BANK OF FINLAND DISCUSSION PAPERS

11 • 2001

Michael Ehrmann – Martin Ellison –
Natacha Valla
Research Department
3.8.2001

Regime-dependent impulse response functions in a Markov-switching vector autoregression model

Suomen Pankin keskustelualoitteita
Finlands Banks diskussionsunderlag

http://www.bof.fi

Suomen Pankki
Bank of Finland
P.O.Box 160, FIN-00101 HELSINKI, Finland
☎ + 358 9 1831

BANK OF FINLAND
DISCUSSION PAPERS

11 • 2001

Michael Ehrmann* – Martin Ellison** – Natacha Valla***
Research Department
3.8.2001

Regime-dependent impulse response functions in a Markov-switching vector autoregression model

The paper was finalized while Martin Ellison visited the Bank of Finland. We would like to thank Mike Artis, Stephen Cecchetti, Roger Farmer, Soren Johansen and Frank Smets for helpful comments.

* European Central Bank, michael.ehrmann@ecb.int. This paper has been written before the author joined the ECB.
** European University Institute, ellison@iue.it.
*** European University Institute, valla@iue.it.

Suomen Pankin keskustelualoitteita
Finlands Banks diskussionsunderlag

http://www.bof.fi

ISBN 951-686-723-5
ISSN 0785-3572
(print)

ISBN 951-686-724-3
ISSN 1456-6184
(online)

Suomen Pankin monistuskeskus
Helsinki 2001

Regime-dependent impulse response functions in a Markov-switching vector autoregression model

Bank of Finland Discussion Papers 11/2001

Michael Ehrmann – Martin Ellison – Natacha Valla
Research Department

Abstract

In this paper we introduce identifying restrictions into a Markov-switching vector autoregression model. We define a separate set of impulse responses for each Markov regime to show how fundamental disturbances affect the variables in the model *dependent* on the regime. We go to illustrate the use of these regime-dependent impulse response functions in a model of the U.S. economy. The regimes we identify come close to the "old" and "new economy" regimes found in recent research. We provide evidence that oil price shocks are much less contractionary and inflationary than they used to be. We show furthermore that the decoupling of the US economic performance from oil price shocks cannot be explained by "good luck" alone, but that structural changes within the US economy have taken place.

Key words: vector autoregression, regime switching, shocks, new economy

Tilasidonnaiset impulssivasteet markovilaisessa kahden regiimin autoregressiomallissa

Suomen Pankin keskustelualoitteita 11/2001

Michael Ehrmann – Martin Ellison – Natacha Valla
Tutkimusosasto

Tiivistelmä

Tutkimuksessa kehitetään talouden toimintaa kuvaavia aikasarjamalleja siten, että niillä voidaan käsitellä talouden rakenteiden tilamuutoksia. Tämä tehdään liittämällä identifioivia rajoitteita markovilaiseen usean tilan vektoriautoregressiomalliin. Kullekin talouden tilalle määritellään omat impulssivasteensa, niin että talouden reaktiot häiriöihin riippuvat tilasta, jossa talous on. Näitä tilasidonnaisia impulssivasteita havainnollistetaan Yhdysvaltain taloutta kuvaavalla mallilla. Tutkimuksessa tunnistettavat talouden tilat vastaavat läheisesti ”vanhan talouden” ja ”uuden talouden” regiimejä, jotka ovat tulleet esiin viimeaikaisissa tutkimuksissa. Tulosten mukaan öljyn hintasokkien vaikutukset eivät nykyisin ole yhtä kontraktiivisia ja inflatorisia kuin aiemmin. Tutkimuksessa osoitetaan myös, että Yhdysvaltain talouden aiempaa vähäisempi herkkyys öljyn hintojen vaikutuksille ei riipu pelkästään sattumasta, vaan johtuu taloudessa tapahtuneista rakennemuutoksista.

Asiasanat: vektoriautoregressio, regiiminmuutokset, sokit, uusi talous

Contents

1 Introduction

Vector autoregression models have rapidly established themselves as the dominant research methodology in empirical macroeconomics. Their popularity rests partly on their attractiveness as a unifying framework in which to analyse alternative theories and hypotheses but also in the significant theoretical developments which enhance the basic unrestricted model. One increasingly popular and important advance is to allow for Markov switching effects in the model. Since the seminal contribution of Hamilton (1989), it has become common to model economic time series as vector autoregressive processes subject to regime switches.

In a Markov-switching vector autoregression there are many parameters which can switch across regimes. Since means, intercepts, autoregressive parameters, variances or covariances may all be regime-dependent, the number of switching parameters is potentially very high. For example, in a four-variable two-lag model in which intercepts, autoregressive parameters, variances and covariances all switch, there is a total of forty-six switching parameters. Interpreting the estimation results from such a model is inevitably difficult. Whilst mean switches may have a natural interpretation such as the business cycle (like in Hamilton, 1989, or Krolzig and Toro, 1999), there is often no obvious explanation for switches in other parameters.

The problem of interpreting the parameter estimates from Markov-switching models is analogous to that of interpreting the parameter estimates from a simple unrestricted vector autoregression. A now standard approach suggested by Sims (1980) to the latter problem is to impose identifying restrictions on the parameter estimates and derive a structural form for the model that has economic intuition. Central to this structural vector autoregression approach is the use of impulse response analysis, which traces out how fundamental disturbances affect the variables in the model.

In this paper we argue that imposing similar identifying restrictions on Markov-switching models can give significant insights into the characteristics of the economy.[1] Our tool is the regime-dependent impulse response function, which traces out how fundamental disturbances affect the variables in the model *dependent* on the regime. Instead of one set of impulse response functions we have a set for each regime. The regime-dependent impulse response function conveniently summarises the information in the autoregressive parameters, variances and covariances of each regime, making interpretation much easier than solely on the basis of the switching of individual parameters. Asymmetries in terms of magnitude, persistence and significance of the impulse responses are easily revealed by comparing the regime-dependent functions.

Our approach combines Markov-switching and identification in a two-stage procedure of estimation and identification. In the estimation stage a Markov-switching unrestricted vector autoregression model is estimated, allowing means, intercepts, autoregressive parameters, variances and covariances to switch or be constant as desired. In the identification stage restrictions are

[1] For lovers of acronyms we might say we take a Markov-switching structural vector autoregression (MS-SVAR) approach.

imposed on the parameter estimates to derive a separate structural form for each regime, from which the regime-dependent impulse response functions can be calculated.

The possibility of regime-dependent impulse responses has recently been recognised by Krolzig and Toro (1999). However, they emphasise the response of the economy to changes in regime rather than the impulse responses within regimes. Our approach is also similar in spirit to that of Koop, Pesaran and Potter (KPP, 1996). Their Generalised Impulse Response Functions (GI) differ from the usual impulse responses in several ways. Traditional impulse response analysis reports the results of an experiment where a shock hits a system at time t, with no further shocks hitting afterwards, compared to a benchmark case where the system stays unperturbed all of the time. KPP, on the other hand, treat the future differently by allowing for future shocks, which are averaged out eventually. The GI do not represent the responses to a shock of a certain size and sign, but instead treat the shock itself as a random variable. The GI are therefore reported in terms of density functions, rather than time trajectories as usual. Our approach sticks to the traditional definition of impulse responses in these respects. We are much closer to KPP's GI in the treatment of history: both approaches recognise that the impulse response can depend on the point in time and the state of the economy at which the shock occurs (eg, whether the economy is hit during a recession or expansion). In general, KPP therefore treat not only the disturbance, but furthermore the history as a stochastic process, which again leads to a density function for the impulse responses. They also mention (p. 130) the possibility to condition "on particular subsets of the history or the shock. For example, one might condition on all histories where the economy was in recession in the most recent period." This idea is very close to the approach suggested here, in that what KPP call 'history' in our case is a Markov switching regime.

To illustrate the use of regime-dependent impulse response functions we apply our approach to a model of the US economy. We are able to identify regimes that come close to the "old economy" and "new economy" regimes identified in recent research. We analyse how oil price shocks affect the economy within the two regimes, and find that they are much less contractionary and inflationary than they used to be. We show furthermore that the decoupling of the US economic performance from oil price shocks cannot be explained by "good luck" alone, but that structural changes within the US economy have taken place. The impact of the 1999/2000 oil price increases on the current economic slowdown can therefore not have been of major importance.

2 Procedure

2.1 A Markov-switching vector autoregression model

Equation (1) describes a general Markov-switching vector autoregression model. The K endogenous variables X_t are explained by an intercept v_i, autoregressive terms of order p and a residual $A_i u_t$. In this general specification all parameters of the autoregression are allowed to switch between regimes[2] so each of the m regimes is characterised by an intercept v_i, autoregressive terms $B_{1i}, \ldots, B_{pi}$ and a matrix A_i.

$$X_t = \begin{cases} v_1 + B_{11}X_{t-1} + \ldots + B_{p1}X_{t-p} + A_1 u_t & \text{if } s_t = 1 \\ \vdots & \\ v_m + B_{1m}X_{t-1} + \ldots + B_{pm}X_{t-p} + A_m u_t & \text{if } s_t = m \end{cases} \tag{1}$$

$$u_t \sim N\left(0; I_K\right)$$

u_t is a K-dimensional vector of fundamental disturbances which are assumed to be normally distributed and uncorrelated at all leads and lags. The variance of each fundamental disturbance is normalised to unity to give the identity variance-covariance matrix. However, in equation (1) the fundamental disturbances are premultiplied by a regime-dependent matrix A_i. Equation (2) shows how consequently the variance-covariance matrix $\mathbf{\Sigma}_i$ of the residuals $A_i u_t$ will also be regime-dependent.

$$\mathbf{\Sigma}_i = E(A_i u_t u_t' A_i') = A_i E(u_t u_t') A_i' = A_i I_K A_i' = A_i A_i' \tag{2}$$

In Markov-switching models the regime s_t is assumed to follow a hidden m-state Markov chain. The probability of being in regime j next period conditional on the current regime i is assumed to be exogenous and constant. In an m-state model there are $m \times m$ such conditional transition probabilities. Equation (3) defines the conditional transition probabilities and collects them into an exogenous transition matrix P.

$$\Pr(s_{t+1} = j \,|s_t = i) = \rho_{ij}$$

$$P = \begin{bmatrix} \rho_{11} & \rho_{12} & \cdots & \rho_{1m} \\ \rho_{21} & \rho_{22} & \cdots & \rho_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{m1} & \rho_{m2} & \cdots & \rho_{mm} \end{bmatrix} \tag{3}$$

2.2 Estimation

Estimation of a Markov-switching model entails joint estimation of all the parameters and the hidden Markov chain followed by the regime. Since the Markov chain is hidden, the likelihood function has a recursive nature: optimal inference in the current period depends on the optimal inference made in the previous period. Under such conditions the likelihood cannot be maximised using standard techniques. The model can however be estimated by applying the Expectations-Maximization (EM) algorithm, see Hamilton (1990) and Krolzig (1997). The first *expectations* step optimally infers the hidden Markov

[2] In the terminology of Krolzig (1997) this is an MSIAH(m)-VAR(p) model.

chain for a given set of parameters. The second *maximisation* step then re-estimates the parameters for the inferred hidden Markov chain. These steps are continued until convergence.

Applying the EM algorithm to the Markov-switching vector autoregression in equation (1) gives estimates of the parameters associated with each regime, the transition probability matrix and an optimal inference of the hidden Markov chain followed by the regime. The estimation results therefore consist of parameters $\{\hat{v}_i; \hat{B}_{1i}, \ldots, \hat{B}_{pi}; \hat{\Sigma}_i\}$ for $i = 1, \ldots, m$, transition matrix $\hat{P}$, and the optimal inference $\hat{\xi}_{i,t} = \Pr(s_t = i)$ for $i = 1, \ldots, m$ and $t = 1, \ldots, T$. The latter are known as smooth probabilities.

2.3 Identification

We are interested in deriving the relationship between the fundamental disturbances and the endogenous variables for each regime. The identification problem arises because the EM algorithm gives only estimates of the variance-covariance matrices $\Sigma_1, \ldots, \Sigma_m$ and not the matrices $A_1, \ldots, A_m$. To identify these matrices we have to impose restrictions on the parameter estimates from the unrestricted model. The choice of which restrictions to impose is the subject of the structural vector autoregression literature, see *inter alia* Sims (1980), Blanchard and Quah (1989) and King, Plosser, Stock and Watson (1991). To derive the conditional impulse response functions we are free to impose a variety of identification schemes. As an example we describe below the recursive form identification by Sims (1980).

Each matrix A_i has K^2 elements to be identified so K^2 restrictions have to be imposed. The identity $A_i A_i' = \Sigma_i$ from equation (2) naturally imposes $K(K+1)/2$ restrictions because of the symmetry of the variance-covariance matrix Σ_i. This leaves $K(K-1)/2$ missing restrictions. Sims (1980) derives the additional restrictions by imposing a recursive structure on the model. The endogenous variables are ordered and it is assumed that the fundamental disturbance to a variable has only contemporaneous effects on the variable itself and on variables ordered below it. For example, in a four-variable system the third disturbance has only contemporaneous effects on the third and fourth endogenous variables. Under this identification procedure the matrix A_i is lower triangular and is exactly identified. It can be easily recovered from a Choleski decomposition of the matrix Σ_i.

2.4 Regime-dependent impulse response functions

Standard impulse response analysis shows how the endogenous variables in the model react to the fundamental disturbances. The impulse response functions summarise expected changes in the endogenous variables after a one standard deviation shock to one of the fundamental disturbances. For the Markov-switching model we introduce an analogous concept, the regime-dependent impulse response function. This describes the relationship between endogenous variables and fundamental disturbances *within* a regime. Regime-dependent impulse response functions are conditional on a given regime prevailing at the time of the disturbance and throughout the duration of the response.

The validity of regime conditioning depends on the time horizon of the impulse response and the expected duration of the regime. As long as the time horizon is not excessive and the transition matrix predicts regimes which are highly persistent then the conditioning is valid and regime-dependent impulse response functions are a useful analytical tool. For a longer time horizon or frequently switching regimes, it would be more attractive to condition on the expected path of the regime throughout the response. Such impulse responses could be calculated for our model and are close to the Generalised Impulse Response Functions of Koop, Pesaran and Potter (1996). However, the additional information they contain is limited when regime switches are exogenous and regimes are persistent.

In our general model there are mK^2 regime-dependent impulse response functions, corresponding to the reaction of K variables to K disturbances in m regimes. Equation (4) mathematically defines the regime-dependent impulse response functions for regime i. It shows the expected changes in endogenous variables at time $t+h$ to a one standard deviation shock to the k-th fundamental disturbance at time t, conditional on regime i. A series of K-dimensional response vectors $\theta_{ki,1}, \ldots, \theta_{ki,h}$ predict the response of the endogenous variables.

$$\left.\frac{\partial E_t X_{t+h}}{\partial u_{k,t}}\right|_{s_t=\ldots=s_{t+h}=i} = \theta_{ki,h} \text{ for } h \geq 0 \tag{4}$$

Estimates of the response vectors can be derived by combining the parameter estimates of the Markov-switching unrestricted vector autoregression with the estimate of the matrix $\hat{A}_i$ obtained through identification restrictions.[3]

The first response vector measures the impact effect on endogenous variables of the k-th fundamental disturbance and is easily estimated. A one standard deviation shock to the k-th fundamental disturbance implies that the initial disturbance vector is $u_0 = (0, \ldots, 0, 1, 0, \ldots, 0)$, i.e. a vector of zeros apart from the k-th element which is one. Premultiplying this vector by the estimate of the regime-dependent matrix $\hat{A}_i$ as in equation (1) gives the impact responses.

The remaining response vectors can be estimated by solving forwards for the endogenous variables in equation (1). Equations (5) and (6) show the solution linking the estimated response vectors with estimated parameters.

$$\hat{\theta}_{ki,0} = \hat{A}_i u_0 \tag{5}$$

$$\hat{\theta}_{ki,h} = \sum_{j=1}^{\min(h,p)} \hat{B}_{ji}^{h-j+1} \hat{A}_i u_0 \text{ for } h > 0 \tag{6}$$

[3] Recall that in our example of the Sims (1980) recursive identification scheme the matrix $\hat{A}_i$ is identified by a Choleski decomposition of the estimated variance-covariance matrix $\hat{\Sigma}_i$.

2.5 Bootstrapping

It is possible to gauge the precision of the estimated response vectors by employing standard bootstrapping techniques. The technique involves creating artificial histories for the variables of the model and then submitting these histories to the same estimation procedure as the data. The artificial histories are created by replacing the parameters in the model with their estimated values, drawing residuals whose moments are determined by the estimated variance-covariance matrix, and then calculating the endogenous variables. Since the artificial histories are typically short samples their estimates will not coincide exactly with those from the original data. By creating a large number of artificial histories we can therefore make a bootstrapped approximation to the distribution of the estimated parameters.

In Markov-switching the bootstrapping is complicated by the presence of the hidden Markov-chain determining the regime. To create an artificial history it is first necessary to create a history for the regimes and to then use this to continue with the endogenous variables. The full procedure consists of five steps for each history.

1) **Create a history for the hidden regime s_t.** This can be done recursively using the definition of a Markov process (3) and replacing the exogenous transition matrix with its estimated value $\hat{P}$. At each time t we draw a random number from a uniform [0,1] distribution and compare it with the conditional transition probabilities to determine whether there is a switch in regime.

2) **Create a history for the endogenous variables.** Again this is done recursively, on the basis of equation (1). All parameters are replaced by their estimated values and new fundamental residuals are drawn from the normal distribution $u_t \sim N(0; I_K)$. Equation (1) can then be applied recursively using the artificial regime history created in step one.

3) **Estimate a Markov-switching vector autoregression,** using the data from the artificial history. Estimation gives bootstrapped estimates of the parameters $\{\tilde{v}_i; \tilde{B}_{1i}, \ldots, \tilde{B}_{pi}; \tilde{\Sigma}_i\}$ for $i = 1, \ldots, m$, the transition matrix $\tilde{P}$, and the smooth probabilities $\tilde{\xi}_{i,t} = \Pr(\tilde{s}_t = i)$ for $i = 1, \ldots, m$ and $t = 1, \ldots, T$.

4) **Impose identifying restrictions.** Applying the same restrictions as to the data provides bootstrapped estimates of the matrices $\tilde{A}_1, \ldots, \tilde{A}_m$.

5) **Calculate the bootstrapped estimates of the response vectors.** Substituting the new parameters $\tilde{B}_{1i}, \ldots, \tilde{B}_{pi}$ and $\tilde{A}_1$ into equations (5) and (6) gives bootstrapped estimates of the response vectors $\tilde{\theta}_{ki,0}, \ldots, \tilde{\theta}_{ki,h}$ for each regime $i = 1, \ldots, m$.

Applying the above five steps for a sufficiently large number of histories gives a numerical approximation to the distribution of the original estimates $\hat{\theta}_{1i}, \ldots, \hat{\theta}_{hi}$. In impulse response analysis this distribution forms the basis for adding confidence bands to the central estimate of the impulse response function.

3 An illustrative example

To illustrate the use of regime-dependent impulse response functions, we analyse how oil price shocks are transmitted to output and inflation in the US economy. Oil prices increased sharply during 1999 and 2000 to a level that had not been reached since the Gulf war. Yet, unlike during other previous periods of similar oil price increases, this time there was only little concern that the effects of this oil price shock would impinge on the US economy. Three reasons have been suggested why oil price shocks should be less of a concern nowadays than they used to be.[4] One explanation hinges on the "new economy" – the recent increase in productivity due to the development in information and production technologies is believed to have insulated the economy: the profitability of production now seems to be much less dependent on labour and non-labour input costs. The second reasoning proceeds in a similar way. The oil price shocks of the 1970s have led to massive improvements in the energy efficiency of the US economy. Therefore, overall economic performance is much less vulnerable to oil price increases than it used to be. A third explanation relates to the insulation of the economy by better monetary policy. With the Fed acting more aggressively to keep inflation expecations under control, oil price increases will not lead to a similar rise in inflation expectations, as has already been seen during the Gulf war oil price shock of the 1990s.

However, at the moment there are signs of a considerable slowdown in the US economy. It is therefore useful to investigate whether this could be a consequence of the oil price increase, and thus whether the old relationship between oil price increases and recessions has reappeared. To do so, we estimate a small VAR model including capacity utilisation, the consumer price index excluding energy prices and the spot price for Saudi Arabian light oil (in US$ per barrel). The data are monthly and span 336 observations from 1973:1 to 2000:12.

We begin by estimating an unrestricted Markov-switching VAR.[5] We assume the existence of two distinct regimes and allow intercepts, autoregressive parameters, variances and covariances to all switch between regimes. The lag length was chosen to be three, to ensure that the residuals are serially uncorrelated (see figures A1 and A2 in the appendix). This lag length is supported by one of the specification tests.[6] The results of our analysis remain basically unchanged, whether we use two or three lags.

	AIC	HQ	SC
lag length =4	–4.418	–3.998	–3.364
lag length =3	–4.609	–4.271	–3.763
lag length =2	–4.572	–4.317	–3.933

The estimation results support the following transition matrix for the two regimes.

[4] See, eg, Kliesen (2001).

[5] Estimations are performed using the MSVAR 0.99 package for Ox 2.10.

[6] Non-stationarity of the variables does not impose problems with the estimations. Our residuals are reasonably behaved, as shown in figures A1 and A2 in the appendix. The interpretability of test statistics on the regression parameters is achieved by using bootstrap techniques.

$$\hat{P} = \begin{bmatrix} 0.96 & 0.04 \\ 0.03 & 0.97 \end{bmatrix}$$

The regimes are estimated to be very persistent, with expected durations of 25 and 40 months respectively. The lower panel of Figure 1 reports the estimated smooth probabilities of being in regime 1. The US economy appears to have been in regime 1 up to 1986 and in regime 2 ever since. Whereas since 1986 no regime switch has occured, the regimes have been switching somewhat more frequently up to 1986.

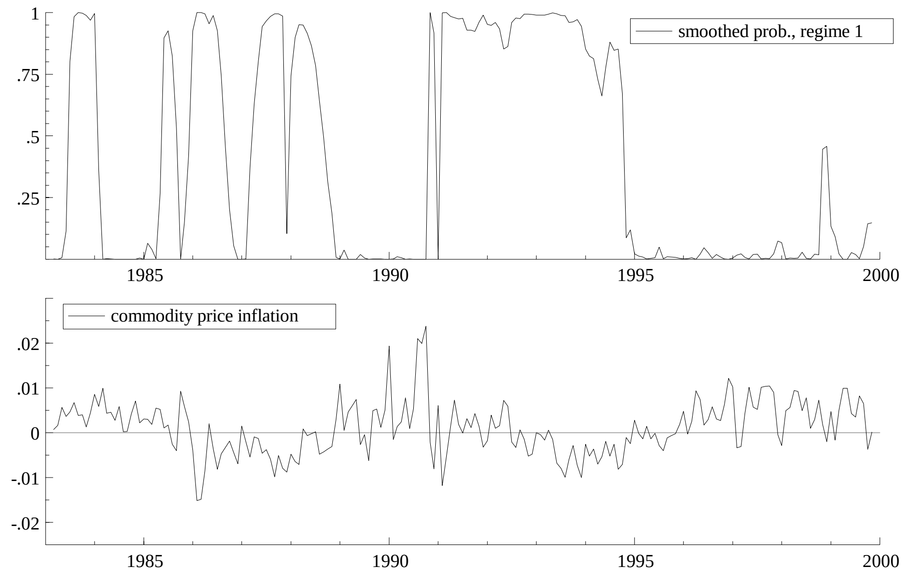

Figure 1: Capacity utilisation and smooth probabilities

The timing of the regimes already gives a first indication as to how they differ from each other. It is well known that the US economy has been much less volatile since mid 1980, and that oil prices themselves have increased their variability around that time.[7] The upper panel of figure 1 plots capacity utilisation, where the volatility of the series decreases strongly around 1985. However, the estimated model allowed all other parameters to switch as well, so it is not immediately obvious what other changes might have occured. A simple first step towards characterising the regimes is therefore to investigate which of the parameters in the model switch significantly between the regimes. Table 1 lists all parameters that switch at a 95% level of significance.

[7] On the decreased volatility in the US economy, see McConnell and Perez-Quiros (2000) and Kim and Nelson (1999). On the increased volatility of oil prices see, eg, Hamilton (1996).

Parameter	state 1	state 2	t-statistic
$\hat{b}_{111}$	1.355	1.032	2.36
$\hat{b}_{121}$	−23.424	59.762	−1.98
$\hat{b}_{222}$	0.103	−0.338	2.43
$\hat{b}_{223}$	−0.175	0.095	−2.19
$\hat{b}_{311}$	−0.249	0.875	−2.48
$\hat{b}_{321}$	9.032	412.83	−2.24
$\hat{b}_{323}$	−31.328	−571.636	3.29
$\hat{b}_{333}$	0.115	−0.107	2.13
$\hat{\sigma}_1^2$	0.449	0.169	3.85
$\hat{\sigma}_2^2$	5.3×10^{-6}	9.6×10^{-7}	4.97
$\hat{\sigma}_3^2$	0.449	5.665	−6.86

Table 1: Significantly switching parameters[8]

The most significant switches do indeed occur for the variances of all three variables. Capacity utilisation and prices have a lower variance in regime 2, but the variability of oil prices increases. On top of this, there are many more autoregressive parameters which change markedly across regimes. This makes us believe that regime dependent impulse responses will be a useful tool to draw a more complete picture of the differences between regimes. As in the example in Section 2.3., we identify the impulse responses by assuming a recursive structure for the model. The variables are ordered as capacity utilisation, CPI excluding energy, and the oil price. Positioning the oil price last in the vector autoregression implies that it can react to all variables, but is unable to affect capacity utilisation and prices within the same month. This is in line with the general assumption that prices which are determined in highly liqiuid markets like commodity exchanges react instananeously to the arrival of news, whereas output and consumer prices react more sluggishly.

The responses of all variables to an oil price shock are shown in Figure 2. Responses are regime-dependent: the left-hand diagrams correspond to regime 1 whilst the right-hand diagrams refer to regime 2. To make the two sets of responses comparable, we have normalised the size of the oil price shock to \$1 for both regimes. The dashed lines are one standard deviation confidence bands calculated on the basis of 1,000 bootstrap replications.[9]

[8] The notation is as follows: $\hat{b}_{ijk}$ denotes the parameter in equation i, on variable j, with lag k. $\hat{b}_{123}$ therefore denotes the parameter on the third lag of CPI in the capacity utilisation equation. The variances are written as $\hat{\sigma}_i^2$.

[9] A program in Ox 2.10 to calculate the impulse response functions and associated confidence bands is available from the authors on request.

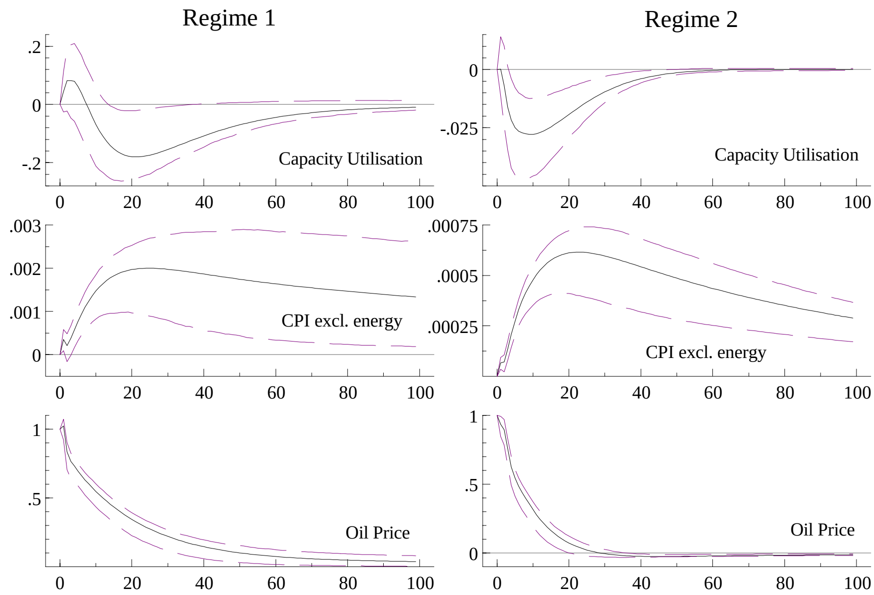

Figure 2: Responses to an oil price shock

In both regimes, the impulse responses show a significant effect of oil price shocks on capacity utilisation and prices. An increase in oil prices leads to a fall in production, and to inflationary pressures. A closer look reveals that there are quite remarkable differences across regimes. Firstly, the scale of the impulse responses differs strongly. An oil price shock of the same magnitude has much lower effects on output and prices in regime 2 than in regime 1. The maximum response of capacity utilisation is around 6 times as large in regime 1 than in regime 2, and approximately three times as large for prices.

Another difference in the responses across regimes is found in the timing of the effects. Wheras an oil price shock in regime 1 contracts capacity utilisation for nearly 20 months, the effects are much more short-lived in regime 2, where capacity utilisation starts to recover already after 10 months. The timing of the price response seems fairly similar across the two regimes. It is interesting to see that prices tend to fall back towards their old level in regime 2. The point estimates for regime 1 show a similar reversion, but due to the increasing width of the error bands this effect cannot be verified.

Oil prices themselves show a different pattern in their return to baseline. In regime 1, they are much more persistent than in regime 2. This might be an explanation why the output response is more protracted in this regime.

In order to evaluate whether the regime-dependence of our impulse responses conveys additional information, it is helpful to compare the results to those of a conventional VAR. Figure 3 reports the impulse responses to an oil price shock, which have been calculated in exactly the same way as before, but without allowing for Markov switching regimes. Obviously, the responses are an average of those for the two separate regimes in figure 2. Several points are worth mentioning here. Firstly, in order to gauge the effects of the 1999/2000 oil price shock on the US economy, a conventional VAR analysis would have been misleading. Not accounting for the differences across regimes results in impulse responses which underestimate the effects in regime 1, whereas they overestimate the consequences of an oil price shock in regime 2, the current regime. In a conventional VAR estimate, Deravi and Hegji (1992, p. 1) found that “a negligible percentage of inflation’s forecast error variance can be attributable to increases in the price of oil” – which can be due to mixed evidence over their sample period 1970 to 1990.

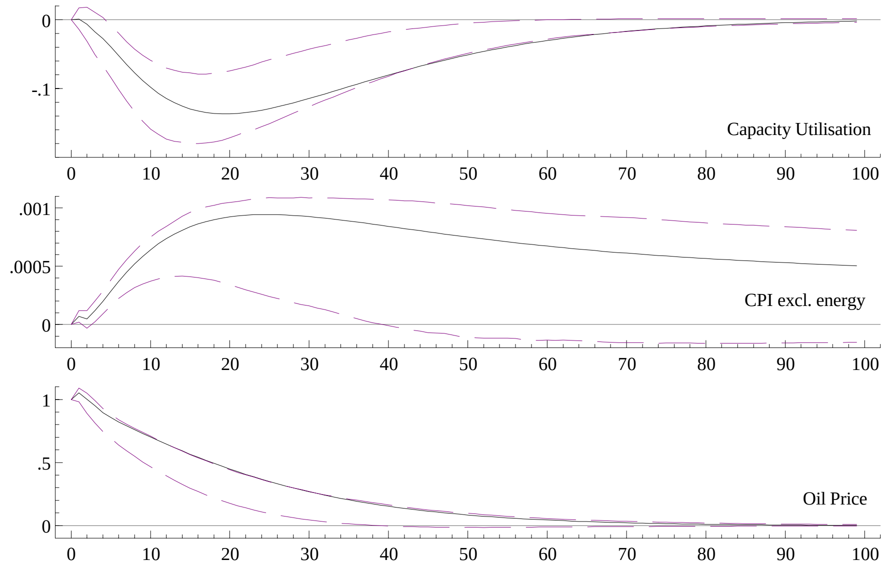

Figure 3: Responses to an oil price shock in a conventional VAR

Another interesting aspect of this figure is that the bootstrapped error bands are not only highly asymmetric, but sometimes do not even wrap the point estimate. This is a sign for seriously non-normal residuals in the estimation, a suspicion which can be substantiated by looking at the residuals in figure A3 in the appendix. The residuals in the conventional VAR suffer from ARCH effects and are by no means normally distributed. On the other hand, the Markov switching residuals are much better behaved, as can be seen by the much more symmetrically bootstrapped error bands and by looking at a plot of the residuals and according test statistics in figures A1 and A2.

The results of our regime dependent impulse responses are very much in line with earlier findings. Hutchinson (1993) estimates that an oil price shock of identical magnitude has only one fourth of the contractionary effect on GDP in the 1980s compared with the 1970s, and similarly so for inflation. Hutchinson achieves these results by splitting the sample in two subperiods. This means that he has to impose the regime shift, whereas we can identify it endogenously. Also, he has to estimate two different models where no direct testing across regimes is possible.

Additionally, our Markov-switching model allows for further regime shifts. We do not a priori restrict the US economy to stay in regime 2, as is done by splitting the sample. The fact that the model could not detect another regime shift for the latest oil price shocks is in itself interesting, because it means that (at least up to now) there are no signs that the US economy would react to this shock in the way it used to do in the 1970s and early 1980s.

Estimating the VAR within one model allows us to perform some tests on a very recent controversy around the stabilisation of the US economy. The marked reduction in volatility of many US macroeconomic variables in the 1980s has be supported by several explanations. On the one hand, it has been argued that this is due to "good luck", that the US economy has been hit less by adverse structural shocks.[10] Another explanation is that the volatility reduction came about because of technological advances that allowed firms to improve their inventory management ("good practices").[11] Thirdly, it has been argued that "good macroeconomic policies" have reduced the business cycle volatility.[12]

In a first experiment, we test for the "good luck" theory, taking it at face value. We use the actual identified shocks that were hitting the US economy, and feed them through the VAR model of regime 1. This way, we assume a constant structure of the economy, but allow for the reduced volatility of shocks that affect the economy. Figure 3 reports the results of this experiment: the solid line represents the actual time path of capacity utilisation, whereas the dotted line shows its predicted time path for a constant economic structure. Had it not been for the change in the structure of the economy, there would have been a recession in 1999 equal in magnitude to that of 1991 – the last official NBER recession. The economic structure we have estimated for regime 2 seems to be in a better shape to cope with the structural shocks that the US economy actually experienced in the last years. Without another ingredient from the "good macroeconomic policy" or "good practices" explanations, the observed reduction in output volatility would at least not have been as strong as it was.

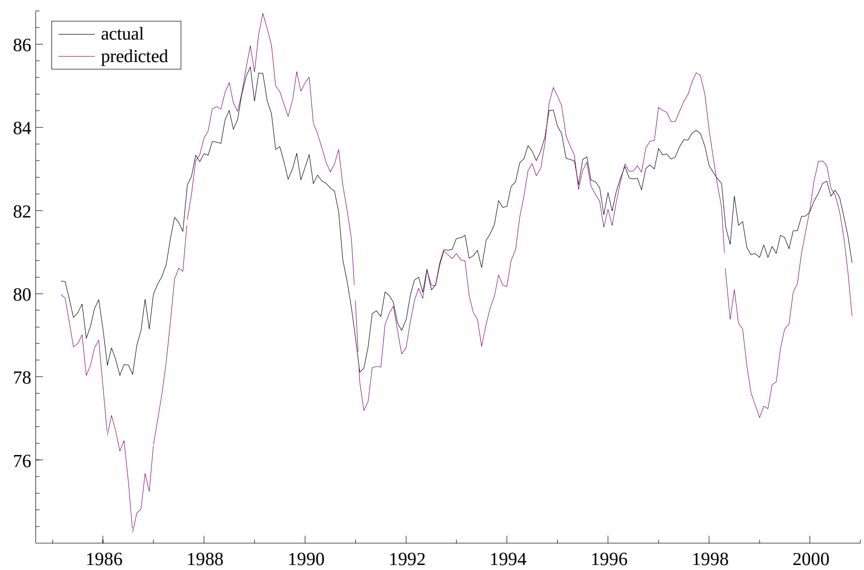

Figure 4: Predicted and actual capacity utilisation

In a similar vein, we conduct a second experiment that forces the oil price shocks in the regime 2 to be identical to those experienced in regime 1. The shocks differ with respect to their persistency, with those in regime 1 being

[10] For an overview of the various hypotheses, see Ahmed et al. (2001).
[11] See McConnell and Perez-Quiros (2000).
[12] See Clarida, Gali and Gertler (2000).

much more persistent than those in regime 2. Assuming that the persistency of shocks is exogenous (hence not determined by, eg, the monetary policy response[13]), what if one of these "bad", persistent shocks would hit the US economy today, in regime 2? Figure 5 answers this question. Naturally, the response of output and prices would be more protracted than they are otherwise in regime 2. However, the magnitude of responses still does not compare to those found in regime 1. We can therefore conclude that the severe output contractions and price increases that follow oil price shocks in regime 1 are not only a consequence of the persistency of the shocks; shocks with the same pattern in regime 2 do still cause much less disruption in output and prices.

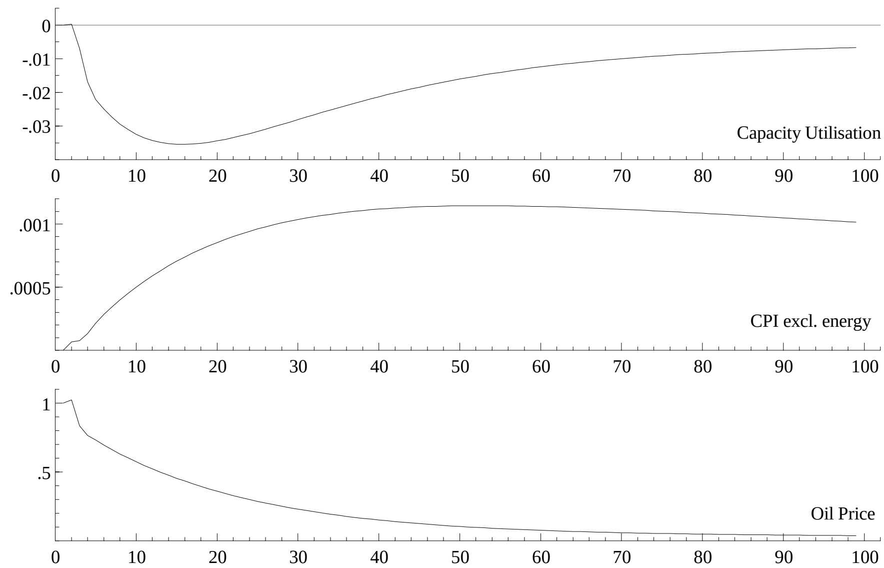

Figure 5: Effects of a regime 1-type oil shock hitting in regime 2

4 Conclusions

This paper has developed a tool for the analysis of Markov switching VAR models. Very often, the Markov regimes in models where many parameters are allowed to switch are very hard to interpret. Our regime dependent impulse responses can help in this respect. By imposing structure spearately on each regime, it is possible to derive two sets of impulse response functions that characterise the different patterns of the economy when in different regimes.

Our illustrative example shows how regime-dependent impulse response functions can give valuable insights into the characteristics of regimes in a Markov-switching model. For the US economy, we are able to identify two regimes which track the recently documented change in volatility of the US economy. The two regimes do not only differ with respect to the variances of variables, but also with respect to several autoregressive parameters. The model shows that the two regimes are distinctly different as to how they absorb

[13]This might be a reasonable assumption for oil price shocks, which most probably depend much more on OPEC's behaviour.

oil price shocks. Recently, these shocks have had a much less disruptive effect on the US economy. Since we can estimate the two regimes within one model, we can show that these effects cannot be explained by "good luck" alone.

Our approach is applicable to a wide variety of situations. If the estimated regimes can be characterised directly in terms of individually switching parameters then regime-dependent analysis can give a deeper understanding of the economy. For example, Ehrmann (2000) identifies business cycle regimes in Germany and then uses regime-dependent analysis to demonstrate that monetary policy tends to be more effective when the economy is in recession. Alternatively, it may be almost impossible to characterise the regimes on the basis of individual parameters. For example, Ellison and Valla (2000) find evidence of persistent regimes in the G7 economies but the regimes have no natural interpretation in terms of individual switching parameters. However, the regime-dependent impulse response functions reveal that the regimes are associated with periods of high and low monetary policy effectiveness.

References

[1] Ahmed, S. – Levin, A. – Wilson, B.A. (2001) **Recent US Macroeconomic Stability: Good Luck, Good Policies, or Good Practices?** Paper presented at the 2001 AEA meeting, New Orleans.

[2] Blanchard, O.J. – Quah, D. (1989) **The Dynamic Effects of Aggregate Demand and Supply Disturbances.** American Economic Review 79: 655–673.

[3] Clarida, R. – Gali, J. – Gertler, M. (2000) **Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory.** Quarterly Journal of Economics CXV: 147–180.

[4] Deravi, K. – Hegji, C.E. (1992) **The Inflationary Impact of Oil Price Shocks: A Vector Autoregressive Study.** Review of Financial Economics 2: 1–16.

[5] Ehrmann, M. (2000) **Firm Size and Monetary Policy Transmission – Evidence from German Business Survey Data.** European Central Bank Working Paper No. 21.

[6] Ellison, M. – Valla, N. (2000) **Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions.** forthcoming in Journal of Monetary Economics.

[7] Hamilton, J.D. (1989) **A New Approach to the Economic Analysis of Non Stationary Time Series and the Business Cycle.** Econometrica 57: 357–384.

[8] Hamilton, J.D. (1990) **Analysis of Time Series Subject to Changes in Regime.** Journal of Econometrics 45: 39–70.

[9] Hamilton, J.D. (1996) **This is what happened the oil price-macroeconomy relationship.** Journal of Monetary Economics 38: 215–220.

[10] Hutchinson, M.M. (1993) **Structural Change and the Macroeconomic Effects of Oil Shocks: Empirical Evidence from the United States and Japan.** Journal of International Money and Finance 12: 587–606.

[11] Kim, C.-J. – Nelson, C.R. (1999) **Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle.** Review of Economics and Statistics 81: 608–616.

[12] King, R.G. – Plosser, C.I. – Stock, J.H. – Watson, M.W. (1991) **Stochastic Trends and Economic Fluctuations.** American Economic Review 81: 819–840.

[13] Kliesen, K.L. (2001) **Rising Oil Prices and Economic Turmoil. Must They Always Go Hand in Hand?** The Regional Economist, Federal Reserve Bank of St. Louis, January.

[14] Koop, G. – Pesaran, M.H. – Potter, S.M. (1996) **Impulse Response Analysis in Nonlinear Multivariate Models.** Journal of Econometrics 74: 119–147.

[15] Krolzig, H.-M. (1997) **Markov Switching Vector Autoregressions: Modelling, Statistical Inference and Application to Business Cycle Analysis.** Lecture Notes in Economics and Mathematical Systems No. 454. Berlin: Springer-Verlag.

[16] Krolzig, H.-M. – Toro, J. (1999) **A New Approach to the Analysis of Shocks and the Cycle in a Model of Output and Employment.** EUI Working Paper ECO No. 99/30.

[17] McConnell, M.M. – Perez-Quiros, G. (2000) **Output Fluctuations in the States: What has Changed Since the Early 1980's?** American Economic Review 90: 1464–1476.

[18] Sims, C.A. (1980) **Macroeconomics and Reality.** Econometrica 48: 1–48.

A Residual analysis

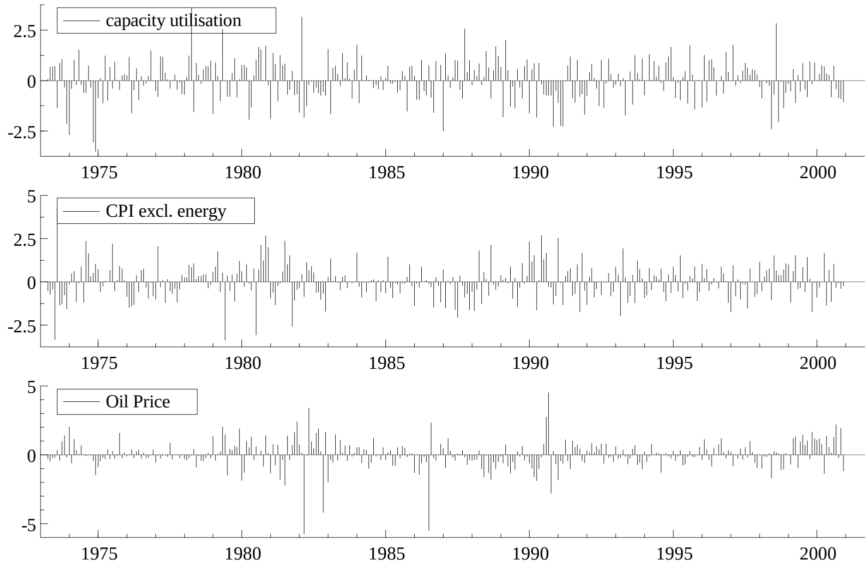

Figure A1: Residuals in the Markov Switching VAR

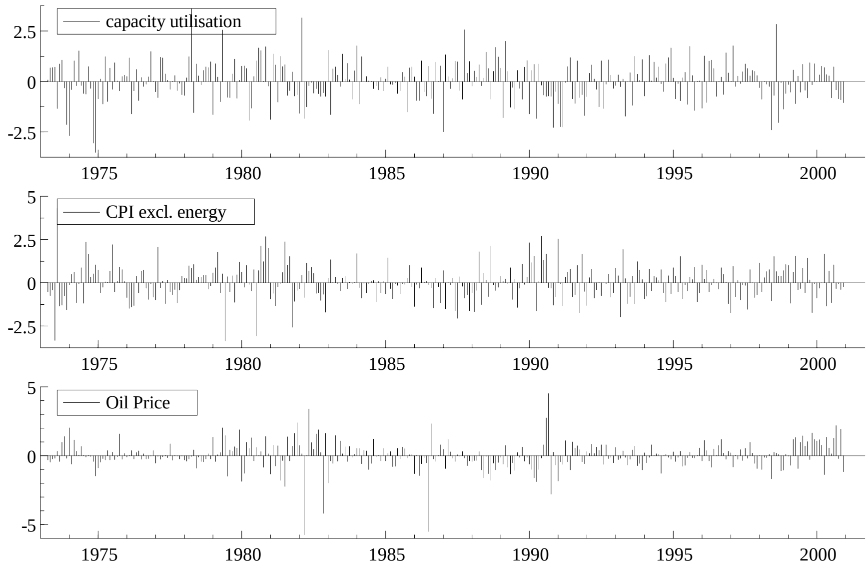

Figure A1: Residuals in the Markov Switching VAR

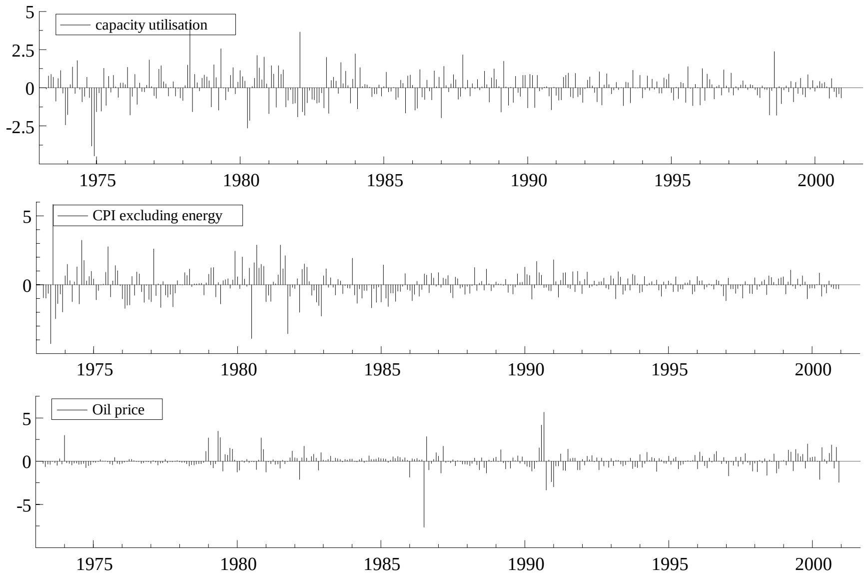

Figure A3: Residuals of a conventional VAR

BANK OF FINLAND DISCUSSION PAPERS

ISSN 0785-3572, print; ISSN 1456-6184, online

1/2001 Risto Herrala **An assessment of alternative lender of last resort schemes.** 2001. 28 p. ISBN 951-686-701-4, print; ISBN 951-686-702-2, online. (TU)

2/2001 Esa Jokivuolle – Karlo Kauko **The new basel accord: some potential implications of the new standards for credit risk.** 2001. 23 p. ISBN 951-686-703-0, print; ISBN 951-686-704-9, online. (RM)

3/2001 Mika Kortelainen **Actual and perceived monetary policy rules in a dynamic general equilibrium model of the euro area.** 2001. 76 p. ISBN 951-686-705-7, print; ISBN 951-686-706-5, online. (TU)

4/2001 Martin Ellison **Stabilisation bias in monetary policy under endogenous price stickiness.** 2001. 20 p. ISBN 951-686-707-3, print; ISBN 951-686-708-1, online. (TU)

5/2001 Erkki Koskela – Rune Stenbacka **Equilibrium unemployment with credit and labour market imperfections.** 2001. 35 p. ISBN 951-686-709-X, print; ISBN 951-686-710-3, online. (TU)

6/2001 Jarmo Pesola **The role of macroeconomic shocks in banking crises.** 2001. 61 p. ISBN 951-686.713-8, print; ISBN 951-686-714-6, online. (TU)

7/2001 Anssi Rantala **Does monetary union reduce employment?** 2001. 36 p. ISBN 951-686-715-4, print; ISBN 951-686-716-2, online. (TU)

8/2001 Tuomas Välimäki **Fixed rate tenders and the overnight money market equilibrium.** 2001. 72 p. ISBN 951-686-717-0, print; ISBN 951-686-718-9, online (TU)

9/2001 Morten L. Bech – Kimmo Soramäki **Gridlock resolution in interbank payment systems.** 2001. 34 p. ISBN 951-686-719-7, print; ISBN 951-686-720-0, online (RM)

10/2001 Antti Ripatti – Jouko Vilmunen **Declining labour share – Evidence of a change in the underlying production technology?** 2001. 43 p. ISBN 951-686-721-9, print; ISBN 951-686-722-7, online (TU)

11/2001 Michael Ehrmann – Martin Ellison – Natacha Valla **Regime-dependent impulse response functions in a Markov-switching vector autoregression model.** 2001. 24 p. ISBN 951-686-723-5, print; ISBN 951-686-724-3, online (TU)