DISCUSSION PAPER SERIES

No. 5790

DO ACTIONS SPEAK LOUDER THAN WORDS? HOUSEHOLD EXPECTATIONS OF INFLATION BASED ON MICRO CONSUMPTION DATA

Atsushi Inoue, Lutz Kilian and Fatma Burcu Kiraz

INTERNATIONAL MACROECONOMICS

ISSN 0265-8003

DO ACTIONS SPEAK LOUDER THAN WORDS? HOUSEHOLD EXPECTATIONS OF INFLATION BASED ON MICRO CONSUMPTION DATA

Atsushi Inoue, University of British Columbia and North Carolina State University
Lutz Kilian, University of Michigan and CEPR
Fatma Burcu Kiraz, North Carolina State University

Discussion Paper No. 5790
August 2006

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 5790

August 2006

ABSTRACT

Do Actions Speak Louder than Words? Household Expectations of Inflation Based on Micro Consumption Data

Survey data on household expectations of inflation are routinely used in economic analysis, yet it is not clear to what extent households are able to articulate their expectations in survey interviews. We propose an alternative approach to recovering households' implicit expectations of inflation from their consumption expenditures. We show that these implicit expectations have predictive power for CPI inflation. They are better predictors of CPI inflation than survey responses, except for highly educated consumers. Moreover, households' implicit inflation expectations respond to inflation news, consistent with recent work on the transmission of information across consumers. The response of consumers' expectations to inflation news tends to increase with their level of education. Our evidence strengthens the case for macroeconomic models with sticky information.

JEL Classification: D12, D84 and E31
Keywords: consumer expenditure survey, Euler equation, inflation expectations, Michigan survey of consumers and survey of professional forecasters

Atsushi Inoue
Dept of Agricultural &Resource Economics
Box 8109
North Carolina State University
Raleigh, NC 27695 8109
USA
Email: atsushi_inoue@ncsu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=158326

Lutz Kilian
Department of Economics
University of Michigan
238 Lorch Hall
Ann Arbor, MI 48109-1220
USA
Email: lkilian@umich.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=132598

Fatma Burcu Kiraz
Department of Agricultural and
Resource Economics
North Carolina State University,
Box 8109,
Raleigh, NC 27695-8109
USA

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=165207

Submitted 13 July 2006

1 Introduction

Survey data on household inflation expectations are routinely used in economic analysis (see, e.g., Thomas 1999; Mankiw, Reis and Wolfers 2003; Souleles 2004), yet there is reason to be skeptical of the reliability of survey data. We provide evidence that households may have difficulty articulating their views about future inflation in response to survey questions. Given this evidence, we propose an alternative approach to measuring household inflation expectations based on household consumption expenditure data. The central idea is that by trading off future consumption against current consumption, households effectively take a stand on inflation expectations, even if they cannot articulate these expectations. Thus, by observing household consumption growth and by assuming that households, controlling for demographic characteristics, on average optimize their consumption decisions, we should be able to construct an implicit measure of households inflation expectations, provided that we are willing to take a stand on the interest rate faced by households, on the functional form of their utility function and on their intertemporal elasticity of substitution. Whether this model-based implicit measure of inflation expectations is a useful alternative to the standard Michigan survey measure is an empirical question to be addressed in this paper.

Our empirical analysis is based on household expenditure data from the Consumer Expenditure Survey (CEX) conducted by the BLS. We construct estimates of consumers' implicit inflation expectations both at the aggregate level and controlling for educational status as a proxy for households' ability to articulate inflation expectations. Our analysis provides a rich set of testable implications. The test results are of interest to users of the Michigan Survey of Consumers as well as to academic macroeconomists.

One set of results pertains to aggregate measures of inflation expectations without controlling for educational status. The first question is whether the implicit measure of household inflation expectations contains useful information about CPI inflation beyond the information conveyed by standard survey measures. We show that the implicit measure of inflation expectations is a better predictor of the realizations of

CPI inflation than the Michigan survey. The reduction in the root prediction mean-squared error (RPMSE) is 3.7 percentage points using quarterly Michigan survey data and 5.0 percentage points using Michigan survey data for the last month of the quarter.

A second question of interest is whether the new implicit measure of inflation expectations proposed in this paper contains useful information about future CPI inflation beyond the information contained in lagged CPI inflation. We confirm that this measure has marginal predictive value for CPI inflation at the one-quarter horizon, although not as much as the Survey of Professional Forecasters (SPF). The marginal predictive power of the implicit measure relative to that of the Michigan survey depends on whether we use quarterly data or last-month of the quarter data from the Michigan survey. In the latter case, the implicit measure has higher marginal predictive power than the Michigan survey measure.

A third question is how the implicit household expectations are affected by news to inflation as measured by the linearly unpredictable component of the inflation forecasts reported in the Survey of Professional Forecasters. We find strong evidence that household inflation expectations are driven by news about inflation, consistent with models of sticky information and rational inattention (see, e.g., Ball, Mankiw and Reis 2005; Barsky and Kilian 2002; Carroll 2003a,b; Mankiw and Reis 2002; Roberts 1995, 1997, Sims 2002, 2005).

The second set of results explicitly controls for the educational status of the household. While aggregate results provide a useful benchmark, the consumer's educational attainment can be expected to be correlated with his ability to articulate expectations in response to survey questions. A natural conjecture is that the predictive power of the implicit measure of inflation expectations will be stronger relative to the survey measure for consumers with lower levels of education. Such a pattern would be consistent with the view that consumers with less education are less able to articulate the beliefs that they base their consumption decisions on.

We show that this conjecture is broadly supported by the data. Specifically, for consumers with at most a high school degree the reduction in RPMSE from using the implicit measure instead of the quarterly Michigan survey measure ranges from 5.4 to

6.7 percentage points. For consumers with some college experience, this number drops to 3.9 percentage points. For consumers with at least a college degree the reduction diminishes to between 1.6 and 1.7 percentage points. Moreover, formal model selection criteria suggest that there is strong statistical evidence that the implicit measure has higher predictive power for CPI inflation than the survey measure for consumers with low levels of education. For all consumers with less than a college degree, the Schwarz Information Criterion (SIC) selects the forecasting model based on the implicit expectations. For consumers with higher education levels, the ranking is reversed in favor of the Michigan survey measure. Results even more favorable to the implicit measure of inflation expectations are obtained using Michigan survey data for the last month of the preceding quarter. The RPMSE gains relative to the Michigan survey measure range from 3.1 to 7.8 percentage points, depending on the level of education, and the SIC favors the implicit measure for all educational groups. Finally, we show that the implicit measure of inflation expectations has higher marginal predictive power than the Michigan survey measure for all but the highest levels of education.

Disaggregate data also shed light on the transmission of news about future inflation. One test of the economic plausibility of our measure is that more educated consumers should respond more strongly to inflation news from the Survey of Professional Forecasters. Using structural impulse response analysis we find evidence supporting that view. Impulse response estimates show that the responsiveness of household expectations to SPF surprises is systematically increasing in educational status. In addition, we construct an inflation news index along the lines of Carroll (2003a) based on information in Lexis/Nexis. Using an econometric model of the conditional heteroskedasticity in inflation expectations, we provide evidence that highly educated consumers update their implicit inflation expectations more, as the intensity of media reports about inflation increases. While the estimated coefficient on the news variable is only slightly positive and insignificant in the aggregate, there are systematic differences across educational groups. We show that the extent of the updating of expectations tends to increase with the educational status of the consumer for all but the highest levels of education. These findings lend further support to models of sticky information.

We conclude that actions indeed speak louder than words, especially for agents with low levels of education, consistent with the conjecture that only the most highly educated consumers are able to articulate their inflation expectations in response to survey questions. Our evidence shows that CEX data in conjunction with crude economic models can provide an effective tool for measuring household inflation expectations.

The remainder of the paper is organized as follows. In section 2 we present evidence that casts doubt on the reliability of the Michigan survey measure of inflation expectations and motivates our alternative approach. Section 3 introduces the model of consumption behavior underlying the econometric analysis. We show how that model motivates regressions that allow us to recover households' implicit inflation expectations. The data are described in section 4. Section 5 contains the analysis of inflation expectations at the aggregate level as well as disaggregated by consumers' educational status. We conclude in section 6.

2 How Reliable are the Michigan Survey Expectations?

Despite recent evidence in Ang, Bekaert and Wei (2006) that survey expectations of inflation tend to be more accurate than term structure forecasts and regression-based forecasting methods including Phillips curve models, there is reason to be skeptical about the accuracy of these household survey data. It is well known that reported survey expectations of inflation may differ systematically from both the inflation forecasts available in the survey of professional forecasters and from actual consumer price inflation rates (see Figure 1). One possible explanation is that households are simply not acting rationally, which has prompted tests of the rationality of household inflation expectations (see, e.g., de Menil and Bhalla 1975; Fackler and Stanhouse 1977; Gramlich 1983; Bryan and Gavin 1986; Grant and Thomas 1999; Mehra 2002; Souleles 2004). This paper considers an alternative explanation. We explore the possibility that some households are unable to communicate accurately their expectations in response to survey questions.

The prima facie evidence for this explanation is strong. The Michigan survey of consumers elicits consumers' inflation expectations in two steps: The first question relates to the direction of future inflation: *'During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?'.* Respondents are then confronted with a more specific question: *'By what percent do you expect prices to go up, on the average, during the next 12 months?'.* The first row of Table 1a shows that on average more than 1% of the respondents to the Michigan survey of consumers are unable to answer the first question. An additional 7% of consumers are able to determine the likely direction of future inflation, but fail to answer the second question because they cannot articulate the expected value of the future inflation rate.

If our alternative explanation were true, one would expect that more highly educated consumers would be better able to articulate their inflation expectations. Indeed Table 1a shows that a systematic decline in the fraction of nonrespondents, as the educational status of the household improves. Whereas 3.12% of the respondents without a high school diploma were completely unable to answer question 1, that fractions falls to 1 % for high school graduates, 0.69% for consumers with some college education, 0.67% for respondents with a college degree and 0.66% for respondents with a graduate degree. Similarly, the fraction of respondents who cannot answer the second survey question drops from 16.12% for consumers without high school diplomas, to 7.74% for high school graduates, 5.31% for consumers with some college experience, 4.34% for college graduates and 4.31% for consumers with graduate degrees.

This evidence, striking as it may be, is likely to understate the problem. It stands to reason that there must be consumers who arbitrarily indicate some range of inflation rather than admit their inability to complete the survey. In addition, there will be respondents who are unable to report their views accurately despite their best intentions. This view is supported by the prevalence of some extreme views of survey respondents that seem at odds with the actual inflation experience over the same sample period (see Table 1b). For example, on average 3.33% of respondents expect implausibly high inflation in excess of 15% and an additional 16.32% of respondents on average expect no inflation at all, of which a quarter goes as far as expecting consumer prices to fall. Thus, the reliability of survey data on inflation expectations cannot be taken for granted.

In this paper, we develop an alternative methodology for measuring household inflation expectations that does not rely on households' ability to articulate expectations. Rather than using survey data on expectations, we rely on survey data on consumption expenditures. In doing so, we tap a previously unutilized source of information on households' inflation expectations. Our starting point is the standard partial equilibrium model of consumption.

3 Model

Suppose that household i maximizes

$$E_0(\sum_{t=0}^{\infty} \beta^t u(c_{i,t}))$$

with respect to consumption $c_{i,t}$ subject to a sequence of budget constraints where β is a discount factor. The utility function embodies the commonly used assumption of constant relative risk aversion:

$$u(c_{i,t}) \;=\; \frac{1}{1-\rho} c_{i,t}^{1-\rho} \exp(\gamma' x_{i,t} + \eta_i),$$

where $1/\rho$ denotes the intertemporal elasticity of substitution, η_i is a fixed-effect preference shifter and $x_{i,t}$ is a vector of time-varying demographic variables that are assumed to be exogenous. Intertemporal optimization yields the Euler equation:

$$c_{i,t}^{-\rho} e^{\gamma' x_{i,t} + \eta_i} \;=\; \beta E_t \left[(1 + r_{t+1}) c_{i,t+1}^{-\rho} e^{\gamma' x_{i,t+1} + \eta_i} \right], \tag{1}$$

where r_{t+1} denotes the real interest rate prevailing in period t. A second-order approximation to the Euler equation yields

$$E_t(\Delta lnc_{i,t+1}) = \frac{ln\beta}{\rho} + \frac{1}{\rho}E_t(r_{t+1}) + \frac{\rho}{2}\sigma_t^2 + \frac{\gamma'}{\rho}\Delta x_{i,t+1}, \tag{2}$$

where

$$\sigma_t^2 = var_t\left(\Delta lnc_{i,t+1} - \frac{1}{\rho}r_{t+1}\right).[1] \tag{3}$$

This quadratic approximation (2) will be exact when $\Delta lnc_{i,t+1}$ and r_{t+1} are jointly normally distributed.

In addition, we make the following three assumptions:

1. We equate $E_t(r_{t+1})$ with the cross-sectional mean of the households' expectations of the real interest rate, r_{t+1}^e.
2. $Cov_t(\Delta lnc_{i,t+1}, r_{t+1}^e) = 0$.
3. $Var_t(r_{t+1}^e)$ is constant.

Assumptions 1 and 2 would hold, for example, if markets were complete such that cross-sectional and time series averages coincide in population. Assumption 1 allows us to estimate $(1/\rho)E_t(r_{t+1})$ by dummy variables.[2] When Assumption 3 does not hold, our empirical measure of expectations actually represents a weighted average of the conditional mean and of the conditional variance of r_{t+1}^e. Although we have no reason to suspect that these assumptions hold literally, we will proceed as if they do. This approach is consistent with the view that in generating expectations (or forecasts) imposing incorrect structure may still be helpful in reducing out-of-sample prediction errors. The usefulness of these assumptions will be judged in section 5 based on the predictive performance of the resulting implicit expectations measure. Clearly, if our assumptions were far from reality, one would not expect the resulting measure of inflation expectations to be a good predictor. The usefulness of these approximations may also be judged by whether the results broken down by educational status are econom-

[1] See equations (2.10) and (2.11) of Deaton (1992, p. 64), for example.

[2] For a related approach see Beaudry and van Wincoop (1996).

ically plausible. Our empirical results are reassuring on both counts, as will be shown in section 5.

In practice, we will proceed as follows: First we estimate the time dummy coefficients $\{\delta_s\}_{s=1}^{T}$ in

$$\Delta lnc_{i,t+1} = \sum_{s=1}^{T} \delta_s 1(s=t+1) + \gamma' \Delta x_{i,t+1} + \frac{1}{2}\rho \left(\Delta lnc_{i,t} - \sum_{s=1}^{T} \delta_s 1(s=t) - \gamma' \Delta x_{it} \right)^2 + u_{i,t+1} \quad (4)$$

by nonlinear least-squares (NLS), where $1(\cdot)$ is an indicator variable chosen such that $1(A) = 1$ if event A is true and $1(A) = 0$ otherwise. The estimates of $\{\delta_s\}_{s=1}^{T}$ are intended to capture the real interest expectations. In practice, these estimates may be contaminated by aggregate shocks that affect all households equally. That possibility will be addressed in the empirical section. To facilitate the exposition we will abstract from aggregate shocks for now.

The second term in equation (4) captures the second-order term of the quadratic expansion. Our specification of the conditional variance term is similar to the ARCH specification of the income process used in Meghir and Pistaferri (2004). Note that we do not include an intercept in (4) because neither the location nor the scale of the expected real interest rate is identified. The expected real interest rate will be an affine transformation of the estimates of the dummy coefficients. This fact does not affect our subsequent statistical analysis because we are only interested in the linear effects of changes in expectations.

Equation (4) can be estimated by NLS because there are no endogenous regressors and the regression disturbance is orthogonal to lagged variables. To the extent that our measure of household consumption includes durables, there could be an MA(1) component in the regression error. If so, the NLS estimates would be inefficient, but consistent, and the presence of an MA component in the regression error would not impair our analysis. The presence of household durables also would call into the question the interpretation of the time dummies as expectations. We will provide additional sensitivity analysis in section 5 that addresses this concern.

Given estimates of $\{\delta_s\}_{s=1}^{T}$, our affine measure of household inflation expectations is defined by

$$\pi^e_{t+1|t} \; = \; i_{t+1} - \rho\widehat{\delta}_{t+1}. \tag{5}$$

where i_{t+1} is the nominal interest available to consumers in quarter t. This rate can be observed in principle. In this paper, we will use the 3-month Treasury bill rate as a proxy for the marginal interest rate faced by consumers. Similar results would be obtained with the average rate of interest charged on credit card accounts. The regressor $\hat{\delta}_{t+1}$ is a generated regressor, the estimation uncertainty of which vanishes as the cross-sectional dimension of the panel increases. This uncertainty can be treated as negligible in practice, allowing standard inference. Given a value for ρ, this relationship implies the date t expectation of inflation from period t to period $t+1$.

As Carroll (2001) points out, it is not clear how reliable regression estimates of ρ will be in general. Many empirical tests of the usefulness of the implicit expectations measure (such as predictive accuracy tests) can be conducted by simply regressing the variable of interest on i_{t+1} and $\hat{\delta}_{t+1}$. The advantage of focusing on unrestricted linear combinations is that one does not require an explicit choice of ρ. Where we do construct an explicit time series for the expectation of inflation in section 5, we will consider a range of alternative values of ρ. One approach to estimating ρ is based on independent survey evidence. Barsky, Kimball, Juster and Shapiro (1997) elicit estimates of the intertemporal elasticity of substitution from households' survey responses to hypothetical situations. The midrange of their elasticity estimates is about 0.2. Since the intertemporal elasticity of substitution is $1/\rho$, this midrange estimate implies $\rho = 5$, which is also consistent with the regression estimates in Hall (1988). Other authors have obtained somewhat lower regression estimates of ρ. For example, Basu and Kimball (2002) arrive at a value of $\rho = 2$. Given the potential imprecision of these estimates, in section 5 we compute results for a grid of values, encompassing estimates implied by independent survey evidence as well as regression estimates. Our qualitative findings are unaffected by the choice of ρ.

4 Data

4.1 CEX Data

The estimation of equation (4) requires household data on consumption expenditures. In this paper, we will use expenditure data from the Consumer Expenditure Survey (CEX). Unlike the Panel Study on Income Dynamics (PSID), the CEX contains not only food consumption data, but also other relevant household consumption expenditures along with the characteristics of the households. The CEX data set contains data from two different types of surveys: an interview survey and a diary survey. We will use the interview survey data only, since the diary survey does not allow construction of time series of expenditures for specific households at monthly or quarterly frequency.

4.1.1 Household Selection

Panel A consists of households that are surveyed in January, April, July and October. Panel B consists of households that are surveyed in February, May, August and November. Panel C consists of households that are surveyed in March, June, September and December. We will drop households

- with missing data relevant to our analysis.
- with negative or zero total consumption expenditures.
- in Panel A when data are not available for at least three consecutive quarters.
- in Panels B and C when data are not available for four consecutive quarters.

4.1.2 Set of Controls

In estimating equation (4) we control for the demographic characteristics, x_{it}, of each household. The following data are obtained from the Consumer Unit Characteristics and Income (FMLY) file:

- Consumption (c_{it}): Total expenditures (TOTEXPPQ and TOTEXPCQ).
- Demographics and Family Characteristics (x_{it}):

- Family size (FAM_SIZE).
- Number of males age 16 and over (AS_COMP1)
- Number of females age 16 and over (AS_COMP2)
- Number of males age 2 through 15 (AS_COMP3)
- Number of females age 2 through 15 (AS_COMP4)
- Number of members under 2 (AS_COMP5)
- Number of children less than 18 (PERSLT18)
- Number of persons over 64 (PERSOT64)
- Age of reference person (AGE_REF)

4.1.3 Data used for the classification of households

In addition, we make use of the following data when classifying households prior to the regression analysis: Consumer unit identification number (NEWID), interview month (QINTRVMO), and interview year (QINTRVYR).

4.2 Other data

Since the CEX consumption data are not seasonally adjusted, we use seasonally unadjusted CPI inflation rates, π_t, for all urban consumers from the Federal Reserve Bank of St. Louis data base, suitably converted to quarterly frequency. Our proxy for the marginal nominal interest rate, i_{t+1}, faced by consumers is the 3-month Treasury bill rate available from the Federal Reserve Board. Similar results would be obtained using the average rate of interest charged on credit card accounts. To conserve space we focus on the results based on the Treasury bill rate. The survey data of inflation expectations used in this paper are described below.

5 Empirical Results

The estimation period is 1983.QIV-2004.QIV. We discard the observation for 1986.QI due to missing survey data, resulting in a pseudo panel with 81 time series observations after accounting for pre-sample observations. In the empirical section, we will present results at the aggregate level as well as by educational status. The CEX data as well as the Michigan survey data allow us to assign each survey respondent to one of five educational groups: (1) Less than a High School Degree, (2) High School Graduate, (3) Some College, (4) College Graduate, (5) Graduate School.

We implenent the regression approach outlined in section 3 using quarterly data from the CEX. We focus on expectations one quarter ahead. Although in principle the same approach could be used for longer horizons, the fact that CEX data include at most four consecutive quarters of data for the same household precludes the estimation of inflation expectations for horizons longer than one quarter. These data have to be matched with the corresponding survey expectations data. Among the quarterly inflation forecast data available in the Survey of Professional Forecasters we select the forecast for the one-quarter horizon. In contrast, the Michigan Consumer Survey expectation of inflation, while available quarterly, are recorded for a horizon of one year. No data for the one-quarter horizon are available. We therefore follow the approach of Roberts (1997) in using a suitably scaled version of the survey expectations data as a proxy for the one-quarter ahead expectations. The baseline results below use the quarterly data from the Michigan survey. In section 5.2.3. we will present alternative results based on Michigan survey data for the last month of the quarter.

5.1 Issues of Model Specification

We begin by addressing some potential concerns regarding the reliability of our implicit expectations measure. One concern is that our expectations measure is based on regressions that implicitly involve ex post realizations of π_{t+1}. Since real consumption growth is constructed as the log difference of nominal consumption growth and consumer price inflation, in the limiting case, if nominal consumption growth were constant, all the variation in the regressand would be due to changes in future inflation. In that situation, one would expect the time dummy regressors to mimic the variation in future inflation by construction. While nominal consumption growth is not constant in practice, lack of variation in nominal consumption growth would still undermine the credibility of our expectations measure. There are two points that can be made in defense of our approach. First, the standard deviation of the time series of cross-sectional averages of nominal consumption growth adjusted for demographics is more than seven times larger than that of consumer price inflation over the same period. Whereas the former standard deviation is 0.0385, the latter is only 0.0053. Second, if our expectations measure were simply picking up variation in future inflation, its predictive performance should be equal across educational groups rather than systematically varying with educational status, as our evidence below suggests.

A second concern is that our econometric model does not allow us to distinguish between aggregate shocks that affect consumption across all households on the one hand and shifts in real interest rate expectations on the other. Both would be picked up by the time dummies. This point has been discussed by Deaton (1992, pp. 146-148) and Mariger and Shaw (1993), among others. We address this concern by constructing proxies for aggregate shocks and removing their effect on the estimated time dummies. More formally, if the aggregate shocks enter additively, we can decompose the error term in equation (4)

$$u_{i,t+1} = a_{t+1} + \varepsilon_{it+1}$$

into an aggregate component (a_{t+1}) and an idiosyncratic component (ε_{it+1}), where the aggregate component a_{t+1} may be thought of as a weighted average of j aggregate shocks. The aggregate shocks are proxied for by forecast errors constructed from linear autoregressions for observable real variables that are likely to impact consumption:

$$a_{t+1} = \sum_j \alpha_j (y_{j,t+1} - y_{j,t+1|t})$$

This model suggests that we regress $\widehat{\delta}_{t+1}$ on a constant and a_{t+1} and define the real interest rate expectation as the residual of that regression, denoted by $\widetilde{\delta}_{t+1}$. In practice, we include five proxies for aggregate shocks: the commonly used net increase measure of real oil price shocks (see Hamilton 2003; Kilian 2005) and forecast errors from autoregressive models of real S&P500 stock returns, real disposable income growth, the Chicago Fed principal components index of real economic activity ($CFNAI$), and the real 3-month Treasury bill rate obtained by subtracting CPI inflation from the nominal rate. The real oil price shock variable is based on data in Kilian (2005). The CFNAI business cycle index is available at http://www.chicagofed.org/economic_research_and_data/cfnai.cfm. The data on real disposable personal income growth are from the BEA. The S&P500 index series has been deflated by the CPI. The lag orders of the forecasting models are selected based on the SIC (see Inoue and Kilian 2006). The SIC suggests a random walk model for real stock returns, AR(1) models for real disposable income growth and for the $CFNAI$, and an AR(4) model for the real Treasury bill rate. No model is needed for the real oil price shock series. All shock measures considered are in real terms, as consumers would not be expected to respond to nominal shocks, unless these shocks are reflected in unanticipated changes in real variables.

Table 2 shows that these variables jointly account for about 9 percent of the variation in the aggregate $\widehat{\delta}_{t+1}$. At the disaggregate level, the R^2 may be as high as 12 percent for some educational groups. Table 2 also shows OLS point estimates for each aggregate shock and standard errors that account for the generated regressor problem (see Newey and McFadden 1994, pp. 2182-2184). Note that the α_j parameters are estimated separately for each educational group, allowing aggregate shocks to affect each group differently. The distinction between $\widetilde{\delta}_{t+1}$ and $\widehat{\delta}_{t+1}$ matters. In general, controlling for aggregate shocks lowers the predictive power of the implicit measure of expectations. In the remainder of the paper we therefore will employ $\widetilde{\delta}_{t+1}$ rather than $\widehat{\delta}_{t+1}$ as our measure of real interest rate expectations.

The evidence in Table 2 also helps address the concern that households' consumption growth may be related to income growth (see, e.g., Campbell and Mankiw 1990, 1991). Implicit in our model specification is the assumption that households are rational and do not respond to current income. If some households did respond to fluctuations in current income, this misspecification might bias the estimates of δ_s. Explicitly controlling for individual-specific income changes in equation (4) would create an endogenous regressor problem, the standard response to which in models of aggregate data would be instrumental variable estimation. Allowing for the presence of such rule-of-thumb consumers also would introduce unobserved heterogeneity into the model in that the coefficient on income growth would be zero for some consumers, but not for others. This fact suggests that we treat the coefficient on income growth as random from the econometrician's point of view. Campbell and Mankiw in their analysis did not face this problem because they estimated Euler equations on aggregate data, in which case the coefficient on income growth may be treated as constant. To estimate similar types of models on micro data by instrumental variable methods would require several additional assumptions, each of which seems highly implausible:

- Income growth is serially correlated (which may be implausible if income follows an approximate random walk causing the well-known weak instrument problem).
- Whether or not consumers act rationally does not depend on the variables used as instruments such as household income growth.
- The measure of household income used is free of measurement error.

Thus controlling for individual-specific income growth in equation (4) does not seem feasible at this stage. We do not view this as a serious problem. If income growth were important on average, one would expect our expectations measure to be highly correlated with shocks to aggregate income growth, which we showed not to be the case.

5.2 Predictive Power for CPI Inflation

5.2.1 Aggregate Results

A simple first test of the ability of alternative expectations measures to explain future CPI inflation is provided in Table 3. Column 1 focuses on the RPMSE of predictive regressions of CPI inflation on a constant and the expectations measure for the same quarter. For the Michigan survey measure we report the RPMSE of the regression

$$\pi_{t+1} = \alpha_0 + \alpha_1 Michigan_{t+1|t} + v_{t+1}. \tag{6}$$

where $Michigan_{t+1|t}$ denotes the mean survey expectation of inflation reported in the Michigan Survey of Consumers as of quarter t. We also experimented with imposing the restrictions that $\alpha_0 = 0$ (unbiasedness) and $\alpha_1 = 1$ (proportionality). These results are not reported because using these restrictions (one at a time or in conjunction) did not systematically improve the RPMSE of the Michigan survey measure and in several cases raised it compared to the unrestricted model. For the implicit expectations measure we report the RPMSE of the regression

$$\pi_{t+1} = \alpha_0 + \alpha_1 i_{t+1} + \alpha_2 \widetilde{\delta}_{t+1} + v_{t+1} \tag{7}$$

where we do not impose any restrictions on α_0, α_1 and α_2. The advantage of this regression is that we can assess the predictive accuracy of the implicit expectations measure without taking a stand on the value of ρ. This feature is appealing given the well-known difficulties of estimating reliably the parameter ρ from regressions (see Carroll 2001) and the imprecision of survey estimates of ρ (see Barsky et al. 1997). Whereas the magnitude of α_1 and α_2 has no intrinsic meaning, the sign does. We find that all our estimates have a positive sign for the nominal interest rate coefficient and a negative sign for $\widetilde{\delta}_{t+1}$, as would be expected.

All RPMSE results in Table 3 are presented as ratios that normalize the RPMSE of the implicit expectations measure relative to that of the Michigan survey measure. A ratio below unity indicates that the implicit inflation expectation measure is a better predictor of actual CPI inflation than the Michigan survey measure. Table 3 shows an improvement in the RPMSE by 3.7 percentage points.

This finding does not necessarily mean that the implicit expectations measure can be expected to be a better predictor out-of-sample because the regression model (7) contains one more regressor than model (6). A common approach to choosing between competing forecasting models is to rank models by an information criterion that involves a penalty term for parameter profligacy. As shown in Inoue and Kilian

(2006), under weak assumptions the Schwarz Information Criterion (SIC) will consistently select the best out-of-sample forecasting model among any finite set of nested or nonnested models.[3] This property is not shared by alternative methods of ranking forecasting models such as the recursive RPMSE criterion, which would be unappealing in any case given our short sample. The lower the SIC value, the more accurate is the forecasting model expected to be out-of-sample. Table 3 shows that the implicit expectations measure has a strictly lower SIC value (-10.454) than the Michigan survey measure (-10.433), despite the greater parsimony of the latter forecasting model. That conclusion is robust to imposing unbiasedness and/or proportionality restrictions on equation (6).

Although evaluations of predictive accuracy provide a stringent test of the validity of the proposed measure of household inflation expectations, note that we do not advocate the use of these expectations measures for real-time forecasting. Not only are the CEX data available only with a considerable delay, but our expectations measure is based on data for $c_{i,t+1}$ and $x_{i,t+1}$that are not available at date t. Rather the point is to show ex post what household expectations at that point in time must have been, given households' consumption choices. The type of expectations measure constructed in this paper is useful for studying the expectations formation of households, which in turn is of central importance for the design of macroeconomic models. Evaluations of predictive performance simply provide a useful check on the realism of the implications of our model-based approach to measuring expectations.

The SIC results in Table 3 constitute strong evidence that even a crude version of our model-based approach to inferring inflation expectations is practically useful as a predictor of CPI inflation at the quarterly horizon. A closely related question is whether the new indirect measure of inflation expectations proposed in this paper contains useful information about future CPI inflation beyond the information contained in lagged CPI inflation. Table 4a summarizes the results of several alternative predictive regressions. The dependent variable is always one-quarter-ahead CPI inflation, π_{t+1}. The baseline

[3] An exception is the comparison of two nonnested regression models with different degrees of parsimony, but exactly identical PMSEs in population. For further discussion see Inoue and Kilian (2006). We abstract from this possibility which seems remote in practice.

model is:

$$\pi_{t+1} = \alpha_0 + \alpha_1 \pi_t + v_{t+1}. \tag{8}$$

In addition, we consider models with the following sets of additional regressors involving expectations as of date t:

$$\begin{aligned}
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 SPF_{t+1|t} + v_{t+1} && (9)\\
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 Michigan_{t+1|t} + v_{t+1} && (10)\\
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 i_{t+1} + \alpha_3 \widetilde{\delta}_{t+1} + v_{t+1} && (11)\\
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 SPF_{t+1|t} + \alpha_3 Michigan_{t+1|t} + v_{t+1} && (12)\\
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 SPF_{t+1|t} + \alpha_3 i_{t+1} + \alpha_4 \widetilde{\delta}_{t+1} + v_{t+1} && (13)\\
\pi_{t+1} &= \alpha_0 + \alpha_1 \pi_t + \alpha_2 Michigan_{t+1|t} + \alpha_3 i_{t+1} + \alpha_4 \widetilde{\delta}_{t+1} + v_{t+1}, && (14)
\end{aligned}$$

where $SPF_{t+1|t}$ denotes the inflation forecast from the Survey of Professional Forecasters, available from the Philadelphia Fed, and $Michigan_{t+1|t}$ denotes the mean survey expectation of inflation reported in the Michigan Survey of Consumers. The predictive value of each of the two survey measures can be assessed by a simple one-sided t-test. The predictive value of the implicit expectations measure can be tested by conducting a Wald test of the null hypothesis that the regression coefficients of $\widetilde{\delta}_{t+1}$ and i_{t+1} are both zero. Note that this test does not require us to take a stand on the value of ρ. The results of this Wald test will be reported in Tables 4a under the column label *Implicit*. P-values based on suitable standard error estimates that account for the generated regressor problem and possible heteroskedasticity are reported in parentheses.

For all but the last regression in Table 4a the Breusch-Godfrey (BG) test results are consistent with the absence of serial correlation in the regression error. While we do not show individual regression estimates for i_{t+1} and $\widetilde{\delta}_{t+1}$, we note that in all cases the estimate of the nominal interest rate coefficient is positive and that of $\widetilde{\delta}_{t+1}$ is negative. Table 4a shows that the implicit measure of inflation expectations is highly significant, as are the Michigan survey measure and the SPF measure. These test results establish conclusively the marginal predictive content of our expectations measure for CPI inflation. The individual statistical significance of the implicit measure is lost, when the implicit measure is combined in the same regression with the Michigan survey measure or with the SPF measure. The same is true for the Michigan survey measure when it is combined with other measures. In contrast, the SPF measure remains significant at the 10 percent level when combined with other predictors.

Even if there is evidence that expectations measures help predict CPI inflation in population relative to models including only lagged CPI inflation, the existence of predictability in population does not guarantee that these regressors also have predictive value out-of-sample. We again assess the out-of-sample predictive power of each regression based on the SIC. The lower the value of the SIC, the higher the predictive power of the regression model for CPI inflation. Table 4a shows that both household expectations measures improve on models with lagged inflation only (-10.381). Adding the implicit inflation expectations raises the predictive power of the forecasting model (-10.440) as does adding the Michigan survey measure of inflation expectations (-10.451) or adding the SPF forecast (-10.489). The relative gain is greatest with the SPF measure, a fact that will help motivate the impulse response analysis further below. Unlike in Table 3, the Michigan survey measure ranks ahead of the implicit measure when combined with lagged inflation. Combinations of alternative measures of expectations have higher SIC values than any one expectations measure alone, reflecting the small T and unfavorable bias-variance trade-off.

An important concern is that our measure of CEX consumption may be contaminated by the inclusion of at least some durables in the CEX consumption measure. Although the predictive performance of our measure is beyond question, its interpretation as an expectations measure hinges on the appropriateness of the theoretical framework discussed in section 3. Our Euler equation approach is explicitly designed for modeling nondurables consumption as opposed to durables. There is no readily available and suitable measure of quarterly CEX nondurables consumption. One way of gauging the validity of this concern is to compute the contemporaneous correlation of our expectations measure, $\widetilde{\delta}_{t+1}$, with future growth in real personal consumption expenditures on durables ($\Delta c_{t+1}^{durables}$), as defined in the National Income and Product Accounts (NIPA). That correlation is 0.263 in the aggregate. At the disaggregate level, these correlations vary between 0.173 for high school graduates and 0.352 for consumers with graduate school degrees. While these results are at best suggestive, given the inherent conceptual differences between CEX and NIPA data, the concern about mismeasurement of the consumption data must be taken seriously.

We address this concern in Table 4b by explicitly controlling for future growth in durables consumption ($\Delta c_{t+1}^{durables}$) in the forecasting equations underlying Table 4a. Table 4b shows that the inclusion of $\Delta c_{t+1}^{durables}$ significantly raises the predictability of inflation and lowers the SIC value from -10.381 to -10.392. Adding the implicit expectations measure further lowers the SIC to -10.439. In fact, the regression involving only lagged inflation and the implicit expectations measure with an SIC value of -10.440 dominates all regressions involving $\Delta c_{t+1}^{durables}$, whether in isolation or in conjunction with other predictors. This result suggests that the predictive power of the implicit expectations measure is not driven by the inclusion of durables in the CEX consumption data.

5.2.2 Results by Educational Status

Both the Michigan survey expectations data and the CEX consumption data are recorded separately for each of the five educational groups listed at the beginning of this section. This allows us to use our model-based approach to construct measures of expected inflation for each educational group and to compare these implicit inflation expectations to the Michigan survey expectation for the same educational group. All results shown below have been obtained by re-estimating all regressions separately for each educational group, thus controlling for possible heterogeneity across groups.

A natural starting point is the first column of Table 3 which shows the reductions in RPMSEs from using the implicit expectations measure. We find that the greatest gains accrue at lower levels of education, consistent with the view that consumers with low educational attainment are unable to articulate their expectations, allowing even crude proxies based on their consumption choices to improve forecast accuracy, whereas implicit expectations cannot improve forecast accuracy for highly educated consumers with no difficulty in accurately responding to survey questions.

Table 3 shows that a reduction of between 5.4% and 6.7% in RPMSE for consumers without college experience, confirming the superior predictive accuracy of the implicit measure; these gains shrink to 3.9% for households with some college training, to 1.7% for college graduates and to 1.6% for consumers with a graduate degree. The SIC ranking favors the implicit measure for all consumers but those with at least a college degree.

Table 4c studies the marginal predictive content of alternative household expectations measures by educational group. There is no evidence of serial correlation. As in the aggregate analysis, the implicit expectations predictor individually is highly significant for each educational group, as is the Michigan survey measure. Combining both measures results in both predictors being insignificant at the 5% level for consumers with at most a college degree, whereas for consumers with a graduate degree only the Michigan survey measure retains its significance.

Based on the SIC, the implicit measure has higher out-of-sample predictive power for CPI inflation than the Michigan survey measure for all consumers who have not earned at least a college degree. For each of these groups, the SIC favors the implicit measure. For consumers with more education, the SIC ranking is reversed in favor of the survey measure. Notably, for college and university graduates the Michigan survey measure is the more accurate predictor. This evidence once again confirms the potential for implicit expectations measures to improve the accuracy of forecasts made by relatively uneducated consumers who have difficulty articulating their inflation expectations. Combinations of expectations measures are generally suboptimal predictors, reflecting the unfavorable bias-variance trade-off, although for the lowest levels of education they still are more accurate than using the Michigan survey alone.

5.2.3 Sensitivity Analysis: How Important is the Timing of the Michigan Survey Data?

The Michigan survey in addition to quarterly data also includes monthly data. These data have advantages as well as disadvantages compared with the quarterly data we used for the baseline analysis. The disadvantage is that the monthly expectations data provided by the Michigan survey offer less detailed information about consumers' educational status. The breakdown available is: (1) at most a high school degree, (2) some college experience, or (3) at least a college degree. The advantage of monthly data is that the last month of the preceding quarter is likely to be a more accurate measure of the household inflation expectations for the current quarter than the quarterly Michigan survey data.

The predictive analysis using these alternative data yields results broadly similar to (and in some cases stronger than) those reported in Tables 3 and 4. Starting with the direct comparison of the Michigan survey measure and the implicit measure in Table 5, we find that the implicit measure reduces the RPMSE ratio by 5.0 percentage points in the aggregate. For consumers with at most a high school degree, the estimated reduction is 7.8 percentage points, for consumers with some college training 5.3 percentage points and for the most educated 3.1 percentage points. Unlike in Table 3, the SIC ranks the implicit measure ahead of the Michigan survey measure for the aggregate and for each educational group. Moreover, unlike in Table 4a, the marginal predictive power of the implicit measure in Table 6a is second only to the SPF forecast in the aggregate. This result is robust to controlling for future real growth in durables consumption. Broken down by educational status in Table 6b, the implicit measure is preferred to the Michigan survey forecast for all groups but consumers with at least a college degree, consistent with the results in Table 4c. These additional results strengthen our case that the implicit measure provides useful information about household inflation expectations over and above the Michigan survey measure.

5.3 The Response of Household Inflation Expectations to Inflation News

5.3.1 Aggregate Results

An important additional test of the plausibility of consumers's implicit expectations is the question of how these expectations respond to news about future inflation. As our analysis in Table 4a demonstrated, SPF forecasts of inflation contain additional information beyond household expectations data. Building on Carroll (2003a,b) we treat linearly unpredictable changes in SPF forecasts of inflation as a proxy for news about future inflation. One would expect that a surprise increase in professional forecasts of inflation would induce consumers to raise their expectations as well. This question may be addressed in the context of a trivariate vector autoregressive (VAR) model with intercept for professional forecasts of inflation ($SPF_{t+1|t}$), the nominal interest rate (i_{t+1}) and households' real interest rate expectations ($\widetilde{\delta}_{t+1}$). Note that i_{t+1} is assumed to be observed at the beginning of period t. In other words, households form real interest rate expectations $\hat{\delta}_{t+1}$, having observed i_{t+1}. In contrast, SPF forecasts for $t+1$ are formed before i_{t+1} is set. The lag order of the VAR is set to one.

By our timing conventions, SPF forecasts cannot respond to innovations in i_{t+1} within the quarter. In addition, we make the following identifying assumptions: First, we impose the assumption that SPF forecasts of inflation do not respond within the same quarter to innovations in household expectations of inflation. Given the delayed availability of CEX data, this assumption seems reasonable. Our second identifying assumption is that i_{t+2} does not respond to $\widetilde{\delta}_{t+1}$ within the same quarter, which again may be motivated by the delayed availability of the CEX data. Third, we impose the assumption that the implicit household expectations do not respond to SPF innovations within the same quarter. This assumption is less obvious since SPF forecasts are released in the middle of the second month of each quarter, leaving households some time to adjust consumption.

Our identifying assumptions thus can be summarized as follows. Let ε_t denote the vector of structural innovations of the VAR model. Then, suppressing the lagged regressors, the structural VAR model may be written as:

$$\begin{pmatrix} SPF_{t+1|t} \\ i_{t+1} \\ \widetilde{\delta}_{t+1} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ 0 & a_{32} & a_{33} \end{bmatrix} \times \begin{pmatrix} \varepsilon_t^{SPF} \\ \varepsilon_t^{i} \\ \varepsilon_t^{\delta} \end{pmatrix} + \ ... \tag{15}$$

where a_{ij} denotes parameters in the impact multiplier matrix that are to be estimated. Table 7 shows the results of a J-test of the overidentifying restriction that $a_{31} = 0$ for the aggregate as well as by educational status. Since it is not straightforward to correct the p-values to account for the generated regressor problem, Table 7 shows the conventional p-values. In no case can the null be rejected that this restriction is valid.

Given the responses of i_{t+1} and $\widetilde{\delta}_{t+1}$, equation (4) allows us to construct the implied responses of the implicit inflation expectations to a one-standard deviation surprise increase in the SPF forecast of inflation for any choice of ρ. We consider values of $\rho \epsilon \{1, 2, 5\}$. The overall shape of the response is unaffected by the choice of ρ. The main difference is in the scale of the response, which is of little interest here. We therefore impose $\rho = 2$ in the results shown below. The central question from our point of view is whether inflation expectations respond to SPF news. The answer is affirmative. The left panel of Figure 2 shows a sharp peak in the response of aggregate inflation expectations after one quarter, followed by decline that gradually levels off. This evidence is consistent with the view that households adjust their expectations in response to inflation news, as postulated in recent models of macroeconomic expectations (see, e.g., Carroll 2003a,b).

5.3.2 Results by Educational Status

By analogy to the aggregate analysis, we can compute the effect of innovations to the SPF forecast of inflation on household inflation expectations for each educational group. A natural conjecture is that the degree of adjustment in response to news about inflation should increase with the level of education. This view is consistent with models of how information is transmitted in the economy (see, e.g. Carroll 2003a,b; Mankiw and Reis 2002; Mankiw, Reis, and Wolfers 2003). The second panel of Figure 2 broadly confirms this conjecture. We again focus on the results for $\rho = 2$, noting that our qualitative results are robust to alternative choices of ρ. All estimated responses in this figure show a peak after one quarter, before leveling off. The magnitude of the response after one quarter is generally increasing in the level of education, except for the highest level of education. There is a distinct gap between consumers with at most a high school degree and consumers with higher levels of education. Consumers with at least a college degree respond nearly twice as much after one quarter as consumers with at most a high school degree. Since we expect that agents with better education are better able to process news about inflation, this differential response is consistent with models that stress the transmission of news as an important source of frictions in the macroeconomy.

While these differences seem large, it is unclear to what extent they merely reflect sampling error. Given the generated regressor nature of the $\widetilde{\delta}_{t+1}$ and the high persistence of i_{t+1}, it is not straightforward to compute confidence intervals for the responses in Figure 2. One way of reducing sampling error is to aggregate across educational groups. The last panel of Figure 2 shows analogous results for consumers with at most a high school degree, for consumers with some college experience, and for consumers with at least a college degree. All responses are positive and there is a clear increase in the degree of responsiveness with rising levels of education.

5.4 The Effect of Increased News Intensity on Inflation Expectations

5.4.1 Aggregate Results

An alternative approach to assessing the importance of inflation news, is to focus on a direct measure of the intensity of inflation news in the media. In this paper, we construct an inflation news index along the lines of Carroll (2003a) based on information in Lexis/Nexis. For each quarter we count the number of news items in the *Washington Post* and in the *New York Times* that involve the word inflation (or any of its derivatives such as the word *inflationary*). We normalize the series by dividing it by its maximum value. The resulting index of inflation news intensity is shown in Figure

3. A natural conjecture is that households are more prone to adjusting their inflation expectations when they are exposed to more inflation news. This conjecture may be verified by regressing the conditional variance, h_t, on a constant and the log-difference of the number of inflation news items, as recorded in Lexis/Nexis. We specify the news variable in differences because of the high persistence of that variable. The conditional variance in turn is obtained from a regression of the implicit expectations measure on a constant and its own lags.

$$\begin{aligned}
\pi^e_{t+1|t} &= \alpha_1 + \alpha_2 \pi^e_{t|t-1} + \alpha_3 \pi^e_{t-1|t-2} + \alpha_4 \pi^e_{t-2|t-3} + \alpha_5 \pi^e_{t-3|t-4} + u_{t+1} \\
u_{t+1} &= \sqrt{h_t}\varepsilon_{t+1}, \text{ where } \varepsilon_{t+1} \overset{iid}{\sim} (0,1) \\
h_t &= \beta_1 + \beta_2 \Delta n_t.
\end{aligned}$$

A value of $\beta_2 > 0$ would be evidence that households update their implicit inflation more, as the intensity of news about inflation increases. In constructing $\pi^e_{t+1|t}$ we considered $\rho \epsilon \{1, 2, 5\}$. The results discussed below are for $\rho = 2$. The other choices of ρ yielded qualitatively similar results.

Table 8 shows the estimated coefficients of the conditional variance equation (multiplied by 100) with p-values in parentheses. There is no evidence of unmodelled serial correlation in the squared residuals. The point estimate for β_2 is slightly positive, but the estimate is not statistically significant. We conclude that there is no evidence that consumers on average adjust their inflation expectations in response to an increase in news intensity. In other words, a substantial fraction of consumers must be unresponsive to changes in news intensity. An interesting question that we will turn to next is why these consumers do not seem to update their expectations. As we will show, how much consumers do adjust, is linked to their educational status.

5.4.2 Results by Educational Status

Table 8 shows the corresponding estimates of the response to changes in news intensity by educational group. Although most point estimates are statistically insignificant, there is a clear pattern in the results. The higher the level of education, the more positive the response of household expectations to news. The estimate of β_2 for consumers with less than a high school degree is negative, but statistically insignificant. For households with at least a high school degree the estimate of β_2 is always positive. While most estimates are imprecisely estimated, that for college graduates is highly statistically significant. The last rows of Table 8 show additional results at a higher level of aggregation across educational groups. Although none of the estimates are significant at the 5% level, both the estimate of β_2 and its t-statistic are strictly increasing in education and the t-statistic for consumers with at least a college degree reaches a p-value of 0.061. While these results are not as sharp as the earlier impulse response results, the overall pattern in Table 8 suggests that the low and insignificant response found in the results for the aggregate is largely driven by the relatively uneducated consumers.

6 Conclusion

We proposed a new method of imputing inflation expectations based on household expenditure data. Our evidence shows that CEX data in conjunction with the use of simple economic optimizing models can provide an effective tool for measuring household inflation expectations. The expectations data we derived complement existing measures of inflation expectations from the Michigan Survey of Consumers. We showed that the new expectations measure contains useful information about future CPI inflation beyond the information contained in survey measures.

The tools developed in this paper also are useful for testing implications of models of sticky information. Using our new measure of inflation expectations, we provided evidence in support of economic models of household behavior such as Carroll (2003a,b) that rationalize the slow response of macroeconomic expectations based on models of information processing and propagation. These economic models provide an important source of monetary nonneutralities, also known as sticky information, that is empirically more plausible than menu costs or other sources of ad hoc frictions (see, e.g., Mankiw and Reis 2002, Ball et al. 2005).

Existing work on the transmission of inflation news has focused on the aggregate behavior of households. Building on this literature, our analysis highlighted the importance of differences in educational attainment for the speed with which news about inflation is reflected in household expectations of inflation. Our results are consistent with the view that the slow adjustment of household expectations reflects at least in part the inability of agents to process news about future inflation. Incorporating these types of heterogeneities into macroeconomic models is an important challenge for future work.

References

1. Ang, A., Bekaert, G., and M. Wei (2006) "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?," forthcoming: *Journal of Monetary Economics.*

2. Ball, L., Mankiw, N.G., and R. Reis (2005), "Monetary Policy for Inattentive Economies," *Journal of Monetary Economics*, 52, 703-725.

3. Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual 2001*, 137-183.

4. Barsky, R.B., Kimball, M., Juster, T., and M.D. Shapiro (1997), "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey ," *Quarterly Journal of Economics*, 112, 537-579.

5. Basu, S., and M. Kimball (2002), "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption," mimeo, Department of Economics, University of Michigan.

6. Beaudry, Paul and Eric van Wincoop (1996), "The Intertemporal Elasticity of Substitution: An Exploration using a US Panel of State Data," *Economica,* 63, 495-512.

7. Bryan, M.F., and W.T. Gavin (1986), "Models of Inflation Expectations Formation: A Comparison of Household and Economist Forecasts. Comment," *Journal of Money, Credit and Banking*, 18, 539–544.

8. Campbell, J.Y., and N.G. Mankiw (1990), "Permanent Income, Current Income and Consumption," *Journal of Business and Economic Statistics,* 8, 269-279.

9. Campbell, J.Y., and N.G. Mankiw (1991), "The Response of Consumption to Income. A Cross-Country Investigation," *European Economic Review,* 35, 723-767.

10. Carroll, C.D. (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics* 1, Article 6.

11. Carroll, C.D. (2003a), "The Epidemiology of Macroeconomic Expectations," forthcoming in: Blume, L. and Durlauf, S. (eds.), *The Economy as an Evolving Complex System, III.* Oxford University Press

12. Carroll, C.D. (2003b), "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118, 269-298.

13. Deaton, Angus (1992), *Understanding Consumption,* Oxford University Press: Oxford, UK.

14. Fackler, J., and B. Stanhouse (1977), "Rationality of the Michigan Price Expectations Data," *Journal of Money, Credit and Banking*, 9, 662–666.

15. Gramlich, E.M. (1983), "Models of Inflation Expectations Formation: A Comparison of Household and Economist Forecasts," *Journal of Money, Credit and Banking*, 15, 155–173.

16. Grant, A.P., and L.B. Thomas (1999), "Inflation Expectations and Rationality Revisited" *Economics Letters*, 62, 331-338.

17. Hall, R.E. (1988), "Intertemporal Elasticity of Substitution," *Journal of Political Economy*, 96, 971-989.

18. Hamilton, J.D. (2003), “What is an Oil Shock?” *Journal of Econometrics*, 113, 363-398.

19. Inoue, A., and L. Kilian (2006), “On the Selection of Forecasting Models,” *Journal of Econometrics*, 130, 273-306.

20. Kilian, L. (2005), ”Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” mimeo, Department of Economics, University of Michigan.

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

22. Mankiw, N.G., Reis, R., and J. Wolfers (2004), “Disagreement about Inflation Expectations,” *NBER Macroeconomics Annual 2003*, 209-248.

23. Mariger, R.P., and K. Shaw (1993), “Unanticipated Aggregate Disturbances and Tests of the Life-Cycle Consumption Model Using Panel Data,” *Review of Economics and Statistics*, 75, 48-56.

24. Meghir, C., and L. Pistaferri (2004), “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72, 1-32.

25. Mehra, Y.P. (2002), “Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality,” *Federal Reserve Bank of Richmond Economic Quarterly*, 88, 17-36.

26. Menil, G. de, and S. Bhalla (1975), “Direct Measurement of Popular Price Expectations,” *American Economic Review*, 65, 169-180.

27. Newey, W.K., and D.L. McFadden (1994), “Large-Sample Estimation and Hypothesis Testing,” in Engle, R.F., and D.L. McFadden (eds.), *Handbook of Econometrics*, Volume 4, Elsevier: Amsterdam, Netherlands, 2111-2245.

28. Roberts, J.M.. (1995), “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit and Banking*, 27, 975-984.

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

30. Sims, C.A. (2002), "Implications of Rational Inattention," mimeo, Department of Economics, Princeton University.

31. Sims, C.A. (2005), "Rational Inattention: A Research Agenda," mimeo, Department of Economics, Princeton University.

32. Souleles, N.S. (2004), "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," *Journal of Money, Credit and Banking*, 36, 39-72.

33. Thomas, L.B. (1999), "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives*, 13, 125-144.

Table 1a. Education Group and Do-Not-Know Responses in Michigan Survey

Average Percentages by Education Level: 1983:Q4–2004:Q4

Education Level	Do Not Know; NA	Will Go Up By; Do Not Know How Much
Full Sample	1.12	6.99
Less than High School	3.12	16.12
High School	1.00	7.74
Some College	0.69	5.31
College Degree	0.67	4.34
Graduate Studies	0.66	4.31

Source: Michigan Survey of Consumers

Table 1b. Education Group and Extreme Responses in Michigan Survey

Average Percentages by Education Level: 1983:Q4–2004:Q4

Education Level	Deflation	No Inflation	Inflation>15%
Full Sample	3.51	16.32	3.33
Less than High School	3.90	16.18	5.33
High School	3.06	16.78	4.21
Some College	3.76	16.18	2.85
College Degree	3.63	16.17	1.96
Graduate Studies	3.38	14.58	1.25

Source: Michigan Survey of Consumers

Table 2. Explanatory Power of Aggregate Shocks for Time Dummies

		Forecast errors					
	Constant	Real Disposable Personal Income Growth	CFNAI	Real S&P500 Returns	Real T-Bill Rate	Real Oil Price Shock	R^2
Full Sample	0.011	0.004	0.001	-2.728	3.846	1.859	0.091
	(0.015)	(0.493)	(0.006)	(1.920)	(1.500)	(1.056)	
Less than High School	0.001	-0.336	-0.002	-2.959	4.719	2.989	0.110
	(0.019)	(0.612)	(0.007)	(2.283)	(1.768)	(1.092)	
High School Graduate	0.009	0.029	0.005	-2.867	4.715	3.532	0.124
	(0.016)	(0.580)	(0.007)	(2.188)	(1.599)	(1.340)	
Some College	0.020	-0.241	0.007	-1.946	2.722	-0.307	0.030
	(0.020)	(0.658)	(0.010)	(2.429)	(1.877)	(1.301)	
College Graduate	0.018	0.464	0.006	-2.118	3.682	2.112	0.063
	(0.019)	(0.661)	(0.008)	(2.420)	(2.260)	(1.709)	
Graduate School	0.025	0.689	-0.021	-2.691	1.787	0.346	0.049
	(0.025)	(0.915)	(0.011)	(2.991)	(2.738)	(1.627)	

Source: The forecasting models are described in the text. The numbers in parentheses in columns (2)-(7) are standard errors. The standard errors account for the generated regressor problem and possible heteroskedasticity. All regressions pass tests for zero serial correlation in the residuals.

Table 3. Predictive Accuracy of Expectations Measures for CPI Inflation Outcomes

	RPMSE Ratio	*SIC*	
	Implicit/Michigan	Michigan	Implicit
Full Sample	0.963	-10.433	-10.454
Less Than High School	0.944	-10.428	-10.490
High School Graduate	0.933	-10.387	-10.471
Some College	0.961	-10.406	-10.431
College Graduate	0.983	-10.444	-10.425
Graduate School	0.984	-10.449	-10.427

Notes: The RPMSEs have been constructed based on regressions of actual inflation on a constant and the expectations measure in question. SIC stands for Schwarz Information Criterion.

Table 4a. Marginal Predictive Content of Expectations Measures for CPI Inflation

Full Sample								
Constant	π_t	SPF	Michigan	Implicit	BG	*SIC*	*N*	*T*
0.008	-0.095				0.078	-10.381		81
(0.001)	(0.111)				(0.780)			
0.002	-0.254	0.957			1.682	-10.489		81
(0.002)	(0.122)	(0.000)			(0.195)			
-0.000	-0.298		1.051		0.539	-10.451		81
(0.003)	(0.128)		(0.000)		(0.463)			
0.004	-0.205			10.920	0.298	-10.440	117814	81
(0.002)	(0.119)			(0.004)	(0.585)			
0.002	-0.258		0.415	3.781	0.496	-10.397	117814	81
(0.003)	(0.140)		(0.130)	(0.151)	(0.481)			
0.002	-0.261	0.893	0.093		1.632	-10.435	117814	81
(0.003)	(0.130)	(0.023)	(0.423)		(0.201)			
0.003	-0.245	0.515		2.056	1.231	-10.406	117814	81
(0.002)	(0.132)	(0.069)		(0.358)	(0.267)			

Notes: The numbers in parentheses in columns 1-2 are suitable standard errors, those in columns 3-6 are p-values. BG stands for the Breusch-Godfrey test for first-order serial correlation and *SIC* for the Schwarz Information Criterion.

Table 4b. Marginal Predictive Content of Expectations Measures for CPI Inflation

Controlling for Future Real Growth in Durables Consumption

Full Sample								
Constant	π_t	$\Delta c_{t+1}^{durables}$	SPF	Michigan	Implicit	*SIC*	*N*	*T*
0.008	-0.095					-10.381		81
(0.001)	(0.111)							
0.009	-0.111	-0.035				-10.392		81
(0.001)	(0.103)	(0.013)						
0.005	-0.235	-0.030			10.178	-10.439	117814	81
(0.002)	(0.114)	(0.012)			(0.006)			
0.004	-0.205				10.920	-10.440	117814	81
(0.002)	(0.119)				(0.004)			
0.003	-0.285	-0.030		0.398	4.121	-10.395	117814	81
(0.002)	(0.132)	(0.012)		(0.117)	(0.127)			
0.004	-0.269	-0.029	0.455		2.613	-10.401	117814	81
(0.002)	(0.125)	(0.012)	(0.083)		(0.271)			

Notes: The numbers in parentheses in columns 1-3 are suitable standard errors, those in columns 4-6 are p-values.

Table 4c. Marginal Predictive Content of Expectations Measures for CPI Inflation by Educational Status of Household

Constant	π_t	Michigan	Implicit	BG	*SIC*	*N*	*T*
Less Than High School							
-0.004	-0.193		14.137	0.210	-10.473	22787	81
(0.002)	(0.119)		(0.001)	(0.647)			
0.002	-0.199	0.554		1.261	-10.411		81
(0.002)	(0.118)	(0.002)		(0.262)			
0.003	-0.208	0.150	0.345	0.061	-10.424	22787	81
(0.002)	(0.124)	(0.210)	(0.140)	(0.805)			
High School Graduate							
0.004	-0.196		12.064	0.284	-10.453	35327	81
(0.002)	(0.115)		(0.002)	(0.594)			
0.003	-0.203	0.610		1.617	-10.365		81
(0.003)	(0.126)	(0.019)		(0.204)			
0.005	-1.189	-0.049	0.391	0.234	-10.399	35327	81
(0.003)	(0.130)	(0.568)	(0.134)	(0.628)			
Some College							
0.005	-0.229		9.427	0.035	-10.428	29521	81
(0.002)	(0.121)		(0.009)	(0.852)			
0.002	-0.226	0.831		0.036	-10.394		81
(0.003)	(0.119)	(0.003)		(0.850)			
0.003	-0.252	0.224	0.350	0.049	-10.378	29521	81
(0.003)	(0.127)	(0.230)	(0.142)	(0.825)			
College Graduate							
0.005	-0.238		9.574	0.111	-10.425	17752	81
(0.002)	(0.122)		(0.008)	(0.739)			
0.001	-0.287	1.062		0.523	-10.460		81
(0.002)	(0.127)	(0.000)		(0.469)			
0.002	-0.299	0.574	0.255	0.003	-10.394	17752	81
(0.002)	(0.133)	(0.053)	(0.148)	(0.958)			
Graduate School							
0.005	-0.235		9.374	0.039	-10.425	12427	81
(0.002)	(0.119)		(0.009)	(0.844)			
0.000	-0.360	1.219		0.838	-10.495		81
(0.002)	(0.127)	(0.000)		(0.360)			
0.001	-0.345	0.796	0.201	0.178	-10.411	12427	81
(0.009)	(0.133)	(0.020)	(0.168)	(0.673)			

Notes: The numbers in parentheses in columns 1-2 are standard errors, those in columns 3-5 are p-values.

Table 5. Predictive Accuracy of Expectations Measures for CPI Inflation Outcomes Based on Michigan Survey Data for the Last Month of the Preceding Quarter

	RPMSE Ratio	*SIC*	
	Implicit/Michigan	Michigan	Implicit
Full Sample	0.950	-10.405	-10.454
At Most High School	0.922	-10.385	-10.493
Some College	0.947	-10.378	-10.431
At Least College Degree	0.969	-10.418	-10.427

Notes: See Table 3.

Table 6a. Marginal Predictive Content of Expectations Measures for CPI Inflation Based on Michigan Survey Data for the Last Month of the Preceding Quarter

Full Sample								
Constant	π_t	SPF	Michigan	Implicit	BG	*SIC*	*N*	*T*
0.008	-0.095				0.078	-10.381		81
(0.001)	(0.111)				(0.780)			
0.002	-0.254	0.957			1.682	-10.489		81
(0.002)	(0.122)	(0.000)			(0.195)			
0.000	-0.307		0.967		0.422	-10.416		81
(0.003)	(0.138)		(0.002)		(0.516)			
0.004	-0.205			10.920	0.298	-10.440	117814	81
(0.002)	(0.119)			(0.004)	(0.585)			
0.002	-0.263		0.349	5.875	0.464	-10.394	117814	81
(0.003)	(0.148)		(0.169)	(0.053)	(0.496)			
0.002	-0.245	0.998	-0.070		1.806	-10.435	117814	81
(0.003)	(0.139)	(0.006)	(0.558)		(0.179)			
0.003	-0.245	0.515		2.056	1.231	-10.406	117814	81
(0.002)	(0.132)	(0.069)		(0.358)	(0.267)			

Notes: See Table 4a..

Table 6b. Marginal Predictive Content of Expectations Measures for CPI Inflation by Educational Status of Household

Based on Michigan Survey Data for the Last Month of the Preceding Quarter

Constant	π_t	Michigan	Implicit	BG	*SIC*	*N*	*T*
At Most High School							
0.004	-0.180		13.737	0.093	-10.471	58114	81
(0.002)	(0.117)		(0.001)	(0.760)			
0.003	-0.230	0.594		0.734	-10.368		81
(0.003)	(0.130)	(0.015)		(0.392)			
0.003	-0.212	0.169	0.364	0.229	-10.420	58114	81
(0.003)	(0.137)	(0.276)	(0.131)	(0.632)			
Some College							
0.005	-0.229		9.427	0.035	-10.428	29521	81
(0.002)	(0.121)		(0.009)	(0.852)			
0.005	-0.186	0.463		0.065	-10.349		81
(0.002)	(0.119)	(0.040)		(0.798)			
0.004	-0.239	0.064	0.382	0.048	-10.374	29521	81
(0.003)	(0.123)	(0.403)	(0.134)	(0.826)			
At Least College Degree							
0.005	-0.236		9.519	0.061	-10.425	30179	81
(0.002)	(0.122)		(0.009)	(0.804)			
0.001	-0.299	1.024		0.918	-10.432		81
(0.003)	(0.138)	(0.002)		(0.338)			
0.003	-0.292	0.428	0.299	0.015	-10.383	30179	81
(0.002)	(0.144)	(0.148)	(0.151)	(0.904)			

Notes: See Table 4c.

Table 7. J-Test of Overidentifying VAR Restrictions

	J-Test of Overidentifying Restriction: a_{31}= 0	p-value
Full Sample	0.022	(0.883)
Less Than High School	0.072	(0.789)
High School Graduate	0.320	(0.571)
Some College	0.053	(0.819)
College Graduate	0.214	(0.644)
Graduate School	0.252	(0.616)

Notes: All results shown are based on $\rho = 2$.

Table 8. The Effect of Changes in Inflation News Intensity on the Volatility of the Implicit Household Expectations

	Intercept β_1	Δn_t β_2
Full Sample	0.173	0.021
	(0.000)	(0.417)
Less than High School	0.542	-0.356
	(0.000)	(0.868)
High School Graduate	0.303	0.043
	(0.000)	(0.401)
Some College	0.368	0.027
	(0.000)	(0.450)
College Graduate	0.486	0.609
	(0.000)	(0.013)
Graduate School	0.731	0.180
	(0.000)	(0.315)
At Most High School	0.298	0.021
	(0.000)	(0.444)
Some College	0.368	0.027
	(0.000)	(0.450)
At Least College Degree	0.384	0.294
	(0.000)	(0.061)

Notes: Point estimates multiplied by 100 with p-values in parentheses. All results based on ρ=2.

Figure 1: U.S. Survey Expectations and Realizations of Inflation

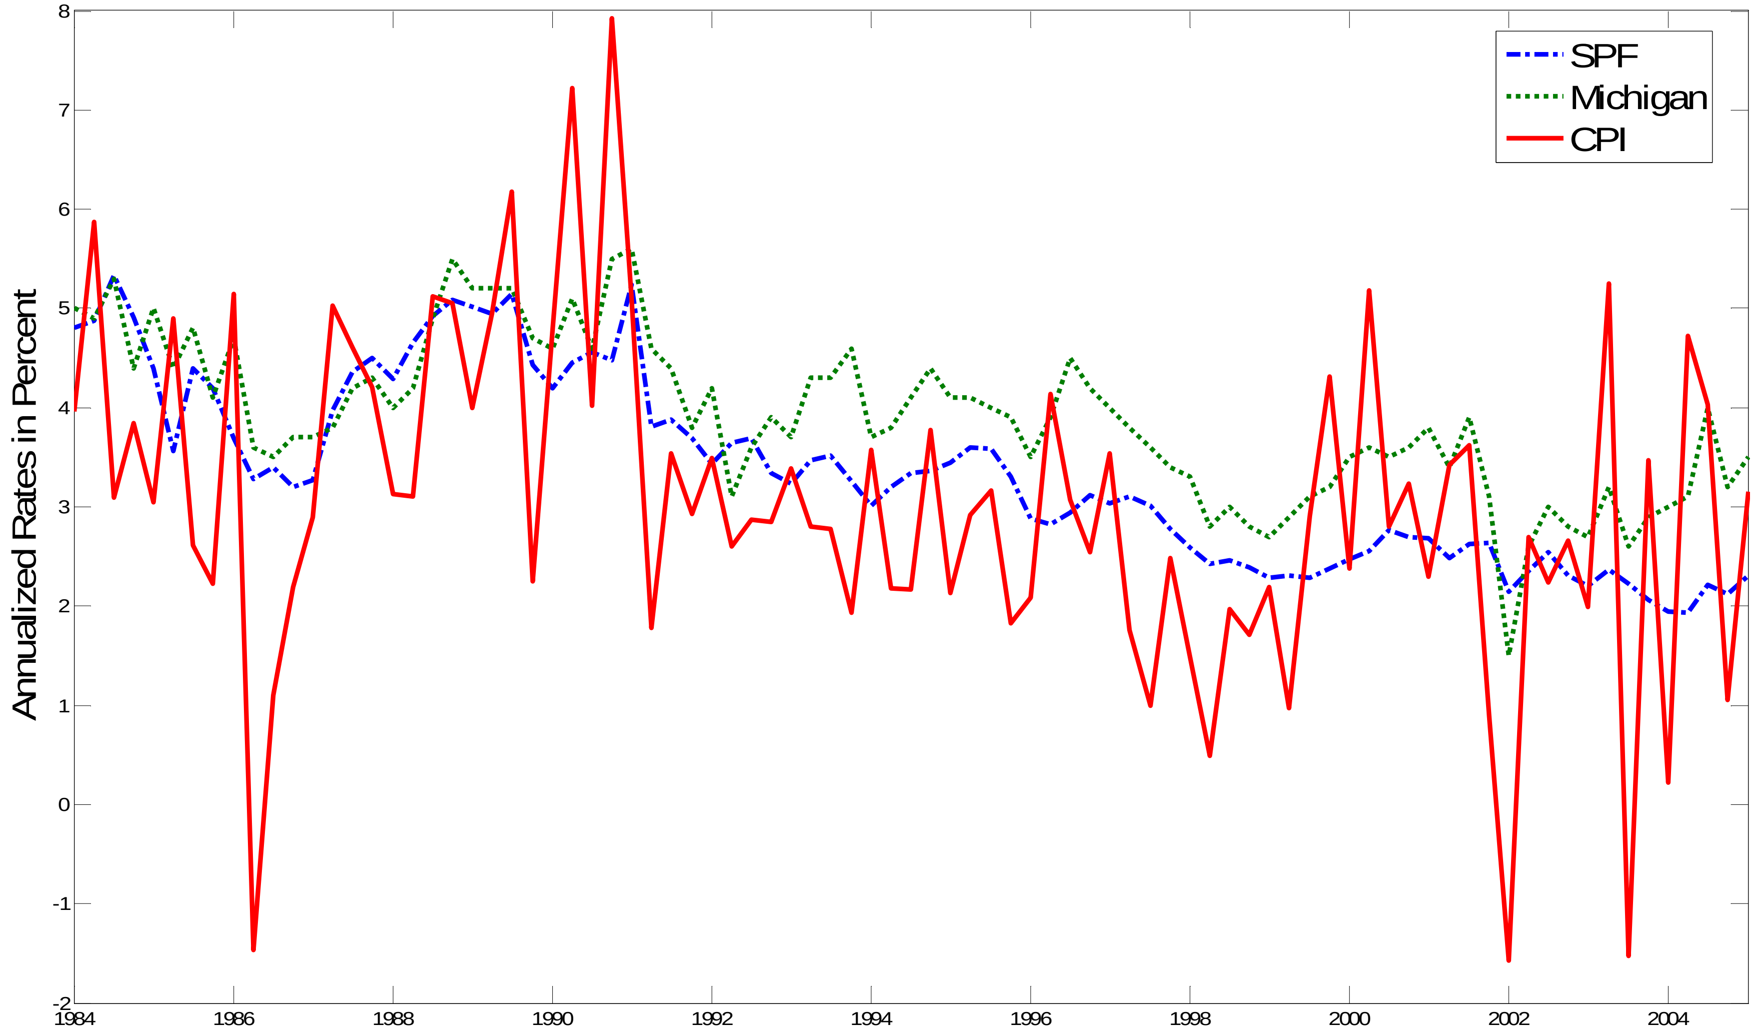

Figure 2:
Responses of Implicit Household Inflation Expectations to an SPF Shock

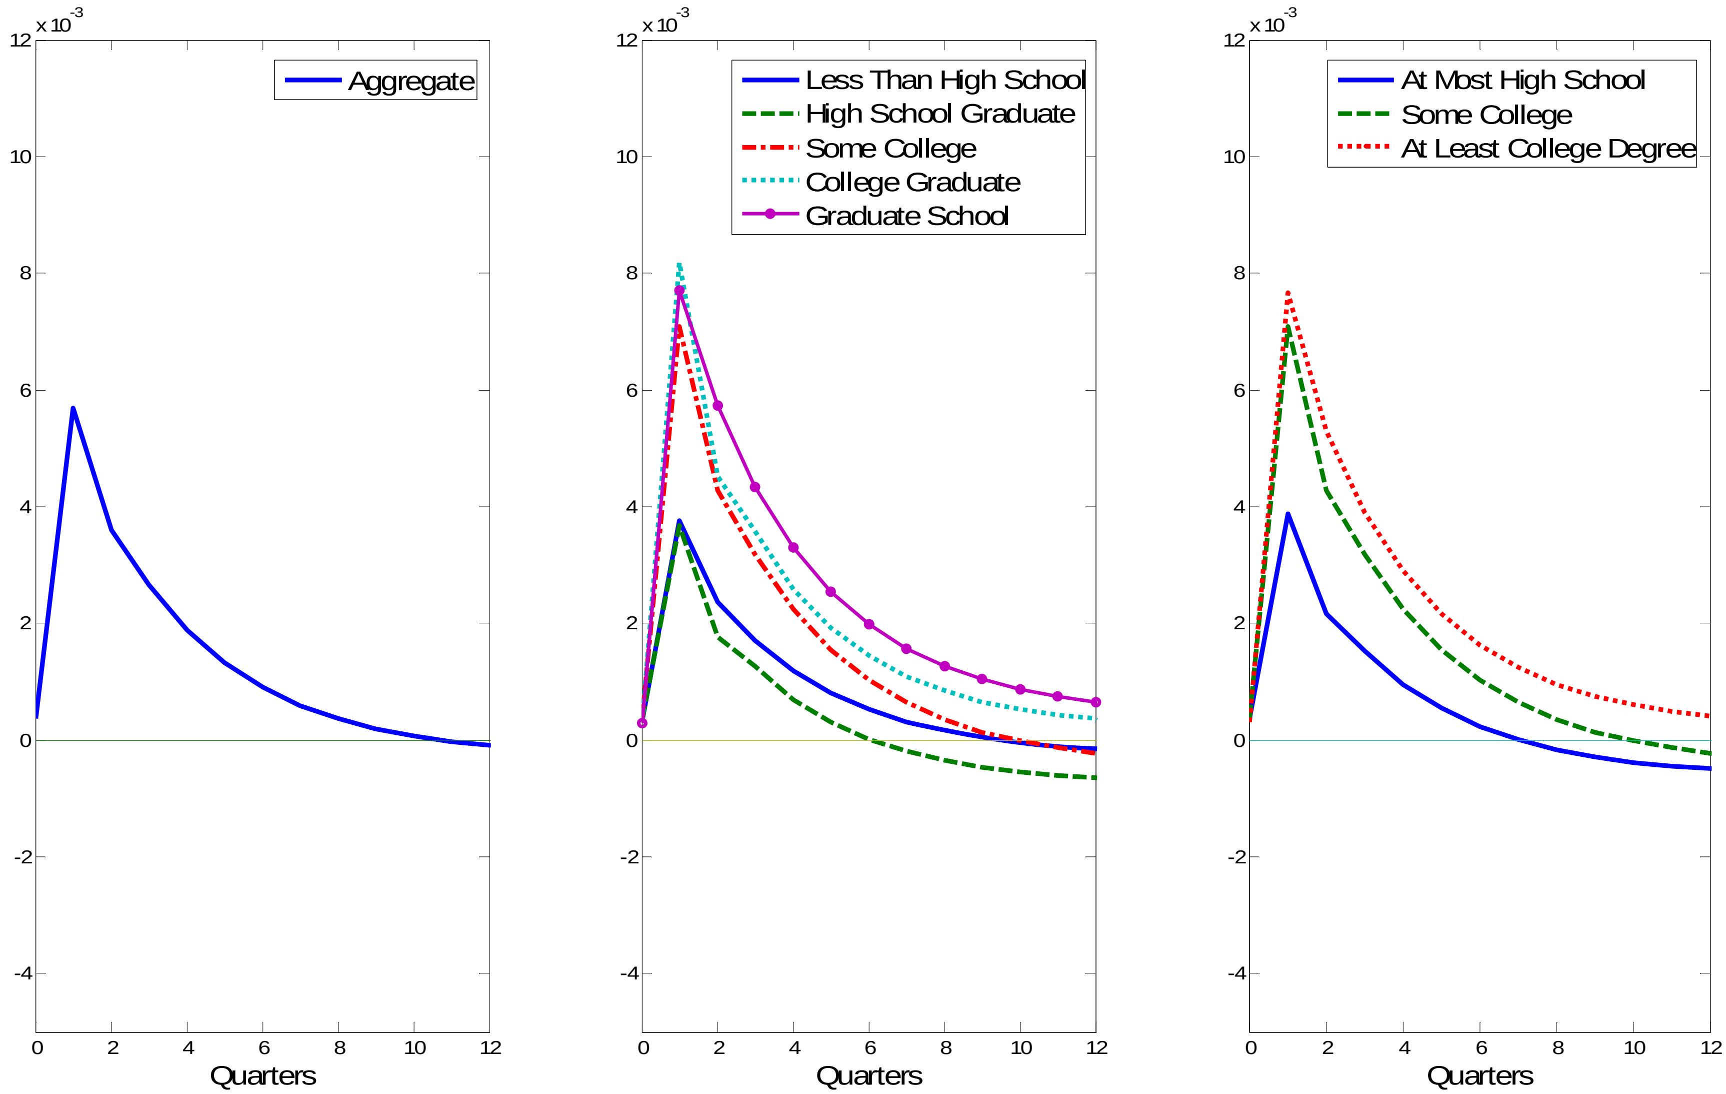

Figure 3: Inflation News Intensity Index

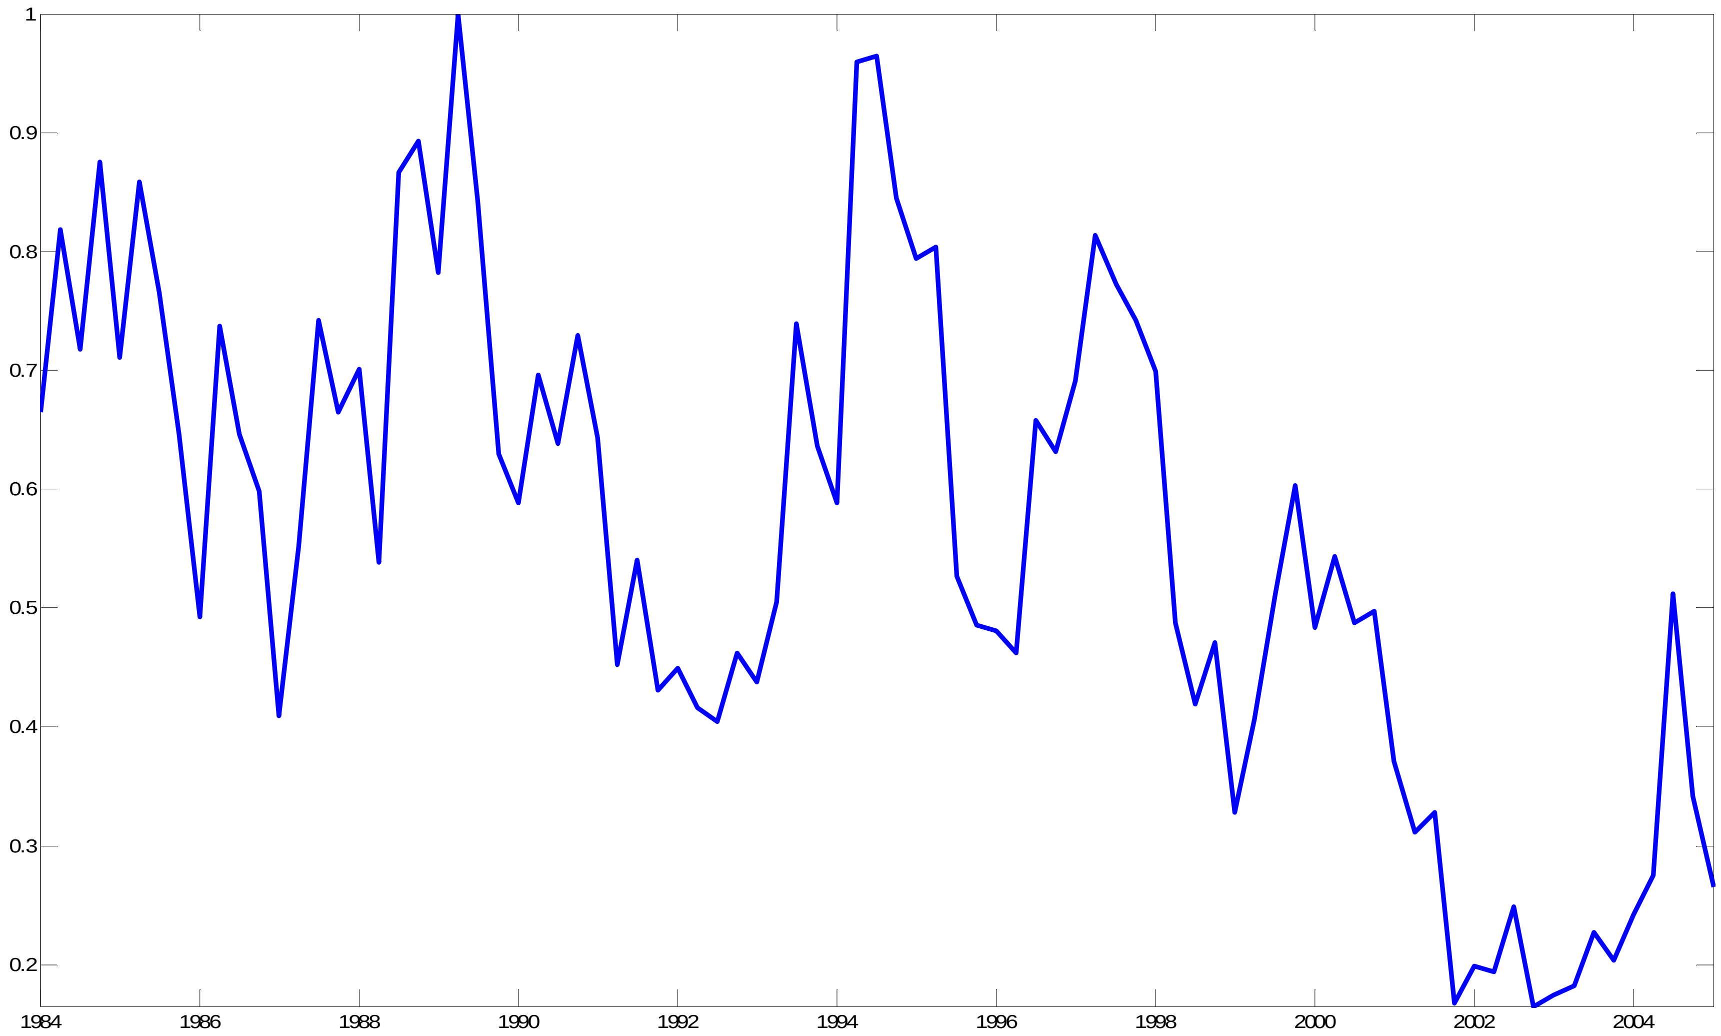

Source: Lexis-Nexis.