International Journal of Forecasting 32 (2016) 271–282

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Using time-stamped survey responses to measure expectations at a daily frequency

Frieder Mokinski

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany

ARTICLE INFO

Keywords:
Surveys
State space models
Kalman filter
Time series
Disaggregation

ABSTRACT

This article addresses one of the drawbacks of survey-based measures of expectations, the fact that updates are relatively infrequent, due to the monthly or quarterly frequency of survey waves. To obtain a more frequent measure, I present a state space method that uses time-stamped survey responses as its daily measurements and the population distribution of expectations as the latent state. Augmenting the method with financial variables that measure expectations indirectly (e.g., bond yields) facilitates measurement at times when no responses are observed. An application to a survey of German financial experts shows that my method is successful for forecasting future responses. Additional analyses show that its daily estimates react in economically plausible ways to both major events, such as the September 2008 crash of Lehman Brothers, and regular releases of economic indicators.

1. Introduction

Survey-based measures of expectations are a valuable source of economic data. As most decisions depend on expectations, these data facilitate a better understanding of the functioning of economies and provide early signals about the economic future. However, a major drawback of such data is their low frequency.

This papers develops a method for extracting a daily measure of subjective expectations from time-stamped survey responses. The method involves a state space model, in which a day's responses are the measurement variables, and the population distribution of expectations is the latent state. The method can be used both to estimate how the unobserved daily distribution of expectations has evolved over time and to forecast its future path. It is possible to update these estimates in real time to account for new responses. Augmenting the model with financial variables that measure expectations indirectly, e.g., bond yields and interest rates, facilitates the measurement at times when no responses are observed.

This novel approach is useful both for practitioners and researchers. First, it allows survey collectors to present survey results in a different way, showing how expectations have changed over a survey period and in response to particular events. Second, it facilitates real-time economic monitoring, which is relevant, for example, for central banks that are interested in inflation expectations. Third, such high-frequency measures provide new possibilities for the study of expectations formation. For example, they may shed new light on the frequency with which respondents revise their forecasts and the underlying causes of these revisions.

Using this new method, I analyze data from a monthly German financial sector survey. First, I employ the method for forecasting future survey releases. In particular, at the end of one survey wave, I generate forecasts of the average response in the following survey wave. Given that the survey collects qualitative forecasts, the forecast target is the balance statistic, which is defined as the share of positive responses minus the share of negative responses.

E-mail address: frieder.mokinski@googlemail.com.

http://dx.doi.org/10.1016/j.ijforecast.2015.06.004
0169-2070/

The empirical results suggest that both time stamps and variables that measure expectations indirectly contribute to the method's forecast accuracy. Second, I carry out additional analyses that show that the time series of daily estimates provides a plausible measure of the impacts of major events such as the September 11 attacks (see Fig. 3) or the crash of Lehman Brothers (see Fig. 4), and that regularly scheduled news announcements such as purchasing manager indexes (PMIs) exert a measurable impact on expectations at the daily frequency.

This paper is closely related to that of Ghysels and Wright (2009), who came up with the idea of measuring expectations at the daily level. Unlike my method, Ghysels and Wright use two alternative approaches that rely on *aggregated* survey data. Specifically, in their first approach, they use a MIDAS regression model to learn how expectations aggregated across a survey period are related to the daily interest rate changes of the preceding few months. Using the estimated model, they predict expectations on days with no survey responses. Alternatively, they use a state space model in the spirit of this paper, which assumes that (1) aggregate expectations are a blend of the latent expectations that were held on the days before the survey deadline date, and (2) interest rate changes convey some information about latent expectations. Thus, this paper extends the approach of Ghysels and Wright (2009) by showing how individual time-stamped responses can be used to identify daily variations in the distribution of expectations. The empirical application demonstrates that this information improves the precision of the measure greatly.

The state space method used in this paper is similar to approaches that extract measures of business conditions, see e.g. Aruoba, Diebold, and Scotti (2009) for a daily measure using indicators observed at different frequencies, or Frale, Marcellino, Mazzi, and Proietti (2011) for a measure based on a large number of indicators. Similar methods can also be found in the nowcasting literature, where the estimate of a variable referring to the ongoing month or quarter is updated continuously using new information that materializes in different indicators (e.g., Camacho & Perez-Quiros, 2010).

The methodological framework presented in this paper is potentially useful for the study of expectations formation. First, it allows the extraction of high-frequency measures of expectations that can be used to analyze the way in which economic agents incorporate new information: either all of a sudden, as would be suggested by full-information rational expectations,[1] or more gradually, as in models of expectations formation that feature information frictions, e.g., Mankiw and Reis (2002), Sims (2003) or Woodford (2001). Second, empirical tests of alternative hypotheses about expectations formation can make use of the freedom to model the state process in alternative ways: for example, Section 4 shows that the full-information rational expectations model requires the state process to be a martingale difference sequence. Alternatively, it is possible to show that models of sticky information acquisition, such as that of Mankiw and Reis (2002), imply a drifting state process. Insights from such analyses could add to the growing body of literature considering the empirical testing of new theories of expectations formation, e.g., Mankiw, Reis, and Wolfers (2004), Coibion and Gorodnichenko (2010, 2012).

[1] Pesaran and Weale (2006) provide a comprehensive survey of this literature.

2. Methodology

This section introduces measurement methods for both quantitative and qualitative expectations data. These methods rely on state space models that consider individual survey responses as draws from a population distribution of expectations. As this distribution changes over time, the goal is to estimate its shape on a given day.

The state space modeling approach is a natural choice for this problem, for various reasons. First, within its framework, it can easily be formalized that the responses on a specific day are noisy signals about the population distribution of forecasts. Second, it facilitates accounting for the variation in the signal precision that arises from the changing numbers of responses per day. Third, the approach naturally overcomes the hurdle of the missing values that are observed when no survey responses are collected. Fourth, it is straightforward to augment the model with additional measurement variables, such as financial asset prices, and to relate them to the dynamics of the population distribution of expectations.

2.1. Quantitative expectations

The model presented in this section consists of a measurement equation that relates the observed responses to their population distribution and a state equation that governs how the population distribution changes over time. Specifically, I assume that the population distribution of responses on day t is

$$y_{it} \overset{i.i.d.}{\sim} \mathcal{N}\left(\mu_t\,,\ \sigma_\varepsilon^2\right).$$

Thus, the average of the n_t responses observed on day t can be represented as

$$\bar{y}_t = \mu_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}\left(0\,,\ n_t^{-1}\sigma_\varepsilon^2\right), \tag{1}$$

which constitutes the measurement equation. The variance of the disturbance being proportional to n_t^{-1} reflects the fact that the signal's precision about the latent mean of the population distribution increases with the number of responses. It now remains only to specify the state equation that governs the dynamics of μ_t. A very simple specification is

$$\mu_{t+1} = \mu_t + \eta_{t+1}, \quad \eta_{t+1} \overset{i.i.d.}{\sim} \mathcal{N}\left(0\,,\ \sigma_\eta^2\right), \tag{2}$$

such that μ_t evolves as a random walk. Although this can be replaced easily, random walk dynamics can be justified on the grounds that alternative specifications, such as a stationary autoregression or a random walk with drift, imply that today's best estimate of the expectations that will be held tomorrow differs systematically from the expectations held today. In contrast, rational expectations (Muth, 1961) require such changes to be unpredictable.