Economics Letters 161 (2017) 112–115

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Consumer forecast revisions: Is information really so sticky?

Carola Binder

Haverford College, Department of Economics, 370 Lancaster Ave., Haverford, PA, United States

HIGHLIGHTS

- Consumer survey respondents revise inflation expectations about four times in six months.
- Low-frequency, rounded data in previous studies leads to underestimates of revision frequency.
- Consumers with high income and education are more likely to "fine-tune" their forecasts.
- Estimates of the stickiness parameter in sticky information models may need to be reconsidered

ARTICLE INFO

Article history:
Received 24 May 2017
Received in revised form 21 September 2017
Accepted 24 September 2017
Available online 12 November 2017

JEL classification:
E31
D83
D84

Keywords:
Inflation expectations
Information rigidities
Sticky information
Consumer surveys

ABSTRACT

Previous studies using consumer survey data on inflation expectations find that consumers revise their inflation forecasts approximately once every eight months, suggesting that information is quite "sticky." However, in the consumer survey data analyzed, respondents take the survey twice with a six-month gap, and responses are reported to the nearest integer. Both the low frequency and the rounding result in overestimation of information stickiness. Higher-frequency unrounded data reveals that consumers revise their inflation expectations far more frequently—about five times in an eight month period.

0. Introduction

Household expectations play an important role in macroeconomic dynamics. Thus, central banks around the world monitor surveys of consumer expectations. Survey data reveals that consumer expectations differ notably from the full-information rational expectations benchmark, and is often used to quantify key parameters in models of information rigidities (Coibion and Gorodnichenko, 2012). A prominent such model is the sticky information model of Mankiw and Reis (2002), in which only a fraction λ of agents updates their information set each period. The model matches certain features of macroeconomic and survey data, such as the slow adjustment of consumption to shocks (Reis, 2006) and heterogeneity in consumer inflation expectations (Mankiw et al., 2004).

One approach to estimating the degree of information stickiness uses aggregate survey forecasts and a set of assumptions about macroeconomic dynamics to estimate λ parametrically. This approach implies that households update their inflation expectations approximately once per year (Carroll, 2003; Mankiw et al., 2004). A second approach uses micro-level survey data, from which a non-parametric estimate of λ_t may be obtained by the fraction of forecasters or households who revise their forecasts at each survey date t (Pfajfar and Santoro, 2013). This approach does not rely on structural assumptions or involve the computation of forecast errors, and allows for the possibility that the stickiness parameter is time-varying. The second approach results in higher estimates of λ, but still suggests that information stickiness is substantial (Drager and Lamla, 2012).

I show that several issues with estimates of information stickiness based on consumer survey microdata lead to substantial *underestimation* of the frequency with which consumers update their expectations. The first issue stems from data frequency. The rotating panel of Michigan Survey of Consumer (MSC) respondents take the survey twice with a six-month gap. A consumer may have the same forecast at months t and $t + 6$ but different forecasts in

E-mail address: cbinder1@haverford.edu.

https://doi.org/10.1016/j.econlet.2017.09.029
0165-1765/

between. The second issue is that responses are reported to the nearest integer. A consumer may update her information, but if the update results in a sufficiently small revisions, it will appear that she has not updated her information.

I use monthly panel data from the Federal Reserve Bank of New York (FRBNY) Survey of Consumer Expectations (SCE), which solicits forecasts that need not be integers. Survey respondents revise their inflation forecasts an average of four times in a six-month period. However, if I round their forecasts to the nearest integer, only around 70% of forecasts are different from the forecast made by the same respondent six months earlier, similar to the MSC data results reported in Drager and Lamla (2012) and Pfajfar and Santoro (2013). I use the SCE data to examine more closely the differences in the frequency and nature of forecast revisions across demographic groups and discuss implications for how we interpret forecast revisions.

1. The frequency of forecast revisions

The FRBNY SCE is a nationally-representative Internet-based survey launched in 2013. Approximately 1300 household heads take the survey each month and may participate for up to 12 months. The survey asks, "Over the next 12 months, do you think that there will be inflation or deflation? (Note: deflation is the opposite of inflation)", followed by, "What do you expect the rate of (inflation/deflation) to be over the next 12 months?" The survey also asks similar questions about the "12 month period between [current date +2 years] and [current date +3 years]". Note that the question wording differs from that of the MSC and was guided by a series of cognitive interviews and experiments (Bruin et al., 2010). Both surveys ask about the one-year horizon, but the MSC asks about the five- to ten-year horizon instead of the two- to three-year horizon.

Let π_{it}^{eh} denote respondent is point forecast at time t for horizon $h \in \{s, l\}$, where $h = s$ for the shorter (one-year) horizon and $h = l$ for the longer horizon. Let $\hat{\pi}_{it}^{eh}$ denote her point forecast rounded to the nearest integer. Let $R_{it,k}^{h} = \pi_{it}^{eh} - \pi_{it-k}^{eh}$ be her point forecast revision from time $t-k$ to t, and $\hat{R}_{it,k}^{h} = \hat{\pi}_{it}^{eh} - \hat{\pi}_{it-k}^{eh}$ the revision of her point forecast rounded to the nearest integer. Define dummy variables $Z_{it,k}^{h}$ and $\hat{Z}_{it,k}^{h}$ indicating that $R_{it,k}^{h}$ and $\hat{R}_{it,k}^{h}$, respectively, are nonzero. The time series variables $Z_{t,k}^{h}$ and $\hat{Z}_{t,k}^{h}$ are the means of $Z_{it,k}^{h}$ and $\hat{Z}_{it,k}^{h}$ across respondents at time t.

In the MSC data, we observe $\hat{R}_{it,6}^{s}$ and $\hat{Z}_{t,6}^{s}$. Drager and Lamla (2012) and Pfajfar and Santoro (2013) find that about three fourths of MSC respondents revise their one-year inflation expectations with respect to their first response six months prior, or in other words, that $\hat{Z}_{t,6}^{s}$ has a mean of 0.75, corresponding to a forecast update approximately every 8 months. In the SCE data, we observe $R_{it,k}^{s}$ and $R_{it,k}^{l}$ for $k \in \{1, 2, \ldots, 12\}$. I construct $\hat{R}_{it,6}^{h}$ and $\hat{Z}_{t,6}^{h}$ for the SCE respondents to see how inferences about updating frequency differ from this measure differ from inferences using $R_{it,1}^{h}$. In particular, if the sum $S_{it}^{h} = \sum_{i=0}^{5} Z_{t-j,1}^{h}$ is at least one, then respondent i has revised her unrounded point forecast at least once in the past six months. The key comparison is between $\hat{Z}_{t,6}^{h}$ and the share of respondents for whom S_{it}^{h} is at least one. Fig. 1 illustrates this comparison, while Table 1 presents relevant summary statistics. The figure plots the share for whom $\hat{\pi}_{it}^{eh}$ differs from $\hat{\pi}_{it-6}^{eh}$ along with the share for whom there has been at least one nonzero revision in π_{it}^{eh} since $t-6$. The mean of $\hat{Z}_{t,6}^{s}$ is 0.7, similar to the estimates of Drager and Lamla (2012) and Pfajfar and Santoro (2013). However, 95% revise their unrounded point forecast for short-run inflation at least once in a six-month period.[1]

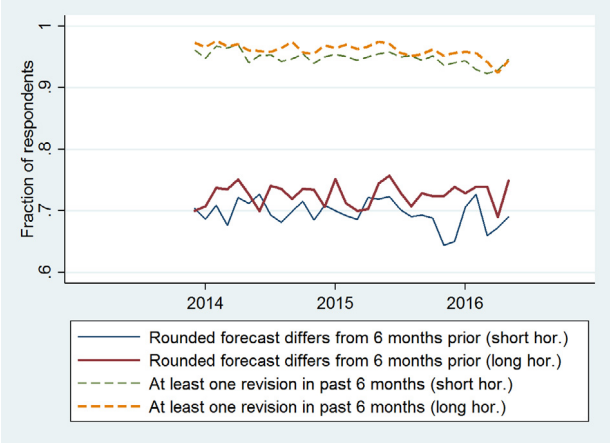

Fig. 1. Share updating at least once in six months.

Table 1
Summary statistics of forecast revision frequency.

Variable	Mean	Std. Dev.	N
$\hat{Z}_{t,6}^{s}$	0.70	0.02	30
$\hat{Z}_{t,6}^{l}$	0.73	0.02	30
$S_{it}^{s} \geq 1$	0.95	0.01	30
$S_{it}^{l} \geq 1$	0.96	0.01	30
S_{it}^{s}	3.72	0.1	30
S_{it}^{l}	3.95	0.08	30
$Z_{t,1}^{s}$	0.64	0.02	35
$Z_{t,1}^{l}$	0.69	0.02	35
$\hat{Z}_{t,1}^{s}$	0.61	0.02	35
$\hat{Z}_{t,1}^{l}$	0.66	0.02	35

The average number of revisions that a respondent makes in a six-month period is around 4 (see Fig. 2). Rather than a single revision per 8 months, the typical respondent makes at least 5. Higher-frequency, unrounded data thus reveals less information stickiness than is inferred from lower-frequency, rounded data. Most of the discrepancy arises from the frequency, rather than the rounding. Table 1 also provides the means of $Z_{t,1}^{eh}$ and $\hat{Z}_{t,1}^{eh}$, $\hat{Z}_{t,6}^{eh}$. The majority of respondents' point forecast – 94% at the shorter horizon and 95% at the longer – are integers, so the mean of $Z_{t,1}^{h}$ is only slightly higher than the mean of $\hat{Z}_{t,1}^{h}$, as illustrated in Fig. 3.

2. Demographic differences in revisions

Inflation expectations vary across demographic groups, possibly reflecting differences in information use (Meyer and Venkatu, 2011; Burke and Manz, 2014). Carroll's (2003) epidemiological foundations for the sticky information model predict that consumers who are more attuned to economic news should revise their inflation expectations more frequently. Likewise, other models with endogenously sticky information predict more frequent revisions for consumers that find it less costly to interpret information about inflation. We might therefore expect households with higher income and education to update more frequently. However, in the MSC and SCE data, $\hat{Z}_{t,6}^{h}$ is *lower* for such households. But the higher-frequency, unrounded SCE data provides a more detailed picture of how the frequency and nature of forecast revisions varies

[1] The SCE data is available for a shorter time sample than the MSC data used by Drager and Lamla (2012) and Pfajfar and Santoro (2013), but this difference does not drive the results, as SCE revisions seen in the high frequency data are more frequent than revisions inferred from the *same* SCE data viewed at lower frequency.