0773)
Middle Income tercile	0.360*** (0.0642)	0.386*** (0.0647)	0.257 (0.178)	0.292*** (0.0869)	0.439*** (0.0685)	0.192*** (0.0649)	0.152** (0.0658)
Education	-0.120*** (0.00973)	-0.118*** (0.00976)	-0.114*** (0.0178)	-0.185*** (0.0164)	-0.103*** (0.0139)	-0.151*** (0.00982)	-0.150*** (0.00990)
Inflation in Survey Month		0.401*** (0.0334)					0.437*** (0.0280)
Realized 1-yr ahead Inflation		-0.0970*** (0.0328)					0.0698*** (0.0260)
Constant	4.916*** (0.144)	3.492*** (0.273)	4.631*** (0.252)	4.990*** (0.250)	4.221*** (0.215)	6.234*** (0.146)	3.948*** (0.230)
Observations	79661	78756	29331	23036	20515	67407	65957
R-squared	0.111	0.113	0.069	0.045	0.045	0.134	0.139

The table reports OLS estimates of the regression of one-year ahead inflation point forecast expectation on various demographics.

^a Young is defined as age < 31; Mid-age is defined as age > 30 & age < 61.

Standard Deviations in Parentheses. ***, **, * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 2: Revisions and Forecast Errors for 1 year Horizon, by Demographics

	Absolute Revision of Point Forecast ^a					Absolute Error ^b	
	All	ALL	1980	1990	2000	1st Survey	2nd Survey
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	1.127*** (0.0453)	0.487*** (0.0374)	0.650*** (0.0697)	0.398*** (0.0586)	0.251*** (0.0545)	1.007*** (0.0369)	0.591*** (0.0322)
Asian	0.828*** (0.186)	0.553*** (0.153)	0.0172 (0.323)	0.924*** (0.228)	0.131 (0.196)	0.594*** (0.1530)	0.602*** (0.1320)
Black	1.764*** (0.0858)	0.600*** (0.0707)	0.821*** (0.129)	0.489*** (0.110)	0.157 (0.106)	1.734*** (0.0688)	0.767*** (0.0606)
Hispanic	1.221*** (0.122)	0.306*** (0.101)	0.519** (0.212)	0.159 (0.144)	0.239* (0.132)	1.387*** (0.0989)	0.648*** (0.0871)
Young ^a	0.125* (0.0724)	-0.0262 (0.0595)	-0.360*** (0.109)	0.00232 (0.0941)	0.156 (0.0961)	0.183*** (0.0578)	-0.0215 (0.0506)
Mid-age	-0.157** (0.0618)	-0.156*** (0.0509)	-0.465*** (0.100)	-0.0642 (0.0805)	0.0670 (0.0656)	-0.0614 (0.0492)	-0.0454 (0.0436)
Lowest Income tercile	0.519*** (0.0758)	0.267*** (0.0624)	0.350** (0.148)	0.448*** (0.0842)	0.289*** (0.0778)	0.399*** (0.0618)	0.263*** (0.0547)
Middle Income tercile	0.0238 (0.0636)	-0.00255 (0.0524)	0.0325 (0.135)	0.0236 (0.0680)	0.0748 (0.0626)	0.0737 (0.0526)	0.0478 (0.0460)
Education	-0.261*** (0.00988)	-0.117*** (0.00815)	-0.141*** (0.0144)	-0.0746*** (0.0131)	-0.0610*** (0.0131)	-0.218*** (0.00789)	-0.134*** (0.00696)
Absolute Error in First Survey		0.648*** (0.00369)	0.653*** (0.00579)	0.676*** (0.00610)	0.591*** (0.00875)		0.245*** (0.00317)
Actual Δ Inflation between Surveys		0.0512** (0.0212)	0.0249 (0.0359)	0.241*** (0.0549)	-0.0786*** (0.0248)		
Constant	6.720*** (0.147)	2.738*** (0.123)	3.462*** (0.211)	1.870*** (0.202)	1.863*** (0.205)	6.200*** (0.116)	4.020*** (0.104)
Observations	62445	61837	22778	18273	15121	76861	71248
R-square	0.081	0.387	0.409	0.438	0.278	0.116	0.18

^a Defined as |1-yr ahead inflation point forecast reported in Second Survey - 1-yr ahead inflation point forecast reported in First Survey|.

^b Defined as |Actual realized 1-yr ahead inflation - Respondent's Expectation of 1-yr ahead inflation| OLS estimates reported of a regression onto various demographics.

Standard Deviations in Parentheses. ***, **, * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 3: Mean expectations

	1-yr Horizon Expectation		5-10 yr Horizon Expectation	
	Coefficient	Std-error	Coefficient	Std-error
Mean expectations, $z_{t'}(t)$ (period 1):				
π_{t+1}	-0.018	0.015	-	-
$\bar{\pi}_{t+4 \times 10, t+4 \times 5}$	-	-	0.397	0.385
$\bar{\pi}_{t-4 \times 5, t-4 \times 10}$	-	-	-0.416	0.404
π_{t-1}	0.484***	0.013	0.139***	0.019
constant	3.901***	0.251	0.194	0.629
Mean expectations, $z_{t'}(t)$ (period 2):				
π_{t+1}	0.048***	0.014	-	-
$\bar{\pi}_{t+4 \times 10, t+4 \times 5}$	-	-	-0.051	0.032
$\bar{\pi}_{t-4 \times 5, t-4 \times 10}$	-	-	-0.134***	0.038
π_{t-1}	0.372***	0.013	0.141***	0.021
constant	3.472***	0.244	-0.989	0.673
log(β):				
Female	0.835***	0.062	0.142***	0.008
Asian	0.162	0.130	0.065**	0.029
Black	0.507***	0.044	0.119***	0.020
Hispanic	0.453***	0.056	0.202***	0.023
Young	0.366**	0.062	0.205***	0.018
Middle Aged	0.347***	0.049	0.091***	0.013
Low Income	0.684***	0.052	0.122***	0.013
Middle Income	0.360***	0.046	0.039***	0.009
education	-0.106***	0.008	-0.022***	0.002
constant	-1.346***	0.175	0.516***	0.073
θ:				
constant	2.414***	0.233	2.683***	0.082

All terms include dummies for half-decade periods.

Robust Huber-White standard-errors. ***, **, * denote significance at the 1, 5, and 10 percent levels, respectively. N=85350 for 1-yr horizon; N = 59371 for 5-10 yr horizon.

Table 4: Unobserved Heterogeneity

	Coefficient	Std-error	Coefficient	Std-error	Coefficient	Std-error
	$\log(\sigma_{u_i^1}^2)$:		$\log(\sigma_{u_i^2}^2)$:		$\text{logit}(\lambda)$:	
<i>1 year expectations</i>						
Female	0.288***	0.005	0.301***	0.005	0.006	0.014
Asian	0.218***	0.021	0.268***	0.021	0.127**	0.062
Black	0.422***	0.010	0.341***	0.010	-0.008	0.024
Hispanic	0.345***	0.014	0.239***	0.014	0.192***	0.036
Young	0.060***	0.008	0.011	0.008	-0.079***	0.022
Middle Aged	-0.030***	0.007	-0.012*	0.007	0.041**	0.020
Low Income	0.176***	0.009	0.197***	0.009	-0.026	0.024
Middle Income	0.050***	0.007	0.062***	0.007	0.006	0.021
education	-0.064***	0.001	-0.059***	0.001	0.008***	0.003
π_{t-1}	0.046***	0.002	0.039***	0.002	0.005	0.005
$ \pi_{t-1} - \pi_{t-2} $	0.041***	0.007	0.090***	0.005	-0.025	0.016
constant	2.249***	0.024	2.077***	0.025	0.547***	0.072
<i>5-10 year expectations</i>						
Female	0.322***	0.008	0.331***	0.008	-0.053**	0.021
Asian	0.305***	0.031	0.192***	0.030	-0.362***	0.070
Black	0.509***	0.014	0.385***	0.014	-0.079**	0.035
Hispanic	0.414***	0.020	0.255***	0.020	0.008	0.048
Young	0.168***	0.012	0.069***	0.012	0.160***	0.034
Middle Aged	-0.004	0.010	-0.032***	0.010	0.129***	0.031
Low Income	0.304***	0.012	0.241***	0.012	-0.055	0.034
Middle Income	0.077***	0.010	0.104***	0.010	-0.101***	0.029
education	-0.069***	0.002	-0.078***	0.002	-0.006	0.005
π_{t-1}	0.033***	0.003	0.033***	0.003	-0.010	0.009
$ \pi_{t-1} - \pi_{t-2} $	0.013	0.012	0.024*	0.013	-0.151***	0.035
constant	2.409***	0.045	2.453***	0.048	1.041***	0.136

All terms include dummies for half-decade periods

Table 5: Evolution of Inflation Expectations over Time

1-yr Expectations			5-10 yr Expectations		
Coeff	Std-error	Coeff	Std-error	Coeff	Std-error
Panel A: Evolution of unobserved heterogeneity over each half-decade					
	$\log(\sigma_{u_i}^2)$:	$\log(\sigma_{u_i}^2)$:	$\log(\sigma_{u_i}^2)$:	$\log(\sigma_{u_i}^2)$:	$\log(\sigma_{u_i}^2)$:
dummy 1978-80	-	-	-	-	-
dummy 1981-85	0.126***	0.012	0.066***	0.016	0.031
dummy 1986-90	-0.066***	0.016	-0.092***	0.019	-0.179***
dummy 1991-95	-0.040**	0.018	-0.111***	0.022	-0.211***
dummy 1996-00	-0.197***	0.019	-0.192***	0.023	-0.391***
dummy 2001-05	-0.106***	0.019	-0.110***	0.024	-0.417***
dummy 2006-09	-0.094***	0.021	-0.053**	0.025	0.000
Panel B: Changing the effect of lagged inflation over each half-decade					
Mean exp, $z_t(t)$			$z_t(t)$		
(period 1):			(period 1):		
π_{t-1} , 1978-80	0.345***	0.021	0.286***	0.019	0.517***
π_{t-1} , 1981-85	0.318***	0.016	0.243***	0.017	0.114***
π_{t-1} , 1986-90	0.356***	0.024	0.209***	0.024	0.048
π_{t-1} , 1991-95	0.242***	0.050	0.180**	0.091	0.170*
π_{t-1} , 1996-00	0.451***	0.041	0.464***	0.038	0.203***
π_{t-1} , 2001-05	0.750***	0.043	0.665***	0.037	0.109
π_{t-1} , 2006-09	0.803***	0.060	0.645***	0.051	0.000
N=85350 for 1-yr horizon; N = 59371 for 5-10 yr horizon.					
			$z_t(t)$ (period 2):		
			0.215***	0.078	0.093
			0.122***	0.019	0.024
			0.077*	0.046	0.034
			0.302***	0.046	0.102
			0.168***	0.038	0.036
			0.057	0.044	0.071
			0.000	0.000	0.000

Table 6: Correlates of Income Growth Expectations, and Changes in Income Expectations

	Income exp in 1st Survey	Income exp in 2nd Survey	Change in inc expectations	Absolute change in in income expectations	
	(1)	(2)	(3)	(4)	(5)
Female	-1.51*** (0.12)	-1.62*** (0.13)	0.068 (0.16)	-0.44*** (0.14)	-0.44*** (0.14)
Asian	0.71 (0.50)	0.17 (0.53)	0.38 (0.62)	-0.56 (0.54)	-0.56 (0.54)
Black	0.51** (0.23)	0.65*** (0.25)	-0.32 (0.29)	0.011 (0.26)	0.011 (0.26)
Hispanic	0.33 (0.33)	0.94*** (0.35)	-0.046 (0.42)	-0.14 (0.37)	-0.14 (0.37)
Young ^a	8.62*** (0.19)	9.06*** (0.21)	-0.61** (0.26)	7.67*** (0.22)	7.67*** (0.223)
Mid-age	4.44*** (0.16)	4.64*** (0.18)	-0.20 (0.22)	4.40*** (0.19)	4.40*** (0.19)
Lowest Income tercile	3.25*** (0.20)	3.13*** (0.22)	0.15 (0.26)	2.18*** (0.23)	2.19*** (0.23)
Middle Income tercile	0.76*** (0.17)	0.65*** (0.182)	0.041 (0.21)	0.013 (0.19)	0.023 (0.19)
Education	0.65*** (0.027)	0.63*** (0.029)	0.053 (0.035)	0.35*** (0.031)	0.35*** (0.031)
Inflation expectation in first survey	0.030*** (0.010)				
Inflation expectation in second survey		0.036*** (0.012)			
Change in Inf expectations between surveys			-0.019 (0.011)		
Actual change in income between surveys (in 000s)			0.011*** (0.0025)	-0.010*** (0.0022)	-0.010*** (0.0022)
Absolute Change in inflation expectations between surveys				0.19*** (0.013)	0.19*** (0.013)
Realized change in inflation between surveys					-0.12 (0.074)
Constant	-7.91*** (0.40)	-8.02*** (0.43)	0.056 (0.53)	-1.75*** (0.47)	-1.79*** (0.47)
Observations	71194	65510	55277	55277	55277
R-square	0.051	0.054	0.001	0.033	0.033

OLS estimates of income growth expectations reported of a regression onto various demographics.

Standard Deviations in Parentheses. ***, **, * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 7: 1-Year Personal Income Growth Expectations

	Coefficient	Std-error
mean expectations, $z_{t'}(t)$ (period 1):		
$\pi_{t+4,i}^p(t)$	0.068***	0.021
π_{t-1}	0.007	0.048
constant	3.494*	1.852
mean expectations, $z_{t'}(t)$ (period 2):		
$\pi_{t+4,i}^p(t+1)$	0.039*	0.022
π_{t-1}	-0.250***	0.052
constant	7.128***	1.863
log(β):		
Female	-0.191***	0.050
Asian	0.015	0.060
Black	0.024	0.030
Hispanic	0.063	0.043
Young	1.473	1.040
Middle Aged	1.198	0.958
Low Income	0.563***	0.144
Middle Income	0.141***	0.050
education	0.102***	0.031
constant	-2.025	1.549
θ:		
constant	-1.249***	0.124

$N = 69681$

All terms include dummies for half-decade periods.

Robust Huber-White standard-errors. ***, **, * denote significance at the 1, 5, and 10 percent levels, respectively.

Table 8: Unobserved heterogeneity of Personal Income Growth expectations

	Coefficient	Std-error	Coefficient	Std-error	Coefficient	Std-error
	log($\sigma_{u_i^1}^2$):		log($\sigma_{u_i^2}^2$):		logit(λ):	
Female	-0.023	0.015	-0.010	0.016	-0.129***	0.044
Asian	-0.021	0.056	-0.029	0.061	0.215	0.170
Black	-0.027	0.026	0.031	0.030	-0.043	0.080
Hispanic	-0.041	0.037	0.052	0.043	-0.024	0.141
Young	0.381***	0.028	0.383***	0.031	-0.203**	0.085
Middle Aged	0.252***	0.027	0.238***	0.030	-0.142*	0.082
Low Income	0.324***	0.023	0.258***	0.025	0.017	0.065
Middle Income	0.076***	0.021	0.032	0.024	0.148**	0.063
education	0.030***	0.004	0.025***	0.004	0.030***	0.010
π_{t-1}	-0.007	0.005	-0.014**	0.006	0.020	0.017
$ \pi_{t-1} - \pi_{t-2} $	0.038**	0.020	0.028	0.018	-0.042	0.050
constant	2.257***	0.084	2.376***	0.090	0.487*	0.277

All terms include dummies for half-decade periods.

Table 9: Evolution of Heterogeneity in Income Expectations over Time

Panel A: **Unobserved heterogeneity over each half-decade (N= 59371)**

	Coefficient	Std-error	Coefficient	Std-error
	$\log(\sigma_{u_1}^2)$:		$\log(\sigma_{u_2}^2)$:	
dummy 1978-80	-	-	-	-
dummy 1981-85	0.066	0.046	0.005	0.052
dummy 1986-90	-0.158***	0.053	-0.200***	0.056
dummy 1991-95	-0.048	0.055	-0.173***	0.062
dummy 1996-00	-0.149***	0.058	-0.223***	0.065
dummy 2001-05	-0.120**	0.059	-0.232***	0.065
dummy 2006-09	-0.239***	0.060	-0.290***	0.071

Panel B: **Effect of expected inflation over each half-decade (N= 69681)**

	mean exp, $z_{t'}(t)$ (period 1):		$z_{t'}(t)$ (period 2):	
$\pi_{t+4,i}^p(t)$, 1978-80	0.199***	0.056	0.097*	0.054
$\pi_{t+4,i}^p(t)$, 1981-85	0.106***	0.041	0.052	0.046
$\pi_{t+4,i}^p(t)$, 1986-90	0.079	0.051	0.050	0.056
$\pi_{t+4,i}^p(t)$, 1991-95	0.029	0.062	-0.100	0.066
$\pi_{t+4,i}^p(t)$, 1996-00	0.079	0.057	0.093	0.065
$\pi_{t+4,i}^p(t)$, 2001-05	-0.064	0.060	0.037	0.059
$\pi_{t+4,i}^p(t)$, 2006-09	-0.062	0.070	0.022	0.071

Table 10: Fit of the Main Model and Alternatives

Model	Efron's R-square			McFadden Pseudo R-square		
	1-year Inflation	5-10 year Inflation	Income Growth	1-year Inflation	5-10 year Inflation	Income Growth
Heterogeneous information and updating	37.6%	23.8%	10.1%	18.1%	31.3%	21.8%
Purely forward-looking info (rational exp.)	25.1%	21.0%	7.8%	1.6%	0.9%	0.1%
Purely last-quarter info (adaptive exp.)	23.6%	22.4%	6.7%	1.8%	1.1%	0.1%
Malmendier-Nagel lifetime experience	31.8%	19.8%	8.4%	0.1%	1.1%	0.2%

Table 11: Effect of Inflation Expectations on Realized Inflation

Panel A: Effect of one-year inflation expectations
on the next quarter's inflation, π_{t+1}
N= 125, quarterly observations

	Coefficient	Std-error
π_{t-1}	0.834***	0.042
$ \pi_{t-1} - \pi_{t-2} $	-0.352***	0.135
$Median(\pi_{t+4,i}^p(t))$	-0.00009	0.00009
$heterogeneity_t^a$	0.577***	0.096
$sticky_{i,t}^b$	0.006**	0.002
constant	-1.075***	0.396

Panel B: Effect of long-term (5-10 years) inflation
expectations on the next quarter's inflation, π_t
N = 71, quarterly observations

	Coefficient	Std-error
π_{t-1}	0.854***	0.067
$ \pi_{t-1} - \pi_{t-2} $	-0.123	0.140
$Median(\pi_{t+4 \times (5:10),i}^p(t))$	-0.107	0.113
$heterogeneity_t$	0.559***	0.182
$sticky_{i,t}$	6.293***	1.75
constant	-2.381***	0.744

***, **, * denote significance at the 1, 5, and 10 percent levels, respectively.

^a $heterogeneity_t$ is the inter-decile range of inflation forecasts in the 2nd survey.

^b $sticky_{i,t}$ is the median λ_i (the private information updating parameter).