Erasmus University Rotterdam
Erasmus School of Economics

Bachelor Thesis Econometrics and Operations Research

The Effect of Uncertainty on Macroeconomic Activity

Richard Hagen
415308

Supervised by
B. Keijsers
Second assessor
X. Gong

Date Final Version:
July 2, 2017

Contents

Abstract

I replicate and extend the research documented in Jones and Enders (2016) and manage to come to similar conclusions. I employ a nonlinear framework, the logistic smooth transition autoregressive (LSTAR) model, in combination with conditional volatility estimated using an EGARCH model, realised volatility and conditional volatility estimated using a MIDAS model as respective uncertainty measures. I find that uncertainty indeed has a nonlinear effect on macroeconomic activity, in this paper proxied by industrial production, and this effect also turns out to be asymmetric in the sense that positive and negative shocks of the same size have different sized effects on industrial production.

Keywords: Nonlinear Models, LSTAR Models, Uncertainty, Uncertainty Shocks, Macroeconomic Activity, Volatility

1 Introduction

Ever since the global financial crisis in 2007-2009 there has been a significant surge in research in uncertainty. According to Bloom (2014) this surge in research interest was mainly caused by the jump in uncertainty in 2008 that helped shape the financial crisis. Others have also linked the increasing uncertainty to the crisis. The Federal Open Market Committee for example has repeatedly named uncertainty a key factor driving the recession, and Stock (2012) stated that "The main contributions to the decline in output and employment during the 2007-2009 recession are estimated to come from financial and uncertainty shocks." Bloom (2009) and Bloom et al. (2012) showed that uncertainty can lead to large drops in economic activity, and can therefore contribute to a financial crisis. Finally, Arellano et al. (2012) found that uncertainty shocks can explain a substantial fraction of fluctuations in output. All these articles agree that uncertainty has an effect on the macroeconomy during recessions and during expansions. The main reason these articles gave to explain this effect of uncertainty is that both consumers and companies base many of their economic decisions on the degree of uncertainty they are experiencing, since they usually postpone investment, hiring and consumption decisions when uncertainty is high until business conditions become clearer (Bloom, 2014).

Mishkin (2011) was the first to suggest that the effect of uncertainty on the macroeconomy may not be linear. Besides this nonlinearity, Grier et al. (2004) and Foerster (2014) provided evidence that this effect is also asymmetric. Jones and Enders (2016) investigated these findings further by estimating several macroeconomic variables as logistic smooth transition autoregressive (LSTAR) processes with uncertainty as a transition variable. They aimed to find out whether uncertainty affects the macroeconomy in a nonlinear and asymmetric fashion. They found that for several key-macroeconomic variables positive uncertainty shocks have a larger effect on the macroeconomy than negative shocks, and that the effect of uncertainty shocks is highly dependent on the state of the economy. They concluded that linear models highly underestimate the consequences of uncertainty during financial crises.

Throughout this report I investigate whether the conclusions drawn in Jones and Enders (2016) hold up by replicating their main results using the exact same methods as they did and data similar to theirs. These main results are based on an LSTAR model for industrial production, with EGARCH conditional volatility of the monthly financial returns on the S&P 500 Index as uncertainty measure. However, the financial returns on an index such as the S&P 500 are available at a much higher frequency, which sparks the idea for my extension of Jones and Enders (2016) in which I use daily financial returns to estimate an uncertainty measure. I use these daily returns to estimate monthly realised volatility, which according to Jurado et al. (2015) is a well-working proxy for uncertainty. Another method I employ to estimate an uncertainty measure using daily returns is called a MIxed Data Sampling (MIDAS) regression, which was first introduced by Ghysels et al. (2005). They created a model which linked daily financial returns to monthly conditional volatility, which is also able to capture asymmetries in the dynamics of conditional variance and can incorporate nonlinearities, just like the EGARCH model can. This MIDAS model allows me to use data collected at a higher frequency than the data used in Jones and Enders (2016), which I anticipate will result in a higher overall accuracy, and possibly more nuanced conclusions. On top of this expected higher accuracy, Ghysels et al. (2007) showed that MIDAS regressions are particularly good at forecasting conditional volatility during periods of high volatility, even more so than most GARCH models. This all shows that using a MIDAS regression to estimate an uncertainty measure may well prove to be an improvement over the uncertainty measures applied in Jones and Enders (2016). Through the use of these methods I aim to find out whether positive and negative uncertainty shocks have asymmetric effects, and whether the effects of uncertainty shocks vary over the business cycle by estimating a nonlinear model - the logistic smooth transition autoregressive (LSTAR) model - for industrial production where I use the earlier mentioned uncertainty measures: EGARCH conditional volatility, realised volatility and MIDAS conditional volatility. For each of the estimated LSTAR models, I create impulse response functions to investigate the response of industrial production to uncertainty shocks of different signs and sizes. These impulse responses show that positive uncertainty shocks have a greater effect on industrial production than negative uncertainty shocks do. They also show that the timing of uncertainty shocks greatly affects the change in industrial production since uncertainty shocks occurring during the crisis have a far greater effect than the same size shocks occurring outside the financial crisis. These findings for the three different LSTAR Models are very similar to those in Jones and Enders (2016), which shows that their research is both replicable and robust.

In the next section I present the data I use to find my results, and in section 3 I explain the methods I use to find these results. Section 4 presents the estimated LSTAR models and the impulse responses created using these models. In section 5 I compare and discuss the different uncertainty measures, and I compare the impulse responses created using these models. In section 6 I present my conclusions, and in section 7 I present some suggestions for further research into this subject matter.

2 Data

The macroeconomic variable I use to proxy macroeconomic activity is the monthly, seasonally adjusted industrial production as obtained from the Federal Reserve Economic Data (FRED). This variable ranges from January 1950 until January 2012. Throughout this report, I use the log difference transformation of industrial production, which can be calculated as $y_t = \ln(IP_t) - \ln(IP_{t-1})$, where IP_t is the value of industrial production at time t. These transformed values of industrial production can be interpreted as growth in industrial production.

One of the uncertainty measures I use is the conditional volatility of the monthly opening prices of the S&P 500 Index. These monthly opening prices are obtained from the *Yahoo Finance* website, and range from January 1950 until January 2012. I transform this data into returns by taking the log differences of the opening prices in the same manner as explained above. The conditional volatility can then be estimated from these returns, which is explained in detail in section 3.3.

I also use the daily opening prices of the S&P 500 Index to estimate both the monthly realised volatility and the monthly conditional volatility of the index. The returns can also be found on the *Yahoo Finance* website, and they range from the 3 January 1950 until 30 December 2012. I again transform these opening prices into financial returns by taking the log differences.

3 Methodology

3.1 The LSTAR Model

In order to estimate the effect of uncertainty on macroeconomic activity Jones and Enders (2016) used a nonlinear framework that allows for the sign and size of uncertainty shocks to have asymmetric effects. This nonlinear framework is a logistic smooth-transition autoregressive (LSTAR) model, which allows for different regimes, one with high uncertainty and one with low uncertainty. These two regimes should account for recessions and expansions, while the smooth transition in between the two regimes should model the periods in between these recessions and expansions. The regime the economy is experiencing is determined by the magnitude of the uncertainty measure used in the LSTAR model.

The LSTAR model as defined by van Dijk et al. (2002) can be written as follows:

$$y_t = \phi_1' x_t + \phi_2' x_t G(s_t; \gamma, c) + \varepsilon_t, \tag{1}$$

where $x_t = (1, y_{t-1}, y_{t-2}, ..., y_{t-p})'$ and $\phi_i = (\phi_{i,0}, \phi_{i,1}, ..., \phi_{i,d})'$, $i = 1, 2$. ε_t in the above equation is assumed to be a martingale difference sequence with respect to $\Omega_{t-1} = \{y_{t-1}, ..., y_{t-p}\}$, such that $E[\varepsilon_t|\Omega_{t-1}] = 0$ and $E[\varepsilon_t^2|\Omega_{t-1}] = \sigma^2$. The smooth transition between the two regimes is defined by the function $G(s_t; \gamma, c)$, which in this case is a first-order logistic function:

$$G(s_t; \gamma, c) = (1 + \exp[-\gamma(s_t - c)])^{-1}, \, \gamma > 0, \tag{2}$$

where s_t denotes the transition variable, γ the smoothness parameter and c the centrality parameter, and $G(s_t; \gamma, c)$ is a continuous function bounded between 0 and 1. The smoothness parameter, γ, determines the smoothness of the transition between the two regimes, while the centrality parameter, c, determines the threshold between the two regimes. The regime at time t is controlled by the value of the transition variable s_t: The transition function $G(s_t; \gamma, c)$ increases monotonically from 0 to 1 as s_t increases, and equals 0.5 when $s_t = c$. When γ tends to zero, the transition function approaches 0.5 and when γ equals zero, the LSTAR model is reduced to a linear autoregressive (AR) model. When γ becomes very large, the transition of $G(s_t; \gamma, c)$ from 0 to 1 becomes almost instantaneous at c, thereby reducing the LSTAR model to a threshold autoregressive (TAR) model, which is a regime-switching autoregressive model with an instantaneous transition between its two regimes. The fact that the LSTAR model can be reduced to a linear AR model and to a TAR model makes it a useful tool for modelling business cycle variables.

The parameters of the LSTAR model can be estimated using nonlinear least squares estimation (Teräsvirta, 1994), whom also argued that the most accurate estimates of γ can be found using a standardised transition variable. This standardising entails dividing the time series of the transition variable by its own standard deviation. Throughout my investigation I therefore use standardised transition variables to estimate the LSTAR models. The parameters of the LSTAR model can then be estimated by solving the following problem:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_{t=1}^{T} [y_t - F(x_t; \theta)]^2, \tag{3}$$

where $F(x_t; \theta) = \phi_1' x_t + \phi_2' x_t G(q_t; \gamma, c)$. This can be solved using various optimisation algorithms, which converges very well as long as correct starting values are chosen. These starting values can be found using a grid search over all possible values of γ and c. For fixed values of γ and c, the LSTAR model is linear in ϕ_1 and ϕ_2 (Franses et al., 2014), and therefore, conditional on γ and c, ϕ_1 and ϕ_2 can be estimated using ordinary least squares (OLS). The parameters which result in the lowest sum of squared residuals (SSR) for this regression are chosen to be the starting values for the optimisation algorithm.

3.2 Tests for Nonlinearity

Before estimating the LSTAR model, it makes sense to investigate whether there is reason to believe the data to be nonlinear. I therefore submit the data to several different tests, which are described in this section. These tests were also performed in Jones and Enders (2016).

3.2.1 Testing for LSTAR Behaviour

Teräsvirta (1994) describes a way to test for LSTAR behaviour, where linearity is tested against the LSTAR model. This test entails approximating the logistic function displayed in equation

(2) with a third-order Taylor expansion, which when combined with a Lagrange multiplier (LM) test results in the following auxiliary regression:

$$y_t = \beta_0' x_t + \beta_1' \tilde{x}_t s_t + \beta_2' \tilde{x}_t s_t^2 + \beta_3' \tilde{x}_t s_t^3 + e_t, \tag{4}$$

where s_t denotes the transition variable, $\tilde{x}_t = (y_{t-1}, y_{t-2}, ..., y_{t-p})'$ and $x_t = (1, \tilde{x}_t')'$. The null hypothesis for this test statistic is $H_0 : \beta_1 = \beta_2 = \beta_3 = 0$, and the test statistic has an asymptotic F distribution. The value for p can be found by inspecting the Akaike Information Criterion (AIC)[1] and Bayesian Information Criterion (BIC)[2] of a simple AR model and choosing for p the number of lags which resulted in the lowest respective AIC and BIC.

3.2.2 Regression Error Specification (RESET) Test

The RESET test tests a null hypothesis of linearity against an alternative of nonlinearity (Heij et al., 2004). This is done by performing an LM test for a regression of residuals of an AR(p) model for the macroeconomic variable on the regressors and on k powers of the fitted values of said AR(p) model. This results in the following regression:

$$y_t = \beta' x_t + \gamma_1 \hat{y}_t^2 + ... + \gamma_{k-1} \hat{y}_t^k + \varepsilon_t, \tag{5}$$

where $\hat{y}_t$ denotes the fitted values of the AR(p) model. The value for p is again found by inspecting the AIC and BIC for several AR models with different orders, where p is set equal to the lag-order of the model with the lowest AIC and BIC. When k is chosen too high, it is likely that some of the added powers are multicollinear, leading to the inaccurate OLS estimates of the parameters. I therefore use $k = 4$, which is small enough not to cause multicollinearity, and also large enough to include enough information in the test regression. The null hypothesis of this LM test is that $\gamma_1 = \gamma_2 = ... = \gamma_{k-1} = 0$, and the test statistic follows an F distribution.

3.2.3 Testing for Threshold Effects

This test for threshold effects was developed by Hansen (1997) and finds out whether there is significant proof of thresholds in the data. This test is performed by choosing a large amount of possible thresholds and then modelling the data as a simple TAR model for each of those thresholds, and then comparing each of these models with a simple autoregressive model of the same order as the TAR model by calculating the F-statistic:

$$F_T(c) = T\left(\frac{\tilde{\sigma}_T^2 - \hat{\sigma}_T^2(c)}{\hat{\sigma}_T^2(c)}\right), \tag{6}$$

where $\tilde{\sigma}_T^2$ denotes the variance of a simple OLS regression of an AR(p) model, and $\hat{\sigma}_T^2(c)$ is the variance of a TAR(p,c) model, where p is the lag order and c is the chosen threshold. These variances are calculated as $\sigma^2 = \frac{1}{T}\sum_{t=1}^{T} \hat{e}_t^2$, where $\hat{e}_t$ denotes the residuals of either the regression for the AR(p) model or the regression for the TAR(p,c) model. The test statistic of this threshold test is then the supremum of all F-statistics. However, since the true threshold γ is not identified, the asymptotic distribution of this test statistic is unknown, and is therefore approximated using a simple bootstrapping procedure (Hansen, 1996). This bootstrapping procedure entails simulating the dependent (macroeconomic) variable from a standard normal distribution (mean 0 and variance 1), and then performing the same regressions for the simulated variable as described earlier, using the same thresholds. This results in another set of these earlier defined F-statistics, from which the p-value of the test can be found by counting the percentage of bootstrapped F-statistics which are exceeded by the supremum of the true F-statistics.

[1] $AIC = T\ln(SSR) + 2r$, with r the number of estimated parameters and SSR the sum of squared residuals
[2] $BIC = T\ln(SSR) + r\ln(T)$

3.3 The EGARCH Model

The uncertainty measure I use to replicate the main results of Jones and Enders (2016) is the monthly conditional volatility of the returns on the S&P 500 Index. One method of estimating this conditional volatility entails the use of an Exponential Generalised Autoregressive Conditional Heteroskedasticity (EGARCH) model. This EGARCH model describes the relation between past uncertainty shocks to the financial returns and the logarithm of the conditional variance (Franses and van Dijk, 2003). The reason I use an EGARCH model instead of a normal GARCH is that it is able to explain the leverage effect in a financial series. This leverage effect refers to the fact that losses have a greater effect on future volatilities than gains do. The accounting for the leverage effect by the EGARCH model ensures that positive and negative shocks to financial returns are allowed to have asymmetric effects. This fits well with one of the goals of this investigation, which is to find out whether uncertainty has asymmetric effects on macroeconomic activity.

In order to estimate an EGARCH model the returns need to be modelled as follows:

$$\Delta \ln(x_t) = a + \varepsilon_{1,t}, \tag{7}$$

where $\Delta \ln(x_t)$ is the log difference of the financial returns of the S&P 500 Index monthly opening price and a is a constant. The error term $\varepsilon_{1,t}$ conditional on the information available up to time $t-1$ follows a standard normal distribution with a mean of zero and a variance of h_t. With the returns properly modelled, the EGARCH specification as defined in Nelson (1991) can be written as

$$\ln(h_t) = \omega + \beta \ln(h_{t-1}) + \gamma \frac{\varepsilon_{1,t-1}}{\sqrt{h_{t-1}}} + \alpha \frac{|\varepsilon_{1,t-1}|}{\sqrt{h_{t-1}}}, \tag{8}$$

where $\varepsilon_{1,t}$ are the errors of the model for the returns displayed in equation (7). The model can be estimated using maximum likelihood estimation (Alexander, 2008).

3.4 Testing for EGARCH Behaviour

Before estimating the EGARCH model, it is necessary to subject the financial returns to several tests to examine whether the conditional volatility model is well specified. These tests are called the *Sign Bias Test*, the *Negative Sign Bias Test* and the *Positive Sign Bias Test*. The sign bias test examines the effect positive and negative shocks have on the volatility not predicted by the GARCH-type model. The positive and negative size bias tests examine the effects that large and small positive respectively negative return shocks have on the conditional volatility not predicted by the GARCH-type model.

In order to find the test statistics for these tests the standardised residuals of equation (7) are required, which are calculated as follows: $v_t = \varepsilon_{1,t}/\sqrt{\hat{h}_t}$, where $\hat{h}_t$ is estimated from a GARCH(1,1) model for the financial returns. Also, two dummy variables are needed: S^-_{t-1}, which equals 1 if $v_t < 0$ and 0 otherwise, and S^+_{t-1}, which equals 0 if $v_t < 0$ and 1 otherwise.

The respective test statistics for the *Sign Bias Test*, the *Negative Sign Bias Test* and the *Positive Sign Bias Test* are calculated by finding the t-scores of the coefficient b for each of the following three OLS regressions (Engle and Ng, 1993):

$$v_t^2 = a + bS^-_{t-1} + e_t, \tag{9}$$

$$v_t^2 = a + bS^-_{t-1}v_{t-1} + e_t, \tag{10}$$

$$v_t^2 = a + bS^+_{t-1}v_{t-1} + e_t. \tag{11}$$

The three tests mentioned above can also be considered jointly through the following regression:

$$v_t^2 = a + b_1 S^-_{t-1} + b_2 S^-_{t-1}v_{t-1} + b_3 S^+_{t-1}v_{t-1} + e_t, \tag{12}$$

where the joint test statistic can be calculated using an LM test for the null hypothesis of correct specification: $b_1 = b_2 = b_3 = 0$.

3.5 Using Daily Returns to Estimate the Uncertainty Measure

Jones and Enders (2016) use monthly returns to estimate conditional volatility as uncertainty measure. However, returns on an index such as the S&P 500 are observed multiple times per day. These returns with a higher frequency most likely contain more information than the monthly returns, and therefore using these daily returns may result in a more accurate uncertainty measure. I therefore propose two methods to estimate monthly uncertainty measures based on daily returns.

The first method I propose is using monthly realised volatility as uncertainty measure, which can be estimated using the daily returns on the S&P 500 Index using the following formula (Corsi, 2009):

$$RV_t = \sum_{j=1}^{N} r_{t,j}^2, \tag{13}$$

where $r_{t,j}$ denotes the j daily returns in month t. An LSTAR model with this realised volatility as uncertainty measure contains more information through the higher number of observations than the LSTAR model with EGARCH conditional volatility as uncertainty measure, leading me to believe that this uncertainty measure may lead to more accurate or more nuanced conclusions than those drawn in Jones and Enders (2016).

The second method I propose is a MIDAS regression to estimate the monthly conditional volatility. This type of regression allows left-hand and right-hand variables of a time series regression to have different frequencies (Ghysels et al., 2007). This means that it is possible to incorporate more information into estimating the monthly conditional volatility using a MIDAS regression with daily financial returns than Jones and Enders (2016) were able to using an EGARCH model with monthly data. The MIDAS model I use looks as follows (Ghysels et al., 2005):

$$V_{M,t} = 22 \sum_{d=0}^{\infty} w_d r_{t-d}^2 + \varepsilon_t, \tag{14}$$

where $V_{M,t}$ denotes the monthly conditional volatility, and w_d denotes the weight of the squared daily return of day t-d: r_{t-d}. The factor 22 forces the variance to be expressed in monthly units, which is done because a month usually has 22 trading days due to no trading being done during the weekends. Ghysels et al. (2007) proposes the weight given to the squared daily returns at day t-d to be equal to

$$w_d(\kappa_1, \kappa_2) = \frac{\exp(\kappa_1 d + \kappa_2 d^2)}{\sum_{i=0}^{\infty} \exp(\kappa_1 i + \kappa_2 i^2)}. \tag{15}$$

This specific weighting function ensures that the sum of all weights for time t equals one, it guarantees that all weights are positive and it can produce a large variety of shapes since it is a function of two parameters. In order to be able to properly estimate the model, the infinite sums in the equations (14) and (15) need to be delimited. Ghysels et al. (2005) chose the number 252 as limit since that number represents the average number of trading days in one year.

The parameters κ_1 and κ_2 in the MIDAS model can be estimated using maximum likelihood estimation. This could only be done after first assuming again that the monthly returns at time

t follow a time-varying normal distribution with mean 0 and variance h_t^M. This assumption together with the models given in equations (14) and (15) allowed for the log-likelihood function to be written as:

$$l(\kappa_1, \kappa_2) = -\frac{T}{2} \ln 2\pi - \frac{1}{2} \sum_{t=1}^{T} \Big(\ln V_{M,t} \Big) - \frac{1}{2} \sum_{t=1}^{T} \Big((V_t^M)^{-1} R_t^2 \Big), \tag{16}$$

where $V_{M,t}$ denotes the MIDAS conditional volatility in month t as defined in equation (14), and R_t denotes the monthly returns at time t, calculated as described in section 2. The estimates for κ_1 and κ_2 in the weighting function used to calculate $V_{M,t}$ are defined as the values maximising this log-likelihood function.

3.6 Impulse Response Functions

In Koop et al. (1996) a framework for estimating impulse response functions from nonlinear models is developed. According to them, an impulse response function measures the time profile of the effect of a shock on a time series. This shock and the resulting behaviour of the time series lead to a better understanding of the time series.

In this report I create several impulse response functions for industrial production, where the shock that occurs is an uncertainty shock, meaning that one of the residuals of the model for the uncertainty measure is set equal to some shock. In order to create an impulse response to a temporary uncertainty shock, the first step is to select a horizon for the impulse response, meaning that the time at which the uncertainty shock occurs and the length of the impulse response are chosen. Let the uncertainty shock occur at time t, then the innovation ε_t^* of the model for the uncertainty measure, which can be for example the EGARCH model, is set equal to the chosen shock. If the duration of the impulse response is chosen to be p months, then the innovations $\varepsilon_{t+1}^*, \varepsilon_{t+2}^*, ..., \varepsilon_{t+p}^*$ of the model for the uncertainty measure still have to be selected. These innovations are sampled from the residuals of the model for financial returns in equation (7), which is done using standard bootstrapping procedures. This means that the innovations are drawn from a uniform distribution with replacement from the residuals of the model for the monthly returns. The shock together with these bootstrapped innovations is then used to produce the uncertainty measure, which for the EGARCH model is $\{h_t^*\} = h_t^*$ through h_{t+p}^*. These $\{h_t^*\}$ values are then substituted in the estimated LSTAR model to generate the recursive values of the log differences of industrial production: $y_t^*, ..., y_{t+p}^*$. The impulse response of industrial production to the chosen shock and the selected set of innovations is then obtained by transforming the log differences of industrial production back to the actual values of industrial production. This process of bootstrapping the innovations, using them to calculate the uncertainty measure, and then calculating the values of industrial production is repeated 1,000 times. The mean value of these 1,000 impulse responses is then obtained, and the 95% confidence interval can be obtained by taking the 2.5- and the 97.5-percentile of all these bootstrapped impulse responses. This exact procedure and some minor variations to it are used throughout this report to create the impulse response functions. Whenever a variation was used, the exact changes I made in order to obtain the impulse responses are reported.

4 Results

4.1 Results of the Tests for Nonlinearity

The results of the three tests for nonlinearity on the macroeconomic variable industrial production are displayed in table 1. The AIC and BIC of a simple AR(p) model for industrial production show that the optimal lag order for these tests is 2.

For the test for LSTAR behaviour the transition variable s_t in equation (4) is set equal to the lagged log-difference of industrial production: y_{t-1}. This test resulted in a p-value of 0.02, which indicates that there is significant evidence of LSTAR behaviour being present in industrial production. The RESET test statistic is insignificant at the 5% significance level, and hence I can say that there is reason to believe that industrial production is nonlinear. Finally, the test for threshold effects showed that there is evidence at the 5% significance level that there are threshold effects present in industrial production. The aforementioned indicates that there is sufficient evidence to claim that the LSTAR model is a good choice to model industrial production.

Table 1: Test results of the tests for nonlinearity in industrial production

	Test statistic	p-value
Teräsvirta (1994)	2.54	0.02
RESET test	2.29	0.06
Threshold Effects	7.04	0.03

Interestingly, the test statistics displayed in table 1 are slightly different from those presented in Jones and Enders (2016). Especially the test statistic and p-value of the test for threshold effects are very different, since in Jones and Enders (2016) it was concluded that industrial production does not exhibit threshold effects, while I concluded otherwise. The differing test results are most likely the result of a change in the indexation of industrial production after the research in Jones and Enders (2016) was conducted.

4.2 EGARCH Tests and the EGARCH Model

The test results of the four tests discussed in section 2 are displayed in table 2. The GARCH(1,1) model I use for these tests is given by $h_t = 0.00009 + 0.11\varepsilon_{1,t-1}^2 + 0.85h_{t-1}$, where $\varepsilon_{1,t}$ are the residuals from equation (7), and h_t denotes the conditional volatility of these residuals.

The significant coefficient of the regression for the sign bias test indicates that positive and negative shocks have asymmetric effects on the conditional variance. The significant coefficients of the negative and positive size bias test also support the use of an asymmetric GARCH model, and hence I feel that there is sufficient proof to estimate the conditional volatility of the S&P 500 Index with an EGARCH model. The estimated EGARCH(1,1) model looks as follows:

$$\ln(h_t) = \underset{(-3.13)}{-0.87} + \underset{(4.32)}{0.21}|\varepsilon_{1,t-1}|/\sqrt{h_{t-1}} + \underset{(-4.09)}{0.89}\ln(h_{t-1}) \underset{(22.27)}{-0.11}\,\varepsilon_{1,t-1}/\sqrt{h_{t-1}}, \tag{17}$$

where the t-scores are displayed in parentheses below their corresponding parameters. The AIC and BIC found for this EGARCH(1,1) model are smaller than those found for the GARCH(1,1) model earlier in this section, which supports the decision to model the conditional volatility using the EGARCH(1,1) model. Note that the negative coefficient of $\varepsilon_{1,t-1}/\sqrt{h_t}$ ensures that negative shocks produce a higher variance than positive shocks of the same size.

Do not however, that the test results I found were again somewhat different from the ones found in Jones and Enders (2016). This is most likely caused by their GARCH(1,1) model being estimated differently. Namely, the ARCH coefficient in my GARCH(1,1) model equals 0.11, whereas theirs equalled 0.01. On top of the differing parameters, the exact implementation of their tests was somewhat vague, resulting in me following the implementation presented in Engle and Ng (1993), which did however ultimately lead to the same conclusions as in Jones and Enders (2016).

Table 2: Test results of the tests by Engle and Ng (1993)

	Sign Bias Test	Negative Sign Bias Test	Positive Sign Bias Test	Joint Test
Test statistic	4.68	-2.80	-4.25	23.56
p-value	0.00	0.00	0.00	0.00

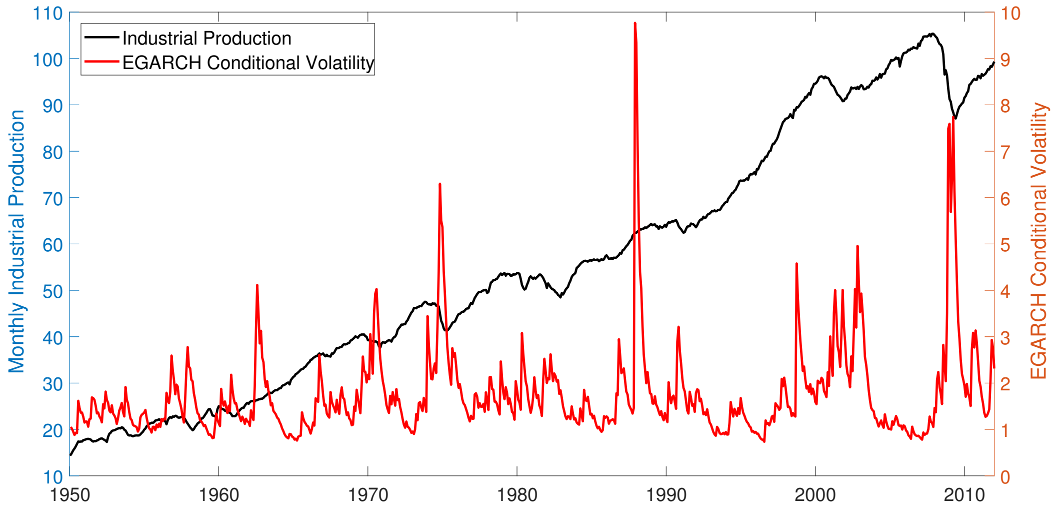

Figure 1: Plot with industrial production on the left y-axis, and standardised EGARCH conditional volatility on the right y-axis.

Figure 1 shows a plot industrial production together with the standardised EGARCH conditional volatility. This figure illustrates the relation between uncertainty and industrial production. A fairly clear correlation can be seen between peaks in the conditional volatility and smaller growth or decreases in industrial production. Especially the financial crisis is a good example of this correlation, as a large peak in uncertainty occurs at the same time as a large decrease in industrial production. However, the largest peak in the figure, which occurs on Black Monday, October 19^{th} in 1987, is not paired with smaller growth or a decrease in industrial production. Since Black Monday did not occur during a recession, and since the financial crisis is one of the worst recessions to date, it seems that the effects of an uncertainty shock could depend on the business cycle at the time of the shock.

4.3 LSTAR Model with EGARCH Conditional Volatility

The LSTAR model used in Jones and Enders (2016) to produce their main results is the model for log differences of industrial production with standardised EGARCH conditional volatility as uncertainty measure, from now on referred to as the EGARCH-LSTAR model. The estimated EGARCH-LSTAR model looks as follows:

$$\begin{aligned} y_t &= \underset{(7.41)}{0.0031} + \underset{(6.69)}{0.2796} y_{t-1} + (\underset{(-4.93)}{-0.0051} + \underset{(3.48)}{0.3547} y_{t-1})[1 + \exp(-4.0314(h_t - 2.1750))]^{-1} + \varepsilon_t \\ \text{AIC} &= -2143.3 \qquad \text{BIC} = -2115.7, \end{aligned} \tag{18}$$

where y_t is the log difference of industrial production at time t and h_t is the standardised EGARCH conditional volatility. The t-scores for the first four parameters are displayed below their corresponding parameters. The parameter variances required to calculate these t-scores can be found by regarding the LSTAR model as if it were a linear autoregressive model. This is done by regressing $X_{t-1} = [1 + G(q_t; \hat{\gamma}, \hat{c}), y_{t-1}(1 + G(q_t; \hat{\gamma}, \hat{c}))]$ on y_t, where q_t denotes the standardised uncertainty sequence and $G(q_t; \hat{\gamma}, \hat{c})$ the transition function as defined in equation (2) computed for the estimated values of γ and c. This, in essence, is the same as estimating industrial production as an AR(1) model. The parameter variances can then be found by calculating the covariance matrix of the regression: $\sigma^2(X'X)^{-1}$, where σ^2 is the variance of the regression and $X = (X_1, X_2, ..., X_T)'$. These t-scores show that the first four parameters all significantly differ from zero. The t-score for the smoothness parameter γ is not reported since the LSTAR model reduces to an AR(1) model when $\gamma = 0$, leading to the parameters of the LSTAR model being undefined. A t-score for the centrality parameter c is also not reported since c is estimated based on the values of the uncertainty sequence, and with this sequence being EGARCH conditional volatility it will always be positive, making a t-test for $c = 0$ unnecessary.

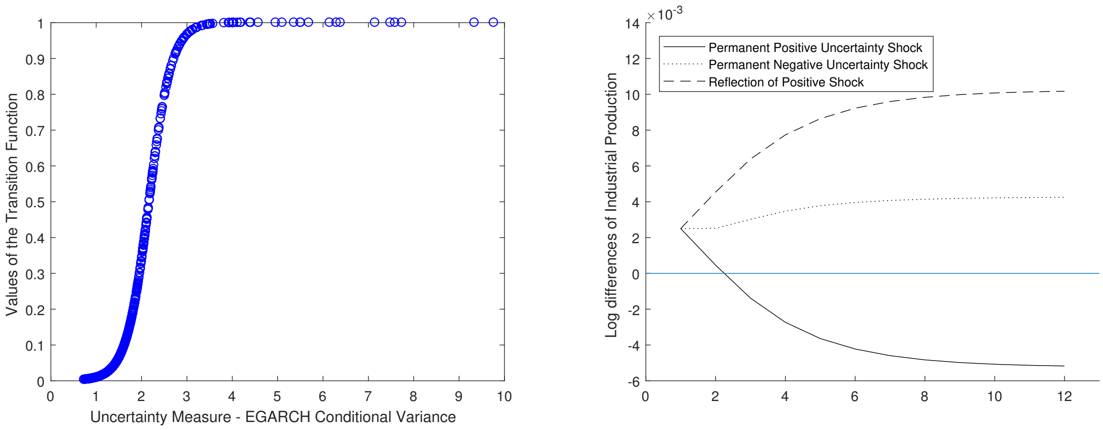

(a) **(b)**

Figure 2: Figure (a) shows the values of the transition function in the estimated EGARCH-LSTAR model; figure (b) shows the asymmetric effects of a continuing positive and negative shock to uncertainty.

Figure 2a shows the transition function of the estimated LSTAR model as a function of the standardised EGARCH conditional volatility h_t. Clearly, the transition between the two regimes occurs around the estimated centrality parameter, $c = 2.1750$. Furthermore, this transition is quite sharp, as a minor increase in uncertainty could lead to a jump from the low-uncertainty regime to the high-uncertainty regime. The EGARCH-LSTAR model together with the transition function show that when the transition function equals zero, meaning that uncertainty is low, the long-run equilibrium of industrial production growth is positive, and that the parameter y_{t-1} has coefficient 0.2796. Likewise, when the transition function equals one, and thus uncertainty is high, the long-run equilibrium of growth in industrial production is negative and the coefficient of y_{t-1} is now much larger at 0.6343. This shows that high values of uncertainty decrease industrial production and are also more persistent than lower values of uncertainty.

In order to further investigate whether the LSTAR model is a good choice for modelling industrial production I compare it with an AR model for industrial production. The optimal lag order for this AR model is 2, as this order resulted in the lowest respective AIC and BIC.

This estimated AR(2) model looks as follows:

$$\begin{aligned} y_t &= \underset{(3.90)}{0.0013} + \underset{(9.97)}{0.36} y_{t-1} + \underset{(3.28)}{0.12} y_{t-2} + \varepsilon_{2,t} \\ \text{AIC} &= -2128.9 \qquad \text{BIC} = -2115.1, \end{aligned} \tag{19}$$

where y_t is the log difference of industrial production at time t. The EGARCH-LSTAR model clearly has a lower AIC and BIC than this AR(2) model, indicating that the former is a better fit for industrial production, even though it estimates three more parameters than the AR(2) model.

As I mentioned in section 3.6, the estimated LSTAR model can be used to create impulse responses, which display the manner in which uncertainty shocks affect industrial production. Note that a positive uncertainty shock is in essence a large negative return; this large negative return results in a large negative innovation in the model for financial returns displayed in equation (7), which due to the EGARCH specification in equation (17) results in a higher conditional variance and hence higher uncertainty.

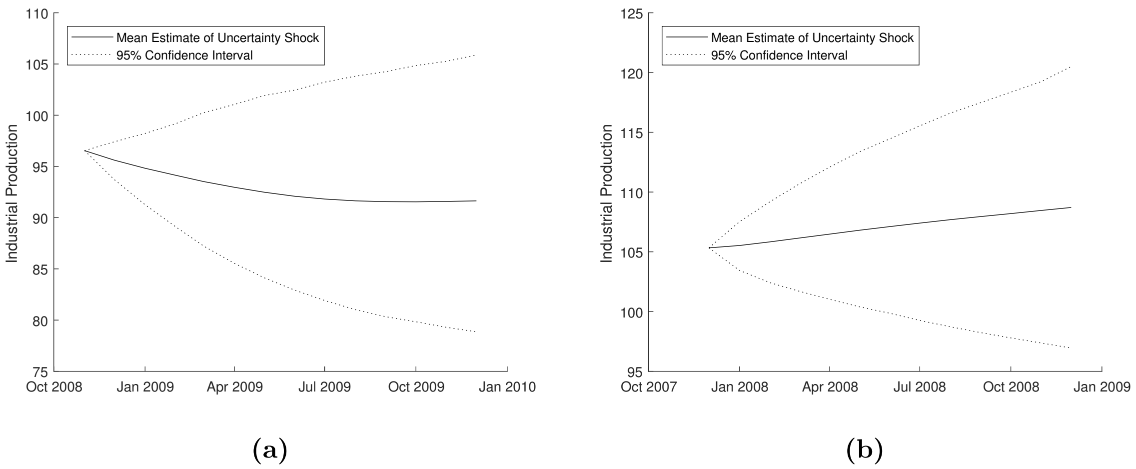

Figure 3: Impulse responses to temporary uncertainty shocks. (a) shows the impulse response to a positive December 2008 one-standard-deviation uncertainty shock occurring in December 2008; (b) shows the the impulse response to an actual December 2008 uncertainty shock occurring in January 2008.

Figure 2b displays the impulse response of industrial production to continuing positive and negative one-standard-deviation uncertainty shocks for a horizon of 12 months. This impulse response was created using a different method from the one explained in section 3.6. The model for this impulse response was initialised by setting the magnitude of uncertainty h_0^* equal to the value of the centrality parameter c, which was multiplied with the standard deviation of the uncertainty sequence to account for the standardising of said sequence. The value of industrial production in period 1 was set equal to the equilibrium of the AR(2) model in equation (19), which is 0.0025. These starting values are chosen to ensure that the positive uncertainty shock forces the estimated LSTAR model to be in the high-uncertainty regime, and that the negative shock forces the LSTAR model to be in the low-uncertainty regime, thereby illustrating how the two regimes differ. For the positive and negative uncertainty shocks the innovations $\varepsilon_{1,1}^*, ..., \varepsilon_{1,12}^*$ in the EGARCH model are set equal to $-\sqrt{h_0^*}, ..., -\sqrt{h_{11}^*}$ and $\sqrt{h_0^*}, ..., \sqrt{h_{11}^*}$ respectively. The h_t^* here are calculated recursively using the initial value of the uncertainty sequence and the innovations as defined above. These specific innovations are chosen because the standard deviation of the residual $\varepsilon_{1,t}$ in the model for financial returns (equation (7)) is $\sqrt{h_t}$. The reflected impulse response to the positive uncertainty shock in figure 2b clearly demonstrates how positive uncertainty shocks to the uncertainty sequence have a more than twice as large effect on the log differences of industrial production than negative shocks of the exact same size do; industrial production falls from 0.0025 to -0.0052 for the continuing positive shock, and rises from 0.0025 to 0.0043 for the negative shock. This clearly indicates that uncertainty has an asymmetric effect on industrial production.

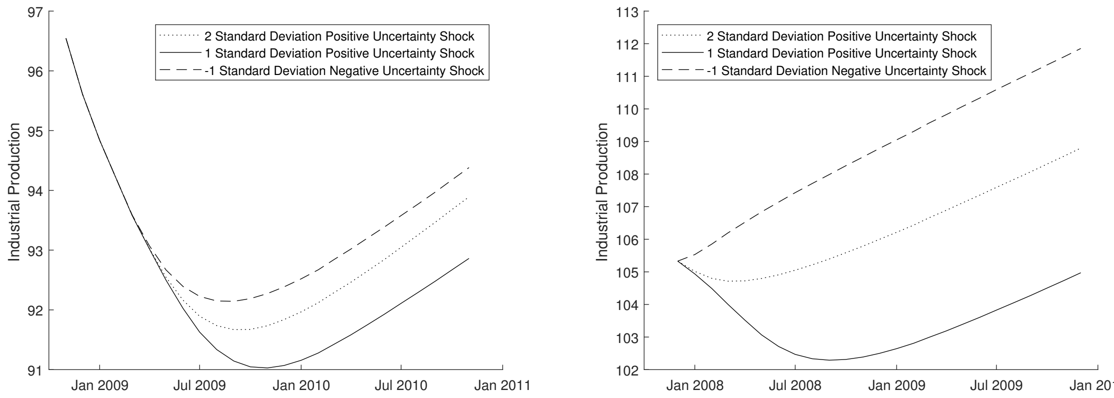

(a) **(b)**

Figure 4: The Asymmetric effects of temporary positive one-standard-deviation, positive two-standard-deviation and negative one-standard-deviation uncertainty shocks. (a) shows the impulse responses to temporary uncertainty shocks occurring in December 2008, which was during the financial crisis; (b) shows the impulse responses to temporary December 2008 uncertainty shocks occurring in January 2008. All lines show mean estimates of the impulse response functions.

Figure 3 shows the impulse responses of industrial production to temporary positive uncertainty shocks for a horizon of 12 months. The procedure used to obtain these impulse responses is slightly different from the procedure explained in section 3.6 in that here there is an additional step: the residuals ε_t of the LSTAR model are also bootstrapped using a uniform distribution with replacement from the actual residuals of the estimated LSTAR model, and are then used to find the recursive values of $\{y_t^*\}$. Figure 3a shows the impulse response of industrial production to a temporary positive one-standard-deviation shock occurring in December 2008, for which the first innovation is set equal to $-\sqrt{h_t}$, where h_t is the EGARCH conditional volatility for December 2008. The remaining innovations after the shock are bootstrapped using the procedure explained in section 3.6. The same method is employed to create the impulse response in figure 3b, yet here the innovation in January 2008 is set equal to $\varepsilon_{1,t}$ from the model for financial returns in equation (7), with t corresponding to December 2008. As before, the remaining innovations after January 2008 are bootstrapped. Figure 3a clearly shows that a positive uncertainty shock occurring during the financial crisis has a negative effect on industrial production, although the series appears to be rising back towards its original level. The figure on the right however shows that a positive uncertainty shock from the crisis has a very small effect on industrial production; in fact, industrial production rises despite the large uncertainty shock, which happens to have twice the size of the shock used for figure 3a. This actual December 2008 crisis shock used for the impulse response function in reality had a large negative effect on industrial production in December 2008, whereas it hardly affects industrial production if it occurs before the crisis, when the economy was strong. This shows how uncertainty shocks which occur when the economy is booming have significantly different effects than same-sized uncertainty shocks which occur during a recession.

Figure 4 shows the impulse responses of industrial production to several temporary uncertainty shocks for a horizon of 24 months. These impulse responses were created by setting an innovation $\varepsilon_{1,t}$ of the EGARCH model in equation (8) equal to a multiple of a one-standard-deviation shock. This means that the first innovation for the impulse response function is set equal to $-\sqrt{h_t}$ for a positive one-standard-deviation shock, $-2\sqrt{h_t}$ for a positive two-standard-deviation shock or $\sqrt{h_t}$ for a negative one-standard-deviation shock, where h_t is the conditional volatility for December 2008 estimated using the EGARCH model. The remaining innovations were bootstrapped as was explained in section 3.6. Figure 4a shows the impulse responses to shocks from December 2008 occurring in December 2008, and figure 4b shows the impulse responses to shocks from December 2008 occurring in January 2008. These figures clearly show that uncertainty shocks occurring during the crisis all have a negative effect on industrial production, even if the temporary shock is a negative uncertainty shock. These same shocks have a very different effect when they occur before the financial crisis; the negative uncertainty shock has a strictly positive effect, and the two negative uncertainty shocks have a much smaller negative effect if they occur before the crisis than if they occur during the crisis. It is therefore clear that the timing of uncertainty shocks makes a significant difference in the effect on industrial production, and that even quite large shocks, such as the positive two-standard-deviation shock, have a small effect on industrial production if they occur when the economy is booming.

4.4 LSTAR Model with Realised Volatility

In section 3.5 I discussed how I would use daily returns to estimate the monthly realised volatility of the S&P 500 Index, which would then be used as uncertainty measure in an LSTAR model for industrial production, which from here on referred to as the RV-LSTAR model. This estimated RV-LSTAR model looks as follows:

$$\begin{aligned} y_t &= \underset{(5.65)}{0.0052} + \underset{(2.49)}{0.2026} y_{t-1} + (\underset{(-4.19)}{-0.0060} + \underset{(2.53)}{0.3250} y_{t-1})[1 + \exp(-5.4397(RV_t - 0.2356))]^{-1} + \varepsilon_{3,t} \\ \text{AIC} &= -2124.9 \qquad \text{BIC} = -2097.2, \end{aligned} \tag{20}$$

where y_t denotes the log difference of industrial production and RV_t the standardised realised volatility of the S&P 500 Index. The t-scores for the parameters are reported between parentheses below their corresponding parameters, and the t-scores for γ and c are omitted for the same reasons as described in section 4.3. The t-scores show that the first four parameters significantly differ from zero. Both the AIC and BIC of this LSTAR model are higher than those of the AR(2) model for industrial production (equation (19)), indicating that the AR(2) model may be better able to model industrial production.

The estimate of the centrality parameter $c = 0.2356$ is somewhat surprising, as it is very close to zero. However, upon closer inspection of the transition function, displayed in figure 5, it becomes clear why this estimate is rather small. Most of the realised volatilities lie very close to zero, indicating that the distribution of these realised volatilities is negatively skewed. Since the centrality parameter is usually estimated to be where the density of the transition variable is the highest, it is logical that c was estimated to be very close to zero. This specific transition function suggests that the low-uncertainty regime lies on the left of the y-axis in the figure, and therefore the model can only be in a low-uncertainty regime for negative values of realised volatility. A negative volatility is of course not possible, meaning that this LSTAR model with standardised realised volatility as uncertainty measure does not allow for a low-uncertainty regime. However, it is still the case that when uncertainty is low, the long-run equilibrium is positive and when uncertainty is high it is negative. Furthermore, the coefficient of y_{t-1} increases as uncertainty increases, and hence high values of uncertainty decrease industrial production and are more persistent than low values of uncertainty.

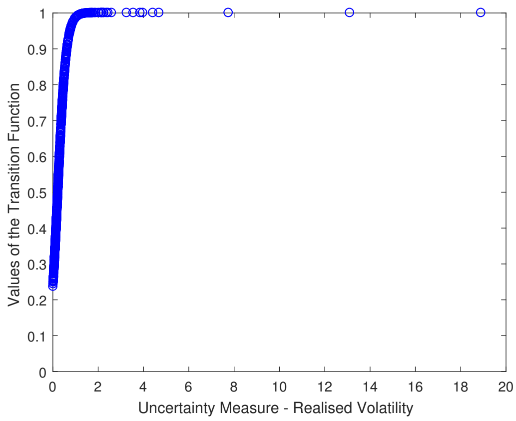

Figure 5: The values of the transition function in the estimated RV-LSTAR model

In order to obtain impulse response functions for this estimated LSTAR model, I model the estimated realised volatility as an AR(2) model so as to obtain a recursive relation for the uncertainty measure through which the temporary shocks can have a long-lasting effect on uncertainty and therefore on industrial production. The optimal lag order for this model is 2 as it resulted in the lowest AIC and BIC. The AR(2) model for realised volatility looks as follows:

$$\begin{aligned} &RV_t = \underset{(2.49)}{0.0010} + \underset{(34.70)}{0.34}\, RV_{t-1} + \underset{(4.40)}{0.14}\, RV_{t-2} + \varepsilon_{4,t} \\ &\text{AIC} = -6157.8 \qquad \text{BIC} = -6144.0, \end{aligned} \tag{21}$$

where RV_t denotes the (non-standardised) monthly realised volatility. Do note that an AR model is a linear model, and hence it incorporates positive and negative shocks in a symmetric manner into the time series. However, the manner in which the time series converges back to its long-run equilibrium does differ for positive and negative shocks.

Using the AR(2) model for realised volatility, I produce impulse responses to temporary uncertainty shocks in the same manner as in figure 4. Note that for the EGARCH model a negative innovation resulted in higher uncertainty, yet this is not the case for an AR(2) model, as here a positive innovation results in a higher realised volatility, and hence a higher uncertainty. Therefore, for a positive one-standard-deviation shock I set an innovation of the AR(2) model equal to $\sqrt{RV_t}$, and for the positive two-standard-deviation and negative one-standard-deviation shocks I set the first innovation equal to $2\sqrt{RV_t}$ and $-\sqrt{RV_t}$ respectively, with RV_t the (non-standardised) realised volatility in December 2008. The innovations following the uncertainty shock are again bootstrapped from the residuals in equation (7). Figure 6a shows the impulse responses to temporary uncertainty shocks occurring in December 2008, and figure 6b shows the impulse responses to temporary uncertainty shocks from December 2008 occurring in January 2008. The impulse responses show that the positive uncertainty shocks occurring during the financial crisis have a large negative effect on industrial production, while these positive shocks have a much smaller negative effect on industrial production when they occur before the crisis. For negative uncertainty shocks, the effect on industrial production is positive, and is slightly larger before the crisis than it is during the crisis. This indicates that reasonable size uncertainty shocks occurring during a recession have a greater effect on industrial production than same-size shocks occurring during an expansion.

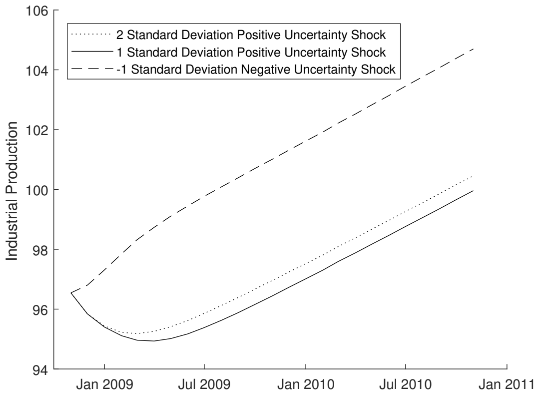

(a)

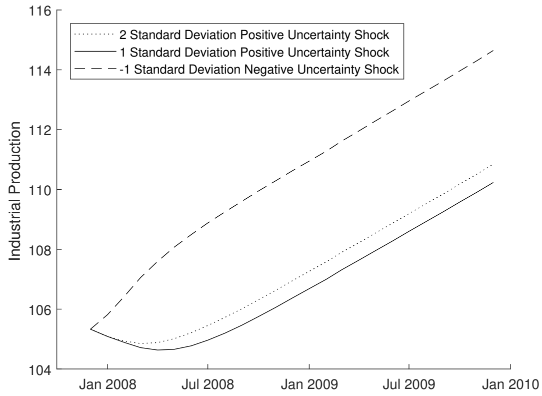

(b)

Figure 6: The Asymmetric effects of temporary positive one-standard-deviation, positive two-standard-deviation and negative one-standard-deviation uncertainty shocks. (a) shows the impulse responses to temporary uncertainty shocks occurring in December 2008; (b) shows the impulse responses to temporary December 2008 uncertainty shocks occurring in January 2008. All lines show mean estimates of the impulse response functions.

4.5 LSTAR Model with MIDAS Conditional Volatility

In section 3.5 I explain how the daily returns on the S&P 500 Index can be used to estimate the monthly conditional volatility. The estimation procedure suggests the following weights for the weighting function displayed in equation (15): $\kappa_1 = -0.0377$ and $\kappa_2 = -3.4979 \times 10^{-11}$. The weighting function resulting from this specific combination of parameters is displayed in figure 7a, where it can be seen how the weights for the first 25 days are significantly larger than the weights for the following days, and how the weights decline over time towards zero. The conditional volatility estimated using this weighting function in equation (14) can, after standardising, be used as uncertainty measure in an LSTAR model for industrial production, which I from hereon refer to as the MIDAS-LSTAR model. Note that the conditional volatility is estimated based on the daily returns from the first 252 days before each month, and since there are no daily returns available from before January 1950, the first estimated conditional volatility is for January 1951. The estimated MIDAS-LSTAR model then looks as follows:

$$\begin{aligned} y_t &= \underset{(6.30)}{0.0044} + \underset{(3.38)}{0.2167} y_{t-1} + (\underset{(-4.62)}{-0.0053} + \underset{(2.64)}{0.2922} y_{t-1})[1 + \exp(-5.3447(V_{M,t} - 0.4046))]^{-1} + \varepsilon_{5,t} \\ \text{AIC} &= -2135.1 \qquad \text{BIC} = -2107.5, \end{aligned} \tag{22}$$

where y_t denotes the log difference of industrial production and $V_{M,t}$ the standardised MIDAS conditional volatility. Again, the t-scores in parentheses below their corresponding parameters show how all parameters significantly differ from zero, and for the same reasons as before I do not report the t-scores for γ and c. The AIC for the LSTAR model is slightly higher than the AIC of the AR(2) model for industrial production, whereas its BIC is smaller than the one for the AR(2) model. Since the BIC punishes more strictly for a higher number of estimated parameters and tends to select more parsimonious models that underfit the data (Burnham and Anderon, 2002), I believe that in this case the LSTAR-MIDAS model is a better fit than the AR(2) model.

The transition function resulting from the estimated γ and c is displayed in figure 7b. Again the estimate for the centrality parameter is very close to zero, which, as can be seen from the figure, is due to most MIDAS conditional volatilities being very close to zero as well, indicating that their distribution is negatively skewed. The estimated transition function indicates that the low-uncertainty regime can only occur for conditional volatilities lower than zero, which is impossible since volatility cannot be negative. Therefore, MIDAS-LSTAR does not allow for a low-uncertainty regime. However, it is still the case that when uncertainty is very low, the long-run equilibrium of the model is positive, and it becomes negative as uncertainty increases, while at the same time the coefficient of y_{t-1} increases. This demonstrates how high values of uncertainty negatively affect the growth of industrial production, and how these higher values are more persistent than lower values.

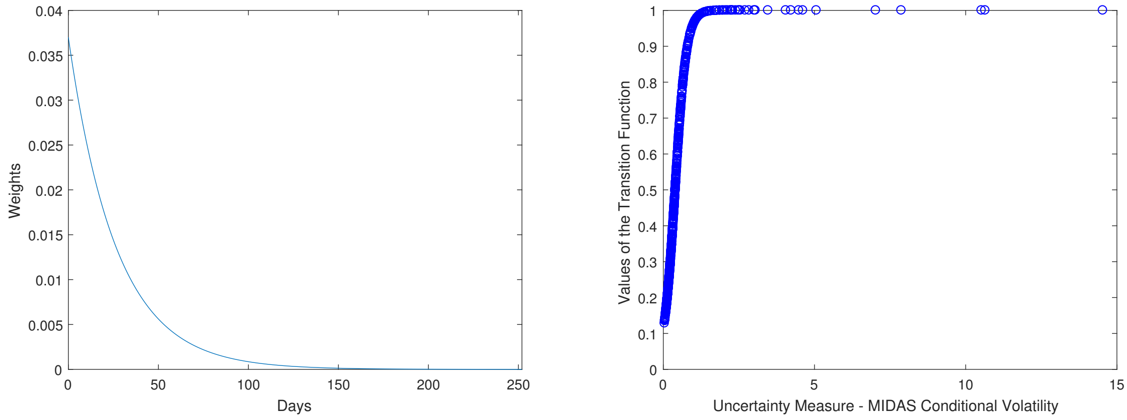

(a) (b)

Figure 7: Figure (a) shows the estimated weighting function for the MIDAS model; figure (b) displays the transition function of the estimated MIDAS-LSTAR model.

In order to obtain impulse responses for this LSTAR model, I again model the uncertainty measure as an AR(2) model. I use lag-order two for the AR model for the sake of consistency and since this order results in a relatively low AIC and BIC. The AR(2) model looks as follows:

$$\begin{aligned} V_{M,t} &= \underset{(2.14)}{0.0005} + \underset{(79.84)}{0.78}\, V_{M,t-1} - \underset{(-2.02)}{0.05}\, V_{M,t-2} + \varepsilon_{6,t} \\ \text{AIC} &= -6803.5 \qquad \text{BIC} = -6789.8, \end{aligned} \tag{23}$$

where $V_{M,t}$ is the (non-standardised) MIDAS conditional volatility. Using this AR(2) model, I produce impulse responses to temporary uncertainty shocks in the same manner as for figures 4 and 6. For a positive one-standard-deviation shock I set an innovation of the AR(2) model in equation (23) equal to $\sqrt{V_{M,t}}$, and for the positive two-standard-deviation and negative one-standard-deviation shocks I set this innovation equal to $2\sqrt{V_{M,t}}$ and $-\sqrt{V_{M,t}}$ respectively, where $V_{M,t}$ denotes the MIDAS conditional volatility for December 2008. The remaining residuals after the shock are again bootstrapped from the residuals of equation (7). Figure 8a then shows the impulse responses to temporary uncertainty shocks occurring in December 2008, and figure 8b shows the impulse responses to temporary uncertainty shocks from December 2008 occurring in January 2008. The impulse responses show that negative shocks to the uncertainty sequence have a positive effect on industrial production, and this effect is slightly larger before the financial crisis than during the crisis. The positive uncertainty shocks all have a negative effect on industrial production, yet this effect is larger during the crisis than it is before the crisis, and on top of that, it takes longer for industrial production to attain its previous level during the crisis than it does before the crisis. This again illustrates that uncertainty shocks during the crisis have a greater effect on industrial production than same-sized shocks occurring during an economic expansion.

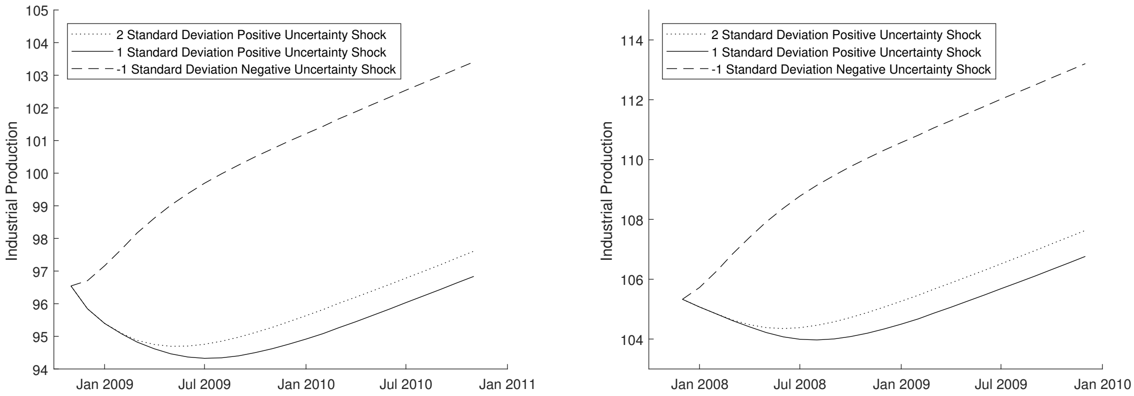

(a) **(b)**

Figure 8: The Asymmetric effects of temporary positive one-standard-deviation, positive two-standard-deviation and negative one-standard-deviation uncertainty shocks. (a) shows the impulse responses to temporary December 2008 uncertainty shocks occurring in December 2008, which was during the financial crisis; (b) shows the impulse responses to temporary December 2008 uncertainty shocks occurring in January 2008. All lines show mean estimates of the impulse response functions.

5 Comparison and Discussion of the Models

In this section I compare the three estimated LSTAR models with each other based on their respective uncertainty measures and based on their respective impulse response functions.

5.1 The Models and Uncertainty Measures

Of the three different uncertainty measures I investigated, the EGARCH-LSTAR model had lowest respective AIC and BIC of the three models, indicating that standardised EGARCH conditional volatility is best able to measure uncertainty when modelling industrial production. Furthermore, both the EGARCH-LSTAR and the MIDAS-LSTAR model outperform the AR(2) model based on their respective information criteria. The RV-LSTAR model performed worst based on its AIC and BIC, and on top of that it was outperformed by the AR(2) model for industrial production, indicating that standardised realised volatility as used in this paper is not as fitting as uncertainty measure as EGARCH conditional volatility and MIDAS conditional volatility.

Figure 9 shows the three standardised uncertainty measures. Realised volatility and MIDAS conditional volatility clearly display sharper and higher peaks than the EGARCH conditional volatility, although the peaks in the latter uncertainty measure occur more often. All three uncertainty measures clearly show large peaks for Black Monday and for the recent financial at the correct times.

Interestingly, the levels of the three standardised uncertainty measures are fairly different, which is the result of the standardising by the standard deviations of the respective measures. The EGARCH conditional volatility is the least volatile of the three, whereas realised volatility is by far the most volatile. This results in the level of the standardised EGARCH volatility being the largest, the level of the standardised MIDAS conditional volatility being the second largest and the level of the standardised realised volatility being the smallest. These low levels show that the distribution of the magnitude of uncertainty for realised volatility and for MIDAS conditional volatility is negatively skewed, and since the centrality parameter is estimated to be where the density of the transition variable is highest it is clear why c was estimated very close to zero for the RV-LSTAR model and MIDAS-LSTAR model.

It appears that there is some evidence of correlation between the level of the uncertainty measure, and therefore the estimate for the centrality parameter c, and the fit of the LSTAR model based on the AIC and BIC. The lower the value of c, the larger the amount of uncertainty sequence observations falling in periods of relatively high uncertainty. Since the long-run equilibrium of industrial production is negative when uncertainty is high, this indicates that the MIDAS-LSTAR model underestimates industrial production during actual expansions and overestimates industrial production during recessions. This underestimating and overestimating is even worse for the RV-LSTAR model.

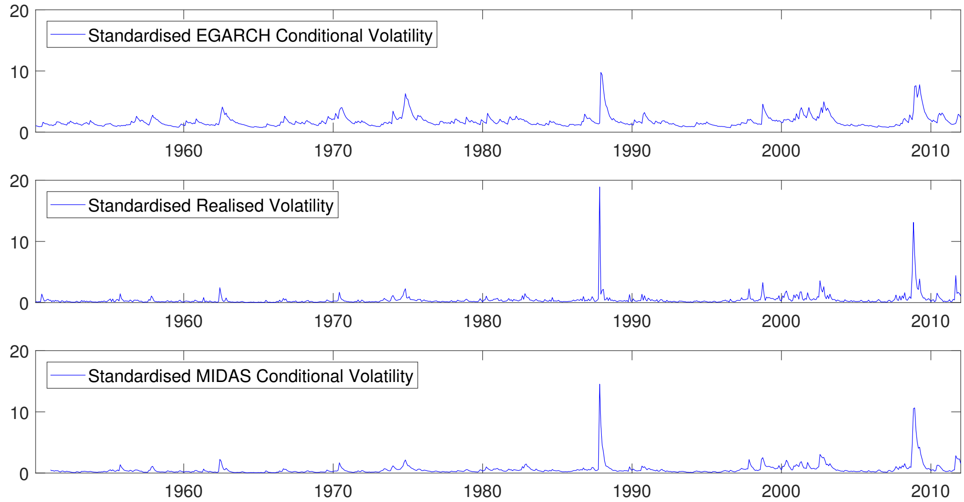

Figure 9: The three standardised uncertainty measures. The first plot shows the EGARCH conditional volatility, the second the estimated monthly realised volatility based on daily returns and the third shows the estimated monthly MIDAS conditional volatility based on daily returns.

The fact that the RV-LSTAR model and the MIDAS-LSTAR model do not allow for a low-uncertainty measure does not necessarily indicate that realised volatility and MIDAS conditional volatility are unfit to be used as uncertainty measures. A logarithmic transformation of both non-standardised uncertainty measures reduces the skewness of both uncertainty sequences significantly and allows for both a low-uncertainty and a high-uncertainty regime in the resulting estimated LSTAR models. On top of that the use of these transformed uncertainty measures results in lower AICs and BICs than were found for the RV-LSTAR model and the MIDAS-LSTAR model. These findings indicate that the transformed realised volatility and MIDAS conditional volatility could be better fit as uncertainty measures in the LSTAR model than their non-transformed versions.

5.2 The Impulse Response Functions

Figure 10 shows the impulse responses of the three LSTAR models to a positive two-standard-deviation crisis shock occurring either during (10a) or before the financial crisis (10b). These impulse responses are the exact same ones as those discussed in section 4. The figures show that large positive uncertainty shocks occurring during the financial crisis have a larger negative effect on industrial production effect than these same shocks occurring before the crisis. Especially the time it takes industrial production to return to its previous level is much longer when the shocks occur during the financial crisis.

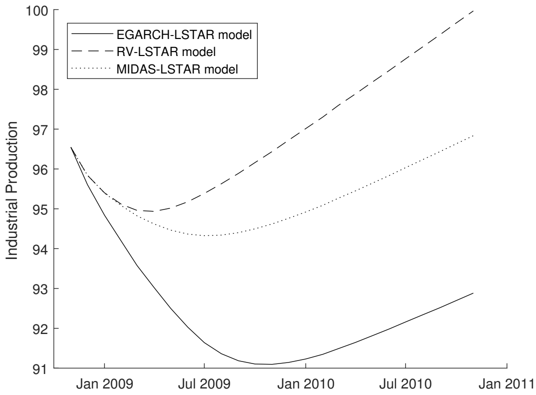

(a)

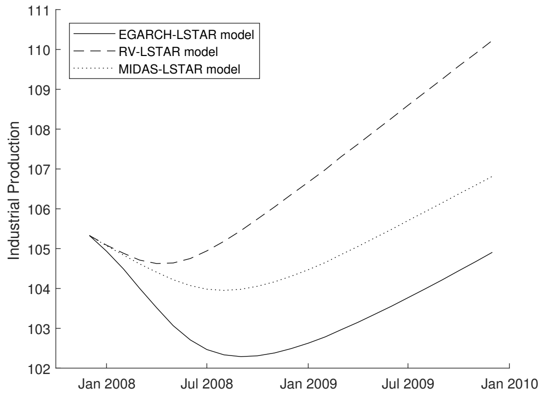

(b)

Figure 10: Figure (a) displays the impulse responses of industrial production to a temporary positive two-standard-deviation December 2008 uncertainty shock occurring in December 2008 calculated for each of the three estimated LSTAR models; figure (b) displays the impulse responses of industrial production to a temporary positive two-standard-deviation December 2008 uncertainty shock occurring in January 2008 calculated for each of the three estimated LSTAR models.

Both figures show that the EGARCH-LSTAR model responds more strongly to a positive uncertain shock than the other two models, while the absolute size of the shocks is larger for the latter two models. This difference can be explained by the different models used for the uncertainty measures; the EGARCH model is an asymmetric model which guarantees that negative uncertainty shocks produce larger conditional variances than positive shocks do, whereas positive and negative shocks of the same size have the same absolute effect in a linear AR(2) model.

Interestingly, the relation between the size of the uncertainty shock and the size of the impulse response of each specific model appears to be inverse. The EGARCH-LSTAR model experiences the smallest shock in absolute terms yet the effect on industrial production is the largest; the MIDAS-LSTAR model experiences a larger shock in absolute terms yet the effect on industrial production is smaller; and the RV-LSTAR model experiences the largest shock in absolute terms yet the effect on industrial production is the smallest. This is in accordance with what I discussed in section 5.1, where I explained that the RV-LSTAR model overestimates industrial production during high-uncertainty regimes, even more so than the MIDAS-LSTAR model does. As both these models assume they are in a high-uncertainty regime too often, as is the case for the shocks used to produce these impulse responses, the underestimating of industrial production shown in these impulse responses makes sense.

6 Conclusion

The purpose of my investigation into the effect of uncertainty on macroeconomic activity was to find out whether positive and negative shocks affect macroeconomic activity in asymmetric manners, and to find out whether the state of the economy at the time an uncertainty shock occurs affects the effect of such a shock.

Using an LSTAR model with EGARCH conditional volatility as uncertainty measure I created several impulse response functions which showed that the effect of positive and negative shocks on industrial production, and therefore on macroeconomic activity, is asymmetric.

I also estimated two other LSTAR models, one with realised volatility and one with MIDAS conditional volatility as uncertainty measure. The impulse response functions of the three different models showed that both positive and negative uncertainty shocks have different effects on industrial production dependent on the state of the economy: both positive and negative uncertainty shocks occurring during the crisis have a larger effect than the same-sized shocks occurring when the economy is booming. This all underlines how the timing of an uncertainty shock greatly influences the effect the shock has on macroeconomic activity. These findings show that monetary policy should aim to reduce and control the magnitude of uncertainty during recessions as much as possible to ensure the macroeconomy keeps functioning properly.

7 Suggestions for Further Research

As I mentioned briefly in section 5.1, the fit of the LSTAR models with realised volatility and MIDAS conditional volatility can be improved by using a logarithmic transformation of these uncertainty measures. The resulting LSTAR models were promising in that they both allowed for the two regimes to be attained by the transition variable. The use of these transformed uncertainty measures should be further researched in order to draw any conclusions on them being fitting proxies for uncertainty.

Also, the AR models for uncertainty I used to produce impulse response functions are linear models which do not allow for asymmetric responses to uncertainty shocks. They therefore may not be the best choice for modelling uncertainty, and hence it may be fruitful to experiment with using different, possibly nonlinear models for uncertainty which do allow for this asymmetry as this may result in more accurate conclusions regarding the effect of uncertainty on macroeconomic variables.

A final suggestion for further research is the use of other proxies for macroeconomic activity in combination with the uncertainty measures proposed in this report, since the conclusions may differ from my conclusions, which are all based on industrial production.

References

Alexander, C. (2008). *Volume II: Practical Financial Econometrics*. John Wiley & Sons, Ltd.

Arellano, C., Bai, Y., and Kehoe, P. (2012). Financial Markets and Fluctuations in Uncertainty. *Federal Reserve Bank of Minneapolis*.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–176.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2012). Really Uncertain Business Cycles. *National Bureau of Economic Research*. Working paper.

Burnham, K. P. and Anderon, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach*. Springer Verlag, New York, 2nd edition.

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Engle, R. F. and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *Journal of Finance*, 48(5):1749–1778.

Foerster, A. T. (2014). The Asymmetric Effects of Uncertainty. *Federal Reserve Bank of Kansas City Economic Review*, 99(3):526.

Franses, P. H. and van Dijk, D. (2003). *Nonlinear Time Series Models in Empirical Finance*. Cambridge University Press, Trumpington Street, Cambridge.

Franses, P. H., van Dijk, D., and Opschoor, A. (2014). *Time Series Models for Business and Economic Forecasting*. Cambridge University Press.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2005). There is a risk-return trade-off after all. *Journal of Financial Econometrics*, 76:509548.

Ghysels, E., Sinko, A., and Valkanov, R. (2007). MIDAS Regressions: Further Results and New Directions. *Econometric Reviews*, 26(1):53–90.

Grier, K. B., Henry, Ó. T., Olekalns, N., and Shields, K. (2004). The asymmetric effects of Uncertainty on Inflation and Output Growth. *Journal of Applied Econometrics*, 19:551565.

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

Hansen, B. E. (1996). Inference When a Nuisance Parameter Is Not Identified Under the Null Hypothesis. *Econometrica*, 64(2):413–430.

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

Jones, P. and Enders, W. (2016). The asymmetric effects of uncertainty on macroeconomic activity. *Macroeconomic Dynamics*, 20:1219–1246.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.

Koop, G., Pesaran, M. H., and Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147.

Mishkin, F. S. (2011). Monetary Policy Strategy: Lessons From the Crisis. *National Bureau of Economic Research*. Working Paper, number 16755.

Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, 59(2):347–370.

Stock, James H.and Watson, M. W. (2012). Disentangling the Channels of the 2007-2009 Recession. *National Bureau of Economic Research*. Working paper.

Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89(425):208–218.

van Dijk, D., Teräsvirta, T., and Franses, P. H. (2002). Smooth Transition Autoregressive Models - A Survey of Recent Developments. *Econometric Reviews*, 21(1):1–47.