H Mumtaz and P Surico
The transmission mechanism in good and bad times
Article

Mumtaz, H and Surico, P

(2015)

The transmission mechanism in good and bad times.

International Economic Review, 56 (4). pp. 1237-1260. ISSN 0020-6598

DOI: https://doi.org/10.1111/iere.12136

Wiley
https://onlinelibrary.wiley.com/doi/abs/10.1111/ie...

The Transmission Mechanism in Good and Bad Times*

Haroon Mumtaz
Queen's Mary University

Paolo Surico
London Business School and CEPR

April 2014

Abstract

Does the transmission of economic policies and structural shocks vary with the state of the economy? We answer this question using a strategy based on quantile regressions, which account for both endogeneous regressors and state-dependent parameters. An application to real activity and the interest rate reveals pervasive asymmetries in the propagation mechanism of economic disturbances across good and bad times. During periods in which real activity is above its conditional average, the estimates of the degree of forward-lookingness and interest rate semi-elasticity are significantly larger (in absolute value) than the estimates associated with below-average periods. Results are robust to alternative estimation strategies to model state-dependent parameters.

JEL classification: E21, E32, E52.
Keywords: state-dependence, asymmetric transmission mechanism, consumption.

*We are grateful to the Editor, Jesús Fernández-Villaverde, and three anonymous referees for very useful comments and suggestions. Financial support from the European Research Council Starting Independent Grant - Agreement 263429- is gratefully acknowledged. Correspondence: h.mumtaz@qmul.ac.uk; psurico@london.edu.

1 Introduction

The great recession of 2007-09 has sparkled renewed interest in the extent to which monetary and fiscal policies can stimulate the economy during bad and good times. On the one hand, a number of empirical contributions have used Vector AutoRegressions (VAR) or Augmented Distributed Lag (ADL) models to show that stabilization policies may have asymmetric effects over the business cycle (see for instance Auerbach and Gorodnichenko, 2011, for fiscal policy and Tenreyro and Thwaites, 2013, for monetary policy). On the other hand, Dynamic Stochastic General Equilibrium (DSGE) models typically assume that the effects of the short-term interest rate and government spending on real activity during expansions are as large as the effects during contractions.[1]

In this paper, we assess the evidence for asymmetries in the consumption-interest rate relationship, which we argue can be interpreted as an Investment-Saving (IS) curve. Relative to the more reduced-form evidence from VAR and ADL models, our strategy accounts for endogenous regressors by using an instrumental variables method. Relative to the more structural evidence from DSGE models, our single equation approach allows us to model explicitly the link between parameters evolution and the state of the economy in a way that is both flexible and computationally feasible.

We propose an empirical model where the coefficients are allowed, but not required, to vary with the (unobserved) state of the economy, which is endogenously determined within the estimation method. Our method is based on Instrumental Variable Quantile Regressions (IVQR) which are designed to handle simultaneously endogenous regressors and state-dependent parameters. To illustrate the potential of using quantile regressions on time-series data, we also present a time-varying coefficient interpretation of our estimates, which complements recent evidence of parameter instability in DSGE models. Furthermore, we show that the evidence against a linear IS curve specification is robust to using instrumental variable Markov-switching and threshold models, though these estimates are

[1] While recent advances have made it possible to solve structural models that feature a zero bound for the nominal interest rate (see Fernández-Villaverde et. al., 2012), we are not aware of contributions that have estimated this type of nonlinearity in the context of a DSGE model.

less accurate than the IVQR estimates.

Our findings on post-WWII U.S. data can be summarized as follows. First, there is strong evidence for state-dependence in the weight of the forward-looking component of the IS curve: periods in the bottom (top) 10% of the conditional distribution of real activity are characterized by fully backward- (forward-) lookingness. Second, there is also a significant extent of nonlinearity in the estimates of the interest rate semi-elasticity with values around -0.005 below the $70th$ percentile and values between -0.02 and -0.08 above that. This suggests that monetary policy is more effective during periods of conditionally high consumption/output. Third, a constant parameter consumption-interest rate relationship significantly over-estimates (under-estimates) the degree of forward-lookingness and the interest rate semi-elasticity during periods of low (high) real activity conditional on covariates. Fourth, mapping the state dependent estimates into time-varying coefficients reveals that periods of *conditionally* low (high) real activity coincide with periods of *unconditionally* below-trend (above-trend) consumption/output. The implication is that our findings can be equivalently cast in terms of phases of the business cycle. Fifth, our results are robust to employing alternative *(i)* filters to isolate cyclical components, *(ii)* measures of real activity, *(iii)* instrument sets, *(iv)* specifications of the lag structure in the transmission mechanism and *(v)* strategy to identify the unanticipated component of movements in the interest rate.

Finally, our findings may help reconcile conflicting evidence from earlier contributions. Estimates of the degree of forward-lookingness in IS curve specifications range from values not statistically different from one (Ireland, 2004) to values not significantly different from zero (Fuhrer and Rudebusch, 2004). The semi-elasticity of interest rate is often statistically insignificant (Lindé, 2005), and when it is not, the point estimates are so small as to imply only modest effects through the transmission of structural shocks or economic policies (Dennis, 2009). The evidence on state-dependent parameters presented in this paper may thus offer a way to rationalize the seemingly contrasting estimates that are based on linear models and different U.S. post-WWII samples.

The paper is organized as follows. The empirical model is presented in section 2. In section 3, we lay out the estimation method to explore state-dependence and account for endogeneity. Section 4 introduces the data and the instrument sets. In section 5, we present the main results of the paper, the robustness to alternative methods to model state-dependent parameters as well as a time-varying coefficient interpretation of the quantile regression estimates. A sensitivity analysis is offered in Section 6. Section 7 compares the forecasting performance of the IVQR model relative to alternative strategies to deal with parameter instability. The appendices report a Montecarlo analysis to assess the small sample bias associated with the quantile regression method presented in section 3 as well as convergence results for our Markov chain Montecarlo algorithm.

2 A state-dependent parameter transmission mechanism

In this section, we lay out a flexible empirical model that will be used in Section 3 to investigate asymmetries in the transmission mechanism. In Section 4, we confirm the robustness of the empirical findings to using two popular methods to estimate state-dependent parameters, namely Markov-switching and threshold models.

Our approach builds on Koenker and Xiao (2006) who develop asymptotic theory and inference tools for quantile autoregressive models. More specifically, suppose that a cyclical measure of real activity c_t evolves according to the following rule $F(\cdot)$:

$$c_t = F\left(c_{t-1}, c_{t+1}, i_t..i_{t-h}, \pi_{t+1}..\pi_{t-h+1}, u_t\right) = F(i_t, d_t) = F(D_t) \tag{1}$$

where i denotes the nominal interest rate, π is inflation, and d refers to leads and lags of consumption and inflation. The unobserved state of the economy, u_t, is the source of heterogeneity.

Our aim is to estimate the shape of (1) using quantile regressions (QR). Above all, this will not assume that the relationship between consumption, its leads and lags and the interest rate is linear. Furthermore, we will consider the possibility that both the ex-ante real rate and future consumption are endogenous variables. The QR approach treats the measure of real activity as a potential latent outcome. It is latent because, given the covariates i_t and d_t, the observed outcome in each unit of time t is only one of the possible realizations in the admissible space of outcomes. The quantiles, Q_τ, of the potential outcome distributions conditional on covariates are denoted by:

$$Q_\tau \left(c_t | i_t, d_t\right) \quad \text{with } \tau \in (0,1). \tag{2}$$

The effect of a change in the real rate on different points of the marginal distribution of the potential outcome is given by:

$$QTE_\tau = \frac{\partial Q_\tau \left(c_t | i_t, d_t\right)}{\partial r} \tag{3}$$

where r is the real rate of interest. The quantile treatment model can then be written as:

$$c_t = q\left(i_t, d_t, u_t\right) \quad \text{with } u_t | i_t, d_t \sim U\left(0,1\right). \tag{4}$$

where $q\left(i_t, d_t, u_t\right) = Q_\tau \left(c_t | i_t, d_t\right)$. Note that we can always work with a suitable monotonic transformation of the underlying measure of unobserved heterogeneity such that u_t is a rank variable, i.e. it measures the relative ranking of states of the economy in terms of potential outcomes. According to this interpretation, QTE_τ measures the causal effect of the real rate on real activity, holding the latent state fixed at $u_t = \tau$.

The model that we propose (and specify explicitly below) is a generalisation of the Quantile Autoregressive model (QAR) introduced in Koenker and Xiao (2006), who consider the following QAR(1) specification:

$$Q_\tau(c_t|\cdot) = \alpha\left(\tau\right) + \mu\left(\tau\right) c_{t-1} \tag{5}$$

where $\mu\left(\tau\right) = \min\left[\lambda_0 + \lambda_1 \tau, 1\right]$, $\lambda_0 \in (0,1)$ and $\lambda_1 > 0$. At higher values of the conditional quantiles, the QAR model implies that c_t displays behaviour consistent with a persistent AR process, while quicker mean reversion occurs at lower conditional quantiles. The QAR model is, therefore, able to describe asymmetric persistence in c_t and provides a useful approximation to non-linear dynamics.

Galvao, Montes-Rojas and Park (2009) extend the QAR model to include exogenous regressors on the right hand side of equation (5). In other words, they consider models of the following form:

$$Q_\tau(c_t|\cdot) = \alpha(\tau) + \mu(\tau)\, c_{t-1} + \varpi(\tau)\, x_t \tag{6}$$

The Quantile Autoregressive distributed lag model (QARDL) in equation (6) accounts for the impact of the exogenous covariates x_t at different values of τ, as well as allowing for the possibility of non-linear dynamics.[2]

While the QARDL model can be used to model reduced form relationships, the exogeneity of x_t rules out applications in a more structural setting such as the one considered in equation (1). Hence, following Chevapatrakul, Kim and Mizen (2009), this paper extends the QARDL model further by explicitly allowing for endogenous covariates. The model that we consider assumes that the empirical specification of the conditional τ-*th* quantile distribution of c_t takes the following form:

$$Q_\tau(c_t|\cdot) = \alpha(\tau) + [1-\mu(\tau)]\, c_{t-1} + \mu(\tau)\, c_{t+1} + \beta(\tau)\left[\frac{1}{\kappa}\sum_{j=0}^{\kappa-1}\left(i_{t+j+m} - \pi_{t+j+m+1}\right)\right] \tag{7}$$

It seems natural to assume that the real interest rate and future consumption may be endogenous for current consumption. Furthermore, the recent evidence discussed in the introduction (and based on VAR/ADL models) suggests that the impact of c_{t+1} and (i_{t+j}-π_{t+j+1}) on c_t tend to vary over the business cycle. A standard QARDL specification that ignores this dependence may result in biased estimates of (7).

The choice of a suitable set of instruments Z_t can provide a source of variation in the endogenous regressors that is independent from the latent state. To the extent that lagged values of c_t, i_t and π_t are uncorrelated with the forecast errors of variables dated at future time periods, then they can be used as valid instruments. We explore various instrument sets below and show that the results are robust across different choices for Z_t.

[2] The authors apply this model to UK house price returns and show that the policy interest rate has a larger negative impact at lower conditional quantiles while the effect of real GDP is smaller at medium values of the conditional quantile.

Given Z_t, the Instrumental Variable Quantile Regression (IVQR) model of Chernozhukov and Hansen (2005) can be used to estimate the parameters of interest. Consider the process for real activity:

$$c_t = q\left(i_t, d_t, u_t\right) \quad \text{with } u_t|Z_t, d_t \sim U\left(0,1\right). \tag{8}$$

where

$$Prob[c \leq q\left(i, d, \tau\right)|Z] = Prob[U \leq \tau|Z, d] = \tau \quad \text{for each } \tau \in (0,1).$$

The IVQR estimator proposed by Chernozhukov and Hansen (2005) relies on two main assumptions: independence and rank invariance/similarity. The first assumption is shared by any instrumental variable estimator and requires that the instruments are uncorrelated with the error term. Interpreting the residuals u_t in (1) as structural shocks (in a combination for instance of preference and fiscal shocks), the assumption of independence corresponds to assume that lagged values of the interest rate, consumption and inflation do not predict demand shocks. As for rank invariance, this requires that the ranking of the unobserved heterogeneity u_t does not vary with potential treatment states. In our application, this amounts to assume that a month whose particular level of interest rate (say below-average) leads to a state of the economy at the τ^{th} percentile of the conditional distribution of (positive) consumption growth would have had consumption growth still at the τ^{th} percentile of the conditional distribution of (negative) consumption growth if another level of the interest rate (say above-average) had instead induced that outcome (i.e. a consumption contraction). In the weaker form of rank similarity, it suffices that the ranks of the states of the economy induced by a particular shock do not to vary systematically with the level of the interest rate or with future consumption. It is important to emphasize that rank similarity does not imply that the states of the economy should not vary with the level of interest rate. Rather, it amounts to the weaker requirement that the ranking in the conditional consumption distribution induced by a shock of a given size is not systematically different from the ranking that the same shock would have generated to if the interest rate was at a different level.

3 Estimation

The parameters of the model in equation (8) can be estimated by solving the following optimisation problem

$$\min_{\Xi}\left\|\frac{1}{T}\sum_{t=1}^{T}\left[1\left(c_t-D_t\Xi\right)-\tau\right]Z_t\right\| \tag{9}$$

where $1(\cdot)$ is an indicator function that takes value one if $(c_t - D_t\Xi) \leq 0$ and zero otherwise and $\Xi = \{\alpha, \mu, \beta\}$ is a $k \times 1$ vector of model parameters.

The objective function in equation (9) is not straightforward to minimise given the discontinuity introduced by the indicator function. Chernozhukov and Hansen (2005) show, however, that minimisation of (9) is equivalent to considering the following Quantile regression objective function

$$Q_\tau(.)=\frac{1}{T}\sum_{i=1}^{T}\rho_\tau\left\{c_t\text{-}\{\alpha(\tau)+[1\text{-}\mu(\tau)]\,c_{t\text{-}1}+\mu(\tau)\,c_{t+1}+\beta(\tau)\,[\frac{1}{\kappa}\sum_{j=0}^{\kappa-1}(i_{t+j+m}\text{-}\pi_{t+j+m+1})]\}\text{-}\lambda Z_t\right\}$$

where $\rho_\tau(u)=(\tau\text{-}1(u<0))\,u$.

The coefficients α and λ can be estimated as $(\alpha(\tau,\mu,\beta),\lambda(\tau,\mu,\beta))=\min_{\alpha,\lambda} Q_\tau(.)$. The coefficients on the endogenous variables μ and β are estimated as those values that make $\hat{\lambda}(\tau,\mu,\beta)$ as close as possible to zero. This latter step involves searching over a grid of possible values for the coefficients. The final parameter vector is given by $\left(\hat{\mu}(\tau),\hat{\beta}(\tau),\hat{\alpha}\left(\hat{\mu}(\tau),\hat{\beta}(\tau),\tau\right)\right)$. The main attraction of this inverse QR approach is its simplicity and the fact that it only requires a series of standard QR regressions to be run at each point in the grid. However, the need for a grid search places a practical limit on the number of endogenous variables in the model. Chernozhukov and Hansen (2005) report that this estimator works well when the number of exogenous regressors is large relative to the number of endogenous covariates.

In our application, instead, the number of endogenous regressors is relatively large and application of the inverse QR method requires a multi-dimensional grid search. We found that in our setting the estimates of the parameters were highly dependent on the limits and the length of the grid. Therefore, we consider an alternative estimation strategy based on a Markov chain montecarlo (MCMC) approach introduced for moment based estimators in Chernozhukov and Hong (2003). The authors define the quasi-posterior of the parameters of moment based models as

$$P_n(\Xi) = \frac{\exp(L_n(\Xi))\,\pi(\Xi)}{\int \exp(L_n(\Xi))\,\pi(\Xi)} \tag{10}$$

where $\pi(\Xi)$ is a prior density and $L_n(\Xi)$ is defined as

$$L_n(\Xi) = -\frac{1}{2}\left(\left(\frac{1}{\sqrt{N}}\sum_{t=1}^{N} m_t(\Xi)\right)' W_n(\Xi)\left(\frac{1}{\sqrt{N}}\sum_{t=1}^{N} m_t(\Xi)\right)\right) \tag{11}$$

with $m_t(\Xi) = (\tau - 1(c_t - D_t\Xi))\,z_t$ and $W_n(\Xi)$ is a weighting matrix. Chernozhukov and Hong (2003) set out the conditions under which a random walk Metropolis-Hastings (MH) algorithm provides valid point estimates and confidence intervals for Ξ. Note that unlike the grid search approach in Chernozhukov and Hansen (2005), the MCMC approach can easily be used for applications involving a moderate to large number of endogenous regressors.

In our application, a prior distribution is not explicitly specified for the parameters and a data driven approach is adopted. However, some of the parameters are required to lie within the bounds implied by economic theory. These are reported in Table 1 and they are imposed by assigning an acceptance probability of zero to any draw that violates these bounds. The steps of the MH algorithm (implemented for a given value of τ) are summarised in algorithm 1.

Following Chernozhukov and Hong (2003) a random walk is used as the candidate density $q\left(\Xi^{new} \diagdown \Xi^{old}\right)$:

$$\Xi^{new} = \Xi^{old} + P'\varepsilon \tag{12}$$

where ε is a $k \times 1$ vector from the standard normal distribution and P is a scaling matrix.[3] The model is estimated using 1, 100, 000 MCMC iterations discarding the first 100, 000 as

[3] More specifically, we conduct a grid search for the scaling parameter $\hat{c}$ over M discrete intervals and store the acceptance rate $\hat{\alpha}$ associated with each run of the MCMC algorithm. We then estimate the following regression model: $\hat{\alpha} = \hat{a}_0 + a_1\hat{c} + a_2\hat{c}^2 + a_3\hat{c}^3 + e$ and find the value of c that solves numerically the polynomial $\hat{a}_0 + \hat{a}_1 c + \hat{a}_2 c^2 + \hat{a}_3 c^3 = \alpha^{\text{target}}$ where α^{target} is the desired acceptance rate. The value of c is used as the scaling factor in the main MCMC run. We set $\alpha^{\text{target}} = 0.3$ and this choice delivers an acceptance rate between 20% and 40%.

Algorithm 1 The random walk Metropolis Hastings algorithm for the QR model

1. **Starting values for the model parameters:** Let $i = 0$. Initialise the draws of the parameters by specifying the initial value of the parameters Ξ^0, set $\Xi^{old} = \Xi^0$ and evaluate the posterior at Ξ^{old} using equation (10). In our application, we set $\Xi^0 = \Xi^{QR}$, where Ξ^{QR} is the standard quantile regression estimator.

2. **Candidate draw:** Draw a candidate value for the parameters Ξ^{new} using the proposal density $q\left(\Xi^{new} \diagdown \Xi^{old}\right)$

3. If Ξ^{new} violates the bounds in Table 1 set $i = i + 1$ and go to step 2. Otherwise go to step 4.

4. **Posterior evaluation**: Evaluate the posterior at Ξ^{new} via equation (10).

5. **Accept/Reject**: Calculate the acceptance probability $accept = \frac{P_n(\Xi^{new})/q\left(\Xi^{new} \diagdown \Xi^{old}\right)}{P_n\left(\Xi^{old}\right)/q\left(\Xi^{old} \diagdown \Xi^{new}\right)}$. Draw a scalar u from the standard uniform distribution. If $accept > u$, set $\Xi^{old} = \Xi^{new}$ and $P_n\left(\Xi^{old}\right) = P_n\left(\Xi^{new}\right)$. Otherwise retain the old parameter draw and the corresponding posterior.

6. **Iterate**: Set $i = i + 1$. Go to step 2 and repeat until $i = R$.

---

burn-in. We retain every $50th$ draw of the remaining $1,000,000$ iterations. As shown in Appendix A, this procedure leads to a chain where the diagnostic statistics of Gelman and Rubin (1992) suggest convergence.

4 Data and instruments

The data were collected from the website of the Federal Reserve Bank of St. Louis. The nominal rate is the three month treasury bill rate. Inflation is measured as the first difference in the logarithm of the personal consumption expenditure (PCE) deflator. In the baseline case, real consumption is measured as the logarithm of the personal expenditure on nondurable goods and services, deflated by the PCE deflator and divided by the civilian noninstitutional population. In one of the sensitivity analyses, we will also consider durable consumption and the monthly estimates of real GDP in Stock and Watson (2010). All series but population are seasonally adjusted at the source.

The Montecarlo analysis in Appendix B reveals that the small sample bias associated with the IVQR method can be large using either a sample size or a sample frequency leading to about 200 observations. In contrast, it appears far more muted using as many as 700 observations. This result encourages us to work with monthly rather than quarterly data as for the latter there would be at most 250 data points available over the post-WWII period. On the other hand, monthly consumption data are available since 1959:1 but they can be extended back to 1948:5 by linearly interpolating quarterly observations. This implies that our sample, which ends in 2010:7, includes about 750 data points.

The cyclical component of real activity c_t is constructed using five methods: (1) HP filter, (2) Band-Pass filter, (3) linear de-trending, (4) quadratic de-trending and (5) the approximated de-trended method proposed by Cogley and Sargent (2005) where the trend c_t^* is computed using the recursion $c_t^* = c_{t-1}^* + 0.075(c_t - c_{t-1}^*)$. The data are displayed in Figure 1.

Our benchmark instrument set includes lags of the cyclical measure of consumption, inflation and the nominal interest rate: $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$. We also consider four alternative sets which include different lags of the endogenous variables, namely $Z1_t = (c_{t-2}..c_{t-6}, i_{t-2}..i_{t-6}, \pi_{t-2}..\pi_{t-6}, 1)$; $Z2_t = (c_{t-2}..c_{t-12}, i_{t-2}..i_{t-12}, \pi_{t-2}..\pi_{t-12}, 1)$; $Z3_t = (c_{t-1}, i_{t-1}, \pi_{t-1}, 1)$; $Z4_t = (c_{t-2}, i_{t-2}, \pi_{t-2}, 1)$.

While the results are similar across the four alternative groups, the benchmark instrument set is preferred when considering instrument strength. The first row of Table 2 shows that the multivariate F-statistics for the baseline instrument set are consistently larger than the critical value of 9.5 determined by Stock and Yogo (2003).[4] The entries for Z_t compare favourably to the entries in the third row of Table 2 for the alternative instrument sets used in the sensitivity analysis.

[4]This critical value depends on the number of endogenous variables, $d = 3$, the number of included and excluded regressors, $k = 5$, and the bias relative to OLS, $b = 0.05$, (see Stock and Yogo, 2003).

5 Results

In this section, we present our main results. These are organized around the different cyclical measures of consumption and the benchmark instrument set. In the following section, we assess the sensitivity of our results to changes in the model specification as well as in the instrument set, the identification strategy and the measure of real activity. An important finding of this section is that the evidence in favour of state dependence in the consumption-interest rate relationship is robust to using three alternative estimation methods: quantile regressions, markov-switching and threshold models.

5.1 Baseline specification

Figure 2 presents the estimates of the state-dependent IS curve parameters using the benchmark instrument set and assuming $\kappa = 1$ and $m = 0$:

$$c_t = \alpha(\tau) + (1 - \mu(\tau))\, c_{t-1} + \mu(\tau)\, c_{t+1} + \beta(\tau)\, [(i_t - \pi_{t+1})] + u_t, \quad u_t \sim U(0,1)$$

The rows of Figure 2 present results for the five filters used in Section 4 to extract the cyclical component of consumption. Our discussion begins with the left column which displays the estimated coefficient on c_{t+1}. The Two Stage Least Square (TSLS) estimates, reported as dotted lines, are centered around the value of 0.5, consistent with the findings in Fuhrer and Rudebusch (2004). On the one hand, the IVQR estimates are significantly smaller than the TSLS estimates for quantiles below 20%. On the other hand, the TSLS significantly under-estimates the extent of forward-looking behaviour at the top 30^{th} percentile of the conditional distribution of consumption. This evidence suggests that periods in which consumption is conditionally low (high) are also periods characterized by higher (lower) persistence. In other words, agents appear more forward looking in good times (as measured by high values of τ). This finding appears robust across the different cyclical measures.

As for the interest rate semi-elasticity, the TSLS estimates in the second column of figure 2 are centered around zero and they are never statistically significant as it is often the case in Fuhrer and Rudebusch (2004). This finding is similar to the IVQR estimates obtained for values of $\tau < 65\%$. In the top 25^{th} percentiles of the conditional distribution of consumption, however, the sensitivity of the IS schedule to the real rate is recorded at values around -0.04. This evidence of nonlinearity suggests that monetary policy (and more generally the effects of real rate on consumption) exerts its maximum impact during periods in which consumption is conditionally high.

In summary, we find strong evidence in favour of state-dependent coefficients: the heterogeneous estimates of the IS curve parameters across the conditional distribution of consumption differ significantly from the average effect estimated using a constant parameter specification.

5.2 Alternative estimation methods

The results in Figure 2 are based on quantile regressions. In this section, we show that the evidence against a linear consumption-interest rate relationship is robust to using two popular methods to deal with state-dependent parameters, namely Markov-switching and threshold models. While the details of these two specifications are postponed until the model comparison in Section 7, we anticipate here that the quantile regression estimates are associated with the best out-of-sample forecasting performance.

Table 3 reports the estimates of the state-dependent parameters obtained from a two-states Markov-switching model (top panel) and a threshold model (bottom panel) using HP-filtered consumption.[5] Both estimation methods reveal that one set of observations, which we refer to as *regime 1*, is characterized by a significantly smaller coefficient on future consumption as well as an insignificant coefficient on the real rate. As for *regime 2*, in contrast, one cannot reject the hypothesis that agents are purely forward-looking and that the impact of the real rate on consumption is significant.

In the threshold model, *regime 1 (2)* corresponds to observations below (above) the estimated threshold of -0.004 and therefore it can be interpreted as the below- (above-)

[5] The results below are robust to estimating a three-states specification. As for the threshold model, we use c_{t-1} as threshold variable but similar findings hold using c_t.

average consumption regime. In the Markov-switching model, this information is conveyed by the regime probability, which is reported in Figure 3 (grey areas) together with the cyclical measure of consumption (blue line). The probability of being in regime 2, which according to Table 3 is characterized by a higher (lower) estimated coefficient on the lead of consumption and the ex-ante real rate is significantly larger during periods of above-average consumption.

In summary, quantile regressions, Markov-switching and threshold models paint a fairly robust picture: state dependent parameters are a pervasive feature of the transmission mechanism in post-WWII U.S. data. In particular, all methods suggest that during periods of above-average consumption agents tend to be more forward-looking and their consumption becomes more sensitive to movements in the real rate.

5.3 From state-dependent to time-dependent estimates

The focus of our analysis is on assessing the assumption of constant parameters that is behind the majority of existing estimates of the IS curve. For the sake of illustration to a broader audience, in this section we offer a time-varying coefficient interpretation of our results. To this end, we link the distribution of consumption conditional on covariates (which is the basis for the IVQR estimation) to the unconditional distribution of consumption. A main goal is to establish whether periods of conditionally low (high) consumption correspond to periods of *unconditionally* below-trend (above-trend) consumption. We find that they do, which allows us to interpret the state-dependent estimates of Figure 2 in terms of phases of the business cycle.

To map the quantile regressions into time-varying estimates, we use a version of the procedure described in Koenker (2005 pp. 295-316). In particular, in this sub-section and *this sub-section only*, we estimate the following version of the benchmark specification:

$$c_t = \alpha^* (\tau) + (1 - \mu (\tau))\, \tilde{c}_{t-1} + \mu (\tau)\, \tilde{c}_{t+1} + \beta (\tau) \left[(\tilde{\imath}_t - \tilde{\pi}_{t+1}) \right] + u_t$$

where the superscript ˜ indicates deviations from the mean. Note that $\alpha^* (\tau)$ now captures the level of c_t in each τ, given the mean of the real rate and lagged consumption. We then use the estimated value of $\alpha^*(\tau)$ in each quantile to link the level of the dependent variable across time to the values of $\mu(\tau)$ and $\beta(\tau)$. More specifically, we define a $T \times 1$ indicator variable $I_i[c_t \leq \alpha^*(\tau_i)]$ for $i = 1$ and $I_i[\alpha^*(\tau_{i-1}) < c_t \leq \alpha^*(\tau_i)]$ for $i = 2..M$ where M indexes the values of τ that we consider. We also specify μ_t and β_t as two $T \times 1$ vectors with all elements initially equal to zero.

For each value of i (and thus the indicator variable $I_i[c_t \leq \alpha^*(\tau_i)]$), we loop through time $t = 1...T$ i.e. through the elements of μ_t and β_t setting $\mu_t = \mu(\tau_i)$ and $\beta_t = \beta(\tau_i)$ if the t^{th} element of $I_i[.] = 1$. Repeating this for $i = 1..M$ fills all elements of μ_t and β_t and produces a time-series for these coefficients.

Our findings are reported in Figure 4. The top (bottom) panel shows the estimates of μ (β) over the post-WWII sample together with our baseline measure of cyclical consumption. Whenever consumption is above trend, consumption decisions appear (i) more forward-looking and (ii) more sensitive to movements in the real rate. On the other hand, during the troughs of the cycle, and especially during the recession episodes of the 1970s, 1981, 1992 and 2008, agents tend to be significantly more backward-looking and their expenditure becomes far less sensitive, if any, to the real rate. While this finding is not based on any specific assumption on the way IS curve parameters evolve over time, it is interesting to note that the evidence in Figure 4 appears more reminiscent of the type of time-variation estimated using a regime switching model than using random walk drifting coefficients.

6 Sensitivity Analysis

In this section, we perform five robustness checks. First, we consider alternative measures of real activity. Second, we allow for four different instrument sets. Third, we experiment with different lags of the real rate structure. Fourth, we estimate the IVQR model without imposing the restrictions that the coefficient on the real rate should be non-positive and that the coefficients on consumption should sum up to one. Five, we use the monetary policy shocks series constructed by Romer and Romer (2004) in place of the real rate. Our findings are robust to all these modifications of the baseline case.

6.1 Measures of activity

In the baseline case, we have used per capita real expenditure on non-durable goods and services. In this section, we first consider non-durable goods and services consumption separately, and then perform the IVQR estimation using the real expenditure on durable goods and the monthly estimates of GDP produced by Stock and Watson (2010).

The results are reported in Figure 5 and they show a similar pattern to Figure 2 with some important qualifications. First, for both IS curve coefficients, the evidence of state-dependence is significantly stronger using services consumption than using non-durable goods consumption. Second, there is little evidence of heterogeneity across phases of the business cycle using durable goods consumption. Third, using Stock and Watson's monthly GDP measure the estimates of degree of forward-lookingness at the upper tail of the conditional distribution are larger than the estimates at the lower tail while the evidence for the interest rate semi-elasticity is more muted. The corresponding F-statistics for the four measures are reported in the second row of Table 2. Only services consumption and Stock and Watson's monthly GDP pass comfortably Stock and Yogo's test of instrument strength. Interestingly, these are also the only two measures for which we find strong evidence of state-dependent IS curve parameters.

6.2 Instrument set

The results of the F-tests using the different filters together with the Montecarlo analysis in Appendix A suggest that our baseline estimates are unlikely to suffer from a weak instrument problem. Nevertheless, it is useful to assess the sensitivity of our findings to using instruments set with a different number of lags for each endogenous variable. The results are reported in Figure 6 for the instrument sets $Z1_t = (c_{t-2}..c_{t-6}, i_{t-2}..i_{t-6}, \pi_{t-2}..\pi_{t-6}, 1)$; $Z2_t = (c_{t-2}..c_{t-12}, i_{t-2}..i_{t-12}, \pi_{t-2}..\pi_{t-12}, 1)$; $Z3_t = (c_{t-1}, i_{t-1}, \pi_{t-1}, 1)$; $Z4_t = (c_{t-2}, i_{t-2}, \pi_{t-2}, 1)$. The conclusions one can drawn upon the alternative instrument sets are very similar to the ones for the baseline case. It should be noted, however, that according to the F-statistics in the third row of Table 2 the alternative sets are less likely to qualify as strong instruments.

Indeed, this is the reason behind our choice of focusing on the baseline set Z_t in Section 5.

6.3 Timing of the transmission mechanism

Different specifications of the lag structure for the real interest rate also appear unable to overturn our evidence of state-dependence. Specifying [k=0, m=1], [k=12, m=0] and [k=12, m=1] for $\sum_{j=0}^{\kappa-1}(i_{t+j+m} - \pi_{t+j+m+1})$ in the first, second and third row of Figure 7 respectively makes our conclusions even stronger. (Fully) backward-looking behaviour is predominant at the bottom (end) 20% of the conditional distribution while (fully) forward-looking behaviour is predominant at the top (end) 20%. As for the interest rate semi-elasticity, we confirm that only the top 25% observations above the consumption average are associated with significantly negative estimates, ranging from values around -0.03 in the first row to values around -0.06 in the last row. The F-statistics for these specifications, reported in the last row of Table 2, are always above Stock and Yogo's critical value.

6.4 Unrestricted estimates

The estimates above restrict the coefficient on the real interest to be non-positive. Furthermore, the coefficients on backward-looking and forward-looking terms are restricted to sum up to one. In this section, we assess the robustness of our finding on state-dependence to relaxing both restrictions.

The estimated coefficients on future and lagged consumption in Figure 8 are very similar to the benchmark results and their sum is never statistically larger than one. The pattern of the coefficient on the real interest rate across quantiles is very similar to the pattern for the benchmark case: the estimates of β become more negative and statistically significant at the right tail of the consumption distribution. Note, however, that for quantiles $\tau < 30\%$, the positive coefficient is hard to justify from an economic perspective. The Stock and Yogo's F-statistics for this case is 18.6.

6.5 Monetary policy shocks

In this section, we use an alternative identification strategy to assess the extent to which the IVQR estimates reported above genuinely reflect the causal effect of interest rate on consumption and therefore can be interpreted as describing an IS curve type of relationship. The alternative strategy is based on the measure of monetary policy shocks proposed by Romer and Romer (2004) and available on their websites. These are the residuals of a regression of changes in the federal funds rate around FOMC meetings on the level of the federal fund rate as well as the levels and the changes of the Greenbook forecasts for inflation, real output growth and unemployment in the previous quarter, contemporaneously and two quarters ahead.

Following Romer and Romer (2004), we use thirty six lags of the monetary policy shock and twelve lags of consumption to generate a quantile regressions version of their estimated transmission mechanism specification:

$$c_t = \delta(\tau) + \sum_{j=1}^{24} \rho_j(\tau) c_{t-j} + \sum_{i=1}^{36} \gamma_i(\tau) MP_{t-i} + v_t, \quad v_t \sim U(0,1) \tag{13}$$

The cumulated impulse responses for each quantile of the conditional consumption distribution estimated on the basis of equation (13) are reported in Figure 9. Two main results emerge from this exercise. First, for any given quantile, the profile of the expected consumption adjustments to a monetary policy shock resembles on average the shape of the impulse response of real activity reported by Romer and Romer (2004), sharing both the hump shape and the peak between two and three years after the shock across most quantiles. Second, consistent with the results in Section 5, for any given forecast horizon, the consumption reactions tend to be close to zero at the left tail of the conditional distribution and then become more negative (and significant) at higher quantiles. Together with the results from the previous section, we conclude that our quantile regression estimates are likely to reflect the causal effect of interest rate movements on real activity.

7 Model comparison

In this section, we perform a statistical comparison of the performance of the IVQR method relative to alternative strategies to model parameter instability in the consumption-interest rate relationship. To this end, we consider four competing specifications: (i) IV Markov switching regimes, (ii) an IV threshold model, (iii) an IV specification augmented with an interaction term between the endogenous variables and a dummy that takes value of one during NBER recession months and (iv) an IV random walk drifting parameter model. The first three specifications explicitly model state dependent parameters with endogenous regressors. The fourth specification postulates that the coefficients of the structural equation evolve smoothly over time.

Following Sola, Psaradakis and Spagnolo (2005), the Markov switching model is specified as

$$c_t = \alpha_s + (1 - \mu_s)\, c_{t-1} + \mu_s c_{t+1} + \beta_s \left[(i_t - \pi_{t+1})\right] + \sigma_s u_t$$

where s=0, 1 denotes the unobserved regime which is assumed to follow a two state Markov chain with fixed transition probabilities. Sola, Psaradakis and Spagnolo (2005) show how to modify the filter in Hamilton (1994) to allow for endogenous variables, which are dealt with via regime dependent instruments.

As for the threshold model, we adopt the strategy proposed in Caner and Hansen (2004) and specify aggregate consumption as:

$$\begin{aligned} c_t &= \alpha_1 + (1 - \mu_1)\, c_{t-1} + \mu_1 c_{t+1} + \beta_1 \left[(i_t - \pi_{t+1})\right] + \sigma_1 u_t \quad \text{if } q_t \leq \gamma \\ c_t &= \alpha_2 + (1 - \mu_2)\, c_{t-1} + \mu_2 c_{t+1} + \beta_2 \left[(i_t - \pi_{t+1})\right] + \sigma_2 u_t \quad \text{if } q_t > \gamma \end{aligned}$$

where γ is the threshold parameter and q_t is the threshold variable. Results below refer to $q_t = c_{t-1}$ but we have verified robustness to using c_{t-2}. Carner and Hansen (2004) generalize the regression case in Hansen (2000) to estimate threshold models with endogenous variables but an exogenous threshold variable.

The time-varying parameter model with endogenous regressors is borrowed from Kim (2006) who proposes a Heckman-type two-step procedure to estimate the following system of equations:

$$
\begin{aligned}
c_t &= \alpha_t + (1-\mu_t)\, c_{t-1} + \mu_t c_{t+1} + \beta_t \left[(i_t - \pi_{t+1})\right] + \sigma u_t \\
\Theta_t &= \Theta_{t-1} + \varepsilon_t \;\; \text{with} \;\; \Theta_t \equiv [\alpha_t, \mu_t, \beta_t] \\
x_t &= Z_t' \delta_t + v_t \;\; \text{with} \;\; \delta_t = \delta_{t-1} + e_t \;\; \text{and} \;\; x_t \equiv [c_{t+1}, (i_t - \pi_{t+1})]
\end{aligned}
$$

where Z_t is the instrument set and the errors are i.i.d. and normally distributed with variance matrices Σ_ε, Σ_v and Σ_e, respectively.

To evaluate the performance of IVQR vis a vis the performance of the alternative strategies (i) to (iv), we focus on a measure of out of sample predictability for each model *relative* to the predictability of the IVQR specification which is used as benchmark. In all models, consumption is measured relative to the HP trend and the instrument set is the same as in Section 4. The first estimation sample is 1948:5-1969:12. After that, we re-estimate each model recursively from j = 1970:1 to 2009:7 adding one observation at a time. At each recursion j and for each model m, we compute a pseudo out-of-sample forecast at the monthly horizon h (with h=1, 6 and 12) and use this to compute the root mean squared forecast error (RMSE) as $RMSE_{m,h} = \sum_{j=1}^{J} \frac{1}{J} \sqrt{\frac{1}{T_j} \sum_{t=1}^{T_j} \left(\hat{c}_{t+h} - c_t\right)^2}$ where T_j is the length of the time series in the j^{th} recursion for model m and J is the total number of recursions. The measure of relative predictability is then the ratio between $RMSE_{m,h}$ and the $RMSE_{IVQR,h}$ of the IVQR model: values larger than one denotes a forecasting deterioration relative to the benchmark model. Statistical differences are assessed using the test of equal predictive accuracy proposed by Diebold and Mariano (1996).

The results are presented in Table 4 and they reveal a number of regularities. First, the forecasting performance of the IVQR model tends to be at least as good as the performance of the other non-linear models. Second, most of the relative RMSE statistics are larger than one and often statistically significant (denoted by asterisks) according to the Diebold-Mariano statistics. Third, the largest gains for the IVQR specification are recorded relative to the Markov-switching and the NBER dummy models. Fourth, the performances of the time-varying parameters and the threshold models tend to be similar to the performance of the IVQR specification, resulting in forecasts that appear less accurate at shorter horizons but marginally more accurate at longer horizons. In summary, the IVQR specification fits U.S. data on consumption well in a way that tends to compare favourably to other nonlinear alternatives, both out-of-sample as shown here and in-sample as shown in Mumtaz and Surico (2011).

8 Conclusions

This paper provides empirical evidence in favour of significant nonlinearity in the dynamic relationship between real activity and the interest rate. In periods of below-average consumption/output, households' decisions tend to be more backward-looking and less sensitive to movements in the real rate of interest. Periods at the other end of the conditional distribution of real activity are associated with fully forward-looking behaviour and with the maximum impact of the real rate. The average effect estimated on the basis of a constant parameter specification, in contrast, points toward an insignificant interest rate semi-elasticity and towards roughly equal weights received by backward-looking and forward-looking components.

Our results offer empirical support for the notion that the dynamics of consumption and output during expansions are qualitatively and quantitatively different from the dynamics during contractions, suggesting that monetary policy (and any other shock channelled through the real rate of interest) has asymmetric effects over the business cycle. One caveat is that our estimates are not obtained in the context of a fully specified structural model. While we are not aware of contributions estimating a DSGE model with parameters that are explicitly allowed to vary with the state of the economy, in future research it would be interesting to generalize our method to the multivariate case in a way that would retain the advantages of a general equilibrium analysis while providing a flexible strategy to model state-dependent parameters.

References

Bianchi, Francesco, 2013, Regime switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics, *Review of Economic Studies*, forthcoming.

Caner, Mehmet and Bruce E. Hansen, 2004, Instrumental Variable Estimation of a Threshold Model, *Econometric Theory* 20, pp. 813-843.

Canova, Fabio and Luca Gambetti, 2009, Structural Changes in the US Economy: Is There a Role for Monetary Policy?, *Journal of Economic Dynamics and Control* 33, pp. 477-490.

Chernozhukov, Victor and Christian Hansen, 2005, An IV Model of Quantile Treatment Effects, *Econometrica* 73, pp. 245-261.

Chernozhukov, Victor and Han Hong, 2003, An MCMC approach to classical estimation, *Journal of Econometrics* 115, pp. 293–346.

Chevapatrakul, Thanaset, Tae-Hwan Kim and Paul Mizen, 2009, The Taylor Principle and Monetary Policy Approaching a Zero Bound on Nominal Rates: Quantile Regression Results for the United States and Japan, *Journal of Money, Credit and Banking* 41, pp. 1705-1723.

Cogley, Tim W., and Tom J. Sargent, 2005, Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII US, *Review of Economic Dynamics* 8, pp. 262-302.

Crossley, Thomas and Hamish Low, 2011, Is the Elasticity of Intertemporal Substitution Constant?, *Journal of the European Economic Association* 9, pp. 87-105.

Dennis, Richard, 2009, Consumption Habits in a New Keynesian Business Cycle Model, *Journal of Money, Credit and Banking* 41, pp. 1015-1030.

Fernández-Villaverde, Jesus, Grey Gordon, Pablo A. Guerrón-Quintana and Juan Rubio-Ramírez, 2012, Nonlinear Adventures at the Zero Lower Bound, NBER wp No. 18058.

Fuhrer, Jeffrey C. and Glenn D. Rudebusch, 2004, Estimating the Euler equation for output, *Journal of Monetary Economics* 51, pp. 1133–1153.

Galvao, Antonio F. Jr., Gabriel Montes-Rojas and Gabriel Sung Park, 2011, Quantile Autoregressive Distributed Lag Model with an Application to House Price Returns, *Oxford Bulletin of Economics and Statistics*, forthcoming.

Gelman, Andrew and Donald Rubin, 1992, Inference from Iterative Simulation Using Multiple Sequences, *Statistical Science* 7, pp. 457-472.

Hamilton, James D., 1994, Time Series Analysis, Princeton University Press.

Hansen, Bruce E., 2000, Sample Splitting and Threshold Estimation, *Econometrica* 68, pp.575-603.

Ireland, Peter, 2004, Technology Shocks in the New Keynesian Model, *The Review of Economics and Statistics* 86, pp. 923-936.

Kim, Chang-Jin, 2006, Time-Varying Parameter Models with Endogenous Regressors, *Economics Letters* 91, pp. 21-26.

Koenker, Roger and Gil Bassett, 1978, Regression quantiles, *Econometrica* 46, pp. 33-50.

Koenker, Roger and Zhijie Xiao, 2006, Quantile Autoregression, *Journal of the American Statistical Association* 101, pp. 980-990.

Lindè, Jesper, 2005, Estimating New-Keynesian Phillips curves: A full information maximum likelihood approach, *Journal of Monetary Economics* 52, pp. 1135-1149.

Mumtaz, Haroon and Paolo Surico, 2011, Estimating the Aggregate Consumption Euler Equation with State-Dependent Parameters, CEPR Discussion Papers 8233.

Romer, Christina and David Romer, 2004, A New Measure of Monetary Shocks: Derivation and Implications, *American Economic Review* 94, pp. 1055-84.

Sims, Christopher A, 2002, Solving Linear Rational Expectations Models, *Computational Economics* 20(1-2), pp. 1–20.

Sola, Martin, Psaradakis, Zacharias and Spagnolo, Fabio, 2005, Testing the Unbiased Forward Exchange Rate Hypothesis using a Markov Switching Model and Instrumental Variables, *Journal of Applied Econometrics* 20, pp. 423-437.

Stock, James H. and Mark W. Watson, 2010, Monthly Estimates of GDP and GDI, available at http://www.princeton.edu/~mwatson/mgdp_gdi.html.

Stock, James H. and Motohiro Yogo, 2002, Testing for Weak Instruments in Linear IV Regression, NBER Technical Working Papers 0284.

Tenreyro, Silvana and Gregory Thwaites, 2013, Pushing on a String: US Monetary Policy is Less Powerful in Recessions, mimeo, London School of Economics and Bank of England.

Table 1: Sign Restrictions in the benchmark model

	Lower	*Upper*
$\mu(\tau)$	0	1
$\beta(\tau)$	-1	0

Table 2: Stock and Yogo multi-variate F-statistics

specifications using different:

filters	*Hodrick-Prescott*	*band pass*	*linear DT*	*quadratic DT*	*approx. DT*
	18.6	310.7	23.3	23.4	17.6
measures of real activity	*non-durable cons.*	*services cons.*	*durable cons.*	*SW's GDP*	
	1.7	25.0	5.2	12.3	
instrument sets	$Z1$	$Z2$	$Z3$	$Z4$	
	9.9	7.0	0.1	0.8	
real rate lag structures	$k=0, m=1$	$k=12, m=0$	$k=12, m=1$		
	18.6	18.2	18.2		

Note: DT stands for de-trending and SW for Stock and Watson (2010). The instruments are $Z1_t=(c_{t-2}..c_{t-6}, i_{t-2}..i_{t-6}, \pi_{t-2}..\pi_{t-6}, 1)$; $Z2_t=(c_{t-2}..c_{t-12}, i_{t-2}..i_{t-12}, \pi_{t-2}..\pi_{t-12}, 1)$; $Z3_t=(c_{t-1}, i_{t-1}, \pi_{t-1}, 1)$; $Z4_t=(c_{t-2}, i_{t-2}, \pi_{t-2}, 1)$. k and m refers to the lag structure in the specification (7). The critical value for which the bias relative to OLS is no more than 5% is 9.5. Above this value, Stock and Yogo (2003) deem the instruments strong and the inference based on the IV estimator reliable.

Table 3: Alternative state-dependent parameter estimates

Markov-switching	c_{t+1}	$i_t - \pi_{t+1}$	*intercept*
regime1	0.404*	-0.007	-0.07*
	(0.14)	(0.006)	(0.03)
regime2	0.929*	-0.02*	0.009*
	(0.06)	(0.009)	(0.03)
threshold model	$E_t c_{t+1}$	$i_t - E_t \pi_{t+1}$	*intercept*
regime1	0.120	0.005	-0.001*
	(0.10)	(0.008)	(.0002)
regime2	0.866*	-0.03*	0.001*
	(0.11)	(0.01)	(.0002)

Note: the instruments are Z_t=(c_{t-2}..c_{t-3},i_{t-2},π_{t-2},1). In the Markov-switching model, the probability to stay in regime 1 (2) is 0.83 (0.92) with standard error 0.02 (0.01). The estimated threshold for c_{t-1} is -0.004 with standard error 0.005 and regime 1 refers to observations below the threshold. * denotes significance at 1% confidence level.

Table 4: Out-of-sample relative RMSE

	monthly forecast horizon (h)		
MODELS	$h = 1$	$h = 6$	$h = 12$
Markov Switching	1.42*	1.48*	1.58*
Time-varying parameters	1.22*	1.04	0.97
NBER dummy	2.18*	1.58*	1.38*
threshold	1.13	1.01	0.94*

Note: the Root Mean Squared forecast Error (RMSE) of each model is reported relative to the instrumental variable quantile regressions model. Asterisks denote significant differences at the 5% level as measured by a statistic greater than 1.96 for the test of equal predictive accuracy proposed by Diebold and Mariano (1996).

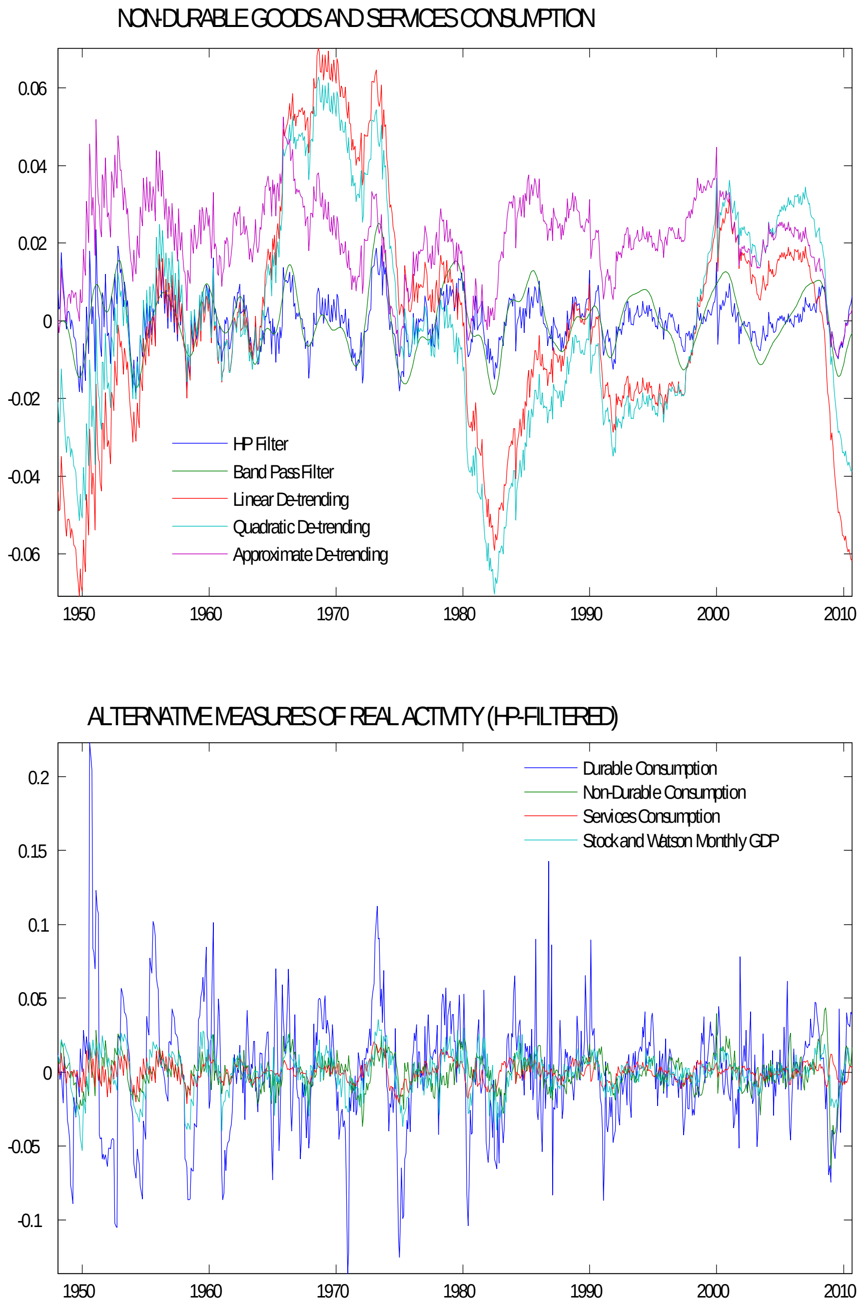

Figure 1: data

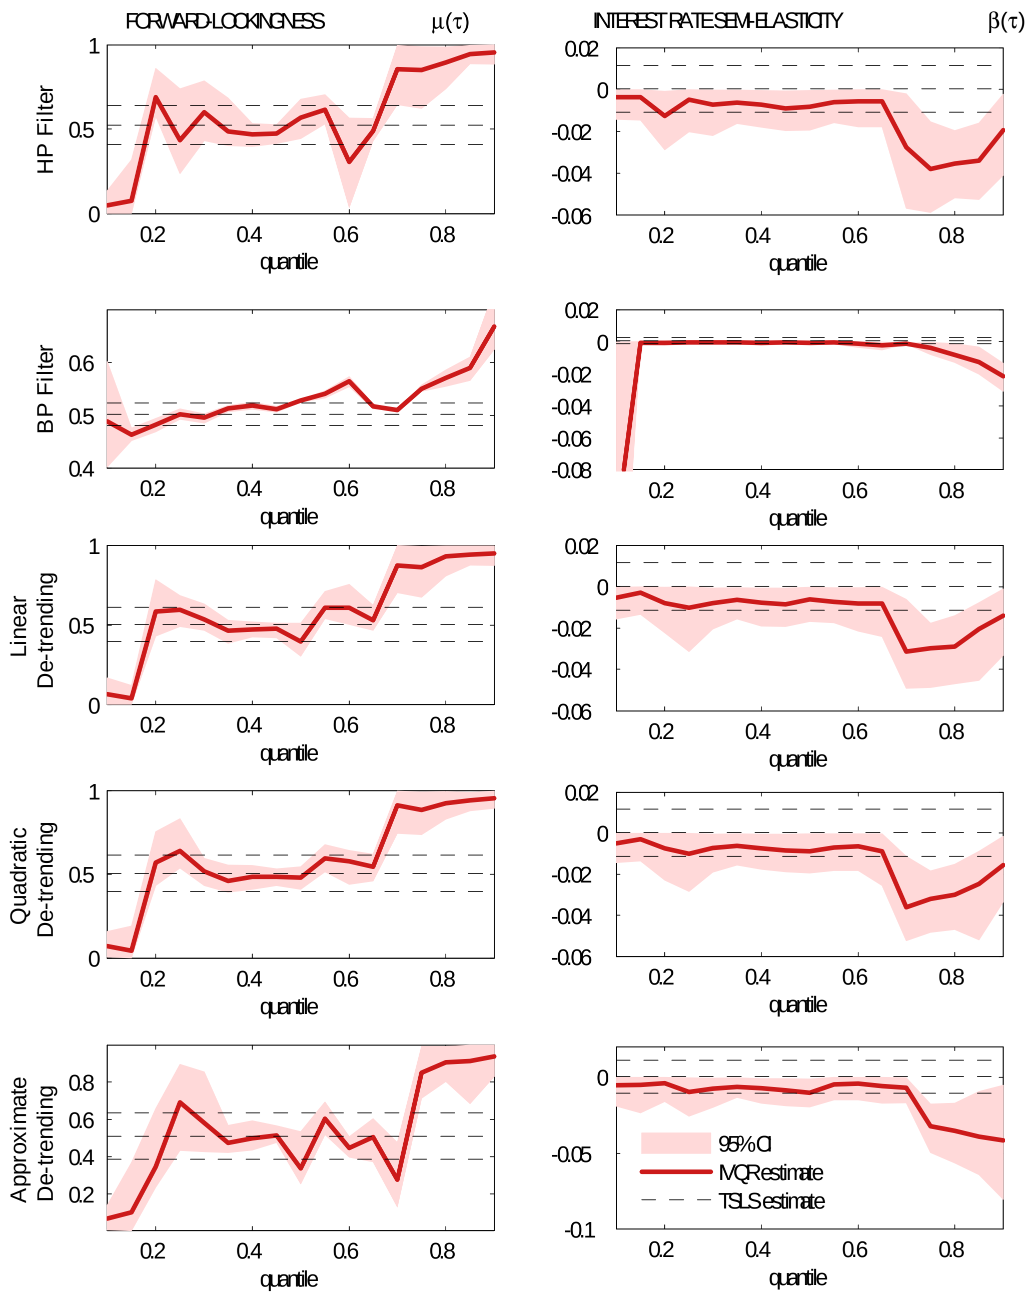

Figure 2: IVQR estimates of the parameters of equation (7) with $\kappa = 1, m = 0$ using the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$ and alternative filters on the per-capita real personal consumption expenditure of non-durable goods and services. Dotted lines represent to TSLS estimates and 95% credible sets. Sample: 1948:5-2010:7.

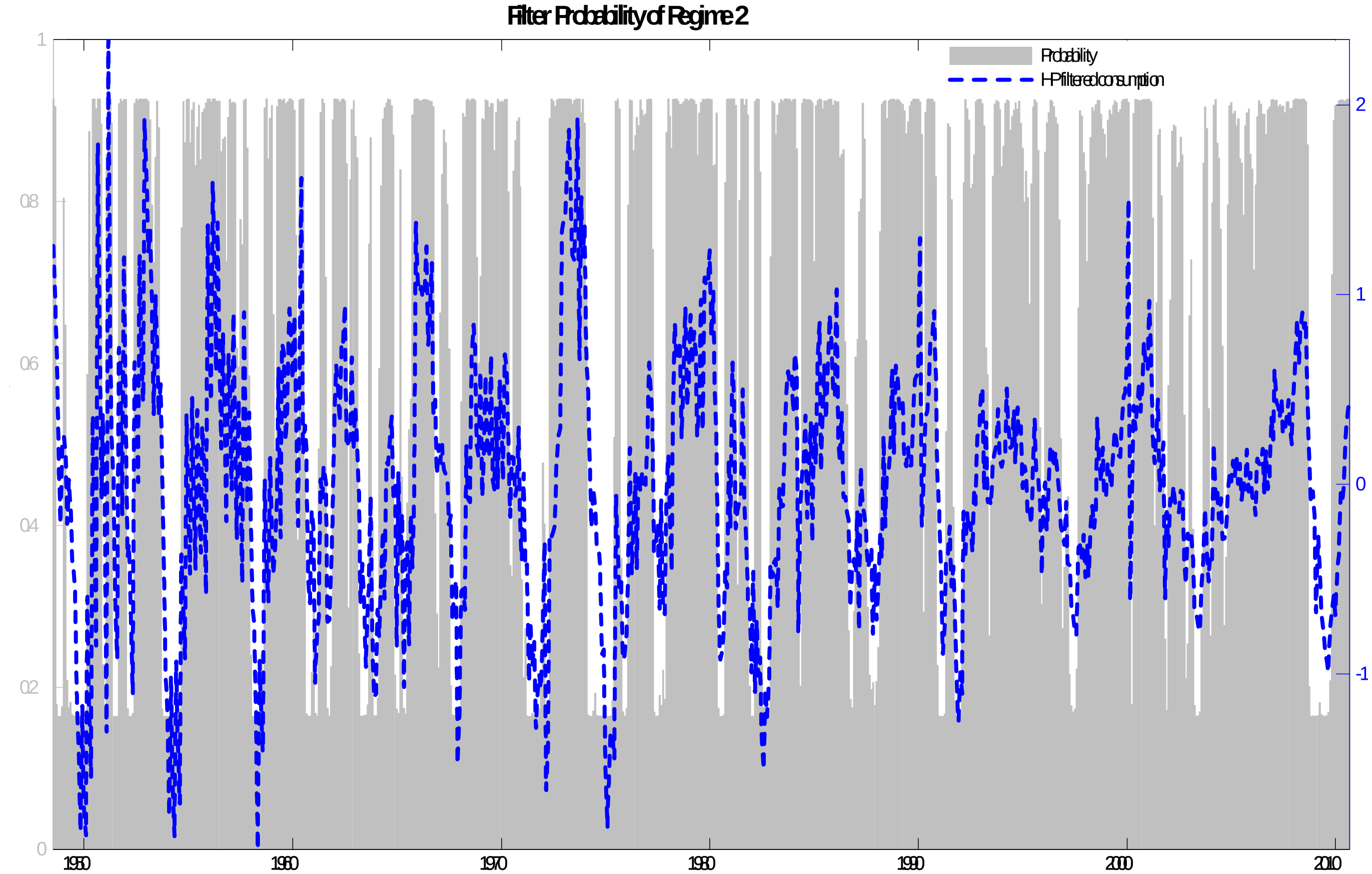

Figure 3: probability of regime 2 in a two-state Markov-switching estimates of the parameters of the IS curve (shaded area in grey) with $\kappa = 1, m = 0$ using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services (dotted line in blue) and the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$. Sample: 1948:5-2010:7.

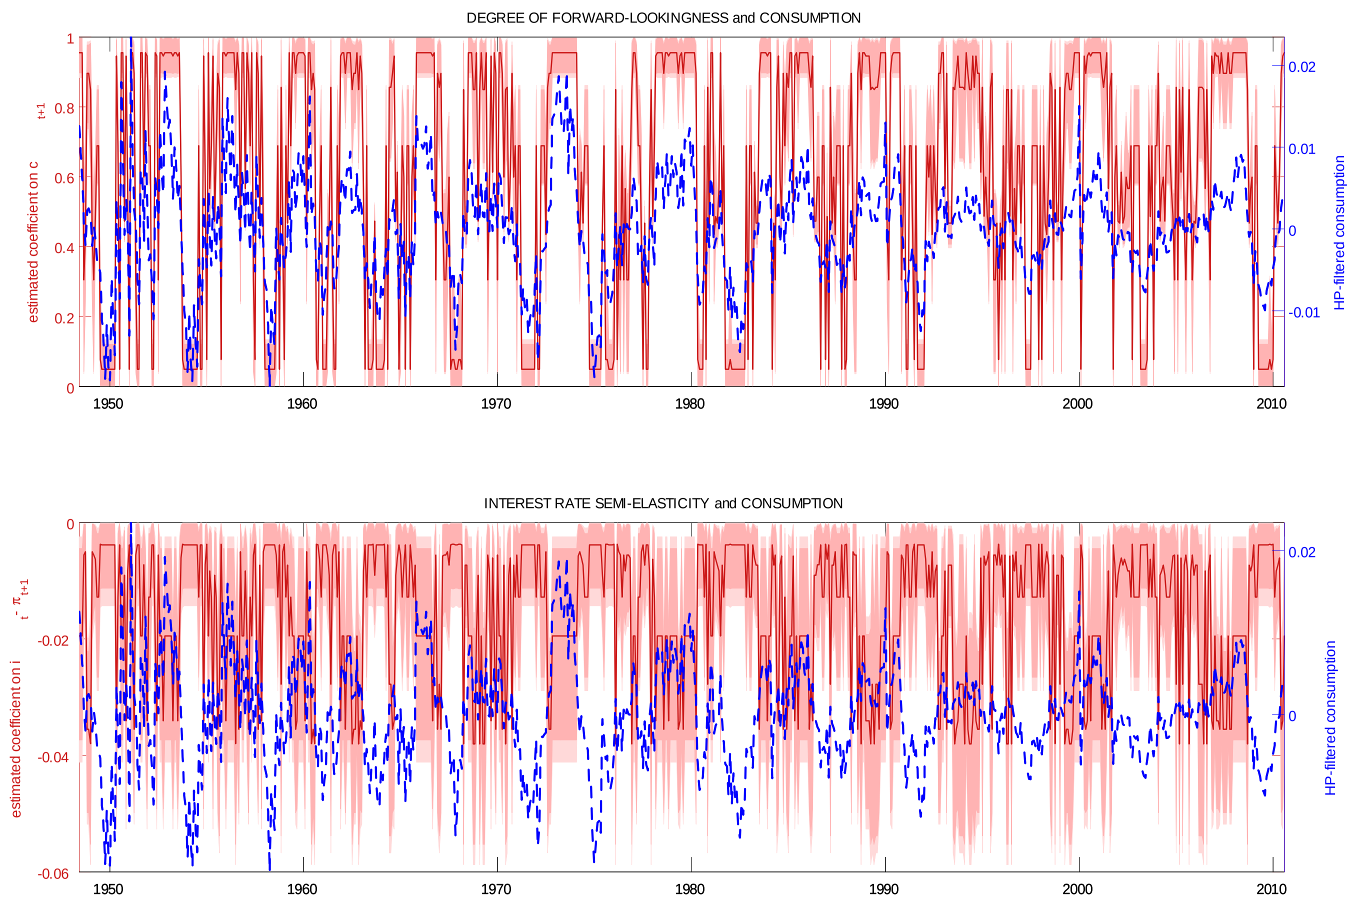

Figure 4: time series mapping of the IVQR estimates of the parameters of the IS curve (solid lines in red with 95% credible set shaded area) with $\kappa = 1, m = 0$ using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services (dotted line in blue) and the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$. Sample: 1948:5-2010:7.

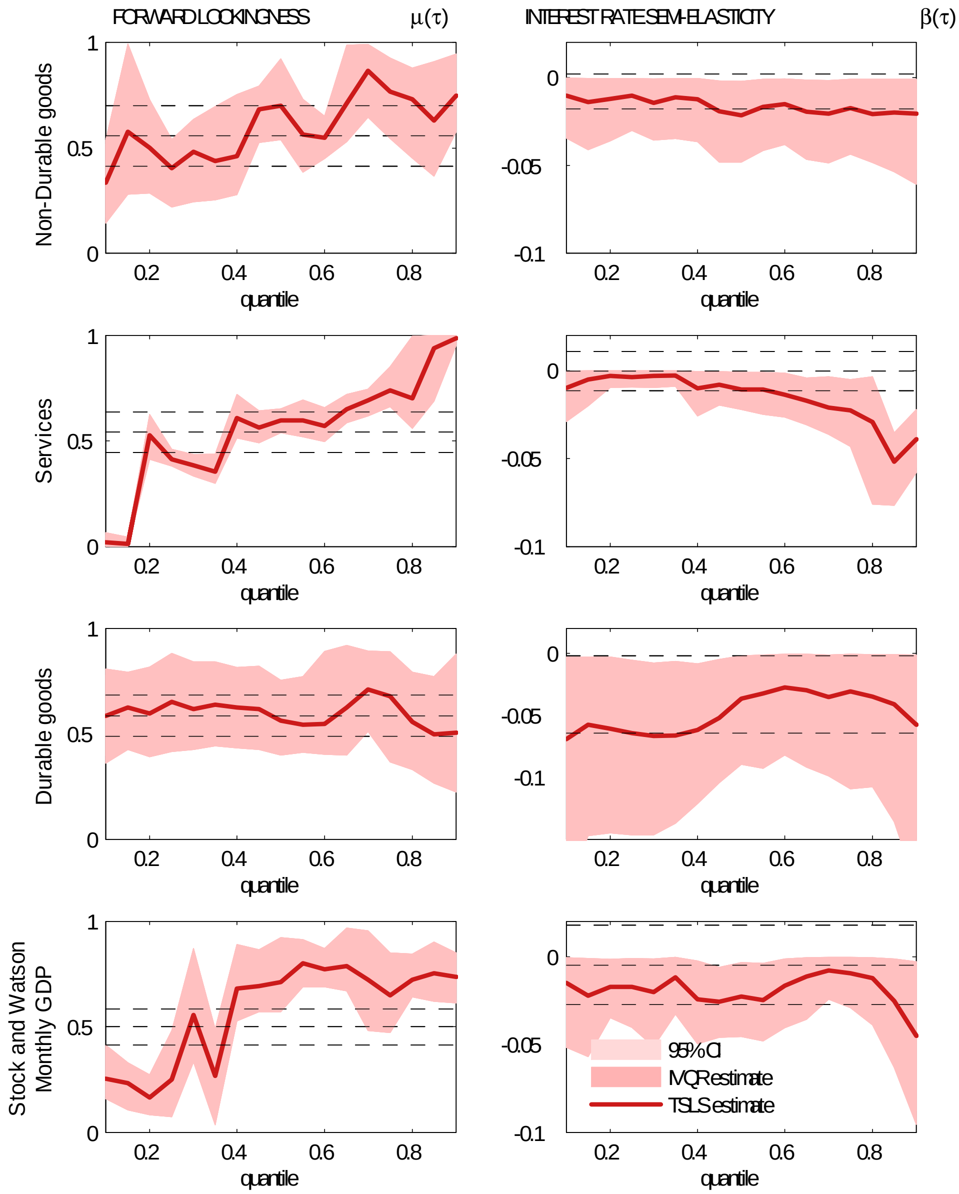

Figure 5: IVQR estimates of the parameters of equation (7) with $\kappa = 1, m = 0$ using the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$, the HP filter and alternative measures of real activity. Dotted lines represent to TSLS estimates and 95% credible sets. Sample: 1948:5-2010:7.

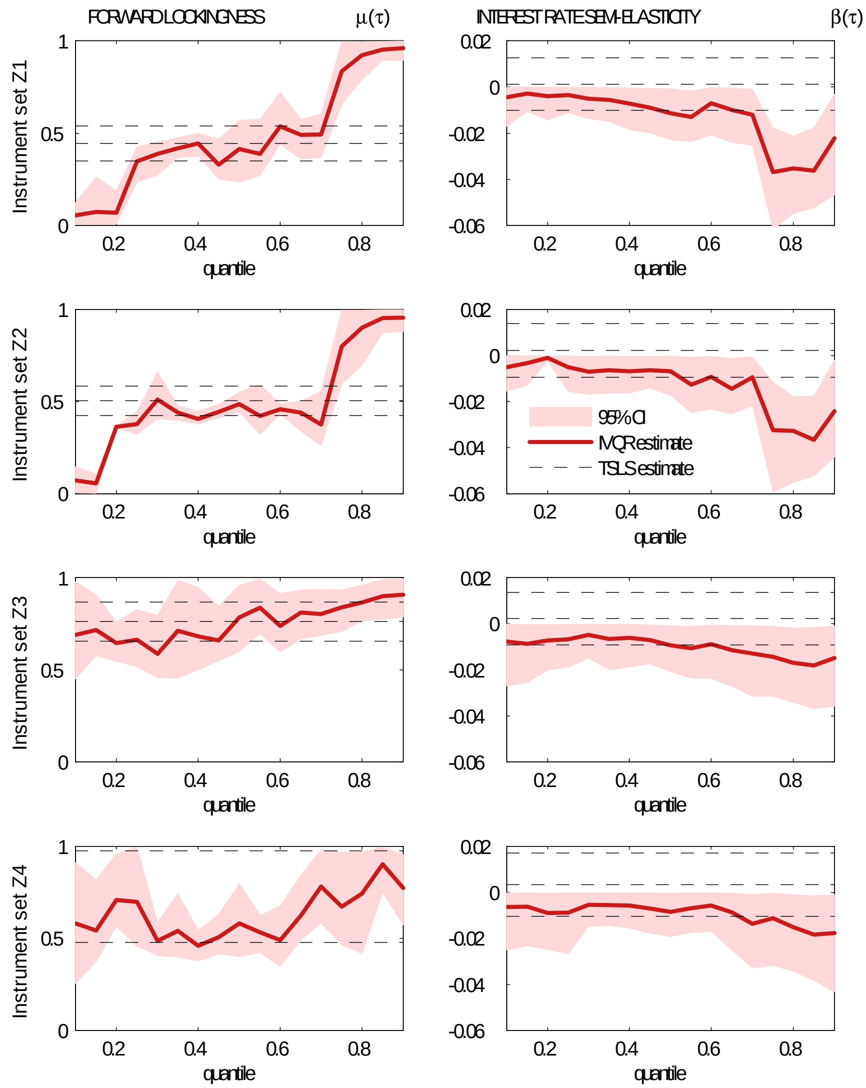

Figure 6: IVQR estimates of the parameters of equation (7) with $\kappa = 1, m = 0$ using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services, and the alternative instrument sets $Z1_t = (c_{t-2}..c_{t-6}, i_{t-2}..i_{t-6}, \pi_{t-2}..\pi_{t-6}, 1)$; $Z2_t = (c_{t-2}..c_{t-12}, i_{t-2}..i_{t-12}, \pi_{t-2}..\pi_{t-12}, 1)$; $Z3_t = (c_{t-1}, i_{t-1}, \pi_{t-1}, 1)$; $Z4_t = (c_{t-2}, i_{t-2}, \pi_{t-2}, 1)$. Dotted lines represent to TSLS estimates and 95% credible sets. Sample: 1948:5-2010:7.

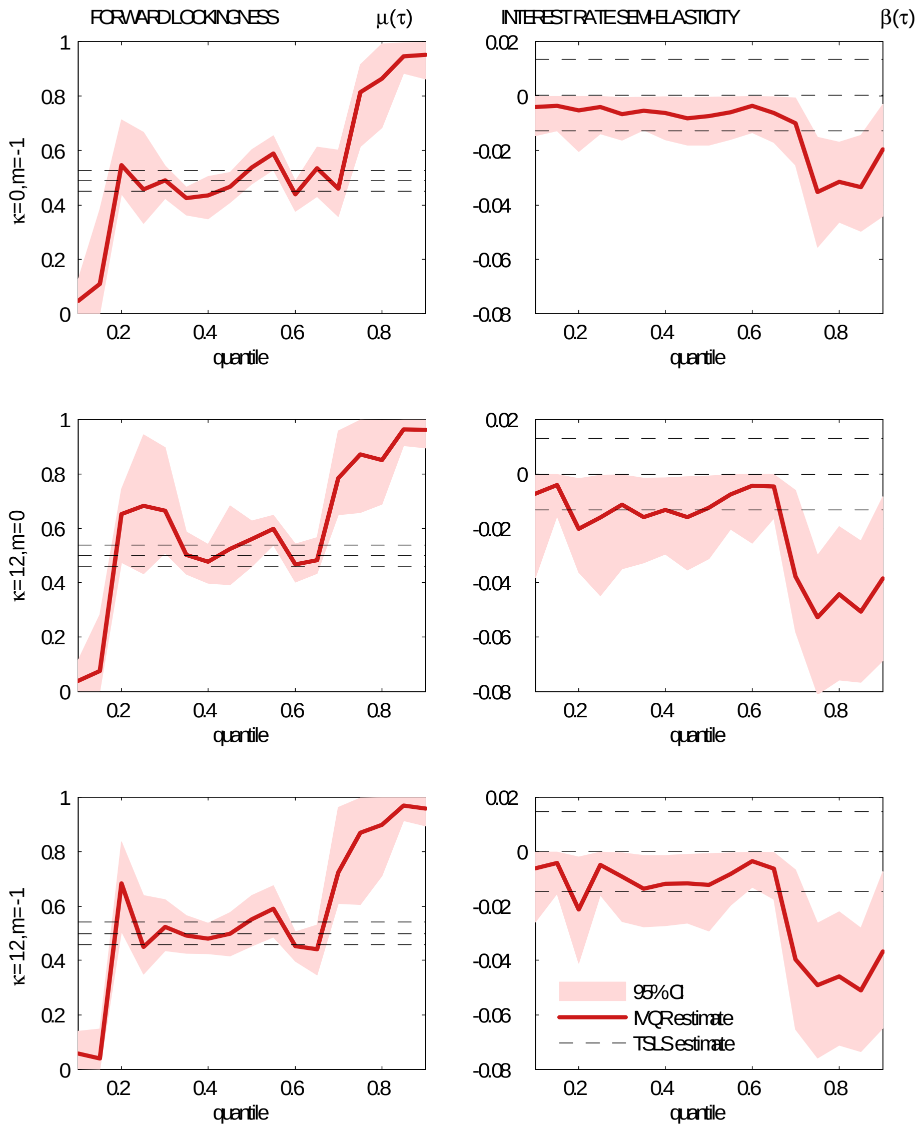

Figure 7: IVQR estimates of the parameters of equation (7) using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services, the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$ and alternative values of κ and m for the timing of the effect of the expected real interest rate. Dotted lines represent to TSLS estimates and 95% credible sets. Sample: 1948:5-2010:7.

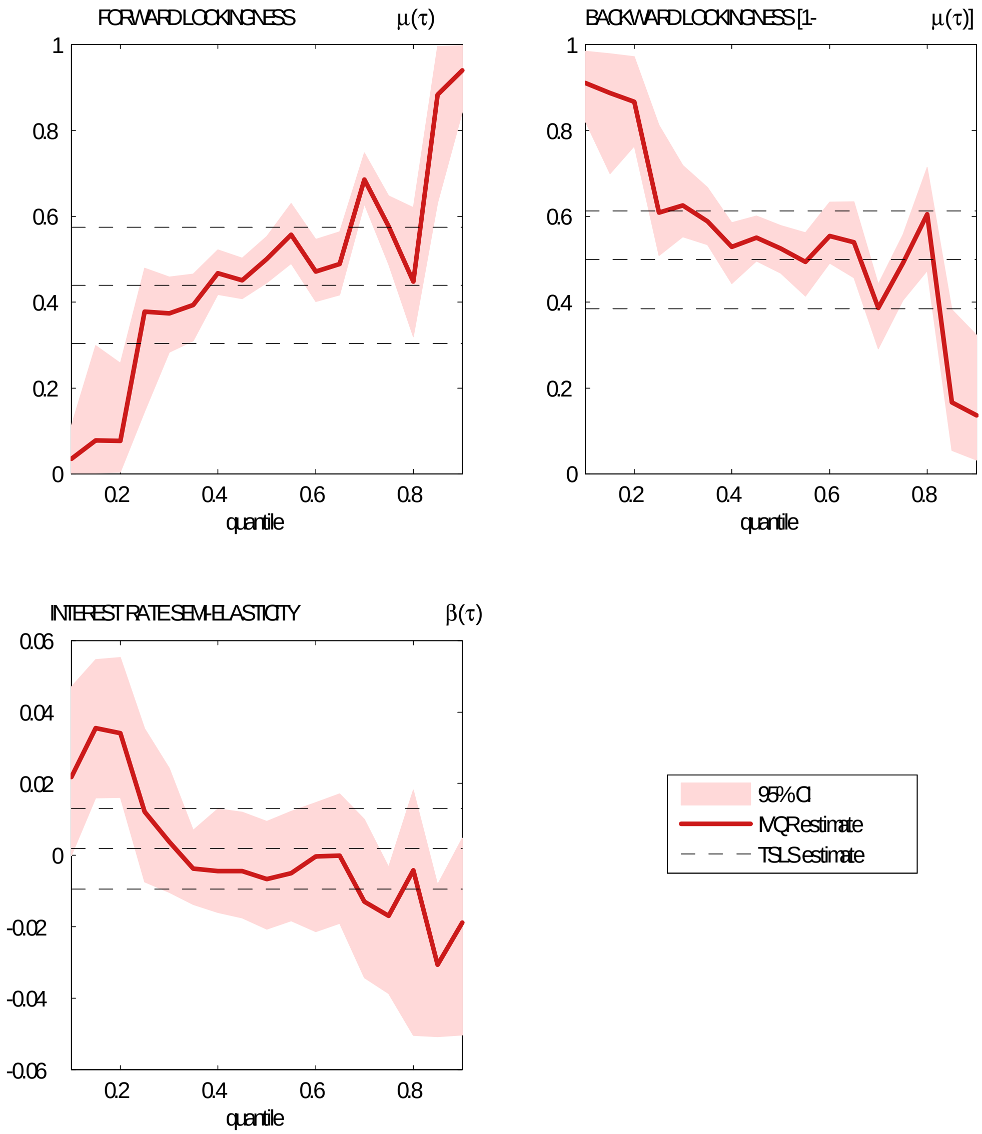

Figure 8: IVQR estimates of the parameters of equation (7) with $\kappa = 1, m = 0$ using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services, and the benchmark instrument set $Z_t = (c_{t-2}, c_{t-3}, i_{t-2}, \pi_{t-2}, 1)$, without imposing (i) the coefficient on expected real interest rate to be non-positive and (ii) the coefficients on backward- and forward-looking components to sum up to one. Dotted lines represent to TSLS estimates and 95% credible sets. Sample: 1948:5-2010:7.

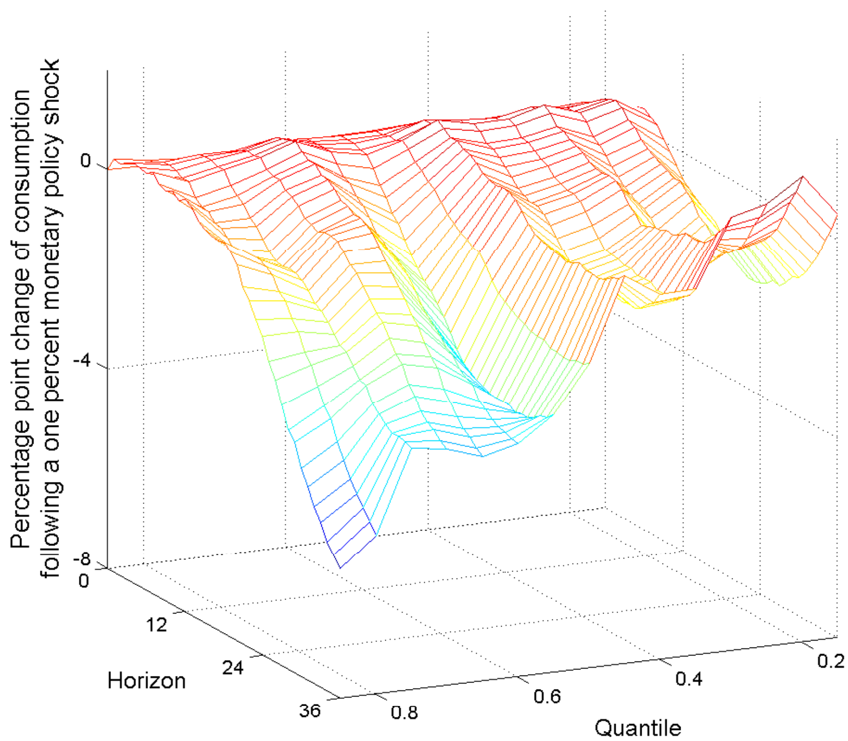

Figure 9: point estimates of the cumulated impulse responses based on (13) using the HP filter on the per-capita real personal consumption expenditure of non-durable goods and services and the measure of monetary policy shocks proposed by Romer and Romer (2004).

Appendix A: convergence

For each quantile τ, we report below the statistics proposed by Gelman and Rubin (1992). These statistics are based on multiple runs of the MCMC algorithm from overdispersed starting values.[6] The statistic compares the within and between estimate of the variance of parameters. Note that the former does not take the overdispersed starting values into account and will under-estimate the variance before convergence. If these two estimates are close together, then the MCMC algorithm has probably converged. Gelman and Rubin (1992) suggest a statistic that compares the two estimates. A value of this statistic near 1 is regarded as evidence in favor of convergence.

Table 5: Gelman and Rubin's convergence statistics

quantile τ	$\mu(\tau)$	$\beta(\tau)$	$\alpha(\tau)$
0.1	0.99995	0.99995	0.99995
0.15	1.0207	1.0001	1.0016
0.2	1.0036	1.0001	1.0017
0.25	1.0021	1.0001	1
0.3	1.0014	1.0002	0.99996
0.35	1.0026	1.0001	1
0.4	1.0067	1.0001	0.99996
0.45	1.0058	0.99996	1
0.5	1.0055	0.99999	1.0008
0.55	1.0157	1	1.0006
0.6	1.0073	0.99999	1.0002
0.65	1.0018	0.99996	1
0.7	1.0054	1.0001	1.0004
0.75	1.0049	1.001	1.0015
0.8	1.0031	1.0003	1.0003
0.85	1	0.99995	0.99995
0.9	1.001	1.0006	1.0006

Note: statistics based on Gelman and Rubin (1992)

[6] More specifically, we run the MCMC algorithm for our benchmark model using four different starting values. We base the starting values for $\tau = 0.1$ to 0.9 on the benchmark posterior estimates at different quantiles. For example, the first run of the MCMC algorithm uses the estimates of the parameters obtained at $\tau = 0.1$ and fixes this starting value for all quantiles. We then repeat this run using the posterior estimates at $\tau = 0.25, 0.75, 0.9$ as starting values. This gives a wide range of starting values of the algorithm.

Appendix B: montecarlo analysis

Chernozhukov and Hong (2003) present montecarlo evidence to show that the MCMC estimator performs well when considering a simple quantile regression model with exogenous regressors. Here we extend their simulation to instrumental variables by considering a state-dependent parameter version of the output/consumption equation studied by Fuhrer and Rudebusch (2004). In particular, we consider a version of the output/consumption equation where the interest rate semi-elasticity is assumed to vary. This DGP allows us to assess the ability of the IVQR estimator to recover coefficients that are fixed or vary across conditional quantiles. The artificial data are generated by the following process:

$$c_t = \alpha_0 (u) + (1 - \alpha) c_{t+1} + \alpha c_{t-1} - \beta (u) [i_t - \pi_{t+1}] \tag{14}$$

where $\alpha = 0.5, \beta (u) = 0.1 + 0.5u$ and $\alpha_0 (u) = \Phi^{-1} (u)$ with Φ^{-1} denoting the inverse normal cumulative distribution function (see Koenker and Xiao, 2006). We augment equation (14) with the following bivariate VAR(1) model which describes the dynamics of $z_t = \{i_t, \pi_t\}$

$$z_t = \tilde{c} + \tilde{B} z_{t-1} + \tilde{\Omega}^{1/2} e_t \tag{15}$$

where the coefficients $\tilde{c}$ and $\tilde{B}$ and the error covariance matrix $\tilde{\Omega}$ are estimated via OLS using HP filtered data for z_t over the sample period $1980m1$ to $2010m7$.

We generate three samples of $T + 50$ observations and discard the first 50 to reduce dependence from initial conditions. The length of the artificial samples are 100, 200 and 700 with the latter reflecting the typical number of observations used in the empirical investigation in the main text. An important goal of this montecarlo analysis is to assess the role played by the number of observations and the span of the data in determining the accuracy of the estimates. In order to isolate the impact of the data span, we report an additional experiment where we sample every $3rd$ observation from the largest generated sample to create a new 'quarterly' dataset. This dataset retains the same span as the original sample but has a $3rd$ of the available observations because of the different sample frequency. The estimates based on this sample can be used to gauge the role played by data span relative to data frequency (i.e. the total number of available observations for a given sample span).

For each observation, we draw the latent state u from a standard uniform distribution. We then solve the system given by equations (14) and (15) using the gensys.m solution algorithm proposed by Sims (2002). The reduced form representation of the structural model is used to generate data on c_t, i_t and π_t.

The model is estimated on this artificial data using $100,000$ replications of the MCMC algorithm for the 20% to 80% quantiles (with incremental steps of 5%). In line with the benchmark instrument set used in the analysis on actual data of Section 5, we use the second and third lags of c_t, and the second lag of i_t and π_t as instruments. The experiment is repeated 1000 times and the results are presented in Figure 10. The panels show the mean estimates across montecarlo replications (red line), the 90% confidence interval across the replications (shaded area) and the (sorted) mean of the true parameter values across the 1000 replications (black line).

The left panel of the figure reveals that the estimator is able to recover the true parameter even with a small number of observations as long as there is no parameter variation across quantiles. The right panel shows, however, that the number of observations matter when the underlying parameter varies across quantiles. Using 700 observations, the estimates track closely the underlying distributions with the true value always within the confidence interval. On the other hand, when the number of observations is lower, the point estimates tend to diverge from the true values. For example, the small sample bias associated with 100 observations could be as large as 150%.

Figure 11 reports the results from the additional experiment that considers the same span of data as the largest sample, but reduces the data frequency to 'quarterly'. In other words, this DGP has around 230 observations but retains the same span of data as the DGP with 700 observations. The right panel of figure 11 shows that the bias in the estimates is substantially larger when compared to the case when all 700 observations

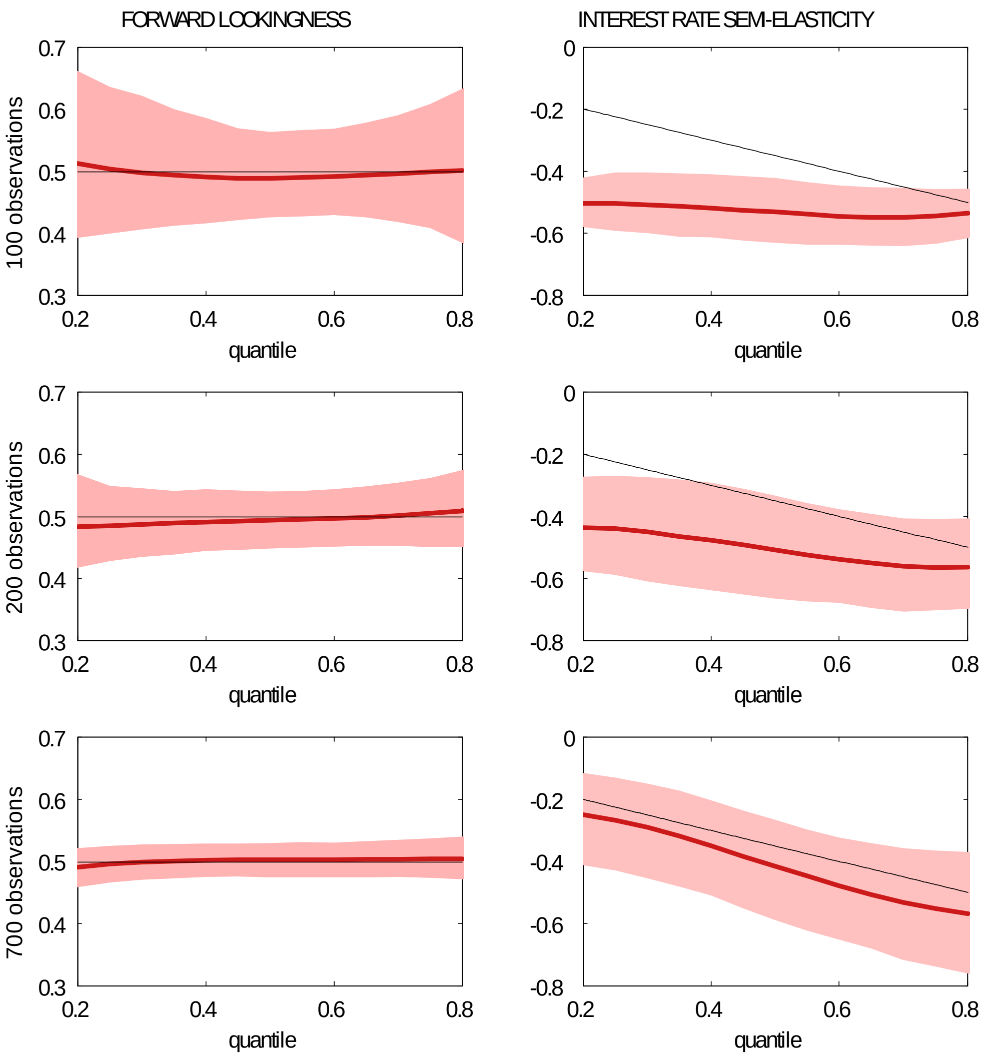

Figure 10: Montecarlo experiments based on different sample sizes.

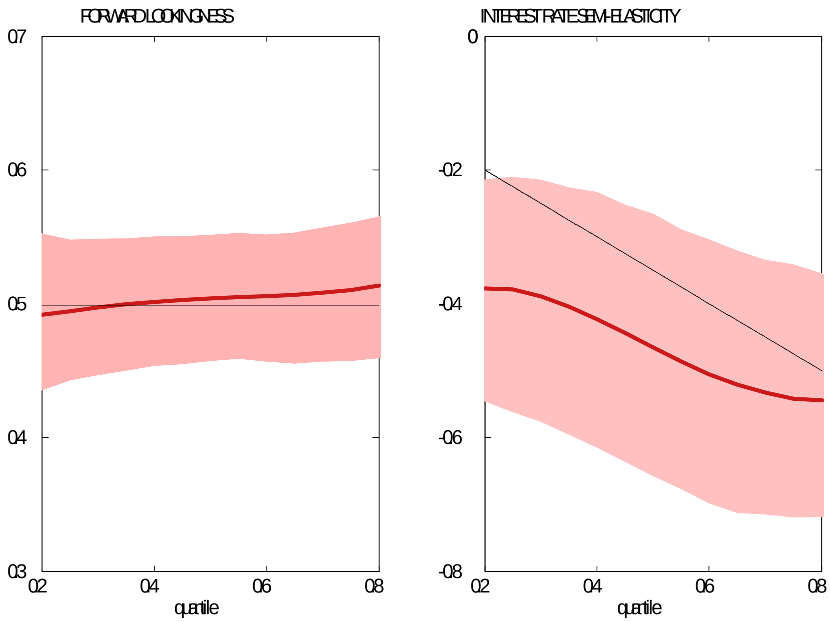

Figure 11: Montecarlo experiment based on quarterly sampling (i.e. one in three observations) the full-length data span of 700 observations.

are used. For example, the mean estimate is almost twice as small as the true value for quantiles below 0.5. This suggests that the frequency (and not exclusively the span) of the sample is important in ensuring good performance of the IVQR estimator.