A Novel Approach to Measuring Consumer Confidence

Bert de Bruijn[1]

Rene Segers[1,2]

Philip Hans Franses[1]

[1] *Econometric Institute and Tinbergen Institute, Erasmus University Rotterdam*

[2] *Gibbs Quantitative Research & Consulting*

Econometric Institute Report 2014-30

Abstract

This paper puts forward a new data collection method to measure daily consumer confidence at the individual level. The data thus obtained allow to statistically analyze the dynamic correlation of such a consumer confidence indicator and to draw inference on transition rates. The latter is not possible for currently available monthly data collected by statistical agencies on the basis of repeated cross-sections. In an application to measuring Dutch consumer confidence, we show that the incremental information content in the novel indicator helps to better forecast consumption.

Keywords: Consumer confidence, Randomized sampling, Markov transition model, consumption.

JEL Classification Codes: C33, C42, C81, E20

1. Introduction

Consumer Confidence Indicators (CCIs) are often regarded as useful variables to measure the current state of the economy as well as to forecast its future states at reasonably short horizons, see Ludvigson (2004) for an assessment. Most industrialized countries report such indicators at a monthly level. Typically consumer confidence is measured by surveying one thousand or more individuals each month. The individuals are asked whether they believe that their situation has improved in the previous period or will improve in the next period, while focusing on their financial situation, employment, and, for example, their purchases of durable and more expensive products in particular. The answer categories are (very) positive, neutral, and (very) negative, and their origin goes back to Katona (1951). The total indicator is constructed by subtracting the percentage of negative answers from the percentage of positive answers. Many countries also report more specific indicators, which are confined to the financial position or employment only. Publicly available data are published in original format as well as after seasonal adjustment.

Despite their widespread use and interpretation, it can be of interest to investigate if the way consumer confidence is measured can be improved. One research angle can concern the questions asked and the way indicators are constructed from these questions. One may for example consider replacing the traditional qualitative questions by probabilistic questions inquiring about more well-defined events, as suggested in Dominitz (2004). Also the fact that consumer confidence data show signs of seasonality can be viewed as inconvenient, and perhaps a rephrasing of the questions can overcome this potential drawback.

A second angle for potential improvement of consumer confidence indicators would be to better understand how consumer confidence varies across individuals with different socioeconomic and demographic characteristics. These insights could be exploited to reduce sampling error due to the use of small and possibly unrepresentative samples in the data collection stage, which improves the reliability of the indicators. We believe that improvement in these two directions can be relevant, but the third research angle to be discussed next seems more promising.

A third angle is that consumer confidence data are usually so-called repeated cross-sectional data. That is, each month approximately one thousand individuals are interviewed, but each month this concerns one thousand *different* individuals. A major consequence of this way of collecting data is that developments over time are difficult to interpret. Basically, when an indicator is, say, -18 in December, while it was -21 in November, we must conclude that the average fraction of more negative answers in

December was smaller than in November. We could even say that in December consumer confidence has increased with 3 points, but we must be aware that this does not concern the same individuals. Hence, an interpretation of a sequence of monthly consumer confidence levels is prone to the so-called ecological fallacy. This fallacy concerns the situation where we seek to derive micro behavior from aggregated data. In the literature there are various suggestions to circumvent or solve this problem, see King (1997), Moffitt (1993), Sigelman (1991), and the collection of papers in King (2004), among many others. In the present paper we seek to do that, but now by applying an alternative method of data collection.

In this paper we put forward a method to collect high-frequency consumer confidence data at the individual level. We keep the Katona-type questions intact, but we merely focus on the collection and analysis of the data, trying to prevent ecological fallacy. To that end, we collect data such that we have the same (though not all) individuals being interviewed from one week to another, to prevent respondents from becoming annoyed or uninterested. In order to statistically analyze the dynamic correlation of our CCI and to draw inference on transition rates, we develop a Markov transition model. The model describes the persistence in an individual's confidence level. We exploit the Markov transition properties of the model to estimate an expected response for each individual on the days on which the individual did not complete the questionnaire. This enables us to compute a daily consumer confidence indicator as if the entire panel was surveyed every day. To illustrate the usefulness of our approach, we employ the indicator to forecast Dutch consumption. We show that the incremental information content in the novel indicator helps to improve forecasting accuracy.

The outline of the paper is as follows. In Section 2 we present our method of data collection, and we argue that it has various convenient properties for the purpose of measuring consumer confidence. Next, in Section 3 we introduce the Markov transition model that will be used to describe longitudinal developments in consumer confidence at the individual level. In Section 4 we illustrate the usefulness of our method by forecasting Dutch household consumption using the novel consumer confidence indicator. In Section 5 we conclude with an outline of various areas for further research.

2. Methodology

To measure developments in consumer confidence over time it is desirable to conduct a longitudinal or panel study where the same individuals are surveyed at multiple points in time. This allows us to study developments in confidence at the individual level. However, surveying the very same individuals frequently likely deteriorates the quality of the survey. People get irritated and they disconnect from the

panel, thereby making the panel less efficient. Or perhaps worse, respondents' (reported) confidence levels may change due to being a member of a panel, which is called panel conditioning. Most statistical agencies therefore collect repeated cross-sections instead of panel data. This amounts to surveying a new group of individuals at each survey occasion, which implies that individuals are surveyed only once. The design is illustrated in Figure 1 (a). Here we index time by t, where $t = 1, ..., T$, individuals by i, where $i = 1, ..., N$ and groups of individuals by g, where $g = 1, ..., G$. A grey square in row g and column t indicates that group g is requested to be surveyed at time t. While repeated cross-sections reduce respondent burden and eliminate potential panel conditioning bias, developments at the individual level cannot be derived without making many assumptions, see the excellent treatment in Moffitt (1993). Therefore it seems promising to collect longitudinal data nevertheless, but to choose the design of the panel carefully such that the adverse effects of repeated interviewing are negligible or, at least, manageable.

To design a panel for a specific purpose, three decisions have to be made. Firstly, as individuals cannot be surveyed continuously, one has to decide on the total time-span that a panel member is requested to join the panel, to be denoted by T^*. To keep the total number of panel members constant, one may decide to invite new individuals to join the panel when existing panel members disconnect from the panel. This strategy is referred to as rotation, see Patterson (1950) and Kish & Hess (1959). Naturally, the next step is to decide upon the number of survey requests within this period, to be labeled n. Note that T^* and n together constitute the sampling frequency $f = n/T^*$ of the survey, which is equal to the reciprocal of the time between subsequent survey occasions, or waves. Thirdly and finally, one needs to decide when to conduct the n surveys within the time-span T^*. We will refer to this aspect as date selection. A natural way is to divide the time-span T^* into n equally long time periods, and to survey around the beginning of each sub period. Typically in this case the implied sampling frequency f is lower than the desired data frequency. Again one may therefore apply rotation, such that at each point in time t a new group of panel members is surveyed and the data are collected continuously. The above strategy is mostly referred to as time sampling. As an alternative, Segers and Franses (2014) proposed to choose the n survey occasions at random, independently for each panel member. They show that in this case of randomized sampling, data is collected to measure every possible autocorrelation up to $T^* - 1$ lags, where the lower lag orders are sampled most frequently. This facilitates the identification of any type of individual dynamics in the data and it allows for efficient estimation.

To the best of our knowledge, the only consumer confidence indicator that is not obtained from repeated cross-sections is the Index of Consumer Sentiment of the University of Michigan. Michigan adopts a rotating panel design in which the respondents are requested to be re-interviewed six months

after the first interview, see Curtin (1982) for details. This design is illustrated in Figure 1 (b). In our terminology, we would characterize the Michigan panel as a rotating panel where $T^* = 12$ months and $n = 2$ survey occasions per individual and time sampling is applied.

-- Insert Figure 1 about here --

An example of our preferred randomized rotating panel, where two new individuals are invited to join the panel in each time period, is shown in Figure 1 (c). In this example, we set the maximum time-span that a panel member is requested to join the panel, T^*, equal to 8 and the number of survey requests, n, equal to 4. As a consequence, the sampling frequency f is 0.5. Each dotted area encloses all survey requests assigned to one particular cohort of individuals.

3. Modeling Consumer Confidence

Notation and motivation

Typically, consumer confidence surveys are composed of 3 to 10 questions. Consumers are asked whether they believe that the economic conditions in their country have improved in the previous period or will improve in the next period. Often, similar questions are then posed related to the respondents' believes about their private financial situation. In practice, the number of answer categories varies from 3 (negative, neutral or positive) to 201 (-100, 0, +100). To obtain one score per individual, the negative answers are then deducted from the positive answers, or the answers are simply averaged, depending on the design of the survey.

To convert the survey data to a daily measure of consumer confidence, we will need to use an econometric model, as we will argue below. We aim to keep the model specification flexible, such that the model can be easily adopted to match a wide range of consumer confidence surveys. First, let us denote the response of individual i on the j-th survey question on day t by $z_{i,j,t}$, with i and t as before, and $j = 1, \dots, J$, where J denotes the total number of questions. Second, let $z_{i,t}$ be the overall confidence score of an individual, which is generally obtained by simply taking the sum or the average of the J answers provided by the individual i at time t. We assume that $z_{i,t}$ is scaled such that all scores are in the closed interval $[-1, \dots, 1]$, with -1 being the most negative response and 1 being the most positive response

5

possible. Finally, but not necessarily, the data is then further reduced by classifying an individual as being in a negative, neutral or positive state of confidence, using

$$q_{i,t} = 1 \; if \; z_{i,t} > \tau$$
$$q_{i,t} = 0 \; if \; |z_{i,t}| \leq \tau \tag{1}$$
$$q_{i,t} = -1 \; if \; z_{i,t} < -\tau$$

where $q_{i,t}$ is the final classification of the individual and τ is a threshold parameter. We collect the values of $q_{i,t}$ in the $I \times T$ matrix Q, defined as

$$Q = \begin{bmatrix} q_{1,1} & q_{1,2} & \cdots & q_{1,t} & \cdots & q_{1,T} \\ q_{2,1} & q_{2,2} & \cdots & q_{2,t} & \cdots & q_{2,T} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ q_{i,1} & q_{i,2} & \cdots & q_{i,t} & \cdots & q_{i,T} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ q_{I,1} & q_{I,2} & \cdots & q_{I,t} & \cdots & q_{I,T} \end{bmatrix} \tag{2}$$

To obtain a daily consumer confidence index c_t, one may calculate the average of the collected $q_{i,t}$'s. This would be equal to the proportion of positive responses minus the proportion of negative responses. The matrix of responses Q is sparse because of two reasons. First, by design, in each wave only a limited number of respondents is requested to be surveyed. Second, respondents may decide not to participate in the survey. This implies that taking the average of $q_{i,t}$ would yield estimates with large measurement error. We can produce more accurate daily averages if we impute the missing values in Q, before taking the averages over all respondents. To impute the missing values in Q, we employ a Markov transition model.

The Markov Transition model

As we want to focus on the number of positive and negative responses while also incorporating dynamics, we use a Markov Transition model, see Ross (2007) among others, for a detailed assessment. This model uses transition probabilities to reflect the probability of transferring from one state of confidence to the next. We jointly denote these probabilities by the transition matrix P. For the situation with a positive, neutral and negative state, we define P as

$$P = \begin{bmatrix} p_{1,1} & p_{1,0} & p_{1,-1} \\ p_{0,1} & p_{0,0} & p_{0,-1} \\ p_{-1,1} & p_{-1,0} & p_{-1,-1} \end{bmatrix} \tag{3}$$

The elements $p_{l,m}$ correspond to the probability of transitioning in one day from state l to state m: $\mathbb{P}(q_{i,t} = m | q_{i,t-1} = l)$, with $l, m \in \{-1, 0, 1\}$. The k-day transitioning matrix P_k can now simply be obtained by taking the k-th power of P: $P_k = P^k$. We denote its elements by $p_{k,l,m}$. Note that it is straight-forward in this setup to include more states. This allows to also capture slightly positive and negative states of confidence.

Given the elements of the matrix P, we can estimate the probability that individual i is in state l at time t, to be denoted by $\pi_{i,t,l}$. For this purpose, we need to administer two additional variables, for each $q_{i,t}$. The first is the number of days that have passed since the last survey response of individual i at time t, to be denoted by Δt. Further, denote the day of the last response by $s = t - \Delta t$. The second is the value of the state $q_{i,s}$. Then

$$\mathbb{P}(q_{i,t} = l) = \pi_{i,t,l} = p_{\Delta t, q_{i,s}, l} = [P^{\Delta t}]_{q_{i,s}, l}. \tag{4}$$

In order to estimate the elements of P, we use a maximum likelihood approach (see Cameron and Trivedi, 2005, among others), which we set up as follows. First, we construct the individual likelihood contribution $\mathcal{L}_{i,t}$. Given an estimate $\hat{P}$, we can calculate an estimate of $\pi_{i,t,l}$, using (4). But for some combinations of i and t, $q_{i,t}$ is known. The likelihood of observing such a single data point is captured by $\mathcal{L}_{i,t} = \prod_{l=-1}^{1} \pi_{i,t,l}^{I[q_{i,t}=l]}$, which is equal to 1 if there is no response for individual i at day t or if the response $q_{i,t}$ is not state l. The latter is equal to the estimated state probability $\pi_{i,t,l}$ if $q_{i,t} = l$. Next, we find the total likelihood by multiplication of the individual contributions, for all time periods and for all individuals, that is,

$$\mathcal{L} = \prod_{i=1}^{I} \prod_{t=1}^{T} \mathcal{L}_{i,t} = \prod_{i=1}^{I} \prod_{t=1}^{T} \prod_{l=-1}^{1} \pi_{i,t,l}^{I[q_{i,t}=l]} \tag{5}$$

For numerical stability and feasibility, we maximize the logarithm of $\mathcal{L}$ for reasons of numerical stability and feasibility, which is given by

$$logℒ = \sum_{i=1}^{I}\sum_{t=1}^{T} logℒ_{i,t}$$

$$= \sum_{i=1}^{I}\sum_{t=1}^{T}\sum_{l=-1}^{1} I[q_{i,t} = l]\, log\, \pi_{i,t,l} \tag{6}$$

We need to maximize $logℒ$ over the elements of matrix P. For each row of P it is the case that if two of the elements of that row are known, the third element can be calculated because the elements must sum to 1. The number of parameters over which we must maximize this function is thus equal to six.

Once consistent estimates of the model parameters are obtained, we impute the missing values in the matrix Q by the difference between the probability of being in the positive and the probability of being in the negative state, at each wave, that is,

$$\hat{q}_{i,t} = \mathbb{P}(q_{i,t} = 1) - \mathbb{P}(q_{i,t} = -1) = \hat{\pi}_{i,t,1} - \hat{\pi}_{i,t,-1} \tag{7}$$

The value of the aggregated consumer confidence index, c_t, is then equal to the average of these state estimates:

$$c_t = \frac{1}{I}\sum_{i=1}^{I}\hat{q}_{i,t}. \tag{8}$$

4. Empirical Illustration

We illustrate the usefulness of our method of data collection and the Markov transition model by showing how the data collection method and model can be employed to produce forecasts of household consumption.

Data collection

To collect consumer confidence data at the individual level using our data collection method, we invited all alumni who graduated as a BSc or MSc at our institute [*Name censored to facilitate the review process*] to join an expert panel specifically set up for the occasion. A first group of alumni was invited to join the panel in January 2010, and a second group was invited to join the panel in September 2011. The alumni

who accepted the invitation were requested to complete a short questionnaire at most four times per year. The invitations were sent out every week on Monday at 9 am. In order to be able to compare our indicator to the official CCI of Statistics Netherlands (SN), we used the same questionnaire as SN. The survey consists of five questions regarding the economic climate in The Netherlands and the respondents' willingness to buy durable goods, see Appendix I for details. The number of answer categories (7 or 21) as well as the sampling strategy (time-sampling or randomized) was varied across the respondents. Over the period January 4, 2010 up to October 13, 2014, 250 alumni participated in the experiment. We sent out 3625 survey requests. The response rate was 33.9%. We fixed the threshold parameter τ in Equation (1) at 0.1, which resulted in three categories of comparable size.

In-sample results

We have estimated the transition probabilities $p_{l,m}$ by maximizing the likelihood function as discussed in Section 3. The results are shown in Table 1. The day-to-day transition matrix is quite similar to the identity matrix, which suggests that a respondent's state of confidence typically does not change on a daily basis. Therefore, in the second panel of Table 1, we also report the transitions over a period of 90 days. 73.2% of the respondents who are optimistic about the economy remain positive after 90 days, while 12.8% becomes neutral and 14.1% becomes pessimistic. Staying in the neutral or the pessimistic state after 90 days is more unlikely than staying in the optimistic state with probabilities of 38.7% and 42.1%, respectively. Finally, in the third panel of the table we report the steady state distribution. In the long-run, 41.8% of our respondents is optimistic, 27.9% is optimistic and 30.4% is pessimistic about the economy. These implications on the longer horizon seem realistic in the context of our data, as the number of responses is approximately the same for each state of confidence.

-- Insert Table 1 about here --

Construction of the index

Next, we discuss how we can use our estimates of c_t to forecast monthly consumption in the Netherlands. We aim to compare our forecasts to the forecasts one would obtain using the official CCI of SN, to be denoted by *CCI_SN*. SN only collects their data during the first ten working days of each month. To allow for a fair forecast comparison, we construct our monthly CCI from the daily CCI by taking the value obtained on the tenth working day of each month only. We denote our monthly CCI by *CCI_MS*.

Besides our two main indicators, CCI_{NS} and CCI_{MS}, we consider two additional indicators, which cannot be observed if the data are collected using repeated cross-sections. The first aims to capture the spread in confidence across individuals, and is defined as the time-varying standard deviation of $\hat{q}_{i,t}$:

$$S_{MS,t} = \sqrt{\frac{1}{I-1}\sum_{i=1}^{I}(\hat{q}_{i,t} - c_t)^2} \tag{10}$$

The second indicator measures the rate at which respondents have recently changed their state of confidence. We define this rate of change as the aggregated absolute 30 day change in $\hat{q}_{i,t}$:

$$A_{MS,t} = \sum_{i=1}^{I}|\hat{q}_{i,t} - \hat{q}_{i,t-30}|. \tag{11}$$

Figure 1a. shows the developments in both the CCI_{MS} and the CCI_{SN} over the period April 2010 to October 2014. First of all, note that the average level of the novel indicator is about 43 points higher than that of SN. This is most probably due to the fact that our panel consists of university alumni only, who are generally positive about their financial situation, while SN invites aims to select a representative sample of the population. However, the cycles of the two indicators are very similar, with the most prominent turning point being the moment when the impact of the European debt crisis became apparent. Finally, we note that there tends to me more short-term variation in the indicator of Statistics Netherlands, compared to the model based indicator. Figure 1b. and 1c. display the spread, $S_{MS,t}$, and the activity, $S_{MS,t}$, of the novel indicator. Both indicators show an increase during the freefall of the index in the summer of 2011 and peak in September 2011, but also distinct fluctuations that do not necessary coincide with developments in the index.

Forecasting comparison

The target variables in our forecast comparison are Dutch national household consumption and its four subcomponents. Household consumption is measured by SN and is subdivided into SERVICES and GOODS, where GOODS are further subdivided into FOOD, DRINK AND TOBACCO, DURABLES and OTHER GOODS. DURABLES are defined as goods that in principle last more than one year, such as clothes and textiles, shoes,

10

furniture, consumer electronics and cars. OTHER GOODS include energy, motor fuels and all other consumer goods that cannot be classified as FOOD, DRINK AND TOBACCO, or DURABLES.

To forecast consumption one month ahead, we consider three basic forecasting models. In the first we only include the level of the consumer confidence indicator of SN, CCI_{SN}, and its first difference, alongside with a constant term and AR(1) dynamics:

$$y_t = \alpha + \beta CCI_{SN,t-1} + \gamma \Delta CCI_{SN,t-1} + \phi y_{t-1} + \varepsilon_t \qquad (12)$$

In the second forecasting model, we replace CCI_{SN} by our index CCI_{MS}:

$$y_t = \alpha + \beta CCI_{MS,t-1} + \gamma \Delta CCI_{MS,t-1} + \phi y_{t-1} + \varepsilon_t \qquad (13)$$

Finally, the third forecasting model is equal to (13) but now also includes the spread and activity indicators, S_t and A_t in (10) an (11):

$$\begin{aligned} y_t = \alpha &+ \beta_1 CCI_{MS,t-1} + \gamma_1 \Delta CCI_{MS,t-1} + \beta_1 S_{MS,t-1} \\ &+ \gamma_1 \Delta S_{MS,t-1} + \beta_1 A_{MS,t-1} \\ &+ \gamma_1 \Delta A_{MS,t-1} + \phi y_{t-1} + \varepsilon_t \end{aligned} \qquad (14)$$

are all estimated using Ordinary Least Squares.

In the first set of columns in Table 2, we display the R^2 values of models (12), (13) and (14). All models for the components FOOD, DRINK AND TOBACCO and TOTAL CONSUMPTION explain less than 10% of the variation in the data. This suggests that confidence indicators are less useful for explaining the latter two components. In contract, SERVICES is explained best by the models, with a R^2 of 0.93 for models (12) and (13) and 0.94 for models (14). This series is relatively persistent, as demonstrated by the estimated value of the AR parameter ϕ (0.89, p-value: 0.08), see columns 4 and 5. We then test whether the additional four parameters of model (14) relative to model (13) are jointly significant using an F-test. The corresponding p-values are shown in column 6. For the components DURABLE GOODS, OTHER GOODS and SERVICES, (13) is rejected in favor of (14) at the 10% significance level. This suggests that the additional components S_t and A_t have indeed explanatory power.

-- Insert Table 2 about here –

Finally, we compare the forecasting performance of models (12), (13) and (14) in Table 3. The first column displays the ratio of the forecast error of (13) to the forecast error of (12). As a result, values below 1 indicate that (13) outperforms (12). Since on average, however, the ratio is close to 1, we conclude that the level of the novel indicator CCI_{MS} performs approximately as good as model (12), which contains the indicator CCI_{SN}. As both models contain the same number of parameters, this indicates that no clear decision between both models can be made. The next column displays the ratio of the forecast error of (14) to the forecast error of (12). Model (14) has relative forecast error variances that are 20% smaller than those of (12), on average. Especially the forecasts for components DURABLE GOODS, OTHER GOODS and SERVICES benefit from the use of the novel indicator and its additional components S_t and A_t.

-- Insert Table 3 about here –

5. Conclusions

In this paper we proposed to collect randomized panel data rather than repeated cross-sections to measure consumer confidence. Randomized panel data allow us to not only observe longitudinal changes in confidence across our respondents, but also to observe changes in the spread in confidence across individuals, as well as changes in rate at which respondents change their state of confidence over time.

We demonstrated the usefulness our approach in an application to measuring consumer confidence in The Netherlands over the period April 2010 to October 2014 with the purpose to produce one-month ahead forecasts of Dutch consumption. We showed that the incremental information content in the novel panel data indicator improves the forecasting accuracy by 20% on average.

We mention various directions for further research. The first is that we can now correlate significant weekly changes with weekly observed macroeconomic variables, in order to study whether consumer confidence has predictive value. Indeed, currently most such studies concern monthly observed cross-sectional data, and it may well be that substantial information is lost. We may also consider enriching the model with various explanatory variables that do correlate with consumer confidence but do not necessarily correlate with the target variable. For example, it has been hypothesized that factors such as the weather and specific events such as terrorist attacks have an impact on consumer confidence.

An indicator that is corrected for one or more of these factors might be a better predictor of the future course of the economy than the uncorrected indicators that are currently in use. Finally, it might be a worthwhile endeavor to extend the Markov transition model. For example, we have assumed constant transition probabilities, while in practice transition probabilities might be time-varying.

References

Ansari, A., C. F. Mela, and S. A. Neslin (2008), Customer Channel Migration, Journal of Marketing Research, 45, 60-76.

Cameron, A.C. and P.K. Trivedi (2005). Maximum Likelihood (Section 5.6). In Cameron, A.C. and P.K. Trivedi (Eds.), Microeconometrics: Methods and Applications (1st ed., pp. 139-146). New York, NY: Cambridge University Press.

Curtin, R. (1982), Indicators of Consumer Behavior: The University of Michigan Surveys of Consumers, Public Opinion Quarterly, 46, 340-352.

Dominitz, J. and C. F. Manski (2004), How Should We Measure Consumer Confidence?, Journal of Economic Perspectives, 18, 51-66.

Erdem, T. and B. Sun (2001), Testing for Choice Dynamics in Panel Data, Journal of Business & Economic Statistics, 19, 142-152.

Franses, P. H. B. F. and R. Paap (2001), Quantitative Models in Marketing Research, Cambridge University Press, Cambridge, UK.

Hansen, R. A. (1980), A Self-Perception Interpretation of the Effect of Monetary and Nonmonetary Incentives on Mail Survey Respondent Behavior, Journal of Marketing Research, 17, 77-83.

Heckman, J. J. (1981), The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time - Discrete Data Stochastic Process, in C. F. Manski and D. McFadden (eds.), Structural Analysis of Discrete Data with Econometric Applications, MIT Press, Cambridge, pp. 179-195.

Katona, G. (1951), Psychological Analysis of Economic Behavior, McGraw Hill, New York.

Keane, M. P. (1997), Modeling Heterogeneity and State Dependence, Journal of Business & Economic Statistics, 310-327.

King, G., O. Rosen, and M. A. Tanner (2004), Ecological Inference: New Methodological Strategies, Cambridge University Press, New York.

King, G. (1997), A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data, Princeton University Press, Princeton.

Kish, L. and I. Hess (1959), A "Replacement" Procedure for Reducing the Bias of Nonresponse, The American Statistician, 13-4, 17-19.

Little, R. J. A. and R. B. Rubin (2002), Statistical Analysis with Missing Data, second edition., John Wiley & Sons.

Ludvigson, S. C. (2004), Consumer Confidence and Consumer Spending, Journal of Economic Perspectives, 18, 29-50.

McKelvey, R. D. and W. Zavoina (1975), A Statistical Model for the Analysis of Ordinal Level Dependent Variables, Journal of Mathematical Sociology, 4, 103-120.

Moffitt, R. (1993), Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross-sections, Journal of Econometrics, 59, 99-123.

Patterson, H. D. (1950), Sampling on Successive Occasions with Partial Replacement of Units, Journal of the Royal Statistical Society Series B, 49, 241-255.

Ross, S. M. (2007). Markov Chains (Ch. 4). In S. M. Ross (Ed.), Introduction to Probability Models (9[th] ed., pp. 185-280). Amsterdam: Elsevier.

Schafer, J. L. (1997), Analysis of Incomplete Multivariate Data, Chapman and Hall, New York.

Segers, R. and P. H. Franses (2014), Panel Design Effects on Response Rates and Response Quality, Statistica Neerlandica, 68, 1-24.

Sigelman, L. (1991), Turning Cross Sections into a Panel: A Simple Procedure of Ecological Inference, Social Science Research, 20, 150-170.

Wooldridge, J. M. (2005), Simple Solutions to the Initial Conditions Problem in Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity, Journal of Applied Econometrics, 20, 39-54

Appendix I. The Consumer Confidence Survey of Statistics Netherlands

As opposed to the consumer confidence indicator measured by the European Commission, the indicator measured by Statistics Netherlands not only concerns consumers' opinions on their financial situation, the economy in general, willingness to save and unemployment in the next twelve months, but also consumers' present situations and their opinions on the previous twelvemonths. The two indicators show roughly the same developments overtime.[1]

Consumer confidence is based on five questions from a more elaborate consumer survey. These questions are subdivided into a section on the economic climate and a section on the respondent's willingness to buy. The questions are formulated as follows:

Economic Climate

1. How do you think the general economic situation in this country has changed over the last twelve months?

Possible answers: At present, it is better (1) / the same (0) / worse (-1)

2. How do you think the general economic situation in this country will develop over the next twelve months?

Possible answers: It will be better (1) / the same (0) / worse (-1)

Willingness To Buy

3. How does the financial situation of your household now compare to what it was twelve months ago?

Possible answers: At present, it is better (1) / the same (0) / worse (-1)

4. How do you think the financial situation of your household will change over the next twelve months?

---

[1] See http://www.cbs.nl/en-GB for details.

Possible answers: It will be better (1) / the same (0) / worse (-1)

5. Do you think that at present there is an advantage for people to make major purchases, such as furniture, washing machines, TV sets, or other durable goods?

Possible answers: Yes, now it is the right time (1) / It is neither the right nor the wrong time (0) / No, it is the wrong time (-1)

The economic climate indicator is computed as the summed score of questions 1 and 2 averaged across all individuals. Similarly, the willingness to buy indicator is computed as the summed score of questions 3 to 5 averaged across all individuals. Finally, consumer confidence is defined as the average of all five scores.

Tables and figures

Figure 1. Panel designs to measure consumer confidence

a. Repeated cross-sections

b. The Michigan-panel

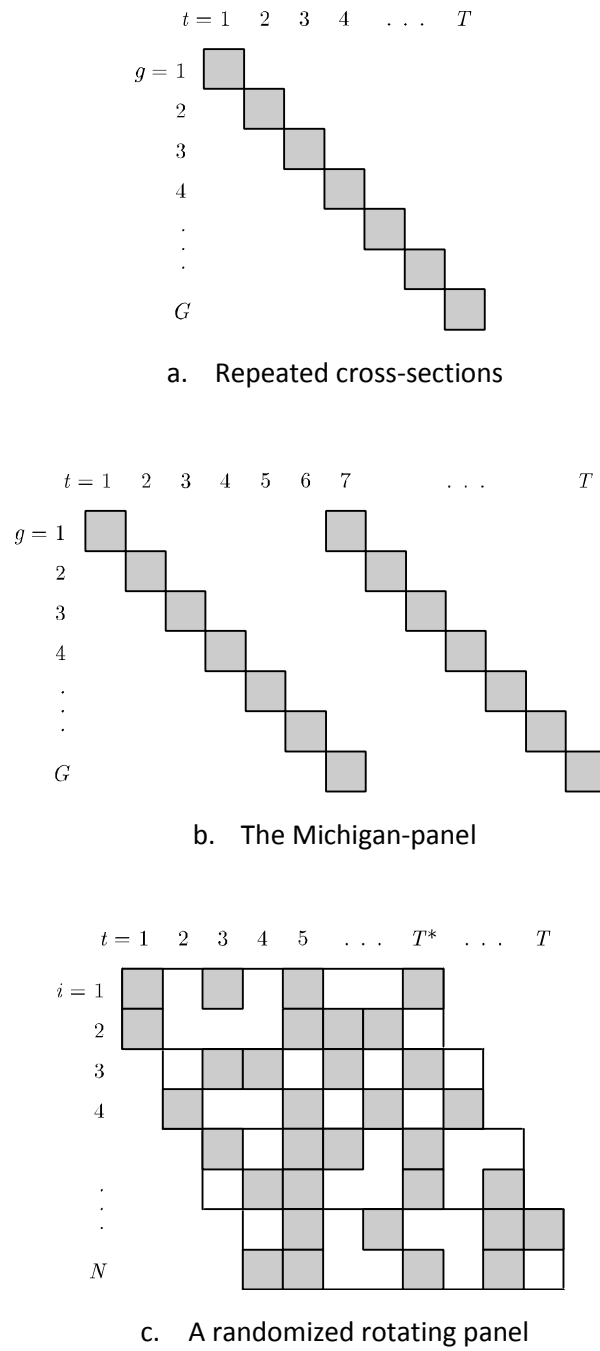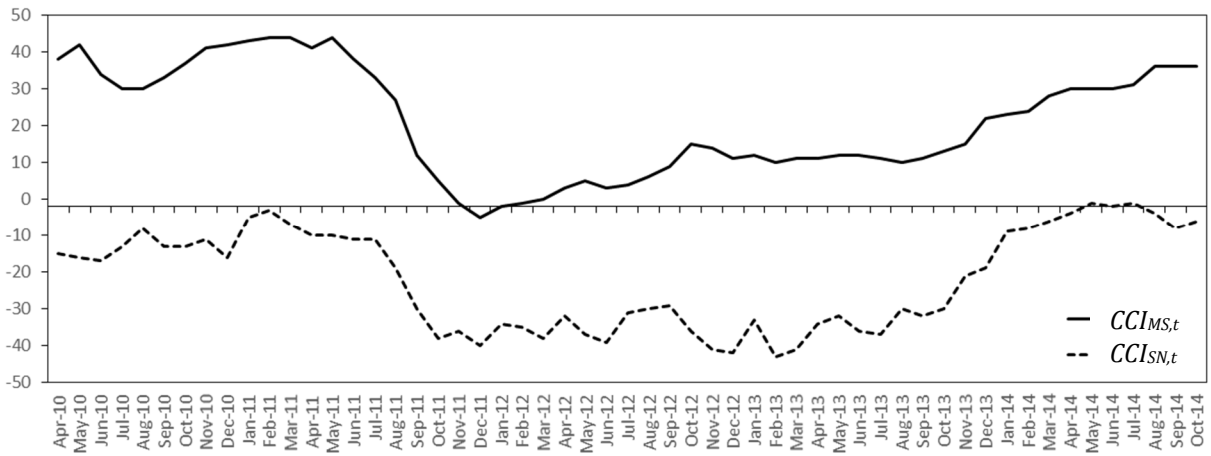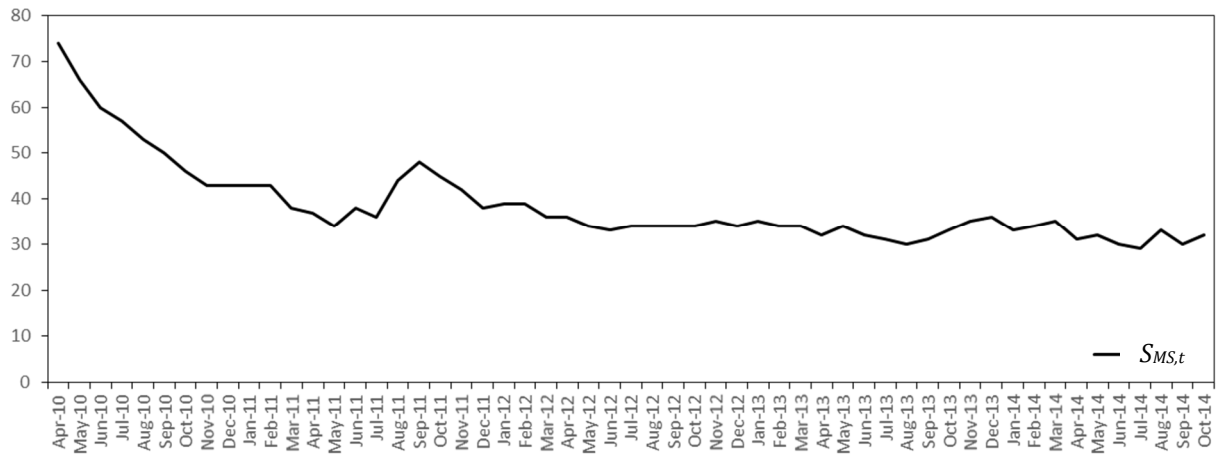c. A randomized rotating panel

Figure 2. Level, spread and activity of our consumer confidence index.
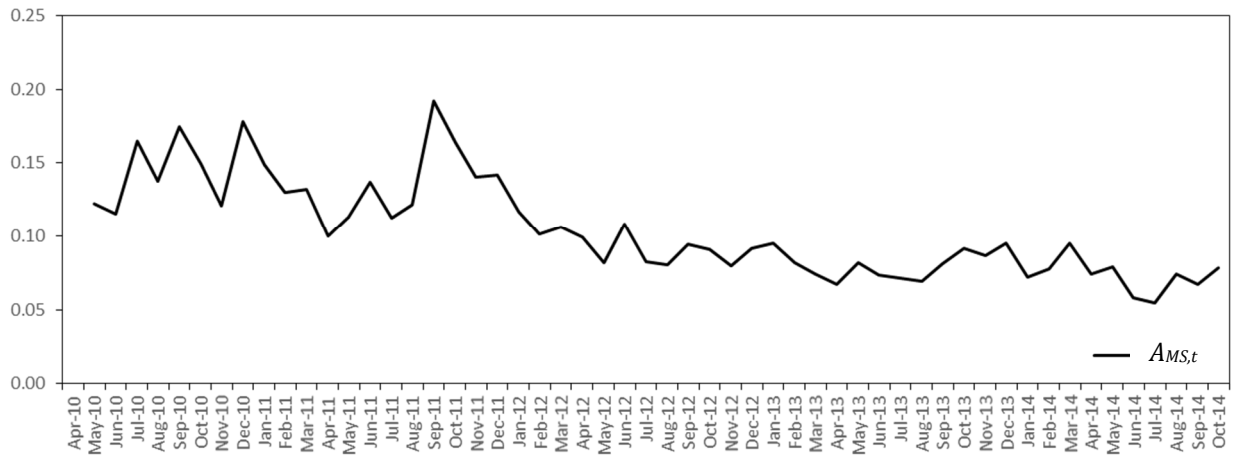
a. Level

b. Spread

c. Activity

Table 1. Estimated transition matrix of the Markov switching model
with its implied long-tern properties

Transition from state	Transition to state (day-to-day)			Transition to state (90 days)			Steady state
	Optimistic	Neutral	Pessimistic	Optimistic	Neutral	Pessimistic	
Optimistic	0.996	0.000	0.004	0.732	0.128	0.141	0.418
Neutral	0.006	0.206	0.788	0.194	0.387	0.420	0.279
Pessimistic	0.000	0.730	0.270	0.191	0.388	0.421	0.304

Table 2. Diagnostics of the estimated forecasting models (12), (13) and (14)

Target variable	R^2 of (12)	R^2 of (13)	R^2 of (14)	φ in (14)	$SE_{\varphi,(14)}$	p-value of (14) vs (13)
Food, drink and tobacco	0.05	0.06	0.17	-0.15	0.17	0.26
Durable goods	0.37	0.25	0.38	-0.07	0.17	0.06
Other goods	0.59	0.56	0.63	0.67	0.10	0.09
Services	0.93	0.93	0.94	0.89	0.08	0.03
Total consumption	0.06	0.06	0.10	0.00	0.17	0.76

Table 3. Relative forecast errors of the forecasting equations (13) and (14), compared to (12)

Target variable[1]	(13): CCI_{SN}, level only	(14): CCI_{SN}, level, spread and activity
Food, drink and tobacco	0.998	0.901
Durable goods	1.127	0.877
Other goods	1.127	0.804
Services	0.878	0.474
Total consumption	0.989	0.957
Average	1.024	0.802

[1] Measured in terms of the total value, where the average value over 2000 is normalized to be 100.

Source: http://www.cbs.nl/en-GB/menu/themas/dossiers/conjunctuur/publicaties/conjunctuurbericht/inhoud/conjunctuurklok/toelichtingen/ck-06.htm