

Testing Changes in Consumer Confidence Indicators

Philip Hans Franses*
Econometric Institute
Erasmus University Rotterdam

Rutger van Oest
Department of Marketing
Tilburg University

April 20, 2006

Econometric Institute Report 2006-18

*We thank Lotje Kruithof and Dick van Dijk for their help with collecting the data. Address for correspondence is Econometric Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR, Rotterdam, The Netherlands, email: franses@few.eur.nl

Testing Changes in Consumer Confidence Indicators

Abstract

The authors propose a statistical methodology to test changes in consumer confidence indicators. These indicators are surveyed monthly and each time concern different individuals. This complicates a straightforward interpretation of changes in the values of the index. The proposed methodology involves estimating the transition matrix which connects the fractions of positive, neutral and negative opinions. The elements of this matrix can be estimated and confidence bounds can be computed. A by-product of the method is a simple tool to correct for seasonality. An illustration to about two decades of Dutch data shows that monthly changes in consumer confidence are not often significantly different from zero.

Keywords

Consumer confidence; Markov process

JEL code

C82, E32

1 Introduction and motivation

Consumer confidence indicators (and also business surveys) are seen as important measures for the perceived current and future state of the economy, and therefore they are treated as leading indicators in many macroeconomic forecasting models. As the indicators are indexed, the variables of common interest are the changes in these indicators. Usually, all values of these changes are included in forecasting models, even in cases where the changes would not have been statistically significant from zero. A key reason for this practice is that a formal statistical methodology for testing such changes is not available. It is the purpose of the present paper to fill in this gap by proposing such a methodology.

The reason why the relevant tests are not available originates from the way the data are collected. For example, each month Statistics Netherlands compiles a consumer confidence index. This index is obtained by interviewing approximately 1000 individuals who are asked to answer five questions¹. The respondents can answer these questions with negative, neutral, or positive. To compile the index, the percentage of respondents with a negative answer is subtracted from the percentage of respondents with a positive answer. The key feature of these data is that the approximately 1000 individuals do not constitute a panel, but that they are so-called repeated cross sections. Of course, Statistics Netherlands seeks to ensure that each time a representative sample of 1000 individuals is drawn, so, overall and on average, the characteristics of these individuals are similar, but still they are not exactly the same individuals. This last aspect makes an evaluation of the indicator over time, that is, changes in the index, less straightforward.

As said, the change in the value of the index is relevant, and it is indeed common practice (particularly in the popular business press) to compare the current month's index to the last month's index to see if consumer confidence has increased or decreased. Although this comparison has great economic relevance, a potentially problematic aspect is that the two compared index values are collected by surveying different respondents. It is thus inappropriate to equate a 5-points decline across index values to a statement like

¹These questions are: economic situation of previous 12 months?, economic situation of next 12 months?, financial situation of previous 12 months?, financial situation of next 12 months?, right time to make big purchases? The overall index is computed as the average of these five components.

“5% of the population has become more negative”. In fact, such a decline really means that the sample of individuals in this month has on average a 5% more negative attitude than the sample of other individuals last month. Moreover, a negative change in the index does not reveal the number of people who have actually changed from positive to negative. This phenomenon of seeking to infer individual behavior from aggregate data is called ecological inference, see King (1997), and it turns out also to be relevant for the consumer confidence measure.

In sum, what one really would want to measure are possible shifts between the three states “negative”, “neutral” and “positive”. Unfortunately, to reveal those dynamics it is necessary to track over time the same respondents rather than different respondents. Such an approach would require a substantially larger effort to collect data than the current methods, and also one may wonder whether each month asking the same individuals would give useful information.

In this paper we propose a statistical methodology that can estimate the elements of the transition matrix in between last month’s three states of negative, neutral and positive and this month’s, while preserving the currently used method of data collection. Subsequently, these transitions can be used to test changes in the consumer confidence index. As it is not possible to collect these individual transitions, they would need to be inferred by means of an econometric model. Although no model is perfect, our approach will be shown to have several advantages. First, changes can be inferred *as if* the same respondents answered the five questions. Second, the underlying model makes it possible to adjust in a simple way for seasonal effects, special events and other factors. Third, as we will demonstrate, confidence bounds can be computed for the elements in the transition matrix, which in turn allow us to make statements like “consumer confidence has increased significantly different from zero”.

The outline of our paper is as follows. In Section 2 we outline the model which should connect the components (negative, neutral and positive) of a consumer confidence index in subsequent months. We also discuss parameter estimation. In Section 3, we propose a methodology to statistically test changes in subsequent values of the index based on this

model. We show how one can compute confidence bounds for the elements of the transition matrix, and we also demonstrate how a simple, hence model-based, method for seasonal adjustment works. In Section 4, we illustrate our methodology for about two decades of Dutch consumer confidence data collected and made available by Statistics Netherlands. One of the results is that changes in the index are rarely statistically significantly different from zero. In Section 5 we conclude with the limitations of our study and with a variety of further research topics.

2 A model for transitions

The key ingredient for our methodology to test changes in consumer confidence is a model that connects the monthly survey observations. This section is based on the model recently put forward by Van Oest and Franses (2005), where these authors have considered it to see which brands gain share from which brands using weekly sales data. That is, retail stores typically observe weekly market shares of brands in a certain category, and they do not observe which customers are switching between brands, while they of course would like to know that. We believe that the model in their study can be used here too. Here we are interested in the transitions from each of the three confidence states in this month to each of these states in the next month. More formally, we want to estimate the elements of a 3×3 transition matrix concerning opinion states. This section first outlines the model and then discusses the method of maximum likelihood to estimate the parameters.

2.1 The model

We denote $S_{1,t}$, $S_{2,t}$ and $S_{3,t}$ as the observed fractions (where we sometimes also use words like percentages and shares) of respondents who are negative, neutral and positive in month t , respectively. Furthermore, we define the fraction of respondents who were in state $l \in \{1, 2, 3\}$ in month $t - 1$ and who move to state $k \in \{1, 2, 3\}$ in month $t \in \{1, \dots, T\}$ by $\lambda_{l,k,t}$. We collect the shares in the vector $S_t = (S_{1,t}, S_{2,t}, S_{3,t})'$, and summarize the transition fractions in the matrix $\Lambda_t = (\lambda_{l,k,t})_{l=1,2,3}^{k=1,2,3}$, where it is the latter matrix that is of focal interest.

To create a model that can be used for our purposes, we, as Van Oest and Franses (2005), need two conditions. The first condition is the redistribution condition, which states that the rows of the transition matrix Λ_t sum to one, that is,

$$\Lambda_t \iota_3 = \iota_3, \quad (1)$$

where ι_3 is a 3×1 vector consisting of ones. The second condition is the well-known Markov condition. The current share of the focal state is the sum of portions carried over to that state, and in our notation this reads as

$$S_t = \Lambda_t' S_{t-1}. \quad (2)$$

Additional to these two conditions, we make a model assumption. We assume that the transition matrix Λ_t can be decomposed into a deterministic component $\tilde{\Lambda}_t$ and a stochastic component E_t with elements $e_{l,k,t}$ independently distributed according to a normal distribution with mean 0 and variance $\sigma^2 S_{l,t-1}^{\gamma_1} S_{k,t-1}^{\gamma_2}$, that is,

$$\Lambda_t = \tilde{\Lambda}_t + E_t. \quad (3)$$

Further, we assume that the elements of $\tilde{\Lambda}_t$ are defined by the logit structure

$$\tilde{\lambda}_{l,k,t} = \frac{\exp(\alpha_{l,k} + x_t' \beta_k)}{\sum_{j=1}^3 \exp(\alpha_{l,j} + x_t' \beta_j)}, \quad (4)$$

where the vector x_t contains potential predictors. Note that each transition has a unique intercept parameter $\alpha_{l,k}$. Furthermore, the explanatory variables contained in x_t have parameters depending on the future state. For example, bad weather, increasing inflation and an increase in unemployment may all “favor” the negative state and may put the positive state at a disadvantage.

The variance parameters γ_1 and γ_2 in the variance term $\sigma^2 S_{l,t-1}^{\gamma_1} S_{k,t-1}^{\gamma_2}$ account for any heteroskedasticity. In fact, it may be expected that $\gamma_1 < 0$ and $\gamma_2 > 0$, which here implies that there is more uncertainty for switching from a small state (a small fraction of respondents) to a large state (a large fraction of respondents), and there is less uncertainty in the opposite direction.

Substituting the decomposition $\Lambda_t = \tilde{\Lambda}_t + E_t$ into (2) gives

$$S_t = \tilde{\Lambda}_t' S_{t-1} + E_t' S_{t-1}. \quad (5)$$

Our model is defined by this system of equations (5) under the condition that the rows of the error matrix E_t sum to zero, that is,

$$E \iota_3 = 0. \quad (6)$$

This latter restriction on the error terms is required to ensure that the redistribution condition of the original transition matrix Λ_t carries over to the deterministic transition matrix $\tilde{\Lambda}_t$, see Van Oest and Franses (2005) for details.

The reduced form is given by

$$\tilde{S}_t = \begin{pmatrix} S_{1,t} \\ \vdots \\ S_{J-1,t} \end{pmatrix} \sim N(\mu_t, \sigma^2 V_t), \quad (7)$$

where

$$\mu_t = \begin{pmatrix} \sum_{l=1}^J \tilde{\lambda}_{l,1,t} S_{l,t-1} \\ \vdots \\ \sum_{l=1}^J \tilde{\lambda}_{l,J-1,t} S_{l,t-1} \end{pmatrix}, \quad (8)$$

and

$$V_t = \sum_{l=1}^J S_{l,t-1}^{2+\gamma_1} \left[-\frac{1}{\sum_{j=1}^J S_{j,t-1}^{\gamma_2}} \begin{pmatrix} S_{1,t-1}^{\gamma_2} \\ \vdots \\ S_{J-1,t-1}^{\gamma_2} \end{pmatrix} \begin{pmatrix} S_{1,t-1}^{\gamma_2} & \cdots & S_{J-1,t-1}^{\gamma_2} \end{pmatrix} + \text{diag}(S_{1,t-1}^{\gamma_2}, \dots, S_{J-1,t-1}^{\gamma_2}) \right], \quad (9)$$

where in our application the number of states J is of course equal to three (negative, neutral and positive).

For our later purposes it is useful to write

$$\tilde{S}_t = \begin{pmatrix} S_{1,t} \\ S_{2,t} \end{pmatrix} \sim N(\mu_t, \sigma^2 V_t), \quad (10)$$

where

$$\mu_t = \begin{pmatrix} \tilde{\lambda}_{1,1,t} S_{1,t-1} + \tilde{\lambda}_{2,1,t} S_{2,t-1} + \tilde{\lambda}_{3,1,t} S_{3,t-1} \\ \tilde{\lambda}_{1,2,t} S_{1,t-1} + \tilde{\lambda}_{2,2,t} S_{2,t-1} + \tilde{\lambda}_{3,2,t} S_{3,t-1} \end{pmatrix}, \quad (11)$$

and

$$V_t = \left(S_{1,t-1}^{2+\gamma_1} + S_{2,t-1}^{2+\gamma_1} + S_{3,t-1}^{2+\gamma_1} \right) \begin{pmatrix} -\frac{S_{1,t-1}^{2\gamma_2}}{S_{1,t-1}^{\gamma_2} + S_{2,t-1}^{\gamma_2} + S_{3,t-1}^{\gamma_2}} + S_{1,t-1}^{\gamma_2} & -\frac{S_{1,t-1}^{\gamma_2} S_{2,t-1}^{\gamma_2}}{S_{1,t-1}^{\gamma_2} + S_{2,t-1}^{\gamma_2} + S_{3,t-1}^{\gamma_2}} \\ -\frac{S_{1,t-1}^{\gamma_2} S_{2,t-1}^{\gamma_2}}{S_{1,t-1}^{\gamma_2} + S_{2,t-1}^{\gamma_2} + S_{3,t-1}^{\gamma_2}} & -\frac{S_{2,t-1}^{2\gamma_2}}{S_{1,t-1}^{\gamma_2} + S_{2,t-1}^{\gamma_2} + S_{3,t-1}^{\gamma_2}} + S_{2,t-1}^{\gamma_2} \end{pmatrix} \quad (12)$$

Note that the restriction on the error matrix implies that the attitude shares are negatively correlated, and that only two equations need to be estimated. Indeed, like the multinomial logit model, one equation is redundant due to the restriction that the shares (fractions) sum to unity. This completes our model.

2.2 Parameter estimation

The parameters of the model can conveniently be estimated using maximum likelihood [ML]. The parameter estimates result from maximization of the log-likelihood function

$$\begin{aligned} \ln L = & -\frac{(T-1)(J-1)}{2} [\ln(2\pi) + \ln(\sigma^2)] + \frac{1}{2} \sum_{t=2}^T \log \det(V_t^{-1}) \\ & - \frac{1}{2\sigma^2} \sum_{t=2}^T (\tilde{S}_t - \mu_t)' V_t^{-1} (\tilde{S}_t - \mu_t), \end{aligned} \quad (13)$$

where μ_t and V_t are defined by (8) and (9), respectively.

We note that the inverse of V_t is given by

$$V_t^{-1} = \frac{1}{\sum_{l=1}^J S_{l,t-1}^{2+\gamma_1}} \left[S_{J,t-1}^{-\gamma_2} \iota_{J-1} \iota_{J-1}' + \text{diag}(S_{1,t-1}^{-\gamma_2}, \dots, S_{J-1,t-1}^{-\gamma_2}) \right], \quad (14)$$

see Van Oest and Franses (2005) for a full derivation. Standard errors are obtained by taking the square roots of the diagonal elements of the estimated covariance matrix, which in turn can be computed as minus the inverse of the Hessian of (13) evaluated in the optimal parameter values.

Numerical techniques, such as the BFGS algorithm or the Newton-Raphson algorithm, have to be used to get the ML parameter estimates. We have programmed this routine in Ox, and it is our experience that estimation takes no more than a few minutes.

3 Testing changes in consumer confidence

The statistical model in the previous section has merits in its own right, but it can also be incorporated in a methodology to statistically test the changes in subsequent values of the

consumer confidence index. In this section we outline various aspects of the methodology, including the computation of confidence bounds and the way one can handle seasonality and extraordinary events that might have an unwanted impact on consumer confidence on the state of the economy.

3.1 Changes in consumer confidence

The parameter estimates obtained using maximum likelihood can be used to make inference on the transition matrix $\tilde{\Lambda}_t$, which is row-conditional in the sense that the elements of a single row sum to one. A related switching matrix can be defined such that the elements of the *entire* matrix sum to one. This would allow for statements about unconditional switching, that is, in terms of percentages of the total population. The elements of this *unconditional* transition matrix can be computed as

$$\bar{\lambda}_{l,k,t} = S_{l,t-1} \tilde{\lambda}_{l,k,t}, \quad (15)$$

see Van Oest and Franses (2005) for technical details. Note that this transition matrix can provide detailed information on how consumer confidence has *changed*.

In practice one is interested in a single index indicating whether consumer confidence has increased or decreased from month $t - 1$ to t . Therefore, we define a new index for a change in consumer confidence, which is given by

$$\Delta CC_t = \sum_{l=1}^3 \sum_{k=1}^3 (k - l) \bar{\lambda}_{l,k,t}. \quad (16)$$

This measure is a weighted sum of the elements of the *unconditional* transition matrix, where the weights correspond to the directions and sizes of the changes in confidence. For example, switching from the neutral state to the negative state corresponds with one step in the negative direction and hence it gets assigned a weight of -1 , while switching from the negative state to the positive state corresponds to two steps in the positive direction and hence gets assigned a weight of $+2$. We wish to emphasize again that this measure indicates *changes* in consumer confidence, rather than confidence levels. For example, a positive value means that individuals have become more positive as compared to one month ago, but of course it does not necessarily imply that the absolute level of confidence is high.

3.2 Computing confidence bounds

Besides having point estimates for the elements of the transition matrix and for the changes in the confidence index, it is also useful to have an impression of how reliable these quantities of interest are, in a statistical sense. Confidence bounds (for example at the 95% level) can provide such insights. Obtaining confidence bounds for the parameter estimates is relatively straightforward by relying on asymptotic principles, but it is easy to see that bounds for transition rates are a bit more complicated.

There are various possible solutions. First, an approximate method is to transform standard errors of the parameter estimates to standard errors of the transition rates via the well-known Delta method, see for example Greene (1993) and Heij et al. (2004). This would result in symmetric confidence bounds, which do not necessarily obey the feasible but bounded range of values, that is, the range of 0 to 1.

A second approach, and which is one we recommend, is a simulation-based approach. This works as follows. We rely on the asymptotic properties of ML estimators. We draw the parameters from their corresponding multivariate distribution, which is asymptotically normal with mean equal to the ML parameter values and a covariance matrix equal to minus the inverse of the Hessian of (13) evaluated at the ML values. By transforming the simulated parameter values to transition rates and confidence indicators, the corresponding distributions can be obtained. Confidence bounds can now be computed such that 2.5% of the draws are located to the left and 2.5% of the draws are located to the right.

3.3 Seasonal adjustment

Even though the questions in the consumer confidence surveys try to make individuals compare the last 12 months with the next 12 months, it is well known that business and consumer survey data show signs of seasonality. It may now be desirable to filter out such seasonality effects to make the quantities of interest more comparable over time, and to detect the “real” underlying pattern.

As a first step to reduce seasonal variation, one can set the seasonal variables in $\tilde{\lambda}_{l,k,t}$

at their own average values. However, as the *unconditional* transition rates (and hence the changes in the confidence indicator) depend on last month's attitude shares, and these shares may in turn depend on the season, this solution is not sufficient. To remove seasonality from the previous market shares of each of the three states, we propose to subtract the month-specific mean and, next, we add the overall mean. Using the de-seasonalized transition matrices and the de-seasonalized opinion shares, one can compute the unconditional transition rates via the relationship $\bar{\lambda}_{l,k,t} = S_{l,t-1} \tilde{\lambda}_{l,k,t}$. Hence, our model allows for an easy to use method to account for unwanted seasonality in changes in the consumer confidence indicator.

4 An illustration

In this section we illustrate the various components of our methodology for a large sample of Dutch data on consumer confidence. Fortunately, Statistics Netherlands reports, via their website and the included program Statline, the fractions of individuals who vote for negative, neutral and positive. So, we have access to the relevant percentages for these three categories. In the section we first outline the data we have, and next we present the results.

4.1 Data

We use monthly data containing the percentages of negative, neutral and positive opinions, taken from Statistics Netherlands. The considered data start in January 1987 and they end in December 2005. The observations in 2004 and 2005 are retained for out-of-sample model validation. Hence, we use $17 \times 12 = 204$ months for parameter estimation and we leave 24 months in a hold-out sample. In the model for the transition matrix, we consider monthly inflation², unemployment³ and temperature⁴, and we include zero-one dummy variables for the various months. Some preliminary testing revealed that the most relevant months are January, July, August and October, so our final estimation results

²source: Datastream

³source: Datastream

⁴source: <http://www.knmi.nl/klimatologie/maandgegevens/datafiles/mndgeg-260.tg.txt>

will be confined by including only the corresponding dummy variables. The inflation and unemployment rates are considered after seasonal adjustment, as they are reported by Statistics Netherlands. Furthermore, these two variables as well as temperature are considered after a first differences transformation, as we deal with changes over time and hence the explanatory variables should be expressed in terms of changes as well.

4.2 Results

We now turn to the estimation results. Before we can interpret the parameter estimates, we first see if the model for the transition matrix makes sense. We therefore start by discussing model performance in terms of in-sample and out-of-sample fit. Next, we provide parameter estimates, and we translate them into meaningful quantities like transition rates and changes in the confidence indicator.

Model performance

Our model provides a good fit to the attitude shares of the three states. Figure 1 provides an illustration, where it again should be noted that the years 2004 and 2005 concern the hold-out sample. The first row of Table 1 contains a few summary statistics, that is, the correlations between actual and predicted shares for the negative, neutral and positive states, and the in-sample and out-of-sample values of the Root Mean Squared Error [RMSE] measure. Clearly, the out-of-sample forecast fit is about equally good as the in-sample model fit.

To provide a further benchmark for these values, we also estimate a so-called attraction model with the same explanatory variables (although inflation, unemployment and temperature now correspond to absolute levels rather than to first differences). An attraction model is a popular model in the marketing literature to describe and forecast market shares. It bears resemblance with the familiar multinomial logit model, see Fok et al. (2002) for a recent summary account of this attraction model. The model can, in the present situation, be seen as a kind of an agnostic model as it does not include any statement about transitions from one state to another. On the other hand, if our model would do worse than this one, one may wonder what the incremental value of our

methodology would be. We also include lagged market shares as explanatory variables in the attraction model to make the two models even more comparable. Comparing the two rows in Table 1 suggests that our model and the attraction model have a similar performance, and that no model dominates the other.

Parameter estimates

It seems that our model for the transition matrix has an adequate fit, both in and out of sample, so now we turn to the interpretation of the estimated parameters. Table 2 contains the parameter estimates and the corresponding standard errors in parentheses. We mention here that we set the parameters $\alpha_{1,3}$ and $\alpha_{3,1}$ at -100 , as it turns out that their estimators walk away to values as $-\infty$ implying a corner solution in which switching from the negative state to the positive state and vice versa never occurs.⁵

As can be seen from Table 2, all response parameters have the expected signs. The results in this table further suggest that recent increases in inflation and recent increases in unemployment have a negative impact on consumer confidence, while more pleasant weather conditions (represented by the change in temperature in the considered month), have a positive effect. Furthermore, and as noted earlier, we find that the months January, July, August and October have a substantial influence on consumer confidence. The first three months provide a positive boost, while confidence seems to decline quite dramatically in October. Explanations could be that January is the beginning of a fresh new year providing inspiration and new opportunities, that July and August are holiday months, while October marks the transition from Summer to Winter. We note that even after accounting for these monthly effects, changes in temperature (which is obviously related to the month) still has explanatory power. The signs of the two heteroskedasticity parameters γ_1 and γ_2 are also as expected, that is, there is more uncertainty around switching from a small state to a large state, and less uncertainty in the opposite direction.

⁵Those parameters receive very large standard errors, as for example both values like -50 and -1000 imply that corresponding switching will always have zero probability. Hence, the transition rates $\tilde{\lambda}_{1,3}$ and $\tilde{\lambda}_{3,1}$ will always be equal to zero without any variation. This justifies fixing both $\alpha_{1,3}$ and $\alpha_{3,1}$ at a large negative value like -100 .

Transitions of opinions

Starting from these parameter estimates and associated covariance matrix, we compute all attitude transition rates and 95% confidence bounds using the simulation-based approach. The results are based on 5.000 simulation runs. Temperature and the four month dummies are the variables we want to adjust for by replacing them by their averages. Table 3 contains estimates of the *row-conditional* transition rates $\tilde{\lambda}_{l,k,t}$ accompanied by confidence bounds. The numbers correspond to averages over time. It seems that about 95% of the respondents does not switch attitudes in two subsequent months. Furthermore, the table suggests that, after correcting for seasonal effects, attitude preservation is even a little bit higher. This makes sense as changing seasonal conditions may stimulate switching.

Table 4 contains the *unconditional* transition rates, again averaged over time. It indicates that we can infer attitude switching quite accurately, as the confidence bounds are quite narrow, and that switching to more positive and more negative states is quite balanced in the long run. As an illustration, Figure 2 displays the simulated distribution of the monthly “overall retention rate”, that is, the distribution of the time average of $\bar{\lambda}_{1,1,t} + \bar{\lambda}_{2,2,t} + \bar{\lambda}_{3,3,t}$ (after seasonal adjustment, in percentages). The distribution is highly skewed and virtually all density mass is located in between 90% and 99%. Hence, even after accounting for parameter uncertainty, we can be confident that on average the percentage of stayers is more than 90%.

To get insight into changes in consumer confidence, one is not only interested in *averages*. It would now be particularly useful to track how the transition rates *develop over time*. Figure 3 contains the graphs of the unconditional transition rates after seasonal adjustment. Confidence bounds are shown as well. We note that the patterns on the diagonal (staying in the previous state) follow the original patterns in Figure 1 quite closely.

Changes in net consumer confidence

Figure 4 displays the change-in-confidence index (16) with corresponding confidence bounds. It also shows the measure currently used by Statistics Netherlands, which is the change in

confidence with the latter being operationalized as the percentage of positive respondents minus the percentage of negative respondents. The two graphs indicate that the Statistics Netherlands measure is rather “bumpy”, while our measure reveals “a kind of business cycle”. Furthermore, by looking at the numbers along the vertical axis one can see that our measure displays much less variation than that of Statistics Netherlands does. For example, the latter suggests that changes in the monthly consumer confidence may be more than five percentage points. This seems to be quite substantial. One explanation is that the index, as is shown here, still includes seasonal effects inducing additional switching. Furthermore, a second explanation is that the index is constructed from the answers in two different months from *different* respondents. Even though the considered number of respondents is reasonably large, there might still be some variation accountable to those different groups. On the other hand, our (seasonally-adjusted) change-in-confidence variable indicates that a change in consumer confidence by more than two percent is already exceptional. Another interesting point is that the associated 95% confidence intervals do not always include zero. For example, at the end of 1991 there has been a significant increase in consumer confidence, while confidence decreased significantly in all months from the end of 1999 to the end of 2000.

Without adjusting for seasonality the correlation between our variable (not shown here without seasonality corrections) and the CBS measure is 0.55, but after correction all correlation has disappeared (0.00). The correlation between our seasonally adjusted and not adjusted changes equals 0.16.

Finally, in Table 5 we report the months in which significant changes of the consumer confidence indicator occurred. A first impression from this table is that there are not that many months with changes that are significantly different from zero. Interestingly, though, is that the significant values seem to precede periods of real decline and growth in the Dutch economy.

5 Discussion and conclusion

In this paper we have put forward a statistical methodology to test changes in the consumer confidence index. It appears easy to apply. When we illustrated it for a large sample of Dutch data, we explored its potential for drawing useful and interesting conclusions.

A serious but inevitable limitation of our approach is that it hinges upon a few conditions and assumptions on econometric models. As always, these assumptions may not be valid. Here the additional drawback is that the accuracy of the model cannot be tested with formal diagnostics, as the true transitions data are not available, so that only out-of-sample fit can be insightful.

Probably the best way to validate the usefulness of our proposed measure for changes in consumer confidence is to see if it establishes better forecasting properties for business cycle variables. Indeed, one may wish only to include in the forecasting model only those changes that are statistically significant, and hence not all changes. We relegate this exercise to further work. Then we will also look at data for other countries and data from business surveys.

References

Fok D., P.H. Franses and R. Paap (2002), “Econometric Analysis of the Market Share Attraction Model” in P. H. Franses and A. L. Montgomery (eds.), *Advances in Econometrics Vol. 16: Econometric Models in Marketing*, Elsevier Science Publishers.

Greene W.H. (1993), *Econometric Analysis*, third edition, Prentice Hall.

Heij, C., P.M. de Boer, P.H. Franses, T. Kloek and H.K. van Dijk (2004), *Econometric Methods with Applications to Business and Economics*, Oxford: Oxford University Press.

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

Van Oest R. and P.H. Franses (2005), “Which Brands Gain Share from Which Brands? Inference from Store-Level Scanner Data”, *Quantitative Marketing and Economics*, 3, 281-304.

Table 1: Performance of our model for the transition matrix, relative to the attraction model, which does not consider transitions. The first three columns measure correct classification.

	Negative	Neutral	Positive	RMSE in-sample	RMSE out of sample
our model	0.980	0.929	0.969	0.018	0.019
attraction	0.980	0.930	0.971	0.018	0.021

Table 2: Parameter estimates and associated standard errors.

$\alpha_{1,1}$	3.967***	(0.747)
$\alpha_{2,1}$	-4.187***	(0.460)
$\alpha_{2,3}$	-4.119***	(0.551)
$\alpha_{3,3}$	-2.811***	(0.410)
inflation (neutral)	-0.301	(0.238)
inflation (positive)	-0.554*	(0.324)
unemployment (neutral)	-0.185***	(0.067)
unemployment (positive)	-0.221**	(0.092)
temperature (neutral)	1.042**	(0.525)
temperature (positive)	1.074*	(0.588)
January (neutral)	2.044***	(0.787)
January (positive)	3.374***	(1.053)
July (neutral)	0.497	(0.426)
July (positive)	1.100*	(0.615)
August (neutral)	1.253*	(0.662)
August (positive)	1.616**	(0.781)
October (neutral)	-0.956**	(0.401)
October (positive)	-1.646***	(0.607)
variance parameter σ	0.025*	(0.013)
variance parameter γ_1	-1.549*	(0.891)
variance parameter γ_2	0.618***	(0.226)

* significant at 10%.

** significant at 5%.

*** significant at 1%.

Table 3: Row-conditional percentages and 95% confidence bounds, seasonally adjusted (s.a.) and not seasonally adjusted (not s.a.).

		point estimate			lower bound			upper bound		
<i>not s.a.</i>		to			to			to		
		neg	neu	pos	neg	neu	pos	neg	neu	pos
	neg	96.16	3.84	0.00	91.85	1.93	0.00	98.07	8.14	0.00
from	neu	1.69	96.08	2.23	0.85	92.88	0.99	3.37	97.91	4.57
	pos	0.00	5.77	94.23	0.00	2.57	88.63	0.00	11.35	97.42
<i>s.a.</i>		to			to			to		
		neg	neu	pos	neg	neu	pos	neg	neu	pos
	neg	97.02	2.98	0.00	92.26	0.77	0.00	99.23	7.74	0.00
from	neu	1.35	96.60	2.05	0.42	92.96	0.71	3.25	98.57	4.76
	pos	0.00	5.45	94.55	0.00	2.19	88.37	0.00	11.62	97.80

Table 4: Unconditional percentages and 95% confidence bounds, seasonally adjusted and not seasonally adjusted.

		point estimate			lower bound			upper bound		
<i>not s.a.</i>		to			to			to		
		neg	neu	pos	neg	neu	pos	neg	neu	pos
	neg	21.77	0.86	0.00	20.83	0.44	0.00	22.19	1.80	0.00
from	neu	0.93	53.40	1.24	0.47	51.62	0.55	1.86	54.41	2.54
	pos	0.00	1.26	20.54	0.00	0.57	19.33	0.00	2.46	21.23
<i>s.a.</i>		to			to			to		
		neg	neu	pos	neg	neu	pos	neg	neu	pos
	neg	21.98	0.65	0.00	20.93	0.17	0.00	22.46	1.70	0.00
from	neu	0.74	53.68	1.14	0.23	51.66	0.39	1.80	54.78	2.65
	pos	0.00	1.18	20.61	0.00	0.48	19.28	0.00	2.52	21.32

Table 5: Months in which confidence has decreased/increased significantly.

year	decreased	increased
1987	2	-
1988	-	1
1989	3 4 6 7 8	-
1990	2 4 9	-
1991	-	5 8 10 11 12
1992	-	4 6
1993	-	4 8 11
1994	9	3 4 5
1995	8 9 11	-
1996	7	4
1997	5 7 8	-
1998	1 2 6 7 8 9	-
1999	1 2 5 7 8 11 12	-
2000	1 2 3 4 5 6 7 8 9 11	-
2001	2 3	10
2002	1	4 10
2003	-	2 4 5 6 8 10

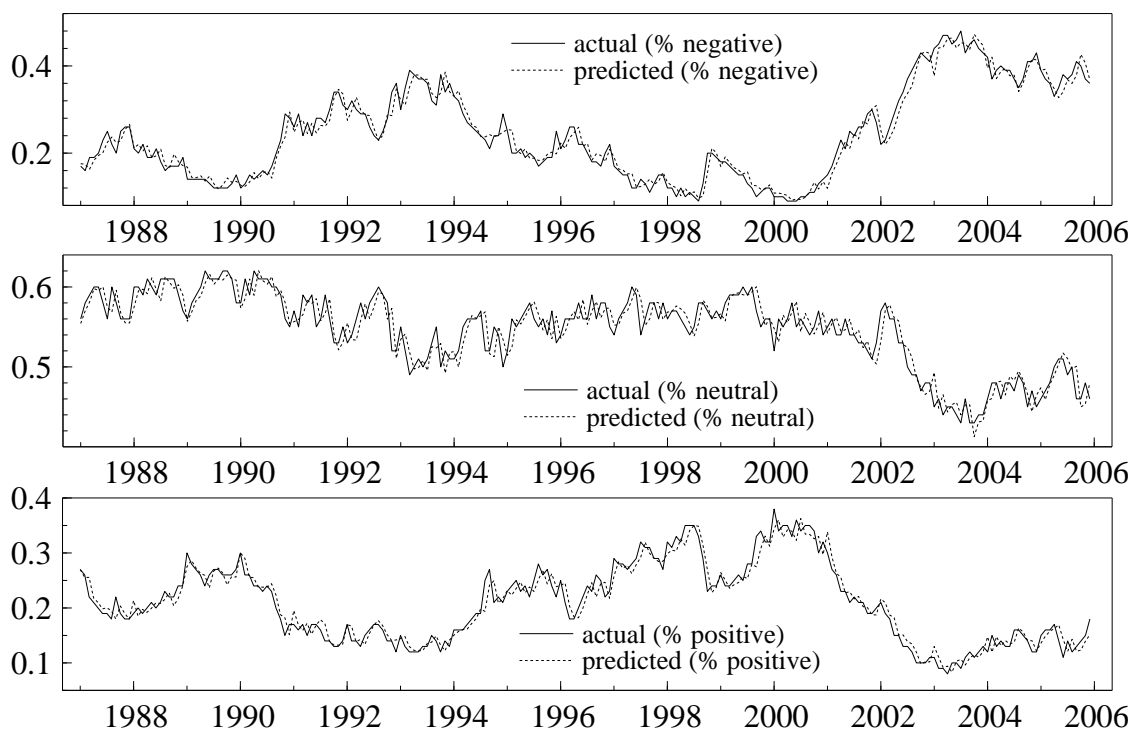


Figure 1: Actual and predicted fractions of negative, neutral and positive opinions.

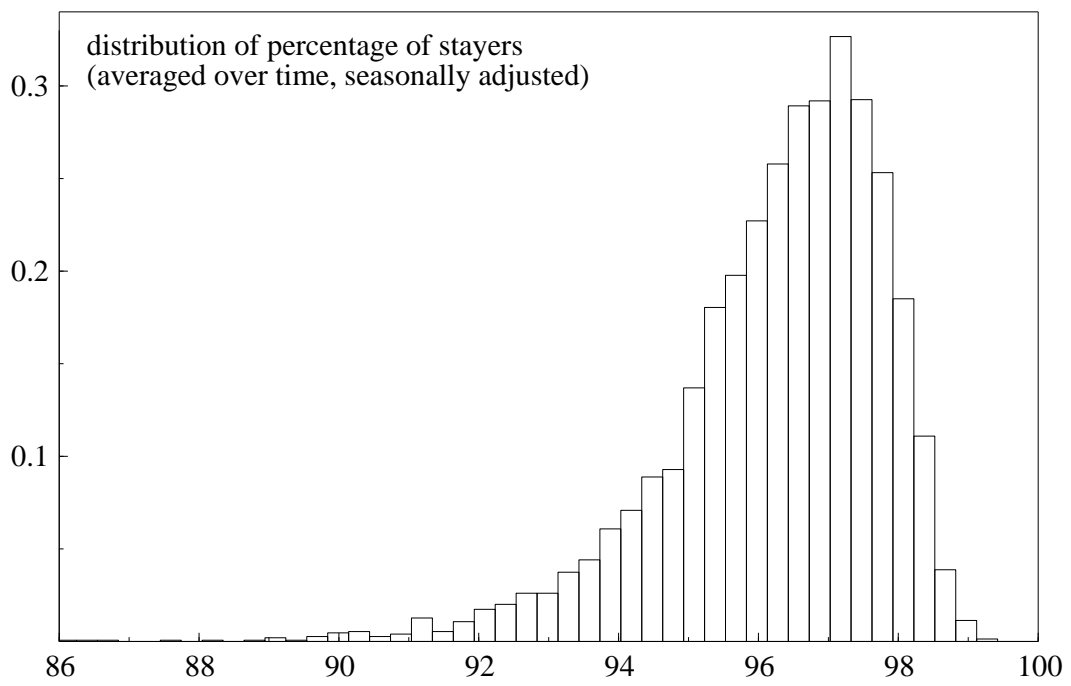


Figure 2: Distribution of the percentage of stayers.

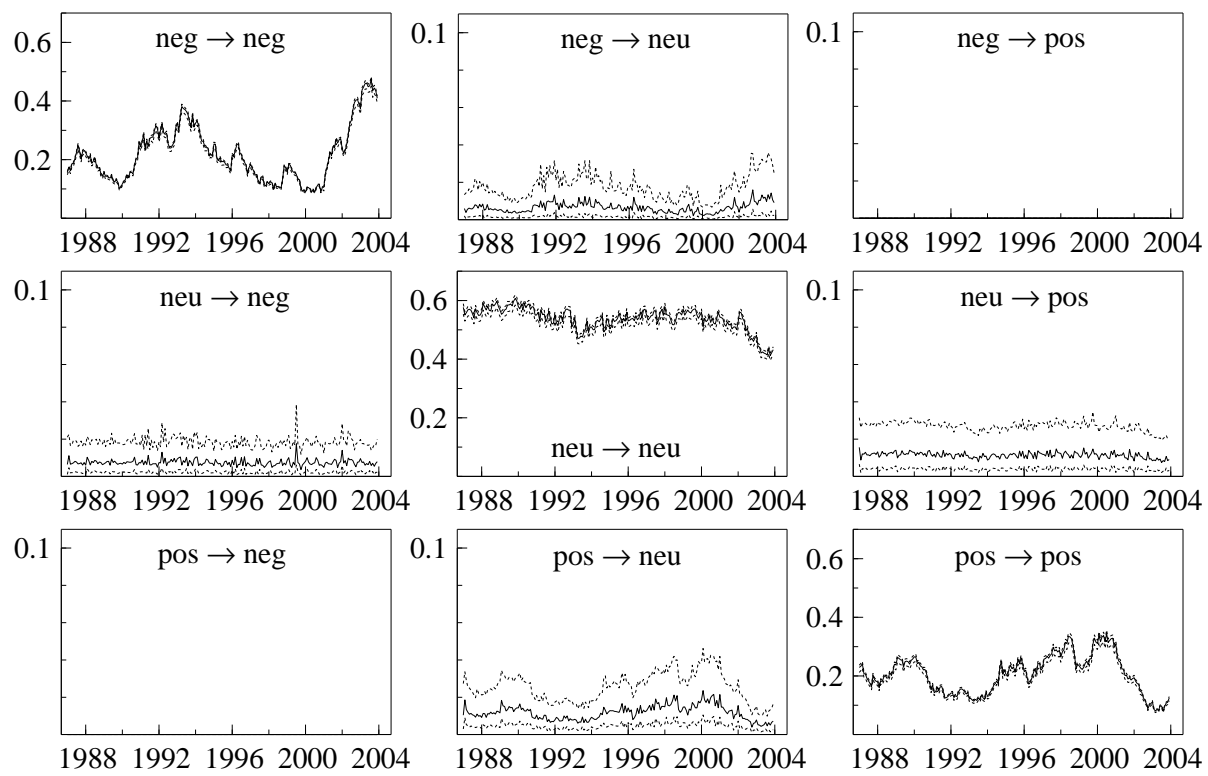


Figure 3: Unconditional attitude transition rates, seasonally adjusted.

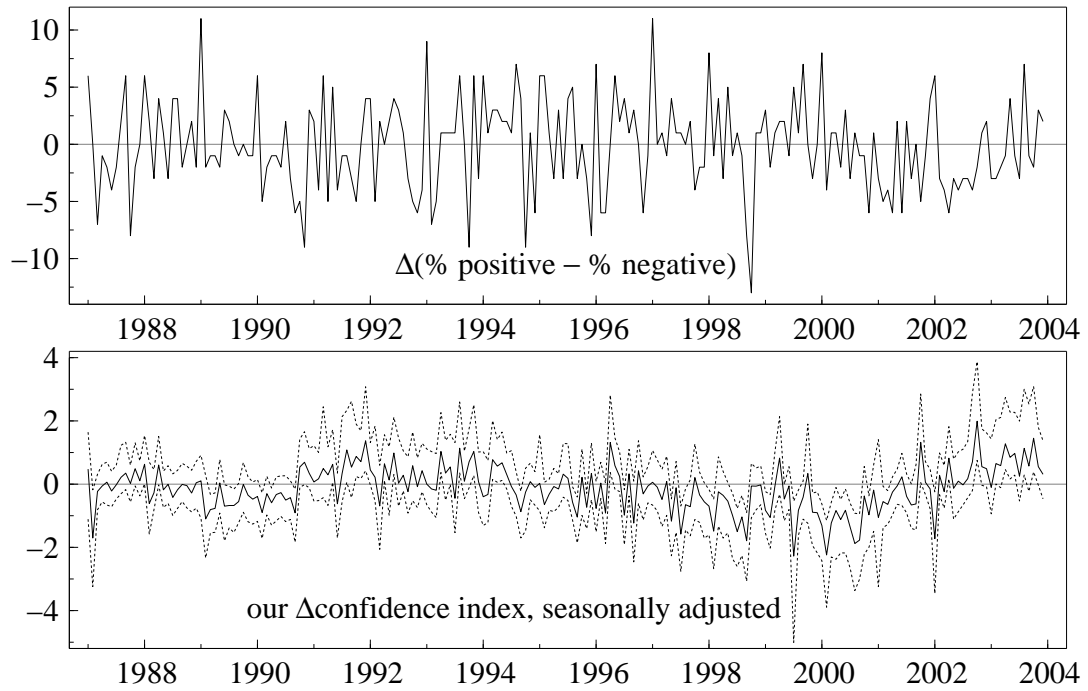


Figure 4: Change-in-confidence measure from Statistics Netherlands (top) and our seasonally-adjusted measure (bottom). The bottom panel also displays the 95% confidence bounds.