STOCKHOLM SCHOOL OF ECONOMICS

5350 Master Thesis in Economics

FORECASTING HOUSEHOLD CONSUMPTION

A NON-LINEAR APPROACH

Eric Ramstedt*
Student no. 40446

Abstract

This paper aims at improving the forecasting capabilities of the household consumption model currently in use by the National Institute of Economic Research. The mortgage discount, calculated as the difference between the banks' official mortgage rates and the rates that the households actually pay, is added in both a linear and a non-linear fashion as an explanatory variable to an existing model based on the permanent income hypothesis. The result from the estimations are rather mixed and clear conclusions are hard to find. No improvements to forecasting performance are made, but the non-linear model suggests there might be a relationship between a decreasing mortgage rate discount and drops in household consumption whereas an increase in the mortgage rate discount does not seem to lead to an increase in household consumption.

Supervisor: Rickard Sandberg
External Supervisor: Pär Stockhammar
Examiner: Örjan Sjöberg
Discussant: Benjamin Rosenthal

JEL classifications: D1, D9, E27
Keywords: Household consumption, mortgage rate discount, non-linear forecasting

Date of submission: 06-01-2014
Date of presentation: 14-01-2014

*40446@student.hhs.se

Acknowledgements

This master thesis is the result of a collaboration between the National Institute of Economic Research (Konjunkturinstitutet) and the Stockholm School of Economics.

Firstly, I would like to thank Pär Stockhammar at the National Institute of Economic Research for both taking me in on this project, and for providing valuable advice and help throughout the process. I am also very grateful for the help, assistance and advice of Pär Österholm and Bengt Assarsson (both at the National Institute of Economic Research as well).

Furthermore, Ronny Andersson at Statistics Sweden deserves a mentioning. He has been very helpful in the data acquiring process, and has clearly taken a personal interest in this project that goes beyond the scope of his regular duties.

Lastly, my main thanks goes out to Prof. Rickard Sandberg at the Stockholm School of Economics. In the role as supervisor he has been more than helpful, and it is no overstatement to say that the completion of this project would not have been possible without him.

Eric Ramstedt, January 2014

Table of contents

1. Introduction 1
2. Current state of knowledge 3
  2.1 Previous research 3
    2.1.1 The Life-cycle / Permanent income hypothesis 3
    2.1.2 Consumer Sentiment Index 5
    2.1.3 Current NIER model 7
  2.2 Non-linear time series modeling 8
3. Method 10
  3.1 Empiric strategy 10
    3.1.1 Specification 10
    3.1.2 Estimation 11
    3.1.3 Evaluation 12
  3.2 Data 13
    3.2.1 Household consumption 13
    3.2.2 Real interest rate 14
    3.2.3 Mortgage rate discount 15
    3.2.4 Variable correlations 16
4. Results 19
  4.1 Interpretation and internal validity 19
  4.2 Implications 21
  4.3 Limitations 23
5. Conclusion 24
6. References 26
Appendix A 28
  A.1 Maximum Likelihood Estimation 28
  A.2 TAR estimation in MATLAB 29
Appendix B 34
  B.1 Pre-testing the data 34
  B.2 Thresholds and non-linearity testing 35
    B.2.1 Finding ‘c’ 35
    B.2.2 Non-linearity 36

1. Introduction

At the very heart of the economic science lies the challenge of forecasting. Forecasts within economics are of essence in both the public and the private sector; and are of great importance not only for the government's budgetary planning, but also for firms and households. Among the national accounting variables, household consumption is a particularly hard one to forecast.[1] This is becoming increasingly problematic for Swedish GDP forecasts, given that household consumption is rising in importance for Swedish GDP.[2] There is far from any consensus within the academic society on how to best make prognoses of household consumption, and the methods differ as widely as from standard Keynesian consumption functions to household surveys.

At the same time as the problem above has surfaced, another phenomenon has reached the attention of the Swedish media: the discount on mortgage loan rates. The mortgage rate discount, defined as the difference between the banks' official mortgage rates[3] and the rates that the households actually pay (or in other words: how much you haggle), emerged as a topic of interest during the previous financial crisis when the seemingly big differences between households became apparent.[4]

To this author's knowledge, the mortgage rate discount has never been properly examined before. When modeling household consumption, the real interest rate is typically one of the relevant variables, as a lower real interest rate should increase consumption and a higher real interest rate should make us defer consumption. The interest rate variable is often thought to incorporate the cost of mortgages, as higher interest rates usually lead to higher mortgage payments.

But what if the real interest rate and the mortgage rate discount represent different effects? If the real interest rate and the mortgage rate discount correlate positively, an increase in the real interest rate (which should reduce consumption) could also mean an increase in the mortgage rate discount (which should increase consumption), and valuable information might potentially be ignored when the mortgage rate discount is not included. After all, two thirds of the Swedish population lives in mortgaged property – and it is fair to say that how much you pay on your mortgage will in some manor affect how much you can spend on consumption. In the form of percentage points the mortgage rate discount might not vary that much – but even a one percentage point change will bring about a rather significance difference in disposable income if you have a mortgage of two or three million SEK.[5]

[1] Stephens, M. (2008). *The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans*. Review Of Economics And Statistics 90(2): p 241-252.
[2] Ekonomistyrningsverket (2013). *Statens budget och de offentliga finanserna.* Stockholm: Publikationsservice, p 10.
[3] ”Official mortgage rates” refer to the list prices found e.g. in newspapers and on websites.
[4] Zachrison, O. (2012, April 23rd). *Unik räntekarta med din hjälp.* Svenska Dagbladet, online edition.
[5] Finansinspektionen (2013). *Den svenska bolånemarknaden 2013.* Stockholm: Publikationsservice, p 4.

As commissioned by the National Institute of Economic Research (henceforth NIER), this paper therefore has two aims. The first aim, an underlying aim, is to properly investigate the Swedish mortgage rate discount. Mainly, it is interesting to see if the discount is something that varies over time, or if it is relatively constant. If it does indeed vary over time, it could possibly add explanatory power to consumption models.

The second aim, the main aim, of this paper is to try to ameliorate the forecasting properties of NIER's current consumption model, which builds on the permanent income hypothesis, by adding the mortgage rate discount as an explanatory variable in both a linear and a non-linear setting. When using both a linear and a non-linear model one allows for the possibility that the mortgage rate discount can potentially have different effects in different regimes, for instance by having one type of effect when the mortgage rate discount is negative and one type of effect when the mortgage rate discount is positive. It is of course possible that consumption models could benefit from non-linearity in other variables apart from only the mortgage rate discount (for instance in the real interest rate), but as agreed upon in collaboration with NIER this paper will focus on the effect of the mortgage rate discount.

The rest of this paper is organized as follows. Section 2 describes the current state of knowledge both with regards to household consumption forecasting and to non-linear time series modeling, and it also describes the current model in use by the National Institute of Economic Research. Section 3 presents the method and model of this paper, as well as the data used. Section 4 presents the results of the estimation, and the analysis. Section 5 concludes.

2. Current state of knowledge

2.1 Previous research

The previous research on the particular topic of this thesis, the addition of a (non-linear) mortgage rate discount variable to household consumption forecasting models, is more or less non-existent, and one of the main motivations of the undertaking of this project. There exists, however, a vast literature on the more general topic of forecasting household consumption, as described below.

There are two main schools of economic thought with regards to consumption forecasting: using the Life-Cycle / Permanent income hypothesis and other macroeconomic variables, or using "Consumer Sentiment Indices". These two (sometimes conflicting) theories will both be elaborated on, in turn, together with a presentation of the model currently in use by the National Institute of Economic Research.

2.1.1 The Life-cycle / Permanent income hypothesis

The most widely used theories when forecasting household consumption are the 'Life-Cycle' and 'Permanent income' hypotheses (henceforth LC-PIH), as developed in parallel over 60 years ago by Modigliani and Friedman, respectively. They both refuted Keynes idea that consumption mostly depends on current income, and claimed that it rather depends on the expected income over the whole life span.[6]

Keynes sprung the modern debate by arguing that consumption decisions mainly are dependent on the level of income, and that households then have a marginal propensity to consume somewhere between 0 and 1. In addition to this he believed in a marginal propensity to consume that was diminishing in income, which meant that extra money should matter less for consumption the higher your income is. This was a step away from consumption models based on interest rate levels, which had typically dominated the economic literature up to this point.[7] Simon Kuznets, however, refuted Keynes' idea by showing that consumption patterns in US data seemed rather stable in the long run, even though income varied.[8]

The argument of Modigliani and Friedman is, instead, that with reasonably well functioning financial markets and rational agents, households will try to smoothen (by borrowing/lending) their consumption given both their current and their expected future income. Putting this another way, it means that given a constant real interest rate, well functioning financial markets and perfect foresight, no changes to real consumption will ever occur as households are already maximizing their utility by smoothing consumption over the

[6] Dion, D.P. (2006). *Does Consumer Confidence Forecast Household Spending? The Euro area case.* MPRA paper 911: p 3f.

[7] Keynes, J.M. (1936). *General Theory of Employment, Interest and Money*. London: Macmillan.

[8] Kuznets, S. (1942). *Uses of National Income in Peace and War.* New York: National Bureau of Economic Research.

life cycle. This in turn, which is the main point, means that the state of the world today is summarized in the consumption decision made today – and thus means that no other variables from the past should be relevant when explaining future consumption.[9]

The two theories (today mostly referred to as one single theory), can be summarized in accordance with e.g. Campbell and Deaton through the following equation

$$\Delta C_t = \frac{r}{1+r}\sum_{k=0}^{\infty}(1+r)^{-k}(E_t - E_{t-1})Y_{t+k} \qquad (1)$$

where C is consumption, r is the real interest rate, Y is income and E is the expectation (thus stating that changes in consumption will only be affected by changes in the real interest rate or in the expected future income, as household chose their level of consumption based on the expected lifetime earnings).[10]

Ever since the introduction of the LC-PIH there has been a great debate as to whether the postulated theory actually holds, and over 50 years later there is still no consensus among economists worldwide. Most do however agree that predicting consumption and household expenditure is far from easy. Marjorie Flavin (among others) has tried to estimate consumption using simple univariate models with lagged consumption or lagged income, but with little success. She concludes that simple univariate models, although appealing, will most likely never work as the households themselves decide upon expenditures and investment with much more information at hand than the econometrician trying to perform a forecast.[11]

One effort to deduce whether the LC-PIH holds was undertaken by Davis and Palumbo, where they used the stock market boom in the later part of the 1990s in the US to try to estimate how a sudden increase in wealth (for many) affects household consumption. The LC-PIH, as outlined above, predicts that agents are perfectly rational and that consumption will be rather smooth, even though wealth can vary from one period in life to another. If the hypothesis holds, Davis and Palumbo claim, there would be no significant increase in consumer spending following the stock market boom, as the rational households would have incorporated swings in the stock market into their level of expenditures.[12]

Using an error-correction framework they conclude that there is a certain increase in consumer expenditures following the increase in wealth. They notice a gradual adaption between wealth and consumption, but observe neither sharp increases nor a status quo. The authors do, however, admit that the results of their error-correction framework are very

[9] Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton: University Press, p 7.
[10] Campbell, J.Y. & Deaton, A. (1989). *Why is consumption so smooth?* Review Of Economic Studies 56(3): p 357-373.
[11] Flavin, M. (1993). *The Excess Smoothness of Consumption: Identification and Interpretation*. Review Of Economic Studies 60(3): p 651-666.
[12] Davis, M. A. & Palumbo, M.G. (2001). *A primer on the economics and time series econometrics of wealth effects*. Divisions of Research & Statistics and Monetary Affairs: Federal Reserve Board, p 2ff.

sensitive to model specification, and conclude by conceding that their paper is most likely not the last one on the topic of the LC-PIH.[13]

The main problem of the approach described above is that it is more or less impossible for the economist to know if the households have already anticipated a stock market increase when they decide on their consumption level. To reconcile this problem, Melvin Stephens Jr. approach the problem of the LC-PIH from another angle by using data over repayment of vehicle loans in the US. The idea of the paper is to track, on an individual basis, the consumption of households during the repayment scheme of the loan. Since the final date of the payment of the loan is known from the beginning, a rational consumer under the LC-PIH should not display an upswing in expenditures once the final payment is made. What he observes is on the other hand that consumption expenditures increases with roughly 2-3 % after the final payment is made, thus suggesting that households and agents are not as perfectly rational as suggested by the LC-PIH.[14]

In addition to the discussion of whether the LC-PIH approach is the correct one with regards to consumption forecasting, there has been an upswing in literature focusing on the fact that consumption forecast are harder to make now than they were before the massive deregulation of the financial system in the western economies in the late 1980s.

As examples, both Carruth & Henly and Eitrheim, Jansen & Nymoen in separate studies conducted in the US and in Norway, respectively, conclude that their old models lost considerable predictive power during the end of the 1980s.[15] [16] This more general problem for the traditional consumption models, based on the LC-PIH and other "hard" macroeconomic variables such as inflation, income and wealth, to make reliable forecasts has however resparked the interest for forecasting via a method that is as much psychology as it is economics: the method of consumer sentiment indices.

2.1.2 Consumer Sentiment Index

To complement the forecasting models based on traditional macro variables, the method of 'Consumer Sentiment Index' (henceforth CSI) surveys was developed at the University of Chicago in the 1950s. The consumer sentiment index is, in essence, a set of questions spanning from current and future outlook on the economy to current and future plans about expenditures, which are then conveyed to a sample of households in the economy. The answers to the different questions are then put together into one index: the CSI.

James Wilcox describes the benefits of CSI by arguing that there is always a delay or lag from the time of changed household outlooks to the time when it is actually noticeable in the national accounting. For example, the oil price shock of 1973, or the appointment of a new Federal Reserve chairman known for advocating an expansive monetary policy, will affect households' future expectations about the economy long before a change in household

[13] ibid, p 40.
[14] Stephens (2008), p 241-252.
[15] Eitrheim, O., Jansen, E.S. & Nymoen, R. (2002). *Progress from Forecast Failure--The Norwegian Consumption Function*. Econometrics Journal 5(1): p 40-64.
[16] Carruth, A. & Henley, A. (1990). *Can existing consumption functions forecast consumer spending in the late 1980's?* Oxford Bullentin Of Economics And Statistics 52(2): p 211-222.

consumption shows up in the national accounting. With a CSI survey, Wilcox claims, these changed expectations can be noticed much quicker, thus allowing a better forecast to be made.[17]

In his own review of the current literature on CSI, Wilcox notes that there is substantial variation with regards to the marginal benefits of adding CSI to consumption forecast models. He does claim, however, that the most recent literature tends to point towards rather small but still significant contributions made by the CSI. In a paper of his own, he finds evidence much in line with the recent literature. By using a baseline consumption function based on disposable income, non-home equity, home-equity (the last two representing wealth), inflation, interest rate level and expenditures on durables, non-durables and vehicles, he then adds the CSI in both a one-quarter and a four-quarter forecast. He finds that the 4-quarter forecast is significantly better (ex post) than the baseline model, but he also finds that the 1-quarter forecast is only as accurate as a model using solely the traditional income and wealth variables.[18]

Dion reaches a similar conclusion in his review of the current state of empirical knowledge with regards to CSI, i.e. that the evidence is very mixed. According to him, some authors have tried to prove the non-contribution of the CSI by estimating a regression such as equation 2

$$CSI_t = \lambda + \delta X_t + \varepsilon_t \tag{2}$$

where λ is a constant and X_t a vector of macroeconomic variables. If the R^2, the explained variation, of the regression were sufficiently high that would be evidence of the non-contribution of the CSI (as the variation then is captured by the macroeconomic variables). Again, however, the empirical evidence is rather mixed and conclusive evidence is hard to find.[19]

Dion also concludes that some types of questions within the CSI survey (the exact questions vary a lot) have greater predictive power than others, and that maybe the compounded CSI is not the most useful.[20] Juster makes a similar observation when he decomposes the index and investigates the forecasting power of the different types of question. His paper dates a few years back, but his conclusion – that it is the questions regarding consumer buying intentions that have the greatest forecasting capacity – is yet to be refuted.[21]

Gàbor Vadas, at the time at the Hungarian central bank, follows a rather similar train of thought in his research when he investigates which questions about consumer confidence that have a good predictive power. He relies on a simple empirical strategy, where he uses a baseline LC-PIH econometric model and then measures adjusted R^2 after adding different subquestions of the survey (using time series from both the US and Hungary). The findings

[17] Wilcox, J.A. (2007). *Forecasting Components of Consumption with Components of Consumer Sentiment*. Business Economics 42(4): p 2.
[18] ibid, p 5ff.
[19] Dion (2006), p 10.
[20] ibid, p 12f.
[21] Juster, F.T. (1960). *Prediction and consumer buying intentions*. American Economic Review 50(1): p 14.

seem to converge to other literature, in that he finds little conclusive evidence neither from the question about the general economic outlook, nor from specific consumer questions such as e.g. about planned vehicle purchases. He does however find a significant predictive power in the questions about planned expenditures in the upcoming half-year or year; indicating that the households' outlook about the future can best be read in how they plan their expenditures.[22]

2.1.3 Current NIER model

The conclusion to be drawn from the literature review and the current state of knowledge is that there is far from any consensus on how to predict household consumption. The current consumption model in use by the NIER, which will serve as the base line model of this paper, builds directly on that state of knowledge in the sense that it is a very parsimoniously specified model – a result of the very mixed evidence on many of the features outlines above.

Bengt Assarsson, head of the household consumption forecasts at NIER, has – using a simple LC-PIH model as the foundation – tried to improve forecasting by adding variables such as wealth, income, inflation and car purchases, but with little result.[23] The parsimonious design is thus in line with e.g. Carruth & Henly or Eitrheim, Jansen & Nymoen, who also concluded that the "hard" macroeconomic variables (such as wealth or income) seem to have lost predictive power. The NIER has also tried to augment and improve forecasting behavior by the use of CSI:s, but again with little improvements made.[24] Some improvements might occasionally be made by CSI:s, but as pointed out by e.g. Wilcox and Dion it is hard to find consistent results.

Given the background provided on the LC-PIH provided above, describing NIER's current model is fairly simplistic. Building on the work of Robert E. Hall, the model is based on the LC-PIH and contains only last period's consumption and changes in the (real) interest rates. Hall concludes, after having empirically tested to include variables such as income and wealth, that *"[…] no variable apart from current consumption should be of any value in predicting future consumption"*, thus following directly in line of Modigliani and Friedman.[25] He does, however, also conclude that *"A higher expected real interest rate makes consumers defer consumption, everything else held constant".*[26]

In essence then, the NIER model is based on the fact that today's consumption should be the only variable affecting future consumption (as today's consumption incorporate knowledge of the expected lifetime income), but also on the fact that a higher real interest rate may make households defer consumption. This gives the following NIER model, and baseline model of this paper, just as specified by Robert E. Hall already in 1988:[27] [28]

---

[22] Vadas, G. (2005*). Beyond macro variables: consumer confidence index and household expenditure in Hungary*. Microeconomics 0512006: p 9ff.
[23] Assarsson, B. Personal interview (20130807).
[24] ibid.
[25] Hall, R.E. (1978). *Stochastic implications of the Life Cycle-Permanent Income Hypotheis: Theory and Evidence.* Journal of Political Economy 86(6): p 971.
[26] Hall, R.E. (1988). *Intertemporal Substitution in Consumption.* Journal of Political Economy 96(2): p 339.
[27] ibid, p 341.

$$C_t = \alpha_0 + \alpha_1 \Delta r_t + C_{t-1} \tag{3}$$

or

$$\Delta C_t = \alpha_0 + \alpha_1 \Delta r_t \tag{4}$$

Following the reasoning above, changes in (log of) consumption, ΔC_t, are explained by changes in the real interest rate, r_t, and by a constant α_0. The specification of the model also further motivates the inclusion of the mortgage rate discount variable. The real interest rate, r_t, is thought to represent the general level of different interest rates in the economy, including mortgage rates. The model predicts that when the real interest rate increases, so do mortgage rates – and household consumption should therefore decrease. If, however, increasing mortgage rates also mean an increasing mortgage rate discount, it could be the case – as mentioned in the introduction – that important information and explanatory power is lost if the model were to not include the mortgage rate discount.

The model in equation (4) will, for the rest of this paper, be referred to as the 'baseline model', on which the augmentations with the mortgage rate discount will build.

2.2 Non-linear time series modeling

While the reader is assumed to possess a general understanding of econometric time series modeling, the features of non-linear time series might not be considered common knowledge. Alas, the following subsection will present the basic theory needed with regards to the models used in this paper. For readers already familiar with non-linear econometric time series theory, this subsection can easily be skipped. For the attentive reader further elaborations are made in Appendix A, where both the method of maximum likelihood estimation as well as the MATLAB code performing the estimations are presented.

That not all economic time series can be adequately described by linear models is nowadays to be seen as generally accepted. For instance, there are several examples of how unemployment rises steeply in a recession, but only declines rather slowly once the economy expands again. The same thinking will be applied in this paper, but in the sense that the mortgage rate discount could have different effects on household consumption in different 'regimes' (an example of two different regimes could be an economy that is booming and an economy that is busting). The non-linearity in this paper is therefore not motivated in theory itself, but rather on the empirical observation that some variables, such as unemployment, seem to display a non-linear behavior.

The most common, and the simplest, model that can account for such regime-switching behavior is the Threshold Autoregressive (henceforth, TAR) model. The TAR model, much as the name implies, allows for a non-linear estimation given a threshold value. Using a simple AR(1) process and the arbitrary threshold c, a TAR can be represented as follows:

[28] Some of the notation is changed for consistency with this paper, but the equation remains the same.

$$y_t = \begin{cases} a_1 y_{t-1} + \varepsilon_t \; if \; y_{t-1} < c \\ a_2 y_{t-1} + \varepsilon_t \; if \; y_{t-1} \geq c \end{cases} \qquad (5)$$

This means that y_t will follow the upper process when y_{t-1} is smaller than c, and that y_t will follow the lower process when y_{t-1} is larger than or equal to c. The threshold c can then be set to almost anything, the threshold does not need to depend on the dependent variable and there could even be more than two thresholds – but the equation above outlines the TAR models in a general manner.[29]

There is much more to write on the topic of non-linear time series, especially with regards to (i) identifying the threshold and (ii) diagnostic checks on non-linear models. These points will be addressed in section 3, once the models in use have been thoroughly presented.

[29] Enders, W. (2010). *Applied Econometric Time Series.* 3rd ed. Danvers: John Wiley & Sons, Inc., p 439.

3. Method

3.1 Empiric strategy

Following the standard approach to time series modeling, the empiric strategy will be based on the following three steps: (i) specification, (ii) estimation and (iii) evaluation. In the first step (as the name suggests), the specification of the model in use is decided upon. With the specification given, a suitable method for estimation is then chosen – and lastly the models estimated are evaluated in order to the find the best model(s).

3.1.1 Specification

Based on three different model specifications, a total of nine estimations will be performed in this paper. It is important to stress that these models will be based on the LC-PIH hypothesis and NIER's empiric experience, rather than being purely 'data driven' models. Testing other (more data driven) models might be of interest for future researchers, but falls outside the scope of this particular paper.

As all models are based on first differences, it means that the time series in levels need to be non-stationary and that the time series in first differences need to be stationary (as the model would otherwise be wrongly specified). The stationary condition also applies to the non-linear model, where the time series need to be stationary in both regimes.[30] The fulfillment of these criteria, in the form of Augmented Dickey-Fuller tests, is presented in Appendix B.

The first subcategory of these models is the baseline model, just as specified in the former section. The model below is thus the same model as the one in use by the NIER (equation 4)

$$\Delta C_t = \alpha_0 + \alpha_1 \Delta r_{i,t} + \varepsilon_t \tag{6}$$

where $i \in \{1,2,3\}$ represent the three different data sets of the real interest rate that will be utilized (more on this in the following subsection), and where ε_t is assumed to be $iid \sim N(0, \sigma_\varepsilon^2)$. This specification gives that the baseline model, in time series jargon, is a random walk with a drift term. For such a model to be appropriate, the (log of) level of the data series needs to display a trend, or, equivalently, that the first difference of the time series should not have a zero mean. Figure 1 confirms this on page 14, where the consumption data is plotted.

Following the baseline model, a linear augmentation with the mortgage rate discount is the next subcategory. The mortgage rate discount variable is simply added in a linear fashion, giving the model as

[30] Enders, W. & Granger, C.W.J. (1998). *Unit-Root Tests with an Asymmetric Adjustment With an Example Using the Term Structure of Interest Rates.* Journal of Business & Economic Statistics 16(3): p 304.

$$\Delta C_t = \alpha_0 + \alpha_1 \Delta r_{i,t} + \alpha_2 \Delta disc_t + \varepsilon_t \tag{7}$$

where $i \in \{1,2,3\}$ again represent the three different real interest rates, and where ε_t is still assumed to be $iid \sim N(0, \sigma_\varepsilon^2)$.

Specifying the last subcategory of models, the non-linear ones, is a bit more complicated. Given that the dataset in use is made up of quarterly data, a TAR model is preferred over a more "sensitive" model such as the STAR model (where the regime change happens gradually rather than instantaneous, as in the TAR model). The use of a STAR model for this research question would have been very interesting, but given that household consumption is only observable on a quarterly basis the frequency of the data is not high enough to motivate a STAR model as you would then typically need weekly or even daily data.[31]

Having decided on the TAR model, two specification problems remain: to decide on the number of thresholds, and to decide on the value of the threshold(s). This paper will only use one threshold in order to keep the models parsimonious, but even then the task of finding the threshold remains. This is done by creating a grid of values ranging from the minimum to the maximum of the mortgage rate discount, and then to estimate the TAR model for each and every one of the suggested grid points. For every estimation the sum of squared residuals is calculated – and the models, one for each real interest rate, with the best fits (i.e. lowest sum of squared residuals) will be the ones carried on to the estimation and evaluation phase.[32] A table presenting the selected threshold as well as a graph over the mortgage rate discount time series split up according to the two regimes are reported in Appendix B.2.

With the threshold identified at –0.3 , the model becomes

$$\Delta C_t = \begin{cases} \alpha_{10} + \alpha_{11} \Delta r_{i,t} + \alpha_{12} \Delta disc_t + \varepsilon_t \; if \; \Delta disc_t < -0.3 \\ \alpha_{20} + \alpha_{21} \Delta r_{i,t} + \alpha_{22} \Delta disc_t + \varepsilon_t \; if \; \Delta disc_t \geq -0.3 \end{cases} \tag{8}$$

where $i \in \{1,2,3\}$ remains the same as in the previous models and where ε_t is, again, assumed to be $iid \sim N(0, \sigma_\varepsilon^2)$.

3.1.2 Estimation

Given the specification above, the estimation is carried out by maximizing the likelihood function with regards to the model coefficients. The derivation of the likelihood function is found in Appendix A.1, and examples of the MATLAB code used to perform the estimations are found in Appendix A.2.

As an insurance against potential misspecification of the model, robust standard errors (in the sense that they are correct even if $\varepsilon \nsim N(0, \sigma_\varepsilon^2)$) are estimated.[33]

---

[31] Enders (2010), p 457f.

[32] Franses, P.H. & van Dijk, D. (2008). *Non-linear time series models in empirical finance.* 6th ed. Cambridge: University Press, p 89.

[33] White, H. (1982). *Maximum Likelihood Estimation of Misspecified Models*. Econometrica 50(1): p 1-25.

3.1.3 Evaluation

Having done the estimation as outlined above, the evaluation is done in by using (i) t-tests and F-tests, (ii) the Akaike and Bayesian information criterions and lastly (iii) in sample forecasting.

The t-test is used to test for statistical significance in the estimated coefficients, and is applicable both to the linear and the non-linear models.[34] In the TAR model, the F-test will be used to test if the non-linear estimates are statistically different from the estimates in the linear model (if they are not, it implies the model should be specified in a linear fashion instead).[35]

Following the statistical inference, the Akaike (AIC) and the Bayesian (BIC) information critera is also calculated. The AIC and BIC are measures of how well the model fits the data, but still favors a parsimonious model by imposing a penalty on adding more parameters. The AIC and the BIC are, respectively, computed as

$$AIC = 2k - 2\ln(L) \tag{9}$$

$$BIC = -2\ln(L) + k\ln(n) \tag{10}$$

where k is the number of parameters estimated, n is the number of observations and L is the value of the likelihood function at its maximum. The lower the AIC or BIC, the better the model fits the data in relation to the number of parameters.

Last of all, and most importantly, in-sample forecasting abilities of the nine models are tested. This is done by estimating the model using only part of the data (in this case roughly 75 % of the data, which means that the forecasting starts in the fourth quarter of 2008) to make a 1-quarter ahead forecast, calculating the difference between the estimated and the real value, and then repeating this by increasing the number of observations used in the estimation until one reaches the last observation.

If the difference between the actual household consumption and the forecasted household consumption in period T+h is

$$e_{T+h} = cons_{T+h} - E_{T+h}(cons) \tag{11}$$

then the mean square prediction error (MSPE) is given as

[34] Chan, K. S. (1993). *Consistency and Limiting Distribution of the Least Squares Estimator of a Threshold Autoregressive Model.* The Annals of Statistics 21(1): p 520-533.
[35] Hansen, B.E. (1997). *Inference in TAR models.* Studies in Nonlinear Dynamics and Econometrics 2(1): p 3.

$$MSPE = \sqrt{\frac{1}{H}\sum_{i=1}^{H} e_i^2} \quad (12)$$

where H is the number of forecasted periods. You want, quite intuitively, the MSPE to be as small as possible. [36]

An F-test will also be used to test if the MSPE:s of the different models are statistically different from each other. Under the null hypothesis of equal forecasting performance and (H, H) degrees of freedom, the F-statistic is calculated as:[37]

$$F = {\sum_{i=1}^{H} e_{1i}^2}\Big/{\sum_{i=1}^{H} e_{2i}^2} \quad (13)$$

If the null hypothesis can be rejected there is evidence of unequal forecasting performance, and if the null hypothesis cannot be rejected there is no statistical significance in the forecasting behavior of the tested models.

3.2 Data

In order to estimate the models described in 3.1, data over the following three variables is needed: household consumption, the real interest rate and the mortgage rate discount. All data is taken from Statistics Sweden, unless otherwise stated.

3.2.1 Household consumption

Data over the (seasonally cleansed) household consumption is a part of the Swedish national accounting data and are thus reported on a quarterly basis. The data is collected via the tax authorities, and is available from the first quarter in 1980 (even though this paper will only utilize the data from the first quarter of 1996 and onwards). Following the work of Hall, data over consumption will be in (natural) logarithms. This means, given that the independent variables will not be in log form, that the estimates should be interpreted as a semi-elasticity.[38]

The (log of) household consumption in both levels and in first differences can be seen in figure 1:

[36] Enders (2010), p 86.
[37] ibid.
[38] Wooldridge, J.M. (2012). *Introductory Econometrics: A Modern Approach.* 5th ed. Mason: Cengage Learning, p 333.

Figure 1 – Swedish household consumption

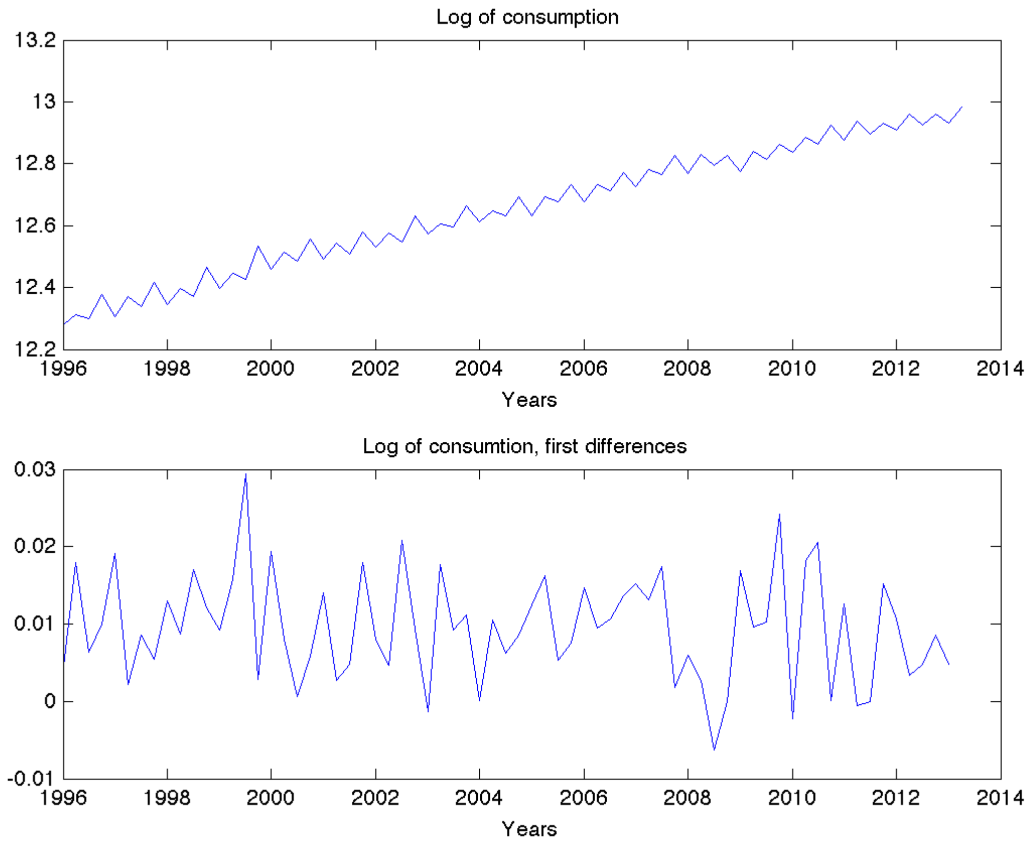

3.2.2 Real interest rate

To compute the real interest rate is a bit more complicated than it first appears, mainly due to two reasons. Partly it is difficult because the real interest rate is depending on expected inflation (which is hard to measure), and partly (mainly) because the real world deviates from economic text books in the sense that there is often (read: always) more than one interest rate in the economy. Both these issues must thus be addressed in order to generate time series over the real interest rate.

The fact that there are many different interest rates in the Swedish economy is hard to do anything about, and will simply have to be mended by this paper experimenting with several, different, interest rates. In order to do so, three different interest rates will be used. This is done mainly as a robustness check, were the results for the different real interest rates ideally will not vary that much.

The first interest rate used will be the Swedish central bank's repo rate, thus representing an interest rate with a short duration (1 week). Following that, this paper will also utilize interest rates from Swedish government bonds with 3 months (medium) and 5 years (long) duration, respectively. Time series over these are available from the data central of the Swedish central bank. As the interest rates are usually reported on a monthly basis, they have been transformed to quarterly data.

Having acquired the time series over the three nominal interest rates, the expected inflation needs to be subtracted, given that the real interest rate (ex ante) is defined as

$$r_t = i_t - \pi^e \tag{14}$$

As aforementioned, data on expected inflation is typically not available. As some kind of proxy is needed, the standard procedure of NIER is followed by letting past inflation act as an approximation of expected inflation. This means, following the reasoning that inflation has typically been rather stable in Sweden during the observational period, that the real interest rate in this paper is calculated (for each of the three nominal interest rates, respectively) as

$$r_t = i_t - \pi_{t-1} \tag{15}$$

Note that time series of inflation is reported on a monthly basis, and has thus been transformed to quarterly data. The data of both the interest rates and the inflation is collected from the first quarter of 1996 and onwards.

The real interest rate variables will be labeled r_1 (based on the repo rate), r_2 (based on the government bond with 3 months duration) and r_3 (based on the government bond with 5 years duration), respectively.

3.2.3 Mortgage rate discount

To calculate the mortgage rate discount is, with regards to data collection, the most difficult task of this paper. The mortgage rate discount is defined as the difference between the official mortgage rates and the mortgage rates that the households actually pay, and time series over these two variables are thus needed. The main problems of this data collection are that, similar to the case of the real interest rate, that there is (i) not just one mortgage rate institute and (ii) not just one mortgage rate duration in the economy.

Starting with the official mortgage rates, the two problems referred to above have been remedied by using the mortgage rate index currently in use by Statistics Sweden's CPI (consumer price index) computations. In the CPI computations, where changes in prices are computed from a "bundle" of goods, the official mortgage rate component represents about 5 %. This index weights the different institutes and the different durations by studying the stock of outstanding mortgage loans.[39]

The weight on the respective mortgage loan institute is their share of all the outstanding loans, and the weight on the respective duration is the share of that duration of the total mortgage loan stock. In the case of the official mortgage rate index the exact calculations were not of essence since Statistics Sweden could provide the index transformed to percentages, but the weighting method is important given that the calculations will have to

[39] Lundin, O. (2008). *Räntekostnader och KPIF.* Örebro: Statistics Sweden.

be done manually in order to generate the actual mortgage rates paid by the households (more on that below). The official mortgage rate index series on a monthly frequency is available from 1987 and onwards, out of which data (transformed to quarterly observations) from 1996 will be used.

To calculate the actual mortgage rates paid by the households is a bit more complicated, given that part of the weighting now needs to be done manually. Part of the weighting is already done in the financial market statistics (of which mortgage loan payments is one part) from Statistics Sweden, in the sense that they have already weighted the different mortgage institutes together.

The troublesome part is then to weigh the different durations of the mortgage rates actually paid together into something that is compatible with the time series of the official rate. The actual mortgage rate payments in the financial market statistic are divided into four durations (<3 months, 3 months – 1 year, 1 – 5 years and > 5 years), whereas the weights provided from the CPI calculations are divided into six different durations (<3 months, 1 year, 2 years, 3 years, 5 years and 8 years). In order to create a time series of the actual payments in accordance with the time series on official mortgage rates, the “matching” is done by weighting according to the following table:

Table 1 – CPI weight matching

CPI weight duration	Actual payment duration
< 3 months	< 3 months
1 year	3 months - 1 year
2, 3, 5 years	1 – 5 years
8 years	> 5 years

When there is more than one weight per mortgage rate duration (i.e. the CPI weigh for 2, 3, and 5 years), an arithmetic average has been applied. Once weighed together, one final time series of the actual mortgage rate payments is acquired. The data of actual payments is available on a quarterly basis from 1996 until 2005, and then on a monthly basis from 2006 till present (meaning some transformations to quarterly data were made).

Having acquired time series over both official mortgage rates and the actual mortgage rates, the mortgage rate discount is simply calculated by subtracting the latter form the former. It should be noted that there, of course, are other ways of weighting the different mortgage rates and the different mortgage rate institutes together, but that the standard set by Statistics Sweden was a logical choice in this case.

3.2.4 Variable correlations

One of the underlying aims of this paper is, as described in the introduction, to examine how the mortgage rate discount correlates with the real interest rate and to examine how it has evolved over time. As can be seen below in figure 2a, the mortgage rate discount seems to vary quite a bit in comparison to the real interest rates, and it is also far from constant over time (rather, it seems to display a slight downward trend). This is also confirmed when looking as the correlation matrix in table 2, where the real interest rates that correlates with the mortgage rate discount the best, $\boldsymbol{r_3}$, have a correlation coefficient of only 0.66 – as compared to the correlations among the real interest rates themselves (which are between 0.87 and 0.99). Rather unsurprising, $\boldsymbol{r_1}$and $\boldsymbol{r_2}$ have the highest correlation of the variables, most likely due to the fact that they are both interest rates with rather short durations.

More surprising is the fact that the mortgage rate discount correlates negatively with household consumption. This appears odd at first – given that an increase in the mortgage rate discount should infer a higher level of consumption – but is most likely explained by the fact that the mortgage rate discount in general has decreased during the period of observation, whereas household consumption has been trending upwards. That the general level of interest rates has been declining during the observation period could be a potential concern, which is elaborated on further in the subsequent section.

Table 2 – Correlation matrix

VARIABLES	$cons$	r_1	r_2	r_3	$disc$
$cons$	1.0000				
r_1	-0.7894	1.0000			
r_2	-0.8178	0.9903	1.0000		
r_3	-0.8938	0.8749	0.8964	1.0000	
$disc$	-0.8095	0.6082	0.6045	0.6653	1.0000

Figure 2a - Real interest rates and the mortgage rate discount (levels)

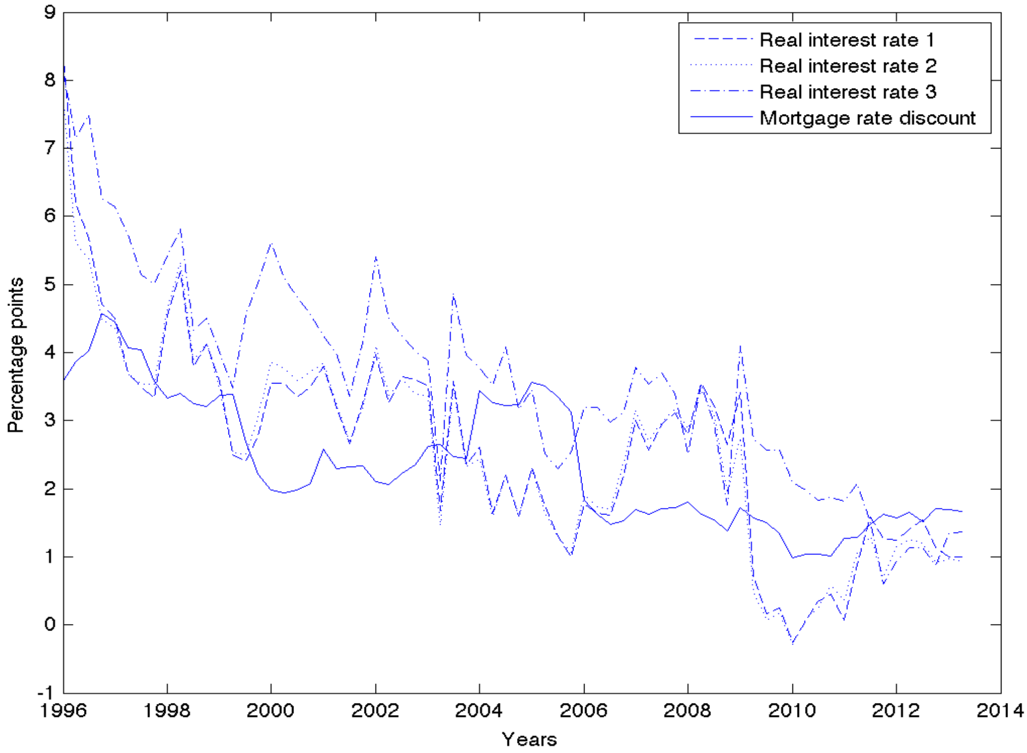

Figure 2b – Real interest rates and the mortgage rate discount (first differences)

4. Results

4.1 Interpretation and internal validity

In table 3 (below), the output from the estimations is presented. As the estimations are carried out in log-linear form, the coefficients should be interpreted as a semi-elasticity. Thus, this means that a one unit (in this case, percentage points) change in the independent variables leads to a percentage change in the dependent variable equal to the estimated coefficient. As a simple example from the output of model (1), a one percentage point increase in the real interest rate implies a 8.3369e-04 percent increase in household consumption (which in turn is measured in millions SEK).

The first things to be noticed in table 3 are the few significant coefficients. This is hardly surprising, and is in line with the current state of literature. It was pointed out already in the first paragraph of this paper, but one should keep in mind that household consumption is hard both to model and forecast.[40] The constant, or the 'drift' part, is on the other hand always positive and significant, which makes intuitive sense: consumption is more or less always increasing.

Out of the real interest rates, r_3 seems to be the real interest rate that fit the data best based on statistical significance (even though the coefficients on r_3 and r_2 in many instances are not significantly different on the 95 % level). That the real interest rate based on the longest duration turns up significant is in many ways intuitively pleasing as changes in longer interest rates are more likely to have an effect on the real economy than fluctuations in the short duration interest rates. It should also be noticed that the other estimated coefficients (the drift term and the coefficients on the mortgage rate discount term) do not vary within the same model specification depending on the choice of real interest rate, which thus indeed serve as a robustness check.

The most problematic segment of the results is the signs on the significant estimates on r_i, which are all positive. This implies that an increase in the real interest rate would lead to a higher level of household consumption, whereas economic theory dictates that an increase in interest rates should deter consumption (as it is then more profitable to save). The explanation to this phenomenon is most likely that both consumption and interest rates with longer durations tend to go up in a booming economy, and thus the first differences correlate in a positive manor (the relationship in levels, on the other hand, is negative, as can be seen in table 2). This explanation, however, raises questions with regards to the model at hand. If the LC-PIH holds, past consumption should be the only control needed – but if that is not the case then it could be that the coefficient on the real interest rate has the "wrong" sign simply because something else would have to be controlled for. This will be elaborated on further in 4.3.

[40] Stephens (2008), p 241-252.

Table 3 – Estimation output

VARIABLES	(1) $\Delta cons$	(2) $\Delta cons$	(3) $\Delta cons$	(4) $\Delta cons$	(5) $\Delta cons$	(6) $\Delta cons$	(7) $\Delta cons$	(8) $\Delta cons$	(9) $\Delta cons$
Constant	0.0096***	0.0097***	0.0097***	0.0094***	0.0095***	0.0096***			
	(0. 0008)	(8.376e-04)	(8.422e-04)	(8.198e-04)	(8.055e-04)	(8.110e-04)			
$\Delta r_{1,t}$	8.3369e-04			7.2168e-04					
	(0. .0011)			(0.0010)					
$\Delta r_{2,t}$		0.0017			0.0014				
		(0.0011)			(0.0011)				
$\Delta r_{3,t}$			0.0020*			0.0017*			
			(0.0012)			(0.0010)			
$\Delta disc_t$				-0.0056*	-0.0053*	-0.0049			
				(0.0033)	(0.0034)	(0.0034)			
Constant ($\Delta disc_t$< -0.3)							0.0339***	0.0314***	0.0283***
							(0.0111)	(0.0113)	(0.0060)
Constant ($\Delta disc_t$≥ -0.3)							0.0090***	0.0091***	0.0091***
							(8.259e-04)	(8.061e-04)	(8.103e-04)
$\Delta r_{1,t}$ ($\Delta disc_t$< -0.3)							0.0089		
							(0.0077)		
$\Delta r_{1,t}$ ($\Delta disc_t$≥ -0.3)							0.0010		
							(0.0010)		
$\Delta r_{2,t}$ ($\Delta disc_t$< -0.3)								0.0071	
								(0.0086)	
$\Delta r_{2,t}$ ($\Delta disc_t$≥ -0.3)								0.0017	
								(0.0012)	
$\Delta r_{3,t}$ ($\Delta disc_t$< -0.3)									0.0040
									(0.0061)
$\Delta r_{3,t}$ ($\Delta disc_t$≥ -0.3)									0.0018*
									(0.0011)
$\Delta disc_t$ ($\Delta disc_t$< -0.3)							0.0261**	0.0234*	0.0191*
							(0.0129)	(0.0141)	(0.0083)
$\Delta disc_t$ ($\Delta disc_t$≥ -0.3)							-0.0041	-0.0040	-0.0035
							(0.0034)	(0.0034)	(0.0036)
AIC	-7.0048	-7.0112	-7.0167	-5.0204	-5.0255	-4.9977	0.9342	0.9289	0.9284
BIC	-2.5366	-2.5430	-2.5485	1.6819	1.6768	1.7046	14.3389	14.3335	14.3653
MSPE	0.0082	0.0082	0.0079	0.0086	0.0085	0.0082	0.0123	0.0117	0.0104
Observations	69	69	69	69	69	69	69	69	69
Adj-R^2	-0.0219	0.0012	0.0206	0.0186	0.0365	0.0470	0.1270	0.1440	0.1457

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The reasoning above can be applied to the estimated effect of the mortgage rate discount in the linear model as well. The coefficient on the mortgage rate discount is always, when significant and thus different from zero, negative. Just like the positive sign on the estimates of the real interest rate, this makes no intuitive sense at first: why should an increase in the mortgage rate discount lead to a lower consumption? The answer to this question could be either that (i) the model needs more controls or (ii) an increase in the mortgage rate discount inclines us to save more, for instance through amortizations on housing property, instead of consuming. To say which is impossible without more data, but following the reasoning regarding r_3, it could be that the mortgage rate discount goes down when the economy is booming, and that household consumption would be negatively associated with an mortgage rate discount increase (even though the consumption is mainly driven by other factors).

Interestingly enough, the non-linear estimation shed some light on this matter. In model (7), (8) and (9), given the lower threshold (i.e. $\Delta disc_t < -0.3$), the effect on household consumption is indeed negative when the household discount is decreasing (due to the positive coefficient). At the same time, the coefficients in the upper threshold are now insignificant instead, implying that the mortgage rate discount does not affect household consumption when the mortgage rate discount difference is larger than −0.3. Not too much faith should be put into the non-linear model given the unequal distribution of observations in the two regimes, but the F-test conducted to test if the non-linear parameters are statistically different compared to the parameters of the baseline and linear model all turn out significant (see Appendix B for details).

It could of course be the case that these drops in the mortgage rate discount just happen to correlate with something else that drives consumption down – but at least it points towards the important finding that the non-linear approach seem to better capture the effect of the mortgage rate discount as compared to the linear model, and thus suggests that a decrease of the mortgage rate discount smaller than −0.3 reduces consumption whereas increases of the mortgage rate discount larger than −0.3 has no effect at all.

Furthermore, it is also interesting to note that the few observations in the lower regime are not occurring during the worst of the economic crises, given that you would potentially expect big drops in the mortgage rate discount e.g. when the dot-com bubble burst in 2000-2001 and during the Lehman Brothers crisis in 2008. It is of course possible that events which seemingly had smaller effect on the Swedish economy, such as the Russia crisis 1997-1998 and the mortgage crisis in the Netherlands 2006, might have effected the mortgage rate discount – but no obvious pattern seem to appear (see graph B.1 in Appendix B).

4.2 Implications

This paper is by its very nature not relevant per se for policy, but rather relevant for forecasting. An improvement in forecasting behavior (i.e. a reduction in the MSPE, the forecast error) of household consumption could mean a better forecast of GDP, and thus – as mentioned in the first paragraph of this paper – that the budgetary process in (mainly) the public sector is facilitated.

As observed by the MSPE output in table 3, no improvements are made with regards to forecasting behavior when comparing with the baseline model. The F-tests of the MSPE suggest that there is no statistical difference between the MSPE in the baseline models and the linear/non-linear models (see Appendix B.2), but, either way, no forecasting improvements are made. These results are also robust in the sense that changing the number of forecasted periods does not seem to alter the result. The fact that no forecasting improvement is made is also confirmed by the Akaike and Bayesian information criteria, which is lowest for the baseline model (the likelihood value is roughly the same for all models, and the AIC and BIC then punishes the less parsimonious models). Interestingly enough the adjusted R^2 (the explained variation within the model) improves in both the linear and the non-linear model, but then again adjusted R^2 is rather low overall and less important than forecasting properties.

One important aspect of econometric modeling, which is often overlooked in the economic literature, is the size of the significant coefficients. If the sizes of the coefficients are not put in relation to the unit of measure of the dependent variable, it is very hard to conceive whether a significant coefficient in reality has an effect. In this case, the coefficients on $\Delta disc_t$ ($\Delta disc_t < -0.3$) imply that Swedish household consumption is reduced by roughly 2 % if the mortgage rate discount is reduced from e.g. −0.3 to −1.3. Two percent of Swedish household consumption is as of 2012 just below one percent of Swedish GDP, which is definitely an effect that would gain the interest of policy makers. It also means – which is from a policy perspective very important – that household consumption can drop quite severely without any chance of fiscal policy interfering. The mortgage rate discount is most likely by and large hard to steer, even for monetary policy.

As for implications with regards to external consistency, the findings in this paper link in rather well with the current state of knowledge. It shows that much of the somewhat old literature – claiming that household consumption is hard to forecast – is still valid, and it also confirms that the same findings apply in Sweden (for the most part of the current state of literature, USA data has been used). The findings also partly fall in line with the 1988 Hall paper, in the sense that the augmentations of the baseline model did not lead to any forecasting improvements. They do, however, contradict Hall's result since the coefficients related to the real interest rate have a positive sign. Subsequently, this paper conforms to the current literature seeing that it brings the validity of the LC-PIH hypothesis to the spotlight. As elaborated on below, a model is only as strong as the assumptions it is relying upon.

What the paper really adds to the current state of knowledge is that the mortgage rate discount now is not only a vague term used by media, but also something on which there is available data. It is also new information that the non-linear model suggests a link between the mortgage rate discount and household consumption, but only during one of the two regimes. Using this, it might be possible to improve forecasting capabilities in the future by including only the lower threshold.

4.3 Limitations

The results presented above are in themselves very interesting, but a few limitations need to be brought up. First of all – just as pointed out previously – the assumptions postulated by the LC-PIH will of course limit the conclusions that can be made. If the LC-PIH holds the results are very interesting in a number of ways, not only because of the mortgage rate discount but also with regards to the fact that higher real interest rates did not infer a reduction in household consumption. If instead the LC-PIH assumptions are wrong, the model is misspecified and the estimates are most likely biased. If past consumption is not the only covariate needed to find the causal effect of the real interest rate and the mortgage rate discount, the results cannot be trusted. Even though adding more controls has been tried with only modest success in the past, it would of course be interesting to see how the results would alter given that you could control for income, wealth, inflation, economic outlook and so on.

Quite naturally, the number of observations and the frequency must also be brought up as a possible limitation. The time series are made up of 70 quarterly observations, and more detailed data would of course have been an asset. The problem of making the data set longer stems from two issues, which are both hard to address: (i) that household consumption is only measured on a quarterly basis and (ii) that interest rate data from before 1993 is very hard to use given the paradigm change in monetary policy. Ideally Statistics Sweden would start to gather data on the household consumption on a monthly basis, which could potentially increase our knowledge of how well variation in the real interest rate and the mortgage rate discount affect household consumption. In addition to this, one must also mention the fact that the whole model is in first differences. Given the non-stationary of the level data there are few other options, but having the model in first differences will reduce the variation in the data material.

Apart from the fact that the frequency and length of the time series ideally could be better, the actuality that all the data are aggregates for whole of Sweden also potentially limits the results. The housing market in Sweden is most likely not a homogenous market and how the mortgage rate discount affects consumption might then have different effects in different regions. In some regions where the population do not have large scale mortgage loans a reduction in the mortgage rate discount most likely affects household consumption very little, whereas an increase in the mortgage rate discount hypothetically could matter more in regions where many have larger mortgage loans and thus eventually a higher marginal propensity to consume. These suggestions are, again, speculations, but having this data on a regional or municipal level could improve precision of the estimates.

Lastly, it should be pointed out that the period used for the in-sample forecasting might not be considered “typical”. The period 2008-2013 has been subject to a significant degree of economic turmoil, and it is possible that the model at hand would perform better under more “normal” economic circumstances. On the other hand changing the forecast periods did not alter the results in terms of forecasting performance, and the claim that forecasting behavior could be better during more economic circumstances will have to be regarded as speculative.

5. Conclusion

A first and rather obvious conclusion from this paper is that household consumption is indeed hard to forecast. Even with the baseline model, which has been tried and tested before, the results are far from intuitive. This in many ways confirms the current state of literature, and implies that the end of this research area is not yet in sight. The results also beg to question the validity of the LC-PIH; another discussion that is most likely far from at an end.

The main contribution of the results is that the mortgage rate discount did not have an effect on consumption in the linear case, but that a negative effect of a reduction in the mortgage rate discount could be discovered in the non-linear case. Given the chosen threshold of −0.3, the households did indeed consume less when the mortgage rate discount decreased. In particular, it is interesting that the effect was quite large and that it is a variable that would be very hard to control from a policy point of view. Attention should of course be paid to the number of observations in the lower threshold, but statistical significance was attained and this particular effect is something that would be very interesting to study in a few years time when more observations are available.

The results of this particular paper, limited as they might be, raise several interesting questions that should spawn further research. First of all, it would be interesting to see whether a more data driven approach to the modeling of consumption depending on the real interest rate and the mortgage rate discount would be superior to the more theory based approach of this paper. If the LC-PIH does not hold, a more data driven approach to the time series modeling might lead to improvements.

Further research in many other directions would also be advisable, given that a more detailed data set becomes feasible. For instance, how much does the mortgage rate discount matter for different income groups? It is hardly a secret that income groups with lower incomes have a higher marginal propensity to consume (and vice versa), and having them separated in the data sample could give a more precise estimate of the true causal effect. In line with that reasoning, it would also be interesting to study this phenomenon in different Swedish regions, given that the housing market is most likely far from homogenous. In addition, accepting the premise that an increase in the mortgage rate discount does not lead to an increase in consumption, it would of course be interesting to study where the money goes instead (for instance savings, amortizations, spending abroad or into the mattress).

Another possible extension would be to model the mortgage rate discount as a relative variable, rather than as an absolute value (as in this paper). The last two decades in Sweden have seen an overall decrease in the general interest rate level, which implies that the amount to which you can "haggle" with regards to your mortgage rate most likely has been reduced. In that aspect, a model with the mortgage rate discount in relationship to the level of the overall mortgage rates might be more suitable – especially if the overall interest rate level keeps reducing.

Furthermore, given that the non-linear modeling points towards that the mortgage rate discount has no effect on household consumption when it is increasing and a negative effect when it is decreasing by at least 0.3 percentage points, the driving forces of the mortgage rate discount must be understood. If the mortgage rate discount potentially can have negative effects on household consumption, finding out the determinants of the discount is the natural next step. For instance, the mortgage rate discount in part depends on the official mortgage rate set by the banks and mortgage institutes (which should reflect their costs), and a throughout investigation of the pass-through from monetary policy to mortgage rates might thus be a good stepping stone given that different levels of the official rates might make the mortgage rate institute more or less inclined to give discounts.

Even more importantly, the actual process of acquiring the mortgage rate discount must be understood. Typically the negotiations about the mortgage rate discount are conducted after already having purchased the property, but little is known about the process itself. There is, for instance, no publically available data on how common it is to get a discount, which would of course be important knowledge when it comes to assessing the importance of the mortgage rate discount. Moreover, it could be that banks and mortgage rate institutes are more likely to give a big mortgage rate discount to “loyal costumers”, meaning that there potentially is a substitution effect between the discount and e.g. (retirement) savings. These suggestions are, again, speculative, but a thorough understanding of that process is most likely necessary if the effect of the mortgage rate discount is to ever be fully disentangled.

All in all, it is clear that this paper will not provide the definite answer to all of these questions, but rather mark the beginning of an interesting area of research. Therefore, this work will hopefully spawn subsequent literature and further investigations in this field, as precise forecasting of household consumption is of utmost important for government planning and fiscal policy.

6. References

Campbell, J.Y. & Deaton, A. (1989). *Why is consumption so smooth?* Review Of Economic Studies 56(3): 357-373.

Carruth, A. & Henley, A. (1990). *Can existing consumption functions forecast consumer spending in the late 1980's?* Oxford Bullentin Of Economics And Statistics 52(2): 211-222.

Chan, K.S. (1993). *Consistency and Limiting Distribution of the Least Squares Estimator of a Threshold Autoregressive Model.* The Annals of Statistics 21(1): 520-533.

Davis, M.A. & Palumbo, M.G. (2001). *A primer on the economics and time series econometrics of wealth effects*. Divisions of Research & Statistics and Monetary Affairs: Federal Reserve Board.

Dion, D.P. (2006). *Does Consumer Confidence Forecast Household Spending? The Euro area case.* MPRA paper 911: 1-24.

Eitrheim, O., Eilev S.J. & Ragnar, N. (2002). *Progress from Forecast Failure--The Norwegian Consumption Function*. Econometrics Journal 5(1): 40-64.

Enders, W. (2010). *Applied Econometric Time Series.* 3rd ed. Danvers: John Wiley & Sons, Inc.

Enders, W. & Granger, C.W.J. (1998). *Unit-Root Tests with an Assymetric Adjustment With an Example Using the Term Structure of Interest Rates.* Journal of Business & Economic Statistics 16(3): 304-311.

Ekonomistyrningsverket (2013). *Statens budget och de offentliga finanserna.* Stockholm: Publikationsservice.

Finansinspektionen (2013). *Den svenska bolånemarknaden 2013.* Stockholm: Publikationsservice.

Flavin, M. (1993). *The Excess Smoothness of Consumption: Identification and Interpretation.* Review Of Economic Studies 60(3): 651-666.

Franses, P.H. & van Dijk, D. (2008). *Non-linear time series models in empirical finance.* 6th ed., Cambridge: University Press.

Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton University Press (for the National Bureau of Economic Research).

Gujarati, D.N. & Porter, D.C. (2009). *Basic Econometrics.* 5th ed. New York: McGraw-Hill.

Hall, R.E. (1978). *Stochastic implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence.* Journal of Political Economy 86(6): p 971-987.

Hall, R.E. (1988). *Intertemporal Substitution in Consumption.* Journal of Political Economy 96(2): 339-357.

Hansen, B.E. (1997). *Inference in TAR models.* Studies in Nonlinear Dynamics and Econometrics 2(1): 1-14.

Juster, F.T. (1960). *Prediction and consumer buying intentions*. American Economic Review 50(1): 604-626.

Keynes, J.M. (1936). *General Theory of Employment, Interest and Money*. London: Macmillan.

Kuznets, S. (1942). *Uses of National Income in Peace and War.* New York: National Bureau of Economic Research.

Lundin, O. (2008). *Räntekostnader och KPIF. Örebro:* Statistics Sweden.

Ramanathan, R. (1993). *Statistical Methods in Econometrics.* Bingley: Emerald Group Publishing.

Stephens, M. (2008). *The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans*. Review Of Economics And Statistics 90(2): 241-252.

Vadas, G. (2005*). Beyond macro variables: consumer confidence index and household expenditure in Hungary*. Microeconomics 0512006: 1-16.

Wilcox, J.A. (2007). *Forecasting Components of Consumption with Components of Consumer Sentiment*. Business Economics 42(4): 36-46.

White, H. (1982). *Maximum Likelihood Estimation of Misspecified Models*. Econometrica 50(1): 1-25.

Wooldridge, J.M. (2012). *Introductory Econometrics: A Modern Approach.* 5th ed. Mason: Cengage Learning.

Zachrison, O. (2012, April 23rd). *Unik räntekarta med din hjälp.* Svenska Dagbladet. Retrieved online from http://www.svd.se

Appendix A

A.1 Maximum Likelihood Estimation

In accordance with e.g. *Basic Econometrics*, the method of estimating via maximum likelihood is outlined as follows:[41]

Use a general specification of an econometric (time series) model, such as

$$y_t = f(x_t; \theta) + \varepsilon_t \tag{A1}$$

where ε_t is assumed to be $iid \sim N(0, \sigma_\varepsilon^2)$. If ε_t is *iid*, then so must

$$y_t - f(x_t; \theta) \tag{A2}$$

The density function, which can also be thought of as the likelihood function, of the model is written as

$$D(\Omega) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} * \exp\left(-\frac{1}{2}\varepsilon_t^2 \Big/ \sigma_\varepsilon^2\right) \tag{A3}$$

where $\Omega = \theta, \sigma_\varepsilon^2$ with θ being the vector of parameters to be estimated and σ_ε^2 being the variance of the error term. Taking the natural logarithm of this function, in order to make estimation easier, and then using that $y_t - f(x_t; \theta) = \varepsilon_t$ yields the log-likelihood function

$$L(\Omega) = \ln\left(\frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}}\right) - \frac{1}{2}\frac{(y_t - f(x_t; \theta))^2}{\sigma_\varepsilon^2} \tag{A4}$$

or

$$L(\Omega) = -\frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln(\sigma_\varepsilon^2) - \frac{1}{2}\frac{(y_t - f(x_t; \theta))^2}{\sigma_\varepsilon^2} \tag{A5}$$

which is the expression to be maximized - with regards to Ω - in MATLAB. To obtain robust standard errors, with $\widehat{\Omega}$ being the quasi-maximum likelihood estimator (QMLE) of Ω and $H_t(\widehat{\Omega})$ being the Hessian matrix, the square root is taken from the diagonal of the Huber Sandwich estimator (for further details, see e.g. *Statistical Methods in Econometrics* by Ramanathan).[42]

$$\hat{V}_{QMLE} = \left(\sum_{t=1}^{T} H_t(\widehat{\Omega})\right)^{-1} \sum_{t=1}^{T} S_t(\widehat{\Omega}) S_t(\widehat{\Omega})' \left(\sum_{t=1}^{T} H_t(\widehat{\Omega})\right)^{-1} \tag{A6}$$

[41] Gujarati, Damodar N. & Porter, Dawn C. (2009). *Basic Econometrics*. 5th ed. New York: McGraw-Hill, p 143ff.
[42] Ramanathan, Ramu (1993). *Statistical Methods in Econometrics.* Bingley: Emerald Group Publishing, p 179.

A.2 TAR estimation in MATLAB

```
%%%%%%%%%%%%%%%%%%% Example code non-linear estimation %%%%%%%%%%%%%%%%%%%

%% Estimation
% Estimating parameters

% Inital guess a10, a11, a12, a20, a21, a22, sigma2
v=[0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1] ;

global c
c=-0.3 ;
k=6 % # parameters

% Maximizing the object function
options=optimset('hessian', 'on','TolX', 1e-10) ;

[a,FVAL,EXITFLAG,OUTPUT,GRAD,HESSIAN]=fminunc(@likelihoodtar, v)  ;

% Displaying the point estimates

a10hat=a(1)
a11hat=a(2)
a12hat=a(3)
a20hat=a(4)
a21hat=a(5)
a22hat=a(6)
sigma2hat=a(7)

% Estimating robust standard deviations

H=HESSIAN ;

global n

G=zeros(7,7) ;

for n=1:69

    options=optimset('hessian', 'on', 'TolX', 1e-10) ;

   [a,FVAL,EXITFLAG,OUTPUT,GRAD,HESSIAN]=fminunc(@likelihoodtarrobust, v)

    g=GRAD*GRAD';
    G=G+g ;
end


ROBUSTH=(H^-1)*(G)*(H^-1); % The Huber-Sandwich estimator
ROBUSTSTD=sqrt(ROBUSTH);

stda10hat=ROBUSTSTD(1,1)
stda11hat=ROBUSTSTD(2,2)
stda12hat=ROBUSTSTD(3,3)
stda20hat=ROBUSTSTD(4,4)
stda21hat=ROBUSTSTD(5,5)
stda22hat=ROBUSTSTD(6,6)
stdsigma2hat=ROBUSTSTD(7,7)

%% Evaluation

% Computing t-statistics

tstata10 = a10hat / stda10hat
```

```
tstata11 = a11hat / stda11hat
tstata12 = a12hat / stda12hat
tstata20 = a20hat / stda20hat
tstata21 = a21hat / stda21hat
tstata22 = a22hat / stda22hat

% Computing AKAIKE and BIC information criteria

AIC=2*k - 2*log(-FVAL)
BIC=-2*log(-FVAL)+k*log(n)

% Computing Adj-R^2

DATA = importdata('data.xls') ;
cons=DATA(:,5) ;
r=DATA(:,2) ;
for t=1:69 ;
   if disc(t)<c

    conshat(t)= a10hat+a11hat*r(t)+a12hat*disc(t) ;

   else

    conshat(t)= a20hat+a21hat*r(t)+a22hat*disc(t) ;

   end
end

SSE=(cons - conshat').^2 ;
SSE=sum(SSE) ;

consbar= sum(cons) / numel(cons) ;

SST=(cons - consbar).^2 ;
SST=sum(SST) ;

Rsq= 1 - (SSE / SST) ;

AdjRsq= 1 - (((1-Rsq)*(numel(cons)-1))/(numel(cons)-k-1))

% Computing MSPE (1-step ahead forecast)

global i

v=[0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1] ;
prederr=zeros(69,1) ;
predcons=zeros(69,1) ;
j=1 ;

for i=52:68

[a]=fminunc(@likelihoodtarmspe, v)  ;

a10hatt=a(1) ;
a11hatt=a(2);
a12hatt=a(3) ;
a20hatt=a(4) ;
a21hatt=a(5);
a22hatt=a(6) ;

if disc(i)< c

predcons(i+1,1)=a10hatt+a11hatt*r(i+1)+a12hatt*disc(i+1) ;
prederr(i+1,1)= cons(i+1)-predcons(i+1) ;
```

```
else

predcons(i+1,1)=a20hatt+a21hatt*r(i+1)+a22hatt*disc(i+1) ;
prederr(i+1,1)= cons(i+1)-predcons(i+1) ;

end
end

prederr=prederr.^2 ;
MSPE=sqrt(sum(prederr)/nnz(prederr))
```

```
function [ l ] = likelihoodtar(v)

global c

% Specifying the location of the intial guesses

a10 = v(1);
a20 = v(4);
a11 = v(2);
a21 = v(5);
a12 = v(3);
a22 = v(6);
sigma2=v(7);

% importing the data series
DATA = importdata('data.xls') ;
cons=DATA(:,5) ;
r=DATA(:,2) ;
disc=DATA(:,1) ;

% the likelihood function

for t=1:69 ;

    if disc(t)< c


likelow(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma21))-...
    ((1/2)*(((cons(t)-a10-a11*r(t)-a12*disc(t)).^2)./sigma21)));

    else

likehigh(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma22))-...
    ((1/2)*(((cons(t)-a20-a21*r(t)-a22*disc(t)).^2)./sigma22)));



end
end

% the object function

l=sum(likelow) + sum(likehigh);

end
```

```
function [ x ] = likelihoodtarmspe(v)

global i
global c
```

```
% Specifying the location of the intial guesses

a10 = v(1);
a11 = v(2);
a12 = v(3);
a20 = v(4);
a21 = v(5);
a22 = v(6);
sigma2=v(7);

% importing the data series
DATA = importdata('data.xls') ;
cons=DATA(:,5) ;
r=DATA(:,2) ;
disc=DATA(:,1) ;

% the likelihood function
for t=1:i ;

    if disc(t)< c

likelow(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma2))-...
    ((1/2)*(((cons(t)-a10-a11*r(t)-a12*disc(t)).^2)./sigma2)));

    else

likehigh(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma2))-...
    ((1/2)*(((cons(t)-a20-a21*r(t)-a22*disc(t)).^2)./sigma2)));

end
end

% the object function

x=sum(likelow) + sum(likehigh);

end
```

```
function [ l ] = likelihoodtarrobust(v)

global c
global n

% Specifying the location of the intial guesses

a10 = v(1);
a11 = v(2);
a12 = v(3);
a20 = v(4);
a21 = v(5);
a22 = v(6);
sigma2=v(7);

% importing the data series
DATA = importdata('data.xls') ;
cons=DATA(:,5) ;
r=DATA(:,2) ;
disc=DATA(:,1) ;

% the likelihood function
```

```
for t=n ;

    if disc(t)< c


likelow(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma2))-...
    ((1/2)*(((cons(t)-a10-a11*r(t)-a12*disc(t)).^2)./sigma2)));

    else

likehigh(t)= -((-(1/2)*log(2*pi))-((1/2)*log(sigma2))-...
    ((1/2)*(((cons(t)-a20-a21*r(t)-a22*disc(t)).^2)./sigma2)));



end
end

% the object function

l=sum(likelow) + sum(likehigh);
end
```

Appendix B

B.1 Pre-testing the data

Testing for stationarity is done in accordance with e.g. Enders, and contains the following three steps:[43]

1. Find, using the Akaike and Bayesian information criteria, the best fitting autoregressive model (i.e. determine the number of lags) for the time series at hand.
2. Test the residuals of the autoregressive process chosen in step 1 for white noise, using the Ljung-Box test. The LB-test tests the null hypothesis of no autocorrelation (i.e. the process is white noise), against the alternative hypothesis of autocorrelation (i.e. the process is not white noise).
3. Having established that the residuals are white noise, the Augmented Dickey-Fuller test for the presence of a unit root. The presence of a unit root (the null) implies a stationary process in first differences, whereas the non-presence of a unit root (the alternative hypothesis) implies a stationary process in levels.

The results of the steps above are presented in the table below:

Table B.1 - LB and ADF tests

VARIABLES	Lags	Observations	LB test statistic	LB p-value	ADF test statistic	ADF critical value (5 %)
$cons$	0	70	42.5275	0.1010	-1.998	-3.45
r_1	0	70	4.4193	0.4908	-1.778	-3.45
r_2	0	70	1.4482	0.9840	-1.205	-3.45
r_3	0	70	1.8509	00.9676	-0.732	-3.45
$disc$	0	70	1.8359	0.1754	-1.916	-3.45
$disc\ if\ \Delta disc \geq -0.3$	0	64	42.5504	0.1200	-0.472	-3.45
$disc\ if\ \Delta disc < -0.3$	0	6	0.0026	0.9591	-0.054	-3.60

*** p<0.01, ** p<0.05, * p<0.1

[43] Enders (2010), p 215-219.

B.2 Thresholds and non-linearity testing

B.2.1 Finding ‘c’

Table B.2 – Sum of squared residuals (TAR models) for different thresholds

	(7)	(8)	(9)
c	SSE	SSE	SSE
-1.2	0.003	0.0029	0.0029
-1.1	0.003	0.0029	0.0029
-1	0.003	0.0029	0.0029
-0.9	0.003	0.0029	0.0029
-0.8	0.003	0.0029	0.0029
-0.7	0.003	0.0029	0.0029
-0.6	0.003	0.0029	0.0029
-0.5	0.003	0.0029	0.0029
-0.4	0.0028	0.0028	0.0028
-0.3	**0.0026**	**0.0026**	**0.0026**
-0.2	0.0029	0.0028	0.0029
-0.1	0.0031	0.0030	0.0030
0	0.0029	0.0029	0.0029
0.1	0.003	0.0030	0.0029
0.2	0.0029	0.0028	0.0027
0.3	0.003	0.0029	0.0028
0.4	0.003	0.0030	0.0029
0.5	0.003	0.0030	0.0029
0.6	0.003	0.0030	0.0030
0.7	0.003	0.0030	0.0030
0.8	0.003	0.0030	0.0030
0.9	0.003	0.0030	0.0030

Lowest SSE in bold

B.2.2 Non-linearity

Figure B.1 – Regime graph

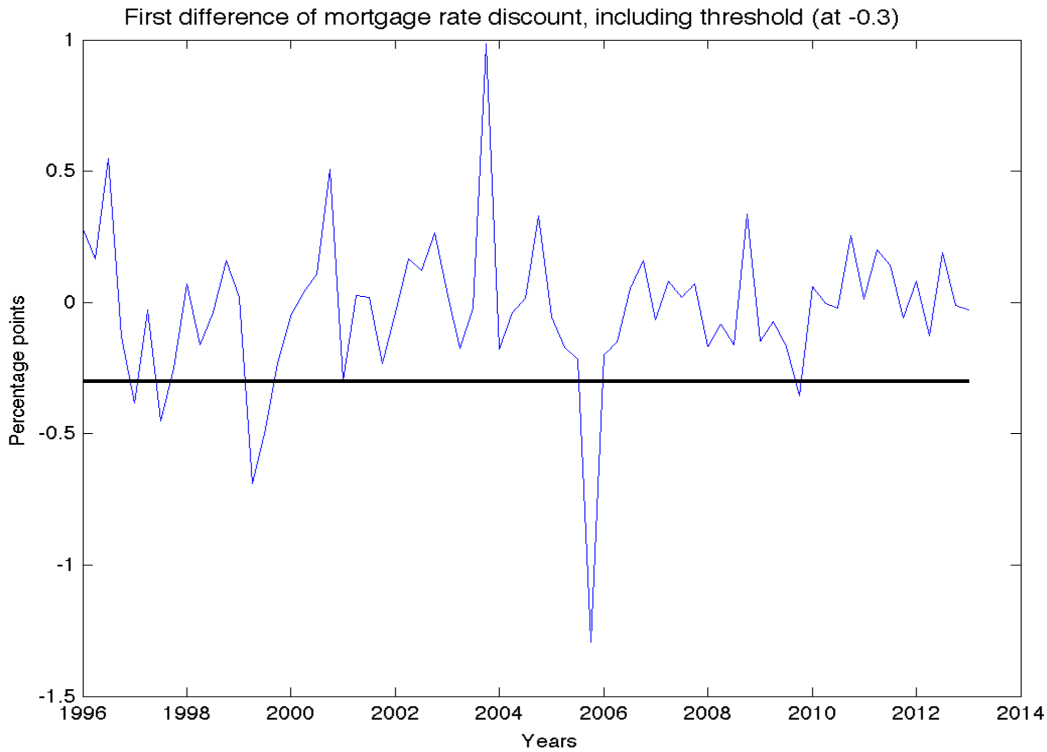

Table B.3 - F-tests

MODEL	Observations	$F_{crit}(5\ \%)$	$F_{crit}(1\ \%)$	F_{test}
Parameters in (1) against (7)	69	2.50	3.60	4.50961***
Parameters in (2) against (8)	69	2.50	3.60	3.86538***
Parameters in (3) against (9)	69	2.50	3.60	3.86538***
Parameters in (4) against (7)	69	2.751	4.109	4.23076***
Parameters in (5) against (8)	69	2.751	4.109	4.23076***
Parameters in (6) against (9)	69	2.751	4.109	3.38461**
MSPE in (1) against (7)	17	2.27	3.24	0.46153
MSPE in (2) against (8)	17	2.27	3.24	0.47826
MSPE in (3) against (9)	17	2.27	3.24	0.61111

*** p<0.01, ** p<0.05, * p<0.1