

Explaining US Consumer Behavior and Expectations with News Sentiment

Working Paper

Author(s):

Uhl, Matthias W.

Publication date:

2010-07

Permanent link:

<https://doi.org/10.3929/ethz-a-010705418>

Rights / license:

In Copyright - Non-Commercial Use Permitted

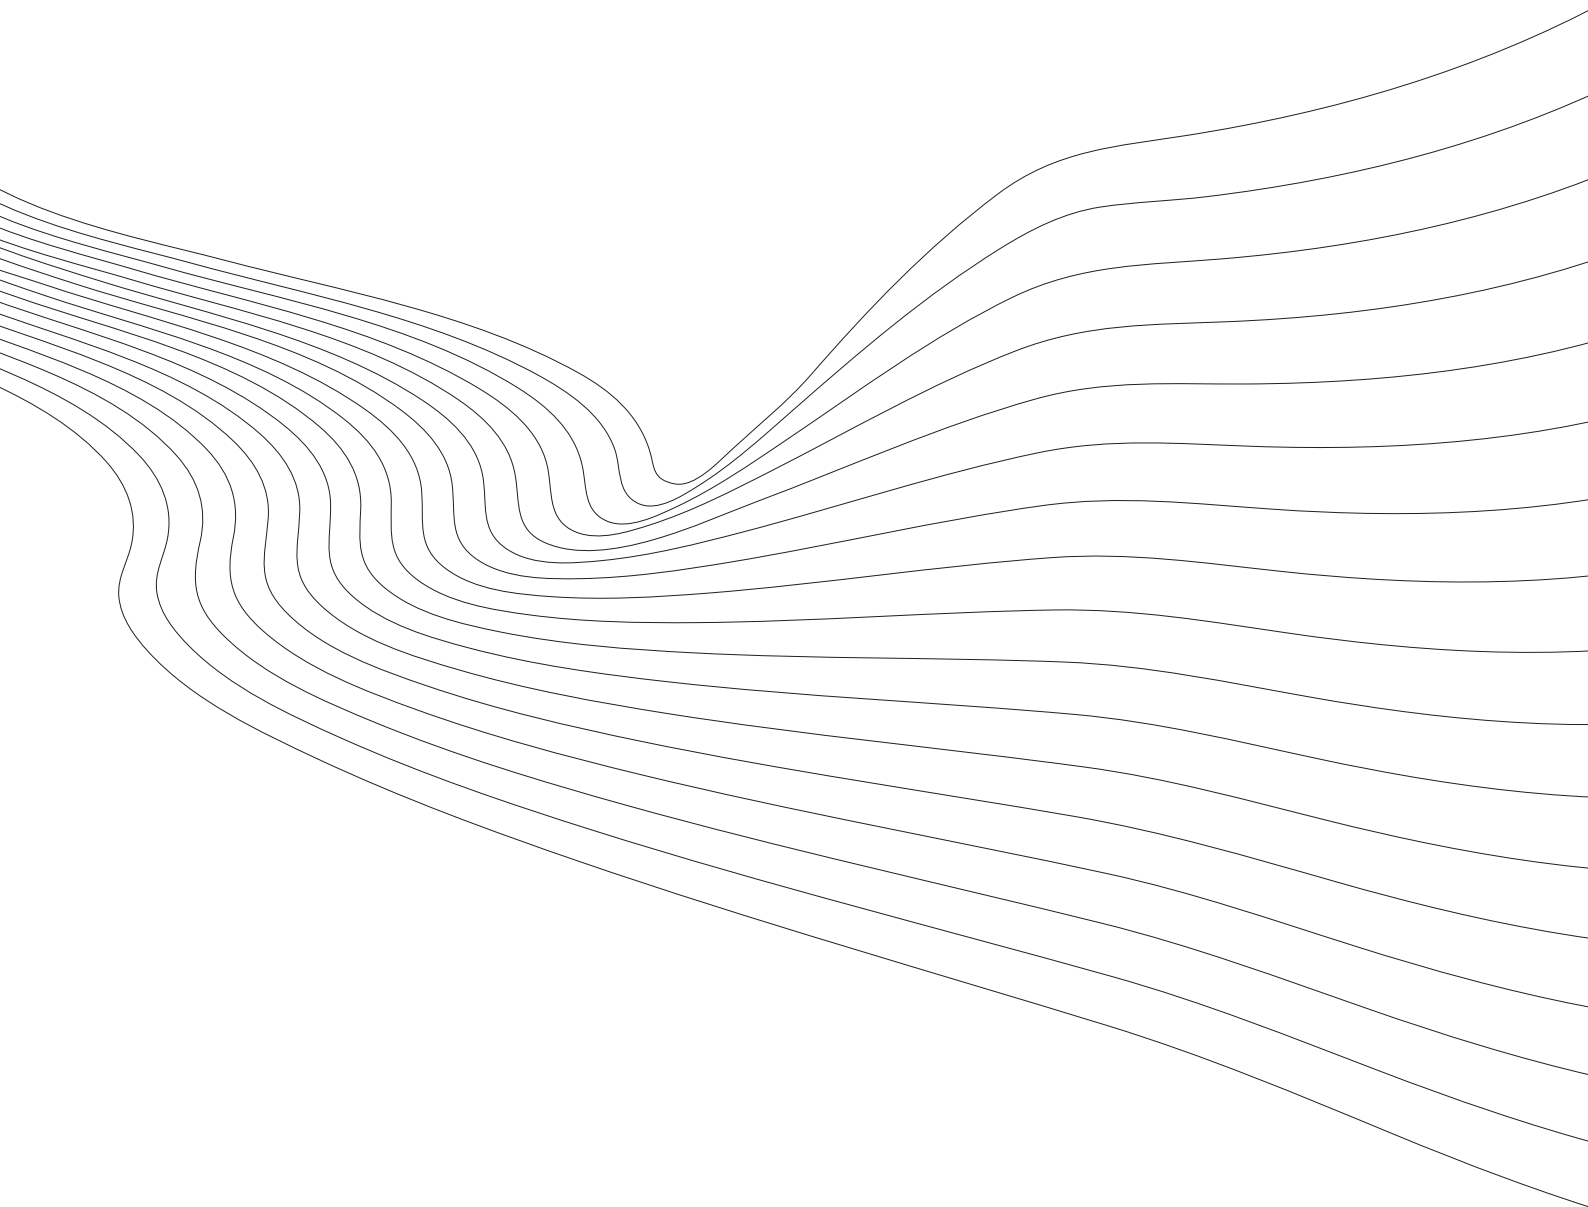
Originally published in:

KOF Working Papers 263

KOF Working Papers

Explaining US Consumer Behavior and Expectations
with News Sentiment

Matthias W. Uhl



KOF

ETH Zurich
KOF Swiss Economic Institute
WEH D 4
Weinbergstrasse 35
8092 Zurich
Switzerland

Phone +41 44 632 42 39
Fax +41 44 632 12 18
www.kof.ethz.ch
kof@kof.ethz.ch

Explaining US Consumer Behavior and Expectations with News Sentiment*

Matthias W. Uhl[†]

This version: September 2010

Abstract

We examine the information content of a newly created news sentiment index from over 300,000 articles from some of the most widely read newspapers in the US to explain changes in the University of Michigan Index of Consumer Sentiment from 1995 to 2009. Using ARMA-models, we show that consumer sentiment is influenced by news sentiment and other variables, such as prices, income and interest rates. While there exists a statistically significant relationship between news sentiment and private consumption, the consumption behavior of private households can best be explained by consumer sentiment combined with changes in personal income and consumer prices. We add news sentiment to the causality chain before consumer sentiment and private consumption.

Keywords: news sentiment, consumer sentiment, private consumption

JEL classifications: D12, E21, E27

*I am grateful for comments and suggestions from Thomas Maag, Michael Lamla, John Leahy, Didier Sornette, and Jan-Egbert Sturm. Thanks go to participants of the Brown Bag Seminar at the KOF Swiss Economic Institute and to participants of the Information, Media and Finance workshop at the IPMZ, University of Zurich for useful input. I further thank Edward Fridael for his valuable help.

[†]*KOF Swiss Economic Institute, ETH Zurich, 8092 Zurich, Switzerland, E-mail: uhl@kof.ethz.ch*

1 Introduction

1.1 Introduction and Summary

We introduce a novel data set with a news sentiment index that was constructed from a selection of over 300,000 newspaper articles from five of the top ten newspapers in the US by circulation. Drawing on the studies of Breeden (1986), Campbell and Mankiw (1989, 1990, 1991) and Carroll et al (1994) among others, we take their idea further by suggesting that news influence consumer sentiment measured by the University of Michigan Index of Consumer Sentiment (ICS), ultimately influencing consumption behavior. By constructing ARMA-models, we show that news sentiment, when combined with other variables, such as personal income, consumer prices and interest rates, achieves statistically significant results to explain changes in consumer sentiment. Consumer sentiment, in turn, has the highest explanatory power in models that explain changes in private consumption. Combined with news sentiment, personal income and prices, the ICS achieves highly statistically significant results and performs best individually.

Section 1 discusses the motivation of the paper and the related literature as well as the novel data set. Section 2 sets out the model and discusses the empirical results, while section 3 concludes.

1.2 Motivation and Literature Overview

Many attempts have been made to explain changes in private consumption of the US economy because it makes up around 70% of its gross domestic product. If economists know how private consumption develops, they have a good understanding of how the overall economy is behaving. According to the Permanent Income Hypothesis (PIH), first formulated by Friedman (1957), consumption patterns of consumers are not determined by current income but rather by their longer-term income expectations. Thus, short-term changes in income have little effect on consumer spending behavior. But what drives the longer-term income expectations of consumers? Hayashi (1982) formulated the basic optimal consumption rule as follows:

$$c_t = \alpha (A_t + H_t), \tag{1}$$

where c_t represents consumption at time t , and A_t is real nonhuman wealth. Real human wealth H_t is defined as the present discounted value of expected future real labor income:

$$H_t = \sum_{k=0}^{\infty} (1 + \mu)^{-k} {}_t y_{t+k}, \quad (2)$$

where μ is the discount rate and ${}_t y_{t+k}$ refers to the household's expectation as of t of real, after-tax labor income at $t+k$. Hayashi (1982) further points out that the rational expectations hypothesis¹ incorporates the idea that ${}_t y_{t+k} = E(y_{t+k} | I_t)$, where I_t is the set of information held by the household at t . It is the information set that each household holds, out of which the consumption behavior is formed. But out of what does this kind of information set consist? By what are consumers' expectations influenced?

In general, let us assume that the ordinary consumer is not a trained economist. This consumer obtains her information about the economy mainly through the news she reads about the economy. This, in turn, shapes her expectations and sentiment about future income and consumption of her household. Each article the consumer reads evokes a certain feeling, opinion, or emotion about the state of the subject, which can be either positive or negative. This is what we call news sentiment in this study. In Fig. 1, we have amended the information flow chart of Doms and Morin (2004) and added news sentiment, postulating that the reader, or the consumer, is influenced by the news and the sentiment portrayed through news that she reads, which form her expectations and sentiment, ultimately driving changes in future consumption.

We thus hypothesize that consumer expectations and sentiment are influenced by news sentiment as well as expectations about future income, personal wealth and general macroeconomic variables. Accordingly, we define I_t as follows

$$I_t = S_{t-i} + \varepsilon_t, \quad (3)$$

where S_{t-i} refers to consumer sentiment and expectations (in this paper, we set $S_{t-i} = \text{ICS}$) at time period t with some lag i .

In their studies, Campbell and Mankiw (1989, 1990, 1991) extend the pure life-cycle / permanent-income hypothesis. As opposed to previous works, they

¹cf. Muth, J.F. 1961. Rational Expectations and the Theory of Price Movements. *Econometrica* 29: 315-335.

distinguish between two kinds of consumers:

$$\Delta c_t^L = \varepsilon_t, \tag{4}$$

$$\Delta c_t^R = \Delta y_t^R, \tag{5}$$

where c_t^L refers to life-cycle consumers, c_t^R to rule-of-thumb consumers, ε_t to news received in period t about lifetime resources, and y_t to current income of private households. A crucial assumption in the Campbell-Mankiw framework is that rule-of-thumb consumers receive a constant proportion λ of total income. Aggregate consumption is then given as follows by the combination of equations (4) and (5):

$$\Delta c_t = \lambda \Delta y_t + \varepsilon_t. \tag{6}$$

Carroll et al (1994) examine the predictive power of consumer sentiment for future changes in consumption spending. They find that lagged consumer sentiment can partly explain current changes in household spending. Drawing on the studies of Campbell and Mankiw (1989, 1990, 1991), Carroll et al (1994) are able to reject their hypothesis that lagged sentiment affects consumption growth only through the income channel, giving room for more variables that might affect consumer behavior. They claim that habit formation should be explored further to identify other channels that could possibly affect consumption growth. News sentiment is tested as a possible new channel that affects consumer expectations and sentiment and ultimately consumption behavior. Carroll et al (1994) also suggest that lagged sentiment might provide incremental information about current consumption growth, as follows,

$$\Delta \log c_t = \alpha_0 + \sum_{i=1}^N \beta_i S_{t-i} + \gamma Z_{t-1} + v_t, \tag{7}$$

where S_t refers to consumer sentiment and expectations (i.e. the ICS) and Z_t is a vector of other variables. They leave room for speculation which other variables are correct to include in the vector Z_t . Their suggestion is key motivator to this study, since they hypothesize that *other variables* can affect consumer behavior. Breeden (1986), for example, found that risk-less interest rates as well as inflation are related to the expected growth rate of aggregate consumption. Thus, we test whether these *other variables* can be personal income y_t , as in

Campbell and Mankiw (1989, 1990, 1991), consumer prices p_t and long-term interest rates r_t , as in Breeden (1986), as well as a novel variable introduced in this study: news sentiment ns_t with some lag i . Thus, in this study Z_t can assume the following values

$$Z_t = y_{t-i} + r_{t-i} + p_{t-i} + ns_{t-i} + u_t. \quad (8)$$

Given equations (7) and (8), we obtain

$$\Delta \log c_t = \alpha_0 + \sum_{i=1}^N \beta_i S_{t-i} + \gamma ns_{t-i} + \delta y_{t-i} + \vartheta r_{t-i} + \eta p_{t-i} + v_t. \quad (9)$$

We thus have a function that incorporates both consumer sentiment and expectations S_t as well as news sentiment ns_t along with other macroeconomic variables that can explain changes in US private consumption.

In another study, Acemoglu and Scott (1994) use UK data to prove that confidence indicators outperform other macroeconomic variables that explain consumer behavior. In this light, they reject the Rational Expectations Permanent Income Hypothesis (REPIH)² and conclude that the predictive ability of confidence indicators is inconsistent with forward-looking behavior. Lloyd (1999) finds that consumer sentiment surveys (including the ICS) perform better than professional forecasters when implemented in forecasting models for inflation and consumer expectations.

Ang et al (2007) find that consumer sentiment surveys (e.g. the ICS) perform best in forecasting models of inflation as opposed to time-series, Phillips curve, and term structure forecasts. They further find little evidence that combining forecasts produces superior forecasts to survey information alone. We thus test each variable individually and jointly against the dependent variable to validate their findings. It is noteworthy to point out that they hypothesize that one possibility for the better performance of survey forecasts is that the survey aggregate information is from many different sources, which are not captured by a single model. They claim that the superior information in median survey forecasts may be due to an effect that is similar to Bayesian model averaging. We thus take this idea further and test whether news sentiment can add value to forecasts based on consumer sentiment surveys, such as the ICS, and if a

²cf. Hall, R. E. 1978. Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of Political Economy* 86: 971-987.

combination of the analysis of widely-read newspapers aids in capturing the effect that news sentiment has on consumers.

Doms and Morin (2004) examine the hypothesis that news media affects consumers' perceptions of the economy. They find that the tone and volume of economic reporting in news affect sentiment of consumers. Further, they identify a short-lived effect of sentiment on consumer spending, lasting only a few months. Given their findings, we want to test whether a positive or negative tone in news reporting (i.e. news sentiment) drives consumer expectations and sentiment, as well as consumption behavior and, if present, how long this effect lasts, so that

$$\Delta S_t = \gamma ns_{t-i} + \varepsilon_t. \quad (10)$$

Given the above, we also want to test whether other variables, such as personal income, risk-free interest rates, and inflation influence and form consumer expectations and sentiment:

$$\Delta S_t = \gamma ns_{t-i} + \delta y_{t-i} + \vartheta r_{t-i} + \eta p_{t-i} + \varepsilon_{t-i}. \quad (11)$$

We thus have two base models set out in this study. First, we test whether news sentiment can explain changes in consumer expectations and sentiment, as well as other variables, such as changes in personal income, interest rates and inflation. Second, we test how consumer expectations can explain changes in private consumption. We further test whether news sentiment can explain changes in private consumption as well, postulating that there might be a direct link between news sentiment and consumption behavior, ignoring consumer expectations and sentiment. And, whether other variables, such as personal income, interest rates and inflation, influence and drive consumption behavior of private households.

1.3 Data

In this study, we introduce a new variable that quantifies news sentiment from the economics section of various newspapers in the US from 1995 to 2009 on a quarterly basis. A sentiment algorithm is used for the analysis of over 300,000 newspaper articles from the *Washington Post* (WP), *USA Today* (UT), the *Houston Chronicle* (HC), the *New York Times* (NYT), and the *Wall Street*

Journal (WSJ). Table 1 shows the average daily circulation of each newspaper and how many newspaper articles were examined for sentiment of each newspaper.³ A news sentiment index was then created from the two newspapers that performed best in the models and that were the most comprehensive graphically: the WP and UT.

The sentiment algorithm distinguishes between positive and negative sentiment of newspaper articles in binary format, namely $\{-1\}$ for negative sentiment and $\{1\}$ for positive sentiment. The sentiment algorithm is based on a broad and complex database of positive and negative words and phrases.⁴ Visual Basic programs were written in order to ease the process of dealing with mass data.⁵ The sentiment algorithm scans each article (headline plus full body) and gives an output file with the respective sentiment rating of each individual article. The article ratings were then aggregated on a quarterly basis.

Quarterly US private consumption data as well as consumer price index data were obtained from the *U.S. Department of Commerce Bureau of Economic Analysis* database.⁶ Long-term interest rates (10-year US-Treasury yields) were obtained from the Federal Reserve Bank.⁷ The University of Michigan Consumer Index data (monthly and quarterly) were downloaded from the University of Michigan and Thomson Reuters public access website.⁸ The ICS is constructed from answers to five questions relating to current economic conditions of consumers as well as consumer expectations.⁹

³The examined newspapers were selected from the top ten list of daily average circulation according to availability in the *LexisNexis* database.

⁴See Appendix A.1 for more information on the sentiment classifier.

⁵See Appendix A.2 for more information on these programs.

⁶See <http://www.bea.gov/>, last accessed 15 April 2010.

⁷See Federal Reserve Bank, http://federalreserve.gov/releases/h15/data/Business_day/H15_TCMNOM_Y10.txt, last accessed 5 June 2010.

⁸See <http://www.sca.isr.umich.edu/>, last accessed 8 June 2010.

⁹A detailed description of the calculation of the index and the individual questions can be found on the homepage of the surveys of consumer from the University of Michigan and Thomson Reuters. See *Index Calculations*, <http://www.sca.isr.umich.edu/documents.php?c=i>, last accessed 8 June 2010.

2 Empirical Analysis

2.1 Modelling

We select the models according to the Box and Jenkins (1979) model selection approach. We thus perform a graphical analysis first. In Fig. 2, we plot the news sentiment index versus the University of Michigan Index of Consumer Sentiment. The news sentiment index ranges between 0.5 and 0.7, indicating that news are - in the period examined - too positive, since the scale is from $\{-1\}$ to $\{1\}$. The identification of a positive bias in news sentiment is consistent with the phenomena that Baron (2006) identified in news media reportings, although this is contrary to the general belief that “bad news sell.” We see that the news sentiment index co-moves nicely with the ICS. Looking at the recent financial crisis of 2008/09, a lagged co-movement becomes apparent. This is in line with our theoretical assumptions from the previous section that news sentiment should influence consumer expectations and sentiment slowly over time. Figures 3 and 4 show changes in personal income and inflation versus the ICS. Again, the recent crisis of the past years becomes evident, in line with the ICS. Figure 5 plots long-term interest rates and the ICS. As we can see, a co-movement is also present. Figure 6 shows the ICS plotted against US private consumption. A co-movement of the ICS and private consumption is somewhat present, especially a co-movement during the recent crisis in 2008/09 stands out. In Fig. 7, the news sentiment index shows some co-movement with changes in private consumption. It shows the economic height in 2000 and the crisis that followed the years after, and it anticipated the recent financial crisis. In Figs. 8 and 9, we plot inflation and personal income versus private consumption. It becomes evident, that these variables move in line with private consumption, especially during the recent crisis. Long-term interest rates also show a similar pattern as changes in consumption behavior in Fig. 10. Thus, according to a first graphical analysis, all variables seem plausible to test in our models, as theoretically laid out in the previous section.

The graphical interpretation shows that the variables taken into consideration might be suited to explain consumer expectations and sentiment as well as changes in private consumption. Nevertheless, we need to test whether these variables hold statistically what they suggest graphically.

We first test each variable for unit roots with Augmented Dickey-Fuller

(ADF) tests according to Dickey and Fuller (1979).¹⁰ We find that private consumption and the ICS as well as interest rates, personal income and inflation have a unit root on the level, whereas the news sentiment index does not have a unit root. To exclude the possibility of spurious regression results and stationarity as Granger and Newbold (1974) noted, we use log differenced values for the dependent variable private consumption and differenced values for the ICS. For the independent variables, such as news sentiment and interest rates we use level data, whereas for personal income and inflation we use differenced values.

We construct a base model that is based on simple autoregressive and moving average models. As in Ang et al (2007), we use the Schwarz criterion (BIC) to determine the order of the autoregression (AR) and moving average (MA) processes. According to these criteria, the two base models have an ARMA (4,1) structure:

$$\Delta S_t = k + \alpha_1 S_{t-1} + \alpha_2 S_{t-2} + \alpha_4 S_{t-4} + \theta_1 \varepsilon_{t-1} + \varepsilon_t, \quad (12)$$

$$\Delta \log c_t = k + \alpha_1 c_{t-1} + \alpha_2 c_{t-2} + \alpha_4 c_{t-4} + \theta_1 \varepsilon_{t-1} + \varepsilon_t, \quad (13)$$

where S_t refers to the ICS, c_t refers to US private consumption, k is the constant term, and ε_t represents the error term. As in Carroll et al (1994), we model the error term, ε_t , with an MA-process because of time aggregation, as consumption expectations and decisions are made continuously, whereas our data set consists of quarterly data.

We then extend our base models by adding the other variables Z_t from equation (8) to equations (12) and (13).

$$\Delta S_t = k + \alpha_1 S_{t-1} + \alpha_2 S_{t-2} + \alpha_4 S_{t-4} + \gamma_i Z_{t-i} + \theta_1 \varepsilon_{t-1} + \varepsilon_t, \quad (14)$$

$$\Delta \log c_t = k + \alpha_1 c_{t-1} + \alpha_2 c_{t-2} + \alpha_4 c_{t-4} + \gamma_i Z_{t-i} + \theta_1 \varepsilon_{t-1} + \varepsilon_t, \quad (15)$$

where i equals the optimal lag length of the independent variable.¹¹

Last, we test whether the ICS and the news sentiment index perform better jointly than individually when attempting to explain changes in private con-

¹⁰See Appendix A.3 for the exact formulation of the ADF test.

¹¹Table 2 shows the optimal lag length for these variables.

sumption. This is in line with what Carroll et al (1994) examined and what we have defined in equation (7) earlier. The equivalent model has an ARMA(4,1) structure as well and is as follows:

$$\Delta \log c_t = k + \alpha_1 c_{t-1} + \alpha_4 c_{t-4} + \gamma_i Z_{t-i} + \rho_i S_{t-i} + \theta_1 \varepsilon_{t-1} + \varepsilon_t. \quad (16)$$

We test the ICS and private consumption against each variable individually and jointly, and apply the Theil Inequality Coefficient for comparison, as according to Theil (1958), as well as the Root Mean Squared Error (RMSE).¹²

2.2 Empirical Results

2.2.1 Explaining the University of Michigan Consumer Sentiment Index

In all models, we derive heteroskedasticity consistent covariance matrices according to White (1980).¹³ Since we utilize ARMA-models, we test for serial correlation with the Breusch-Godfrey Serial Correlation Lagrange Multiplier tests, according to Godfrey (1978) and Breusch and Pagan (1979). We find no serial correlation in any of the models.

Table 3 shows the empirical results of all 7 regressions as in equations (12) and (14) that explain the ICS. The base regression (1) is highly statistically significant and has an adjusted R-squared value of 0.14. In regressions (2) to (5), we test each variable individually against the ICS. In none of these regressions, except for regression (5), is the coefficient statistically significant. The adjusted R-squared values range between 0.12 and 0.21. The signs of the coefficients are as expected: a rising consumer sentiment comes in line with more positive news sentiment. The higher the income of a household, the better their sentiment. In regression (4), we have a negative coefficient sign, which means that if prices fall, consumer are happier, which seems plausible. In regression (5), the case is not so clear cut, although this coefficient is the sole one being statistically significant. The higher long-term interest rates, the higher is consumer sentiment. This can be explained with the assumption that consumers are more content when they receive higher interest rates for their savings. Whether and how this affects

¹²See Appendix A.4 for the calculation of the Theil Inequality Coefficient and the RMSE.

¹³See Appendix A.5 for the exact formulation of the White covariance matrix.

their consumption behavior is examined later. In regression (6) of table 3, we combine all three macroeconomic variables, namely personal income, the consumer price index, and interest rates. The coefficient of inflation is significant at the 10%-level, and the coefficient for interest rates is significant at the 1%-level. The adjusted R-squared value is slightly higher than that in the previous regressions (0.23). Again, higher consumer sentiment can be explained with higher personal income, falling prices and higher interest rates. In regression (7), we add the news sentiment index variable to the variables from regression (6). Interestingly, almost all variables are at least statistically significant at the 5%-level, except personal income. The adjusted R-squared value jumps to 0.28, while the RMSE decreases significantly compared to the previous models. This result suggests that, even though each variable is not statistically significant when tested individually against the ICS, the news sentiment index might be one missing piece of the puzzle that brings the information of all variables together, making them statistically significant. We can thus conclude that changes in consumer expectations and sentiment can best be explained by news sentiment, changes in consumer prices as well as by changes in interest rates.

2.2.2 Explaining Private Consumption

In the next step, we want to explain changes in private consumption by consumer expectations and sentiment (the ICS). Further, we want to test whether any other macroeconomic variables show explanatory power, as described in previous studies and laid out earlier. Since the main focus of this study is on the relationship of news sentiment and consumer expectations and sentiment, we want to go one step further and also test whether there is a direct influence between news sentiment and private consumption. Table 4 shows the regression results. According to equation (13), regression (8) shows the estimation results of the base model, with a statistically significant coefficient and an adjusted R-squared value of 0.45. In regressions (9) to (13), we test each explanatory variable individually against changes in private consumption. The variables news sentiment, consumer prices, and the ICS are statistically significant with adjusted R-squares between 0.44 and 0.58. According to the regression results of (10), personal income is not statistically significant, and neither are interest rates. Results from regression (13) stand out with the ICS as independent variable, as the regression coefficient is highly statistically significant at

the 1%-level with an adjusted R-squared value of 0.58, and much lower RMSE and Theil Inequality Coefficient values. With regards to the coefficient signs, we have a similar picture for the variables than in the previous model when we explained the ICS. Higher private consumption can be explained by higher news sentiment, higher personal income, lower consumer prices, higher interest rates, and higher consumer sentiment. Thus, we can say that consumer expectations and sentiment can explain changes in private consumption best when compared to news sentiment, personal income, consumer prices and interest rates individually. In regression (14), we combine the three macroeconomic variables, personal income, consumer prices and interest rates. Personal income and consumer prices are statistically significant, whereas the coefficient of the variable interest rates is not, which comes in with a negative coefficient sign. The negative sign is contrary to our previous finding from regression (12). We first add news sentiment and then the ICS to regression (14) in (15) and (16), respectively. Comparing the R-squared values of the two regression with the sentiment variables, it becomes obvious that the regression with the ICS (16) has a much higher explanatory power (0.63) than regression (15) with news sentiment (0.53). Again, all explanatory variables are statistically significant in the two regressions except interest rates. Also, the coefficient sign changes in the two regressions, suggesting that long-term interest rates might not be suited to explain changes in private consumption. Last but not least, we test what we set out in equation (16) previously. In regression (17), we thus exclude long-term interest rates and include both sentiment variables. We achieve the highest R-squared value of 0.65 with the lowest RMSE and Theil Inequality coefficients. All coefficient are statistically significant with only personal income not being statistically significant.

Comparing the empirical results in the models with private consumption, we can conclude that the ICS is a much better explanatory variable than news sentiment when combined with macroeconomic variables, such as personal income and consumer prices. Higher private consumption can thus be explained by rising personal income of households, lower consumer prices and higher consumer sentiment and expectations as well as higher news sentiment.

3 Conclusion

Although media coverage has become extremely important in the past decades, the effects of sentiment published in newspaper articles on private consumers have been barely explored in the literature. We introduce a novel data set and procedure by creating a news sentiment index with positive and negative sentiment from over 300,000 newspaper articles of the economics section of some of the most widely-read newspapers in the US from 1995 to 2009, using proprietary tools and a new text mining approach.

We examine empirically the connection and impact of news sentiment on consumer expectations and sentiment, postulating that consumers should form their expectations and sentiment partly based on the news they read, which in turn should ultimately affect their consumption behavior. In accordance with previous research of Breeden (1986), Campbell and Mankiw (1989, 1990, 1991), and Carroll et al (1994), we test other macroeconomic variables such as personal income, inflation, and interest rates, and how they perform in explaining consumer sentiment as well as behavior.

We find that a statistically significant relationship exists between a combination of news sentiment, changes in personal income and consumer prices, and consumer sentiment and expectations. In our ARMA(4,1)-model, a high level of news sentiment, higher personal income, lower consumer prices and higher interest rates explain higher consumer sentiment. Personal income, it seems, does not have the same influence as news sentiment, consumer prices and interest rates on consumer sentiment and expectations, when considering the statistical significance of the coefficients. Consumption behavior can best be explained with a model that comprises news and consumer sentiment, personal income and consumer prices. Higher news and consumer sentiment, higher personal income and lower consumer prices increases consumption behavior in the US. Considered individually, consumer sentiment is the most accurate variable to explain private consumption. Interest rates do not seem to have a significant impact on consumption behavior. When combined with personal income, inflation and interest rates, it becomes evident that the University of Michigan Index of Consumer Sentiment is much better suited than the news sentiment index when explaining changes in private consumption. This finding is in line with our hypothesis that consumer expectations are influenced by news, but form only slowly over time, ultimately influencing the consumption behavior of a household. The causality chain laid out in Fig. 1 is thus confirmed.

We conclude that this first long-term analysis of news sentiment leaves room for future research in order to specify and improve the methods for explaining consumer sentiment and in turn private consumption that are based on news sentiment as well as other variables that affect the ordinary consumer.

References

- [1] Acemoglu, D. and A. Scott (1994) Consumer Confidence and Rational Expectations: Are Agents' Beliefs Consistent with the Theory? *The Economic Journal*, 104, 1-19.
- [2] Ang, A., G. Bekaert, and M. Wei (2007) Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54, 1163-1212.
- [3] Armstrong, D.M. (1961) *Perception and the Physical World*, Humanities Press, London.
- [4] Baron, D. (2006) Persistent media bias, *Journal of Public Economics*, 90, 1-36.
- [5] Bauer, E., and R. Kohavi (1999) An empirical comparison of voting classification algorithms: bagging, boosting, and variants, machine learning, *Machine Learning*, 36, 105-139.
- [6] Breeden, D.T. (1986) Consumption, Production, Inflation and Interest Rates, *Journal of Financial Economics*, 16, 3-39.
- [7] Breusch, T. S., and A. R. Pagan (1979) A Simple Test for Heteroskedasticity and Random Coefficient Variation, *Econometrica*, 48, 1287-1294.
- [8] Campbell, J. and G. Mankiw (1989) Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence, in O. Blanchard and S. Fischer (eds.) *NBER macroeconomics annual 1989*, MIT Press, Cambridge, MA.
- [9] Carroll, C. (2003) Macroeconomic Expectations of Households and Professional Forecasters, *Quarterly Journal of Economics*, 118, 269-298.
- [10] Carroll, C., J. Fuhrer, and D. Wilcox (1994) Does Consumer Sentiment Forecast Household Spending? If So, Why? *American Economic Review*, 84, 1397-1408.
- [11] Dickey, D.A. and W.A. Fuller (1979) Distribution of the Estimators for Autoregressive Time Series with a Unit Root, *Journal of the American Statistical Association*, 74, 427-431.

- [12] Doms, M. and N. Morin (2004) Consumer Sentiment, the Economy, and the News Media, *FRBSF Working Paper*, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=602763.
- [13] Friedman, M. (1957) *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.
- [14] Friedman, N., D. Geiger, and M. Goldszmidt (1997) Bayesian network classifiers, *Machine Learning*, 29, 131-163.
- [15] Godfrey, L. G. (1978) Testing for Multiplicative Heteroskedasticity, *Journal of Econometrics*, 8, 227– 236.
- [16] Granger, C.W.J. and P. Newbold (1974) Spurious Regressions in Econometrics, *Journal of Econometrics*, 2, 111-120.
- [17] Hand, D.J., K. Yu (2001) Idiot’s Bayes - not so stupid after all? *International Statistical Review*, 69, 385-399.
- [18] Hayashi, F. (1982) The Permanent Income Hypothesis: Estimation and Testing by Instrumental Variables, *The Journal of Political Economy*, 90, 895-916.
- [19] Kalman, R.E. (1960) A New Approach to Linear Filtering and Prediction Problems, *Journal of Basic Engineering*, 82, 35-45.
- [20] Kotsiantis, S. B., and P. E. Pintelas (2004) Increasing the Classification Accuracy of Simple Bayesian Classifier, *Lecture Notes in Artificial Intelligence*, 3192, 198-207.
- [21] Lewis, D. (1998) Naive (Bayes) at forty: The independence assumption in information retrieval, in *ECML’98: Tenth European Conference On Machine Learning*.
- [22] Lloyd, B. T. Jr. (1999) Survey Measures of Expected U.S. Inflation, *The Journal of Economic Perspectives*, 13, 125-144.
- [23] Phillips, P.C.B. and P. Perron (1988) Testing for a Unit Root in Time Series Regression, *Biometrika*, 75, 335–346.

- [24] Rish, I. (2001) An empirical study of the naive Bayes classifier, *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence*, available at <http://www.cc.gatech.edu/home/isbell/classes/reading/papers/Rish.pdf>.
- [25] Taylor, S., and R. Bogdan (1998) *Introduction to qualitative research methods: a guidebook and resource*, Wiley Interscience, New York, NY.
- [26] Theil, H. (1958) *Economic forecasts and policy*, North Holland, Amsterdam.
- [27] White, H. (1980) A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, *Econometrica*, 48, 817–838.

Appendix

A.1

The sentiment algorithm is based on one of the most popular classifiers used in machine learning science: the Naive Bayes classifier. The sentiment algorithm java program was obtained from a free web-based sentiment algorithm provider and is tested for accuracy in this study.¹⁴ According to Friedman et al (1997), this classifier learns the conditional probability of each attribute a_i , given its class label c , from training data. The sentiment algorithm was trained to distinguish between *positive* and *negative* sentiment from a pre-defined database of positive and negative words and phrases. The classification is then done by applying Bayes rule to compute the probability of c given the particular instance of a_1, \dots, a_i , and then predicting the class with the highest posterior probability. This means that the computation is rendered feasible by making a strong independence assumption: all the attributes a_i are conditionally independent given the value of the class c . By independence, Friedman et al (1997) further note, probabilistic independence is meant, that is, a is independent of b given c whenever $Pr(a | b, c) = Pr(a | c)$ for all possible values of a, b , and c , whenever $Pr(c) > 0$. This means that in each article, every word and a combination of phrases is checked against the sentiment algorithm and classified as either positive or negative. The sentiment score is then obtained by applying Bayes' rule to the classifications that were obtained for each article individually, so that the output of either positive or negative is generated for each single article.

Lewis (1998) discusses the Naive Bayes approach in historical context by concluding that the algorithm is experiencing a renaissance owing to its broad range of usability. In an empirical study, Rish (2001) concludes that the Naive Bayes classifier is very effective in practice, even though its probability estimates are in theory less accurate than other classifiers. Hand et al (2001) make the case for the Naive Bayes algorithm because of its intrinsic simplicity, which means low variance in the probability estimates and thus greater estimation accuracy. Kotsiantis and Pintelas (2004) show that Naive Bayes is the most flexible learning method. Its accuracy can be boosted over most methods in less time for training.

¹⁴See <http://www.jane16.com>, last accessed 20 September 2009.

A.2

The Visual Basic Programs were written in order to handle the vast amount of articles and process them into a suitable format for the java program that features the sentiment algorithm. When downloading the articles from the *LexisNexis* database, the articles of one day are summarized in one text file. The first program was written to cut the articles into separate text files in order to format them for the java program that runs the sentiment analysis. The output log-file from the java program was then formatted and coded into $\{-1\}$ for negative sentiment and $\{1\}$ for positive sentiment. Neutral values are not coded by the algorithm in order to avoid ambiguity. The daily sentiment data were then aggregated to quarterly values.

A.3

In this paper, the Augmented Dickey-Fuller Unit Root Test is utilized because the model follows a higher order AR-process. Take the following equation

$$\Delta y_t = \alpha_0 + \gamma y_{t-1} + \sum_{i=2}^p \beta_i \Delta y_{t-i+1} + \varepsilon_t,$$

where $\gamma = -\left(1 - \sum_{i=1}^p a_i\right)$ and $\beta_i = -\sum_{j=1}^p a_j$. The Null hypothesis tests whether $\gamma = 0$, and if so, the equation is entirely in first differences and so has a unit root. If $\gamma \neq 0$, then the equation does not have a unit root.

A.4

The Theil Inequality Coefficient is calculated as follows:

$$\frac{\sqrt{\sum_{t=T+1}^{T+h} (\hat{y}_t - y_t)^2 / h}}{\sqrt{\sum_{t=T+1}^{T+h} \hat{y}_t^2 / h} + \sqrt{\sum_{t=T+1}^{T+h} y_t^2 / h}}$$

where the forecast sample is $j = T + 1, T + 2, \dots, T + h$, and the actual and forecasted value in period t is y_t and \hat{y}_t , respectively.

The Root Mean Squared Error (RMSE) is calculated as follows:

$$\sqrt{\sum_{t=T+1}^{T+h} (\hat{y}_t - y_t)^2 / h}.$$

A.5

The White covariance matrix as in White (1980) is given by:

$$\hat{\Sigma}_W = \frac{T}{T-k} (X'X)^{-1} \left(\sum_{t=1}^T u_t^2 x_t x_t' \right) (X'X)^{-1},$$

where T is the number of observations, k is the number of regressors, X is the variable matrix, and u_t is the least squares residual.

Fig. 1: Information Flows to Consumers

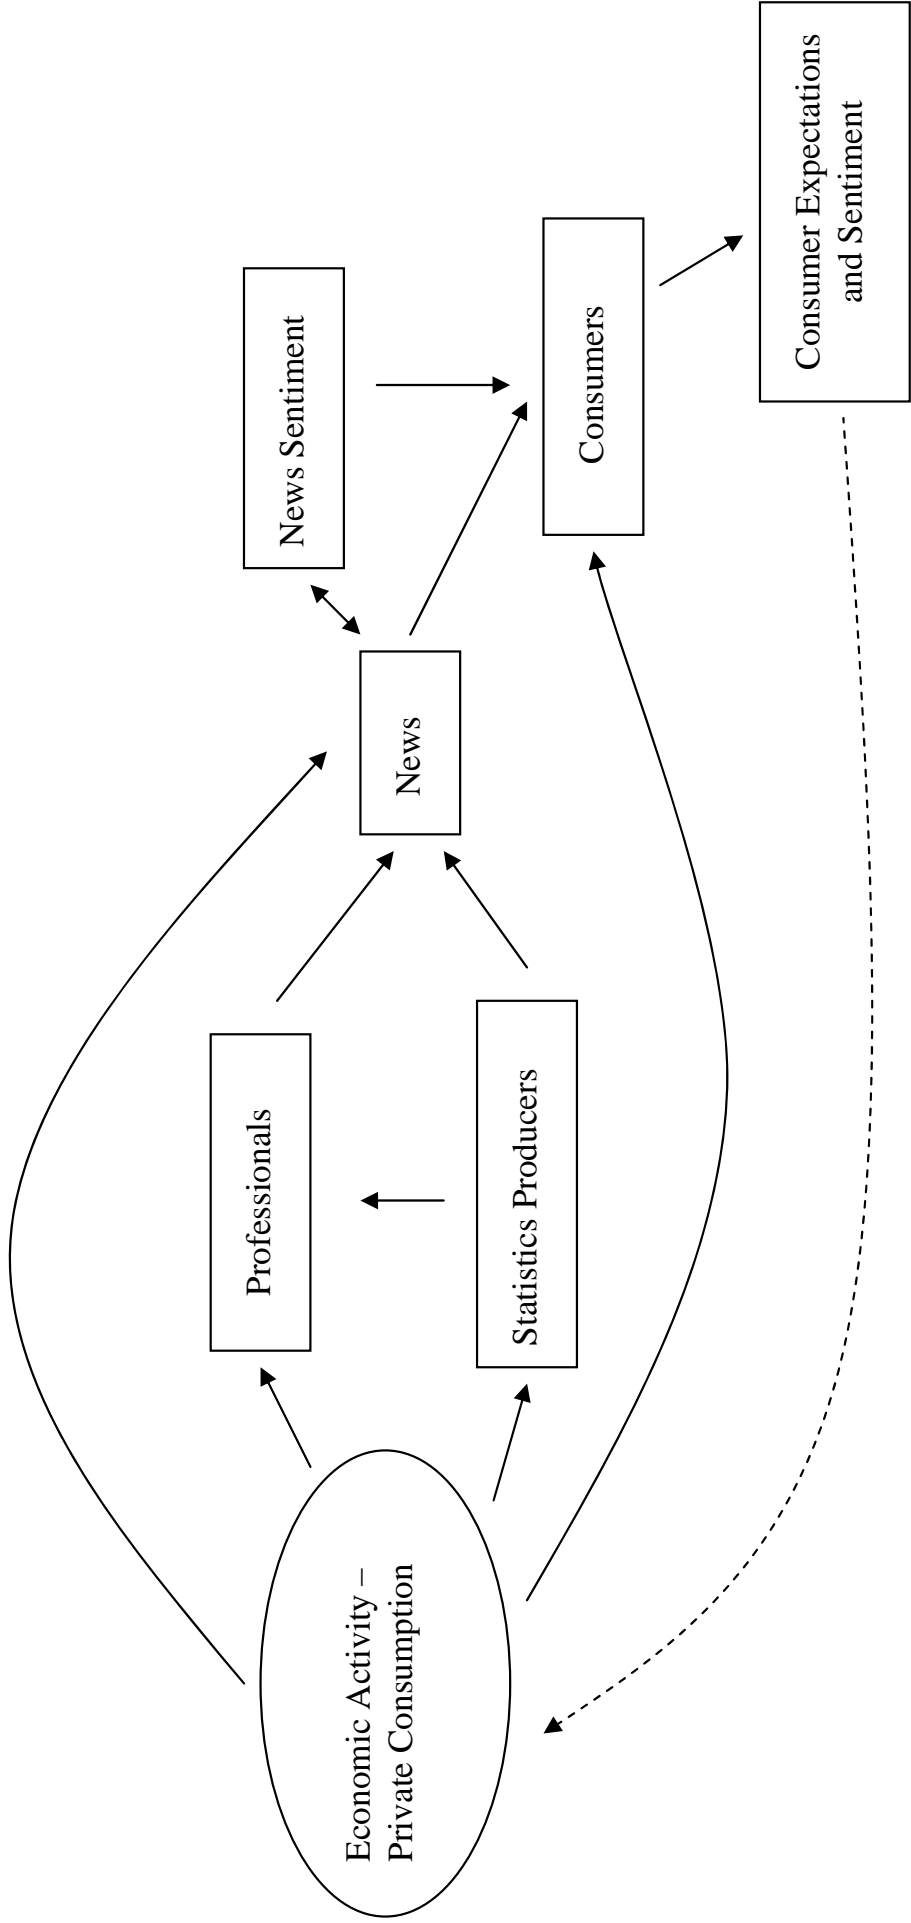


Fig. 2
News Sentiment Index and University of Michigan Index

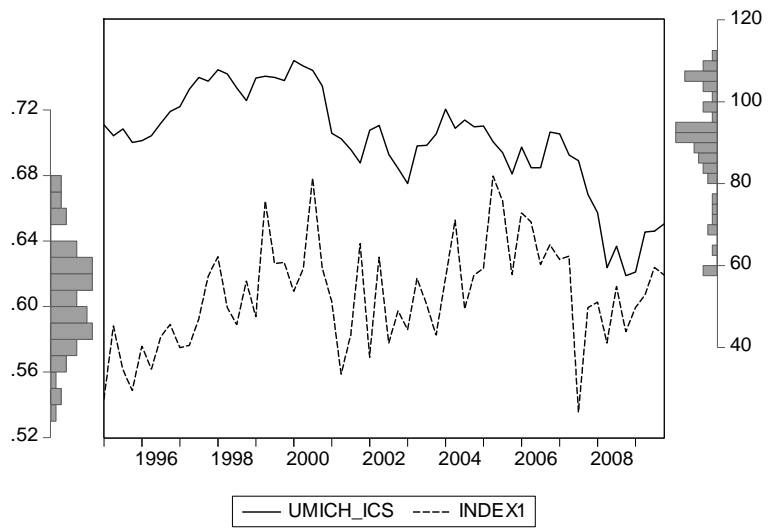


Fig. 3
Personal Income (differenced) and University of Michigan Index

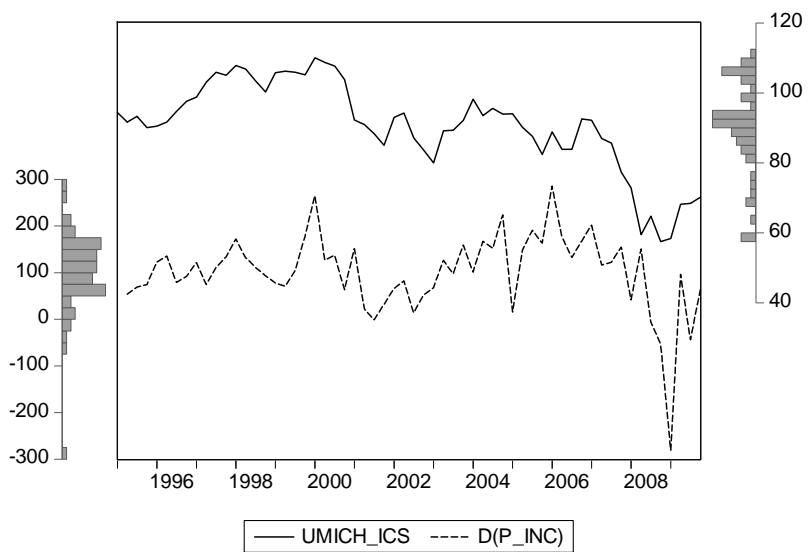


Fig. 4
Consumer Price Index (differenced) and University of Michigan Index

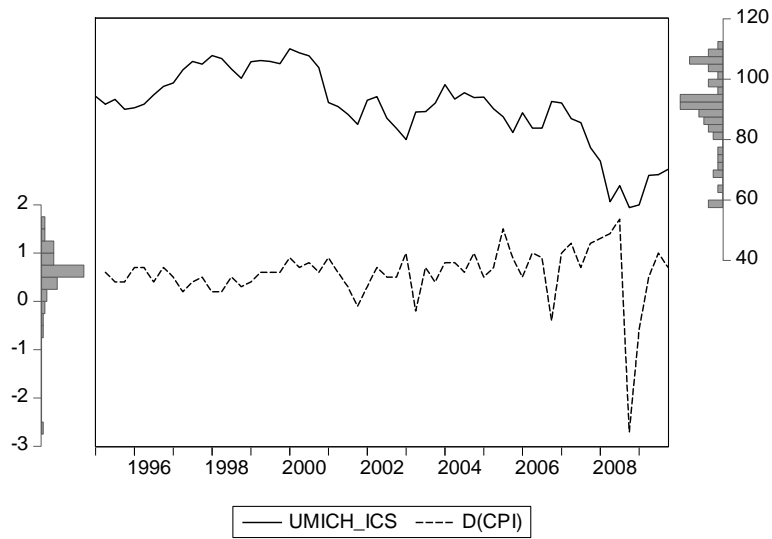


Fig. 5
Long-term interest rates (10-year US-Treasuries) and University of Michigan Index

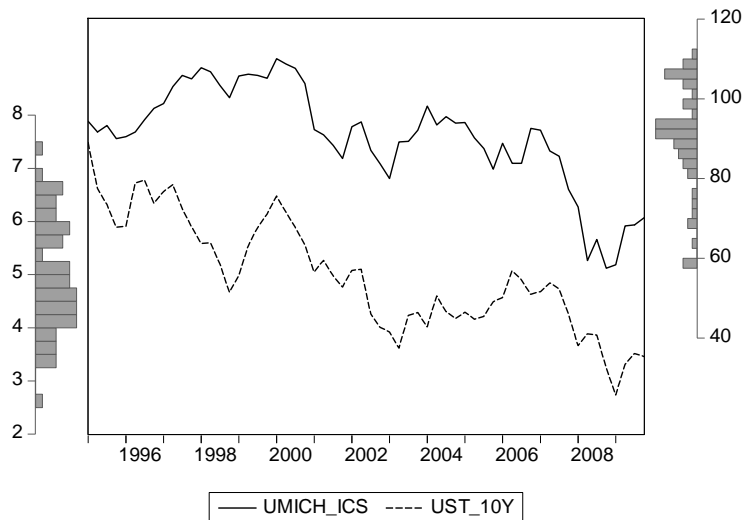


Fig. 6
University of Michigan Index and Private Consumption (log differenced)

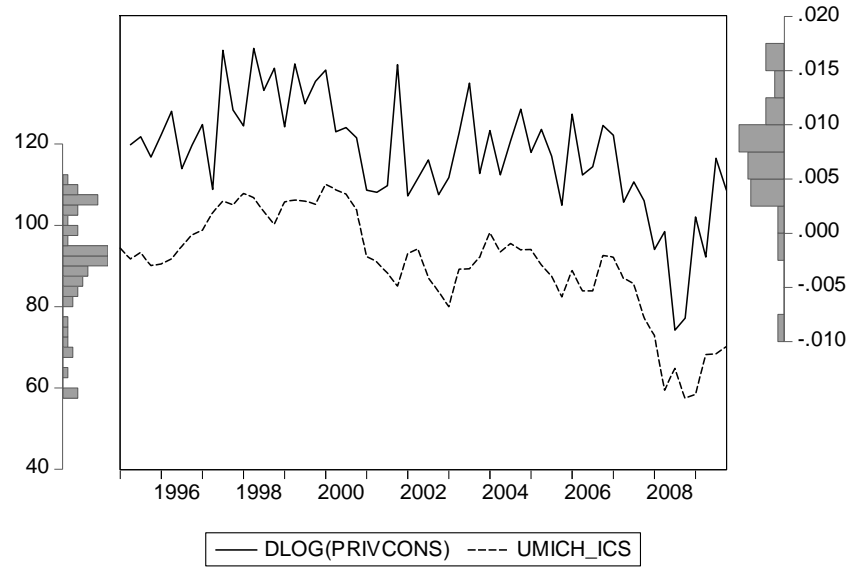


Fig. 7
News Sentiment Index and Private Consumption (log differenced)

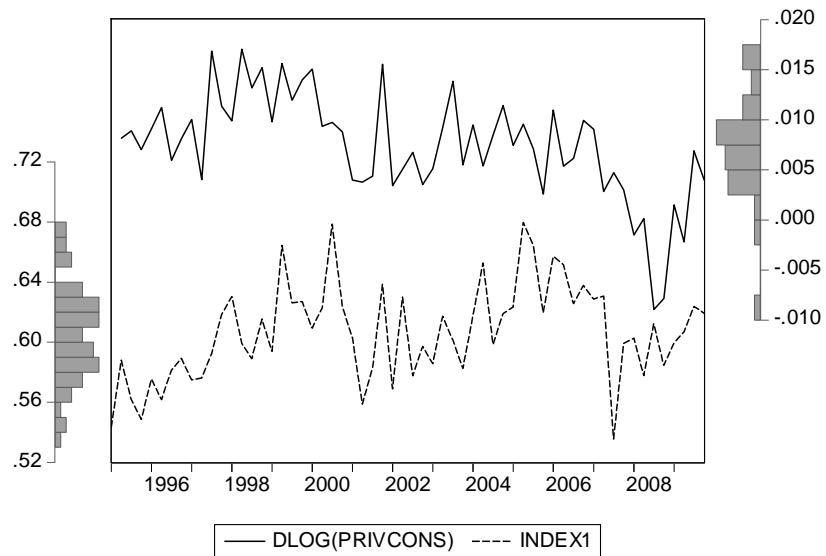


Fig. 8
Consumer Price Index (differenced) and Private Consumption (log differenced)

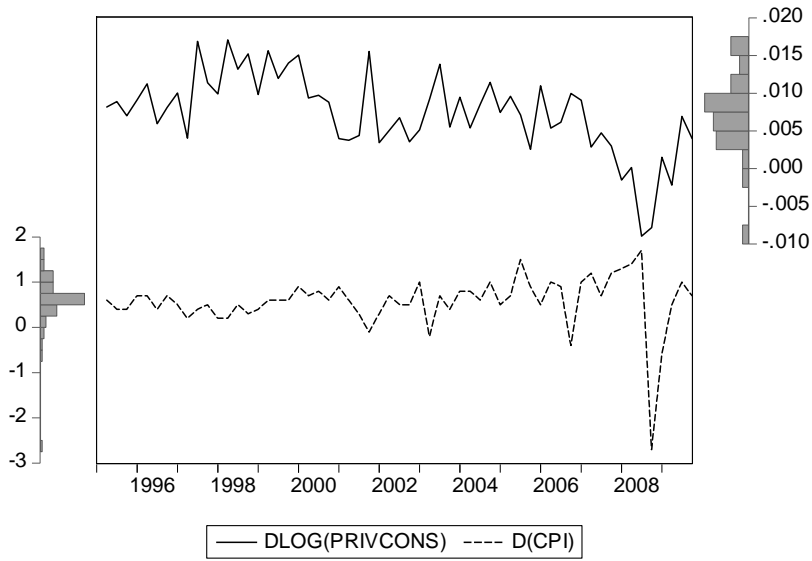


Fig. 9
Personal Income (differenced) and Private Consumption (log differenced)

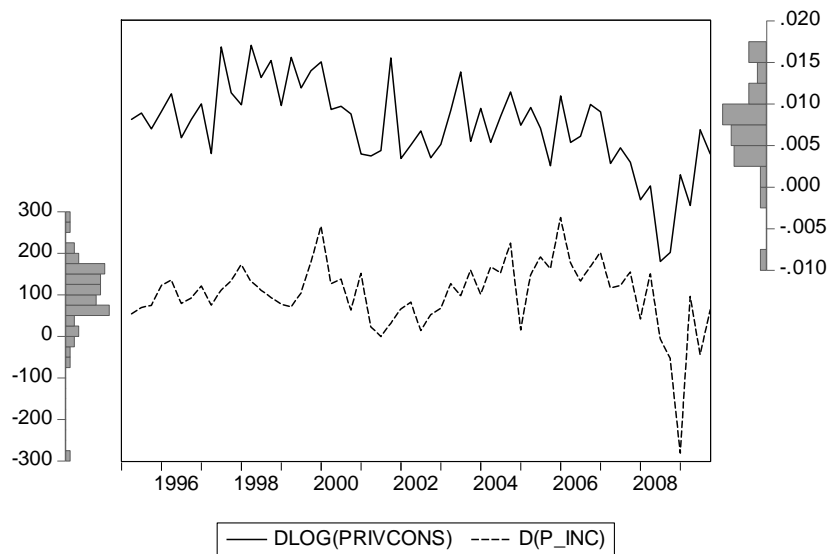


Fig. 10
Long-term interest rates (10-year US-Treasuries) and Private Consumption (log differenced)

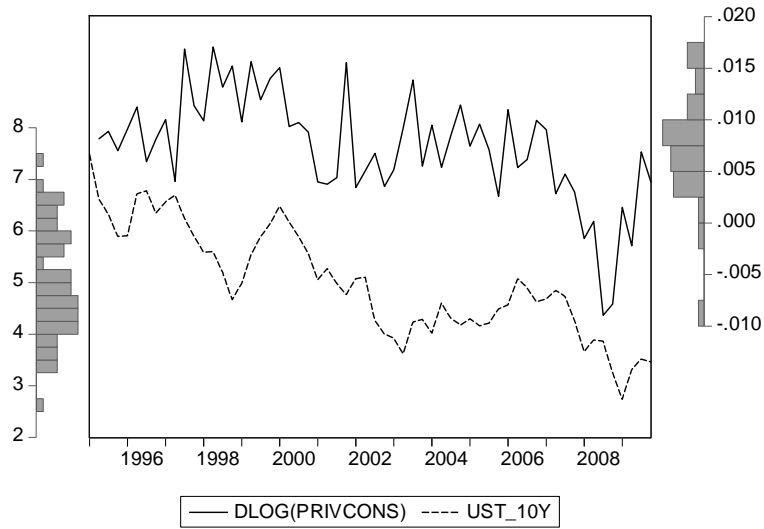


Table 1

Newspaper Statistics

	Number of articles examined for sentiment 1995 - 2009	Average Daily Circulation*
<i>USA Today</i>	28'832	1,826,622
<i>Washington Post</i>	74'206	604'650
<i>Houston Chronicle</i>	30'919	494'131
<i>New York Times</i>	114'454	951'063
<i>Wall Street Journal - Abstracts</i>	74'420	2'092'523
Total	322'831	

* Source: Audit Bureau of Circulation Survey 31/3/2010, (<http://abcas3.accessabc.com/ecirc/newstitlesearchus.asp>), last accessed 9 June 2010.

Table 2

Lag Length Selection Test

	Optimal lag length according to Schwarz info criterion in ARMA(4,1) model with private consumption as dependent variable
<i>News Sentiment Index</i>	0
<i>Consumer Price Index</i>	1
<i>Personal Income</i>	0
<i>Long-Term Interest Rates (10-year US Treasuries)</i>	0
<i>University of Michigan Index of Consumer Sentiment</i>	0

Table 3

Regression Coefficient Estimates of ARMA(4,1) models
(standard errors in parentheses beneath coefficients)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Dependent Variable:</i>							
University of Michigan Index of Consumer Sentiment (differenced)							
<i>Independent Variables</i>							
News Sentiment Index (differenced)		19.52144 (20.26948)					38.90719** (19.17444)
Personal Income (differenced)			0.005395 (0.010652)			0.009711 (0.008802)	0.009118 (0.007846)
Consumer Price Index (differenced)				-0.716099 (1.715577)		-3.138235* (1.632152)	-3.278057** (1.381218)
Long-Term Interest Rates (10-year US Treasuries)(differenced)					3.026012* (1.753438)	4.718697*** (1.570301)	5.221713*** (1.472796)
Constant	-0.764253*** (0.088976)	-0.759040*** (0.081889)	-1.341772) (1.076574)	0.97786 (0.972787)	-0.596541*** (0.138043)	0.423183 (0.914026)	0.58409 (0.810746)
Adjusted R-squared	0.141964	0.14279	0.129859	0.214197	0.165155	0.234055	0.280482
N (after adjustments)	55	55	55	55	55	55	55
Schwarz criterion	6.109009	6.160704	6.175677	6.073728	6.134268	6.152181	6.141006
Root Mean Squared Error (RMSE)	4.277697	4.232664	4.26447	4.052538	4.177085	3.918502	3.757267
Theil Inequality Coefficient	0.023418	0.02317	0.023334	0.022243	0.022862	0.021397	0.02052

*, **, *** denote statistical significance at the 10%-, 5%-, and 1%-level, respectively

Table 4

Regression Coefficient Estimates of ARMA(4,1) models
(standard errors in parentheses beneath coefficients)

	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)
<i>Dependent Variable: Private Consumption (log difference)</i>										
<i>Independent Variables</i>										
News Sentiment Index (level)		0.036414* (0.019793)						0.037696* (0.020854)		0.024616* (0.0123)
Personal Income (difference)			0.0000102 (0.000009)				0.0000205** (0.00000821)	0.0000211** (0.00000810)	0.0000126** (0.00000594)	0.00000775 (0.000006)
Consumer Price Index (difference)				-0.001648** (0.000734)			-0.002429*** (0.000788)	-0.002419** (0.000949)	-0.002648*** (0.000936)	-0.002497*** (0.000927)
Long-Term Interest Rates (10-year US Treasuries)(difference)					0.00026 (0.001370)		-0.000715 -0.001624	-0.000680 (0.001631)	0.00026 (0.001847)	
University of Michigan Index of Consumer Sentiment (level)						0.000323*** (0.0000301)			0.000291*** (0.000266)	0.000257*** (0.000267)
Constant	0.006294** (0.003087)	-0.015127 (0.013234)	0.004555* (0.002659)	0.008125*** (0.001956)	0.006727*** (0.001637)	-0.022020*** (0.002767)	0.004442 (0.005960)	-0.033701 (0.081311)	-0.018866*** (0.002454)	-0.034072*** (0.008334)
Adjusted R-squared	0.446739	0.448027	0.457082	0.467738	0.435658	0.583035	0.504566	0.525854	0.626527	0.654942
N (after adjustments)	55	55	55	54	54	55	54	54	54	54
Lags of independent variables (in order from top to bottom)	n/a	0	0	1	0	0	0,1,0	0,0,1,0	0,1,0,0	0,0,1,0
Schwarz criterion	-7.892251	-7.841924	-7.858465	-7.865425	-7.819763	-8.122421	-7.831945	-7.823974	-8.062641	-8.141775
Root Mean Squared Error (RMSE)	32.35183	32.03475	31.77118	31.28499	32.32521	27.18089	29.51703	28.66371	24.85084	23.89813
Theil Inequality Coefficient	0.001983	0.001964	0.001948	0.00191	0.001982	0.001666	0.001802	0.00175	0.001518	0.001459

***, **, * denote statistical significance at the 10%, 5%, and 1%-level, respectively